Demystifying the Performance of IPv6 Routers

Jean-Marc Barozet

Consulting System Engineer

jbarozet@cisco.com

IPv6 - Key driver for Next Generation Ubiquitous Networking

U.S. DoD Memo Signed June 9, 2003

The DoD goal is to complete the transition to IPv6 for all inter and intra networking across the DoD by FY 2008. To enable this transition it is DoD policy for all Information Technology (IT) and National Security Systems (NSS) which make up the GIG (ref a) that:

- As of October 1, 2003, all GIG assets being developed, procured or acquired shall be IPv6 capable (in addition to maintaining interoperability with IPv4 systems/capabilities). This explicitly includes all acquisitions that reach Milestone C after October 1, 2003. The next version of the Joint Technical Architecture (JTA) will reflect this requirement.

Assistant Secretary of Defense
\- ***John Stenbit***

Broadband Home – A necessity for IPv6 !

So, A REAL Need for IPv6?

During the life cycle of a technology, a new product is often considered to have reached the early majority – or the mass market – after achieving 22 percent penetration.

- Internet population
  - ~945M by end CY 2004—only 10–15% of the total population
  - How to address the future Worldwide population? (~9B in CY 2050)
  - Emerging Internet countries need address space
- Mobile Internet introduces new generation of Internet devices
  - PDA (~20M in 2004), mobile phones (~1.5B in 2003), tablet PC
- Transportation—mobile networks
  - 1B automobiles forecast for 2008—begin now on vertical markets
  - Internet access on planes, e.g. Lufthansa—train, e.g. Narita express
- Consumer, home and industrial appliances

Microsoft Vista

- IPv6 activated by default on Vista

"In Windows Vista and Windows Server "Longhorn," IPv6 is installed and enabled by default. When both IPv4 and IPv6 are enabled on these OSs, the TCP/IP stack prefers to use IPv6 over IPv4."

- All Applications using the new peer-to-peer protocol will run on top of IPv6

"All applications using the Windows Peer to Peer Collaboration Foundation Technologies require IPv6 in this way."

- Windows Core Networking IPv6

  http://blogs.msdn.com/wndp/archive/category/14120.aspx

- Creating IP Agnostic Applications - Part 1

  http://blogs.msdn.com/wndp/archive/2006/08/29/Creating_IP_Agnostic_Applications__Part__1.aspx

ASPECTS OF IPv6 PERFORMANCE

Subtitle

Benchmarking the New IP for Successful Integration

- Deployment planning
  - Identify the architecture of the future network
  - Identity the supported IPv6 services.
  - This architecture translates into a set of functionality and performance requirements for each element of the network.
- Most of these requirements and their implications are well understood due to strong similarities with IPv4, but
  - There is a small but critically important subset that is IPv6 specific
  - It requires a good understanding of the new protocol set.
  - Insufficient coverage of IPv6 specific requirements could lead to operational challenges down the road.
- **Benchmarking IPv6 network element performance becomes an essential guide to requirements definition and to equipment evaluation.**

The Performance of Router Functions

- **Control Plane**
  - Routing Protocols
  - Network Management …
- **Data Plane**
  - Packet Forwarding …
- **Enhanced Services**
  - QoS
  - Tunneling,
  - ACLs, Extended ACLs
  - Encryption
  - Accounting …

Methodology Considerations

- RFC 2544 standardizes the key IP performance metrics and the methodology to measure them
  - Provides guidelines for defining requirements
  - Facilitates comparison of performance data between various platforms.
  - Highlight the importance of evaluating these benchmarks under relevant operational conditions, such as routers with traffic filters applied
- Benchmarks are clearly defined:
  - Throughput, Latency, Frame Loss Rate, System Recovery and Reset
- RFC 2544 is the de facto standard for IPv4 benchmarking.
- What about IPv6? Why Is RFC 2544 insufficient for benchmarking IPv6?

IPv6 Performance Aspects – 1

- RFC 2544 is mostly IP version agnostic
- Certain aspects of IPv6 must be taken into consideration when executing tests and interpreting the data
- The length of the address will impact the lookup speed
    - Address Lookup – 128 bits vs 32 bits
- The fixed 40 bytes long IPv6 header is 20 bytes longer than the typical IPv4 header.
    - Makes the IP packet per second (pps) throughput rates smaller for IPv6 than for IPv4 (IPv6 packets are longer)
    - Most evident at lower packet sizes, where the header represents a significant percentage of the total packet

IPv4 and IPv6 Header Comparison

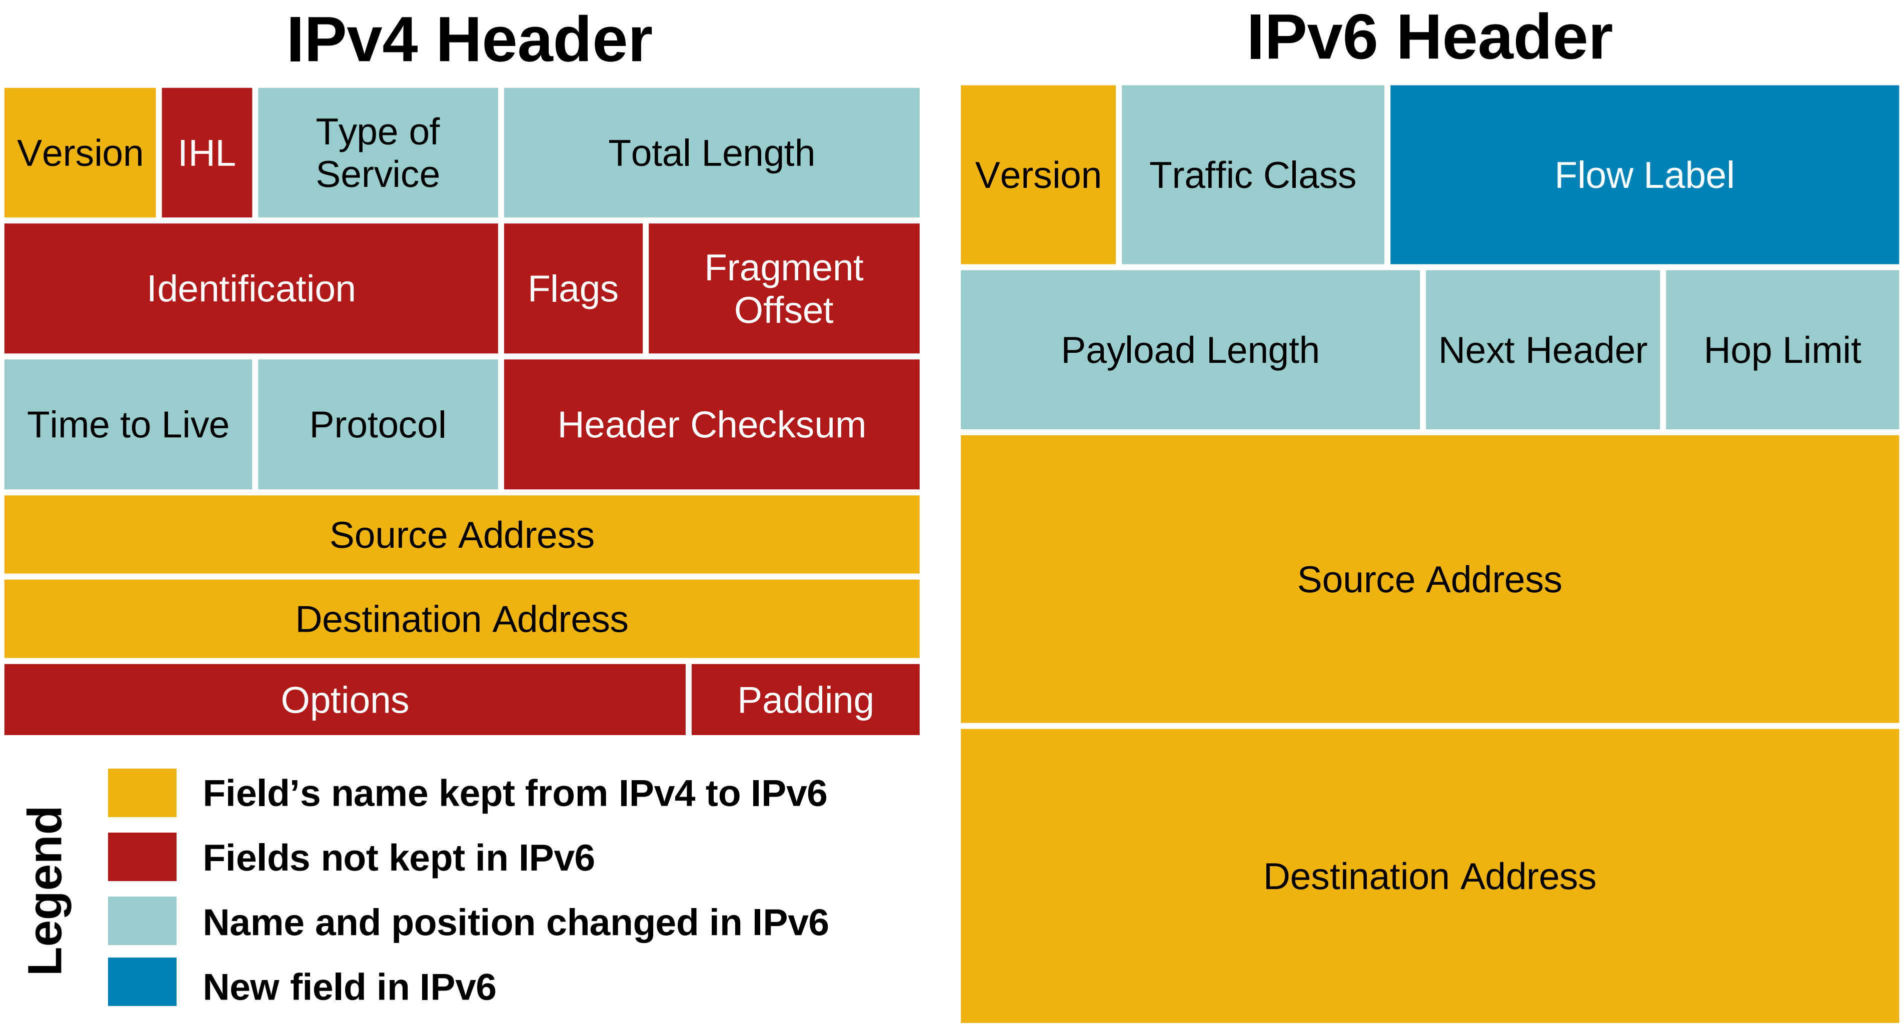

IPv6 Performance Aspects – 2

- More interesting and important, however, is to look at the less obvious aspects of IPv6
  - Aspects that could have a significant impact on performance.
  - A router's handling of these protocol features would indicate whether or not it was designed with IPv6 in mind.

IPv6 Significant Changes

- Packet header structure

  Along with the main header (commonly used in packet forwarding), a set of extension headers was defined that could carry, in a structured way, additional information

  These headers are a strong advantage for IPv6 (Provide protocol extensibility).

- Processing rules for extension headers are designed to improve forwarding

  however, under certain conditions, they can have an impact on performance.

- These are reasonable concerns, since extension headers are commonly used in cases such as: Fragmentation, Mobile IP and Authentication or Encryption of packets.

The Chain of Pointers Formed by the Next Header Field

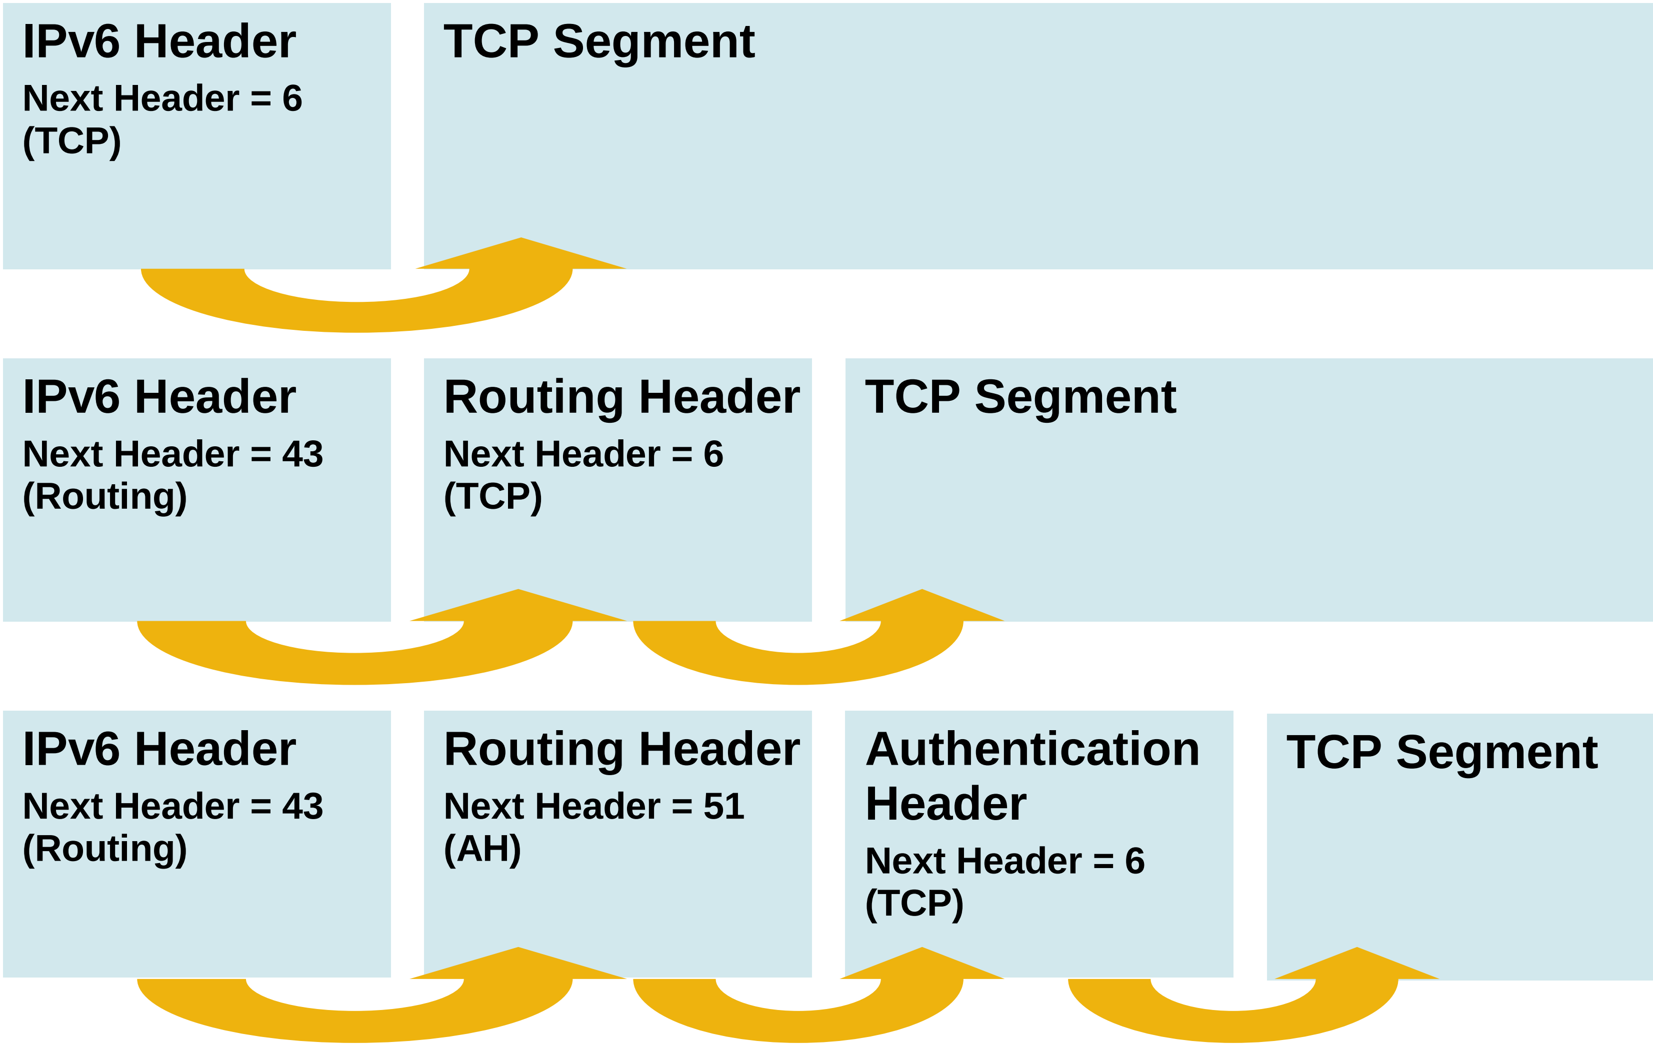

IPv6 Extension Header Types

Header Type	Next Header Value	Description
Hop-by-hop options header	0	Processed by all hops in the path of a packet, when present follows immediately after the basic IPv6 packet header
Destination option header	60	When the destination options header follows hop-by-hop options header, it is processed at the final destination and also at each visited address specified by the routing header. If it follows the Encapsulating Security Payload(ESP) header, it is processed only at the final destination.
Routing Header	43	Used for Source Routing
Fragment Header	44	Used by source when packet is fragmented , fragment header is used in each fragmented packet
Authentication Header (RFC 1826) and ESP Header (RFC 1827)	51	These are used within IP Security Protocol(IPSEC) to provide authentication,integrity and confidentiality of a packet. These headers are identical for IPv4 and IPv6
Upper-layer Header	6 (TCP) 17 (UDP)	These are the typical headers used inside a packet to transport data.

IPv6 Extension Headers Processing

- Routers will not process Extension Headers (EH) except for certain functions (support of Mobile IP)

    Only one extension header must be processed by each hop in the path of the packet, the Hop-by-Hop EH.

- The structure of the Hop-by-Hop header may vary

    Difficult to implement the processing of all its options in hardware

    Can have a performance impact on the router.

- Router's capabilities in processing Hop-by-Hop EH

    Tools to throttle traffic with this extension header type, which is legitimately used in support of Router Alert (for example in the case of Multicast Listener Discovery), and for RSVP or potentially IP Jumbograms, in case of data link layers supporting more than a 64K data payload.

IPv6 Extended ACLs

- One important case to consider

  Traffic with a chain of extension headers going through a router's interface that has packet filtering (access lists) applied to it.

- If upper layer information (TCP or UDP ports) is filtered

  The router must hop from one EH to the other until it gets to it - can impact the forwarding performance

- Network elements that were not designed with IPv6 in min

  will be unable to process the EH chain in hardware and push the traffic in the slow path in order to have the upper layer protocol information extracted

  or, even worse, they may have to drop the packet if unable to handle this case.

IPv6 Benchmarking

- A complete protocol benchmarking is essential to the success of IPv6 deployments.
  - We should not forget that IPv6 is likely to be deployed in existing operational infrastructures
  - so benchmarking its co-existence with IPv4 is equally important.
  - RFC 2544 remains the primary guideline for this process
- But, need for additional IPv6 specifics and co-existence test methodology
  - Work is currently being done on this topic within the IETF.
- The Benchmarking and IPv6 Operations Working Groups have contributed to “IPv6 Benchmarking Methodology”
  - **http://tools.ietf.org/wg/bmwg/draft-popoviciu-bmwg-ipv6benchmarking-02.txt**

MEASURED IPv6 PERFORMANCE

Subtitle

Industry's Broadest Platform Support

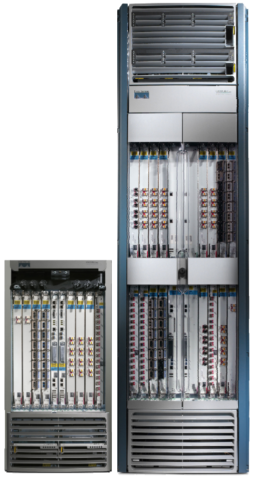

Cisco IOS 12.0S

Cisco 12000 Series Routers

Cisco 10720 Series

Cisco IOS 12.4/12.4T

Cisco 800 Series Routers

Cisco 1700 Series Routers

Cisco 1800 Series Routers

Cisco 2600 Series Routers

Cisco 2800 Series Routers

Cisco 3600 Series Routers

Cisco 3700 Series Routers

Cisco 3800 Series Routers

Cisco 7200 Series Routers

Cisco 7301 Series Routers

Cisco 7500 Series Routers

Cisco IOS-XR

CRS-1, Cisco 12000

Cisco IOS 12.2S & derivatives

Cisco 72/7300 Series Routers

Cisco 75/7600 Series Routers

Cisco 10000 Series Routers

Catalyst 3750/3560 Series

Catalyst 4500 Series

Catalyst 6500 Series

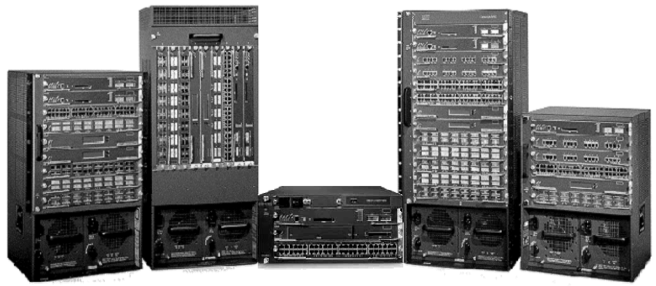

Cisco Product Portfolio

PIX Firewall (7.0), FWSM 3.1, LMS 2.5, MDS9500 series, CNR 6.2, NAM 3.4, NFC 5.x

Proven Performance
Catalyst 6500 Series

Verified by EANTC

Demonstrated 400 Mpps of IPv4

Demonstrated 200 Mpps of IPv6

Verified Interoperability and performance with previous generation modules

Multicast Scalability Over 400,000 mroutes

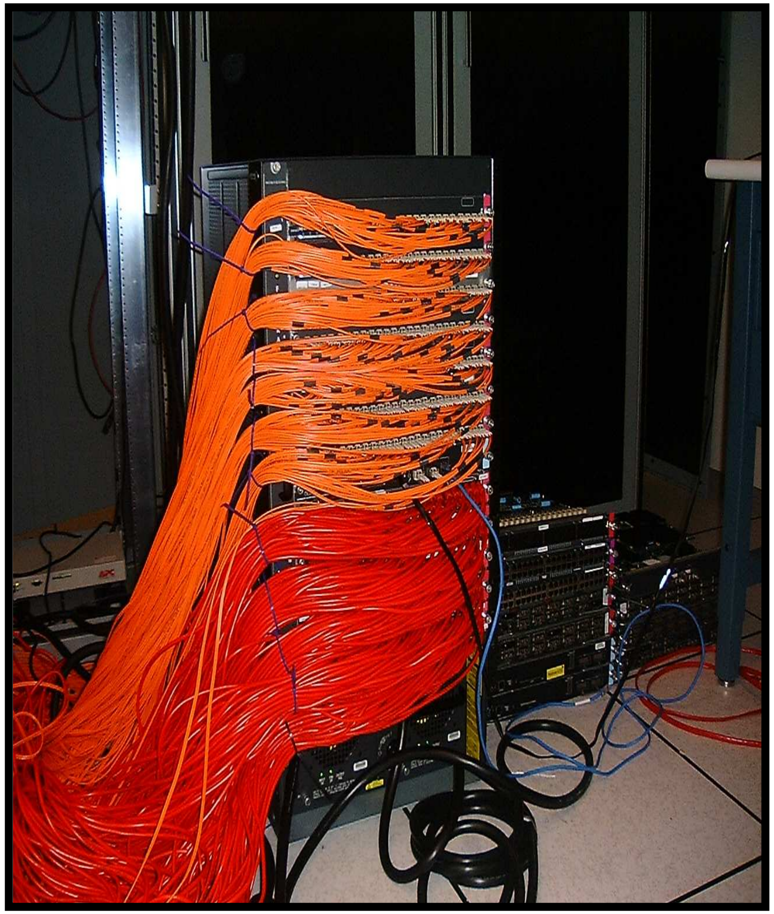

http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont_0900aecd800c9589.pdf
http://www.cisco.com/application/pdf/en/us/guest/products/ps708/c1244/cdccont_0900aecd800c958a.pdf

Proven Performance
Catalyst CEF720 Architecture

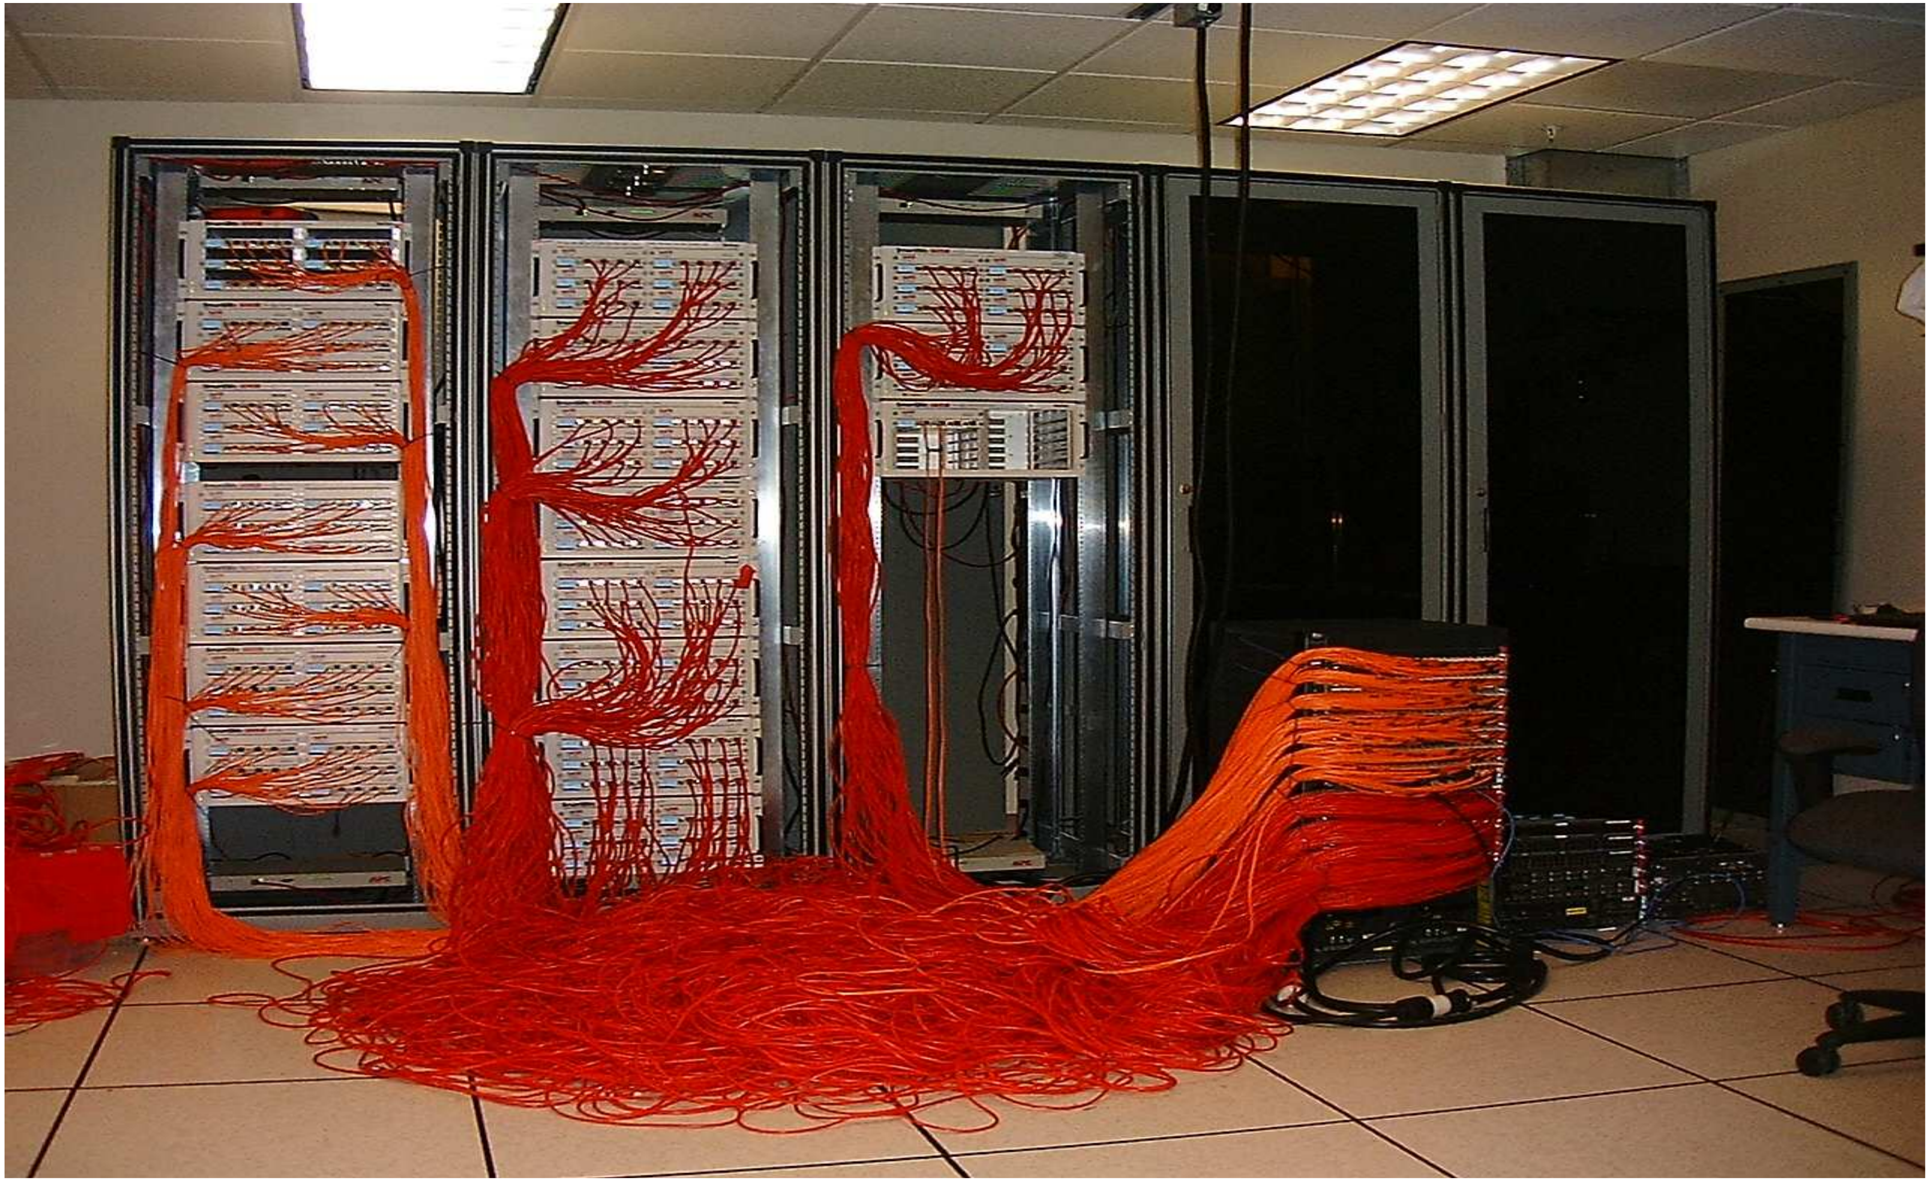

Proven Performance
Cisco 12400 XR Series

- Cisco XR 12000 Series – Service Separation Architecture Tests

  Tests conducted by EANTC

  MIX of IPv4 and IPv6 flows

- The test run

  with a 5,001 entry ACL for IPv4 and another 5,001 entry ACL for IPv6, where 5,000 entries are DENY and the last entry is a PERMIT-ALL did not show any IPv4/IPv6 packet loss at wire-speed IMIX load.

  Forcing the router to inspect the UDP header for access control list processing

Proven Performance
Cisco 12400 XR Series

Verified by EANTC

The tests confirmed the Cisco XR12000 operates with true separation between different entities of logical routers

The Cisco XR12000 demonstrated 100% forwarding rate and low latency for mixed IPv4 and IPv6 traffic with access control lists and logging for unauthorized traffic

- http://www.eantc.com/fileadmin/eantc/downloads/test_reports/2003-2005/EANTC-Summary-Report-Cisco-12kXR.FINAL.pdf

Proven Performance
Cisco CRS-1

- Light Reading, the leading telecom magazine, commissioned EANTC to verify the performance of the Cisco Carrier Routing System (CRS-1) using 10-gig and 40-gig interfaces with a mix of IPv4 and IPv6 flows as well as Services activated

  Tests conducted by EANTC

  First test of 40-Gbit/s Sonet/SDH interfaces

- The Reader's Digest

  “The CRS-1 performed extraordinarily well, demonstrating that it can scale to meet the requirements of service providers far into the future. It scaled to terabits-per-second of bandwidth, millions of routes, and tens of millions of IPv4 and IPv6 flows”

Proven Performance
Cisco CRS-1

Verified by EANTC

The CRS-1 clearly proved that it processes IPv6 completely in hardware

In our mixed scenario, the single-chassis system mastered a packet rate of 820 million pps at line rate

http://www.lightreading.com/document.asp?doc_id=63606

Forwarding Performances with Services

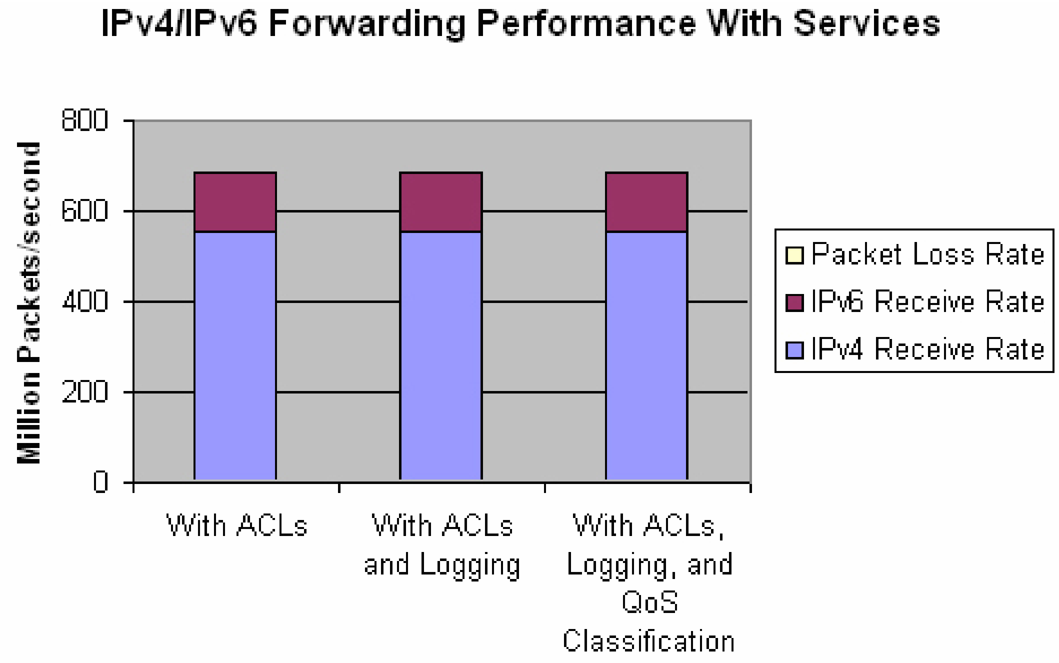

- The test run

  with a 5,001 entry ACL for IPv4 and another 5,001 entry ACL for IPv6, where 5,000 entries are DENY and the last entry is a PERMIT-ALL did not show any IPv4/IPv6 packet loss at wire-speed IMIX load.

  Forcing the router to inspect the UDP header for access control list processing

IPv6 DEPLOYMENT

Subtitle

IPv4–IPv6 Transition/Coexistence

- A wide range of techniques have been identified and implemented, basically falling into three categories:
  1. Dual-stack techniques, to allow IPv4 and IPv6 to co-exist in the same devices and networks
  2. Tunneling techniques, to avoid order dependencies when upgrading hosts, routers, or regions
  3. Translation techniques, to allow IPv6-only devices to communicate with IPv4-only devices
- Expect all of these to be used, in combination

IPv6 Deployment Scenario for Enterprises

	Environment	Scenario	Cisco IOS support
WAN	IPv6 services available from ISP	Dual Stack	Yes
	Dedicated Data Link layers, eg. LL, ATM & FR PVC, dWDM Lambda	Dual Stack	Yes
	No IPv6 services from ISP or experimentation – few sites	Configured Tunnels	Yes
	No IPv6 services from ISP or experimentation – many sites, any to any communication	6to4	Yes
Campus	L3 infrastructure – IPv6 capable	Dual Stack	Yes
	L3 infrastructure – not IPv6 capable, or sparse IPv6 hosts population	ISATAP	Yes

IPv6 Deployment Scenario for ISP

	Environment	Scenario	Cisco IOS support
Core	Core is IPv6 aware – Native IP	Dual Stack	Yes
	Core is IPv6 unaware – MPLS	6PE/6VPE	Yes
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

6PE Routing/Label Distribution

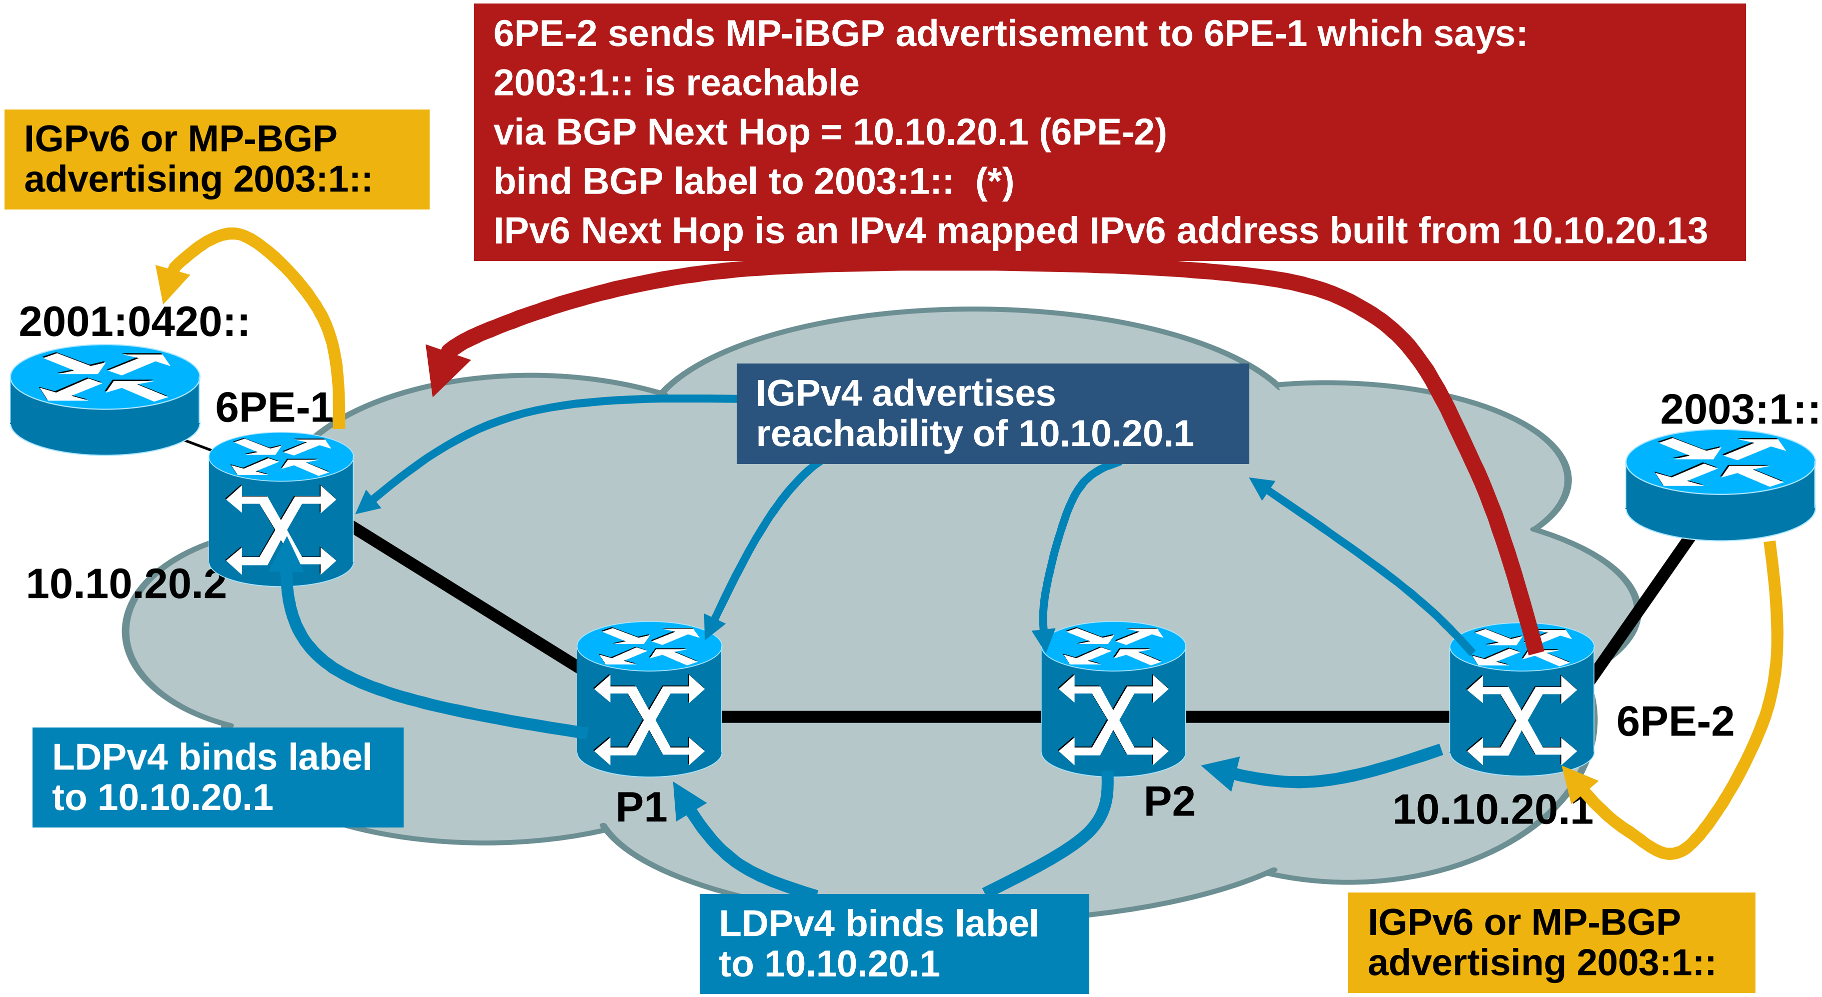

(*) The 2nd label allows operations with Penultimate Hop Popping (PHP)

IPv6 Integration on MPLS VPN infrastructure – 6VPE

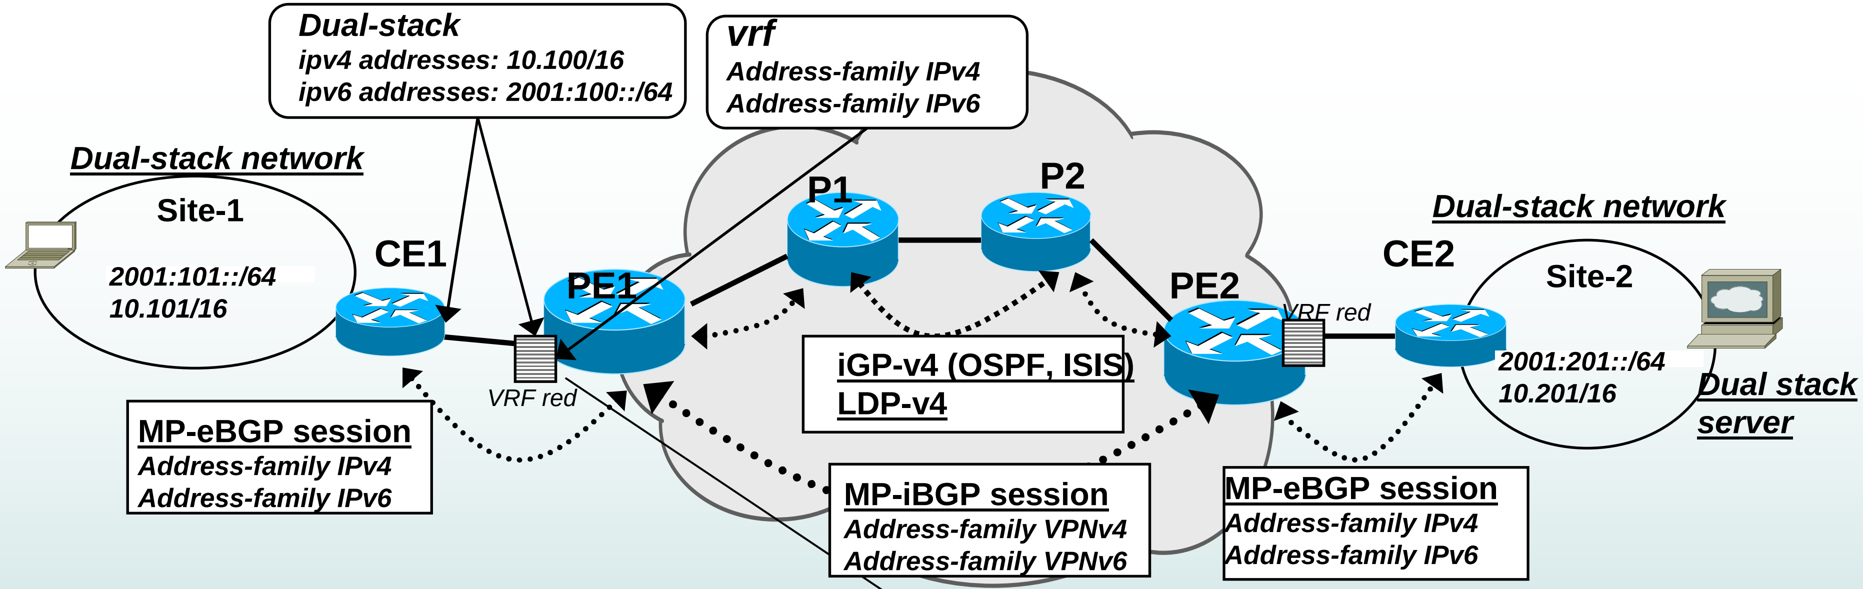


```
vrf definition site1
 rd 100:1
 route-target import 100:1
 route-target export 100:1
  address-family ipv4
  address-family ipv6
!
interface ethernet0/0
vrf forwarding site1
ip address 10.100.1.2 255.255.0.0
ipv6 address 2001:100::72b/64
```

- MPLS/IPv4 Core Infrastructure is IPv6-unaware
- PEs are updated to support Dual Stack/6VPE
- IPv6 VPN can co-exist with IPv4 VPN – same scope and policies
- 6VPE – RFC 4659 – Cisco authored for IPv6 VPN over MPLS/IPv4 infrastructure