Case Study

A Service Provider's Road to IPv6

September 2010

Menog

Amir Tabdili – UnisonIP Consulting

amir@unisonip.com

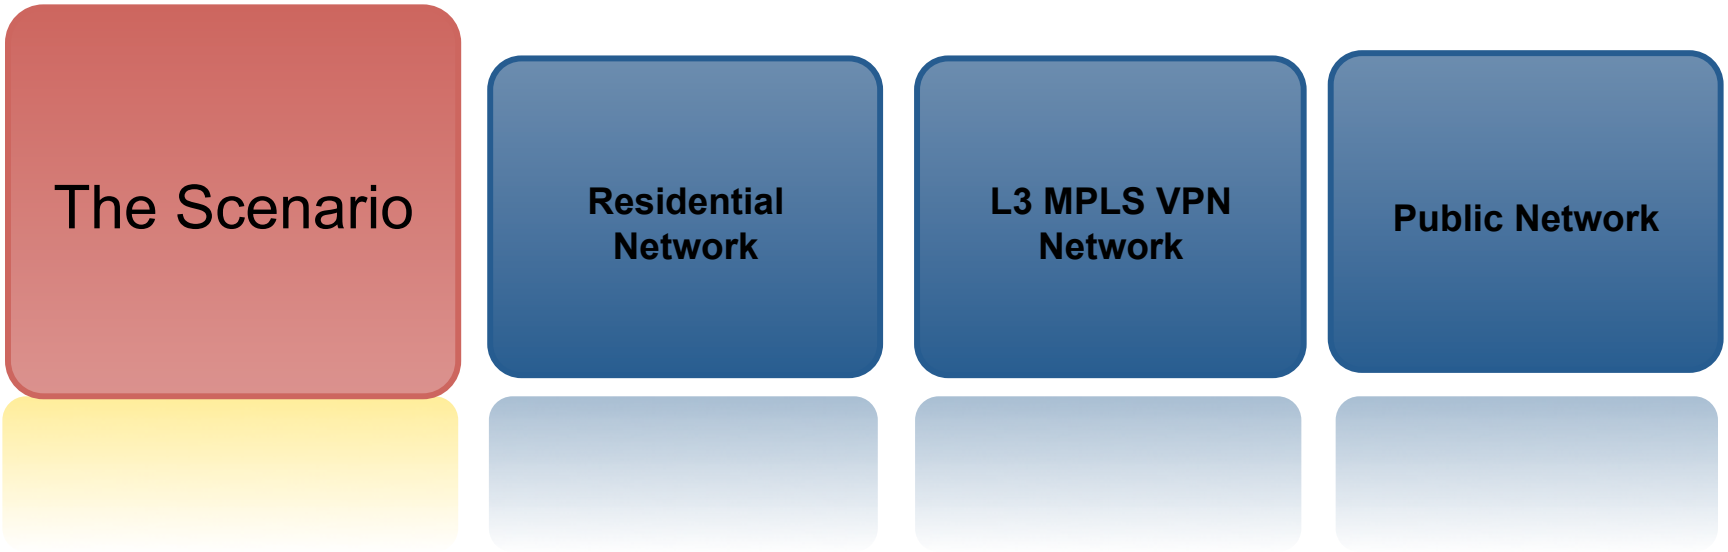
The Scenario
Residential Network
L3 MPLS VPN Network
Public Network

The Scenario

What are we presenting and why?

- **A large European service provider asked us to provide design for IPv6 rollout for various networks it operates**
  - **Residential (BRAS) network (LAC and LNS)**
  - **L3 MPLS VPN network for business customers**
  - **Public network for Internet Access**

- **This presentation shows a typical service provider's dilemma:**
  - ***SPs often operate many multivendor networks***
  - ***Need to focus on all pieces, not just public Internet***
  - ***Inter-dependency of services across networks.***
    - **Regular BRAS customers terminate in public Internet while L2TP wholesale service terminate on MPLS VPN network**

- **What we show is just what one SP decided on**
  - **More than one correct way to go**

The Scenario

Background and Assumptions

- **Deployment Scenario: "Dual-Stack"**
  - **Considered first step**
- **For now, only convert what is visible externally. Example:**
  - **No need to convert network monitoring tools, ssh, telnet, etc.**
  - **Ignore (for now) infrastructure networks (IPTV, etc)**
- **Devices Deployed**
  - **E320/ERX Juniper routers as LAC/LNS**
  - **Juniper T /M series as P And PE MPLS VPN routers**
  - **Juniper T/M series and various Cisco routers as Internet routers**
- **The roll out should be transparent to existing customers**
- **Current IPv4 transport models**
  - **IPv4 only for Public network**
  - **MPLS for the L3 MPLS VPN network**
  - **The Residential Network uses VPLS for aggregation but that is irrelevant to IPv6**

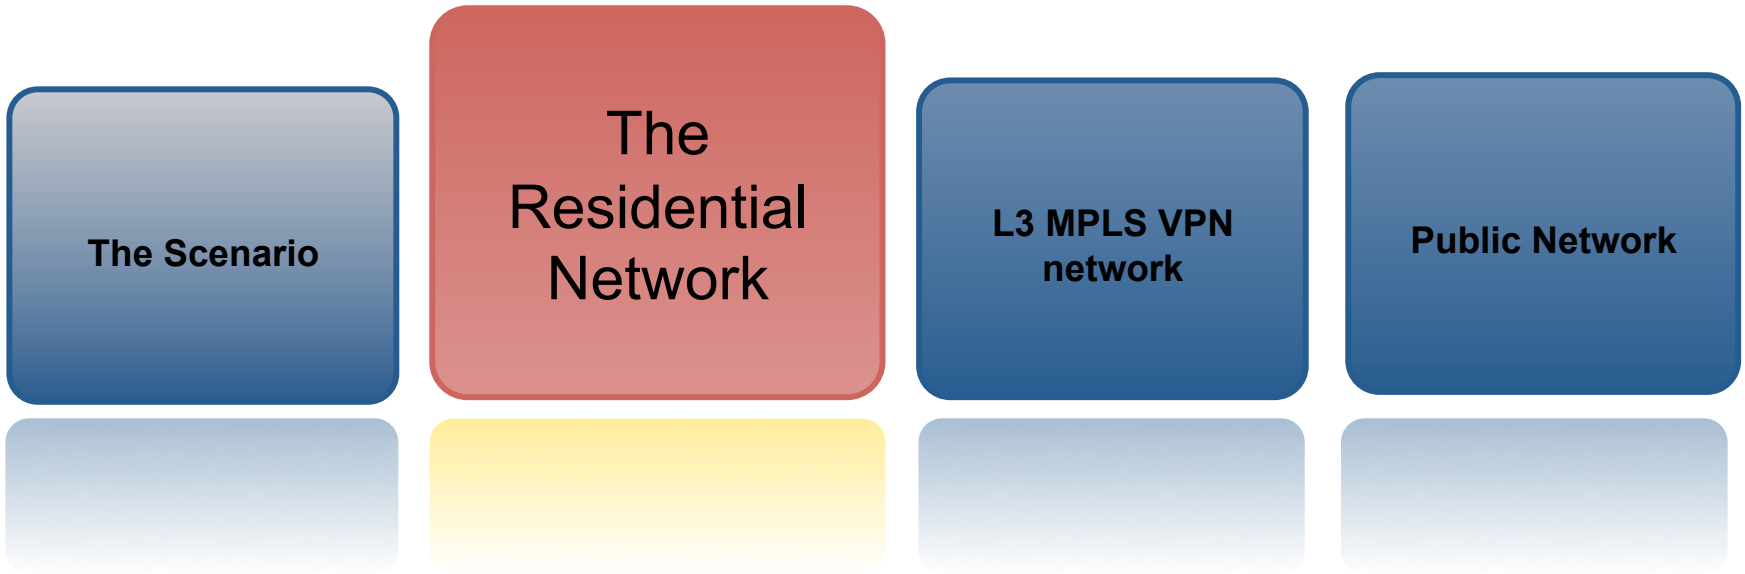
The Scenario
The Residential Network
L3 MPLS VPN network
Public Network

Residential Network

Main Connectivity Models

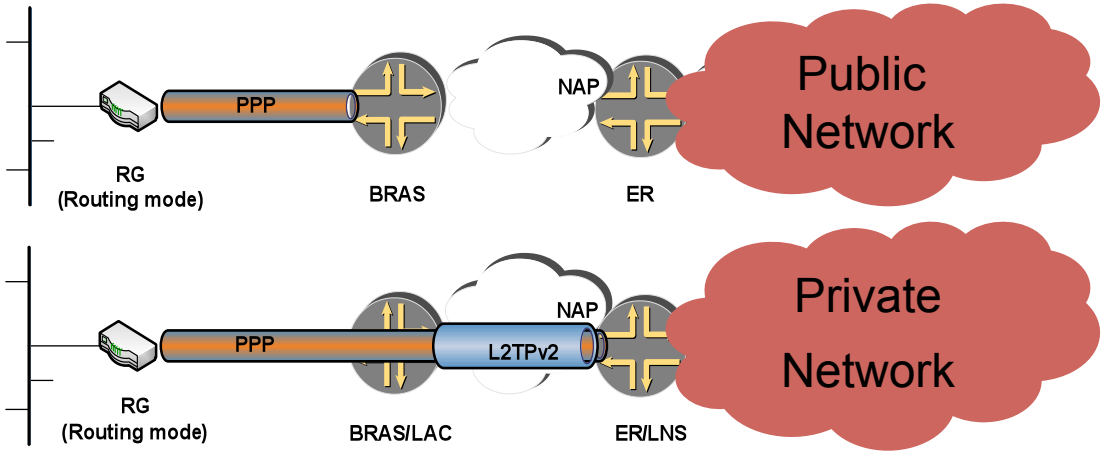

- **Simple Routed Mode: CPE establishes a PPP session to the BRAS**
  - **SP's own customers:**
    - **The PPP session terminates in one of SP's two public virtual routers**
  - **Other ISPs' customers:**
    - **The PPP session terminates in a virtual router allocated to another ISP.**
- **L2TP Backhaul: CPE establishes a PPP session and (LAC) creates L2TPv2 session to an LNS.**
  - **SP's own Enterprise service**
    - **LNS is an ERX and owned by SP, Subscriber terminates in MPLS VPN**
  - **Wholesales**
    - **LNS is owned by another ISP**

Residential Network

Our Choices for PPP Model– Dual Stack or Dual Session

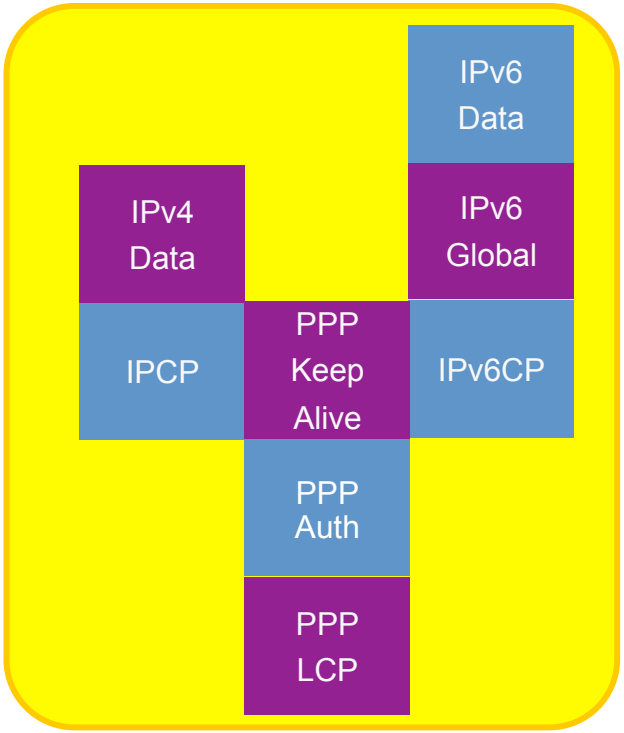

- **One AAA interaction**
- **Most flexible**
- **CPE Driven**

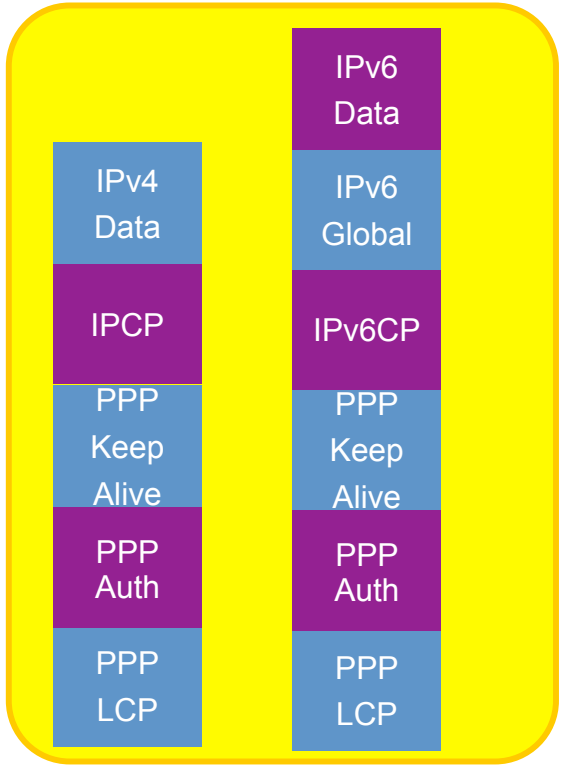

- **Can be interesting for transition:**
  - **IPv6 LNS?**

Residential Network

E320/ERX IPv6 Configuration – Getting Started

- **Many business questions needed to be answered first**
  - **What are the offerings?**
  - **Do all IPv4 services make sense for IPv6?**
  - **Which customers should get an IPv6 address?**
    - **Only new customers**
    - **New and existing customers**
- **Everything is influenced by scaling!**
  - ***In many regards scaling drives many of the design decisions***
- **License Key**
  - **Depending on the vendor**
  - **Enabling of IPv6 in JUNOSe required activation of a license key**
    - **IPv6 needed to be enabled per Virtual Router**

Residential Network

E320/ERX IPv6 Configuration – Interfaces and Routing

- **Backbone interfaces**
  - **Need to have new IPv6 Addresses, /64 netmask**
  - **New IPv6 loopbacks are needed, /128 netmask**
- **ISIS**
  - **Will also carry the IPv6 topology info**
  - **Same SPF calculation**
  - **Will also carry IPv6 loopback as passive interface**
- ***BGP***
  - ***Two Options:***
    - ***Native IPv6 end points***
    - ***IPv4 BGP carrying IPv6 NLRI***
  - ***Solution Picked:***
    - ***New BGP sessions using TCP over IPv6***
    - ***Independent planes***
    - ***Less disruptive to existing customers***
- ***Policy***
  - **The current routing policies achieved through route-maps need to have IPv6 equivalents**

Residential Network

Subscriber Addressing Model

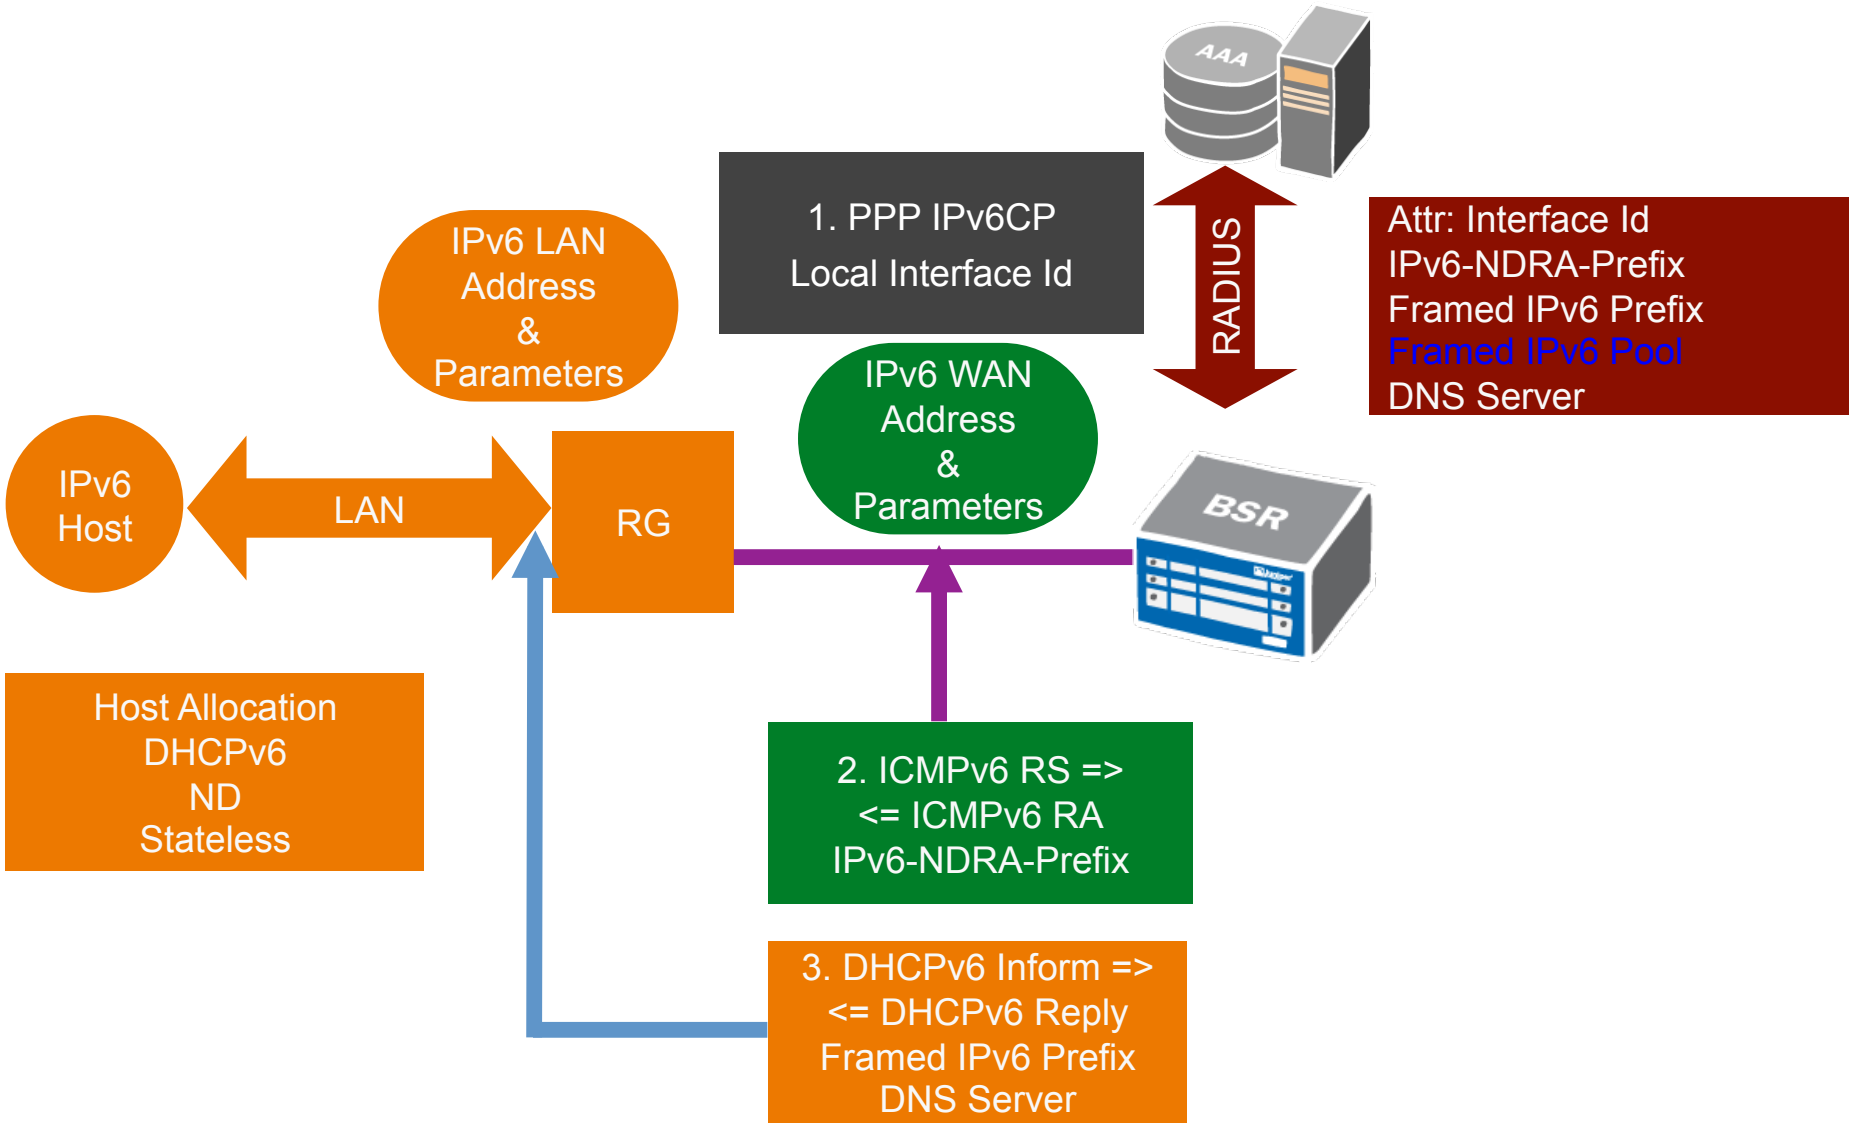

Residential Network

E320/ERX IPv6 Configuration – Address Assignment/Delegation

- **Bigger addresses – /32 => /(64 + 64)**
- **In case that's not enough, we get 3 addresses:**
  - **Link Local**
  - **CPE WAN Side from ICMP ND/RA**
  - **CPE LAN Side from DHCPv6-PD**

- **ICMPv6 (ND)**
  - **May be returned by RADIUS in IPv6-NDRA-Prefix**
  - ***Our Choice: Configured in the profile***
    - ***The /64 ND address is shared for all subscribers on a BRAS***

- **DHCPv6 (PD)**
  - **IF Returned by RADIUS: Two most common RADIUS attributes are in the form of Framed-IPv6-Prefix Attribute, or Framed-IPv6-Pool attribute.**
  - **Our choice:**
    - **Static: Address Pool Assigned in the domain-map**
    - **Intelligence in the provisioning system**
      - **No change to the RADIUS**

Residential Network

E320/ERX IPv6 Configuration – Subscriber Interfaces and Profiles

- **Subscriber interfaces**
  - **No new IPv6 specific configuration is needed since SP used a dynamic PPPoE / PPP stack**

- **Changes are contained in profiles**
  - **Customer preference:**
    - **Correct profiles attached to the subscriber interfaces at provisioning**

- **Profile modifications are needed for the following options**
  - **The loopback interface for IPv6 interface**
  - **Neighbor Discovery (ND)**
  - **IPv6 Policy**
  - **RPF check for IPv6 source validation**
  - **Other optional configurations such as virtual-router assignment**

Residential Network

E320 IP Configuration – DNS Servers

- **On the ERX/E320 Platforms DNS server related configuration can be configured in multiple locations**
  - **using "aaa ipv6-dns" command**
  - **Under local address pools**
  - **Through DHCPv6-LS configuration**
- **Choice driven by**
  - **How many DNS servers are needed**
  - **Is there a need to override the static DNS assignment by RADIUS**
    - **In Junose Using "aaa ipv6-dns" only will give the user choice to override local settings by RADIUS**
    - **Two locally configured DNS servers can be replaced by RADIUS per subscriber**

Residential Network

E320/ERX IPv6 Configuration – Accounting and Counters

- **Accounting**
  - **Accounting of IPv6 services equivalent to IPv4**
  - **Define which attributes are in the records sent to the RADIUS**
    - **All IPv6 attributes included in the Access-Accept from RADIUS can be included in the RADIUS accounting**

- **Counters**
  - **Access to the PPP session counters**
    - **PPP frames and octets**
    - **As IPv4 and IPv6 run on top of one PPP session, this session counters include the IPv4 and IPv6 packets**
  - **Separate counters for IPv6 are supported**
    - **Only count the IPv6 packets**

Residential Network

LNS Specific Configuration - Highlights

- **Most concepts discussed for LAC also apply to the LNS**
  - **The L2TP tunnel end points stay on IPv4**
- **A few tweaks needed.**
  - **Examples:**
    - **Some configurations will also be applied to specific customer VRFs**
      - **All the address-assignment configurations**
    - **We have a virtual router that communicates with L3 MPLS VPN network**
      - **For L3 MPLS VPN access to the corporate networks**
      - **Need to turn on vpnv6 BGP address family**
        - **May be Service Interrupting**

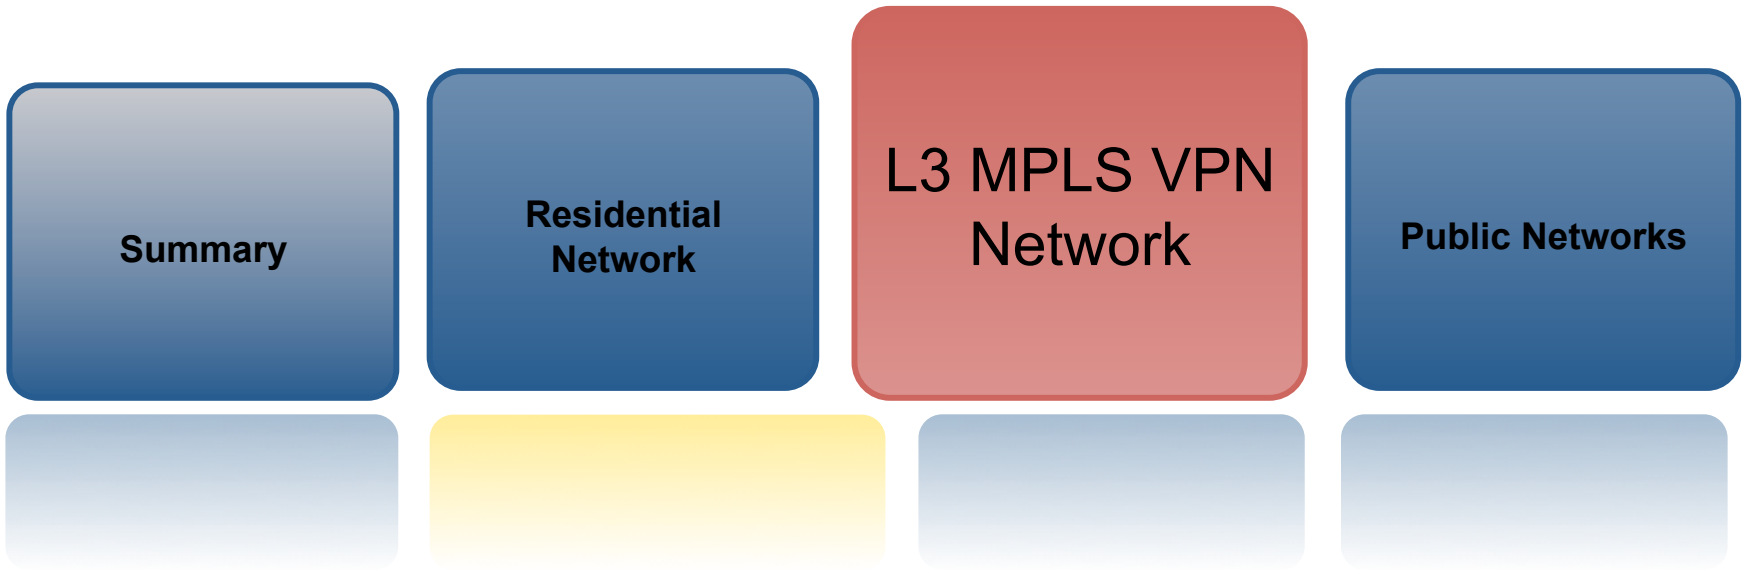
Summary
Residential Network
L3 MPLS VPN Network
Public Networks

L3 MPLS VPN Network

IPv6 Transport Architecture

mplsvpn network:

6VPE Architecture

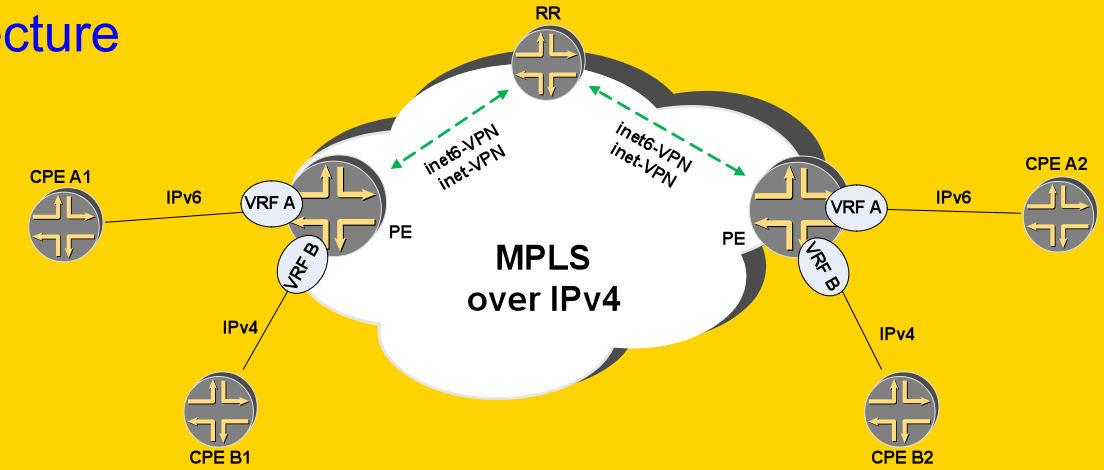

- In principle can be relatively straight-forward
  - Operational model and configuration are very similar to IPv4 VPN
  - Can use same LSPs and same BGP sessions as for existing IPv4 VPNs
  - Simply turn on VPN-IPv6 address family on the BGP sessions
- Same features as for IPv4 VPN can be used:
  - Packet processing features on ingress and egress PE
  - Route Target Filtering
  - Accounting features

L3 MPLS VPN Network

Configuration- Core

- **MPLS**
  - **MPLS used in the core for forwarding**
  - **In Junos, All P and PE routers need "ipv6-tunneling" statement under [protocols mpls]**
    - **So that IPv6 routes are resolved over the LSP tunnels.**
    - **Otherwise no IPv6 traffic will flow through the LSPs**
- **BGP**
  - **BGP needs modifications in the core**
    - **family inet6-vpn added**
    - **Applied to both PE and P routers , to all the iBGP related peer-groups**
  - **Interprovider VPNs**
    - **Same configuration to be followed for Option C peers**

L3 MPLS VPN Network

Configuration- PE-CE Routing

- **All the relevant routing protocols are carried forward into IPv6**
- **We had to generate equivalent templates for EBGP, Static and RIP-NG**
- **BGP**
  - **Configured using a new peer-group**
  - **Set the prefix-limit option for all IPv4 and IPv6 customers**
  - **Idle-timeout forever**
- **RIP-NG/Static**
  - **RIP-NG and Static Configurations are very simple and follow the IPv4 model**
  - **In the case of RIP, similar to BGP routing policies need to be converted to IPv6**

L3 MPLS VPN Network

Configuration- Quality of Service

- **QOS in the core is untouched**
  - **MPLS EXP in the core is blind to IPv6**
- **Customers use BA and MF classifiers**
  - **BA Classification**
    - **New code point alias table.**
    - **Create new equivalent IPv6 classifiers**
      - **Same PHB for the equivalent traffic classes**
    - **New IPv6 classifier needs to be applied to the customer interface.**
  - **MF Classification**
    - **Create new equivalent IPv6 filter or filter-policer**
    - **Apply inbound to interface**

L3 MPLS VPN Network

Configuration- Router Security

- **Core uses 6vPE**
  - **No global IPv6 loopbacks**
  - **The control plane rides on top of IPv4**
  - **No new IPv6 loopback filter required**
- **Edge**
  - **Per customer VRF loopbacks**
  - **New IPv6 filter required.**
    - **This filter should consider also protocols such as OSPFv3 and RIPng, VRRP**
  - **Simple packet filters also used**
  - **RPF check for IPv6 works the same way**

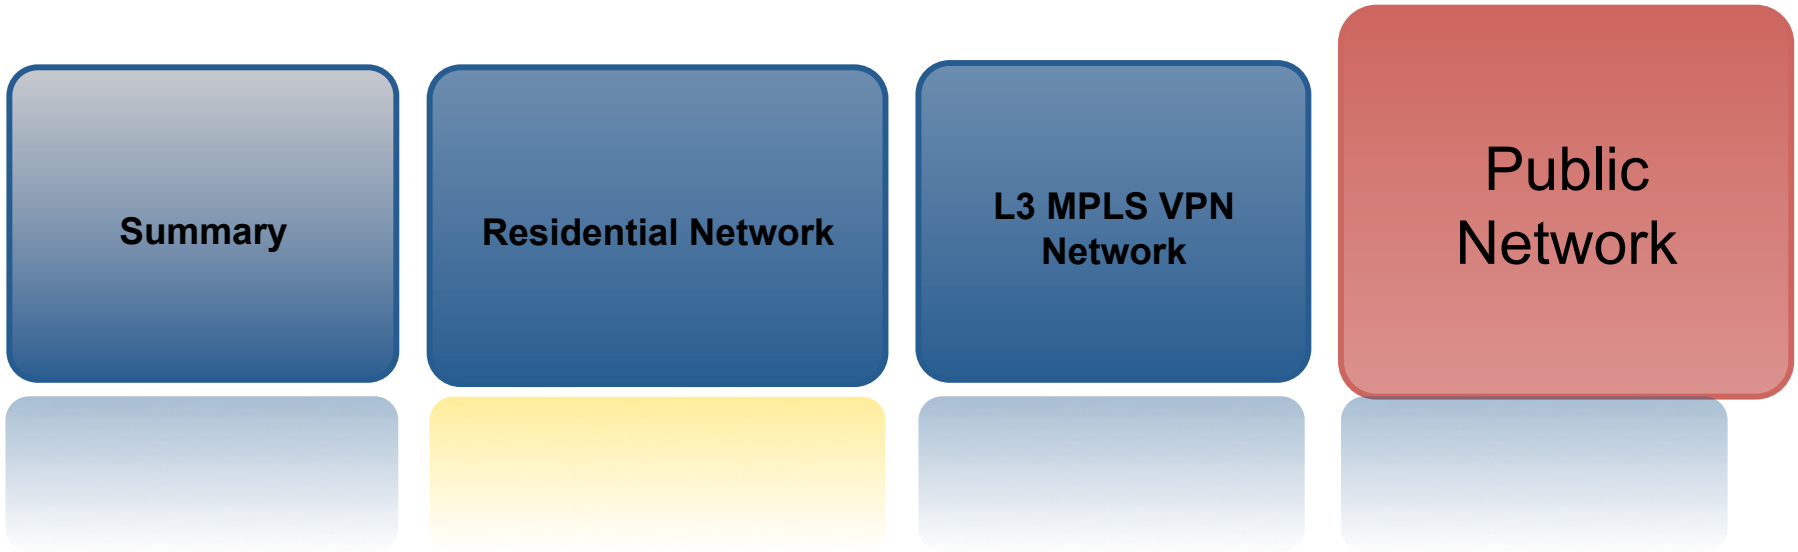
Summary
Residential Network
L3 MPLS VPN Network
Public Network

Public Network

Configuration- Architecture

- **Apply native IPv6 BGP within the core of the public network**
- **ISIS carries both IPv4 and IPv6 prefixes**

Model maintains two separate and parallel IP logical infrastructures

Native BGP over IPv6 end points using TCPv6

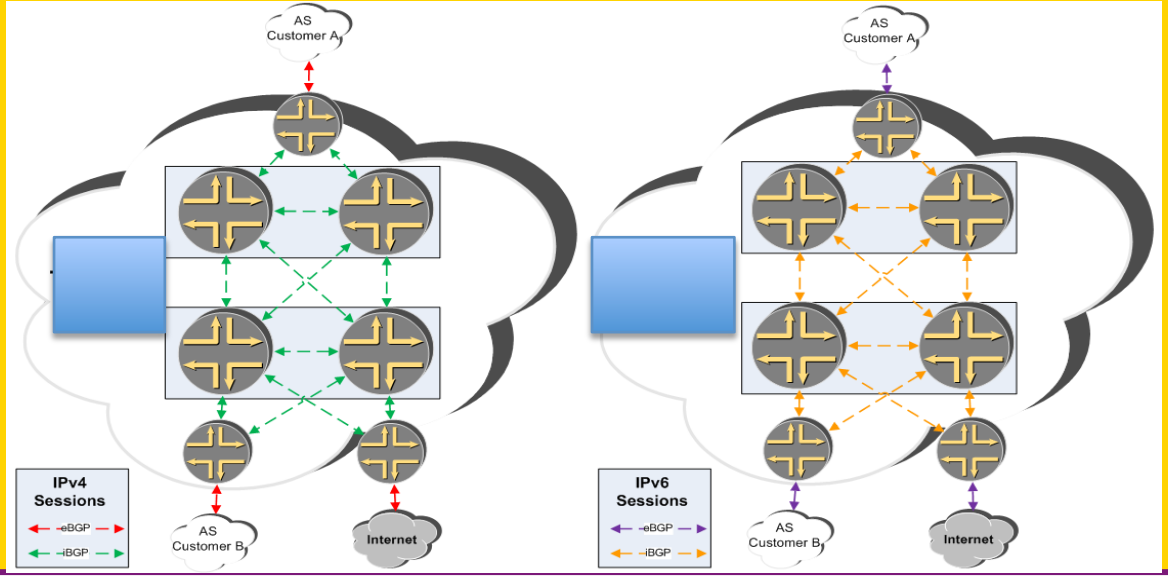

Public Network

Configuration- Routing (ISIS/BGP)

- **ISIS**
  - **Cisco and Juniper have different default behavior**
    - **JUNOS needs no additional configurations to carry IPv6 routes**
- **IBGP**
  - **Configured within a new "IPv6 Specific" peer-group**
    - **The same IPv4 export policies (route-maps) are referenced from the new peer-group**
    - **Export policies (route-maps) for BGP can be re-used if there is no specific reference to IPv6 addressing**
      - **Example: next-hop-self**
- **EBGP can be configured in two ways again:**
  - **Same IPv4 BGP session carrying two address families**
  - **Our Choice: Two separate BGP sessions, one for IPv4 and one for IPv6**
    - **Consistent with the core model**

Public Networks

Configuration- Filters and QOS

- **QOS changes are similar to the those discussed for MPLS network**
  - **One addition: DSCP Re-write to reset customer DSCP settings**
- **Filters are used in every network in the project**
  - **MF classifiers, policing, simple packet filters, etc**
  - **All the filters need to have IPv6 equivalents**
  - **Examples:**
    - **Simple filter used primarily as a security tool**
      - **Meant to deny any illegal addresses from entering the network**
    - **Filter used in order to enforce policing/SLAs**

Public Networks

Configuration- Router Security

- **Core**
  - **New global IPv6 loopback**
  - **Needs a new IPv6 filter**
    - **Move many of the filter terms from IPv4 to IPv6**
      - **BGP, SSH, etc**
    - **As compared to per VRF filters**
      - **No OSPFv3, RIPng, etc but may need other protocols such as NTP**

- **Edge**
  - **RPF Check**
  - **Customers have filters that deny packets to illegal/internal addresses**
    - **New IPv6 filters are defined**

What did we learn?

Dual Stack is only the first step

- Does not really help with IPv4 exhaustion
  - Customers want to know what is next
  - Road is not very clear
- Our Client believes CGN is part of the puzzle
  - Our BRAS architecture had to accommodate future CGN plans
  - Independent routing-domains for subscribers with public and private addresses
  - But how do we make that determination?
    - RADIUS
    - Provisioning Systems
  - May involve moving customers from one virtual router (routing-domain) to another one after authentication
- Customer is evaluating various CGN, DS-Lite solutions

What did we learn?

Service definition is key

- There is often no central list of current IPv4 services
  - At least not a dependable one!
- First, generate an inventory of current IPv4 services
  - This is more time consuming that is sounds
  - Not every SP is equal
- Second, decide what should (or is worth) moving to IPv6
  - May find items that are official IPv4 offering but have almost no customers!
    - Our example: OSPF routing for PE-CE (mplsvpn)
- Then, finally a technical question: Does the model need a new architecture for IPv6?
  - Example: BRAS model of fixed IP address customers being able to log in from any BRAS device
    - Requires all the access-internal subscribers routes (/32 for IPv4 and /128 for IPv6) to float in all the devices
    - ~Doubling the number of routes might not be an option

What did we learn?

The Devil is in the Implementation

- **Scaling, Scaling, Scaling**
  - The single most important issue we dealt with
  - Simple: Know what your device can or can not do
    - Don't be caught off guard
    - Don't' assume that a linear scaling behavior as you move to larger IP addresses, etc!
      - Number of routes in the Control Plane
      - Number of routes in the Data Plane
      - Number of Dual-Stack interfaces
      - Number of DHCPv6 leases
      - On and on…
- **Don't assume features work equally for IPv4 and IPv6**
    - Test carefully: Trust but verify!
    - Our experience:
      - Data plane is where most culprits are!

What did we learn?

IPv6 deployment touches everything

- **Migration to IPv6 requires organizational commitment**
  - **More than just a technical issue**
- **It crosses lines of various organizations**
  - **Provisioning Systems and Order Work Flow**
  - **Billing Systems**
  - **Marketing of the new Service**
  - **Peering Agreements**
  - **Etc.**
- **For a large scale deployment, touching many networks professional project management is very helpful**
- **What drove our project was an actual sense of urgency in all levels of the organization**

Thank You

Amir Tabdili
amir@unisonip.com