Case Study

A Service Provider's Road to IPv6

Amir Tabdili – UnisonIP Consulting
amir@unisonip.com

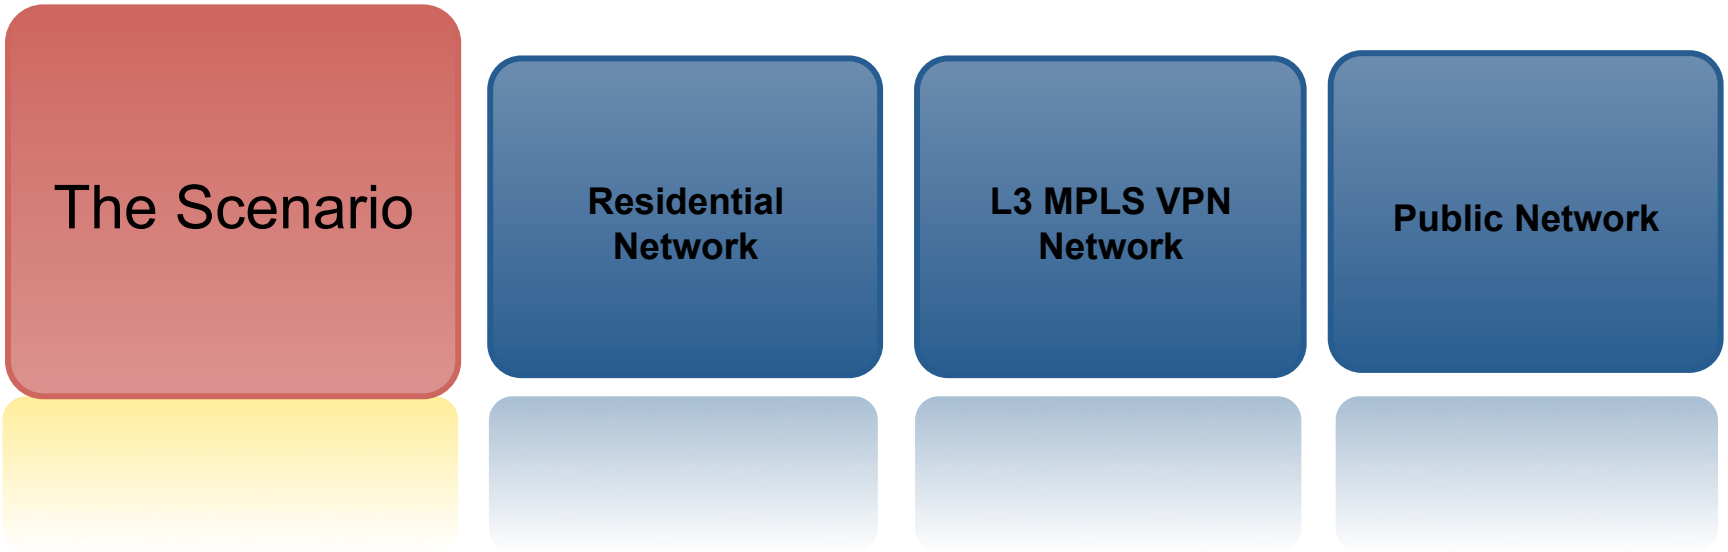
The Scenario
Residential Network
L3 MPLS VPN Network
Public Network

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

- **Many multivendor networks**
- **Not just public Internet**
- **Inter-dependency of services across networks**

The Scenario

Background and Assumptions

- **Deployment Scenario is "Dual-Stack"**
  - **Considered first step**
- **For now, only convert what is visible externally**
- **The rollout should be transparent to existing customers**
- **Current IPv4 transport models**
  - **IPv4 only for Public network**
  - **MPLS for the L3 MPLS VPN network**

The Scenario

Background and Assumptions

- **Devices Deployed**
  - **E320/ERX Juniper routers**
    - **As LAC and LNS**
  - **Juniper T /M**
    - **As P/PE MPLS VPN and Internet routers**
  - **Cisco routers**
    - **As Internet routers**
  - **Alcatel**
    - **As VPLS switches, transparent to IPv6**

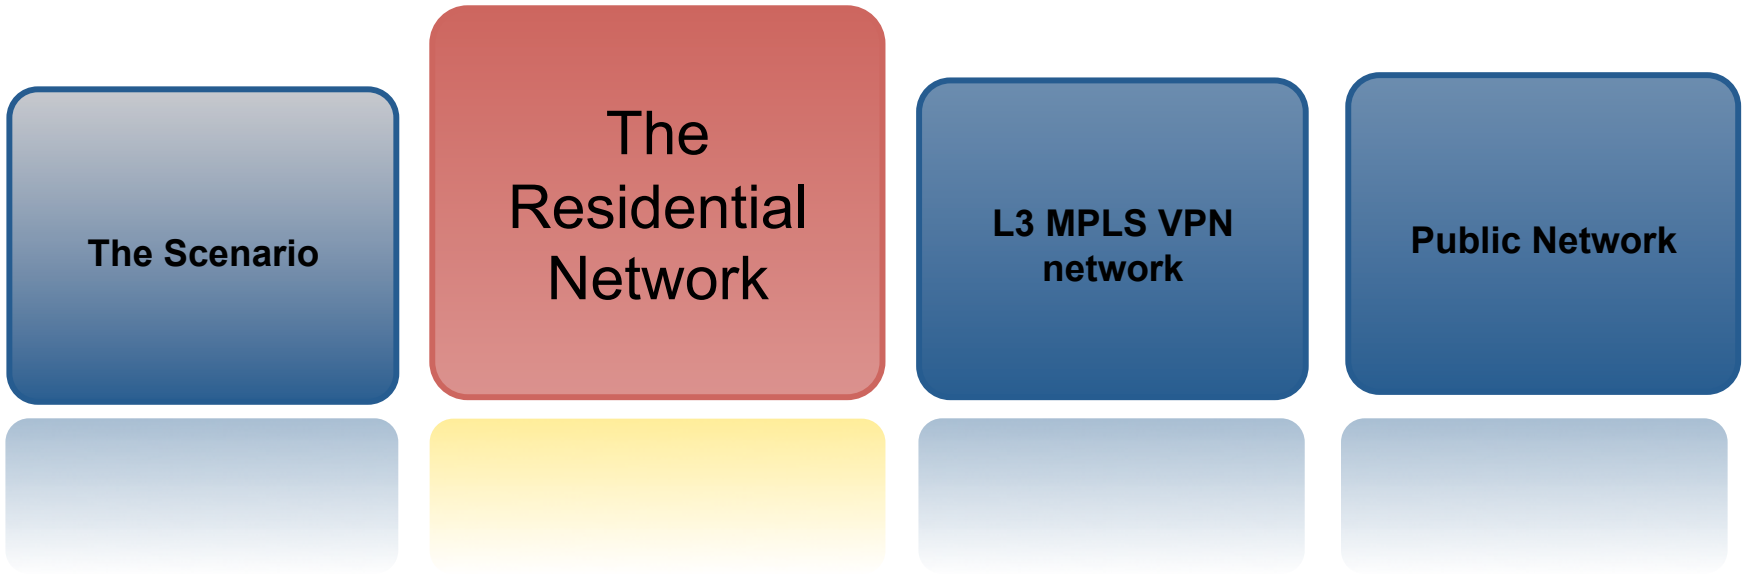
The Scenario
The Residential Network
L3 MPLS VPN network
Public Network

Residential Network

Main Connectivity Models Supported by IPv6

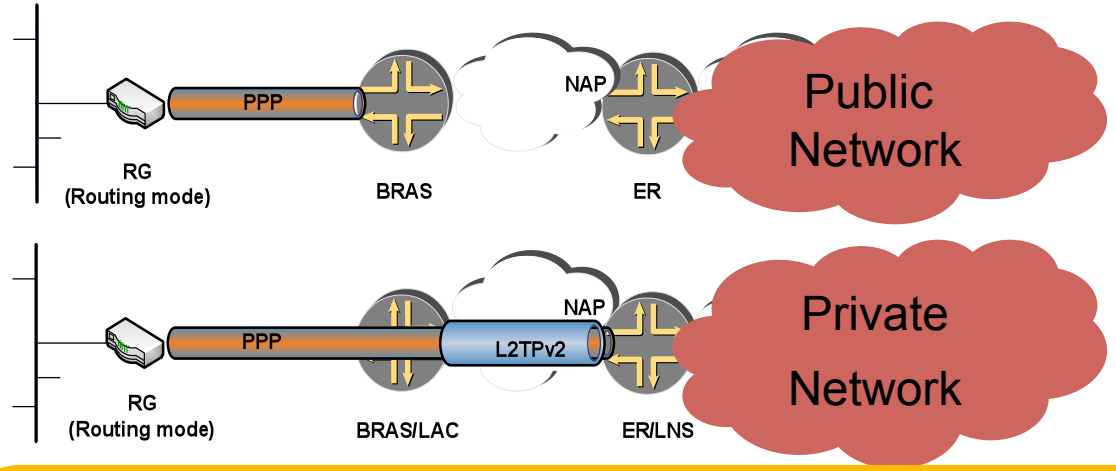

- **Simple Routed Mode**
  - **CPE establishes a PPP session to the LAC**
  - **Support for L3 Backhaul**

- **L2TP Backhaul**
  - **LAC has L2TP session to an LNS**
  - **Enterprise service: Subscriber terminates in MPLS VPN**
  - **L2 Wholesale: LNS is owned by another ISP**

Residential Network

Our Choices for PPP Model– Dual Stack or Dual Session

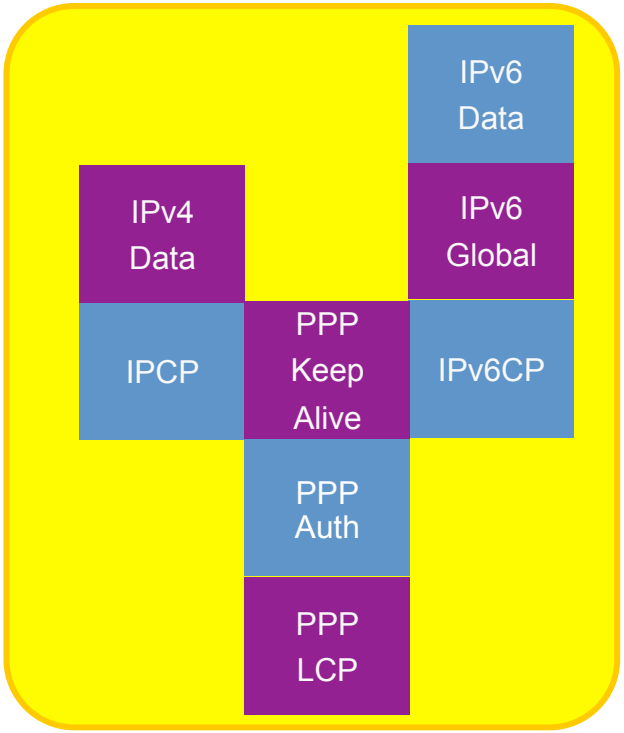

- **One AAA interaction**
- **Most flexible**
- **CPE Driven**

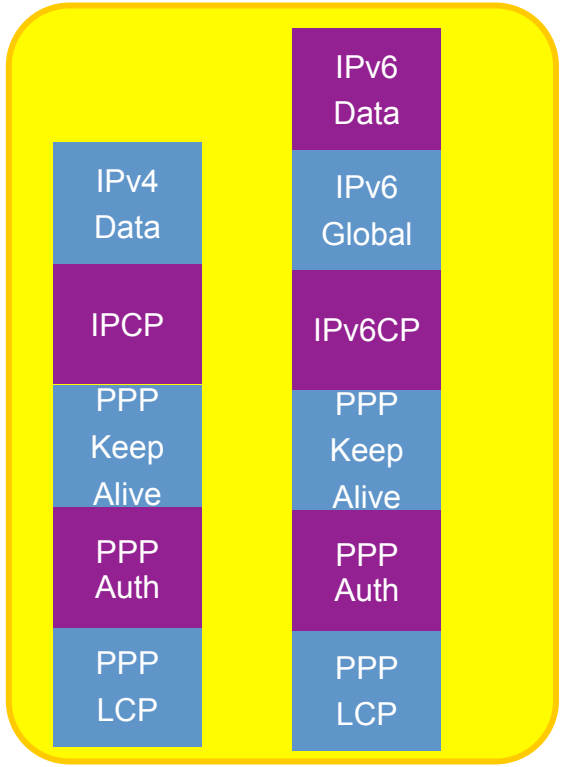

- **Can be interesting for transition:**
  - **IPv6 LNS?**

Residential Network

E320/ERX IPv6 Configuration – Getting Started

- **Many business questions**
  - **What are the offerings?**
  - **Do all IPv4 services make sense for IPv6?**
  - **Which customers should get an IPv6?**
    - **All or just new?**

- **Everything is influenced by scaling!**
  - ***Scaling drives many of the design decisions***

- **License**
  - **IPv6 in JUNOSe required activation of a license key**

Residential Network

E320/ERX IPv6 Configuration – Interfaces and Routing

- **Interfaces**
  - **Need to have new IPv6 Addresses**
  - **Many choices for the subnetting: /126, /127, /64, …**
- **ISIS**
  - **Will also carry the IPv6 topology info**
- ***BGP***
  - ***Two Options:***
    - ***Native IPv6 end-points*** → ***Solution Picked***
    - ***IPv4 BGP carrying IPv6 NLRI***
- ***Policy***
  - **Need to have IPv6 equivalents**

Residential Network

Subscriber Addressing Model

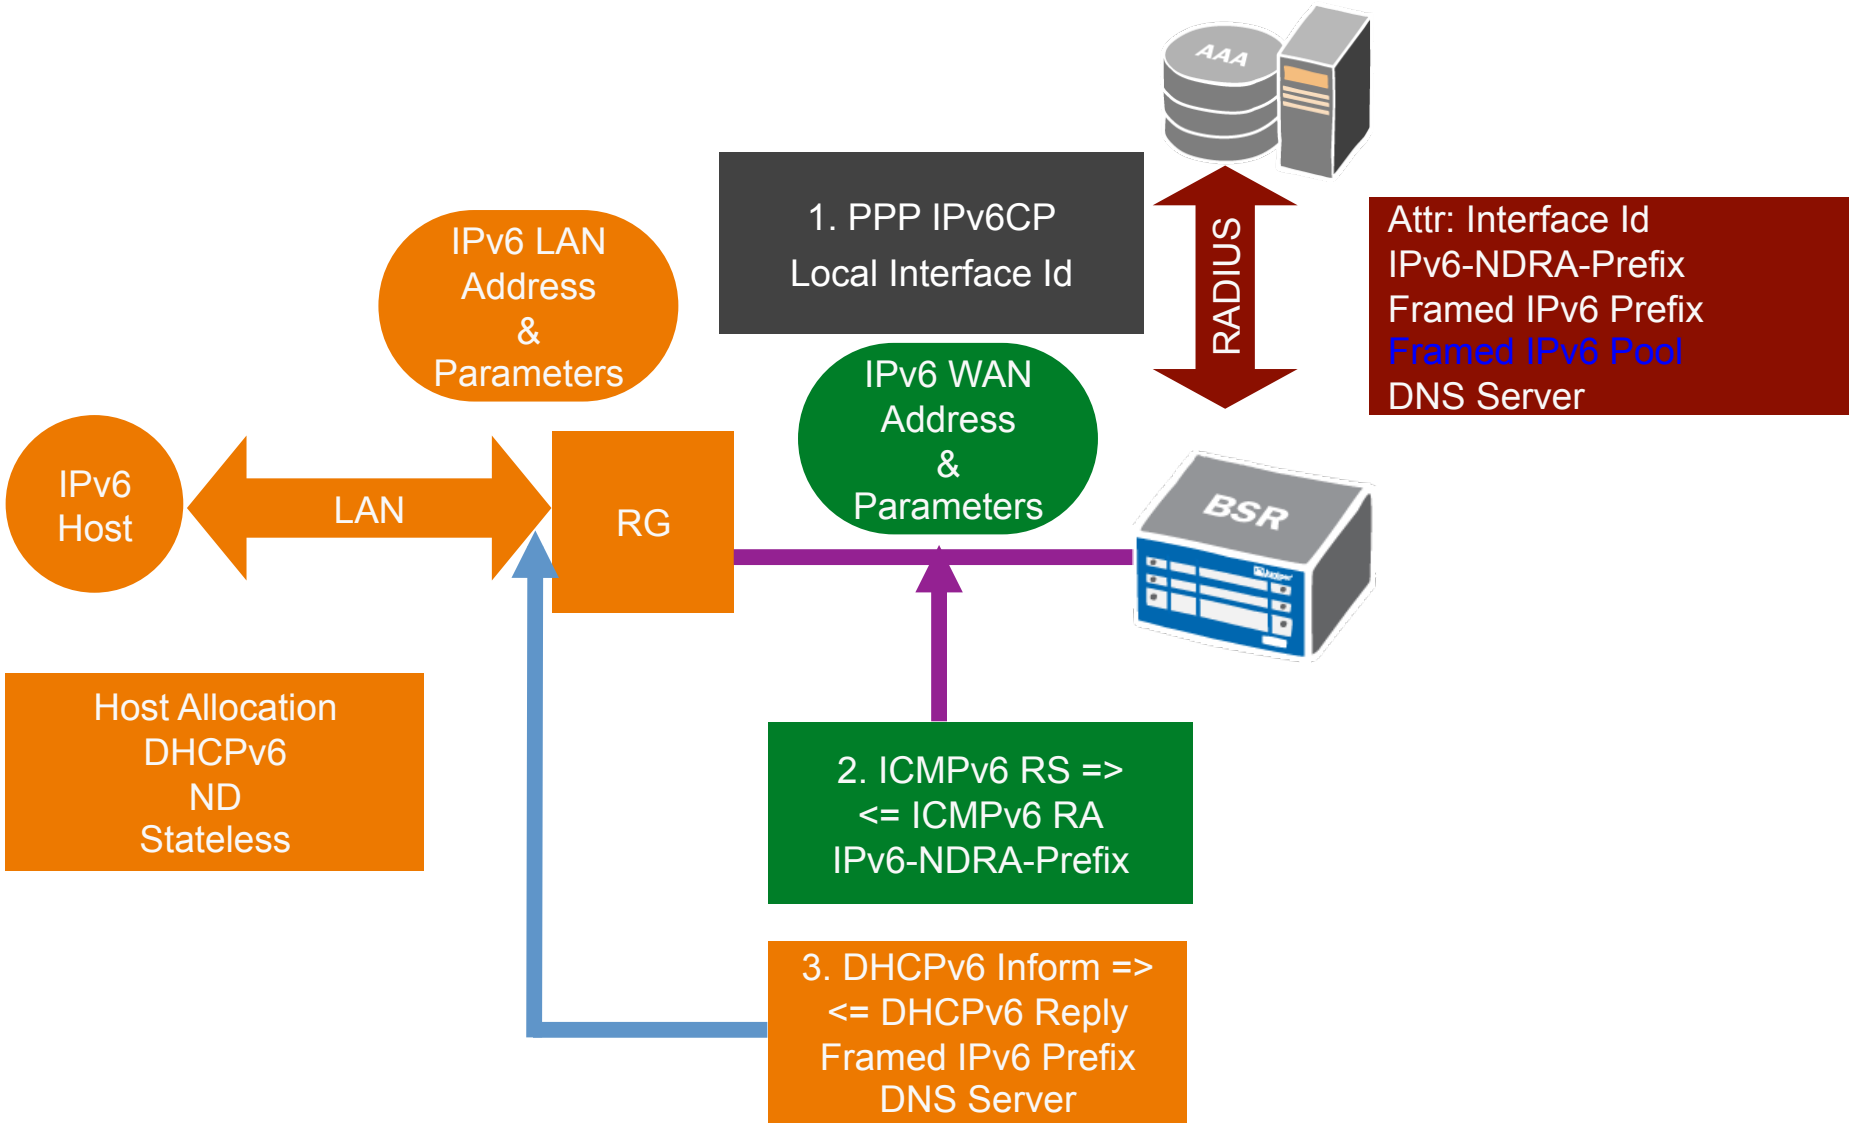

Residential Network

E320/ERX IPv6 Configuration – Address Assignment/Delegation

- **Bigger addresses – /32 => /(64 + 64)**
- **We may get up to 3 addresses**
  - **Link Local**
    - **Is always there!**
  - **ICMPv6 (ND) – CPE WAN SIDE**
    - **IPv6-NDRA-Prefix**
    - ***Our choice: Static, configured in the profile***
  - **DHCPv6 (PD)– CPE LAN SIDE**
    - **Framed-IPv6-Prefix or Framed-IPv6-Pool**
    - **Our choice: Static, assigned in the domain-map**

Residential Network

E320/ERX IPv6 Configuration – Subscriber Interfaces and Profiles

- **Subscriber interfaces**
  - **No new IPv6 specific configuration**
  - **Dynamic PPPoE / PPP stack**
- **Changes contained in profiles**
  - **Attached to the subscriber interfaces at provisioning**
- **Profile modifications**
  - **The IPv6 loopback interface**
  - **Neighbor Discovery (ND)**
  - **IPv6 Policy**
  - **RPF check for IPv6 source validation**
  - **Other optional configurations**

Residential Network

E320 IP Configuration – DNS Servers

- **Specific to ERX/E320 Platforms**
- **DNS server 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**
  - **Need to override the static DNS assignment by RADIUS**

Residential Network

E320/ERX IPv6 Configuration – Accounting and Counters

- **Radius Accounting**
  - **Accounting of IPv6 services is equivalent to IPv4**
  - **All IPv6 attributes included in the Access-Accept from RADIUS can be included**
- **Counters**
  - **Access to the PPP session counters**
    - **IPv4 and IPv6 run on top of one PPP session**
    - **Session counters include the IPv4 and IPv6 packets**
  - **Separate counters for IPv6 are also supported**

Residential Network

LNS Specific Configuration - Highlights

- **Many concepts discussed for LAC also apply to the LNS**
  - **The L2TP tunnel end points stay on IPv4**
- **A few tweaks needed. Examples:**
  - **Some configurations applied to customer VRFs**
    - **All the address-assignment configurations**
  - **Communication with rest of L3VPN MPLS network**
    - **Need to turn on vpnv6 BGP address family**
    - **May be Service Interrupting**

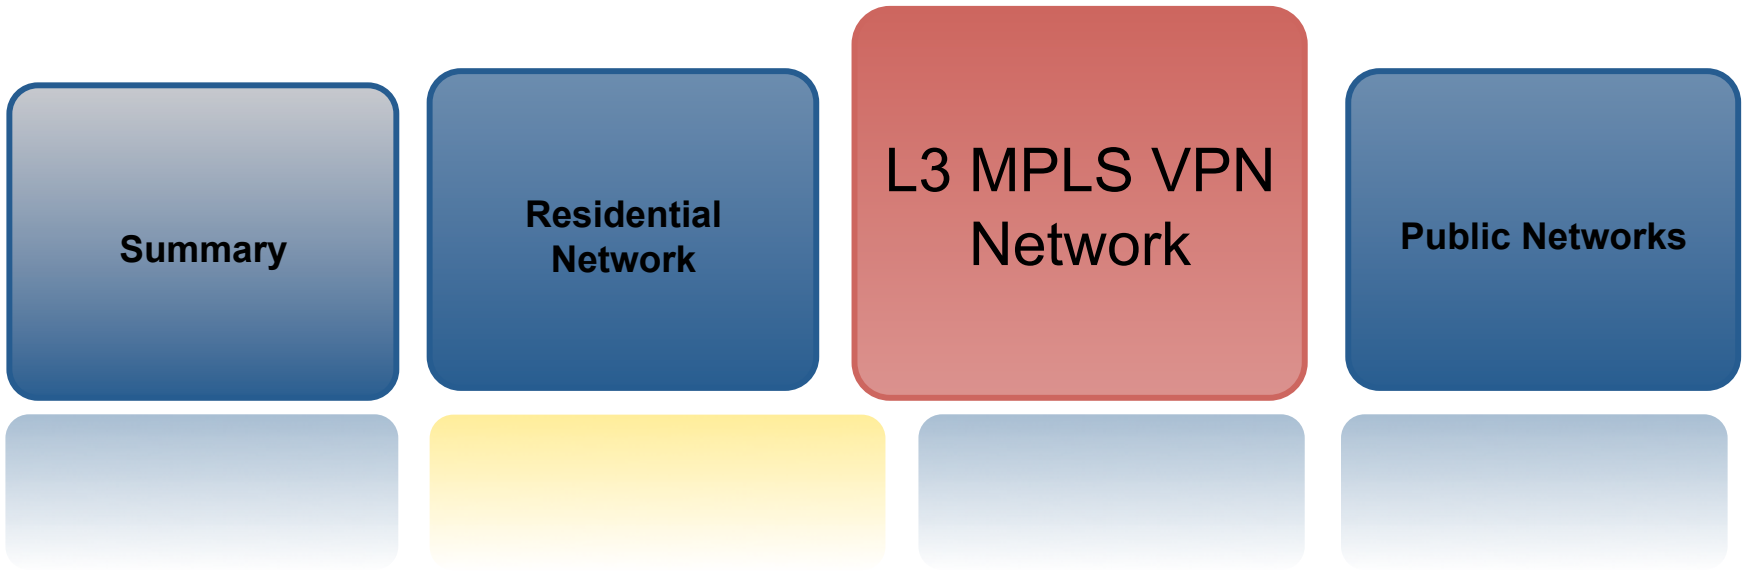
Summary
Residential Network
L3 MPLS VPN Network
Public Networks

L3 MPLS VPN Network

IPv6 Transport Architecture

mplsvpn network:

6VPE Architecture

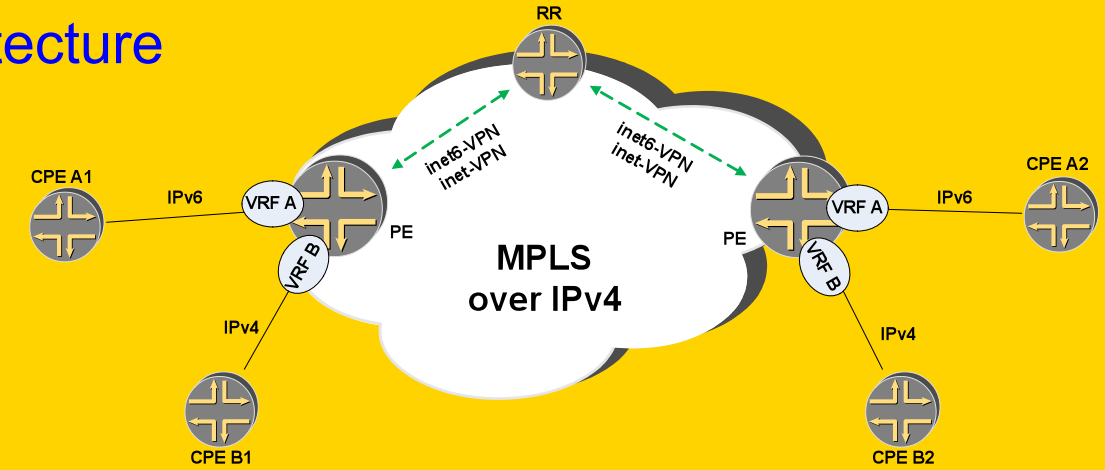

- **Relatively straight-forward**
  - **Operational model and configuration similar to IPv4 VPN**
  - **Can use same LSPs and same BGP sessions**
  - **Same features as for IPv4 VPN can be used**

L3 MPLS VPN Network

Configuration- Core

- **MPLS**
  - **MPLS used in the core for forwarding**
  - **In Junos, All routers need "ipv6-tunneling" under [protocols mpls]**
    - **So that IPv6 routes are resolved over the LSP tunnels.**
- **BGP**
  - **BGP needs modifications in the core**
    - **family inet6-vpn added**
  - **Interprovider VPNs**
    - **Same configuration to be followed for Option C peers**

L3 MPLS VPN Network

Configuration- PE-CE Routing

- **EBGP, Static and RIP-NG**
- **EBGP**
  - **Chose to configure using a new peer-groups**
- **Static**
  - **Very simple**
    - **Follows the IPv4 model and syntax**
- **RIP-NG**
  - **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**
  - **MF Classification**
    - **Create new equivalent IPv6 filter or filter-policer**
- **Watch for hardware restrictions!**

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.**
    - **Considers also OSPFv3 and RIPng, VRRP**
  - **Simple packet filters also used**
  - **RPF check for IPv6 works the same way as IPv4**

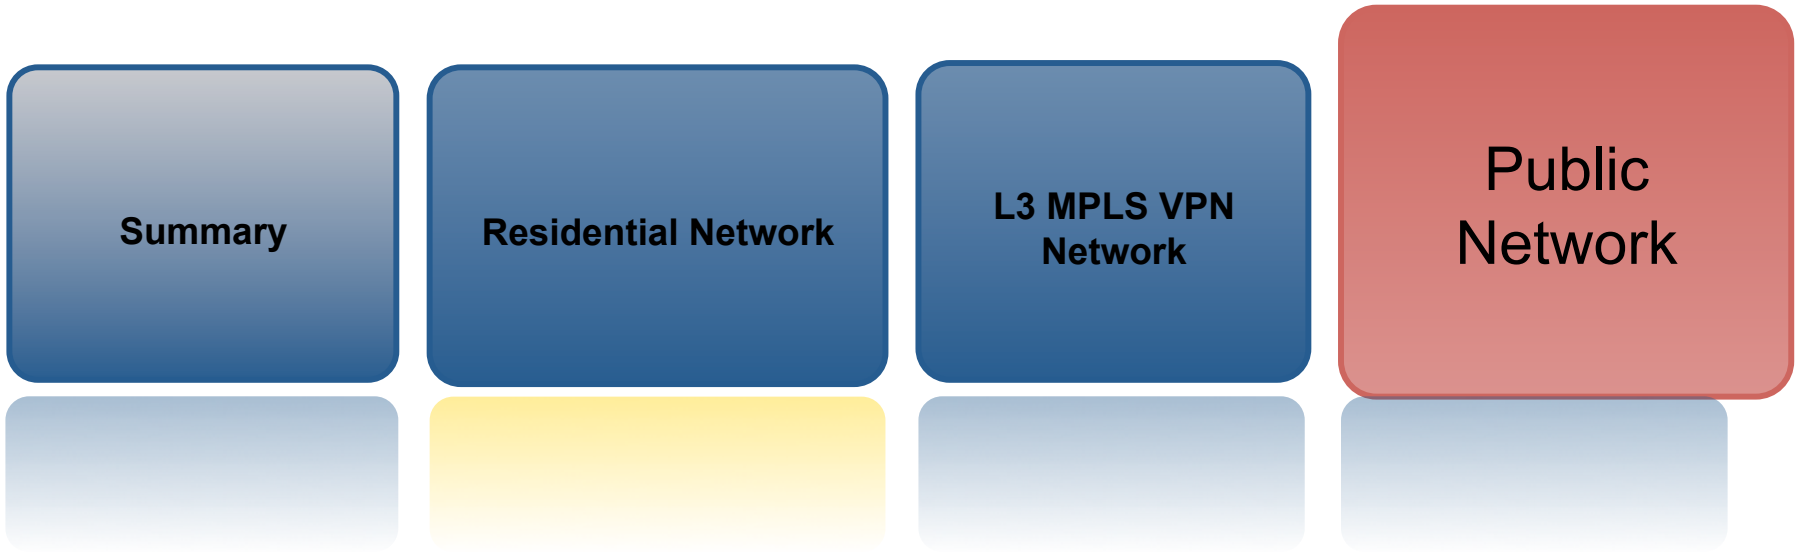
Summary
Residential Network
L3 MPLS VPN Network
Public Network

Public Network

Configuration- Architecture

- **Native IPv6 forwarding in the public network**
  - **BGP over IPv6 end points riding TCPv6**
  - **ISIS carries both IPv4 and IPv6 prefixes**
  - **Shared fate only on the IGP side**

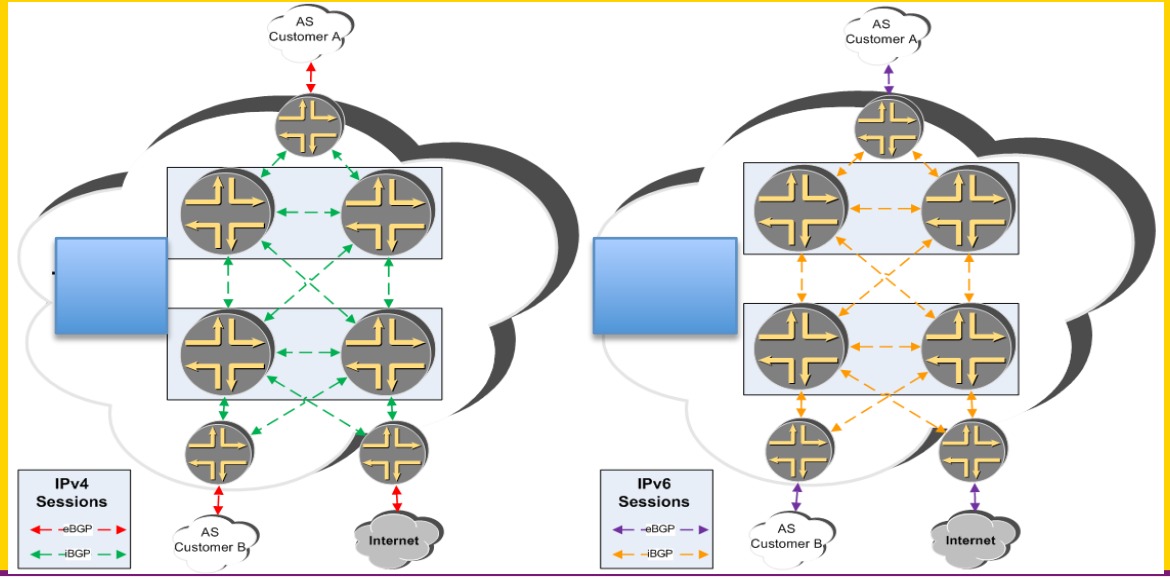

Public Network

Configuration- Routing (ISIS/BGP)

- **ISIS**
  - **JUNOS needs no additional configurations to carry IPv6 routes**
- **IBGP**
  - **Configured within a new “IPv6 Specific” peer-group**
  - **Export policies can be re-used if there is no specific reference to IPv6 addressing**
- **EBGP**
  - **Two separate BGP sessions, one for IPv4 and one for IPv6**
  - **Consistent with the core IBGP model**

Public Networks

Configuration- Filters and QOS

- **Changes similar to the those discussed for MPLS network**
  - **One addition: DSCP Re-write to reset customer DSCP settings**
  - **Created the equivalent for DSCPv6**
- **Filters are used in every network in the project**
  - **MF classifiers, policing, simple packet filters, etc**
  - **All the filters need to have IPv6 equivalents**

Public Networks

Configuration- Router Security

- **Core**
  - **New global IPv6 loopback needs a new IPv6 filter**
    - **Create equivalent terms from IPv4 to IPv6**
      - **BGP, SSH, etc**
    - **As compared to per VRF filters discussed earlier**
      - **No OSPFv3, RIPng, etc**
      - **But may need other protocols such as NTP**
- **Edge**
  - **RPF Check**
  - **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
- Our Client believes CGN is part of the puzzle
  - Had to accommodate future CGN plans
  - Independent routing-domains
    - Used to separate public and private address subscribers
  - May involve moving customers from one virtual router (routing-domain) to another after authentication
- Customer is evaluating various solutions
  - CGN, DS-Lite, others.

What did we learn?

Service definition is key

- First
  - Generate an inventory of current IPv4 services
  - This is more time consuming that is sounds
  - Often no central list of current IPv4 services exists
- Second,
  - Decide what should (or is worth) move to IPv6
- 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
      - Leaking all IPv6 access routes may not be an option!

What did we learn?

The Devil is in the Implementation

- **Scaling, Scaling, Scaling**
  - The single most important issue
  - Simple: Know your device
    - Don't' assume a linear scaling behavior
      - 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**
  - **Billing Systems**
  - **Marketing**
  - **Peering Agreements**
  - **On and on …**
- **What made our project successful**
  - **“Sense of urgency in all levels of the organization”**

Thank You

Amir Tabdili
amir@unisonip.com