IPv6 Transport over IPv4 Technologies and Testing

Steve Jarman

Business Development Manager - EMEA

Agenda

- IPv6 over IPv4 Transport Mechanisms
  - Tunneling (6RD)
  - 6PE
  - 6VPE
- Testing Strategies

IPv6 over IPv4 Transport Mechanisms

Mechanism	Primary Use	Benefits	Limitations
IPv6 over a circuit transport over MPLS	SP with circuit to the CE (ATM, Ethernet, etc.)	Transparent to the SP	Scalability
IPv6 over IPv4 tunnels over MPLS	SP willing to offer IPv6 service on top of an existing IPv4 MPLS service	Impact limited to PE	Tunnel overhead Configuration
IPv6 MPLS with IPv4-based core (6PE/6VPE)	SP willing to offer IPv6 service on top of an existing IPv4 MPLS service	Impact limited to PE	Core is unaware of IPv6: limitations in load-balancing and troubleshouting
IPv6 MPLS with IPv6-based core	SP willing to offer MPLS services in an IPv6-only context	Full MPLS-IPv6 functionality	Impact on entire MPLS Infrastructure Complexity if coexists with an IPv4-MPLS service

SPIRENT

IPV6 OVER CIRCUIT TRANSPORT OVER MPLS

IPv6 over MPLS Tunnels

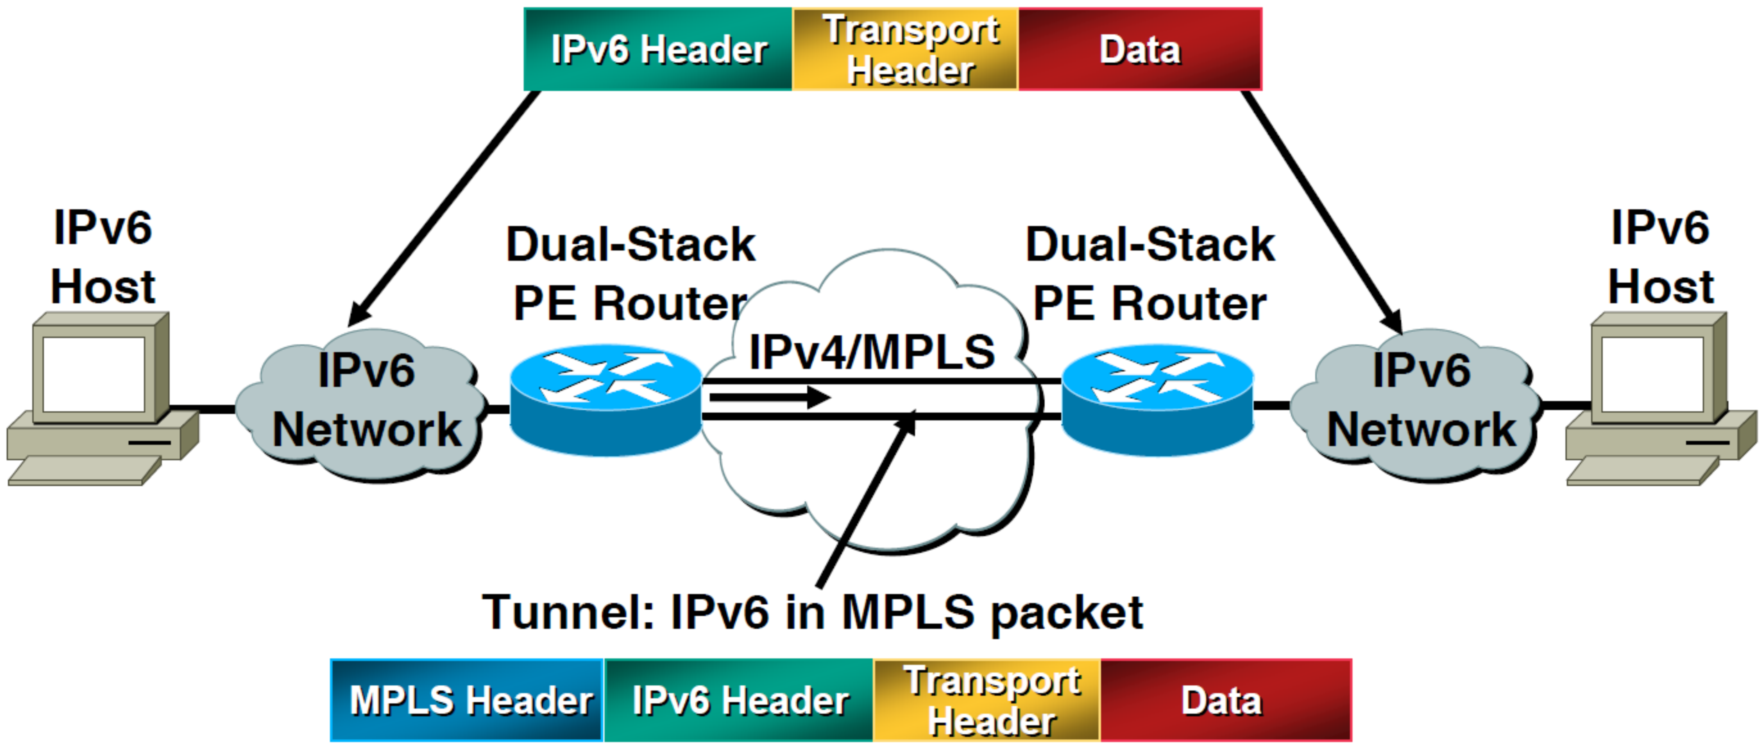

- MPLS Core remains IPv4
- Dual Protocol PE is encapsulating the IPv6 packet into MPLS packets
- Optionally PE provide VPN services for IPv6 (network virtualization)

IPV6 OVER IPV4 TUNNELS (OVER MPLS IF DESIRED)

6RD (IPv6 Rapid Deployment) / 6to4

- For providing IPv6 service over the existing IPv4-only infrastructure, especially when it is technically incapable or costly to upgrade the infrastructure to support dual-stack

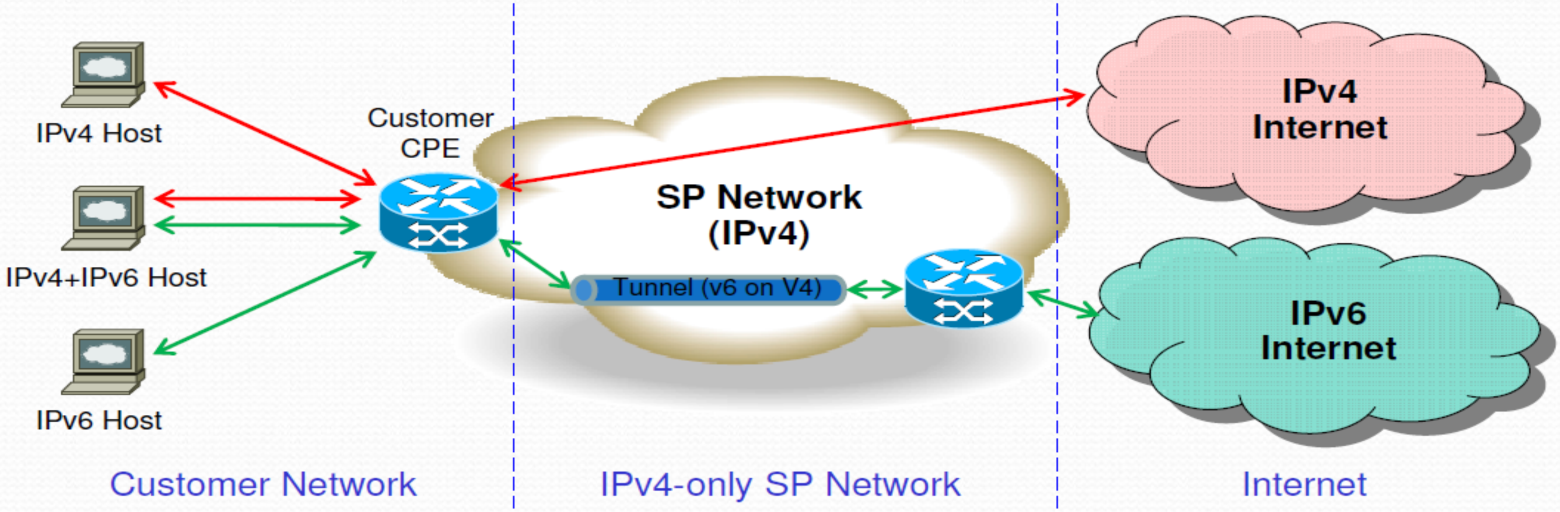

Pros
- Fast IPv6 service provisioning.
- No need to upgrade the whole infrastructure to support IPv6.

Cons
- Still need IPv4 addresses.
- Not pave way for IPv6 migration.
- Need investment on routers supporting tunneling

SPIRENT

IPv6 Rapid Deployment (6rd)

RFC-5969

- 6rd specifies a protocol to deploy IPv6 to sites via a service provider's IPv4 network.
- It builds on 6to4 with the key differentiator that it utilizes an SP's own IPv6 address prefix rather than a well-known prefix (2002::/16)
- 6rd views the IPv4 network as a link layer for IPv6

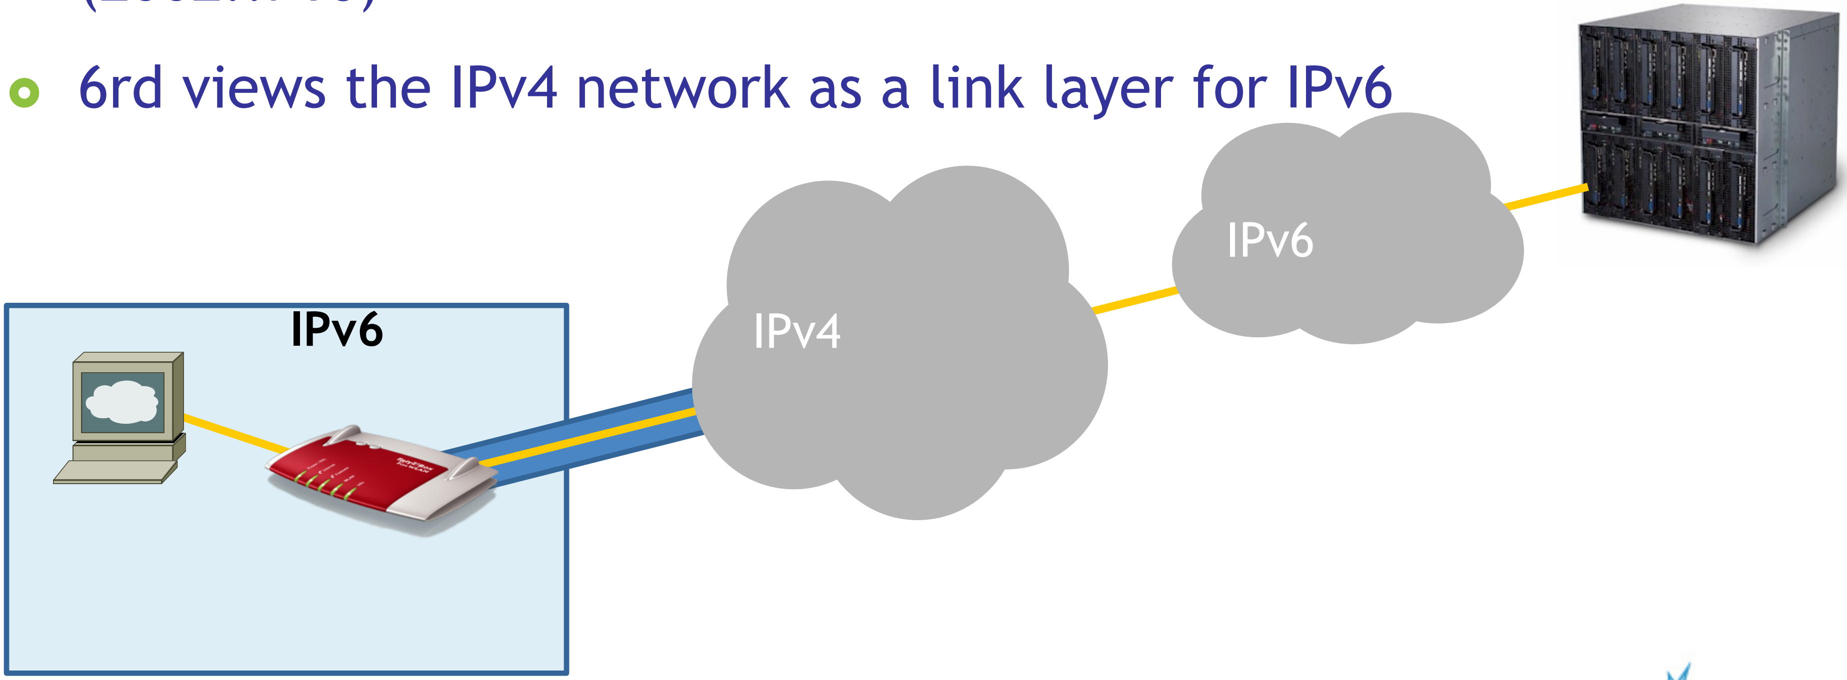

SPIRENT

6rd address structure

6rd Delegated Prefix		Customer IPv6 Address	
6rd Prefix/n bits	CE IPv4 add 0-32 bits	Subnet ID 0-16 bits	Interface ID 64 bits
2001:DB80::/32	6464:0100		

The BR & CE must be configured with the following:

- IPv4MaskLen If 10.0.0.0/8 is used as the CE address, the high order bits will be stripped before constructing the 6rd delegated prefix.
- 6rdPrefix: The 6rd prefix for the given 6rd domain.
- 6rdBRIPv4Address IPv4 address of the 6rd Border Relay for the domain.

Address Construction Example

- IPv4 Address = 10.100.100.1/8
- IPv6 Address = 2001:DB80:64 64:0100::/128

SPIRENT

6rd Example (Customer Edge Example)

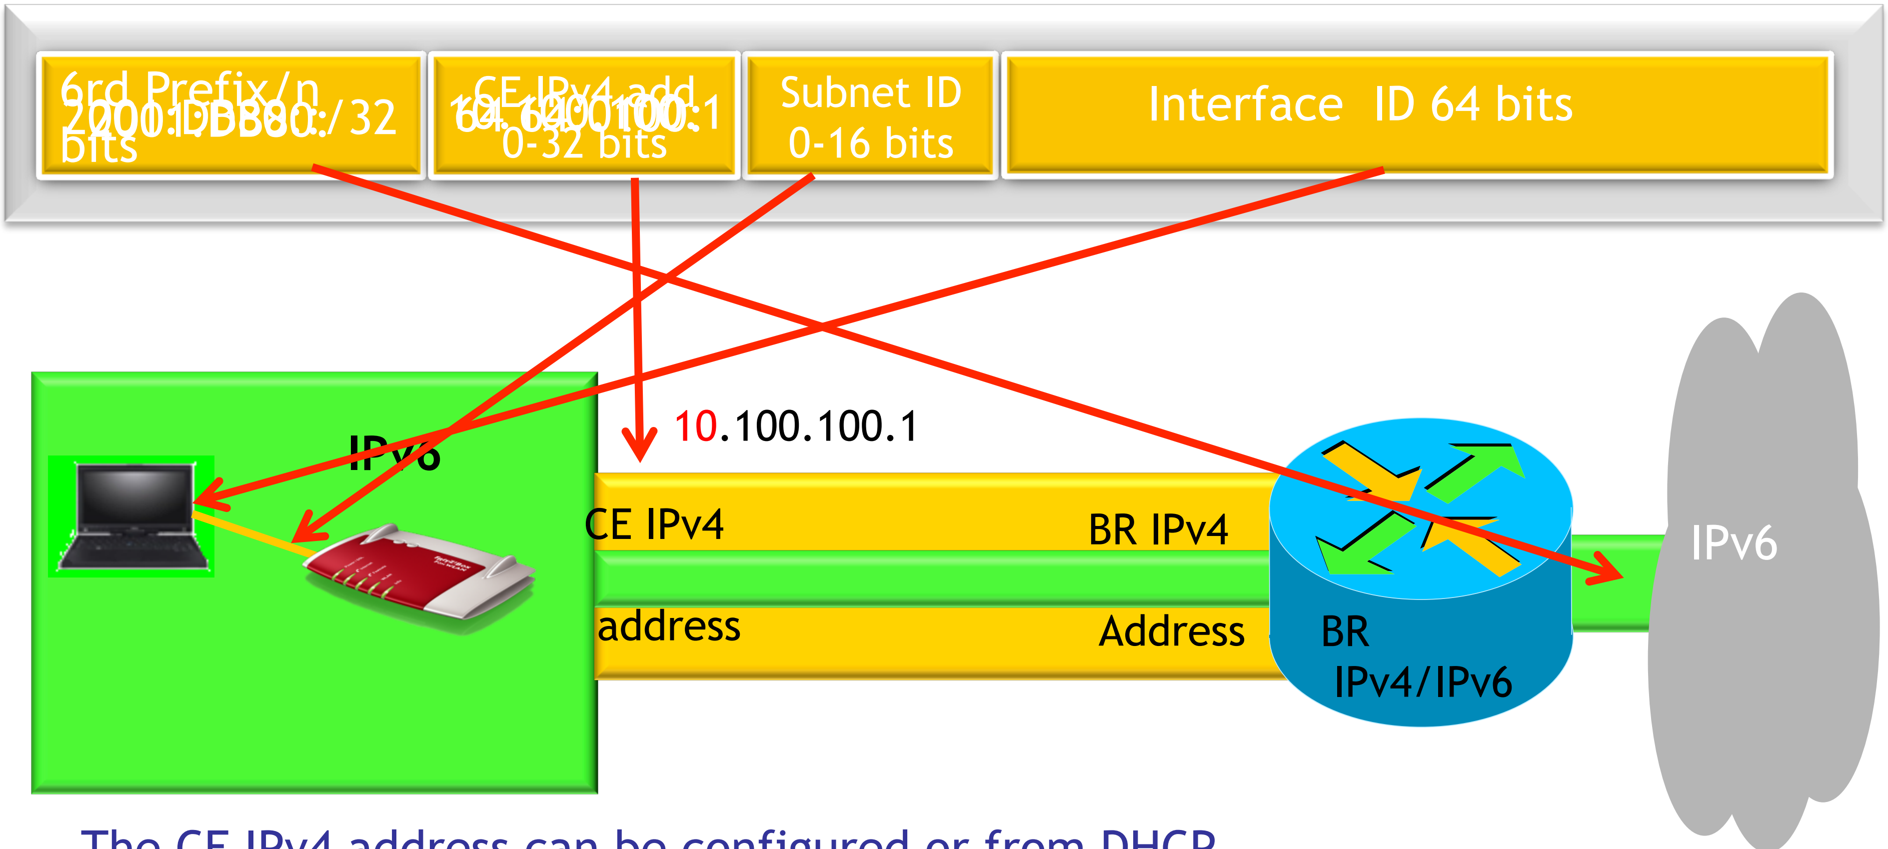

The CE IPv4 address can be configured or from DHCP
The CE IPv4 address can be global or private (RFC 1918)

6rd DHCPv4 Option

Option_6rd	Option-Length	IPv4 Mask len	6rd PrefixLen
6rd Prefix (16 Octets)			
6rdBRIPv4Address(es)			

- Option_6rd Value (212).
- Option-Length Length of DHCP Option (22 with one BR IPv4 Address).
- IPv4MaskLen Number of high order bits that are identical across all CE.
- 6rdPrefixLen Length of SP's 6rd IPv6 Prefix in number of bits.
- 6rdBRIPv4Address One or more IPv4 Address of 6rd Border Relay.

SPIRENT

Security concerns

- All of the popular IPv6 tunneling techniques for carrying IPv6 packets over IPv4 networks raise security concerns.
- IPv6 traffic runs over the IPv4 network unseen because it is disguised as IPv4 traffic.
- This exposes networks to IPv6-based attacks such as botnet command and control.
- Network operators need IPv6-aware firewalls, intrusion-detection systems and network management tools in order to have visibility into encapsulated IPv6 packets.
- BUT – What effect will that have on device and network performance?

SPIRENT

6PE - IPv6 Global Connectivity over IPv4-MPLS core

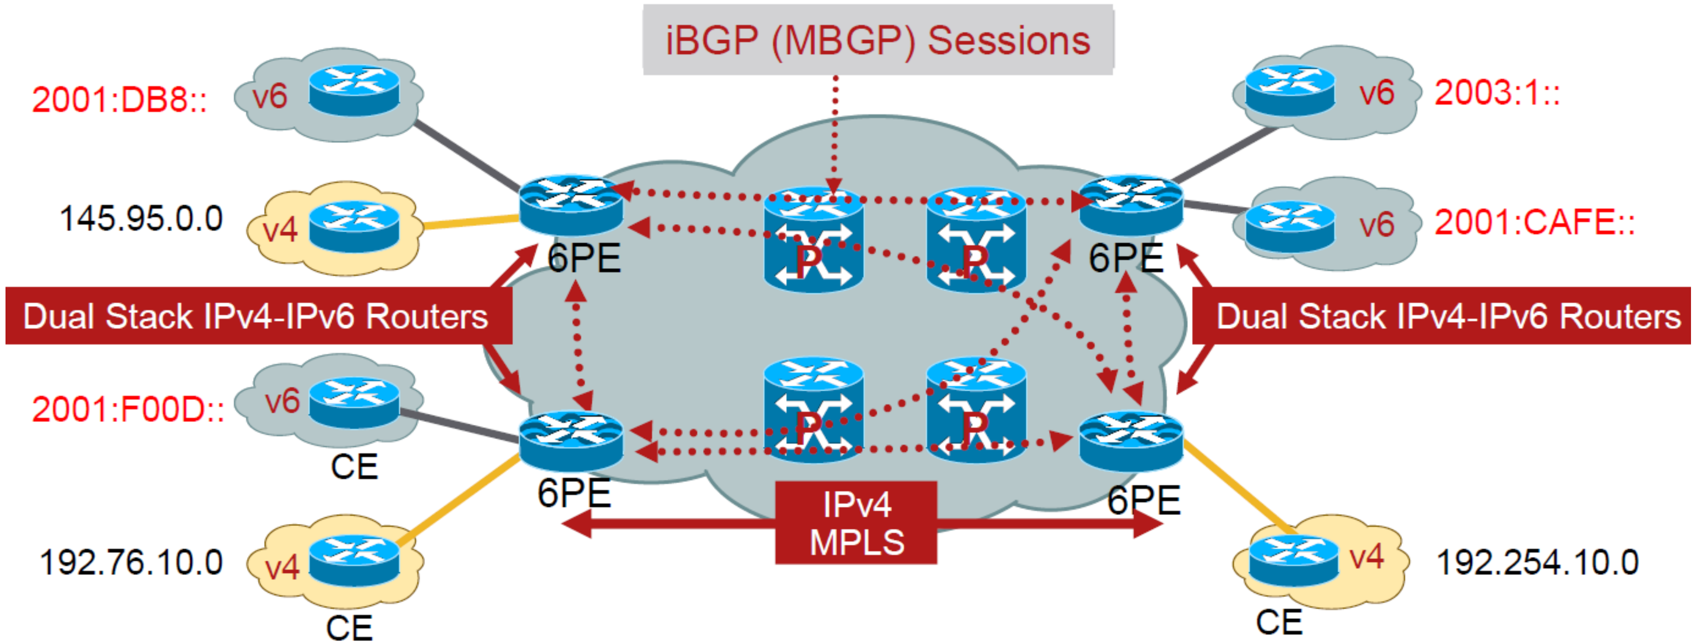

- 6PEs must support dual stack IPv4+IPv6 (6PE)
- IPv6 addresses exist in global table of PE routers only
- IPv6 reachability exchanged among 6PEs via iBGP (MP-BGP)
- IPv6 AF (2) + Label SAFI (4) used to exchange prefixes between PEs
- IPv6 packets transported from 6PE to 6PE inside MPLS (label switching)
- Core uses IPv4 control plane (LDPv4, TEv4, IGPv4, MP-BGP)

SPIRENT

6PE

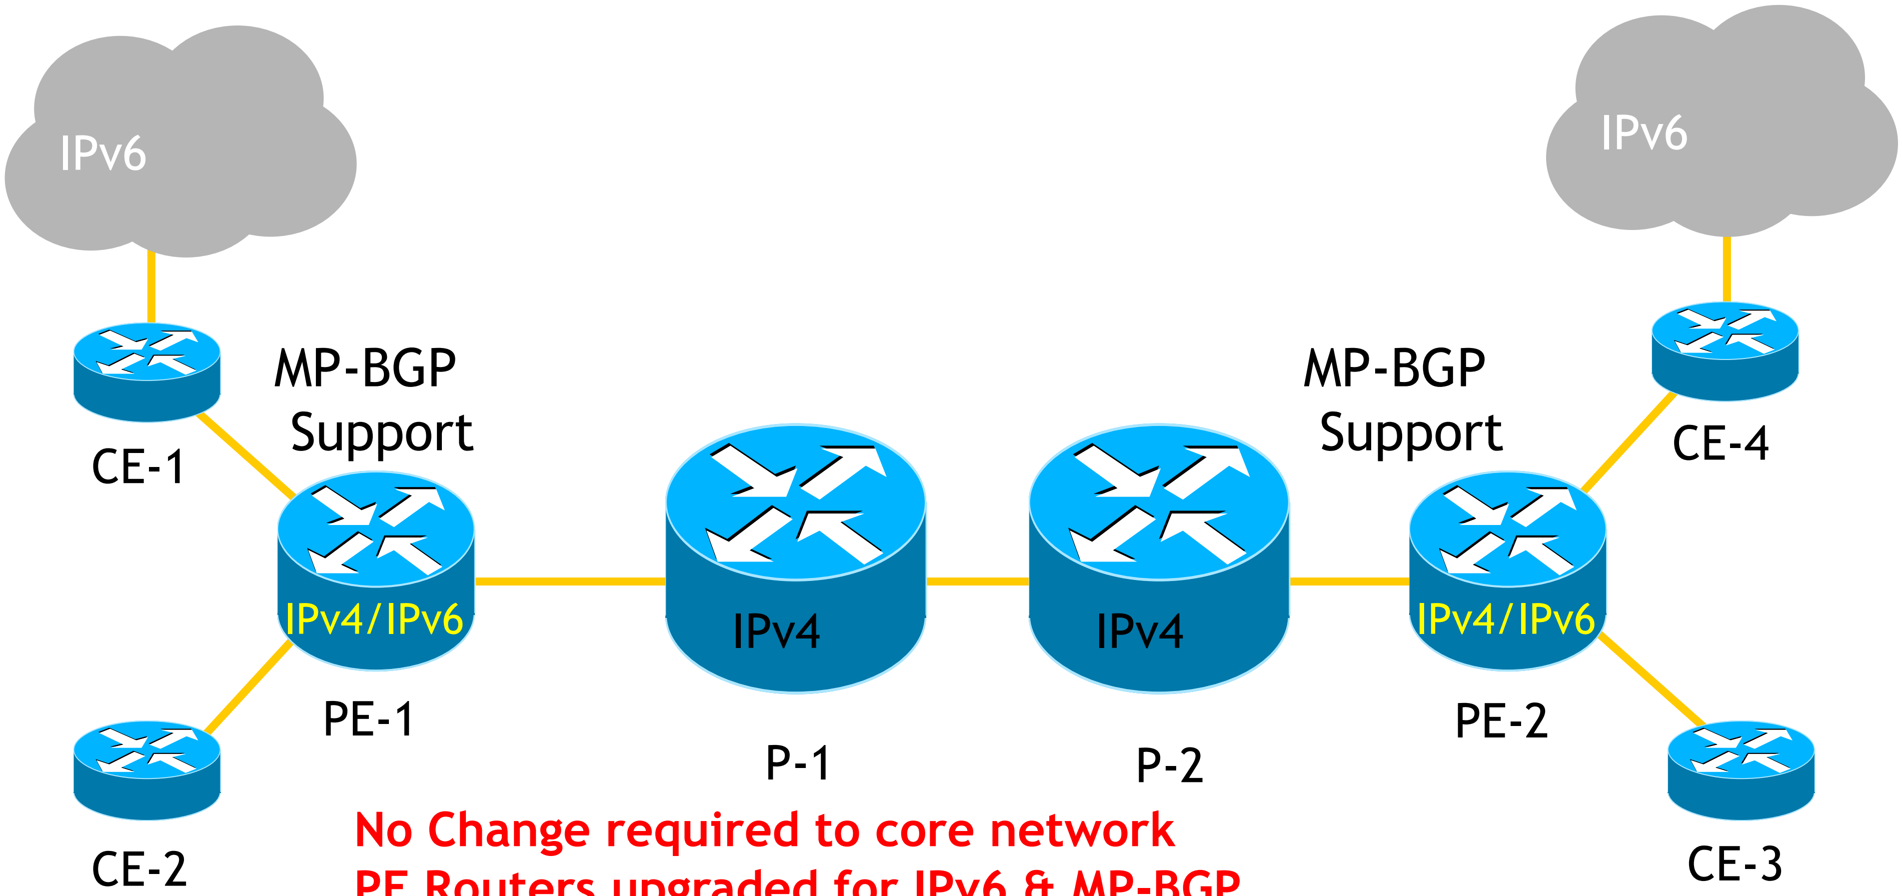

No Change required to core network
PE Routers upgraded for IPv6 & MP-BGP

Label Distribution using 6PE

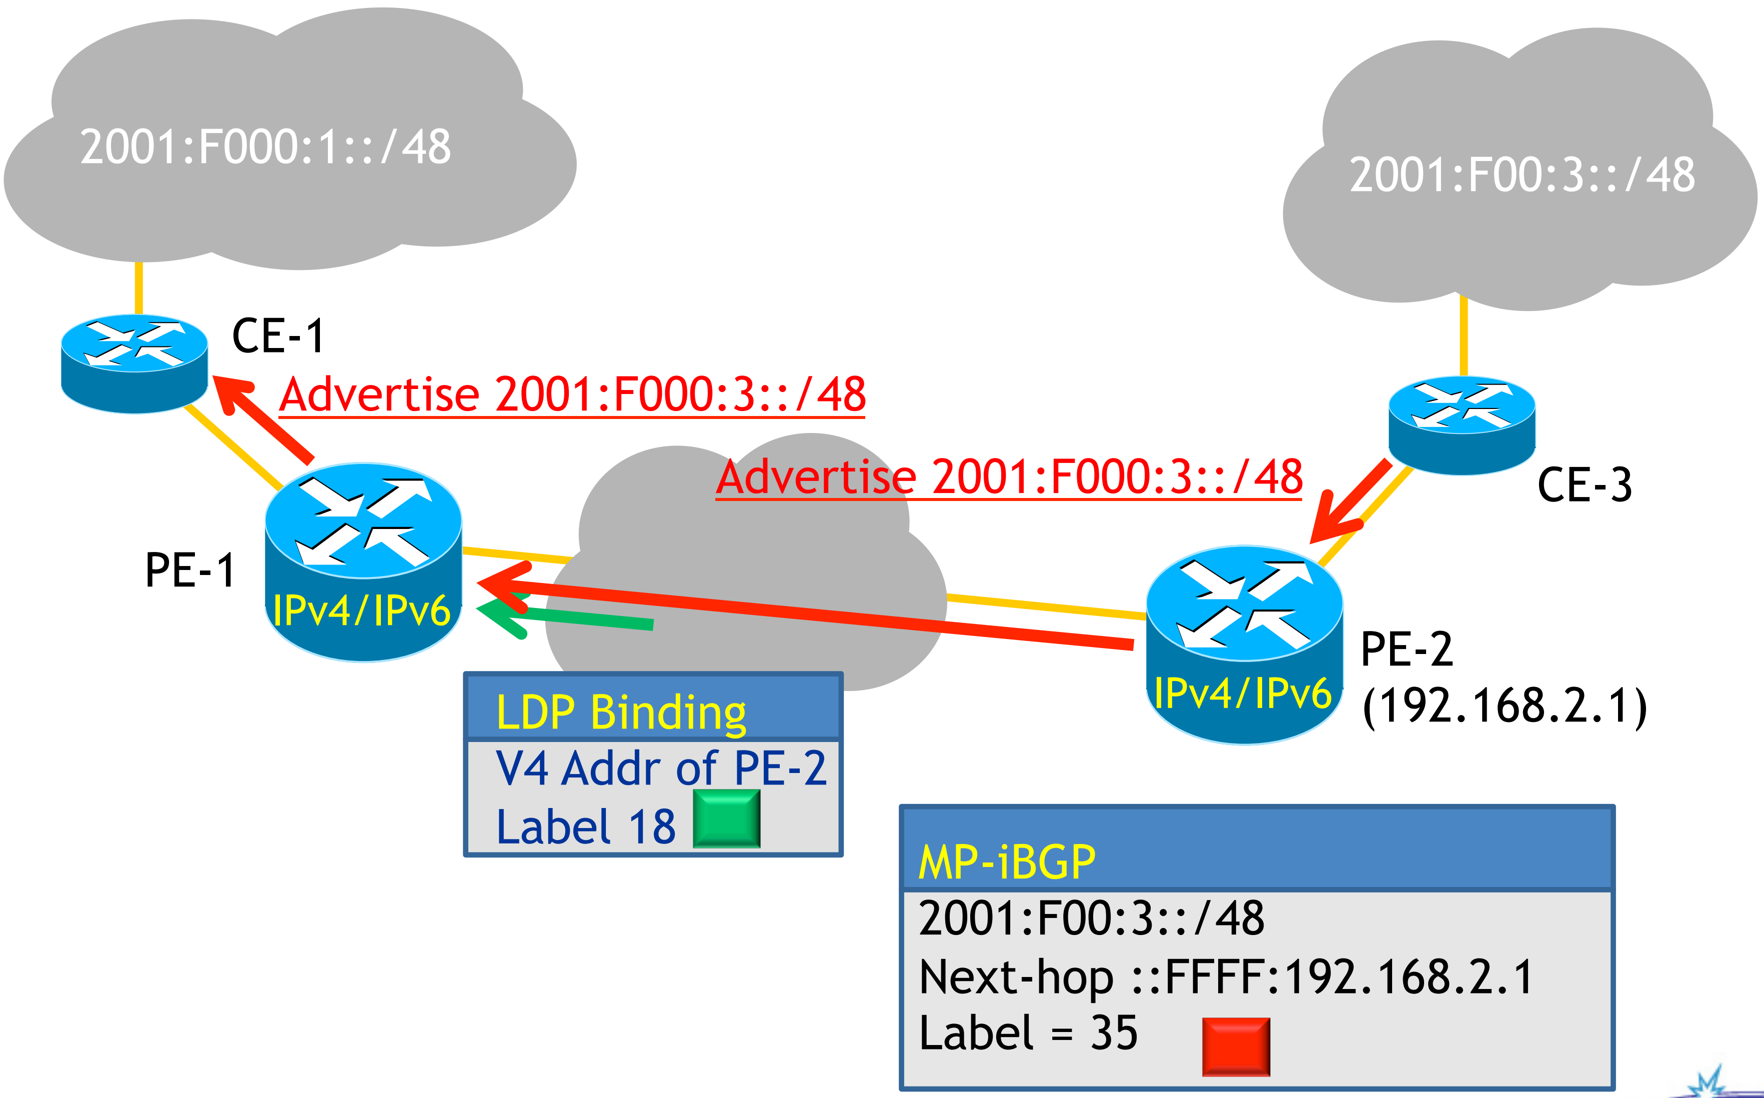

Packet Forwarding

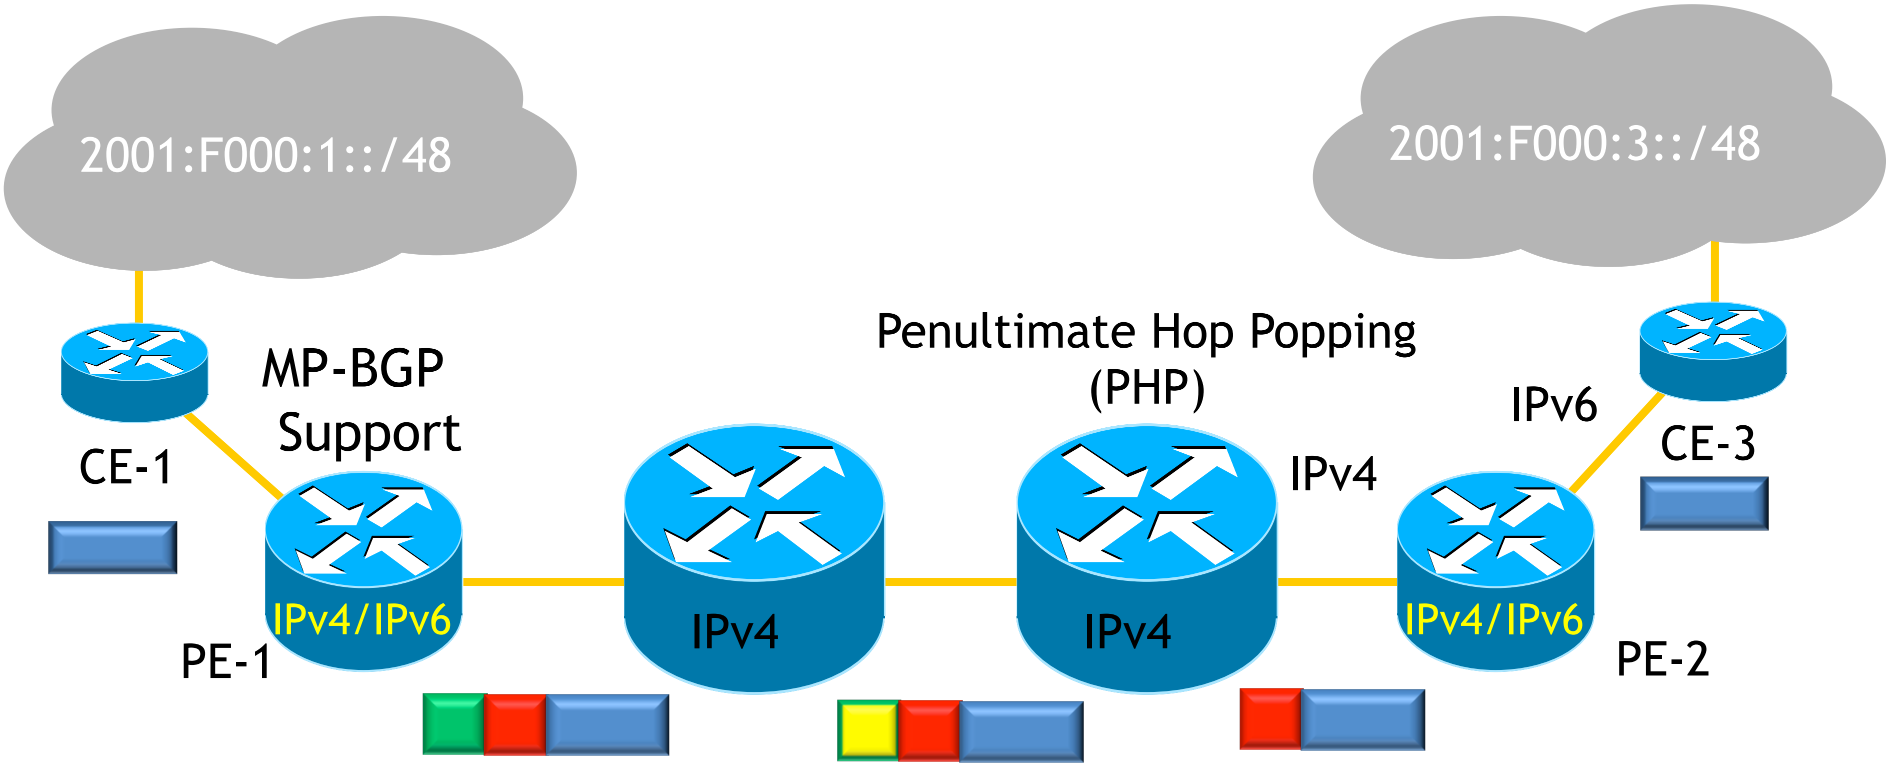

CE-1 sends IPv6 packet to PE-1
Ingress 6PE tunnels pushes Red label
Sends towards next-hop PE2 using Green Label

6VPE - IPv6 VPN Connectivity over IPv4-MPLS core

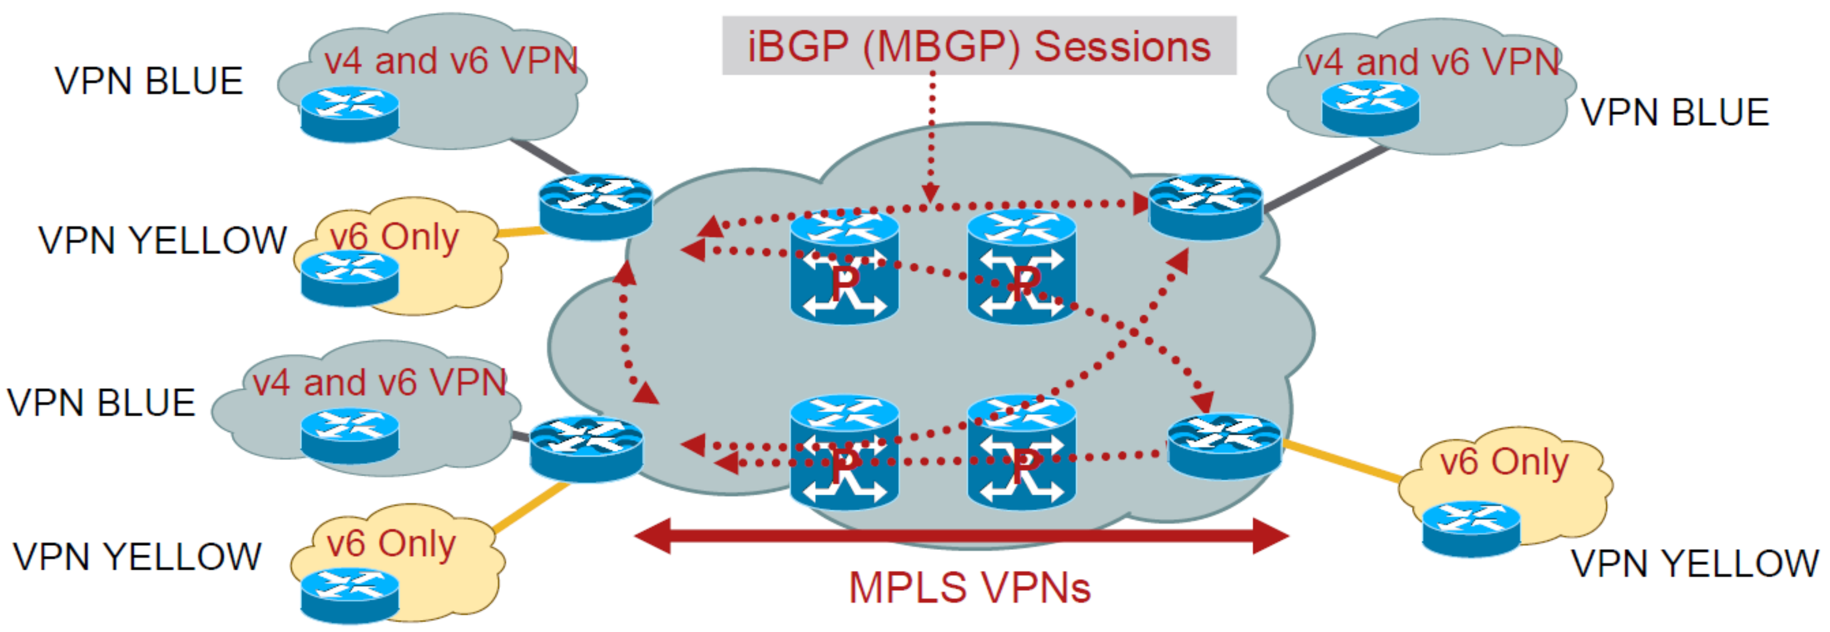

- Apply all RFC4364bis mechanisms to IPv6 VPNs:
- IPv6-VPN reachability exchanged among PEs via MP-BGP
- New BGP address family: AFI=2 (IPv6”), SAFI=128 (VPN)
- NLRI in the form of <length, VPN-IPv6-prefix, label>
- VRFs, RT, SOO, RRs,…operate exactly as with IPv4-VPN IPv6 packets

IPv6 VPN Provider Edge (6VPE)

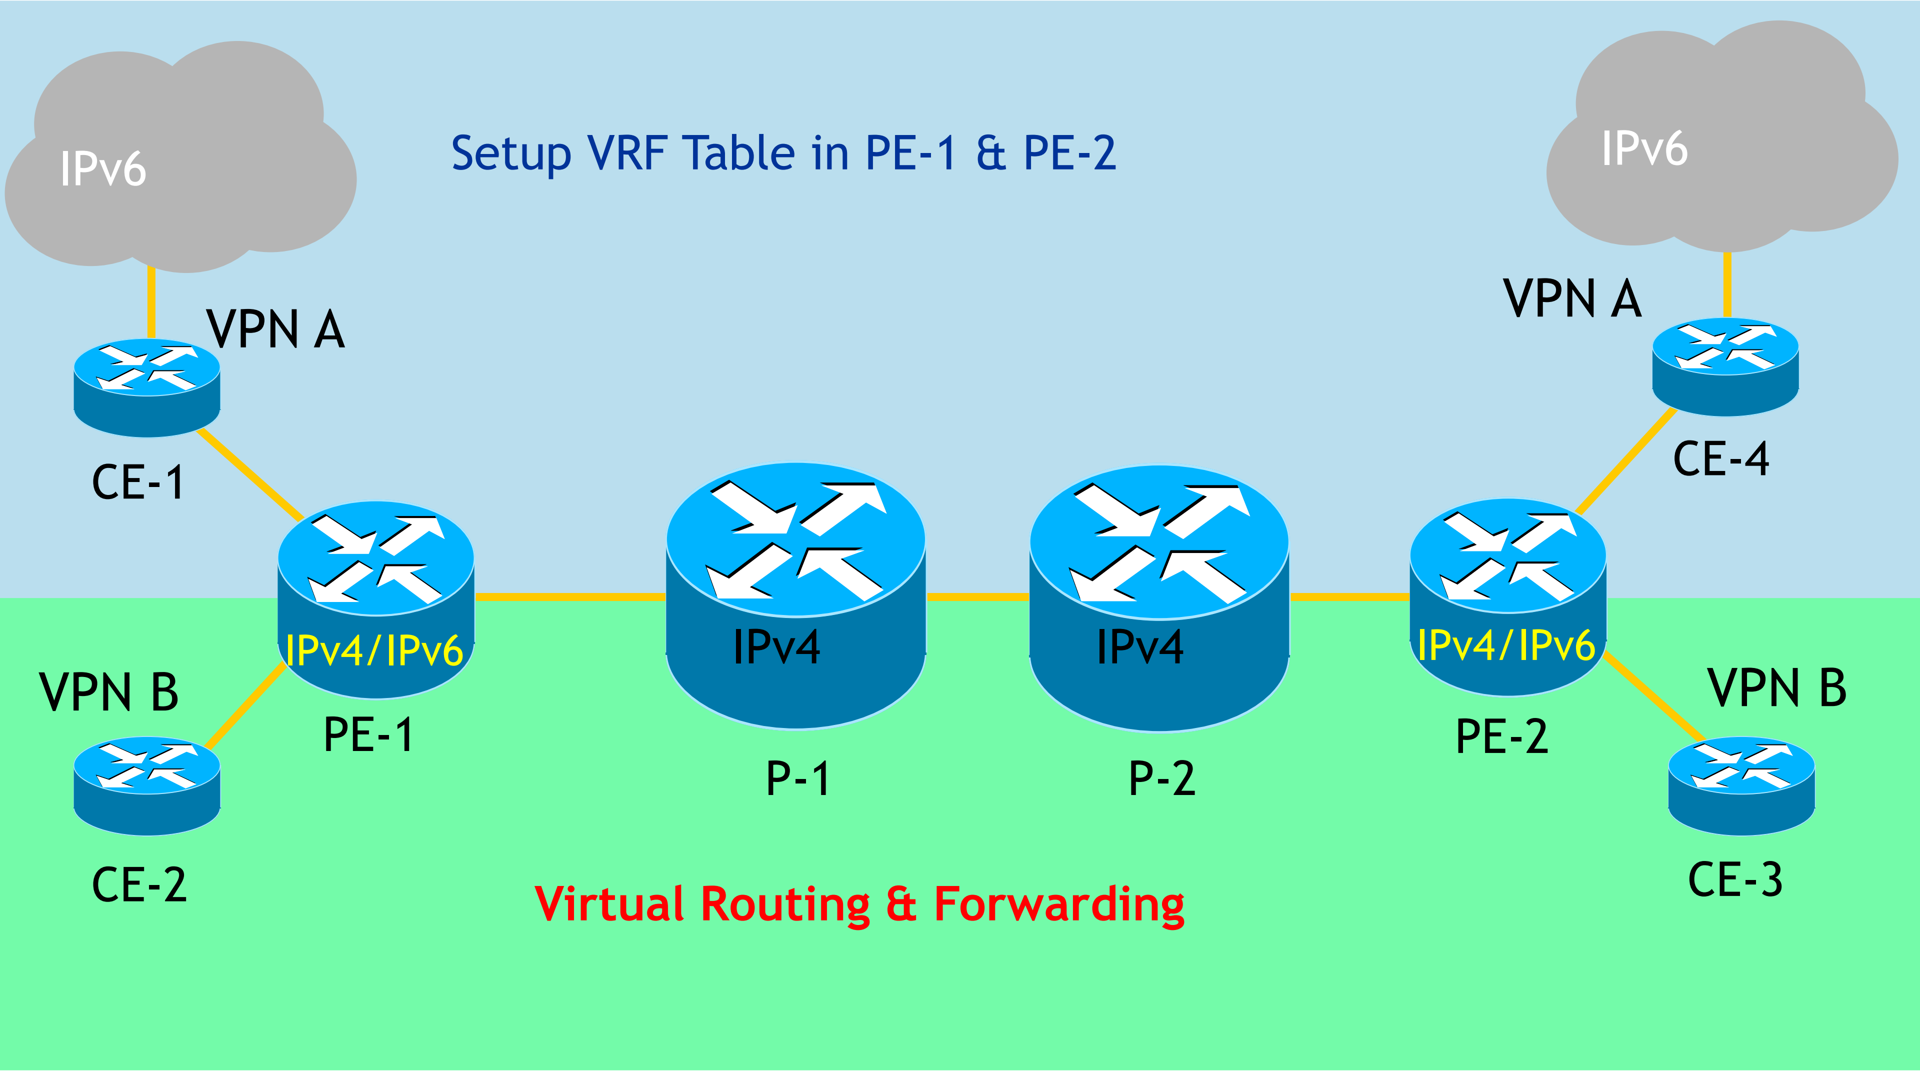

6VPE & VPN Routing & Forwarding (VRF)

1) In PE-1 build VRF Table name = Spirent
2) Associate this Table with e1
3) All addresses on link e1 belong to VPN/Spirent
4) Build a VRF on PE-2 with the same name
5) MP-BGP will advertise routing info between PE-1/PE-2
6) PE-1/PE-2 linked via MPLS Label (LDP or RSVP)
7) All addresses associated with Spirent in second Label

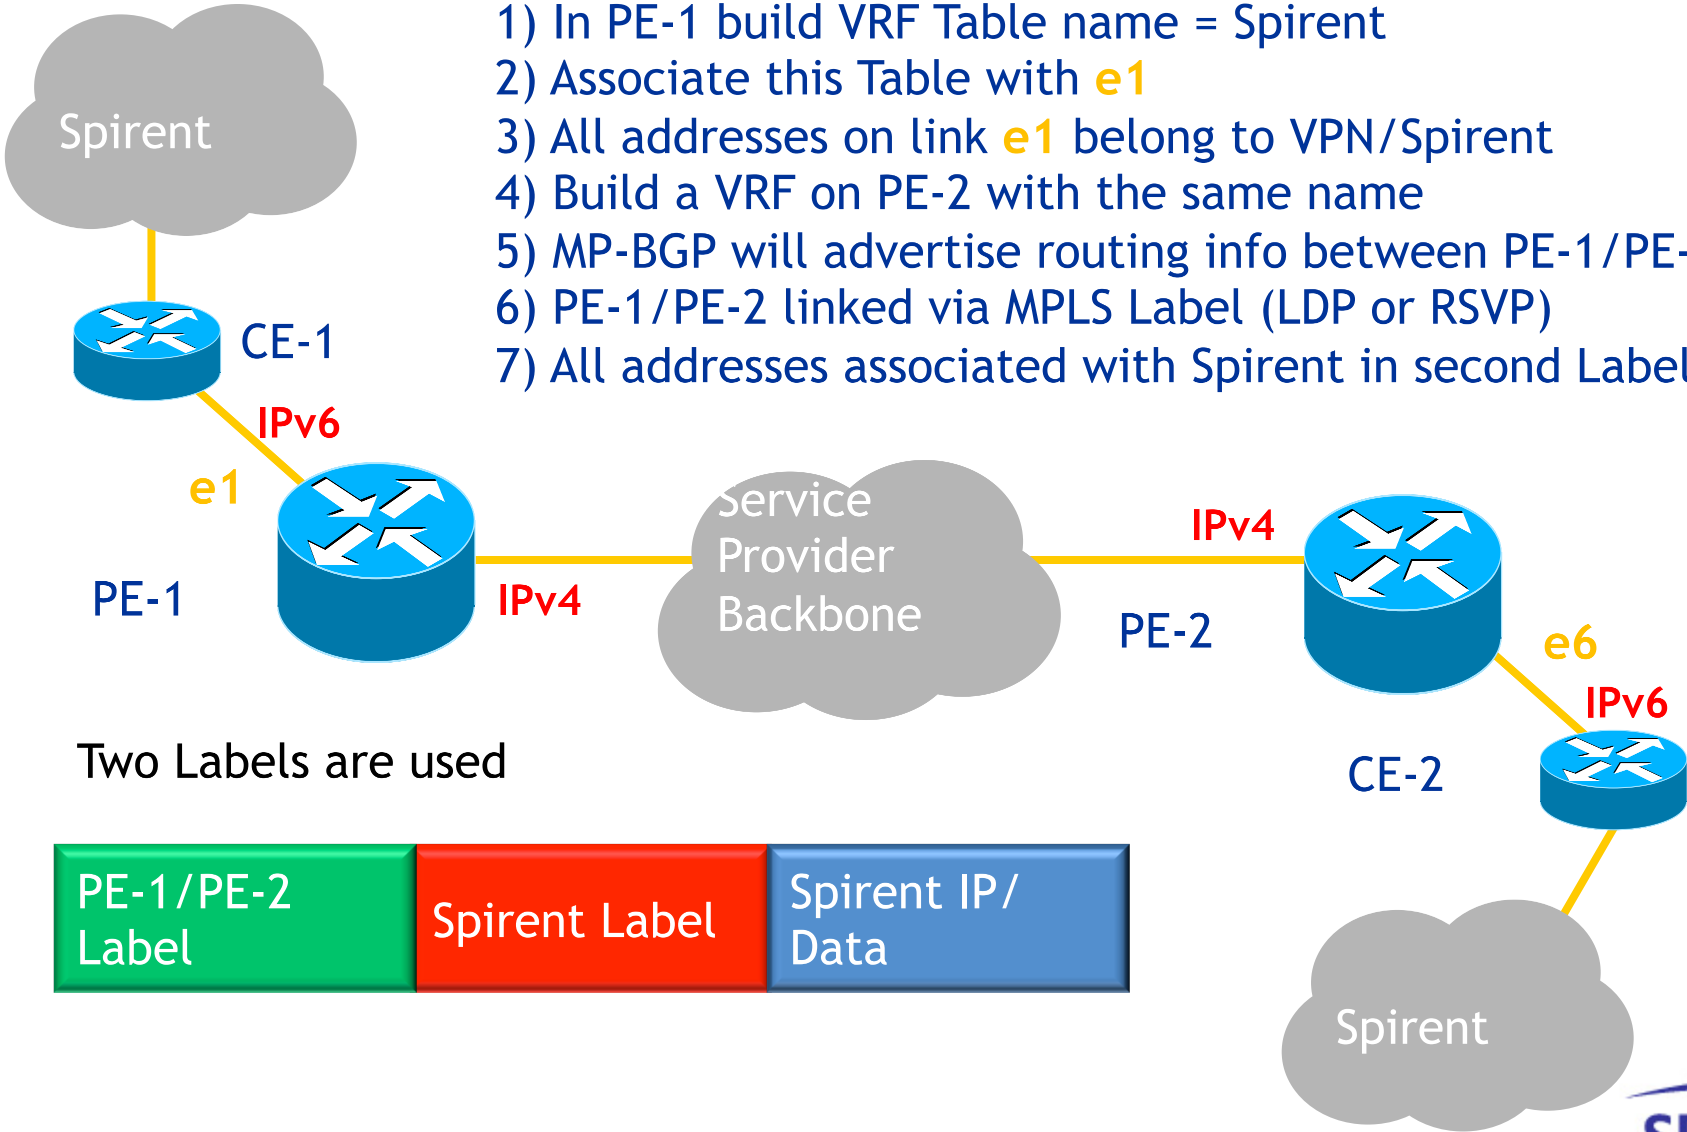

Security Concerns

- Invisible IPv6
- VPN Leakage
- Bandwidth hogging
- QoS / SLA Violations

SPIRENT

TESTING STRATEGIES

One Approach to Testing (not recommended)

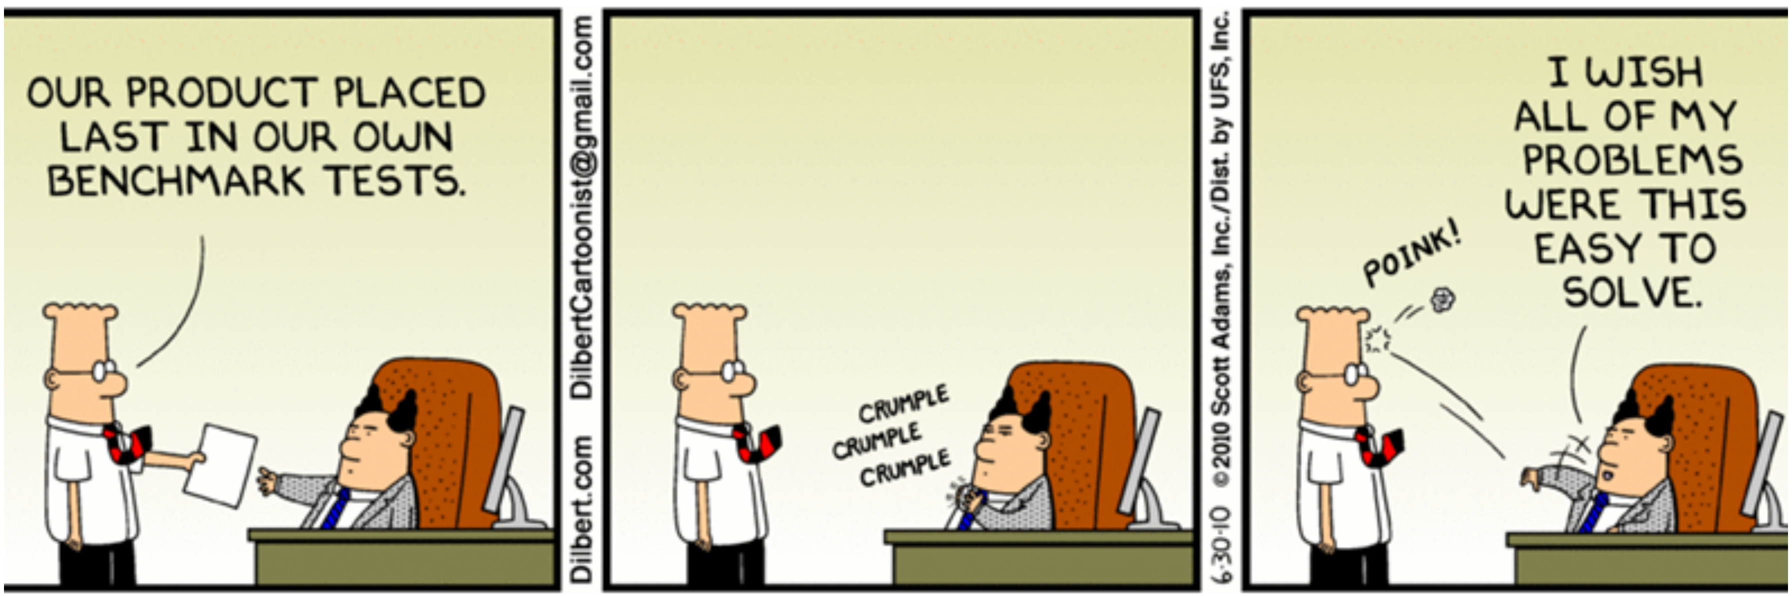

SPIRENT

Testing Strategies - What to Look For

- Does the system conform to relevant specifications?
- Does IPv4/v6 dual stack work correctly?
  - Does IPv6 traffic impact IPv4 traffic or vice versa?
  - How does IPv6 and IPv4 performance compare?
- Does IPv6 tunneling over IPv4 work correctly and does it scale
  - Is there a performance impact versus straight IPv4 or IPv6 forwarding
- Do control protocols function correctly & scale under IPv6?
  - Do they continue to function correctly under high load?
- Do the QoS mechanisms work correctly for IPv6 streams?
- Can IPv6 traffic have an impact on IPv4 and vice versa?
- Is the IPv6/IPv4 tunneling device able to prevent security attacks?
  - What effect does this have on forwarding performance?

SPIRENT

Generic Device Architecture

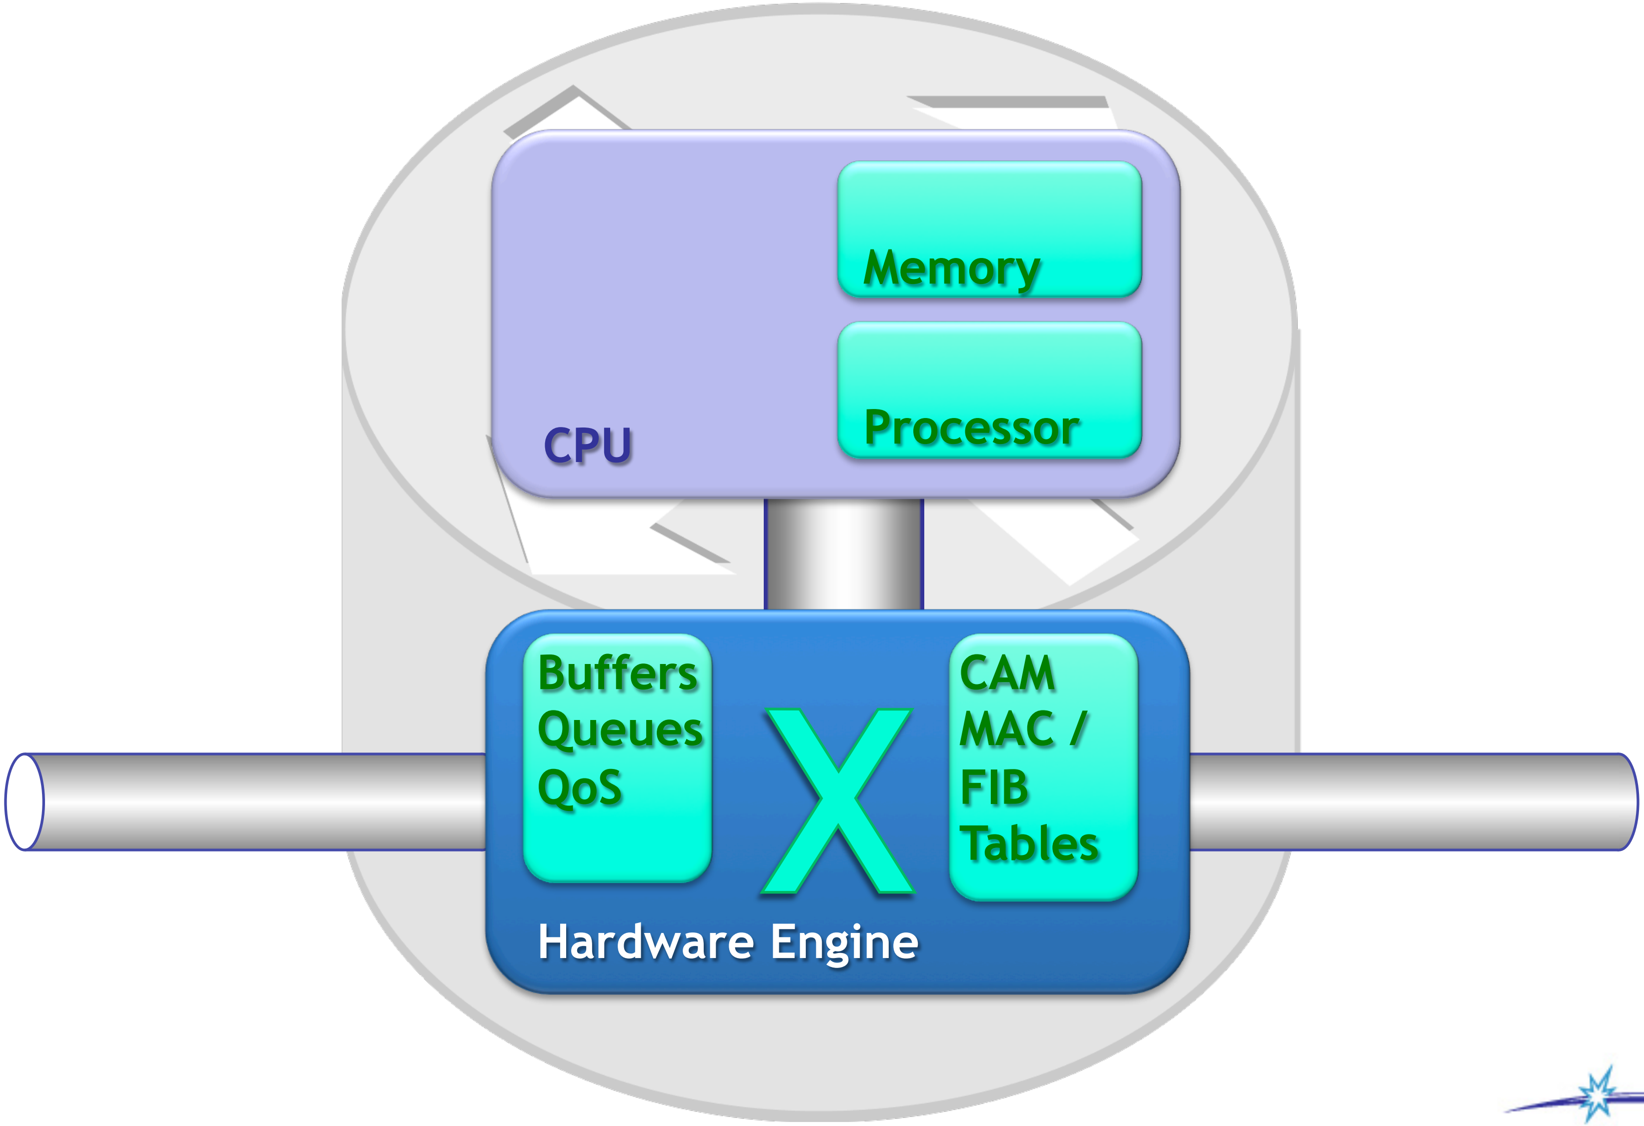

Dual Stack

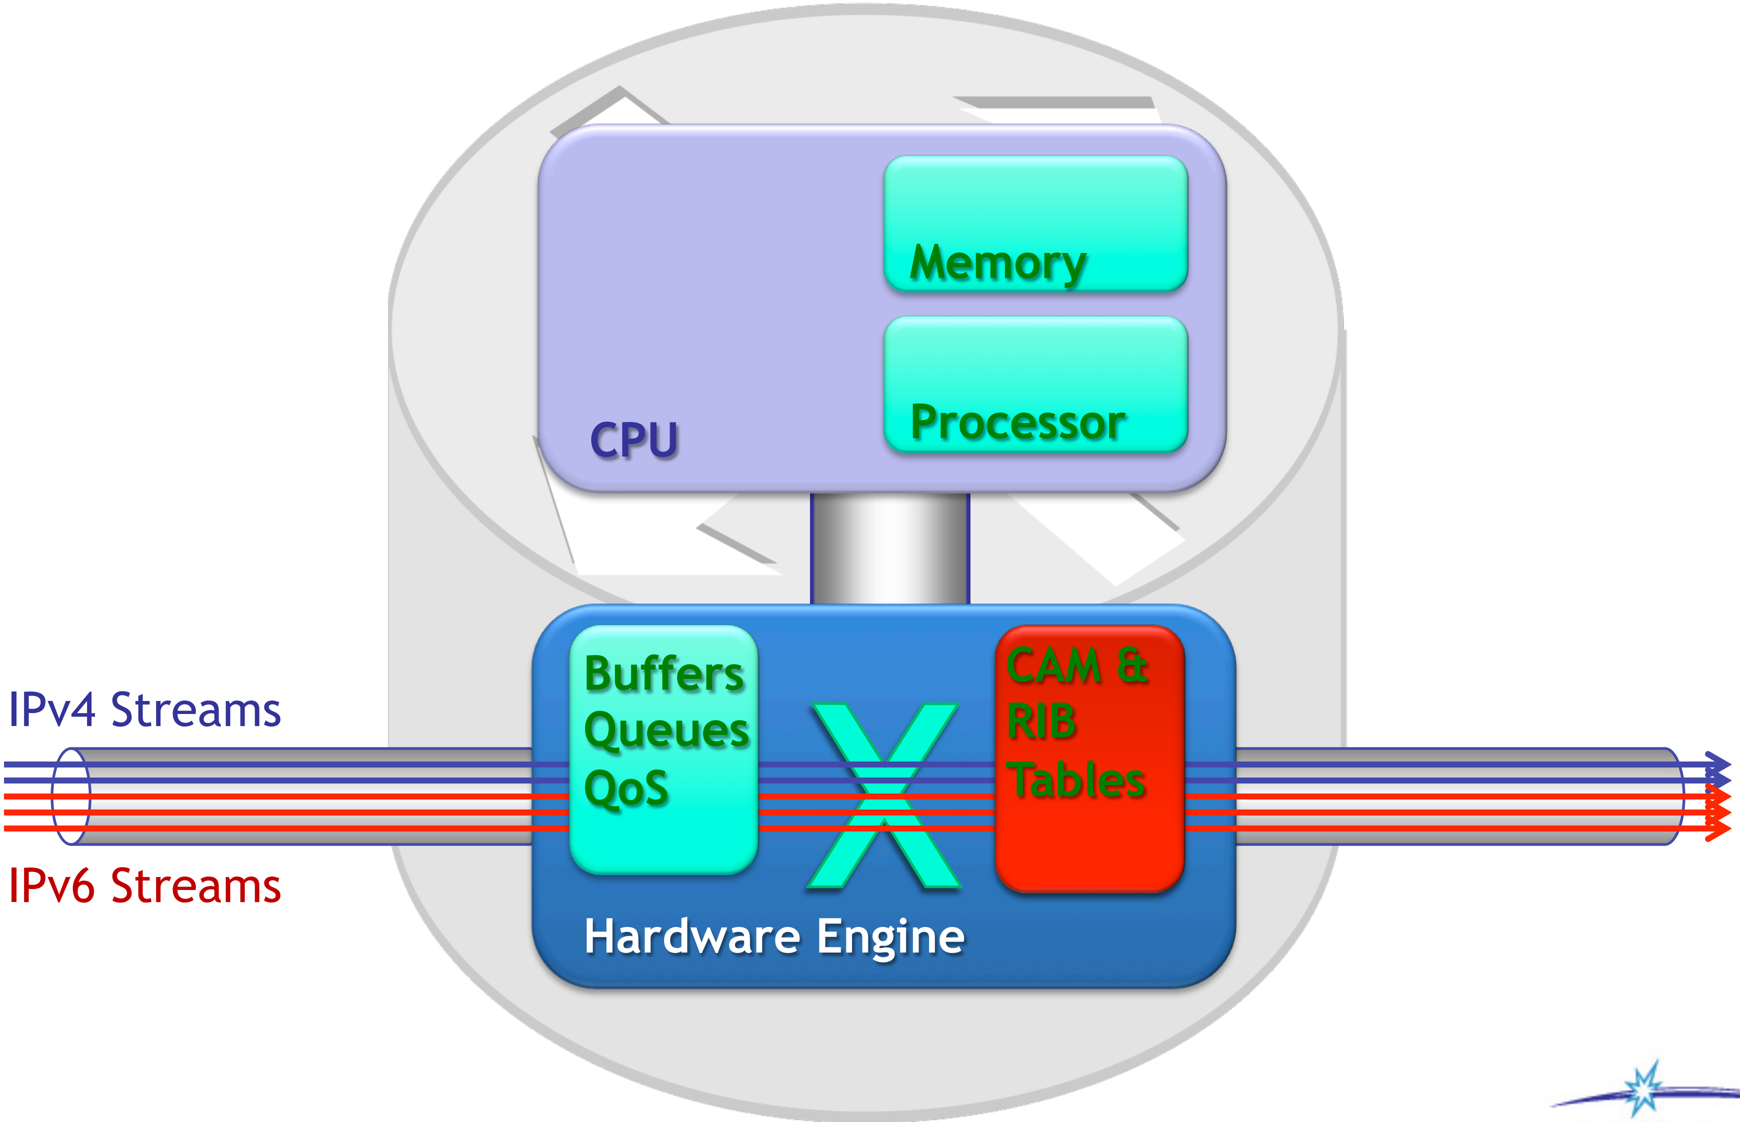

SPIRENT

Dual Stack Testing

Test with IPv4
Test with IPv6
Test with IPv4 & IPv6 (Dual Stack)

IPv6 & IPv4 Traffic Generation

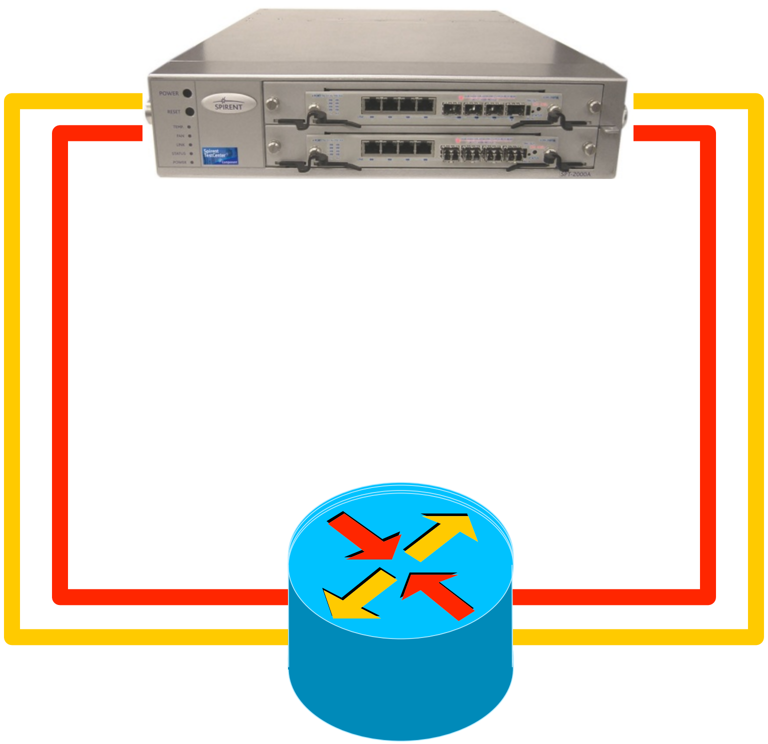

A Good Test Will …

- Use 1000’s of streams of each type
- Use a varied range of addresses and prefix lengths to prevent aggregation in FIB
- Use varied DSCPs to check DiffServ operation across both stacks

Tunnelling - 6RD Border Relay

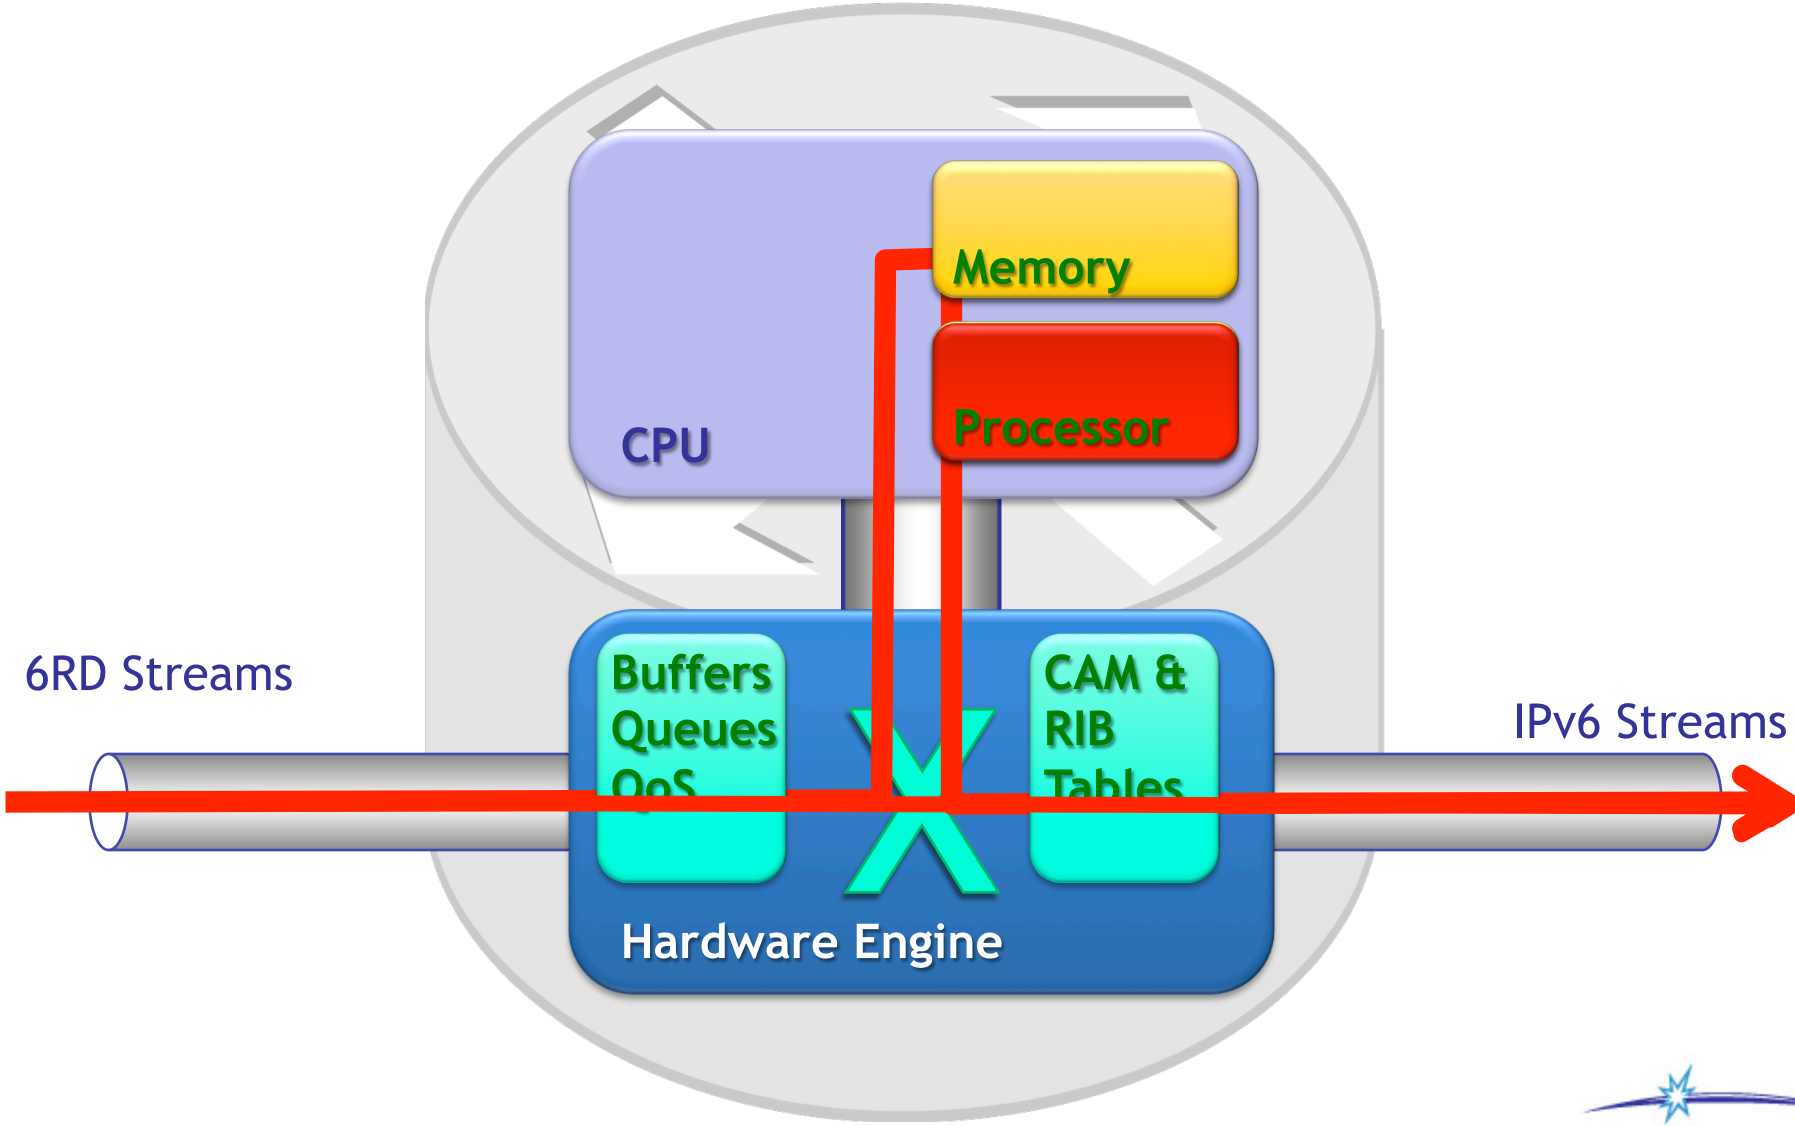

SPIRENT

6to4 and 6RD Testing

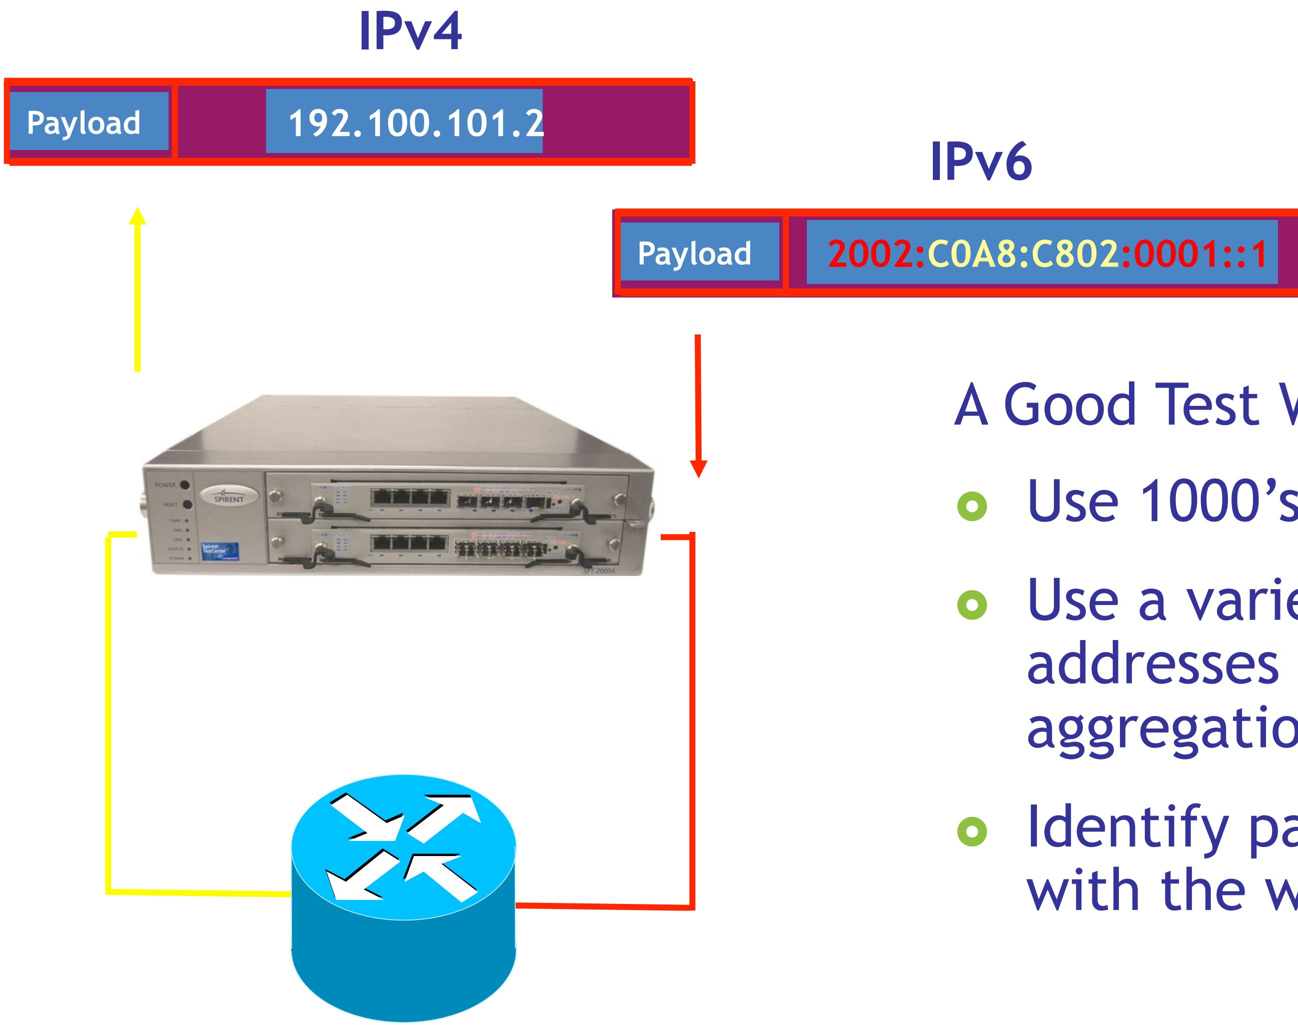

A Good Test Will …

- Use 1000's of streams of each
- Use a varied range of addresses prevent aggregation in FIB
- Identify packets received with the wrong address

6VPE Example Device Under Test

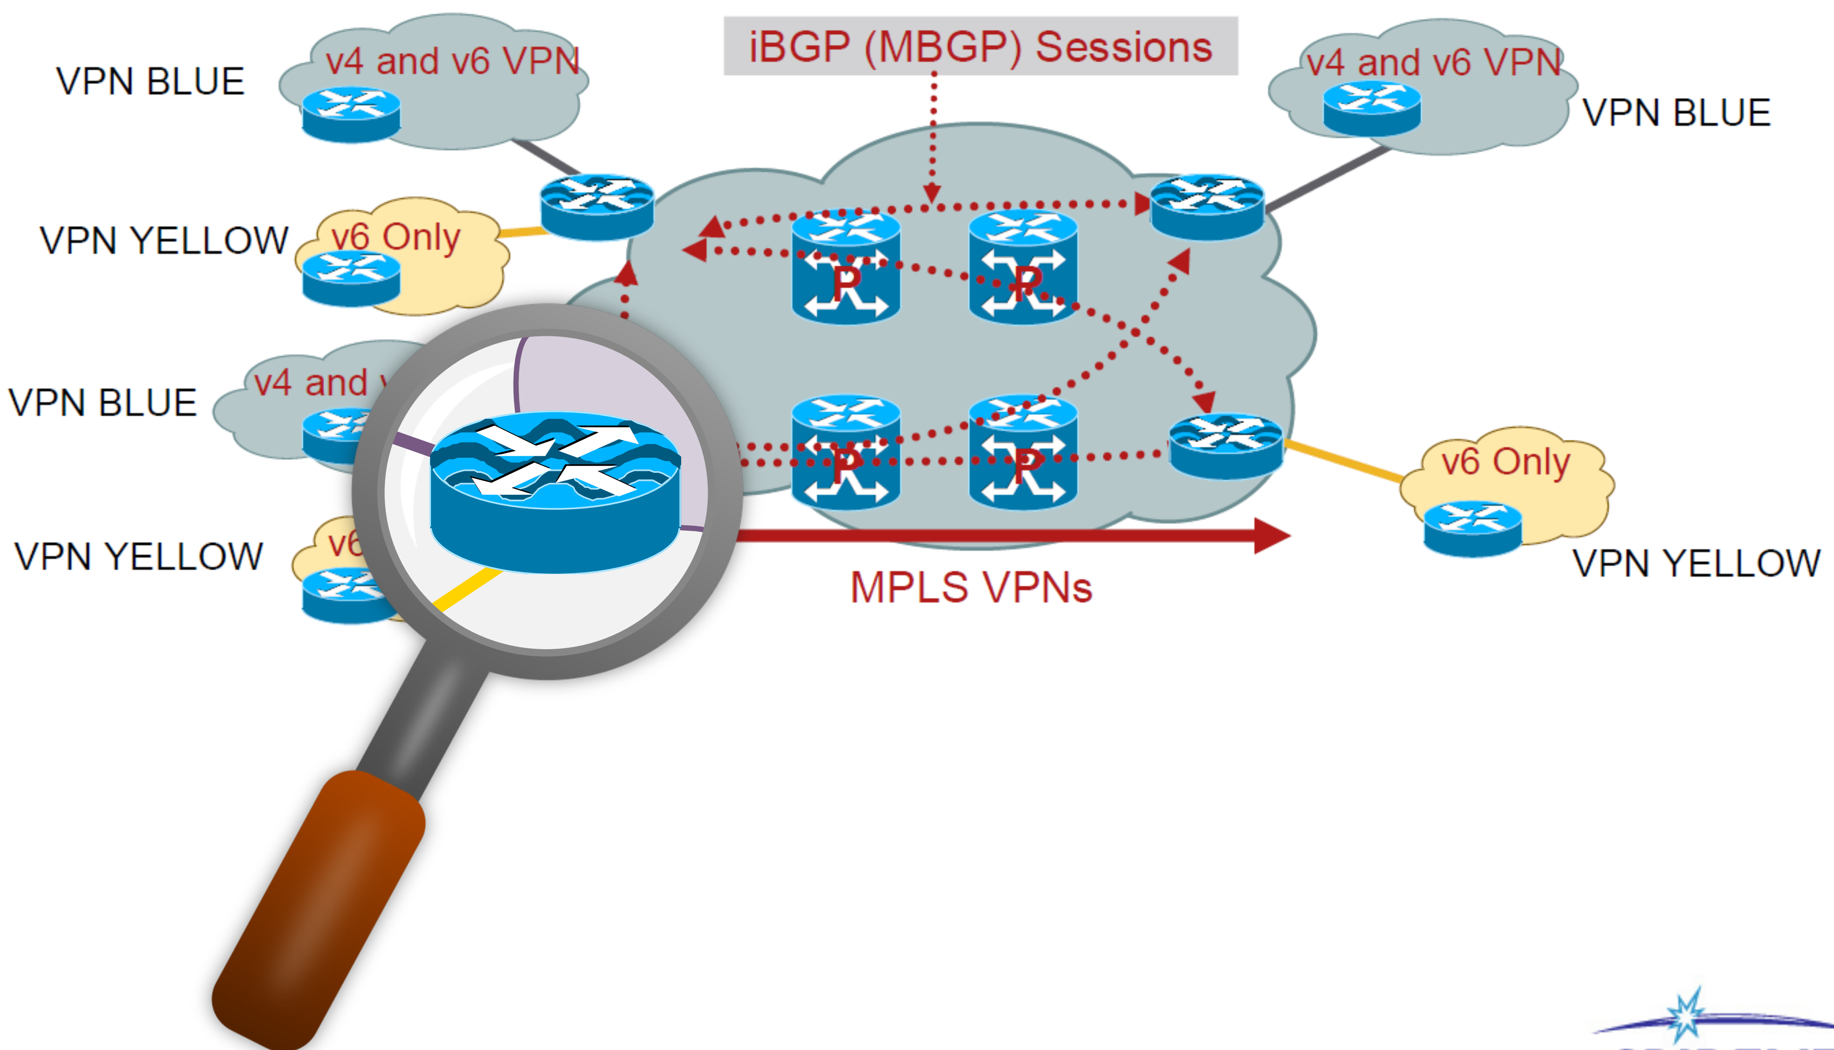

SPIRENT

Complex Environment

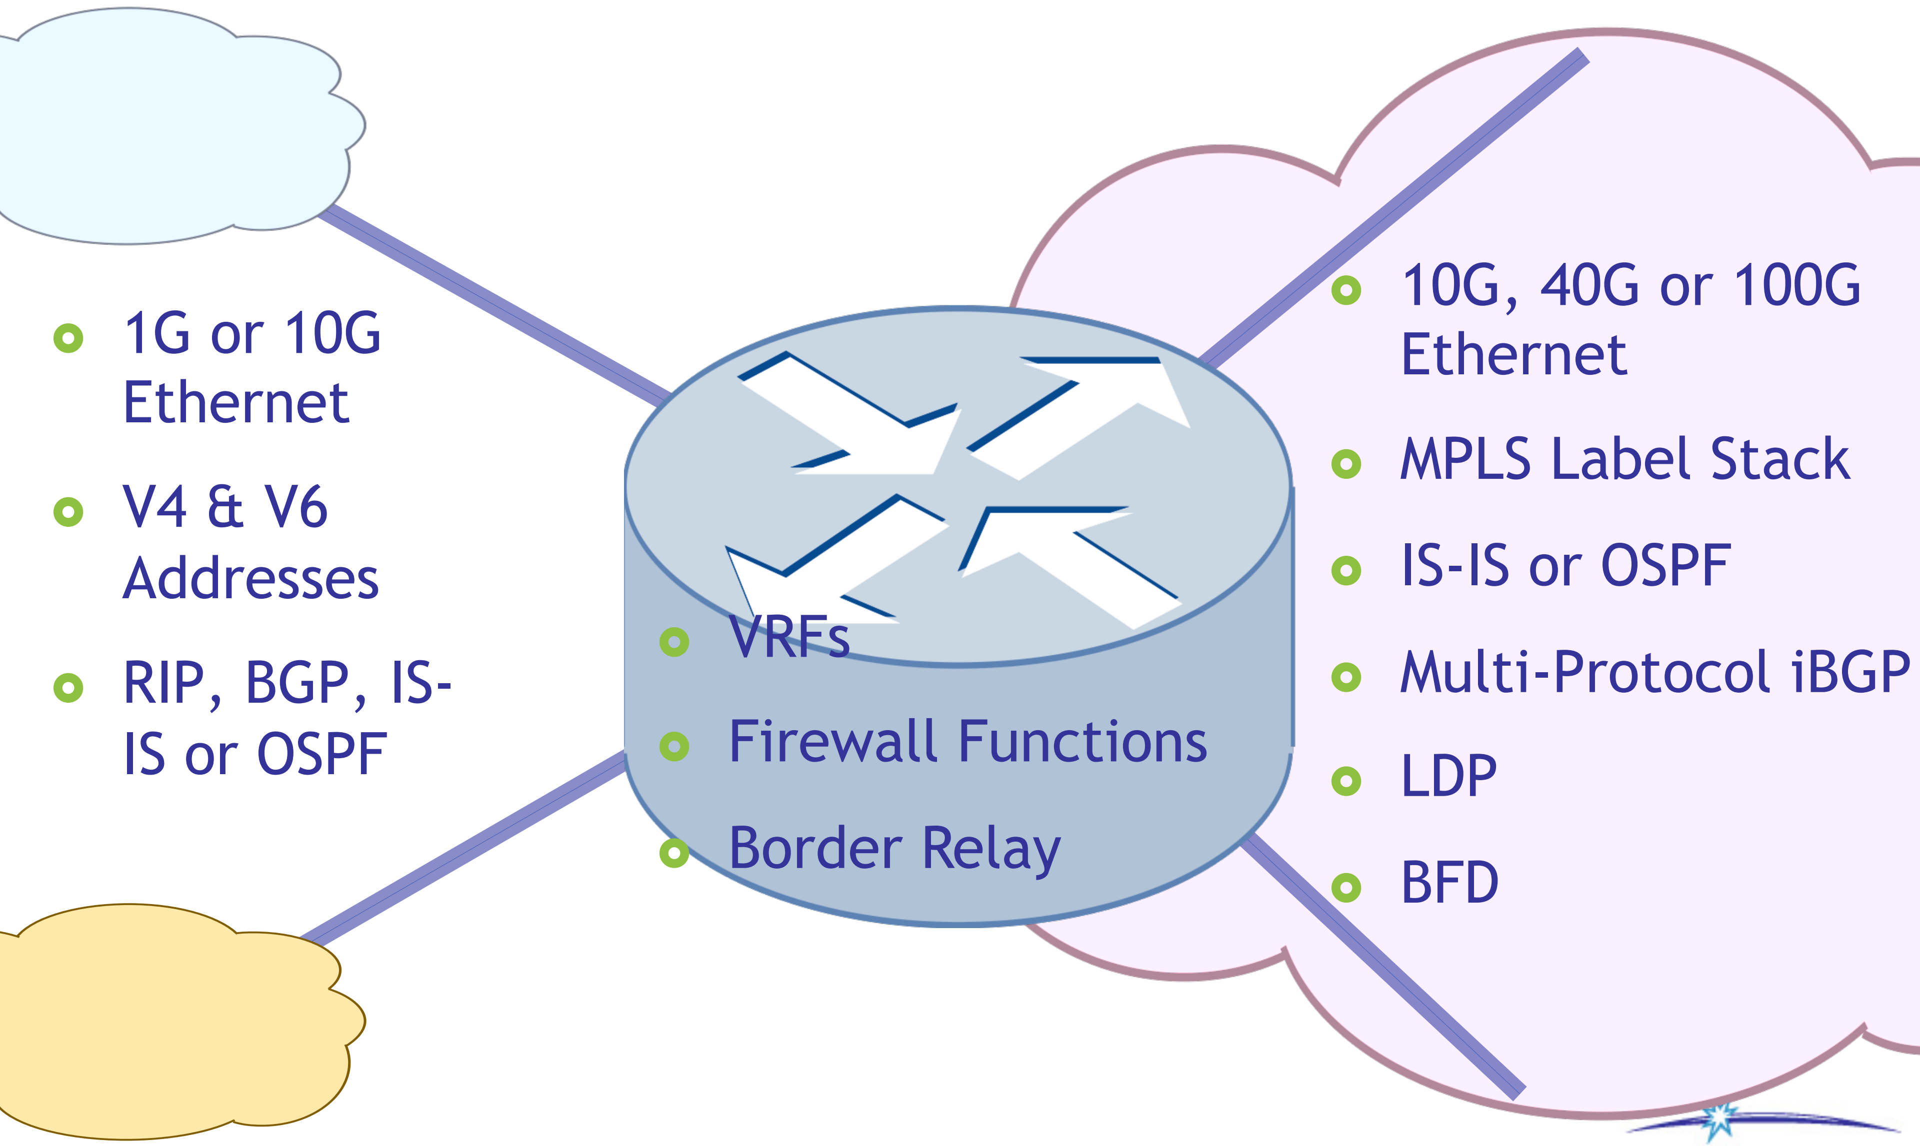

- 1G or 10G Ethernet
- V4 & V6 Addresses
- RIP, BGP, IS-IS or OSPF

- VRFs
- Firewall Functions
- Border Relay

- 10G, 40G or 100G Ethernet
- MPLS Label Stack
- IS-IS or OSPF
- Multi-Protocol iBGP
- LDP
- BFD

6VPE Device - Control Plane Stress

- CPU Load increases with number of peers
- More routes and more protocols requires more memory
- VRFs exacerbate memory problems

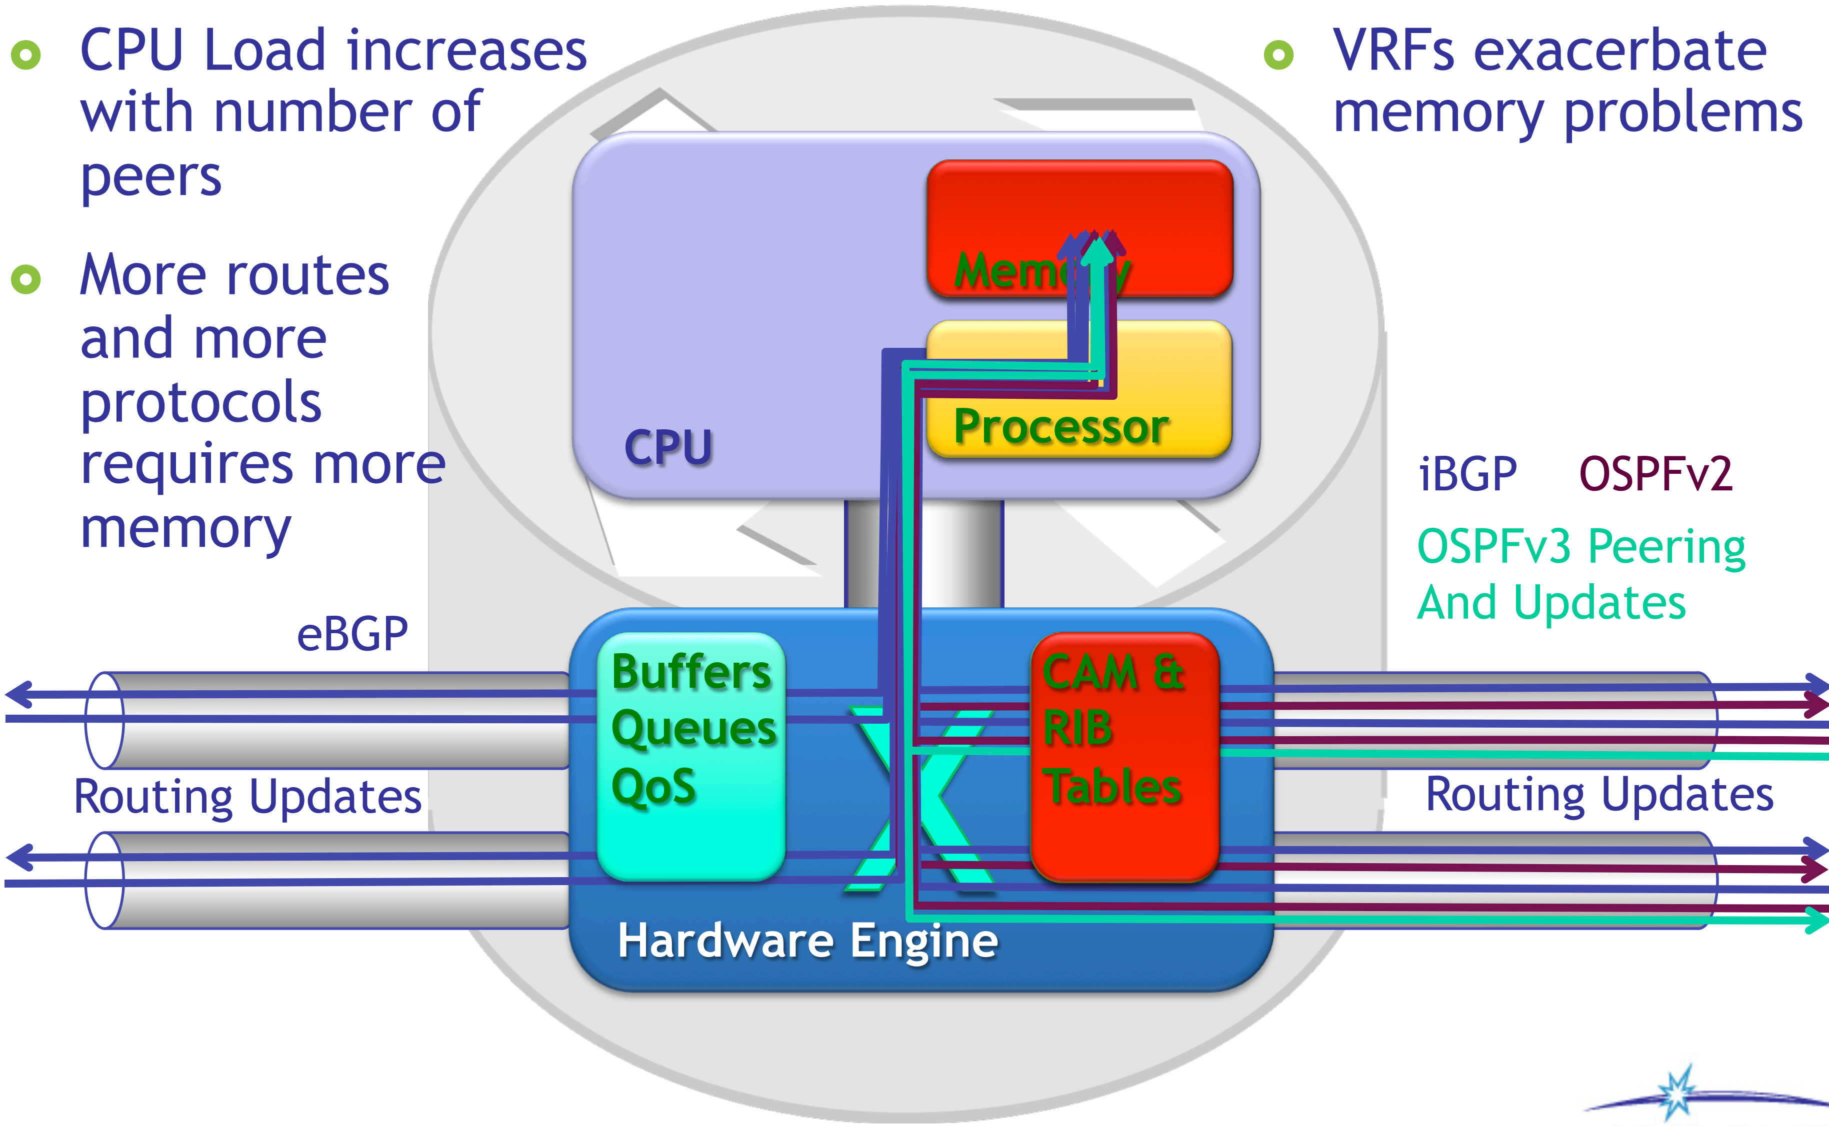

6VPE Device - Data Plane Stress

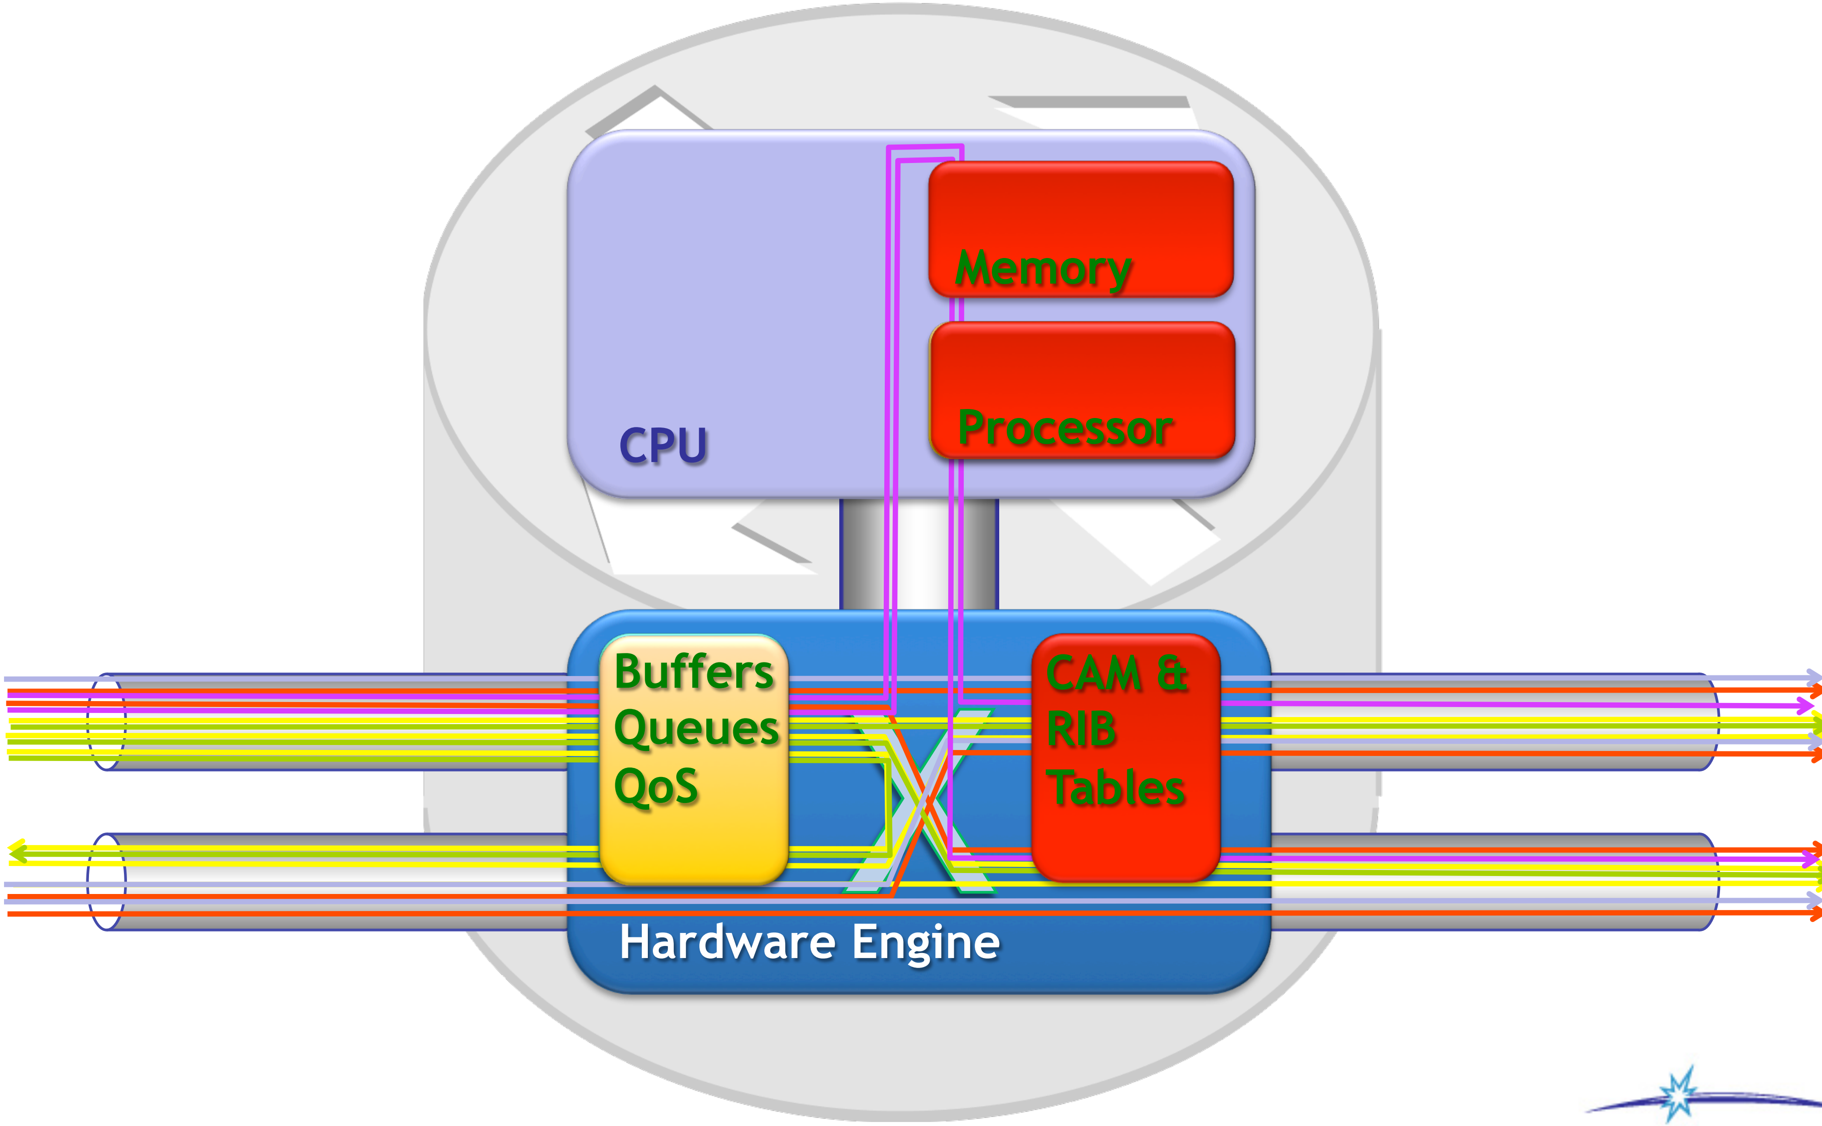

SPIRENT

Example of 6VPE Testing
Test the control plane and data plane

- Set-up BGP Peers on one port and advertise VRF Routes towards the DUT (MPLS core side)
- Transmit data from the second port to the CE side of the DUT using IP addresses advertised above
- Measure the received rate of traffic on the first port
  - Check for latency loss etc.
- Withdraw 50% of routes after 30 seconds
- Measure the effect on the received rate of traffic
- Repeat for different loads

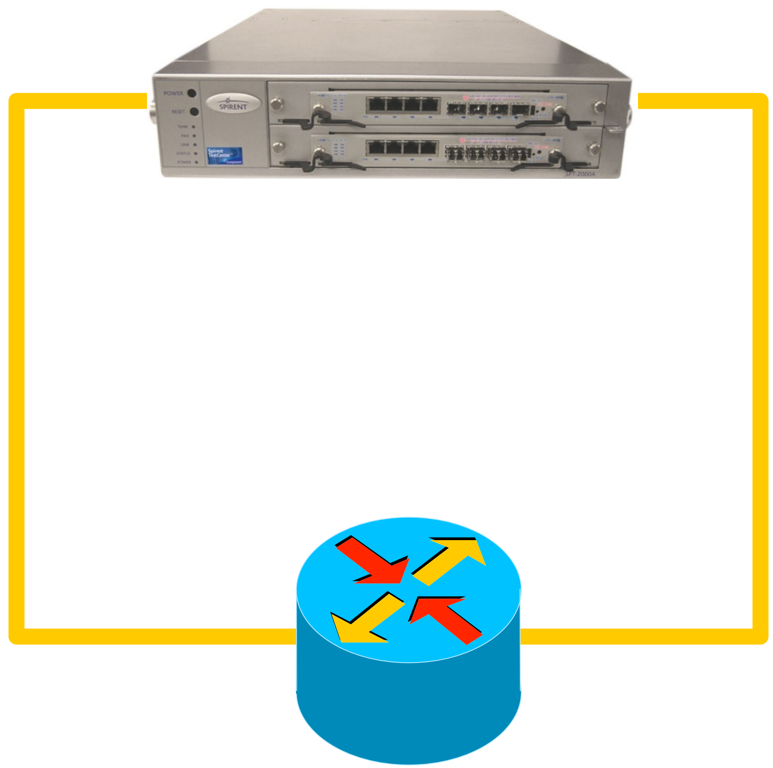

BGP Route Flap

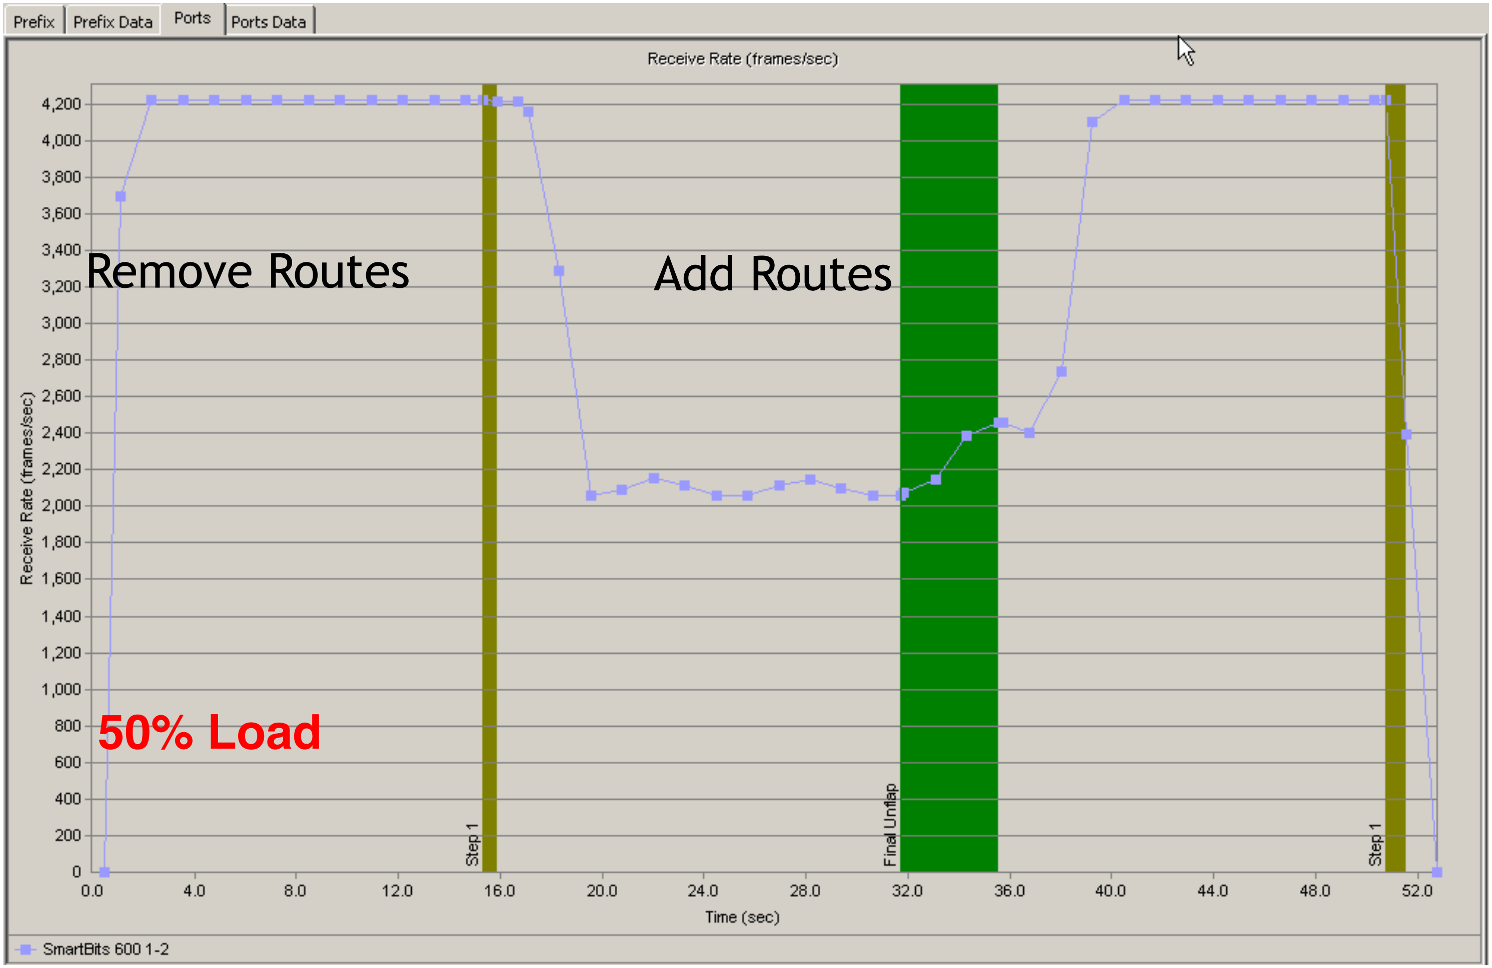

SPIRENT

BGP Route Flap

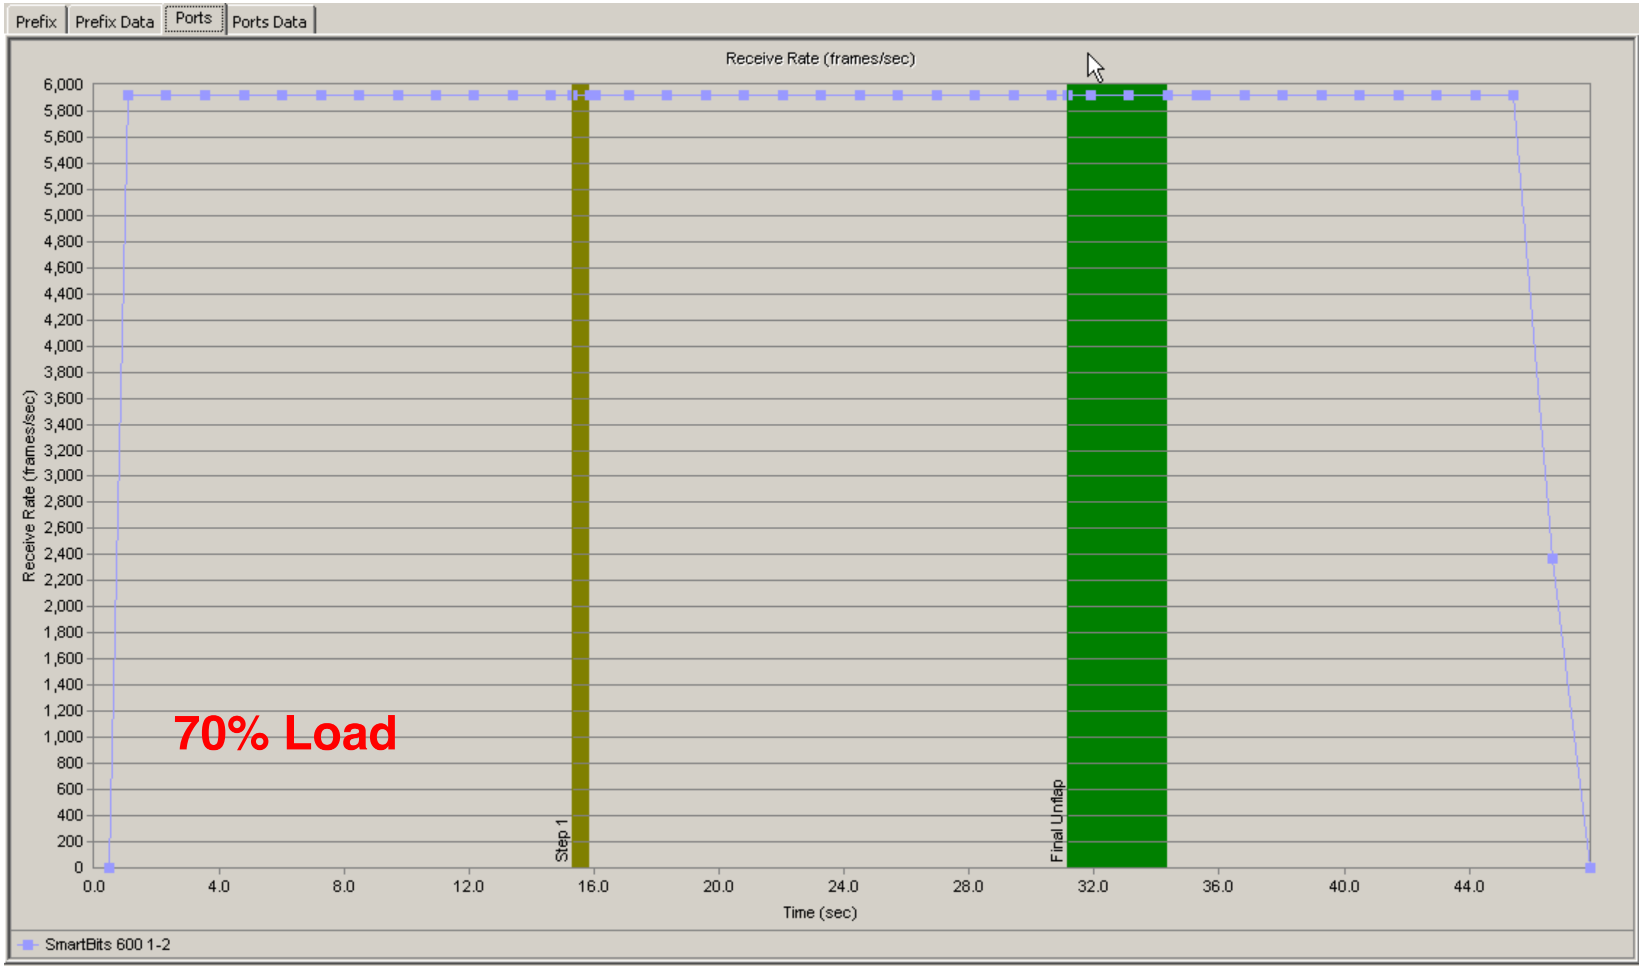

SPIRENT

Summary

- Network devices operate in highly complex environments
- Failures such as VPN leakage tend to happen under stressful network conditions
- In order to find the failure point of the system it is necessary to fully and accurately emulate that environment
- A simple test at 100% load with a few streams will more than likely pass
- Tens of thousands of realistic streams with a highly diverse set of prefixes and prefix lengths should be used.
- **Every device has its limits. Discover what they are via testing and design the network so you never reach them**

Will Your IPv6 Network Pass the Test?

SPIRENT

steve.jarman@spirent.com

THANK YOU

White Papers and other resources available at

www.spirent.com

BACKUP

How can Spirent help?

- Measure performance of Border Gateways
- Measure overall server performance
- Application/Security testing
- IPSec Testing
- Measure performance of IPv6, IPv4 & Dual Stack Routers
- Measure performance IPv6/IPv4 Tunnel Transition Devices

- IPv6 Protocol conformance testing.
- Professional Services

SPIRENT

Service Provider - Why Testing is Important

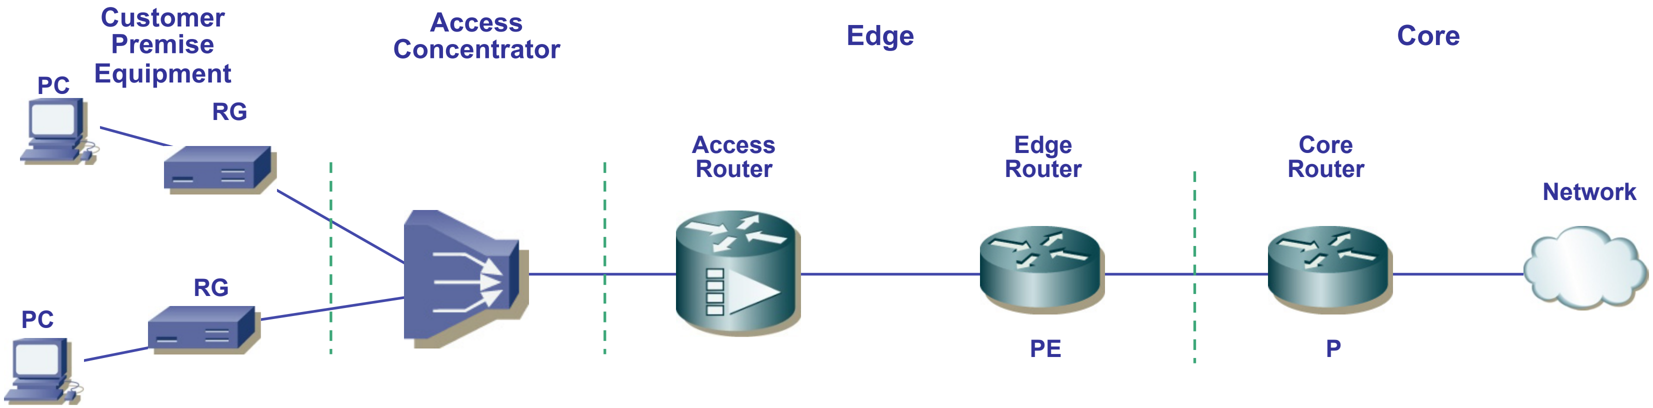

CPE

- Adhere to standards
- Performance
- Vendor interoperability
- Reduce Bad Press

Access

- Subscriber scalability
- Fail over
- Redundancy
- QoS
- Routing & MPLS Functionality

Edge

- Subscriber scalability
- Traffic Management
- QoS/QoE
- Routing & MPLS scale & performance

Core

- Data Performance
- Routing, MPLS performance
- Vendor interoperability

SPIRENT

IPv6 Routing Types

- Static
- RIPng (RFC 2080)
- IS-IS for IPv6
- OSPFv3 (RFC 2740)
- MP-BGP (RFC 2545/2858)

Static Routing

Configured in the same way as with IPv4

There is an IPv6-specific requirement per RFC 2461:

"A router must be able to determine the link-local address of each of its neighbouring routers in order to ensure that the target address of a redirect message identifies the neighbour router by its link-local address."

RIPng

- Features Taken from IPv4:
  - Based on RIPv2
    - Distance-vector
    - 15-hop radius
    - split-horizon
    - poison reverse
    - Etc.
- Features Updated for IPv6:
  - Uses IPv6 for transport
  - IPv6 prefix, next-hop IPv6 address
  - Uses the multicast group FF02::9 for RIP updates
  - Updates are sent on UDP port 521

SPIRENT

IS-IS for IPv6

IS-IS an OSI routing protocol originally designed as an intra-domain routing protocol for Connectionless Network Service (CLNS) traffic,

- Major operation remains unchanged:
  - Level 2 (backbone) device route between Level 1 areas
  - Each IS device still sends out LSP packets
  - Neighborship process is unchanged
- IPv6 support gets added based on RFC 5308 - Routing IPv6 with IS-IS

OSPFv3 - RFC 2740

- Based on OSPFv2, with enhancements
  - Distributes IPv6 prefixes
  - Runs directly over IPv6
- Ships in the night with OSPFv2
  - RFC 5838 - Support of Address Families in OSPFv3 includes IPv4 Unicast and Multicast families
- Adds IPv6-specific attributes:
  - 128-bit addresses
  - Link-local address
  - Multiple addresses and instances per interface
  - Authentication (now uses IPsec)
  - OSPFv3 runs over a link, rather than a subnet

BGP

- To make BGP-4 available for other network layer protocols, RFC 2858 (obsoleted RFC 2283) defined multiprotocol extensions for BGP-4
- Runs over TCP which, in turn, runs over IPv4 or IPv6
- Defines Address Families enabling BGP-4 to carry information of other protocols e.g. MPLS and IPv6
  - Address Family Information (AFI) for IPv6
  - AFI = 2 (RFC 1700)
    - Sub-AFI = 1 Unicast
    - Sub-AFI = 2 Multicast for RPF check
    - Sub-AFI = 3 for both Unicast and Multicast
    - Sub-AFI = 4 Label
    - Sub-AFI = 128 VPN