

Multi-Provider MPLS-VPN Trials

roger.gottsponer@nextra.ch

SWINOG - 2 / 21.3.01

Agenda

- **Lessons learned**
- **Short Introduction to MPLS-VPN**
- **Why do we need Multi-AS VPNs?**
- **Some different Multi-AS scenarios**
- **The Trust Issue**

Agenda

- **Lessons learned**
- **Short Introduction to MPLS-VPN**
- **Why do we need Multi-AS VPNs?**
- **Some different Multi-AS scenarios**
- **The Trust Issue**

**A reload a day
keeps the TAC away**

just kidding...

Lessons learned


- **MPLS gives lots of interesting features**
 - helps to build new customer services (VPN, QoS)
 - and to save money (Traffic Engineering)
- **MPLS makes our life much more difficult**
 - much more difficult for operation people to understand what they do
 - much more difficult for engineers to troubleshoot
- **You often have to wait for new features you would like to use**
- **You loose other features because they are not “vpn-aware” (i. e. NAT)**
- **Sometimes you get a new feature which will have a Bug-ID the next day**
- **It can be very difficult if not impossible to get accurate documentation**
- **If you are one of the first customer using new features like MPLS/VPN in a live-network, the vendors support is greeeeeeaaaaaaaaaat!**
- **The real problems are not technology related...**
- **And after all, it is great fun...**

Agenda

- Lessons learned
- **Short Introduction to MPLS-VPN**
- Why do we need Multi-AS VPNs?
- Some different Multi-AS scenarios
- The Trust Issue


MPLS - Label Switch Path LSP

- LDP and IP routing create **Label Switch Paths** between PE routers
- PE routers 'encapsulate' IP packet with Label header
- somewhat similar to Frame Relay Switching (Label = DLCI, LSP = PCV)
- but LSPs are set up dynamically




MPLS - VPN

- It does not make sense to build a LSP for each Internet route, so BGP is used for external routes
- LSPs are now only needed for the BGP next-hops.
- BGP session between PE routers, Core LSRs do not have any BGP routing
- Core does not see any IP addresses, so we can build VPNs
- But routes and addresses are still visible in PE routing table!
- Solution is to assign a "Virtual Router" to each Customer port
- Routing Tables in Virtual Routers are invisible to each other
- Cisco name: Virtual Routing and Forwarding Instance (VRF)
- PE uses 2nd Level Label (L2) to distinguish between VPNs




MPLS - VPN Routing

- PE Router BGP table is populated by routes exported from VRFs
- Problem: normal BGP cannot tell the difference between identical routes from different VRFs
- Add Route Distinguisher (RD) to BGP routes. RD makes the routes globally unique
- PE Routers cannot tell which BGP routes belong to which VRF
- RD does not have this info - it only makes the routes unique
- Add Route Target (RT) to BGP routes.
- Route is tagged with RT when exported from VRF
- RT tells PE which routes belong to which VRF



MPLS - VPN Forwarding

- Packet arrives on interface E = red VRF
- PE looks up route to 10.0.0.1 in VRF Routing Table
- BGP route from 3.3.3.3 gives L2 = 10
- Recursive lookup on Next Hop 3.3.3.3 gives L1 = 3
- Packet label switched through Core to 3.3.3.3
- L1 removed
- L2 tells PE router to treat packet according to red VRF
- L2 removed, packet forwarded out of interface A



Agenda

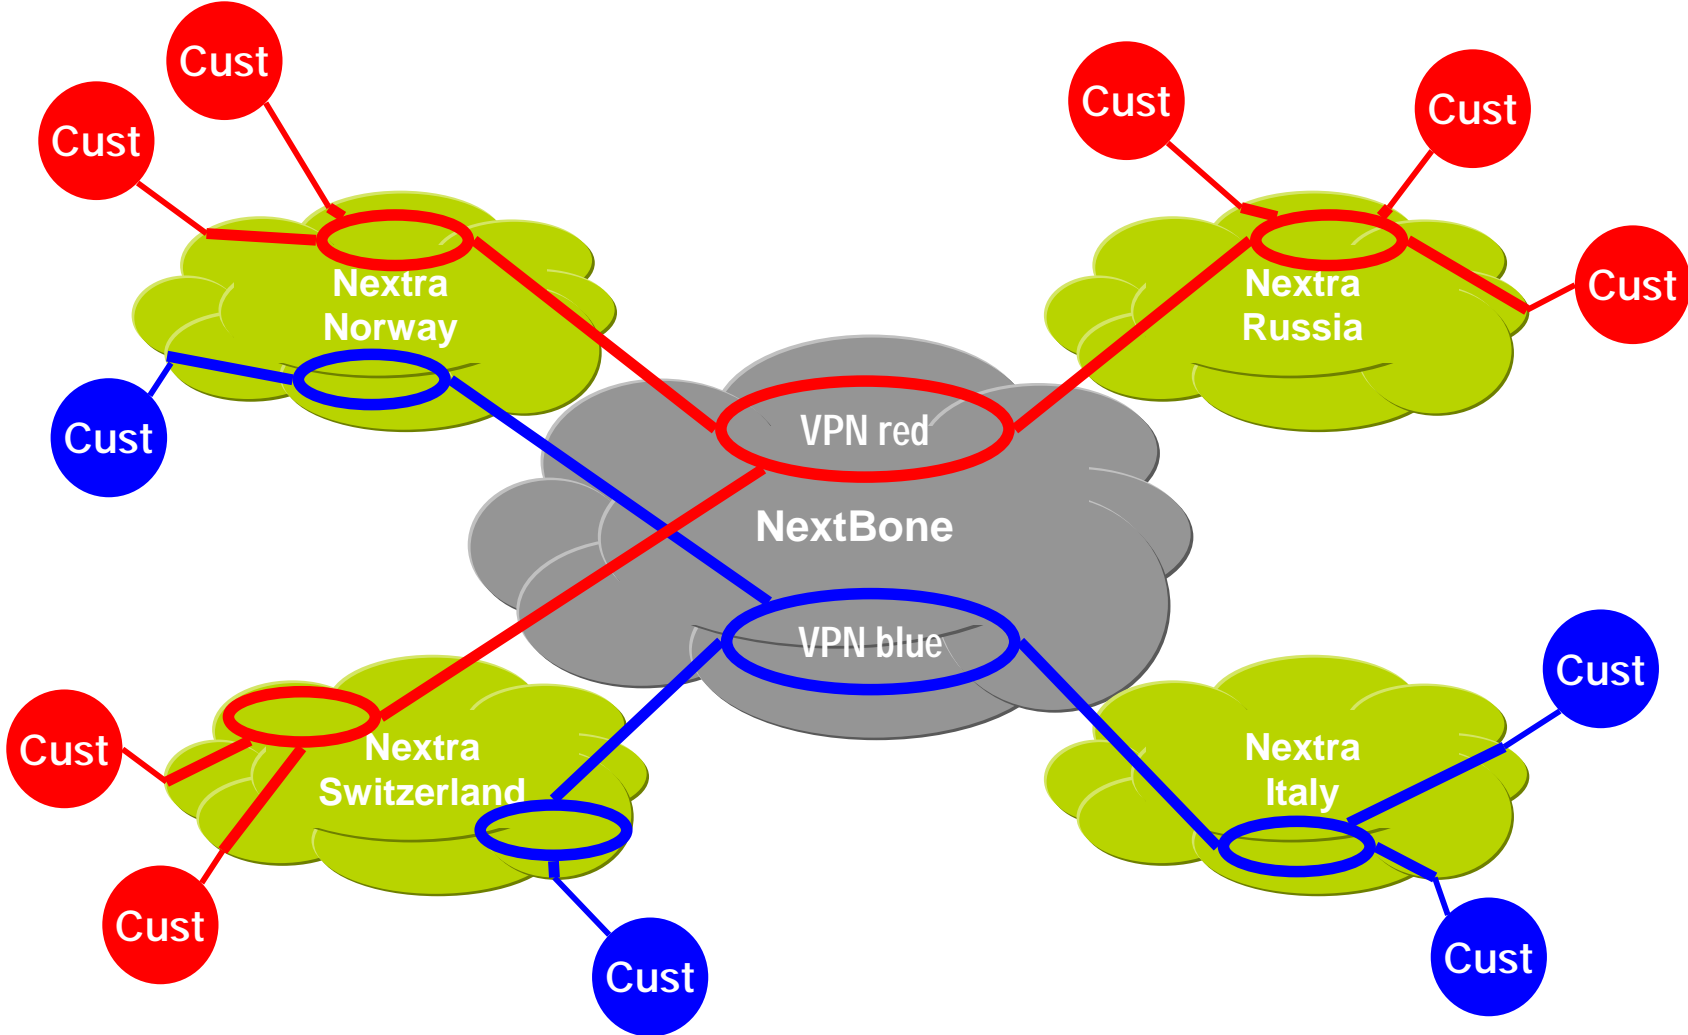
- Lessons learned
- Short Introduction to MPLS-VPN
- **Why do we need Multit-AS VPNs?**
- Some different Multi-AS scenarios
- The Trust Issue

Why do we need Multi-AS VPNs?

job security

just another bad joke...

Why do we need Multi-AS VPNs?



Why do we need Multi-AS VPNs?

- **Plain MPLS-VPN works within one single ASN only.**
- **Sometimes this restriction is not acceptable:**
 - International Provider wants to build an international VPN service between its subsidiaries
 - Provider wants to build VPNs with other Provider (“VPN-Peering”)
 - Provider wants to benefit from BGP confederations
- **Therefore we need Multi-AS MPLS-VPN functionality**

Agenda

- Lessons learned
- Short Introduction to MPLS-VPN
- Why do we need Multi-AS VPNs?
- **Some different Multi-AS scenarios**
- The Trust Issue

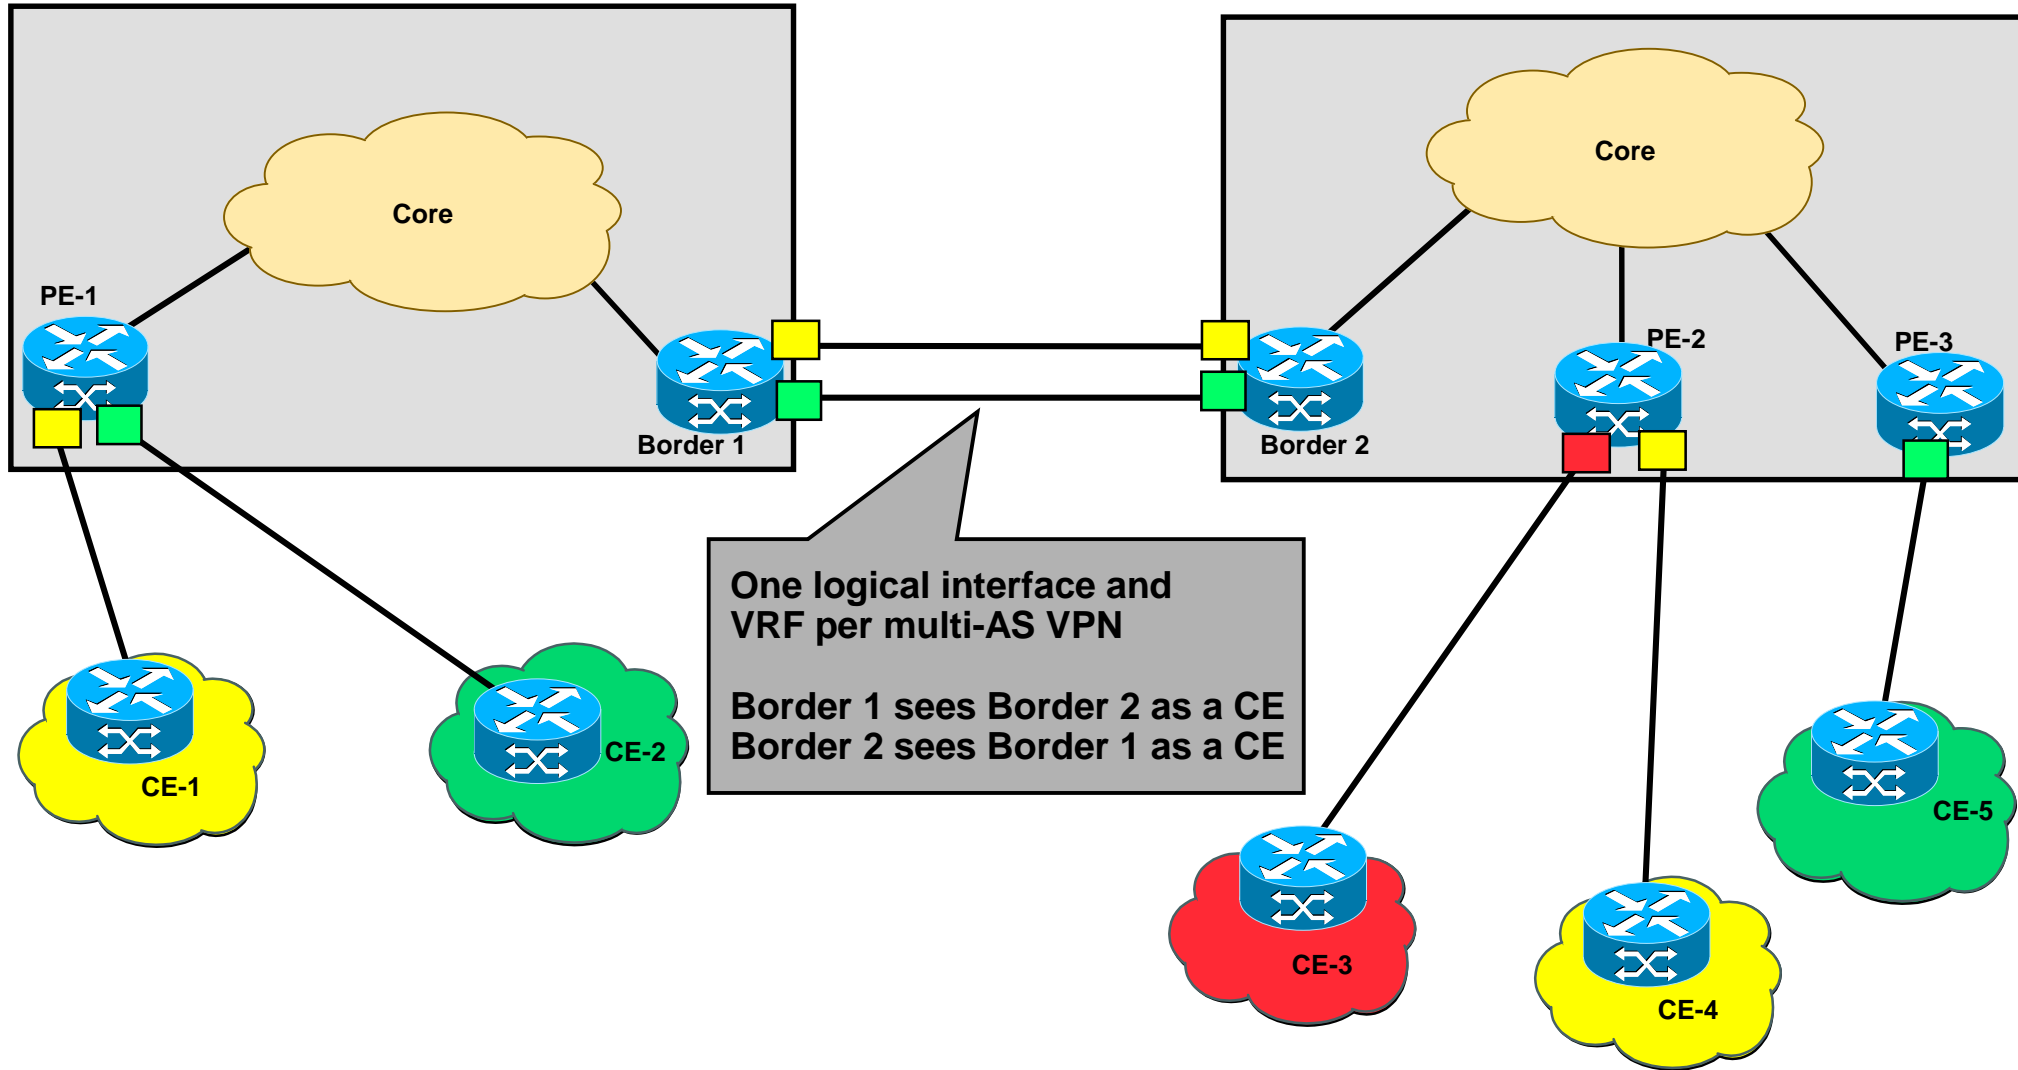
Some different Multi-AS scenarios

- **Providers use VRFs to connect their border routers in a back-to-back fashion**
- **Providers exchange routes between border routers**
- **Providers exchange routes between their Route-reflectors**
- **The Nextra Scenario**

Scenario 1: multiple back-to-back VRFs

- **MPLS-VPN service providers exchange VPN routes using VRFs and logical interfaces (i.e. 802.1q) between border routers**
- **Each border router will consider the peer border router as a CE router**
- **Border routers may use different routing protocols**
 - currently RIP, OSPF, BGP and static are supported

Scenario 1: multiple back-to-back VRFs



Scenario 1: multiple back-to-back VRFs

- **Advantage**

- No MPLS needed between providers**

- Plain IP between providers VRFs
 - Technology we already know
 - **Works today** with MPLS-VPN software

- **Disadvantage**

- Scalability problems if many VPNs**

- One VRF and logical interface per multi-AS VPN
 - Each border router needs to store all multi-AS VPN routes

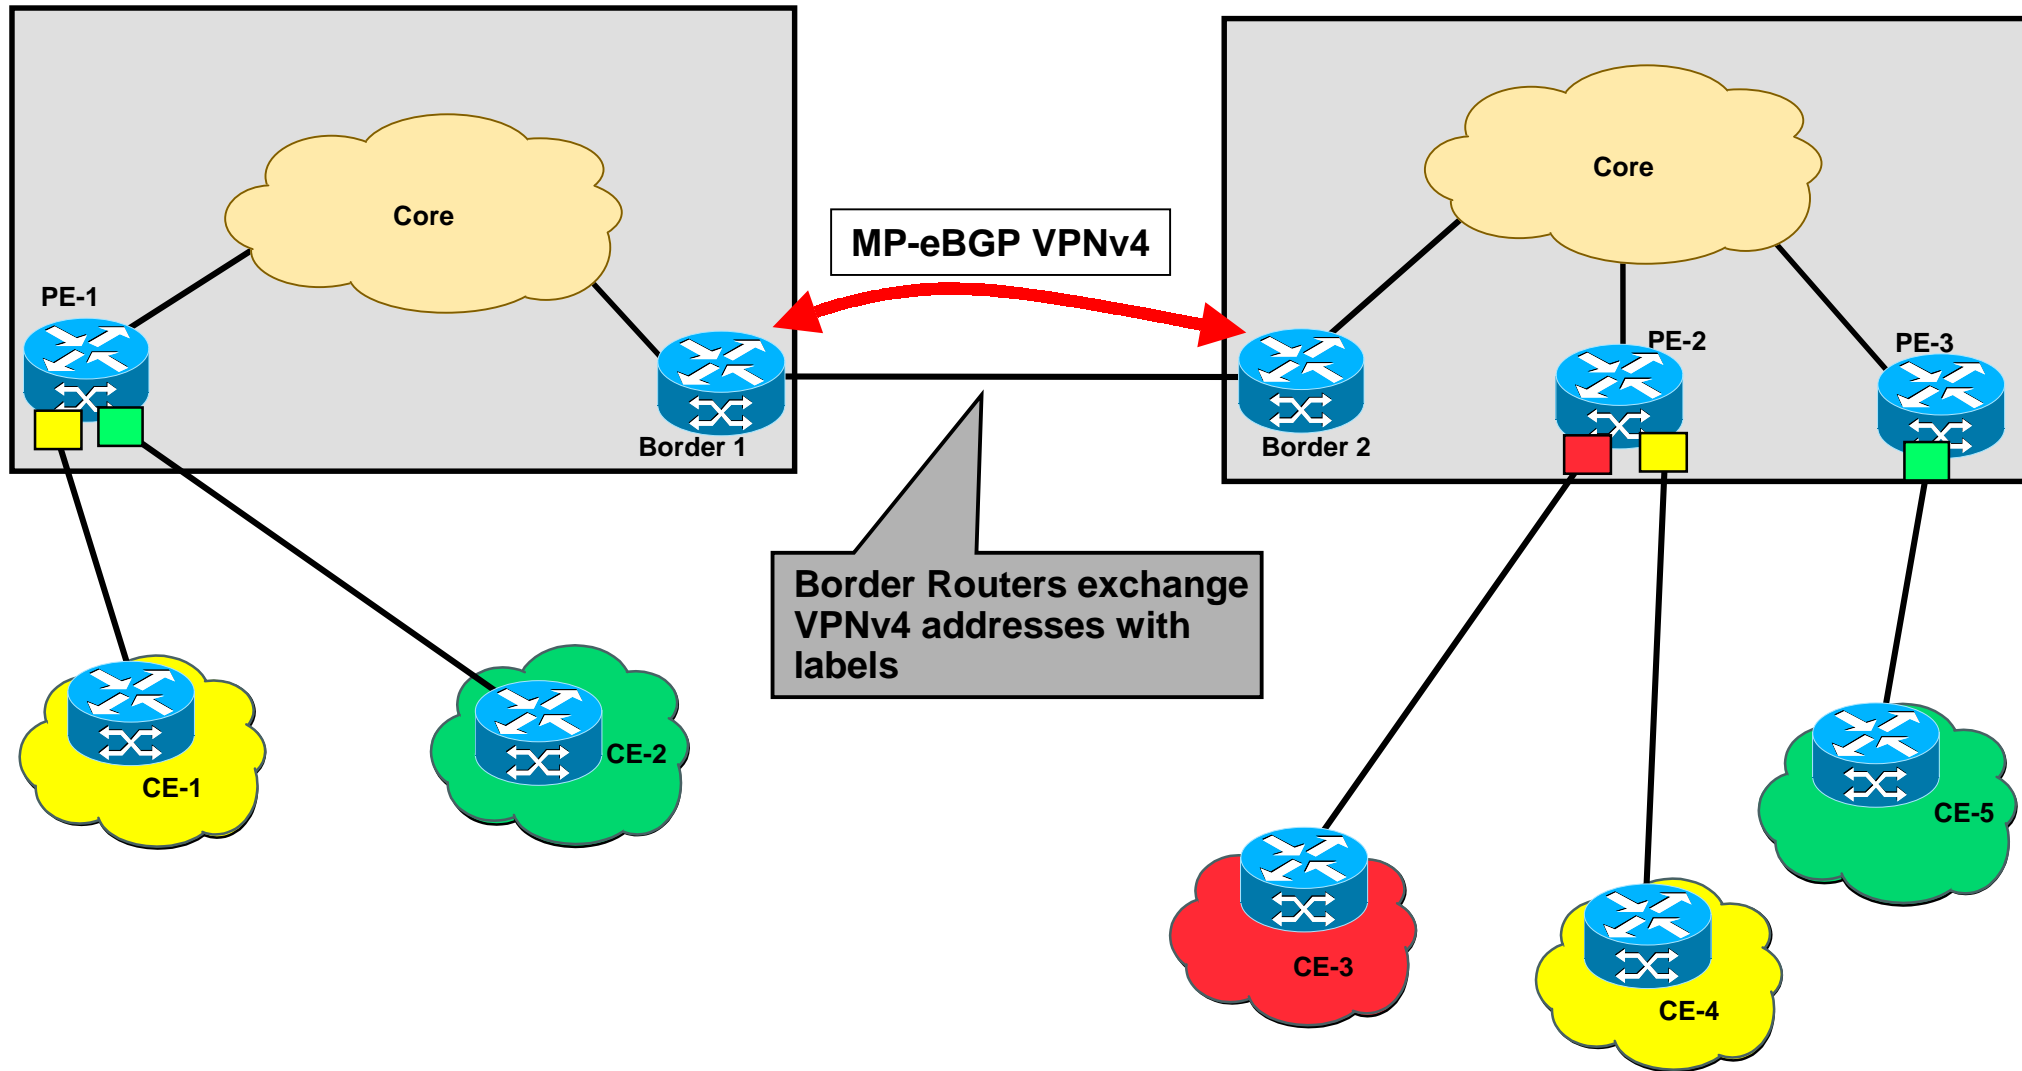
Some different Multi-AS scenarios

- Providers use VRFs to connect their border routers in a back-to-back fashion
- Providers exchange routes between border routers
- Providers exchange routes between their Route-reflectors
- The Nextra Scenario


Scenario 2: VPNv4 routes exchanged between border routers

- **Providers exchange VPNv4 routes between border routers**
- **MP-eBGP for VPNv4 addresses between border routers**
 - Next-hop and labels are re-written by the border router
- **Builds direct LSPs between the PE routers in the different networks**

Scenario 2: VPNv4 routes exchanged between border routers



Scenario 2: VPNv4 routes exchanged between border routers



Scenario 2: VPNv4 routes exchanged between border routers

- **Advantage**

- no VRF and no logical interface per multi-AS VPN
- **Works today** with 12.1T

- **Disadvantage**

Special Configuration between border routers needed

- Technologie is new
- Can be a headache to troubleshoot (traceroute is no longer your friend!)

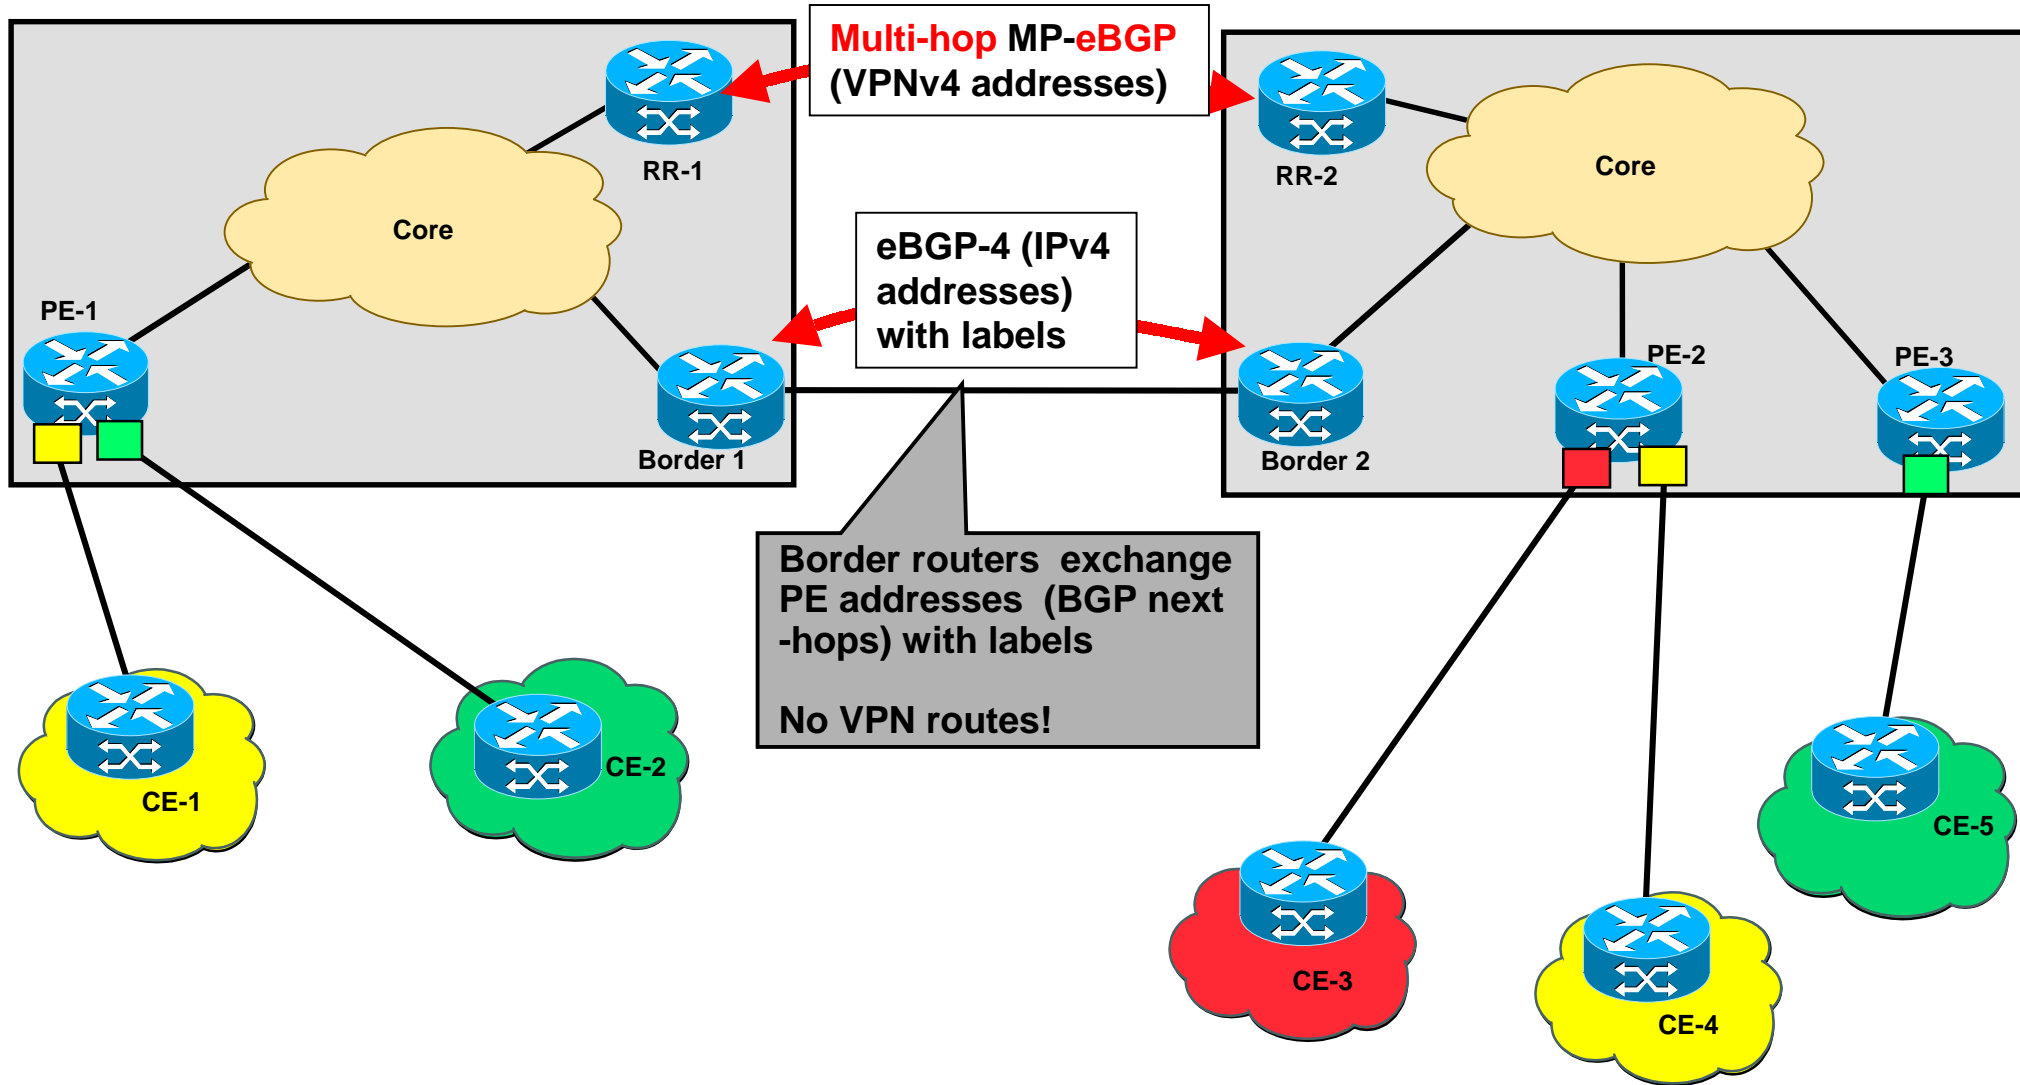
Some different Multi-AS scenarios

- Providers use VRFs to connect their border routers in a back-to-back fashion
- Providers exchange routes between border routers
- Providers exchange routes between their Route-reflectors
- The Nextra Scenario


Scenario 3: VPNv4 routes exchanged between Route-Reflectors

- **Providers exchange VPNv4 routes between their Route-reflectors using multihop MP-eBGP**
 - label and next-hop are preserved
- **Providers exchange IPv4 routes and labels between directly connected Border Routers using eBGP**
 - Include routes to all provider's PEs
- **Builds direct LSPs between the PE routers in the different networks**

Scenario 3: VPNv4 routes exchanged between Route-Reflectors



Scenario 3: VPNv4 routes exchanged between Route-Reflectors



Scenario 3: VPNv4 routes exchanged between Route-Reflectors

- **Advantage**

Improves scalability

- No VPNv4 routes at border routers (Route-reflectors already store VPNv4 routes)
- no VRF and no logical interface per multi-AS VPN

- **Disadvantage**

Special Configuration between route reflectors needed

- Technologie is new (as with scenario 2)

Does not works today.


Some different Multi-AS scenarios

- Providers use VRFs to connect their border routers in a back-to-back fashion
- Providers exchange routes between border routers
- Providers exchange routes between their Route-reflectors
- **The Nextra Scenario**

Scenario 4: The Nextra Scenario

- We use scenario 2 between NextBone and the country networks
- NextBone itself is running MPLS and transports VPNv4 routes, but does not have VPNs directly connected to it
- Country border router and NextBone border router are exchanging VPNv4 routes
- Direct LSPs between all PE routers in all countries are built

Scenario 4: The Nextra Scenario



Scenario 4: The Nextra Scenario

- **Advantage**

Scales quite well

Available today :-) (and it really works)

- **Disadvantage**

Technologie is new

Can be a headache to troubleshoot

– If you already have managed to handle MPLS-VPNs, you will also get used to the Multi-AS stuff

Agenda

- Lessons learned
- Short Introduction to MPLS-VPN
- Why do we need Multi-AS VPNs?
- Some different Multi-AS scenarios
- **The Trust Issue**

The Trust Issue

- **So we have a nice technical solution to build international VPNs**
- **But now we should start to think about security**

- **Who trusts whom? or Who trusts who?**

- **Example 1: Customer has only sites in Switzerland**
 - Will customer trust Nextra Switzerland?
 - *Yes, because there is no way around this*
 - Will customer trust Nextbone or any other Nextra countries?
 - *No, as he has a national VPN, there is no reason for involving other parties.*
 - Other question: What will happen if a configuration error in another country breaks security of his national VPN?
 - *Customer will blame Nextra Switzerland because external configuration errors do influence the Swiss network. We should filter on the borders!*


The Trust Issue

- **Example 2: Customer has only sites in Switzerland + Norway**
 - Whom trusts the customer?
 - *He will trust the directly involved parties: Switzerland, Nextbone, Norway.*
 - Will customer trust another Nextra country?
 - *No, because the other countries are not involved.*
 - Other question: What will happen if an intruder successfully attacks customer network through a security breach in Nextra Italy?
 - *Customer will blame Nextra Switzerland because we did not prevent Italian problems having influence in Switzerland. We should filter on the borders!*

The Trust Issue

- **So we need a per-customer trust model which depends on the customer site locations**
- **For each VPN, we should only accept routing information from the countries where the VPN customer has sites**
- **This means: filtering!**
- **When receiving a VPN route from Nextbone, always check whether the originating country (visible in AS-path) is allowed to send routes with this RT (color)**
- **This is possible to do with today's software, but requires a lot of work to maintain the filters**

The Trust Issue



The Trust Issue

- So we can define per VPN which countries we trust
 - But we still have to trust the countries...
 - A country could send routes belonging to the **blue** VPN with **red** RT
 - We need some authentication to make sure that a route we accept really belongs to the customer we think it belongs to
 - There is no code around today, even there is no clear picture of how to do the authentication
- ➔ A lot of open issues here...

That's it!

Thanks to Phil and to Cisco for the nice drawings

Thank you for your attention

Now, it is time for

- questions
- comments
- corrections
- flames

Multi-Provider MPLS-VPN Trials

roger.gottsponer@nextra.ch

SWINOG - 2 / 21.3.01