IPv6 Management 101
Share Session Anaheim

Laura Knapp
WW Business Consultant
Laurak@aesclever.com

The Past

What network protocols did you run before 1990?

- Bisync
- IPX/SPX
- Appletalk
- NetBios
- DECnet
- XNS
- Others????

What network protocols did you run between 1990 and 2000?

- Above and introduced IPv4

What network protocols are you running now?

- IPv4
- IPv6????
- SNA over IPv4

Managing Fundamentals – Remains

- **FCAPS**
  - Fault
  - Configuration
  - Availability
  - Performance
  - Security
- **Leading to**
  - Service Level Achievement
  - Optimum Resource Utilization
  - Highly available systems
  - High performing systems

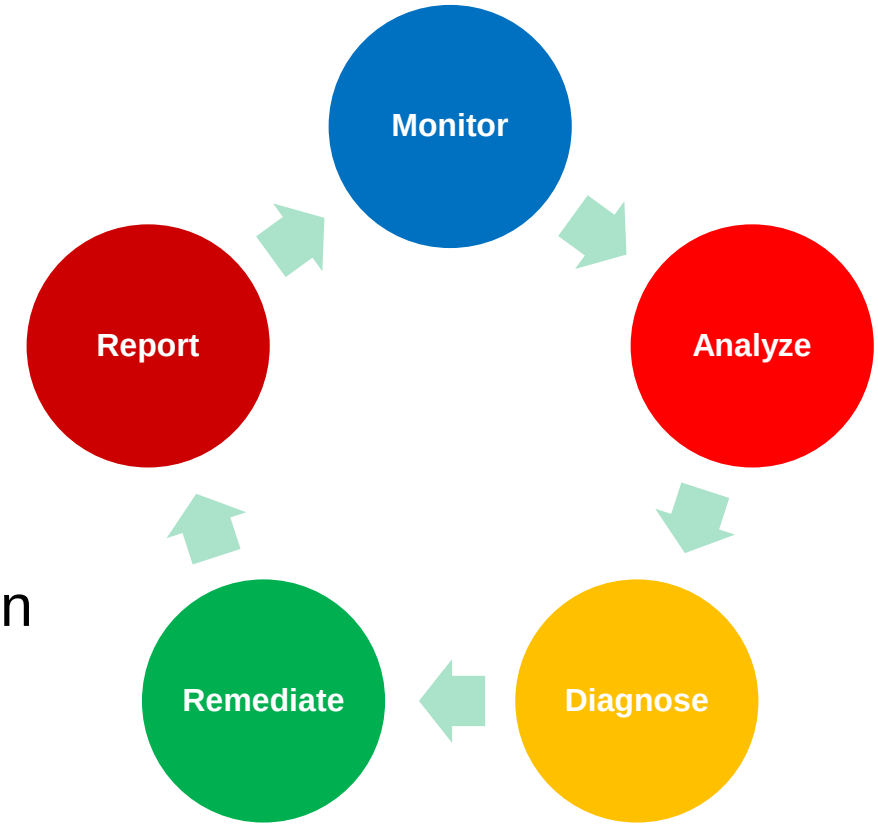

FCAPS

Fault Management

What is the Status?

Configuration Management

What is the configuration?

Availability Management

What's down? What's available?
What's up?

Performance Management

How consistent? How many?
How much? How fast?

Security Management

Who can access? Identify yourself?
Can everyone see it?

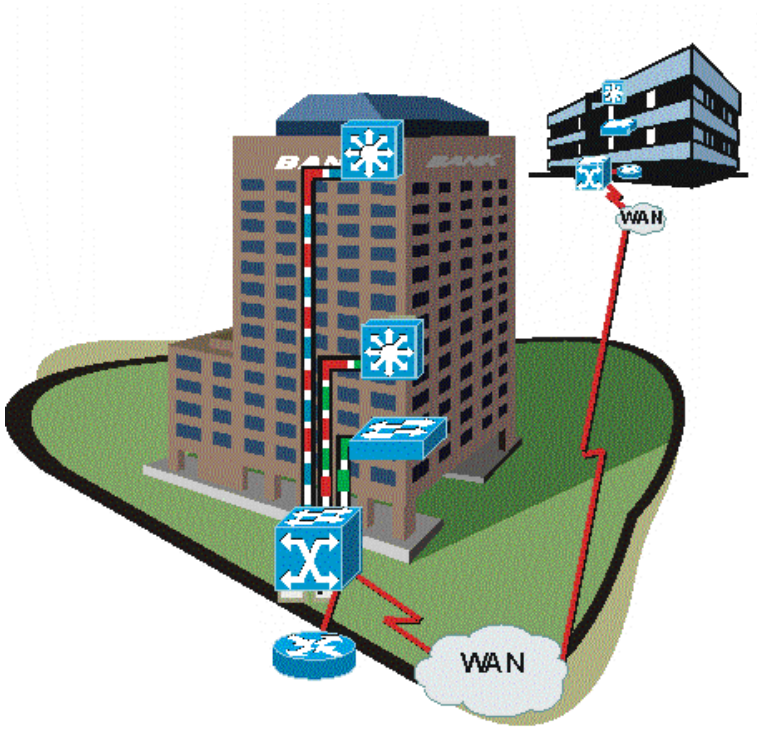

IPv6 the Good News

SNMP infrastructure

Old tools new features

Ping6, Tracert6, netsh,

IPv6 management data can be sent over IPv4 network

Existing management tools must:

- Support larger address space
- Provide screens to view new IPv6 management data
- Provide threshold settings for IPv6 management data
- Understand that an interface can have multiple addresses
- Understand that interfaces can have multiple default routes
- Changing an IP address no longer breaks a connection
- Provide details on new functions like Neighbor Discovery

IPv6 Management Overview

Previously (June '2011):

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP
Cisco										
Brocade										
Juniper										

Now:

	SSH HTTPS	DNS	Syslog	SNMP	NTP	RADIUS	Unified MIB RFC4293	Flow export	TFTP FTP	CDP LLDP	IPv6 MTU	No v4
Cisco[3]								6				
Brocade											9	1
Juniper								5				
ALU	2							4				
A10							8	7				

1. In FESX devices with v4 disabled, still does v4
2. ssh over IPv6 not supported until 10.0R1 (March 2012)
3. 15.2(2)TR
4. R10.4 July 2012
5. 12.3R1 Nov 2012 (beta in August)
6. ASR1K:3.7S (July 2012)
7. 3.0 release, 2012Q4
8. No plans
9. fixed in 7.3.0c (May 2012)

10-Apr-2012

9

Agenda

Introduction and goals

Management planning model

Practices for Management planning

Example

Recognize This?

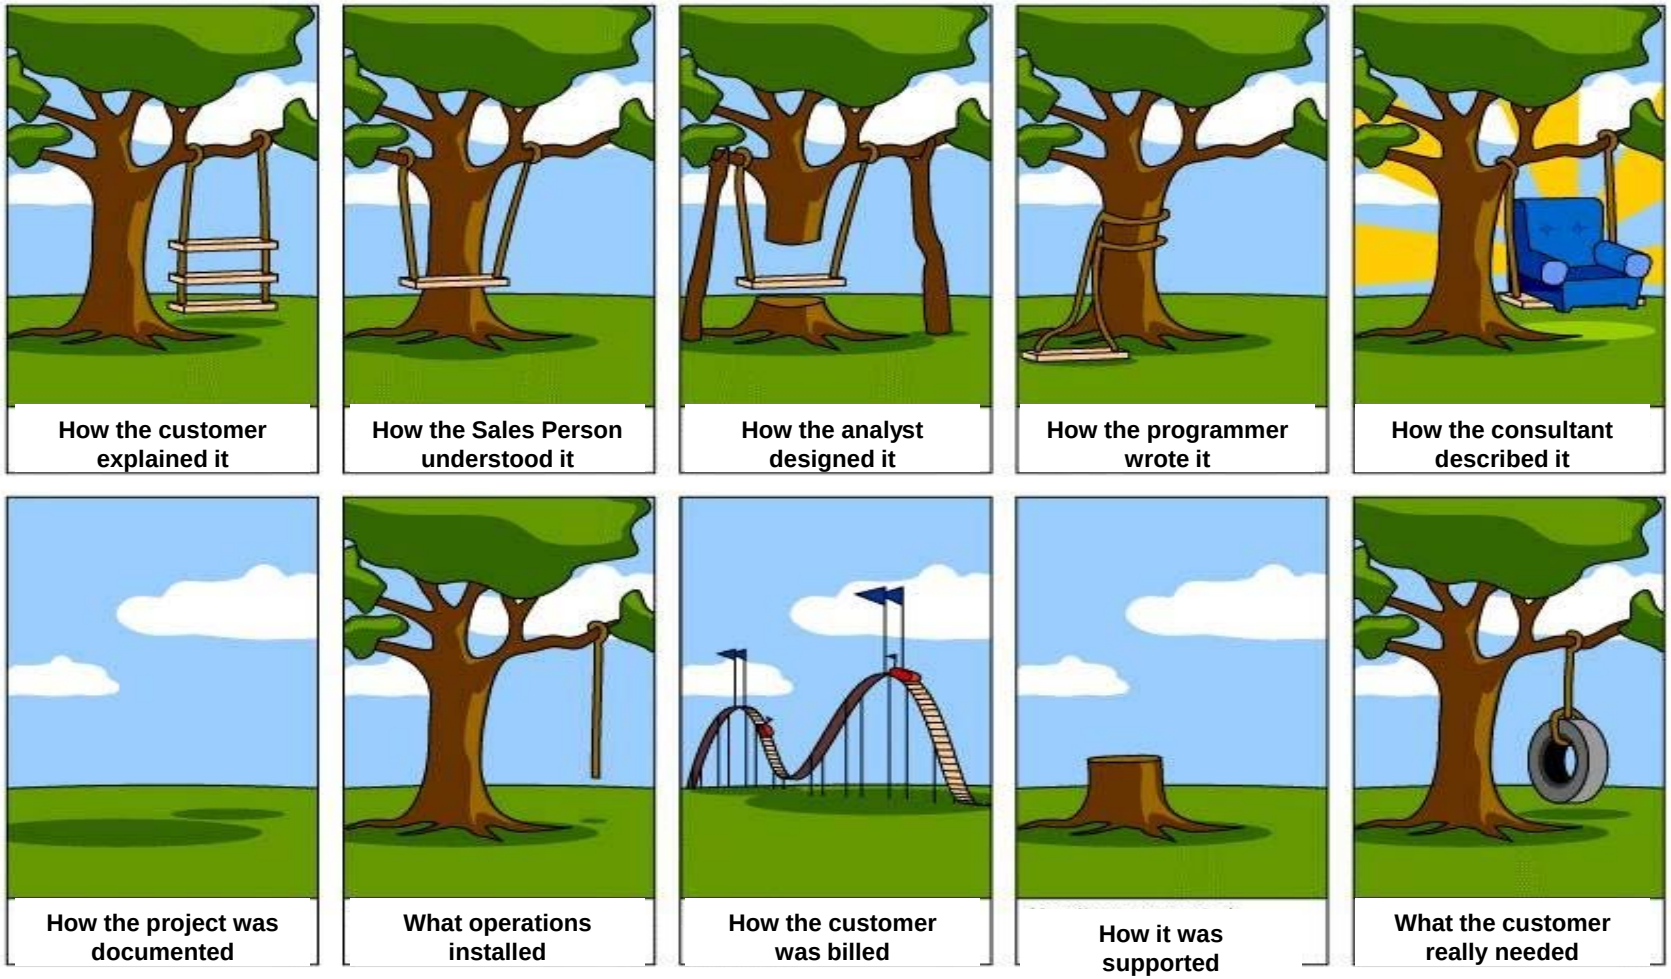

Murphy's Law

If anything can go wrong, it will

If anything just cannot go wrong it will

Left to themselves, things tend to go from bad to worse

If everything seems to be going well, you have obviously overlooked something

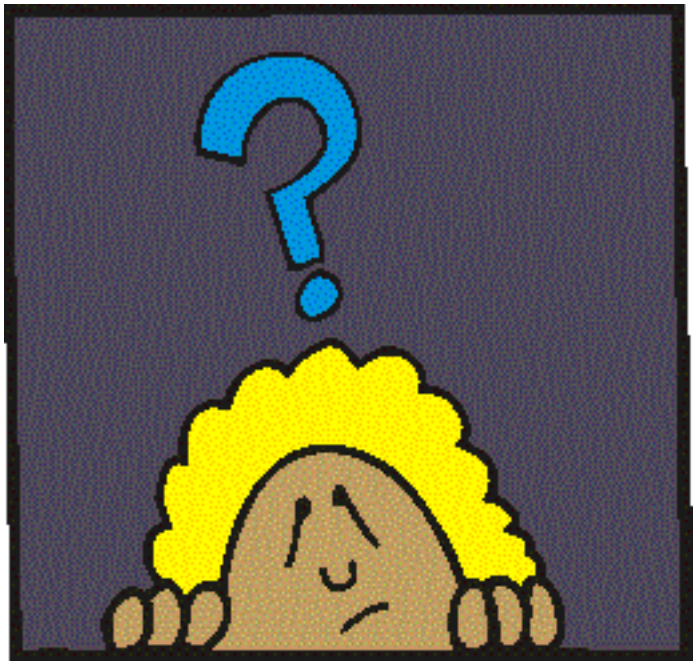

eTOM
- Extends M.3xxx
- Process & Functional Architecture
- Defines processes for providing services

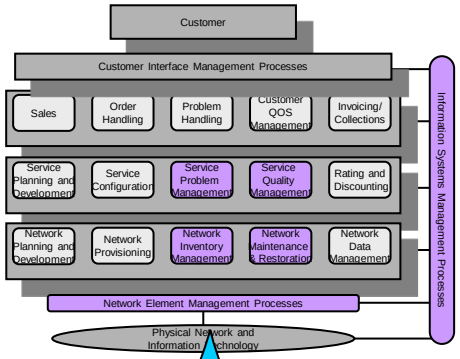

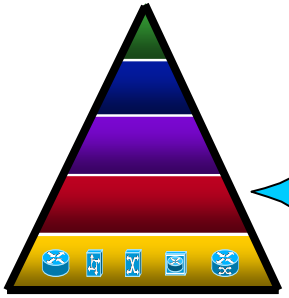

ITU – M.3xxx
- Physical Focus
- Defines interfaces & functions
- M.3400 focuses on functions
- Recommended architecture for TMN
- Recommended interfaces Q_x CMIP

Integrated Service Management

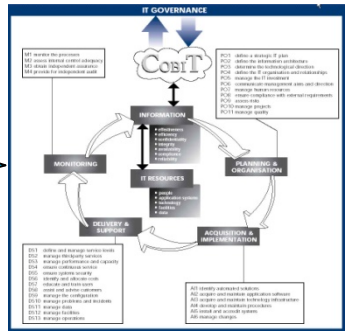

COBIT
-IT Infrastructure management focus
-IT Governance
- Planning
- Investment
- Projects
- Quality
- Delivery
- Support

www.ITIL.org

ITIL
- Process Focus
- IT Service management
- Service level
- Equates to COBIT Dxxx processes

Agenda

Introduction and goals

Management planning model

Practices for Management planning

Example

Effective Management

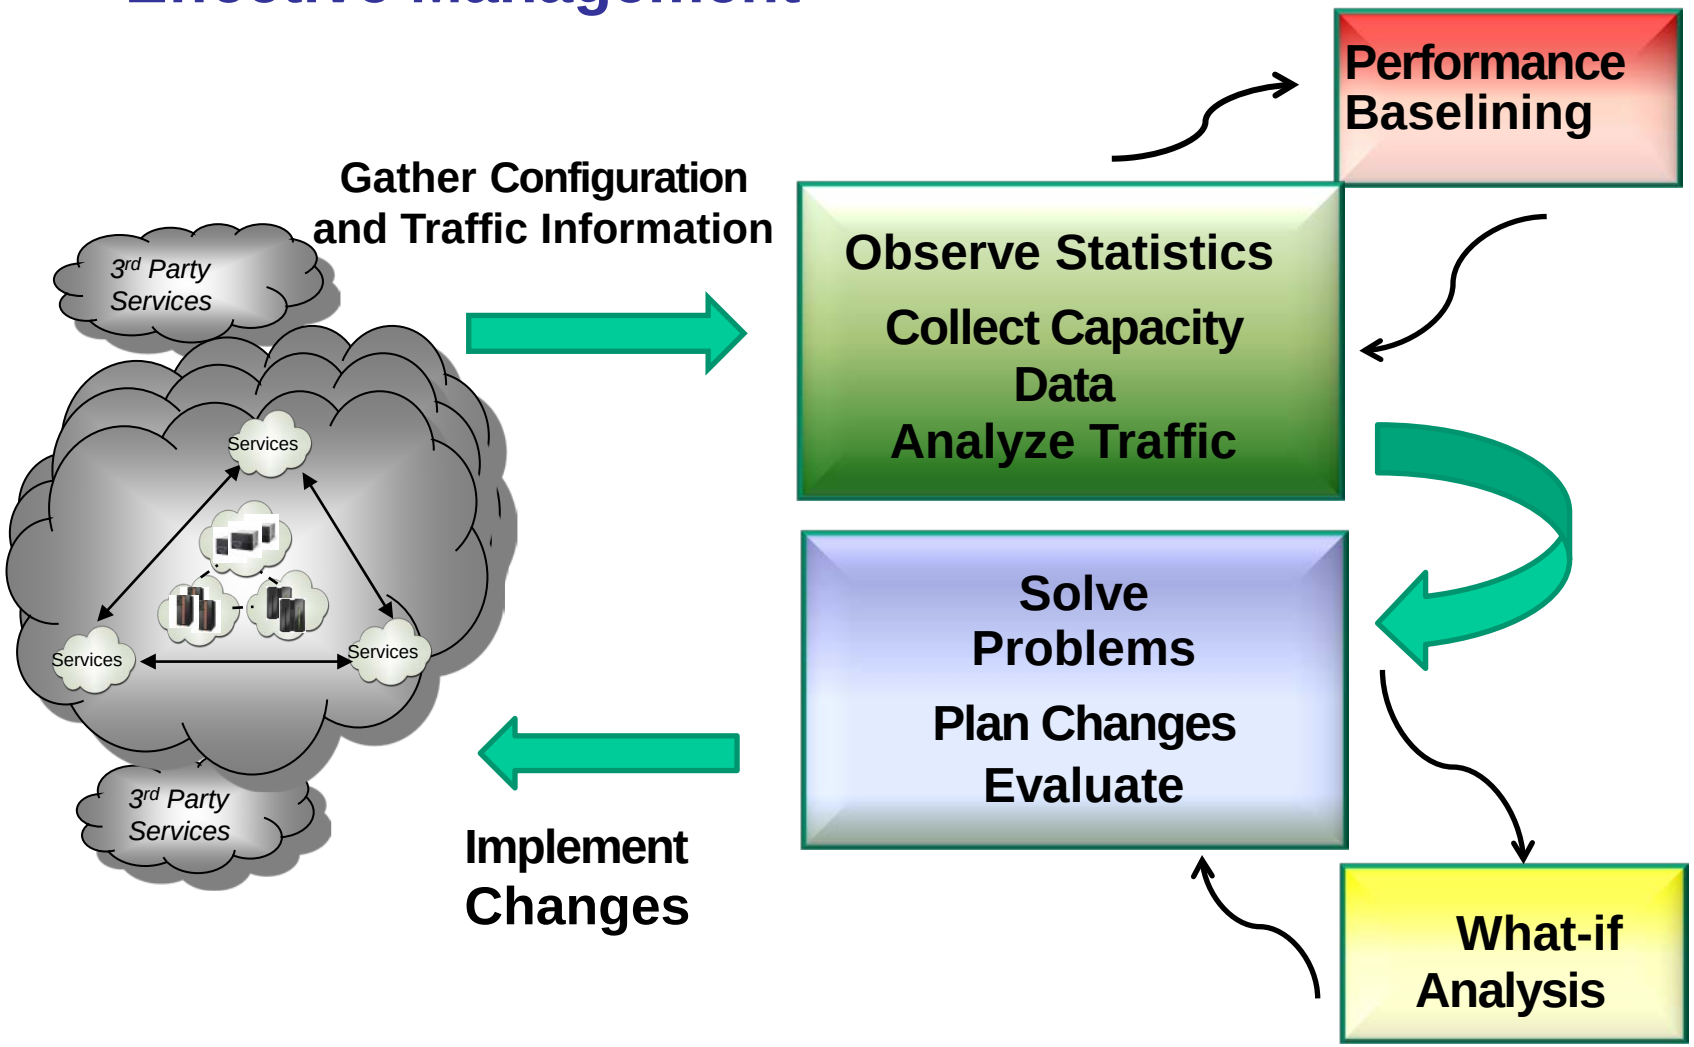

Information to Collect

- Link/segment utilization
- CPU Utilization
- Memory utilization
- Response Time
- Round Trip Time
- Queue/buffer drops
- Broadcast volumes
- Traffic shaping parameters
- RMON statistics
- Packet/frame drop/loss
- Environment specific

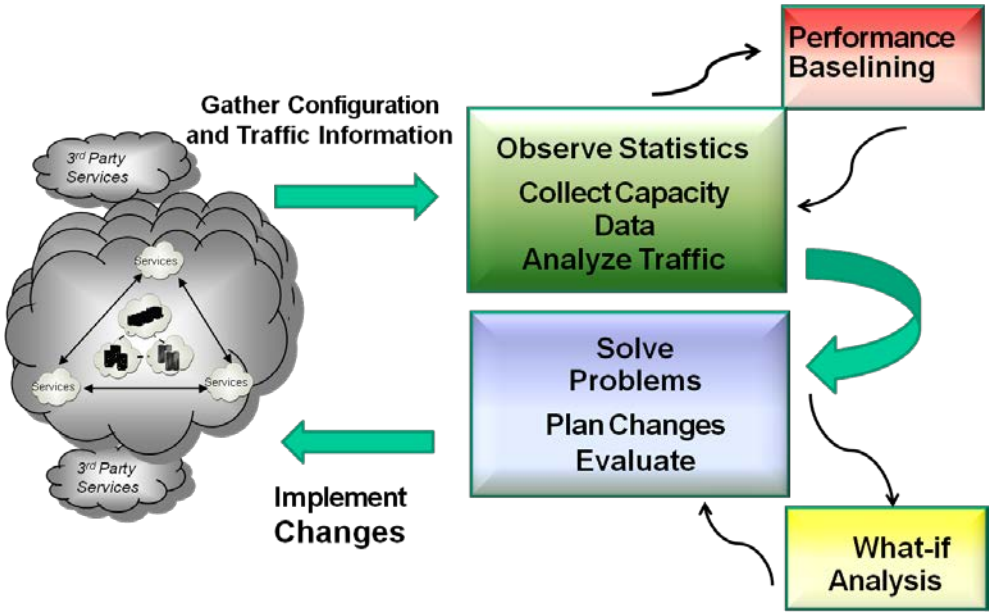

Plan

Develop information collection plan
- Define resources to be monitored for availability
- Define parameters to be monitored/measured and the thresholds
- Acquire proper authority to collect and monitor/measure
- Acquire proper authority to change thresholds
- Determine frequency of monitoring and reporting
- Define parameters that trigger alert mechanism

Define performance areas of interest

Report and interpret results

Determine tools for collecting information

Determine tools for analyzing information

Agenda

Introduction and goals

Management planning model

Practices for Management planning

Example

Management Practices

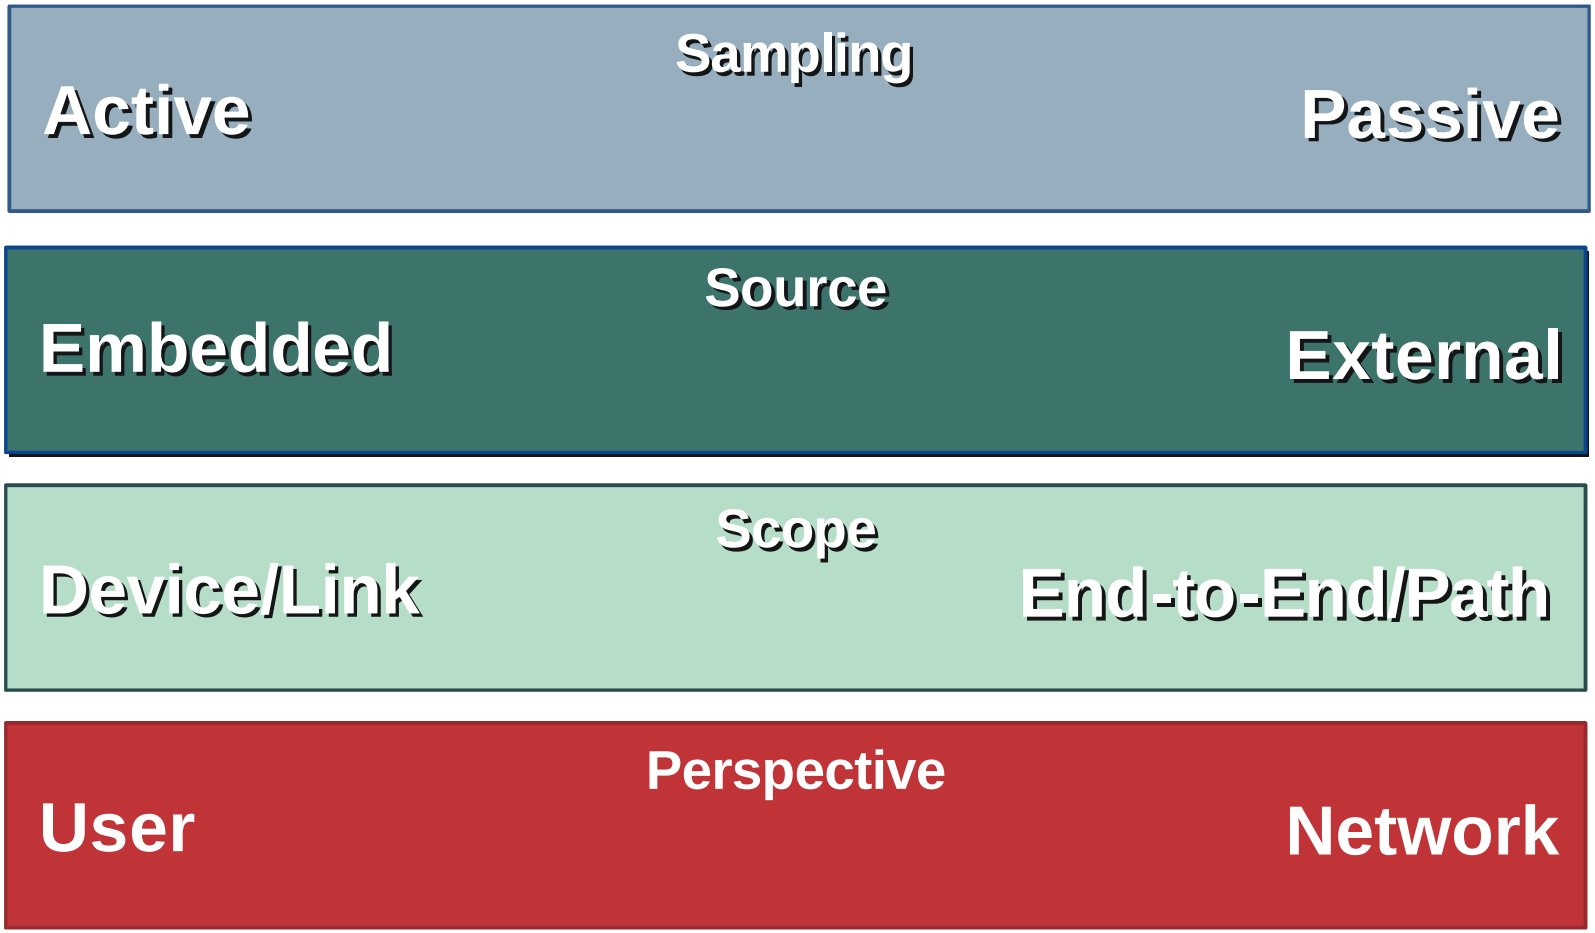

Active and Passive Management

Passive

- **Definition**
  - Actual end-user network traffic where performance is measured by timing specific application traffic flows
- **Advantages**
  - Most accurate for live application traffic on a specified link
- **Disadvantages**
  - Limited to measuring:
    - Existing traffic types, which may not be present on the network at all times
    - Existing traffic patterns, which may not reflect patterns for new or future applications

Active

- **Definition**
  - Network traffic generated strictly for the purpose of measuring a network performance characteristic
- **Advantages**
  - Measures performance:
    - Between any two points in the network
    - Controllable, on a continuous basis
    - By traffic class based on IP Precedence marking
- **Disadvantages**
  - Only an approximation for performance of live traffic

Embedded or External Sourcing

Embedded

- **Definition**

  Mechanisms for collection of network statistics are integrated into the network communication device (e.g., router or switch), itself

- **Advantages**

  Follows network infrastructure

  Gathers metrics that cannot be observed externally

- **Disadvantages**

  Performance monitoring has device-level performance implications

External

- **Definition**

  Mechanisms for collection of network statistics are provided by a stand-alone device specifically designed to collect network performance statistics

- **Advantages**

  Validation of performance performed independent of the devices that transmit network traffic

- **Disadvantages**

  More hardware to administer

  Observed statistics limited to points of deployment

Scoping Practices

Device or Link Oriented

- **Definition**

  Performance measurement based on analysis of specific device or device interface, and typically based on utilization rates

- **Advantages**

  Detailed application performance monitoring of critical network links

- **Disadvantages**

  When network-wide performance problems exist, how does one select which device or link to evaluate?

End-to-End

- **Definition**

  Performance measurement based on analysis of response time across two or more network devices, and typically based on latency

- **Advantages**

  Starting point performance troubleshooting

  Reflects end-user experience

- **Disadvantages**

  Prior knowledge of relevant end-to-end paths is needed

User or Network Perspective

User

- **Definition**

  Measurement based on performance statistics measured at the end-user workstation

- **Advantages**

  Accurate measurement of end-user experience

- **Disadvantages**

  Scale and distribution issues

  Intrusive on the desktop

Network

- **Definition**

  Measurement based on performance statistics measured in network devices

- **Advantages**

  Easy to deploy, and non-intrusive to the desktop

  Identifies network performance issue

- **Disadvantages**

  Imperfect understanding of end-user experience

Steps to Effective Performance Management

Monitor

Monitor over a long period of time to develop baselines and trends

Review and Remediate

Data analysis with no preconceived bias for capacity and performance trend development and any time dependencies

Establish Baselines

Report on trends, changes, and exceptions

Analyze

Baseline re-evaluation and resetting

Baseline Your Network

Gather inventory information

Gather statistics at a given time(s)

Monitor statistics over time and study traffic flows

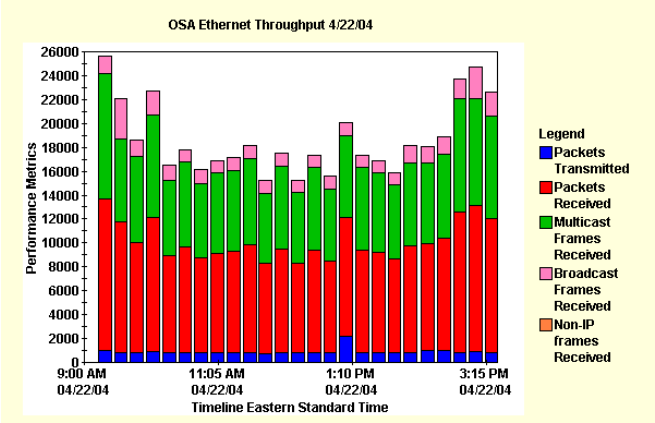

Have logical maps of network, server and application views

Know the protocols and traffic profiles

Document physical and logical network

Document detailed and measurable SLAs

Have a list of variable collected for your baseline

Be part of change control system

Agenda

Introduction and goals

Management planning model

Practices for Management planning

Example

Monitor: ICMPv6 Fields

Field	Description
Node Name	Linux node name
Node Address	IP address associated with the selected node
Last Checked	Time the most current sample was taken: hh:mm:ss
Msg In	Total number of ICMP messages received by an interface which includes all those counted by ipv6IfIcmpInErrors
Errors In	Number of ICMP messages which an interface received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.)
Unreach In	Number of ICMP destination unreachable messages received by an interface
Time Excd In	The number of ICMP time exceeded messages received by an interface.
Parm Prob In	The number of ICMP parameter problem messages received by an interface.
Redirect In	Number of redirect messages received by an interface
Echo In	Number of ICMP echo (request) messages received by an interface

Monitor: ICMPv6 Fields

Echo Reply In	Number of ICMP echo reply messages received by an interface
Admin Prohibs In	Number of ICMP destination unreachable/communication administratively prohibited messages received by an interface
Pkt Too Big In	Number of ICMP packet too big messages received by an interface
Router Solicit In	Number of ICMP router solicit messages received by an interface
Router Advrt In	Number of ICMP router advertisement messages received by an interface
Neighbor Solicit In	Number of ICMP neighbor solicit messages received by an interface
Neighbor Advrt In	Number of ICMP neighbor advertisement messages received by an interface
Group Query In	Number of ICMPv6 group membership query messages received by an interface
Group Resp In	Number of ICMPv6 group membership response messages received by an interface

Monitor: IPv6 Fields

Field	Description
Node Name	Linux node name
Node Address	IP address associated with the selected node
Last Checked	Time the most current sample was taken: hh:mm:ss
Admin Status	Desired state of an interface. When a managed system initializes, all IPv6 interfaces start with ipv6IfAdminStatus in the down(2) state. As a result of either explicit management action or per configuration information retained by the managed system, ipv6IfAdminStatus is then changed to the up(1) state (or remains in the down(2) state)
Oper Status	Current operational state of the interface. The noIfIdentifier(3) state indicates that no valid Interface Identifier is assigned to the interface. This state usually indicates that the link-local interface address failed Duplicate Address Detection. If ipv6IfAdminStatus is down(2) then ipv6IfOperStatus is down(2). If ipv6IfAdminStatus is changed to up(1) then ipv6IfOperStatus should change to up(1) if the interface is ready to transmit and receive network traffic; it should remain in the down(2) or noIfIdentifier(3) state if and only if there is a fault that prevents it from going to the up(1) state; it should remain in the notPresent(5) state if the interface has missing (typically, lower layer) components
Datagrams In	Total number of input datagrams received by an interface, including those received in error
Datagrams In Delivered	Total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP)

Monitor: IPv6 Fields

Multicast In	Number of multicast packets received by an interface
Datagrams Out Delivered	Total number of IPv6 datagrams which local IPv6 user-protocols (including ICMP) supplied to IPv6 in requests for transmission. This counter does not include any datagrams counted in ipv6IfStatsOutForwDatagrams
Multicast Out	Number of multicast packets transmitted by an interface
Headers Errors	Number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, errors discovered in processing their IPv6 options, etc
Too Large	Number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface
No Routes	Number of input datagrams discarded because no route could be found to transmit them to their destination
Address Errors	Number of input datagrams discarded because the IPv6 address in the IPv6 header's destination field is not a valid address to be received at the entity. This includes invalid addresses (e.g., ::0) and unsupported addresses (e.g., addresses with unallocated prefixes). For entities which are not IPv6 routers and therefore do not forward datagrams, this includes datagrams discarded because the destination address was not a local address
Unknown Protos	Number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This increases at the interface to which the datagrams are addressed, which may not necessarily be the input interface for some of the datagrams
Datagrams In Trunc	Number of input datagrams discarded because datagram frame did not carry enough data

Monitor: IPv6 Fields

Datagrams In Discards	Number of input IPv6 datagrams which encountered no problems to prevent their continued processing, but were discarded (e.g., for lack of buffer space). This does not include datagrams discarded while awaiting re-assembly
Datagrams Out Discards	Number of output IPv6 datagrams which encountered no problem to prevent transmission to their destination, but were discarded (e.g., for lack of buffer space). This counter includes datagrams counted in ipv6IfStatsOutForwDatagrams if any such packets met the (discretionary) discard criterion
Datagrams Out Forward	Number of output datagrams which an entity received and forwarded to their final destination. In entities which do not act as IPv6 routers, this includes only those packets which are source-routed via the entity, and the source-route processing was successful
Frag OK	Number of IPv6 datagrams successfully fragmented at the output interface
Frag Creates	Number of output datagram fragments generated as a result of fragmentation at the output interface
Frag Fails	Number of output datagram fragments that failed
Reasm Req	Number of IPv6 fragments received which need to be reassembled at the interface
Reasm OK	Number of IPv6 datagrams successfully reassembled
Reasm Fails	Number of failures detected by the IPv6 re-assembly algorithm (for various reasons: timed out, errors, etc.)

Trace IPv6 and ICMPv6

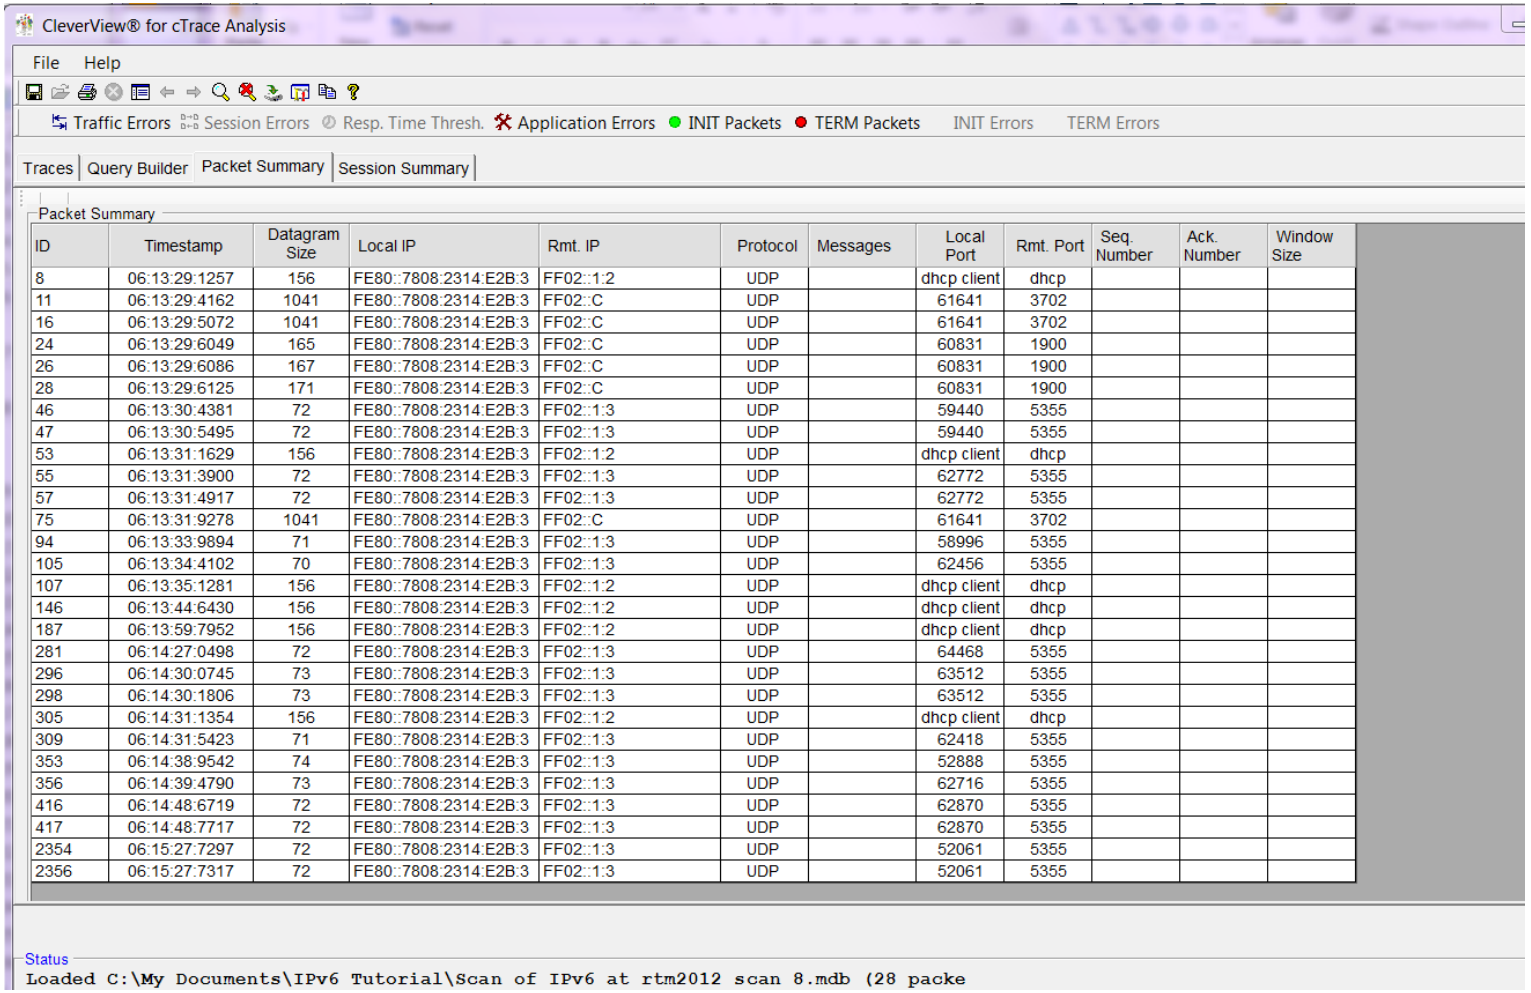

ID	Timestamp	Datagram Size	Local IP	Rmt. IP	Protocol	Messages	Local Port	Rmt. Port	Seq. Number	Ack. Number	Window Size
8	06:13:29:1257	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
11	06:13:29:4162	1041	FE80::7808:2314:E2B:3	FF02::C	UDP		61641	3702			
16	06:13:29:5072	1041	FE80::7808:2314:E2B:3	FF02::C	UDP		61641	3702			
24	06:13:29:6049	165	FE80::7808:2314:E2B:3	FF02::C	UDP		60831	1900			
26	06:13:29:6086	167	FE80::7808:2314:E2B:3	FF02::C	UDP		60831	1900			
28	06:13:29:6125	171	FE80::7808:2314:E2B:3	FF02::C	UDP		60831	1900			
46	06:13:30:4381	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		59440	5355			
47	06:13:30:5495	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		59440	5355			
53	06:13:31:1629	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
55	06:13:31:3900	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62772	5355			
57	06:13:31:4917	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62772	5355			
75	06:13:31:9278	1041	FE80::7808:2314:E2B:3	FF02::C	UDP		61641	3702			
94	06:13:33:9894	71	FE80::7808:2314:E2B:3	FF02::1:3	UDP		58996	5355			
105	06:13:34:4102	70	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62456	5355			
107	06:13:35:1281	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
146	06:13:44:6430	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
187	06:13:59:7952	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
281	06:14:27:0498	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		64468	5355			
296	06:14:30:0745	73	FE80::7808:2314:E2B:3	FF02::1:3	UDP		63512	5355			
298	06:14:30:1806	73	FE80::7808:2314:E2B:3	FF02::1:3	UDP		63512	5355			
305	06:14:31:1354	156	FE80::7808:2314:E2B:3	FF02::1:2	UDP		dhcp client	dhcp			
309	06:14:31:5423	71	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62418	5355			
353	06:14:38:9542	74	FE80::7808:2314:E2B:3	FF02::1:3	UDP		52888	5355			
356	06:14:39:4790	73	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62716	5355			
416	06:14:48:6719	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62870	5355			
417	06:14:48:7717	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		62870	5355			
2354	06:15:27:7297	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		52061	5355			
2356	06:15:27:7317	72	FE80::7808:2314:E2B:3	FF02::1:3	UDP		52061	5355			

Operation of Dual Stack Services

What happens when I retrieve a web image? The web site is dual stack.

IPv6 only
Dual stack IPv6 and IPv4
IPv4 only
IPv6 only with no DNS

V4	V6	Dual	Node Type
✔	✖	V4	V4-Only
✖	✔	V6	V6-Only
✔	✔	V6	V6-Preferred
✔	✔	V4	V6-Capable (V4-Preferred)
✔	✖	✖	Dual-Stack Loss

Look at retrieval rates, behavior, and transaction time
Based on work at www.apnic.net

Dual Stack Preference

As you move to IPv4/IPv6 dual stack investigate which stack is being used

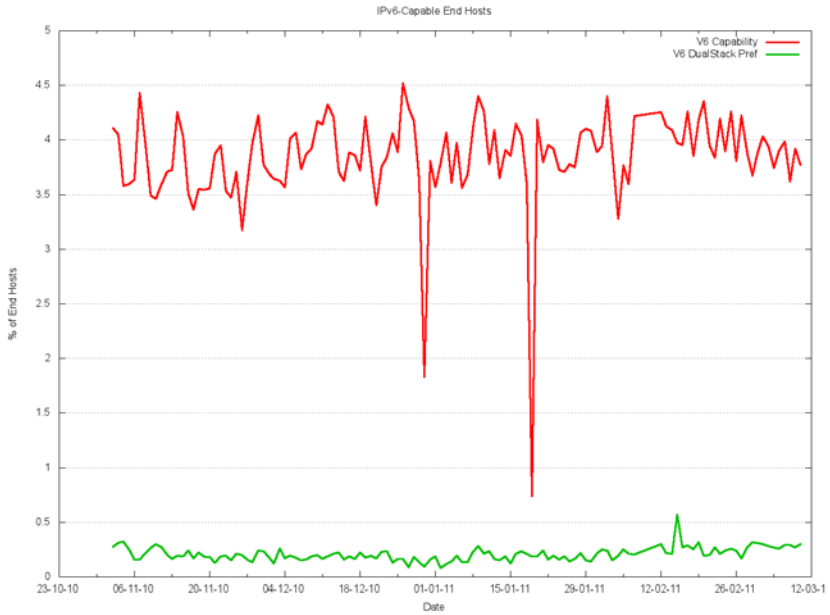

From APNIC Study

Around 4% of the monitored nodes can use IPv6, but only 0.2% are using IPv6

Why is the number so small based on those capable of using IPv6

Impact of Address Type

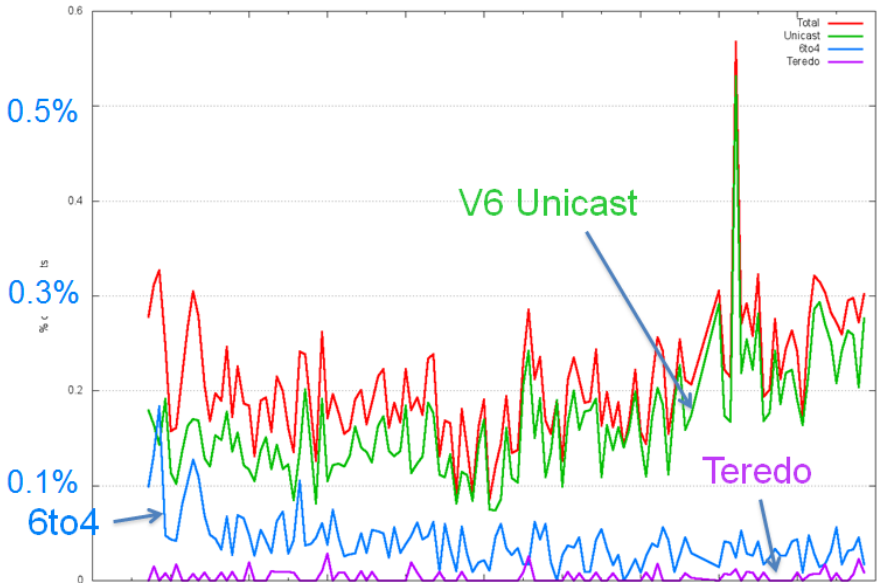

IPv6 Preferred

IPv4 Preferred

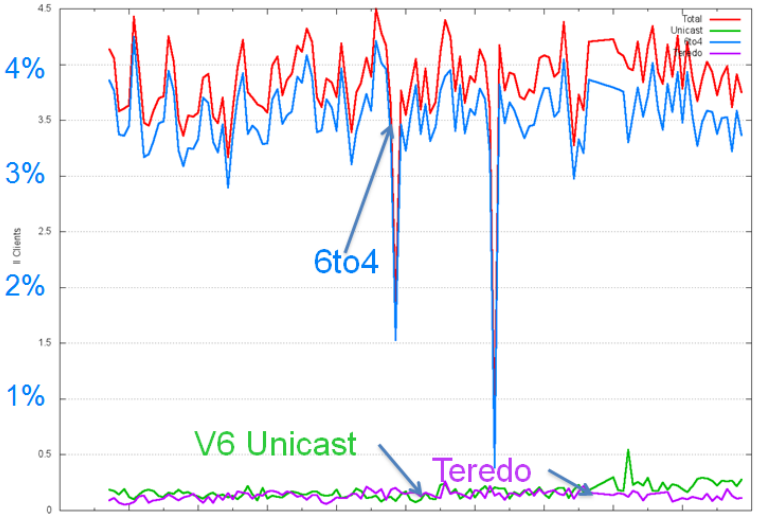

Impact of Address Type

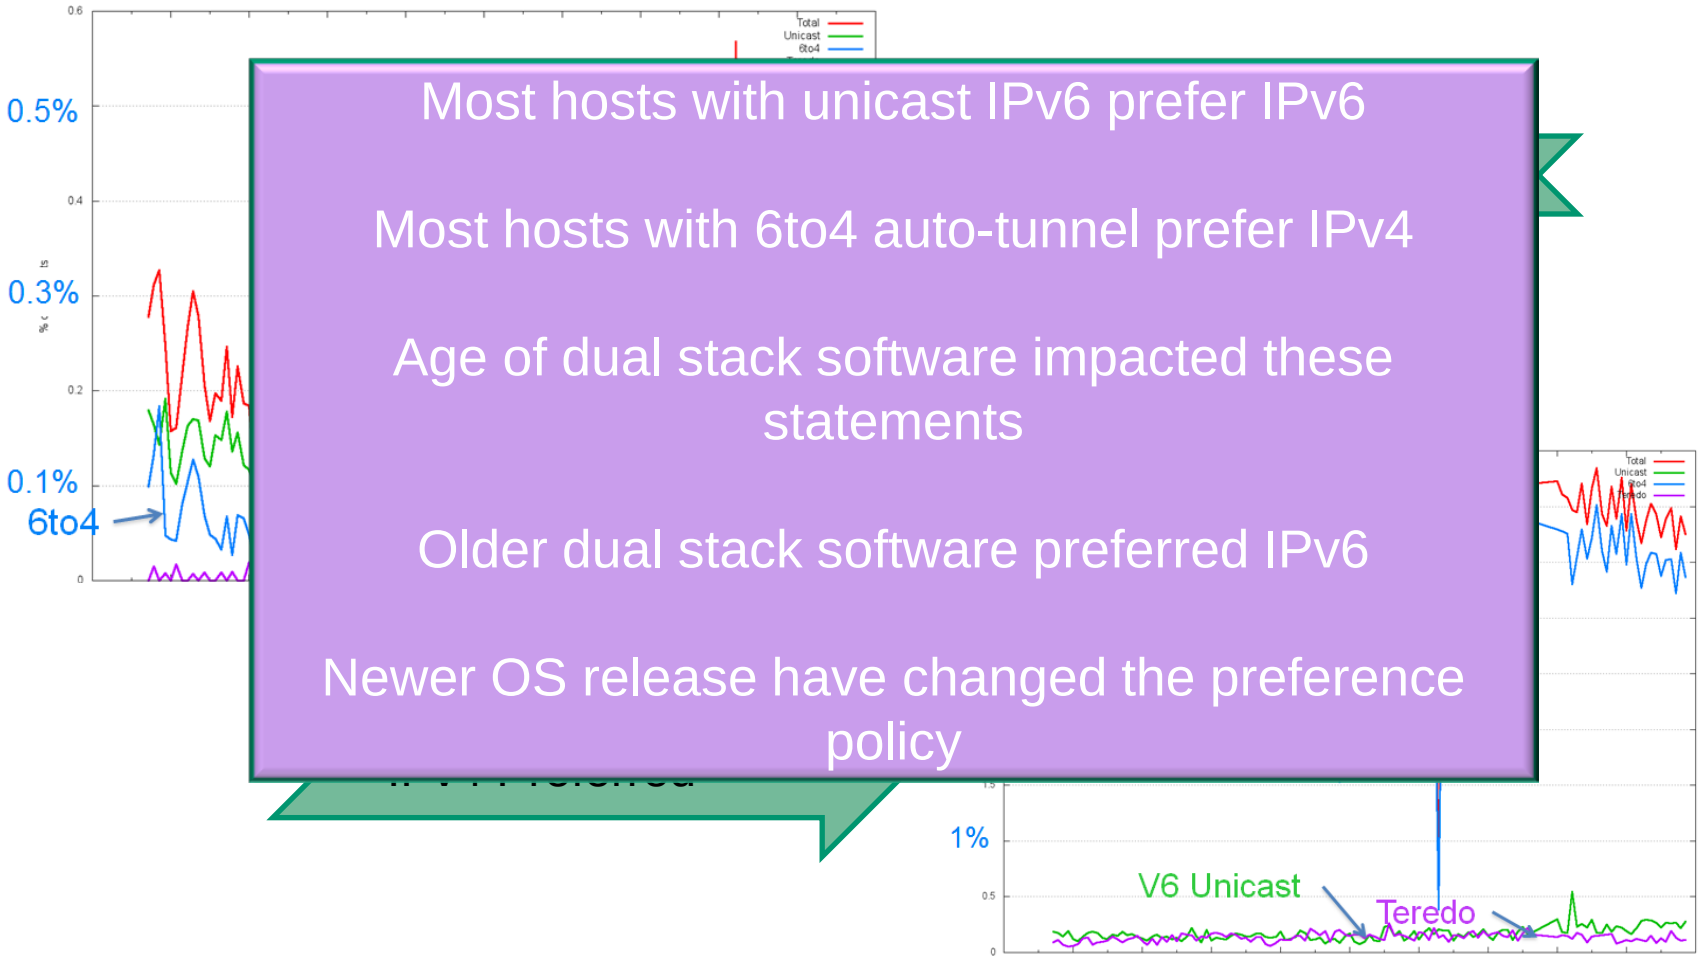

Most hosts with unicast IPv6 prefer IPv6

Most hosts with 6to4 auto-tunnel prefer IPv4

Age of dual stack software impacted these statements

Older dual stack software preferred IPv6

Newer OS release have changed the preference policy

Tunnel Performance

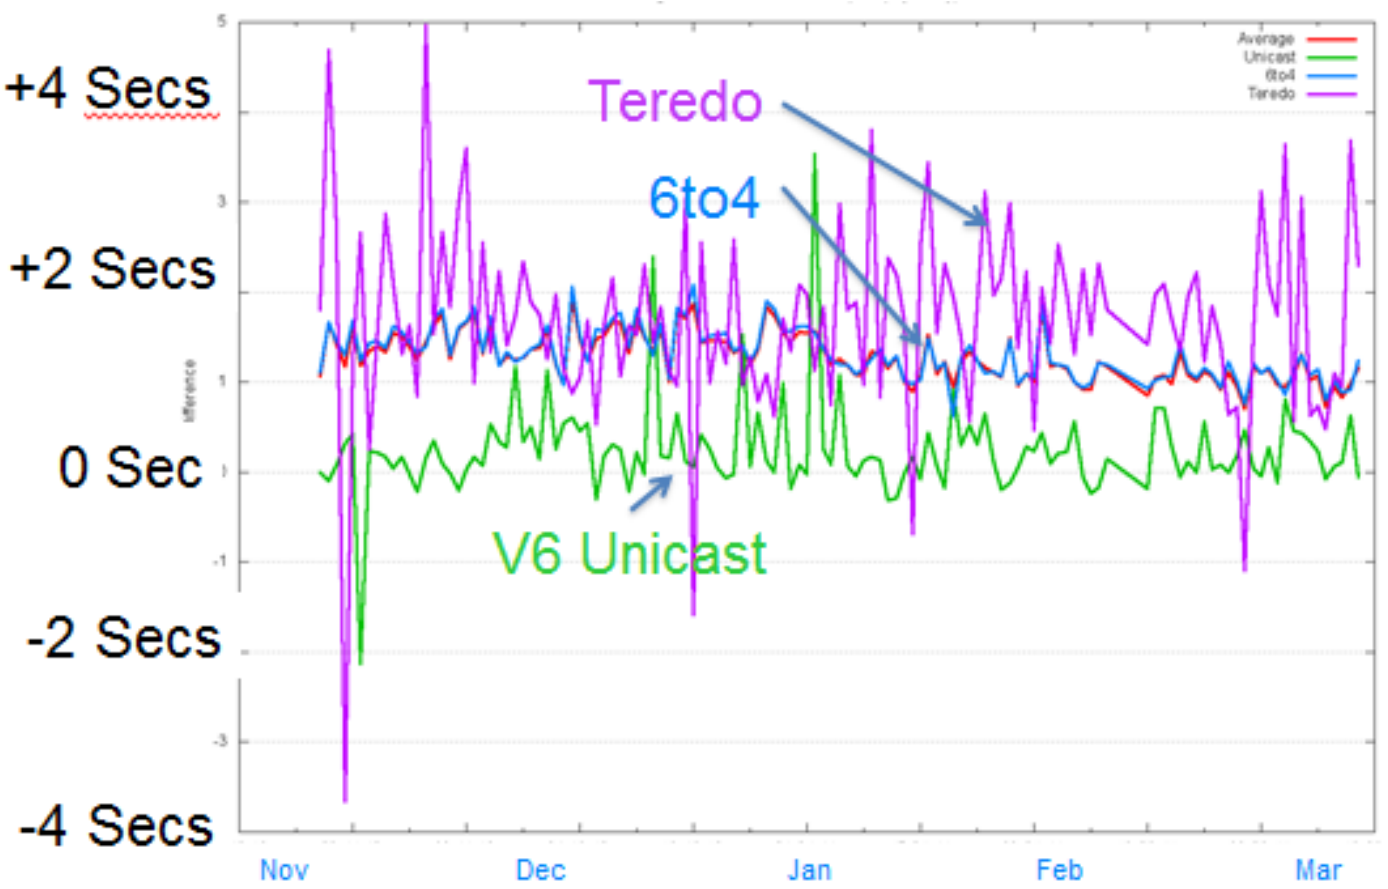

Tunnel Performance

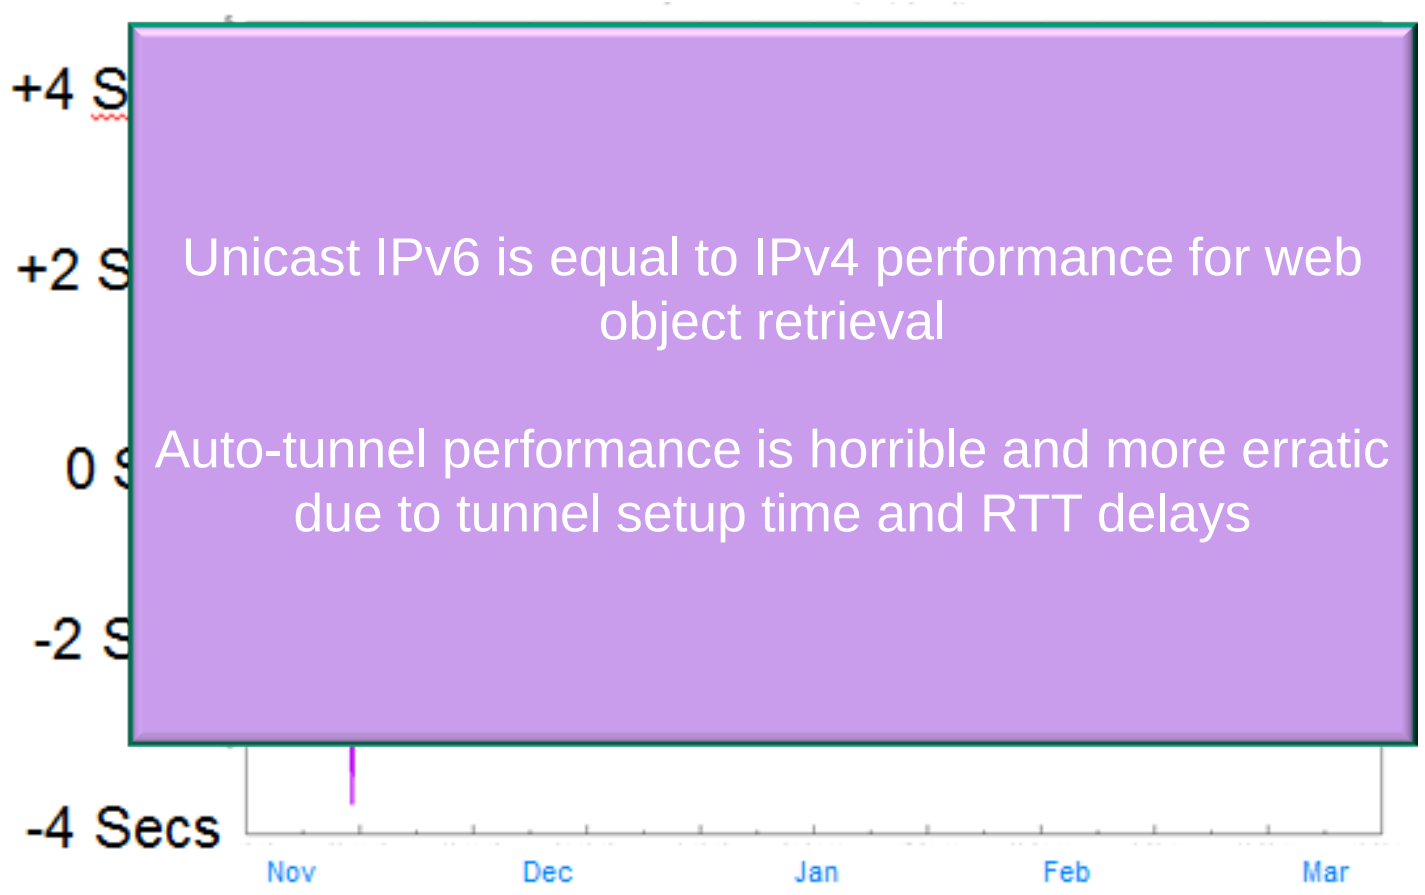

6to4 Packet Path

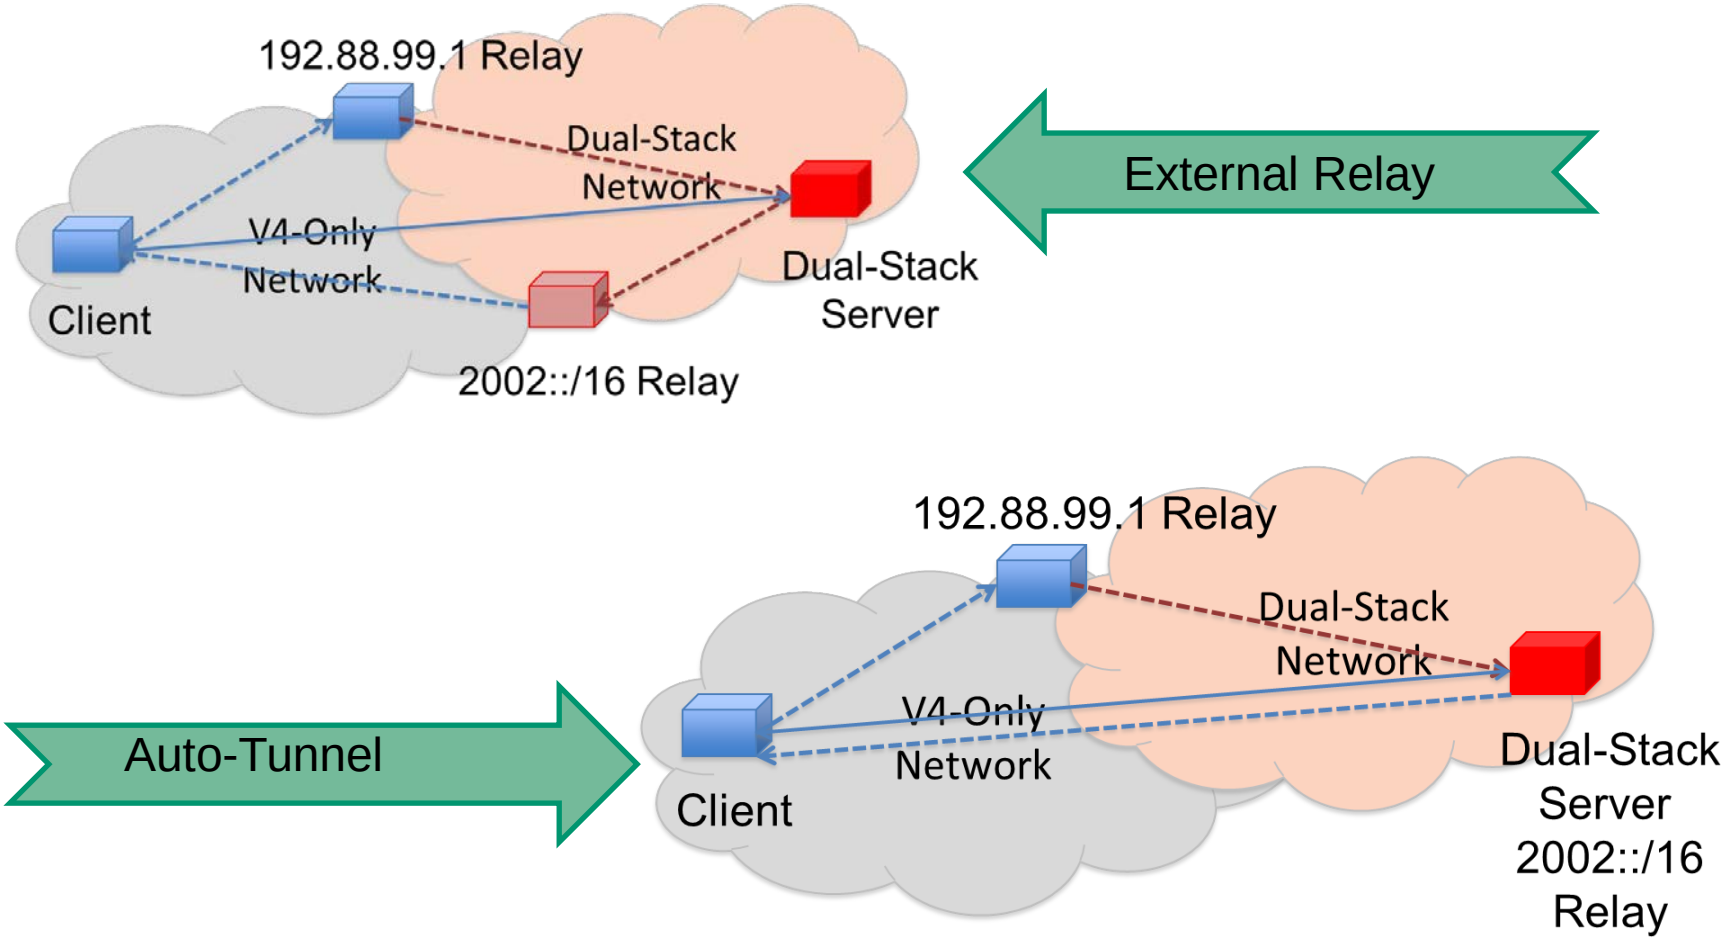

6to4 Performance

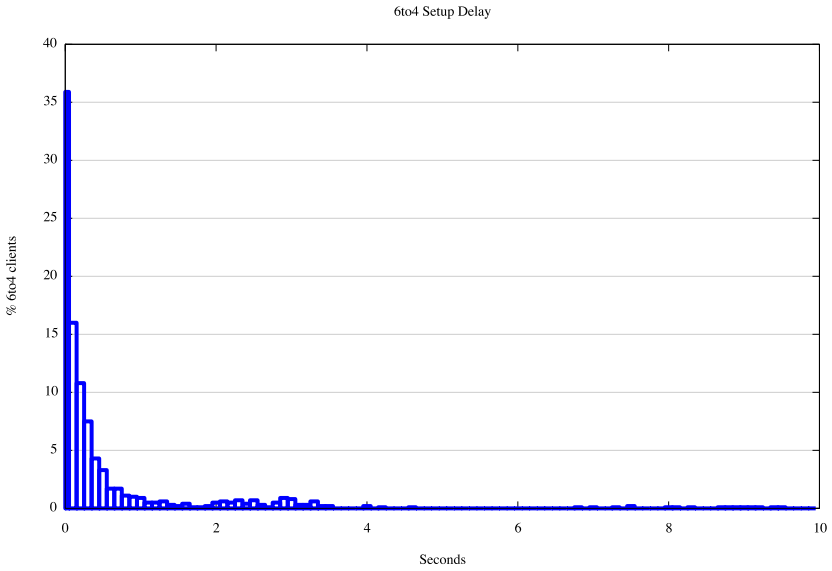

Setup Time

RTT Incremental Delay

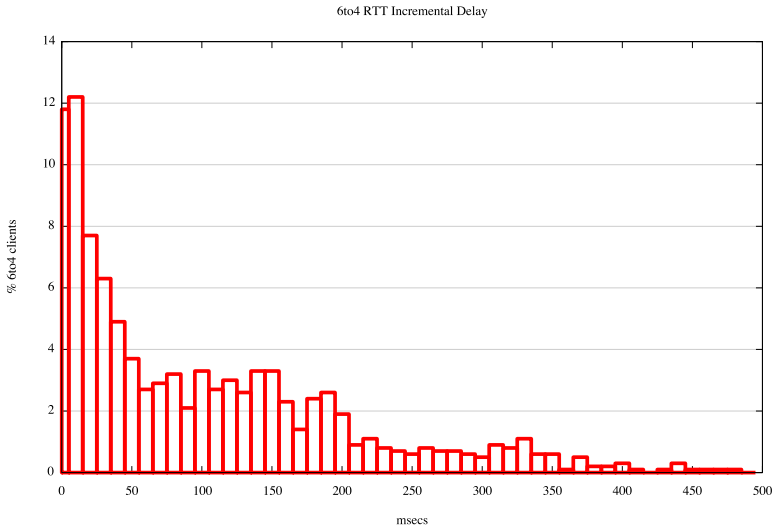

6to4 Packet Path

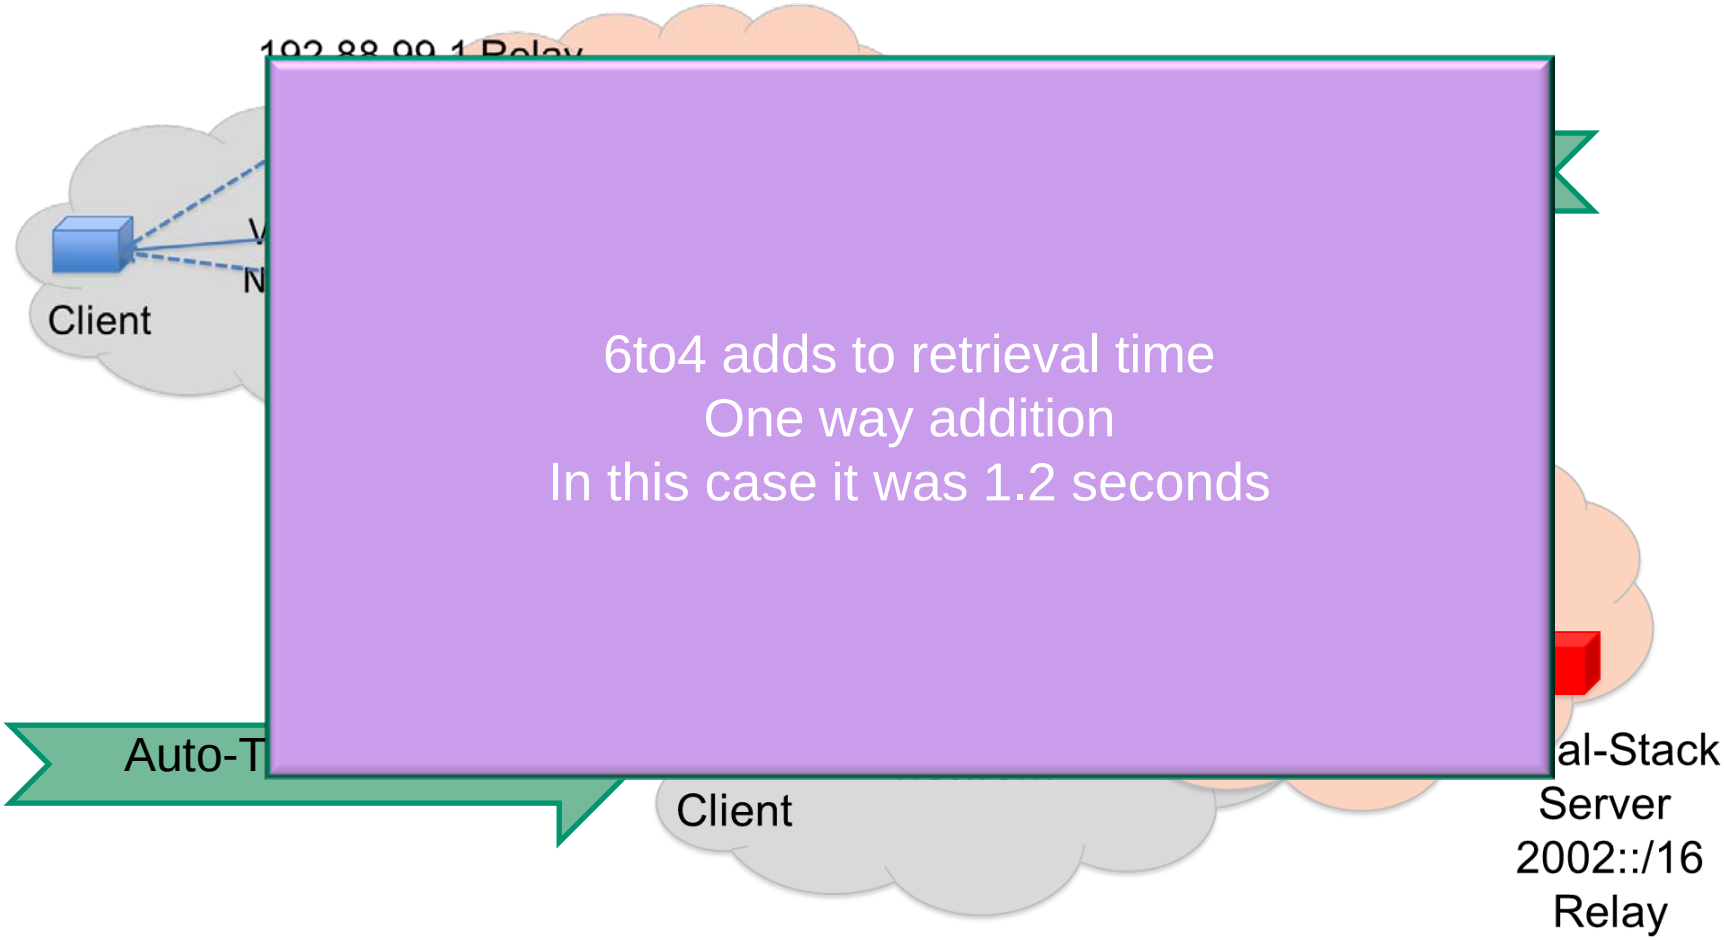

Why so Little Teredo Traffic?

- Windows 7 and Vista will not query for a AAAA record if only local IPv6 traffic is Teredo
- If we change the URL to force IPv6 what changes?

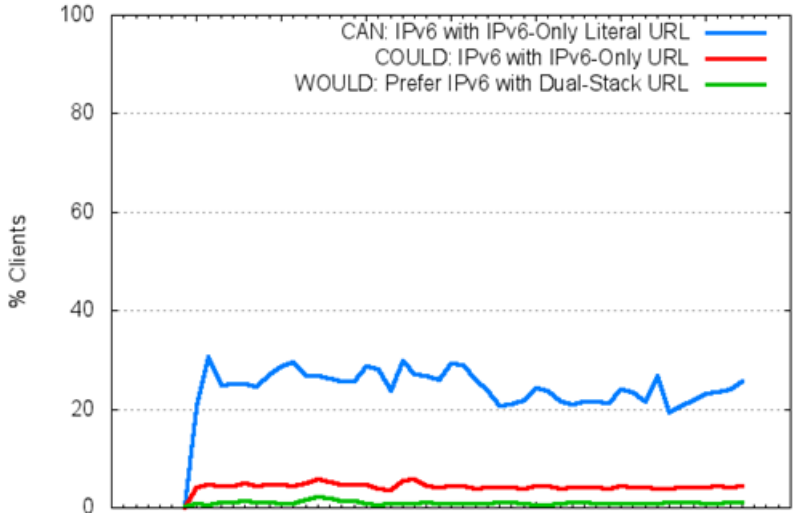

30% fetched URL using Teredo

Teredo Performance

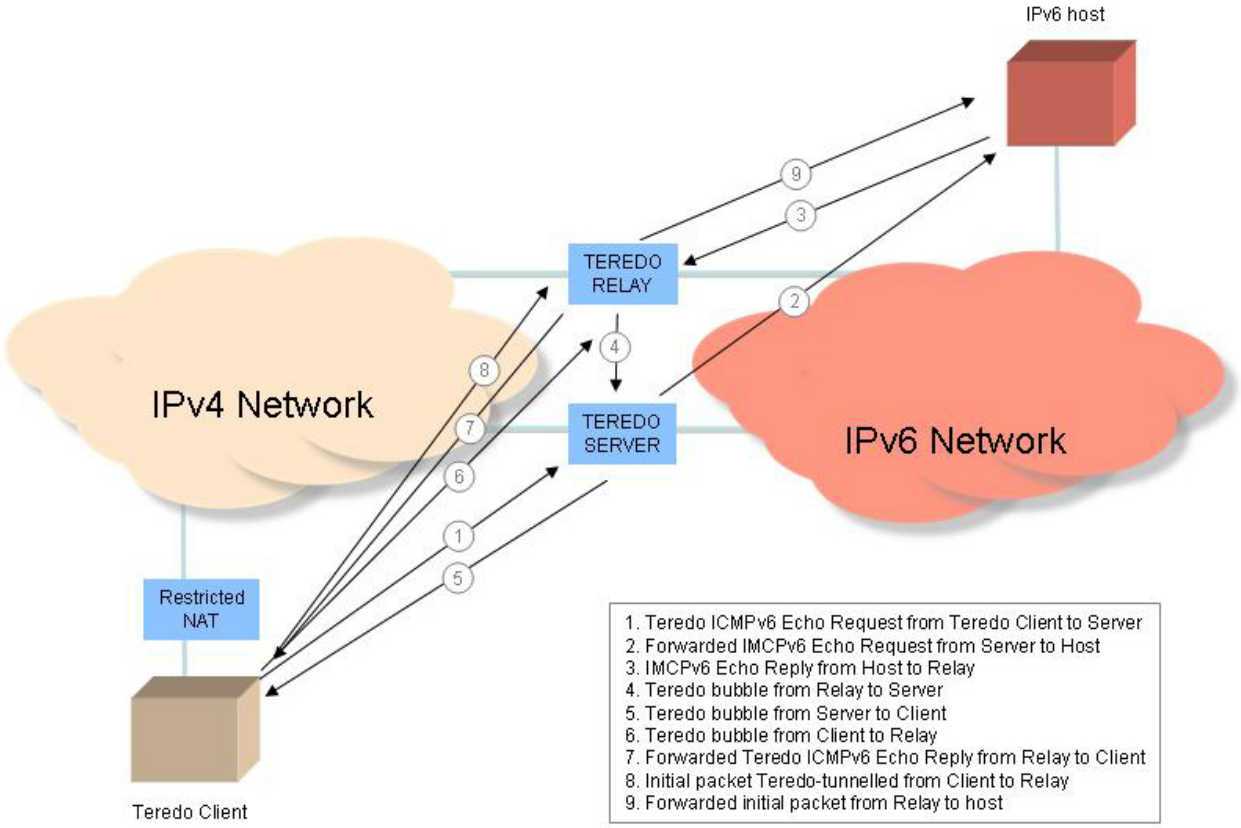

Lots of overhead in setting up Teredo client

Teredo Performance

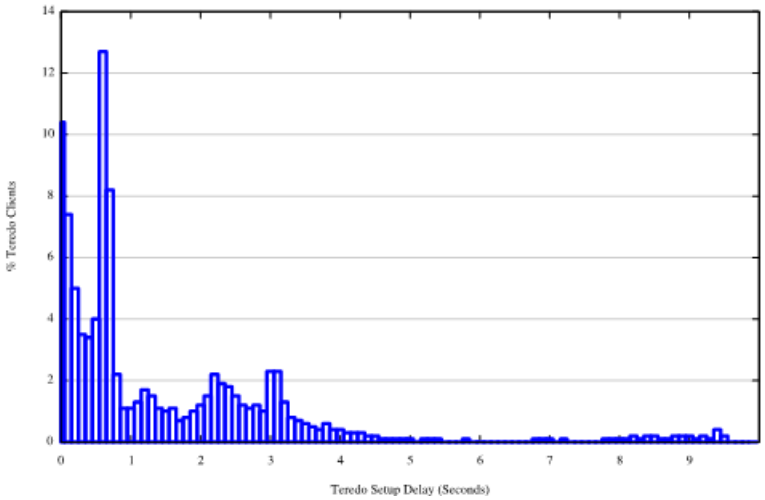

Setup Time

RTT Incremental Delay

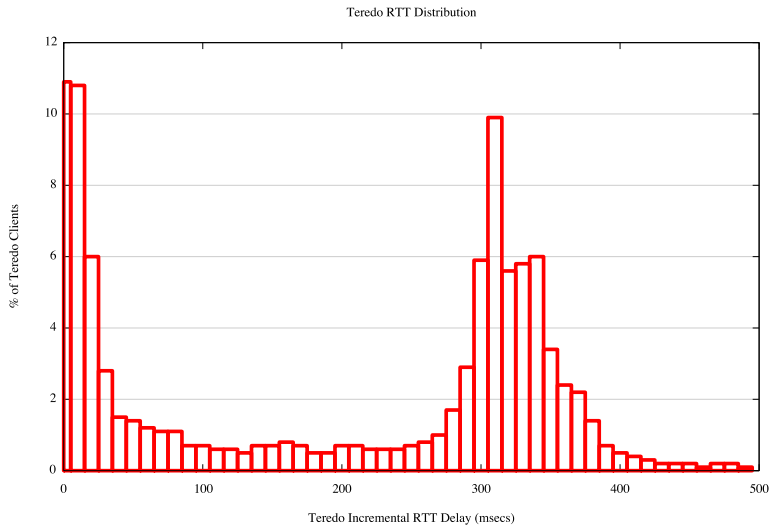

Teredo Performance

Toredo adds 1-3 seconds to the image download
Setup varied between .6 and 3 seconds
Average RTT was 300ms
HIGH OVERHEAD

RTT Incr

Observations

- Unicast IPv6 appears to be as fast as IPv4 for object retrieval
- Auto-tunneling IPv6 attracts some performance overheads
  - these are strongly context dependent
  - widespread deployment of 6to4 relays and Teredo relays and servers would mitigate this, to some extent
  - Dual Stack servers may want to consider using local 6to4 relays to improve reverse path performance for auto-tunneling clients

IPv6 Management Challenges

Double monitoring metrics

New monitoring metrics

Dual stack metric consolidation

Debugging which path?

IPv6 is a ‘flat’ network architecture

Tools still in infancy

Addressing significantly changes network topology

Tunneling can hide critical data

Questions?

laurak@aesclever.com
www.aesclever.com
650-617-2400