IPv6 - Design, Deployment and Challenges

Luyuan Fang
Cisco Systems
lufang@cisco.com

Agenda

- Market Drivers and deployment strategies
- IPv6 Over MPLS Design Considerations
  - Core Design
  - Edge Design
  - RR Design
  - Address Design
  - Lessons learned
- Challenges in IPv6 deployment
  - Scaling VPN Routes
  - QoS and Performance
  - MVPNv6 Support
  - Security Considerations

Market Drivers and Deployment Strategies

IPv6 Drivers

- Next Generation Programs and Capabilities to Defense sectors
- Government Mandates
- IPv4 address depletion concerns
- New services such as content delivery
- New product and operating system introduction - Microsoft Vista, HP, Broadcom, Sony, Panasonic
- Cable consumer market address scaling
- Mobile Services - Quad Play and Fixed Mobile Convergence
- Population densities in Japan and APAC
- China Market 3G deployment

Distribution of IPv4 addresses by /8

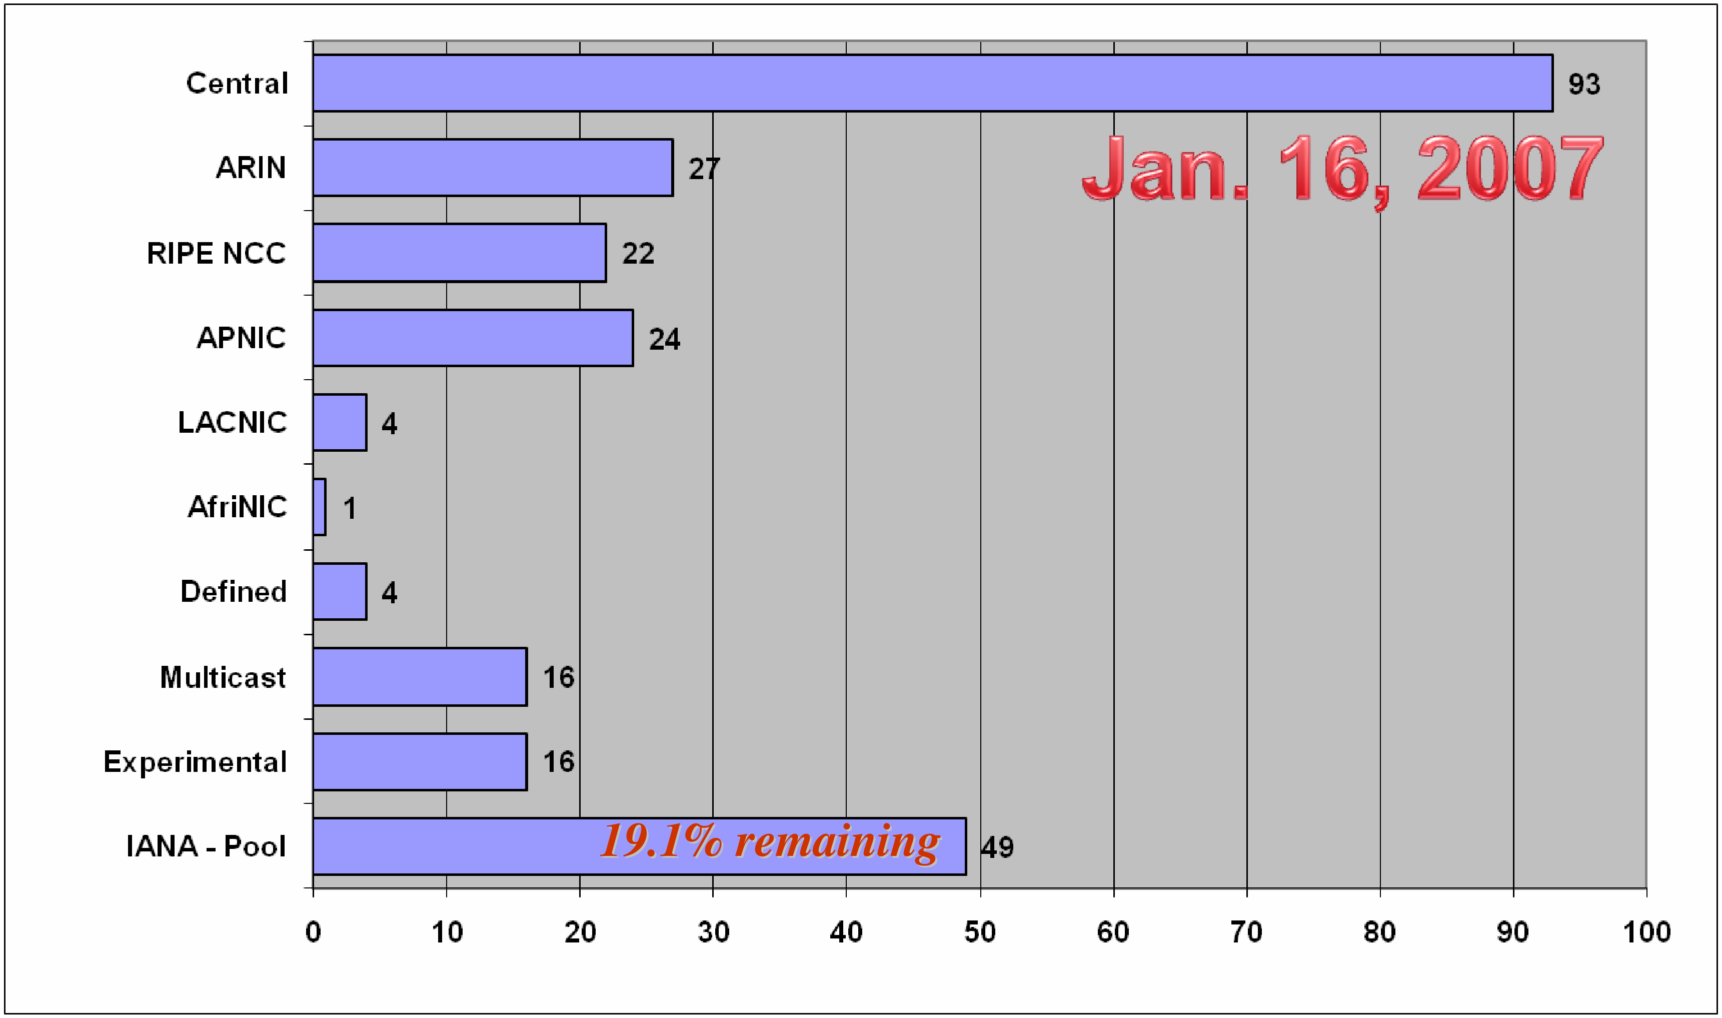

Jan. 16, 2007

19.1% remaining

IP Address Consumption is accelerating

IP Address Allocation History
Full discussion at: **www.cisco.com/ipj**
The Internet Protocol Journal
Volume 8, Number 3, September 2005

- **Consumption is accelerating despite increasingly intense conservation efforts.**
  - PPP / DHCP (temporal address sharing)
  - CIDR (classless inter-domain routing)
  - NAT (network address translation)
  - plus some address reclamation
- Growth is occurring in all regions
  - While growth as seen in the routing system is strongest in Asia, the allocation growth is strongest in Europe.
- **IPv4 Address Report at Potaroo (as of April 10, 2007)** (http://www.potaroo.net/tools/ipv4/)
  - Projected IANA Unallocated Address Pool Exhaustion: **14-Jun-2011**
  - Projected RIR Unallocated Address Pool Exhaustion: **12-Jun-2012**

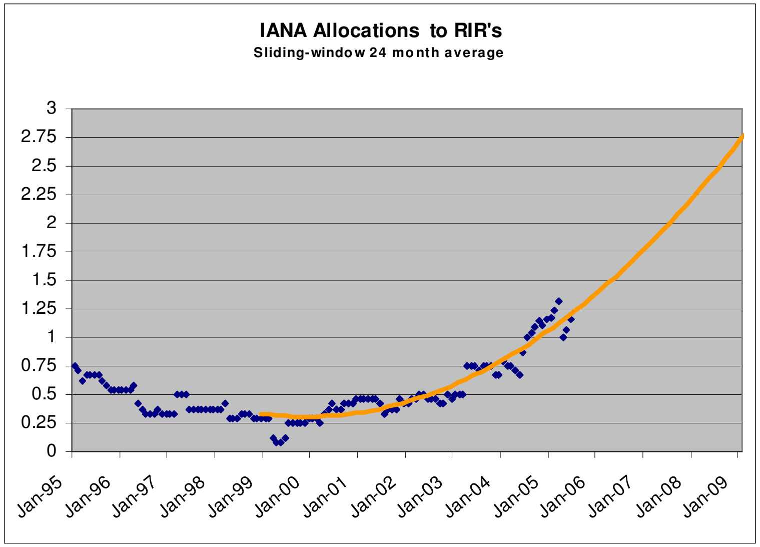

Projection based on IANA* data from 2000

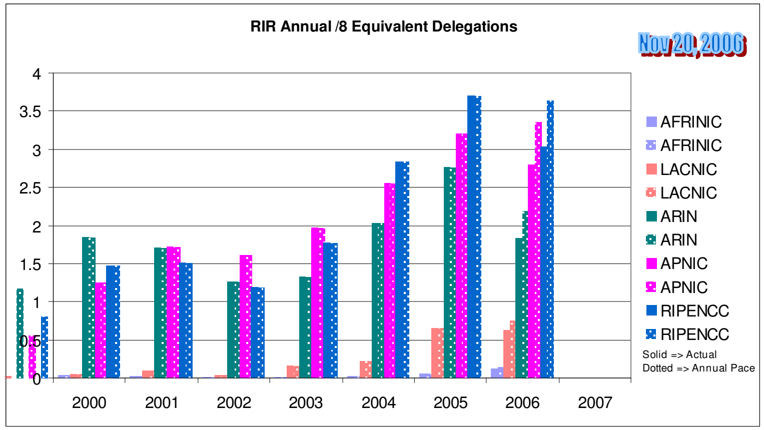

Update to: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipj_8-3.pdf
Exhaustion of the central IANA pool - orange
Exhaustion of the collective RIR pools - magenta
Relative distribution rates between the RIRs
Time depth of collective RIR pools on pub date - white
Time depth between exhaustion events - diff between orange & magenta

Tony Hain

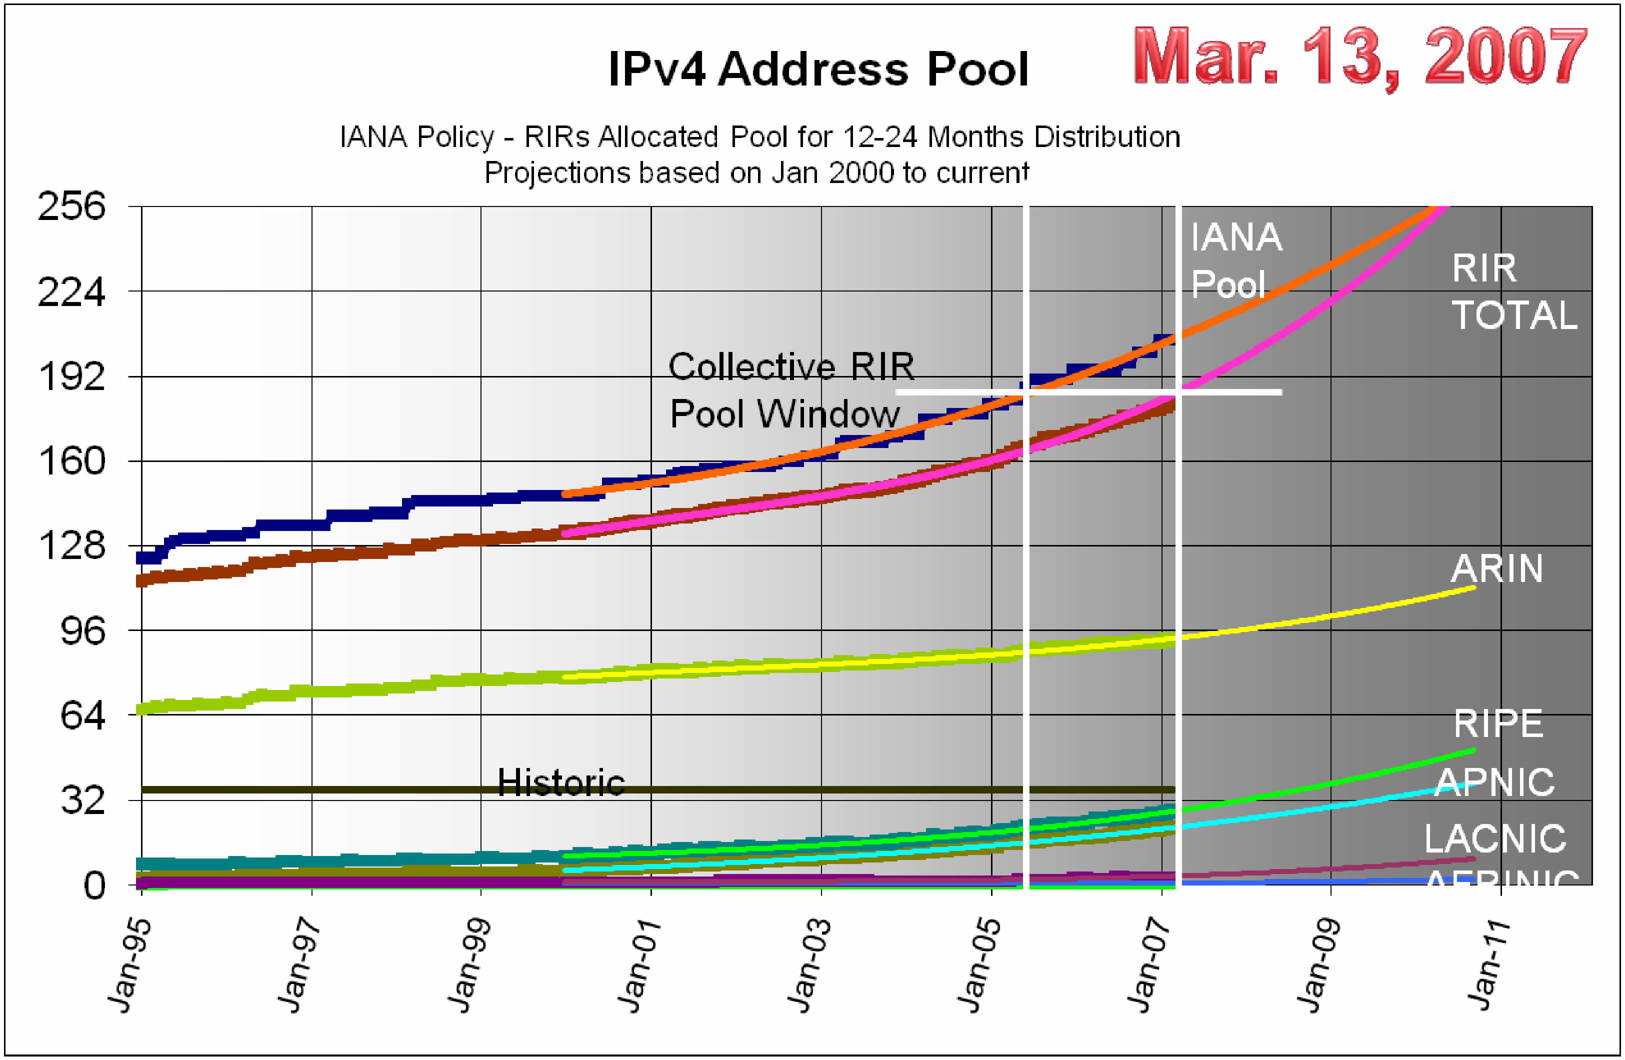

Implications

- Despite the wide-scale deployment of NAT, the consumption of the IPv4 pool continues at an accelerating rate.
- When IANA runs out, existing IPv4 networks still work.

  The only ones that will be immediately impacted are the RIRs when they come back for more space.
- When any RIR runs out, existing IPv4 networks still work.

  The only ones that will be immediately impacted are the LIR/ISP/Enterprise's when they come back for more space.
- When the LIR/ISP runs out, existing IPv4 networks still work.

  The only ones that will be immediately impacted are the people looking for more or new space.
- Any specific network will only *need* IPv6 when they attempt to talk to someone that was unable to acquire enough IPv4 space, or attempt to expand or add new applications and find themselves *unable* to get enough IPv4 space.

IPv6 Deployment Strategies for ISPs

	Environment	*Scenario*	*Solutions availability*
Access	Few customers, no native IPv6 service form the PoP or Data link is not (yet) native IPv6 capable, ie: Cable Docsis	Tunnels	Yes
	Native IPv4-IPv6 services between aggregation and end-users	Dual Stack	Yes
	Dedicated circuits – IPv4 – IPv6	Dual Stack	Yes
Core	Native IP – Core is IPv6 aware	Dual Stack	Yes
	MPLS – Core is IPv6 unaware	6PE/6VPE	Yes

IPv6 Over MPLS Networks Design Considerations

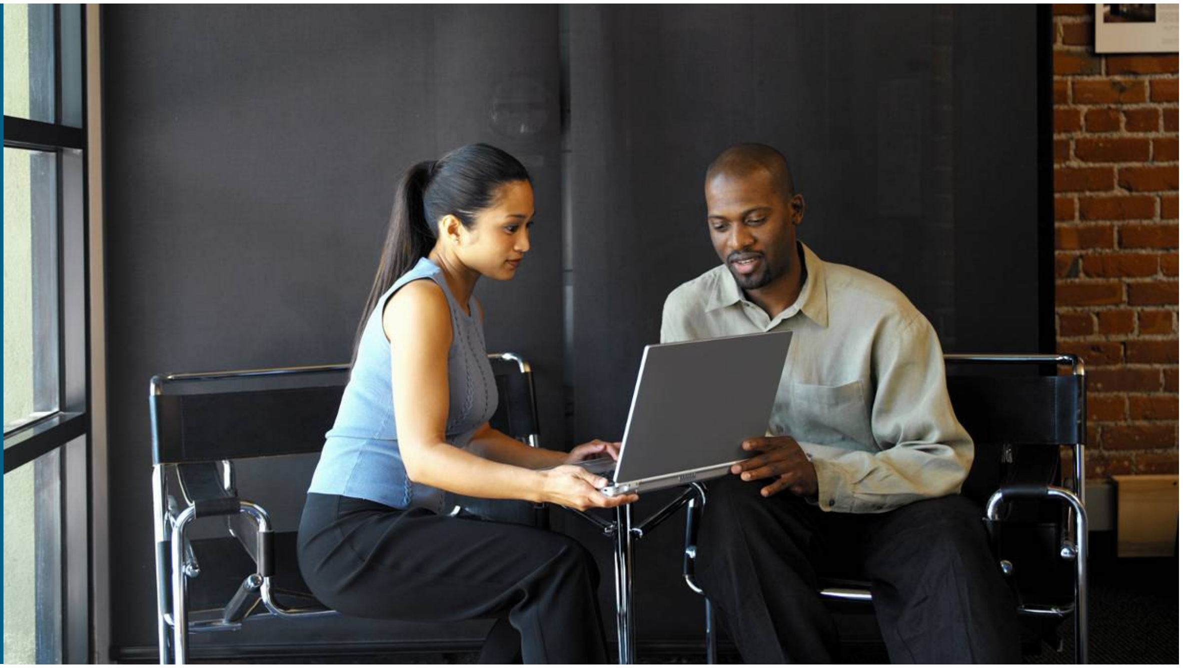

Core Design

Edge design

RR design

Address design

Lessons learned

Building IPv6 Services over MPLS Core

- Today most SPs are running MPLS in their backbones to provide multi-services
- Dual-stack 6PE/6VPE technologies provide easy transition to enable IPv6/IPv6 VPN services without major surgeries in the IPv4 MPLS core

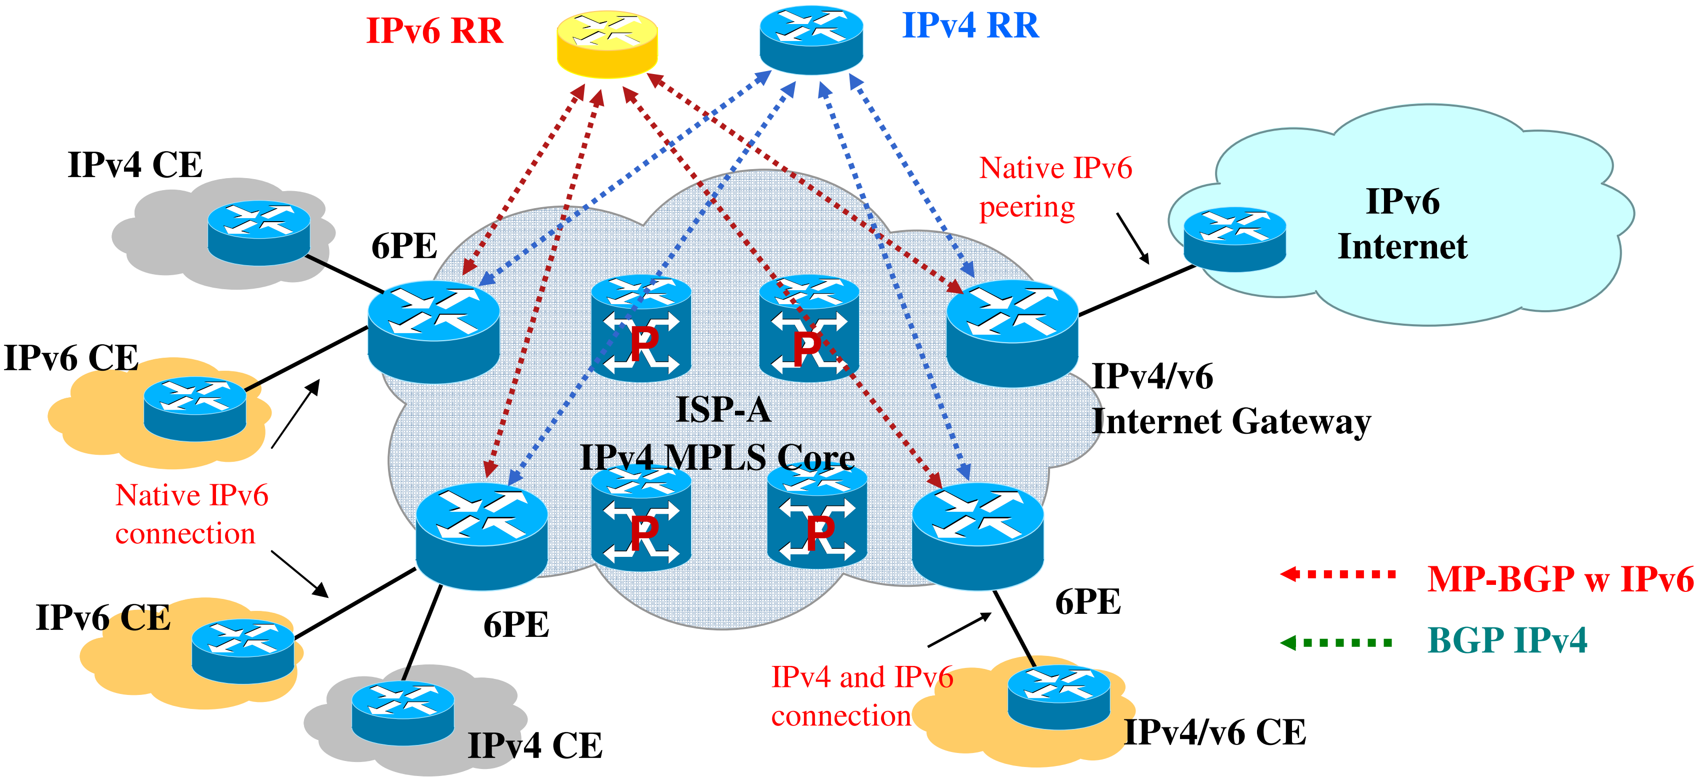

Network Core: Design Considerations

IGP

- Continue run IS-IS or OSPF for enabling IPv4 and VPNv4
- Edges peer using their respective loopback addresses for v4 and will use the same for v6
- New IPv6 RR or PEs will add new entries in the core tables

MPLS

- Continue run LDP or RSVP-TE for enabling IPv4 and VPNv4
- LSPs between loopback addresses
- 6PE and 6VPE will share the same LSPs as IPv4 PE
- No configuration change required

Edge: PE Router Design and Implementation Considerations

- PEs are dual-stack routers
- Core facing interface unchanged
  - IPv4 only
  - IGP is IS-IS or OSPF
  - Label protocol is LDP
- Significant changes in other areas of the PE:
  - IPv6 enabled on CE-facing interfaces
  - IPv6 routing protocol configured on CE-facing interfaces
  - BGP IPv6/VPNv6 peering with remote PEs
  - QoS configuration for both IPv4 and IPv6 if required
  - Management configuration for IPv6
  - Security configurations for IPv6
- PE-CE Design
  - Most PEs (except new IPv6 RRs) are existing IPv4 PEs
  - Assign an IPv6 address on the interface
    - e.g. Use link-local for peering, though this requires to add global address for management purpose

Edge: PE-PE Routing Design

- 6PE (IPv6) PE-PE routing design is similar to IPv4 PE-PE routing design
  - ➢ Same set of routers (Internet Access PEs)
  - ➢ Same protocol (iBGP), different AF (IPv6+label)
  - ➢ Same peering addresses (IPv4 loopback)
  - ➢ Dedicated Route-Reflectors (it is not required, but easy to start)
- 6VPE (IPv6 VPN) PE-PE routing design is similar to VPNv4 PE-PE routing design
  - ➢ Same set of routers (VPN Access PEs)
  - ➢ Same protocol (iBGP), different AF (VPNv6)
  - ➢ Same peering addresses (IPv4 loopback)
  - ➢ Dedicated Route-Reflectors (it is not required, but easy to start)

Route-Reflector Design

- RR are used to scale IPv6 and VPNv6 services
- RR are not part of Label Switched Paths
- IPv6 RR does not need to be dedicated boxes, but using dedicated IPv6 RR can reduce the complexity and risks for the initial deployment
- RRs peers together in a full-mesh topology (among Intra-AS RRs)

VRF Design

All Basic Steps for the VRF Design Have already Been Made for the VPNv4 Service:

1. Definition and configuration of VRF
2. Definition and configuration of RD
3. Definition and configuration of routing policies (import/export)
4. Interaction with the backbone control plane
5. Configuration of CE-facing interfaces
6. QoS policies
7. Existent VRFs would likely need to be modified to the multi-protocol VRF syntax

Address Design
- Address Allocation Model for Aggregation

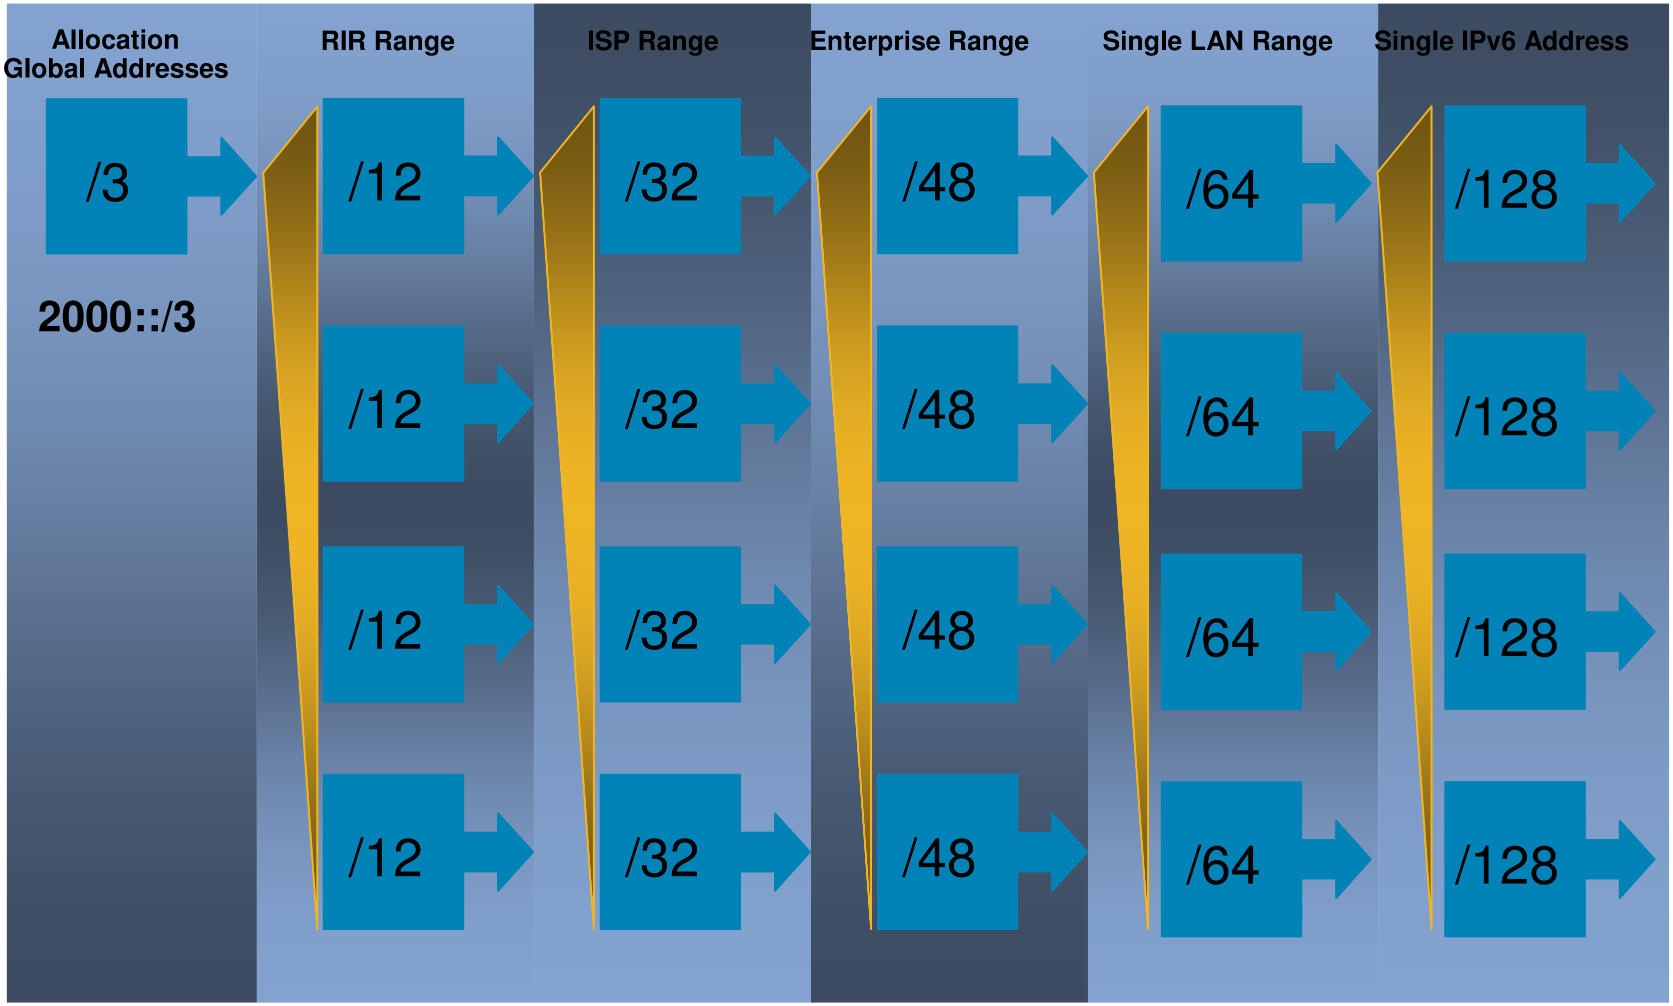

Network Level Address Design Considerations

Address Space Distribution

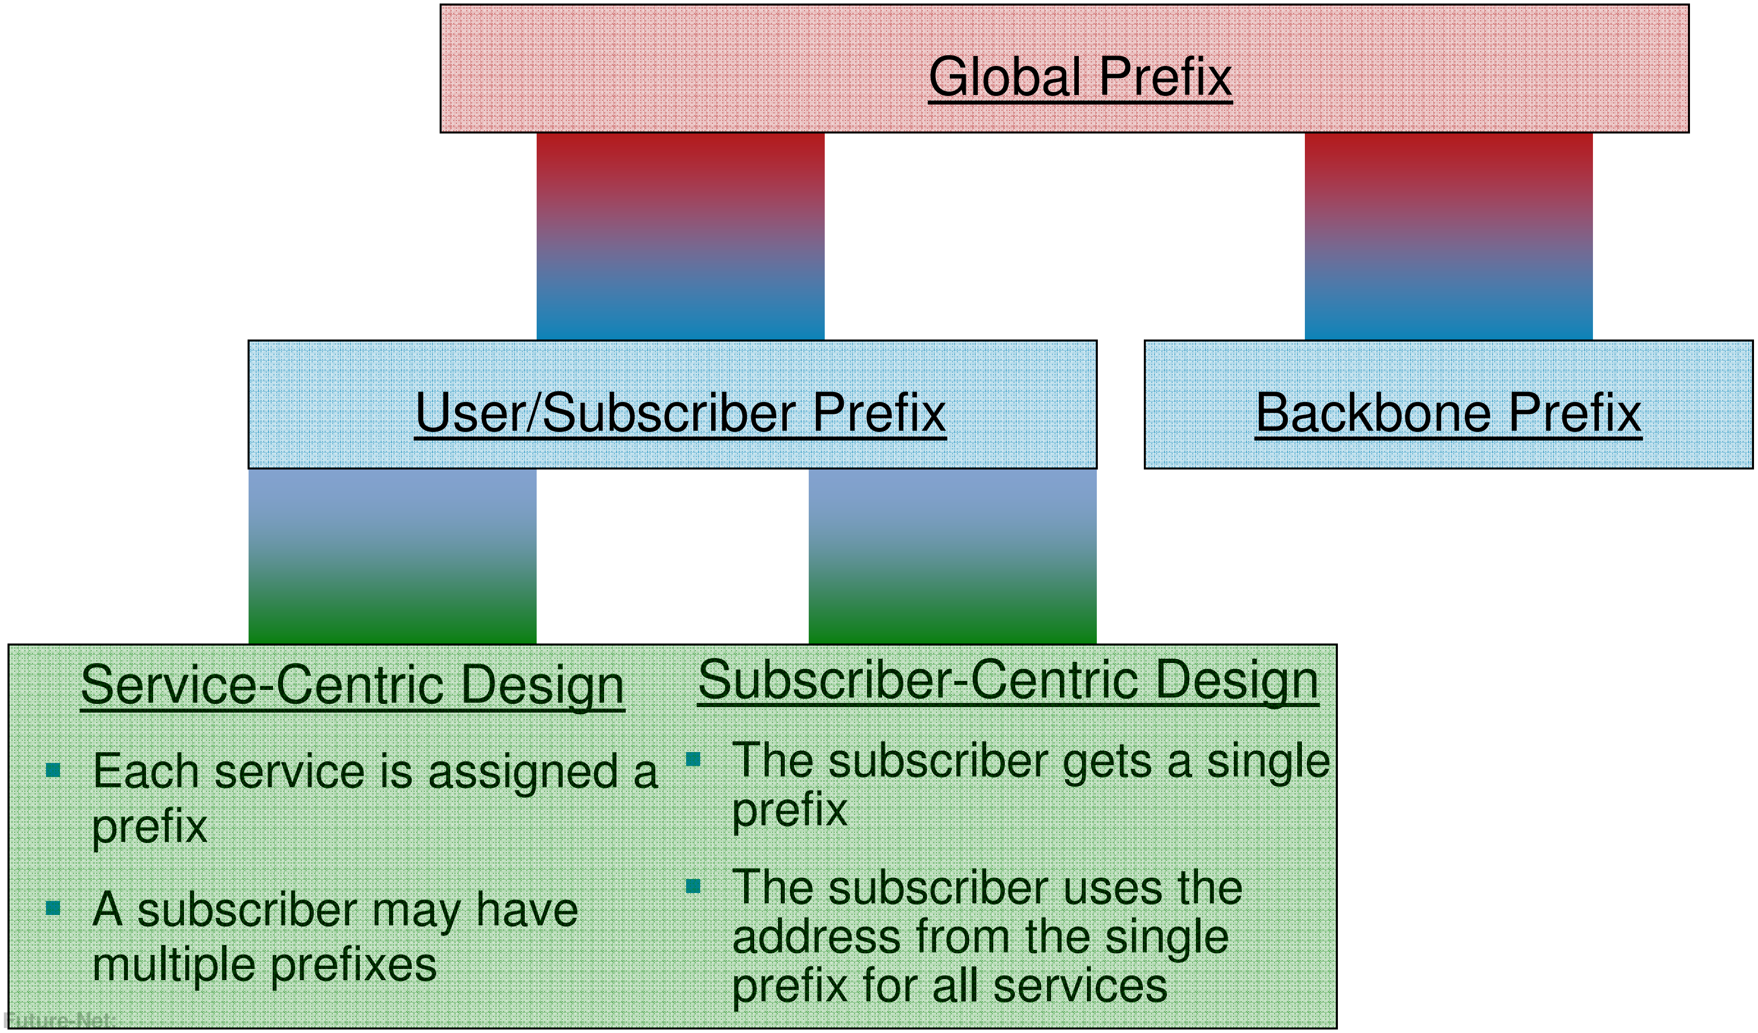

Global Prefix

User/Subscriber Prefix

Backbone Prefix

Service-Centric Design

- Each service is assigned a prefix
- A subscriber may have multiple prefixes

Subscriber-Centric Design

- The subscriber gets a single prefix
- The subscriber uses the address from the single prefix for all services

Network Level Address Design Considerations

- Address Scheme Design Drivers

  Geographical Boundaries - assign a common prefix to all subnets within a geographical area

  Organizational Boundaries - assign a common prefix to an entire organization or group within a corporate infrastructure

  Service Type - reserve certain prefixes for predefined services such as: VoIP, Content Distribution, wireless services, Internet Access, etc.

- Things to Remember

  Prefix Aggregation – to reduce the size of the routing table. Self imposed aggregation is important. It should also be expected to become a constraint imposed by peers (ex: SPs and size of VRF routing tables)

  Plan for Network Growth – reserve resources for growth (RFC 3531)

  Conservation - HD value for IPv6 is 0.94 compared to the current value of 0.96 for IPv4

  Note: IPv6 HD is calculated for sites (i.e. on a basis of /48), instead of based on addresses like with IPv4

Dual Stack IPv4/v6 VPN Deployment Scenario

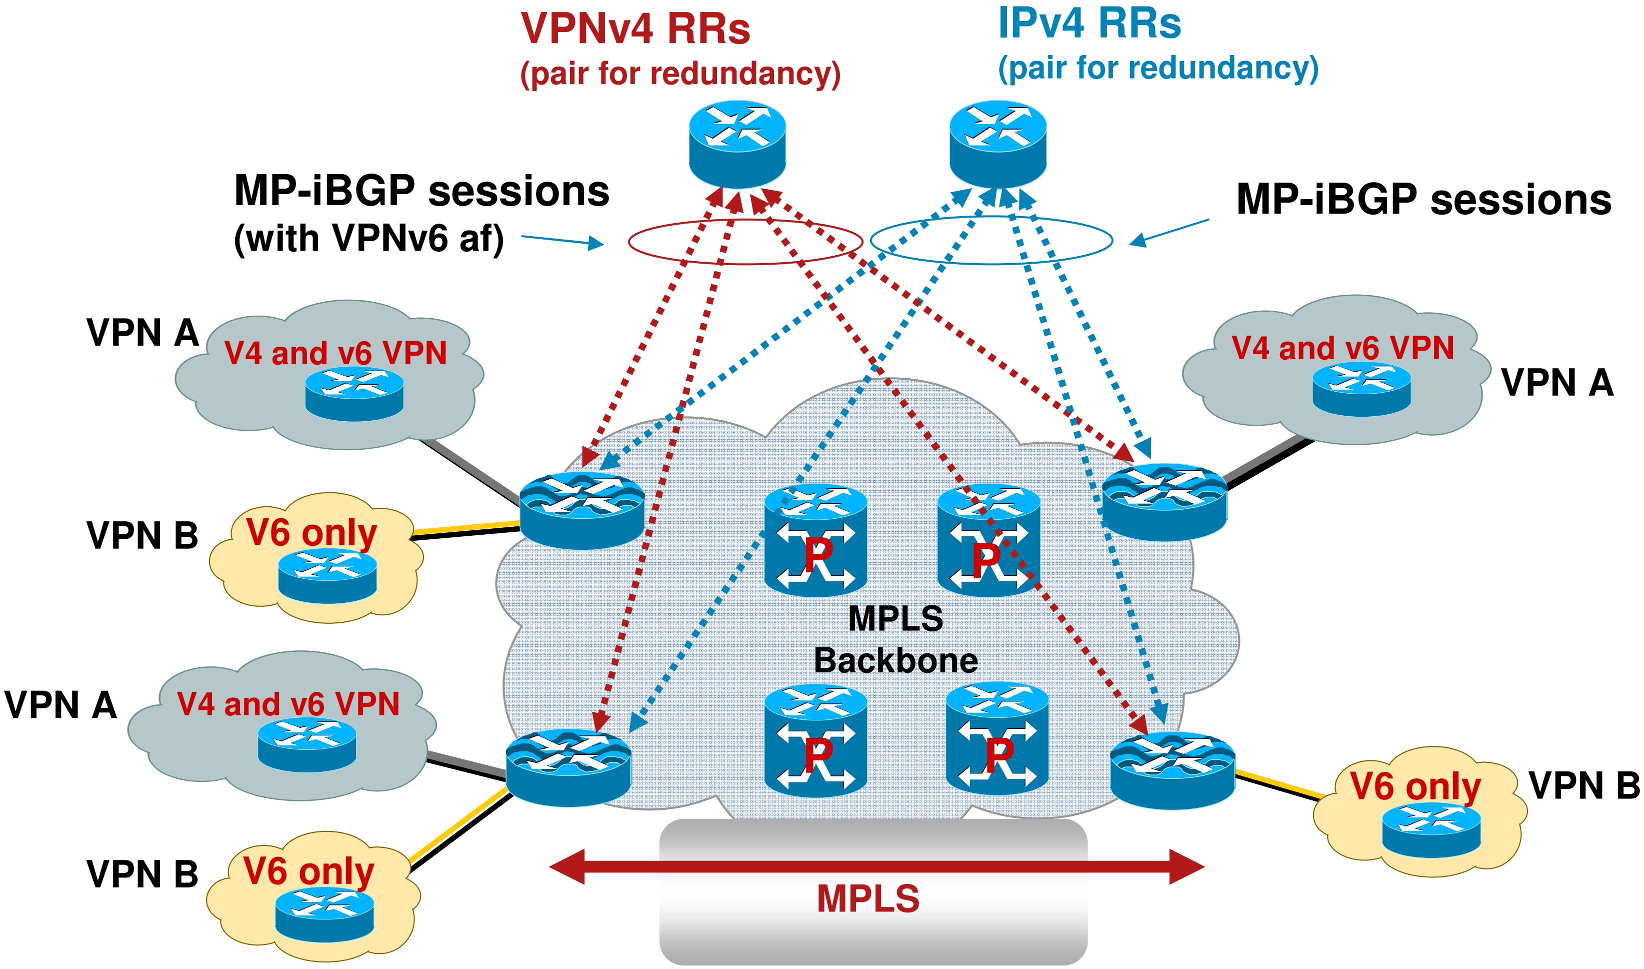

Design Lessons

- When mapped onto existing IPv4 and VPNv4 MPLS services, 6PE and 6VPE offer a low-cost / low-risk deployment strategy
- Link-local peering for eBGP PE-CE session is a useful and safe approach that simplifies the addressing plan
- RR design for IPv6 is strictly identical as IPv4 RR design. Dedicated IPv6 RRs is not required, but minimizes deployment risks
- Whichever QoS mechanism is implemented in the core, and on the edges, it is far easier if it does not differentiate between IPv4 and IPv6. QoS MQC commands such as match precedence are useful because they apply to both protocols
- As IPv6 requirement, MTU MUST be greater or equal to 1280 in the core, the non-IPv6 aware core is likely to black-hole traffic if this requirement is not satisfied. In general, it is better to push MTU negotiation at the PE because even if the Ps understand to reply.
- Feature parity is essential for IPv6/VPNv6 services to be consistent with the existing IPv4/VPNv4 services.

IPv6 Deployment Challenges

Core Design

Edge Design

RR Design

Design Lessons

Scaling IPv4 and IPv6 VPN routes

- VPNv4 routes well passed 1 million for some SPs
- Experiences with VPNv4

	MPLS L3 VPN	Internet
Routes per port	Ave. 2-3	Ave. <3
Routes per customer	Ave. ~300	Ave. ~7
Growth rate	>> 2000/month	~ 2000/month

Note: the recent IPv4 route growth is accelerated, observed 25K – 50K/year growth rate.

- VPNv6 route consumes more memory than VPNv4 route
- Prefix limit should be used as in VPNv4 to encourage aggregation
- Large VPN customers can have thousands of sites
  - 10K sites will result in 20K + VPN routes
- PE shared with Internet and VPN suppose are facing more pressure from the Internet as well as VPN route growth
- More scalable control plane implementation and well planed deployment are needed

QoS and Performance

- QoS
  - QoS should be supported in the same fashion for IPv4 and IPv6
  - IPv4 and IPv6 on dual stack port should support separate and aggregation limit
- Performance
  1. At the very minimum - No performance impact to the existing IPv4 services when implementing IPv6 on the existing PE routers.
  2. No performance degradation on both IPv4/VPNv4 and IPv6/VPNv6 traffic when IPv6/VPNv6 are offered as standard services.
  3. No performance degradation to the dual stack PEs in the presence of advanced features such as ACL, uRPF, etc.

Security Considerations - Common for IPv4 and IPv6

- Alter/disable TTL propagation (core protection - RFC 4111) at the PE

  Make the backbone look like one hop from the outside, Prevent the backbone addresses from being exposed through trace route

  Prevent TTL expiry packets cause ICMP time exceeded replies to consume line-card CPU

- Using ACL against infrastructure attacks

  Control plane protection/policing to protect route processor/routing engine

  Distributed line-card protection

- QoS pollution control

  QoS (MPLS EXP) re-coloring to prevent illegitimate traffic from impacting high priority traffic within the backbone

- eBGP security

  Protect against disruption, redirection of traffic flow

  Route filtering, dampening, maxas-limit, and MD5

  Route limit for VPN

  Control Plane TTL Sanity Check (RFC 3682, GTSM) **-** TTL check on BGP peering packets can effectively block all non-directed BGP spoofing

Attack Vector – Target SP’s Router

1. **Do Not advertise P or PE router’s Loopback**
2. **Do Not advertise PE-CE block to the IGP (Aggregated Null 0 on the Border will drop the attack.)**
3. **iACLs on the Border Router as a backup.**

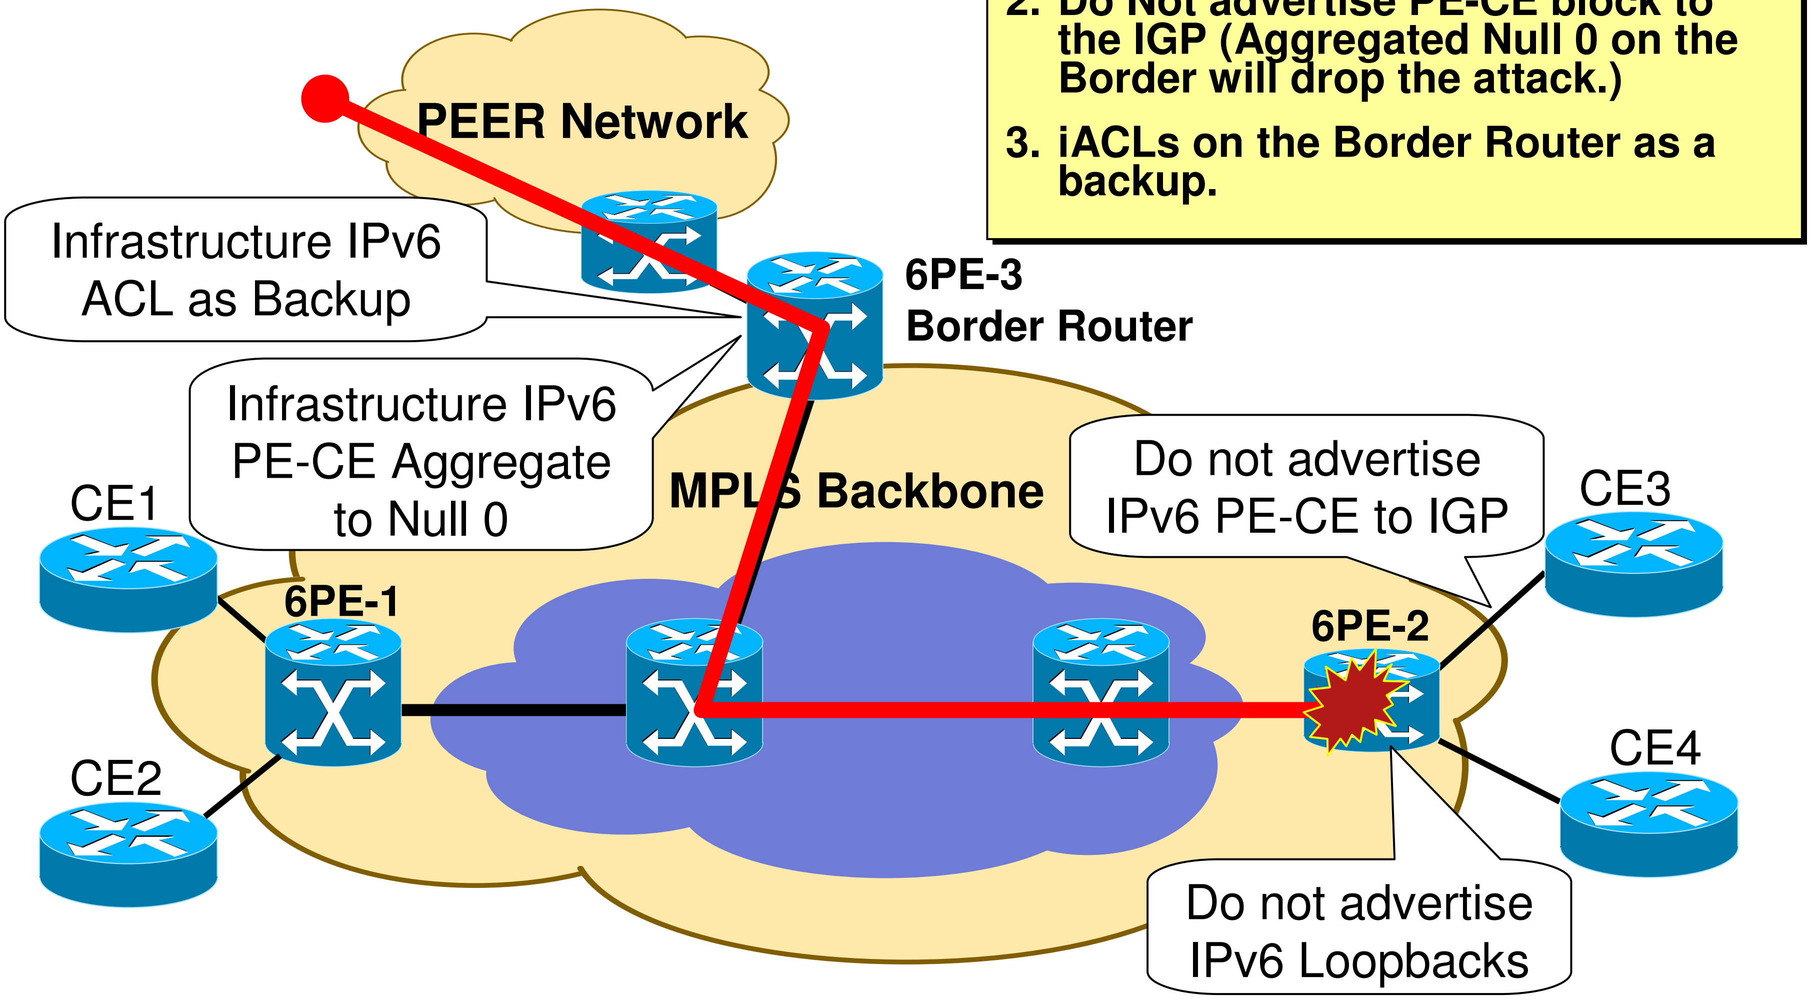

Infrastructure IPv6 ACL as Backup

Infrastructure IPv6 PE-CE Aggregate to Null 0

Do not advertise IPv6 PE-CE to IGP

Do not advertise IPv6 Loopbacks

Attack Vector – Target SP's Router

1. **Failsafe: rACL/CoPP – Point protection on the router in case the primary defenses fail.**
2. **Failsafe: Punt Path Queues. In case all defenses fail – containing the problem to the link and/or the line card.**

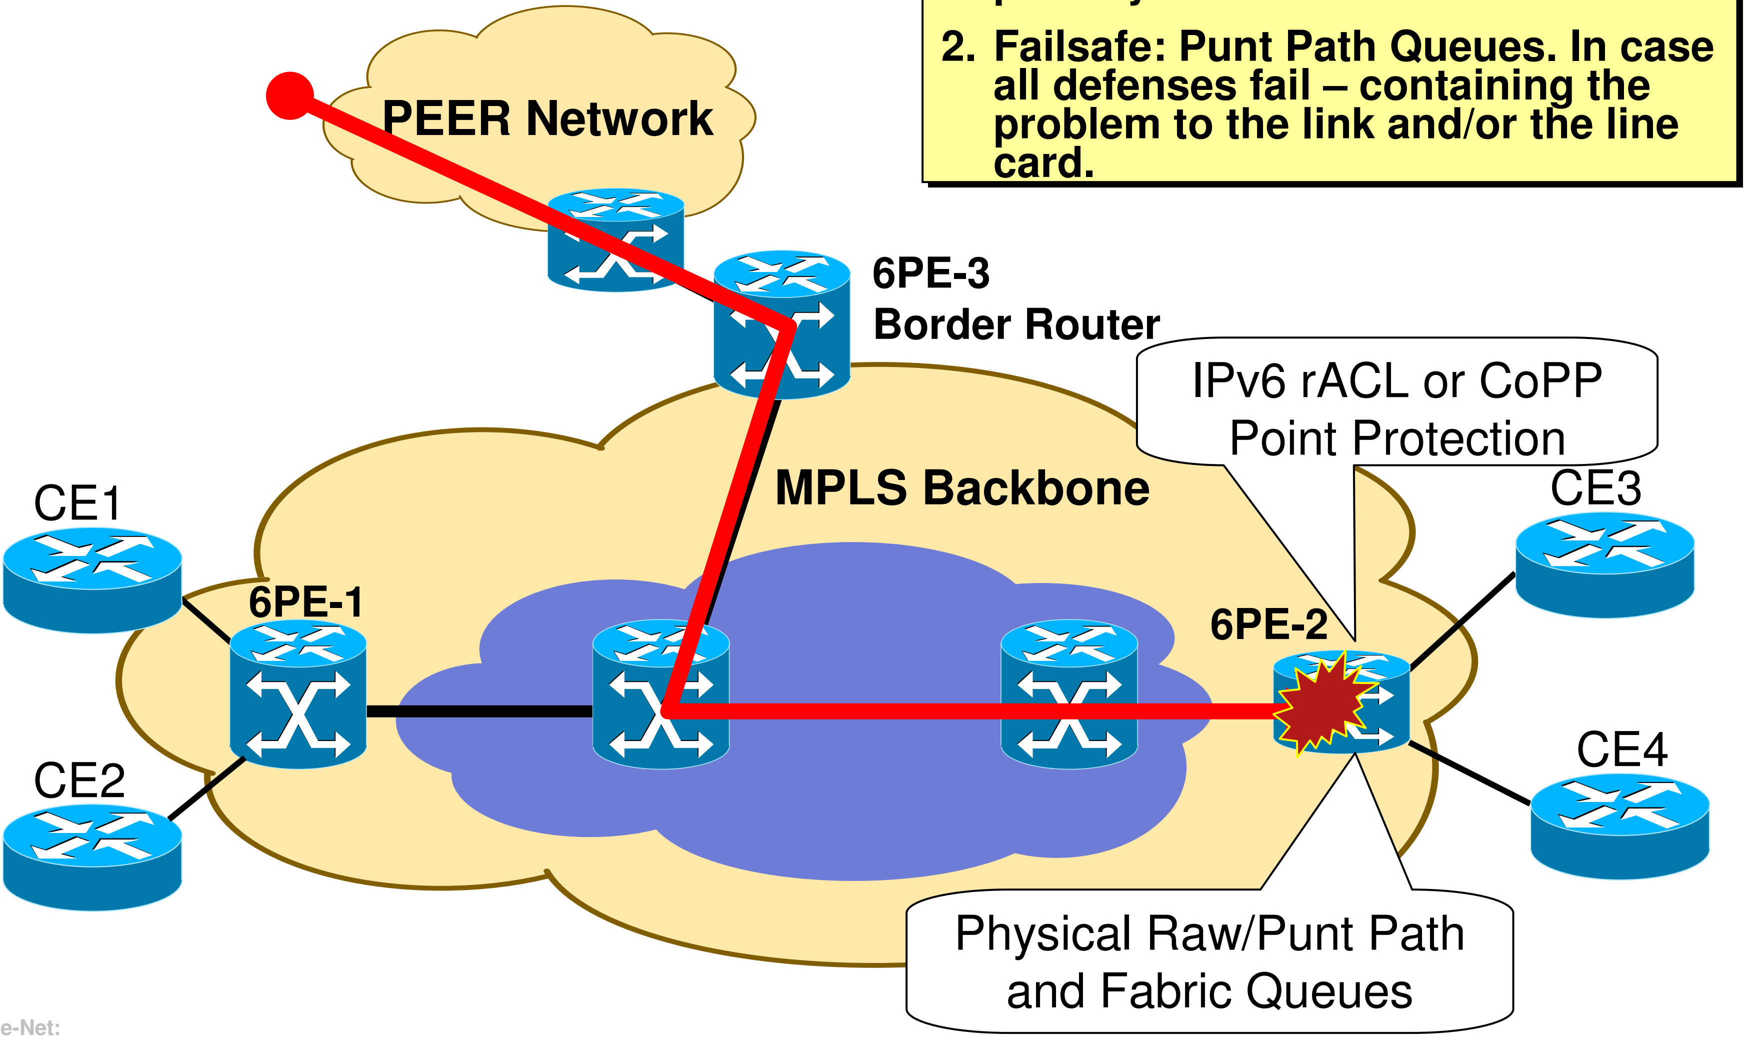

Security Considerations – IPv6 Specific (1) Extension Headers

Order	Header Type	Next Header Code	Use
1	Basic IPv6 Header	-	-
2	Hop-by-Hop Options	0	Jumbograms, MLD (Multicast Listener Discovery Protocol)
3	Destination Options (with Routing Options)	60	Destination
4	Routing Header	43	Source Routing, Mobility
5	Fragment Header	44	Fragmentation
6	Authentication Header	51	
7	Encapsulation Security Payload Header	50	
8	Destination Options	60	Destination
9	Mobility Header	135	
	No next header	59	

Security Considerations – IPv6 Specific (2) Headers

- Headers:

  Extension Headers (EHs) is a major new security concern for IPv6 – you cannot ignore it!

  - EHs can be manipulated with context causing intensive processing by network elements
  - Header chain can be unlimited (per spec) – a large number of EHs can drain the resources of the routers/devices

  Main Header: Flow Label

- Mitigation

  Filter unnecessary Extension Headers

  Put knobs to limit the number of EH

  Understand the capabilities of the network elements and firewalls

Security Considerations – IPv6 Specific (3) L3/L4 Spoofing

Threats

- Address assignment and aggregation makes it easier to implement RFC2827 filtering.
- IPv6 offers more Interface ID options that cannot be filtered based on RFC2827.
- IPv6 address can be spoofed in the EH
- Layer 4 spoofing is the same.

Mitigation

- Implement RFC2827 based filtering.
- Implement uRPF.
- New effort in IETF – “SAVA”?

Security Considerations – IPv6 Specific (4) ICMPv6

- ICMPv6 filtering
  - A large number of functions, message types, and options
  - Security considerations
    - Denial of Service attacks
    - Probing
    - Redirection attacks
    - Renumbering attacks
    - Problems due to ICMPv6 transparency
  - ICMP filtering using IPv6 ACLs, e.g.
    - Rate-limit the number of ICMP error messages generated
    - Re-direct ACL to disable sending redirect packet
  - Best practice guidelines for Filtering ICMPv6 Messages:
    - <draft-ietf-v6ops-icmpv6-filtering-bcp-00.txt>

Summary

- Market drivers for IPv6

  IPv4 address depletion; government mandate; new services for content delivery; ……

- Deploy IPv6 and IPv6 VPN over MPLS networks

  Design considerations

  - Minimize IPv6 deployment impact to existing MPLS networks - 6PE/6PE solutions
  - P, PE, RR design – similarities and differences as IPv4 design
  - Feature parity is essential for IPv6/VPNv6 services

  Challenges

  - Scalability
  - Performance
  - MVPNv6
  - Security

- Next Steps

  Development

  - MVPNv6 development
  - Continuous improvement on scalability

  Design and deployment

  - Taking full advantages of IPv6, innovative design and applications
  - IPv6 addressing
  - IPv6 Multi-homing
  - Expect to see more commercial IPv6 offers/applications – business and consumer

IPv6 Books

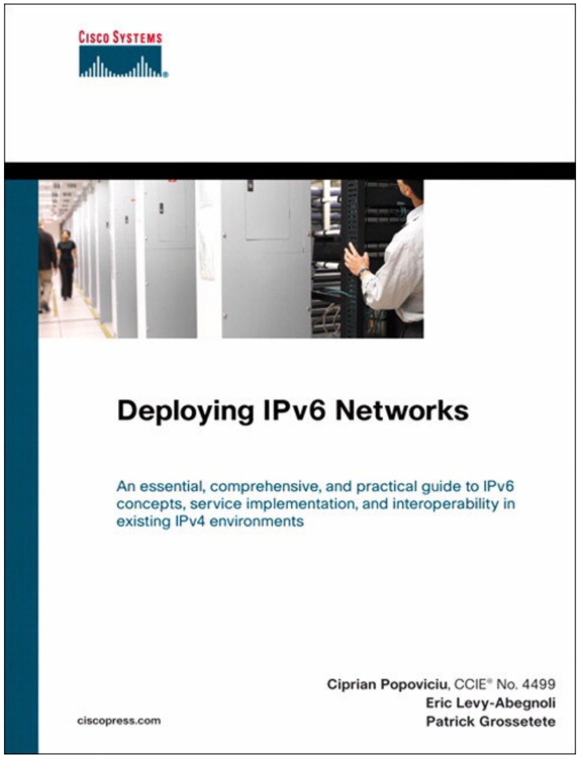

Cisco Systems

Deploying IPv6 Networks

An essential, comprehensive, and practical guide to IPv6 concepts, service implementation, and interoperability in existing IPv4 environments

ciscopress.com

Ciprian Popoviciu, CCIE® No. 4499
Eric Levy-Abegnoli
Patrick Grossetete

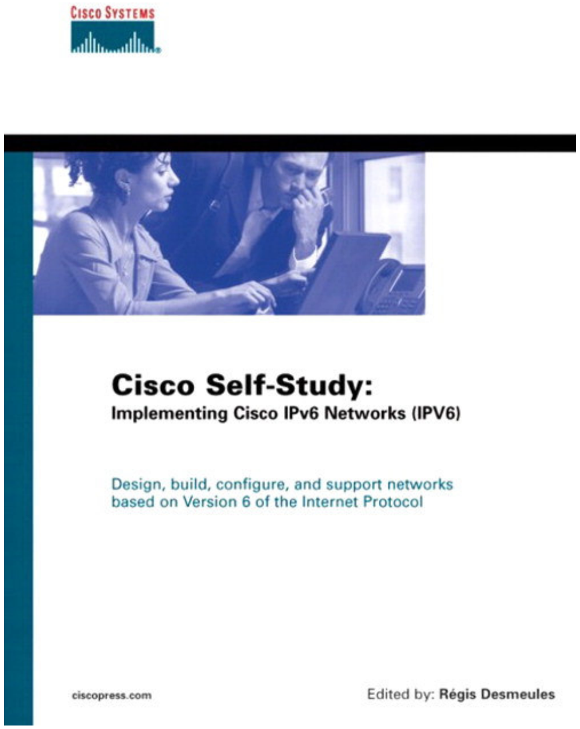

Cisco Systems

Cisco Self-Study:
Implementing Cisco IPv6 Networks (IPV6)

Design, build, configure, and support networks based on Version 6 of the Internet Protocol

ciscopress.com

Edited by: **Régis Desmeules**