

Performance Modeling and Analysis of Networked Application Workloads on S/390 and zSeries

Gus Kassimis
Application Integration Middleware
IBM Software Group
kassimis@us.ibm.com



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

DB2*	Parallel Sysplex*
CICS*	S/390*
e-business logo*	WebSphere*
HiperSockets	z/OS
IBM*	z/VM
IBM logo*	zSeries

* Registered trademarks of IBM Corporation

Intel is a trademark of the Intel Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

SAP, mySAP.com and all other SAP products and services mentioned herein are trademarks or registered trademarks of SAP AG in Germany and several other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Challenges with eBusiness Solutions


- **Dynamic Environment**
 - ▶ Changes must be deployed quickly
 - New applications, changes to existing applications, network infrastructure, workload levels
 - ▶ Without adversely impacting existing application workloads
 - What is the impact on
 - Server platform requirements (CPU, storage, etc.)
 - Network infrastructure (network, routers, server connectivity, etc.)
 - SLA for new and changed workloads must be met
 - Degraded performance is often perceived as "down time"
- **Uncertainty often inhibits change**
 - ▶ For example, we're considering migrating off the 3745 controllers and are looking into SNA/IP migration/integration using Enterprise Extender.
 - What will be the impact of this change?
 - ▶ Modeling/Simulation tools provide an inexpensive solution
 - Changes can be evaluated before they are deployed
 - Detailed performance data for accurate capacity planning
 - Help answer what/if questions while still in the planning phase



Application Workload Modeler (AWM)

■ What is it?

- ▶ A application workload simulation tool for measuring and modeling the performance of the network infrastructure, end-to-end
 - Including the network, systems, and applications



- Allows enterprise or service providers to model the impact of various types of workloads on servers and networks
 - ▶ Users can evaluate the impact of a change **before** the change is deployed in production environment
 - *Do you need to upgrade network components to meet your performance objectives?*
 - *Is the existing network infrastructure sufficient?*
 - *Is this the right communications model for an application?*

Application Workload Modeler for z/OS (AWM)



- Application Workload Modeler for z/OS R1 (5655-J62)
 - ▶ New IBM product
 - ▶ Includes z/OS, Linux for zSeries and Linux/Intel versions of the product
- Worldwide Announce December 10, 2002
- Worldwide General Availability December 20, 2002



Key features

- **Generates application network traffic patterns**
 - ▶ **Simulates large numbers of end-users**
 - Capable of generating high volume, stress load conditions
 - Eliminates the need for manual simulations
 - ▶ **Multiple modes of operation**
 - Client/Server mode
 - Application Client mode
- **Application Workload Modeler provides detailed statistics to help evaluate the performance of these workloads**
 - ▶ Per "session" as well as cumulative averages for Throughput, Response time, Transaction rate
 - ▶ Transient behavior of workloads

Key Usage Scenarios

■ Network Provisioning

- ▶ Using application (client-server) traffic profiles
- ▶ Determine adequacy of network capacity and network-component of response times
 - What-if scenarios for traffic growth and traffic mixes
- ▶ Tune and monitor QoS configurations and policies
- ▶ Monitor performance of network

■ Application provisioning


- ▶ Configure AWM to generate traffic to typical application servers to validate performance for different workloads
- ▶ Complementary to other modeling tools for capacity planning

AWM - Providing answers to key questions

- What is the impact on performance when making changes to the network infrastructure, systems, software and applications?
 - ▶ Network infrastructure changes
 - Enterprise Extender, Virtual Private Networks (VPN), Quality of Service (QoS) deployment
 - Network connectivity updates (Fast Ethernet to Gigabit Ethernet, Channel Attached Routers to OSA Express, router upgrades, network protocol updates, etc.)
 - ▶ Application Changes
 - New/changed TCP/IP application deployment
 - Using Secure Sockets Layer (SSL) for encryption
 - Changes in SNA applications, including migration to TCP/IP
 - ▶ System/Software changes
 - Upgrading hardware or software?
 - Considering a Load Balancing Solution?
 - Server Consolidation?
 - How can I measure the impact of deploying a TN3270E solution?
 - What kind of performance should I expect from FTP transfers?
- IBM Application Workload Modeler can be used to answer all these questions!


Client/Server mode Benchmarks

- Application Workload Modeler simulates network communications for both client and server applications (SNA and TCP/IP applications)
 - ▶ Focus on performance measurements of end-to-end network communication paths
 - ▶ Allows modeling of common application workload patterns
 - Transactional request/response workloads
 - Bulk data transfer workloads
 - ▶ Can be used to model new application "network behavior"
 - **Prior** to application development/deployment
 - ▶ Best/worst case scenarios
 - Application-specific bottlenecks eliminated
 - High volume, stress load, client workload simulation



Application Client mode Benchmarks

- Application Workload Modeler simulates client network communications for well known, standard TCP/IP server applications
 - ▶ Focus on performance measurements of end-to-end network communication paths for key server applications
 - TN3270, FTP, Web Server, CICS sockets, DNS, DHCP, SAP R/3 ICL, SMTP
 - ▶ Customized stress load conditions
 - Number of clients, rates of requests, etc.
 - ▶ Helps answer what/if questions
 - What type of performance can I expect if my client request workload doubles?
 - Can a single server handle the anticipated workload?
 - What are the effects of QoS on specific workloads?




Application Workload Modeler Statistics

- Key output of the Application Workload Modeler
 - ▶ Performance data collection can be customized by user
 - Multiple samples
 - Number of transactions per sample
 - Interim and Final reports
 - ▶ Detailed performance statistics reported for all workload simulation tests
 - Customized reports based on workload type (e.g. SAP vs FTP)
 - Response time metrics (mean, min, max, variances, etc.)
 - Throughput rates, average data transfer rate, etc.


Multi-platform Support

- Additional flexibility with support for multiple platforms
 - ▶ *Application Workload Modeler for z/OS*
 - Supports current z/OS platforms and OS/390 V2R10
 - ▶ *Application Workload Modeler for Linux on zSeries*, which supports the following:
 - Red Hat Linux for S/390, or later
 - SuSE Linux Enterprise Server 7 for S/390 and zSeries, or later
 - ▶ *Application Workload Modeler for Linux*, which supports the following Linux distributions on Intel platforms:
 - RedHat 7.1, 7.2, 7.3, SuSE 7.3
- Allows users to select best platform based on specific workload modeling/simulation needs

AWM Client/Server Mode

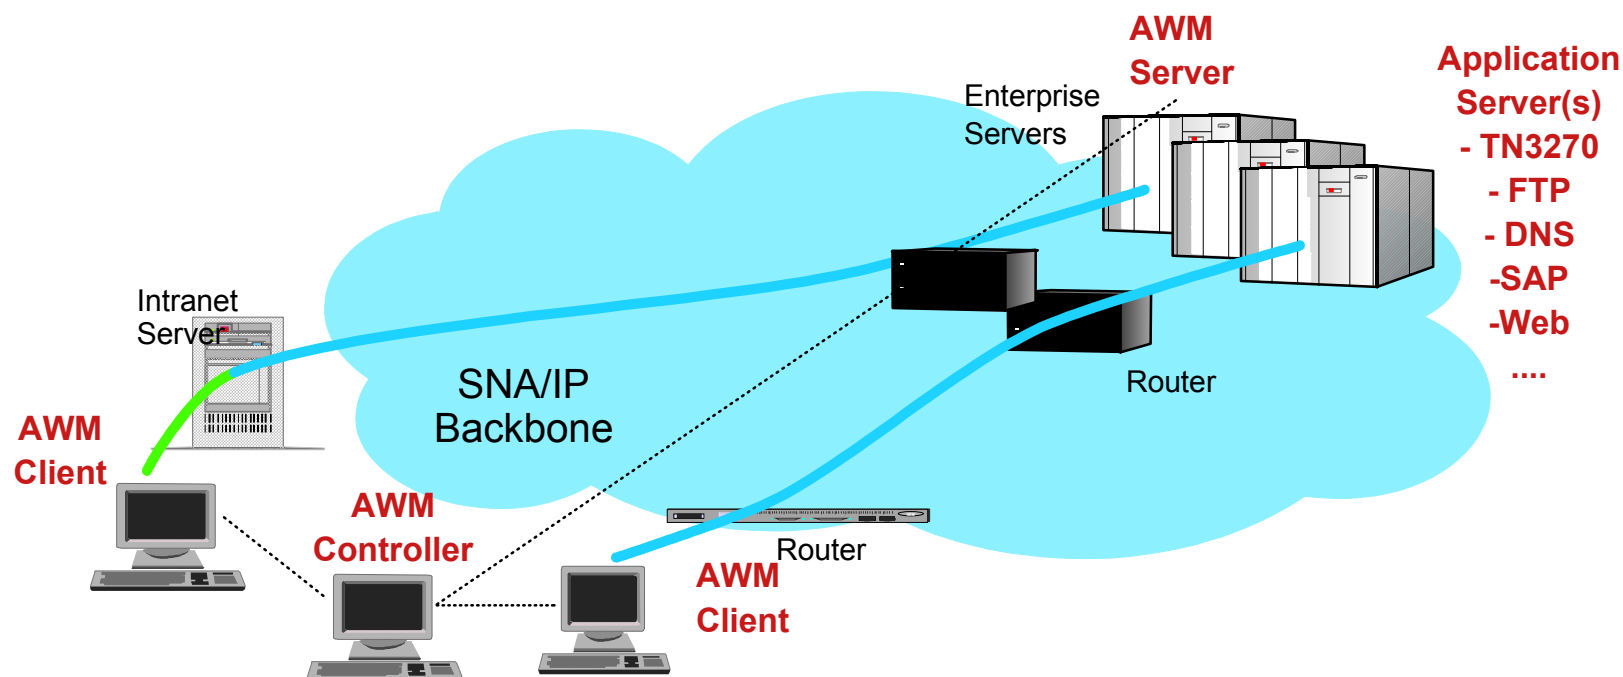


AWM Application Client Mode




Multi-instance Simulation Support

- A single Application Workload Modeler client instance can generate the network traffic equivalent of hundreds/thousands clients
 - ▶ Workload volume bounded only by host and network capacity
- Multiple AWM client instances can collaborate in a simulation test
 - ▶ Level of simulated workload virtually unconstrained
 - ▶ A single AWM instance can act as the *Controller* for the test
 - Single point of control for configuring, executing, monitoring a test across multiple AWM instances
 - Obtains performance statistics from all AWM client instances and provides individual and aggregated reports



Measuring network performance end-to-end




- Flexible Client/Server placement options
 - ▶ Within data center
 - ▶ Between data centers
 - ▶ Between remote users and data center
- Allows testing and measuring of distinct network paths
 - ▶ Comparison of network performance for different sets of users
 - Local and remote end users
 - End users in different branches



Client/Server Benchmarks

AWM Client/Server mode



- ▶ Application specific bottlenecks eliminated
 - Allows true measurement of end-to-end network infrastructure performance
- ▶ AWM simulates both client and server application using standard APIs
 - TCP/IP sockets
 - TCP, UDP sockets
 - Also supports Secure Sockets Layer (SSL)
 - IPv4 and IPv6 support
 - Multicast support
 - SNA APIs (on OS/390 or z/OS only)
 - VTAM RAPI, APPCCMD and HPDT APPCCMD
 - No complicated scripts required!

AWM Client/Server mode...

■ Supported workload models

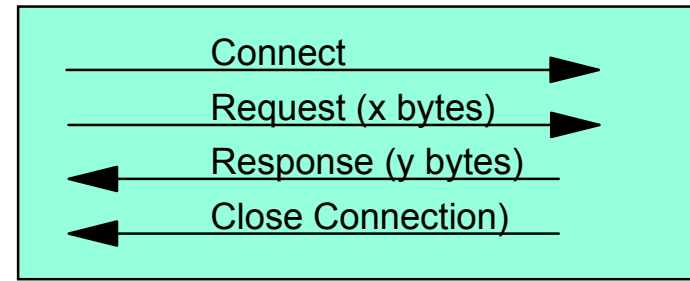
- ▶ Connection-Intensive Workload
 - Connect/Request/Response (CRR)
 - e.g. Web-like traffic
- ▶ Interactive workload
 - Request/Response workload (RR)
 - e.g. Telnet traffic pattern
- ▶ Streaming data
 - Bulk data transfer (STREAMS)
 - in either direction, e.g. FTP traffic

■ Configuration Options

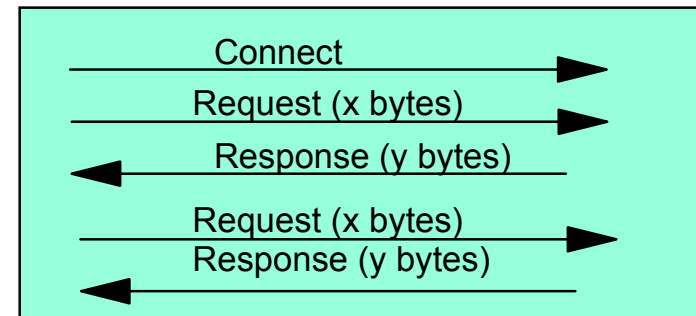
- ▶ Number of clients to be simulated
- ▶ Transaction request rate
 - "think time" for transactional workloads
- ▶ Send/rcv data length

AWM
Client

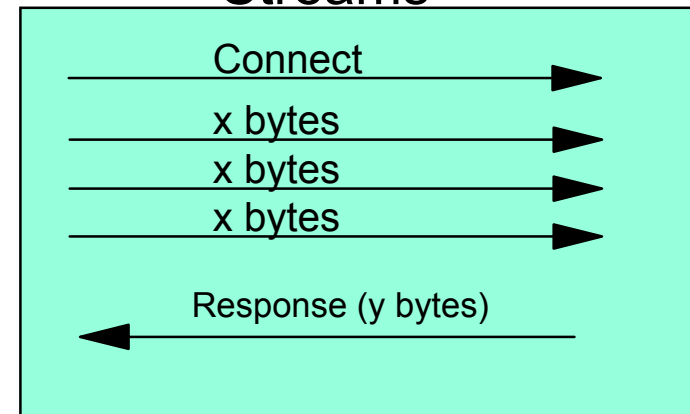
CRR



RR



Streams



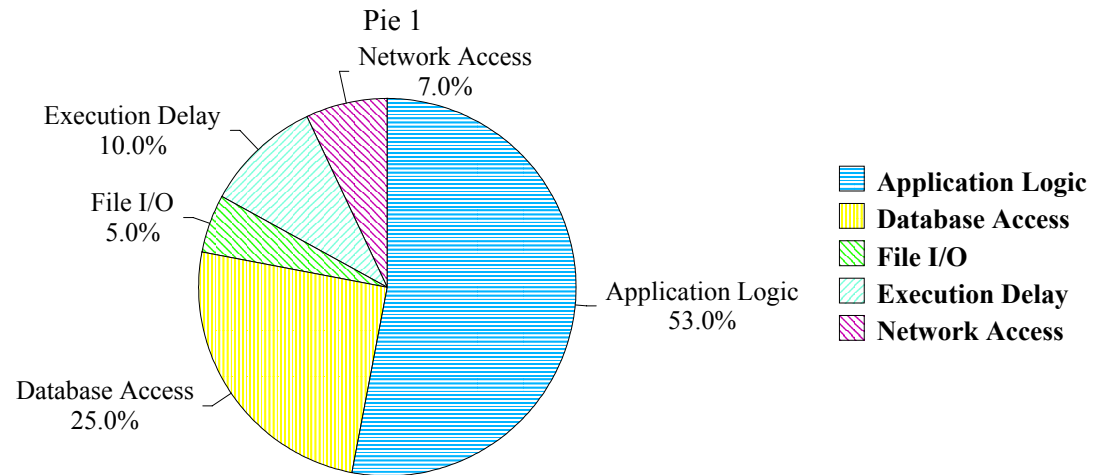
AWM
Server

Client/Server mode...


- Allows measurement of true network infrastructure response time

- ▶ Best/Worst case scenarios:
 - What kind of end-to-end response time is the network capable of?
- ▶ Helps identify bottlenecks
 - Network vs other components (application, CPU, storage, I/O, database access, etc.)
- ▶ Is a network infrastructure upgrade needed?
- ▶ Will a network upgrade have a significant effect on application response time?

Decomposition of application response time



AWM Client/Server Benchmarks Response Time Decomposition



Client/Server mode benchmark reports


SUMMARY OF WORKLOAD PERFORMANCE DATA

Statistics in Time Range: 0 - 86400
Number of samples: 196
Overall Average Transaction Rate: 9911.867521 trans/sec
Overall 95% C.I.of TPS: 9434.94 -10001.9
Overall R.P.of TPS: 2.916977357 %
Overall Maximum Transactions Rate: 10903.10766
Overall Minimum Transactions Rate: 678.2357445
Overall Average Throughput: 11894241.03 bytes/sec
Overall 95% C.I.of Throughput: 1.13219e+07 -1.20023e+07
Overall R.P.of Throughput: 2.916977357 %
Overall Maximum Throughput: 13083729.19
Overall Minimum Throughput: 813882.8934
Overall Average Transaction Response Time: 0.0009984730076
Overall 95% C.I.of TRT: 0.000979699 -0.00101725
Overall R.P.of TRT: 1.880279681 %
Overall Maximum Transaction Response Time: 0.001514039
Overall Minimum Transaction Response Time: 0.000819191

Client/Server mode benchmark reports...

	Transactions per sec	Throughput bytes/sec	Response Time	Relative Precision	Confidence Interval
Overall	9911.87	1.18942e+07	0.000998473	1.88028	1.877409e-05
Session:00	1048.64	1.25836e+06	0.00095365	6.38787	6.091796e-05
Session:01	1034.76	1.24171e+06	0.00096643	3.79586	3.668428e-05
Session:02	1013.16	1.21579e+06	0.000987045	5.49161	5.420467e-05
Session:03	944.128	1.13295e+06	0.00105921	7.32882	7.762762e-05
Session:04	968.268	1.16192e+06	0.0010328	6.75297	6.974468e-05
Session:05	908.377	1.09005e+06	0.0011009	7.63488	8.405207e-05
Session:06	1101.59	1.3219e+06	0.000907825	5.76292	5.231723e-05
Session:07	1068.82	1.28259e+06	0.000935632	5.42444	5.075286e-05
Session:08	1042.79	1.25134e+06	0.000959001	6.16016	5.907593e-05
Session:09	997.592	1.19711e+06	0.00100246	7.75387	7.772938e-05

Enterprise Extender Modeling

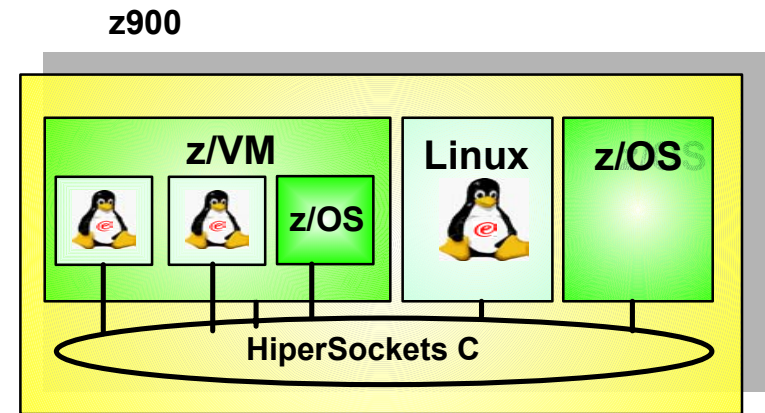


- **AWM can be used to model Enterprise Extender connectivity between data centers or business partners**
 - ▶ Withdrawal of 3745/6 as well as the high speed and reduced networking cost of OSA Express makes migration to Enterprise Extender very attractive
 - AWM can assist in network tuning, ensuring proper Quality of Service for SNA traffic across IP network, and capacity planning
- **Future feature of AWM will be the ability to model Enterprise Extender from the branch to the data center**

Server Consolidation - HiperSockets Modeling


■ HiperSockets Overview

- ▶ High speed, low latency, any-to-any TCP/IP network within a z900 processor
- ▶ Among virtual servers and LPARs (z/OS, Linux and z/VM)
- ▶ Cost savings - no adapters, network boxes, or cabling
- ▶ Transparent to applications

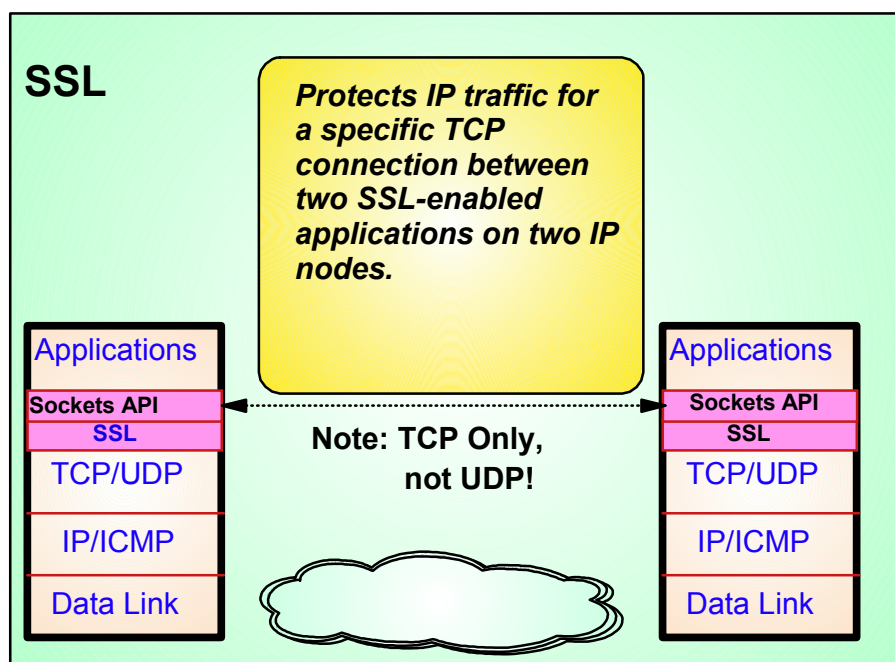
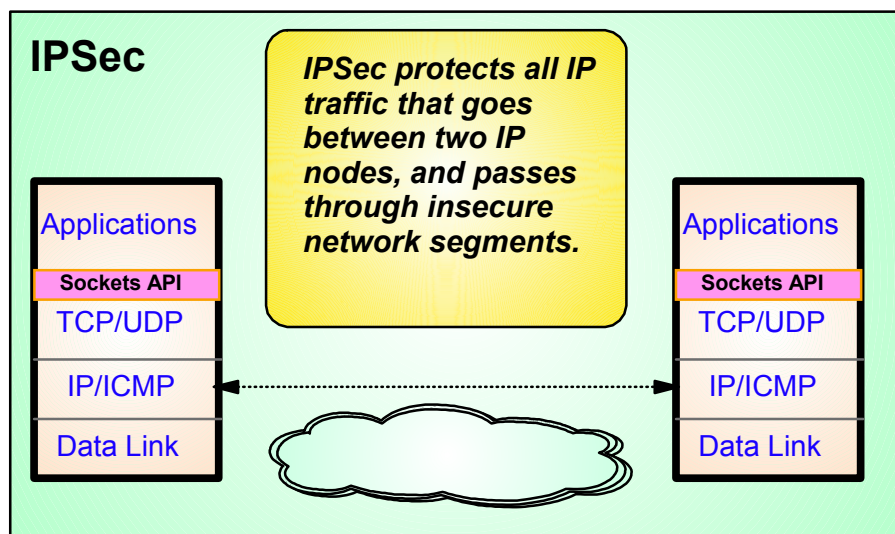


■ AWM can be used to model and measure the performance characteristics of various workloads over HiperSockets

- ▶ Host to host communications within a zSeries processor
- ▶ Server Consolidation Scenarios
 - Multiple Linux for zSeries middle tier servers accessing z/OS enterprise data




Performance Evaluation of IPSec/SSL




- AWM can be used to model and measure the effect of IPSec/SSL prior to production deployment
 - ▶ Generate predictable workload
 - through a VPN (Tunnel or Transport mode)
 - Client/Server or Application Client mode
 - Using SSL
 - Client/Server mode
 - GSSKIT and OPENSSL support (Linux)
 - System SSL support (OS/390 and z/OS)
 - ▶ Measure effects of IPSec/SSL on
 - End user response time
 - Throughput
 - Host Requirements (i.e. CPU)
 - Router utilization


IPv4 to IPv6 Internet Evolution




Today



Stage 1

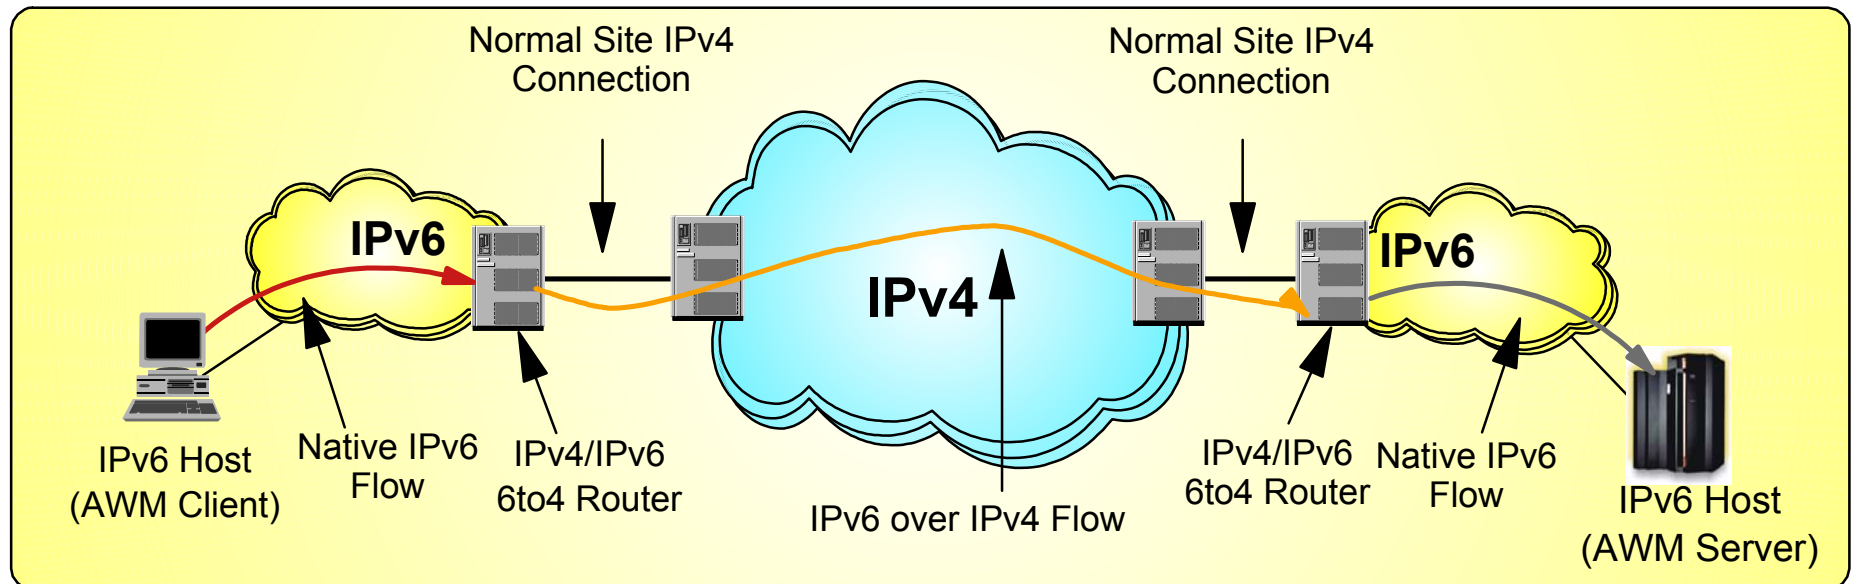


Stage 2



Stage 3

IPv6 - Modeling the performance impact

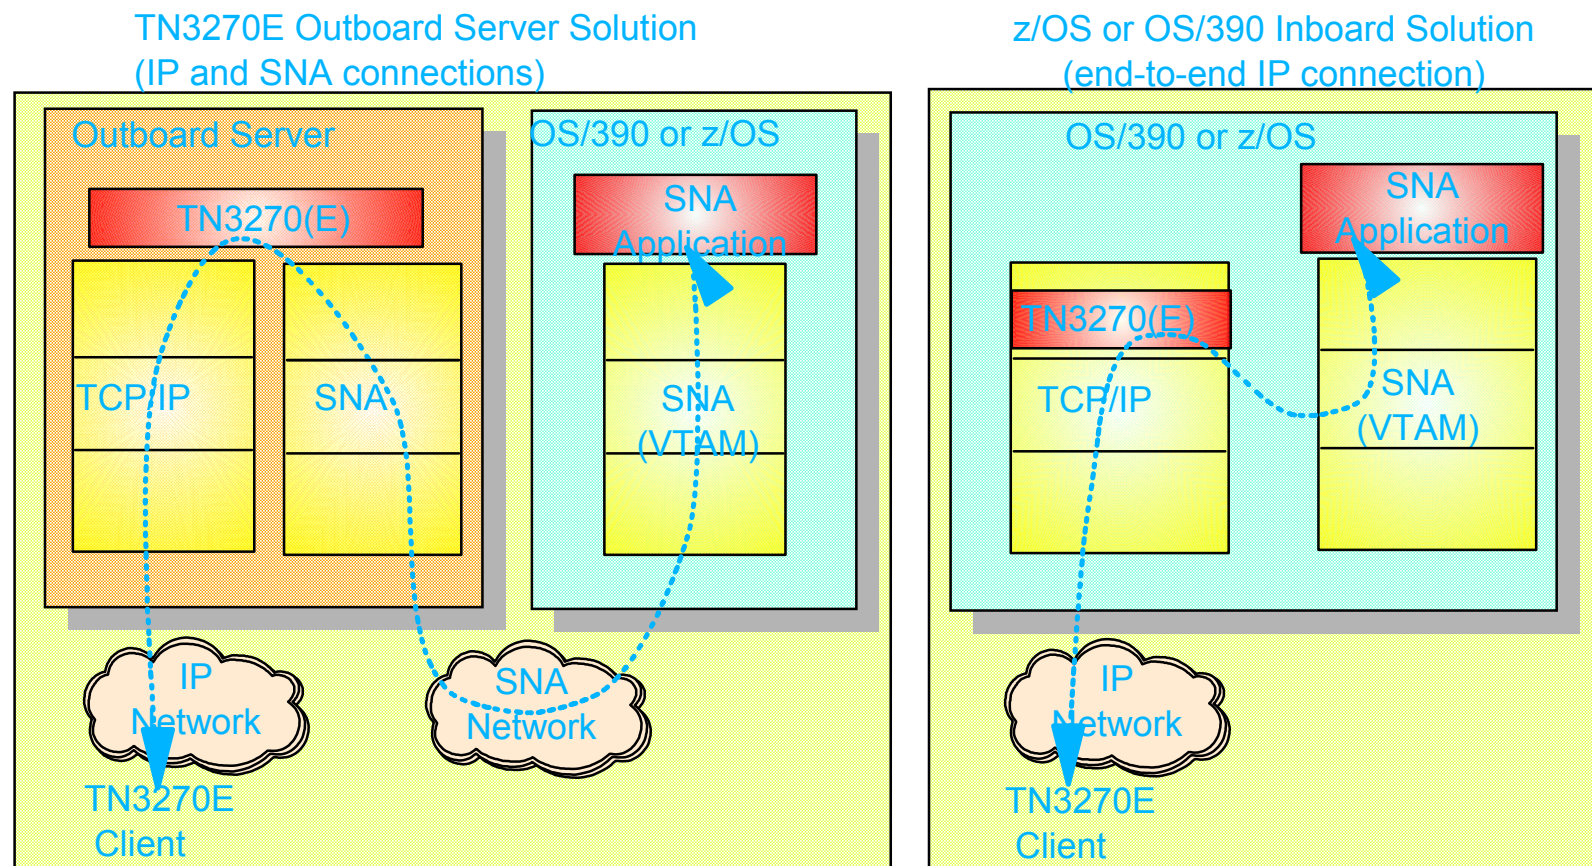


- Tunneling: encapsulating an IPv6 packet in an IPv4 packet and send the IPv4 packet to the other tunnel end-point IPv4 address.
 - ▶ The tunnel endpoint placement depends on connectivity needs
 - Placing endpoints in routers allows entire sites to be connected over an IPv4 network
 - Placing endpoints in hosts allows access to remote IPv6 networks without requiring updates to the routing infrastructure
- AWM can help model and measure the effect of IPv6 network traffic
 - ▶ Network response time for IPv6 application traffic
 - Native and/or through IPv6/IPv4 tunnels
 - ▶ Effect on IPv4 networks and existing IPv4 applications



Application Client Benchmarks


TN3270(E) Server Positioning



- Application Workload Modeler can be used to simulate TN3270(E) client activity
 - ▶ Capable of simulating large number of TN3270(E) clients/activity
 - Used internally to benchmark 60,000 client sessions to a single z/OS TN3270 server
 - ▶ Can be used to benchmark outboard vs inboard TN3270(E) server solution
 - Determine cost/performance/scalability characteristics of each solution
 - ▶ Aids in decision making process for TN3270(E) server placement

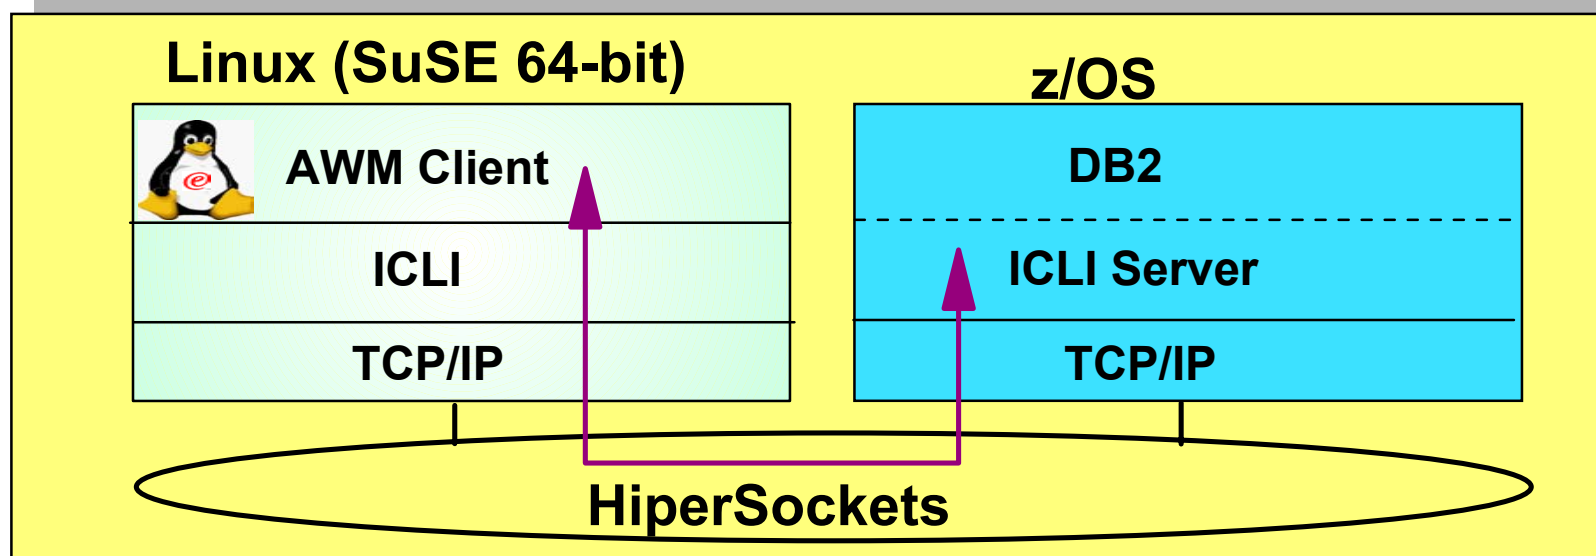
TN3270(E) Client Simulation...

- Additional TN3270(E) Client Simulation Capabilities
 - ▶ TN3270(E) SSL Client Simulation
 - Analysis of SSL performance, cost, overhead
 - ▶ Simulation for IBM TN3270(E) Value Add extensions such as SSL Express Logon Support
 - PKI-Based Identification and Authentication (Single Signon Support)
 - Certificate provides SNA session verification and Logon
 - Supported in HOD V5, PCOMM V5.5
 - ▶ Allows evaluation of these technologies and proper capacity planning prior to deployment



SAP R/3 ICLI benchmarks

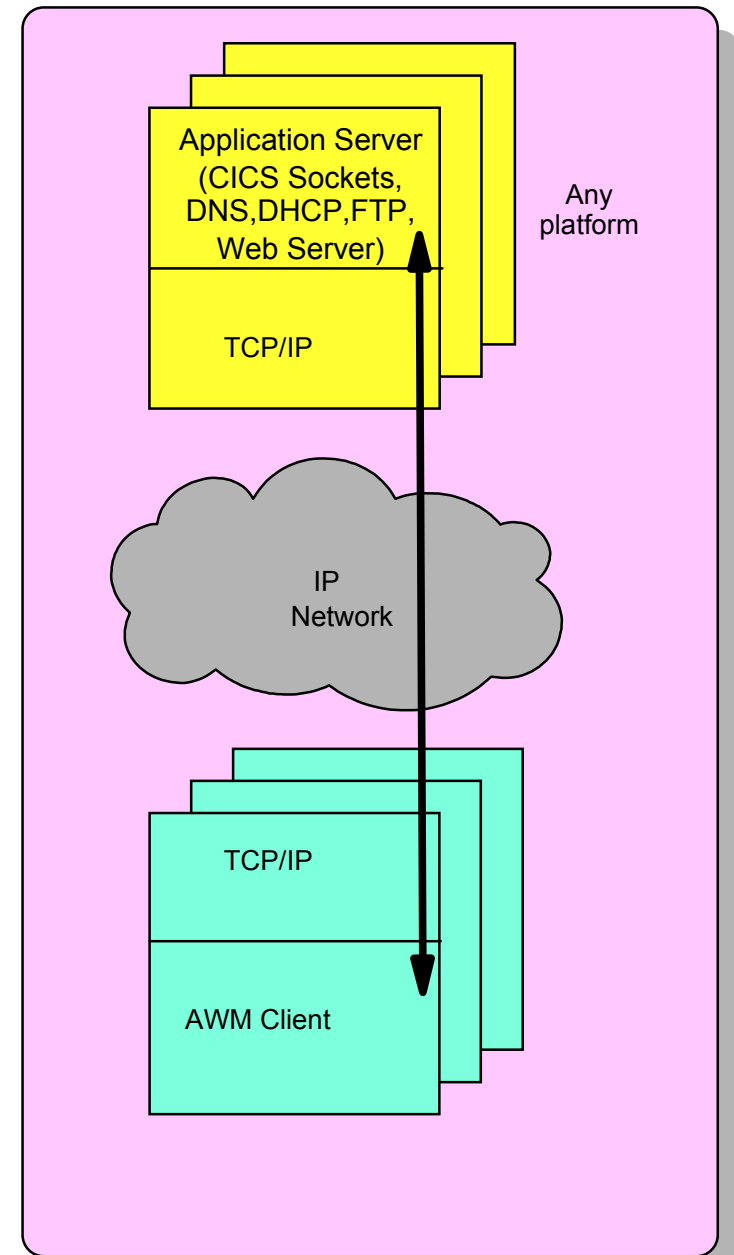
zSeries



- SAP R/3 communications modeling
 - ▶ Model Application Server to Database Server communications
 - Using SAP R/3 Integrated Call Level Interface
 - Drives the entire Communications path
 - ▶ Allows simulation of multiple Application Servers
 - Each servicing hundreds/thousands Client requests
 - ▶ Useful in evaluating SAP R/3 consolidation on zSeries
 - Detailed response time metrics, throughput rates


Additional Application Client Mode Functions

- CICS Sockets
 - ▶ Simulate client traffic to TCP/IP CICS sockets application
- Domain Name System (DNS)
 - ▶ Drive hostname resolution processing
 - ▶ Customized list of hostnames
- Dynamic Host Configuration Protocol (DHCP)
 - ▶ Simulate client DHCP requests
- File Transfer Protocol (FTP)
 - ▶ Simulate FTP Client requests
 - inbound/outbound file transfers
- Simple Mail Transfer Protocol (SMTP)
 - ▶ Simulate SMTP client requests
- Web Server
 - ▶ Simulates Web Server client requests
 - Customized list of URLs
- Customized statistics produced for each workload type



Policy / QoS Modeling


- Applications/workloads have unique SLA objectives
 - ▶ Priority should extend beyond Server resources (CPU, storage, etc.)
 - ▶ Network traffic can be prioritized using Differentiated Services (Quality of Service - QoS)
- AWM can be used to model and measure the effect of a QoS policy prior to deployment
 - ▶ Does the QoS policy have the desired effect?
 - ▶ What is the impact on traffic that is not included in QoS policy?
 - ▶ Helps answer what/if questions when tuning a QoS policy



QoS WebSphere Measurement Example 1

Transient Behavior of User Response Time (WebSphere)

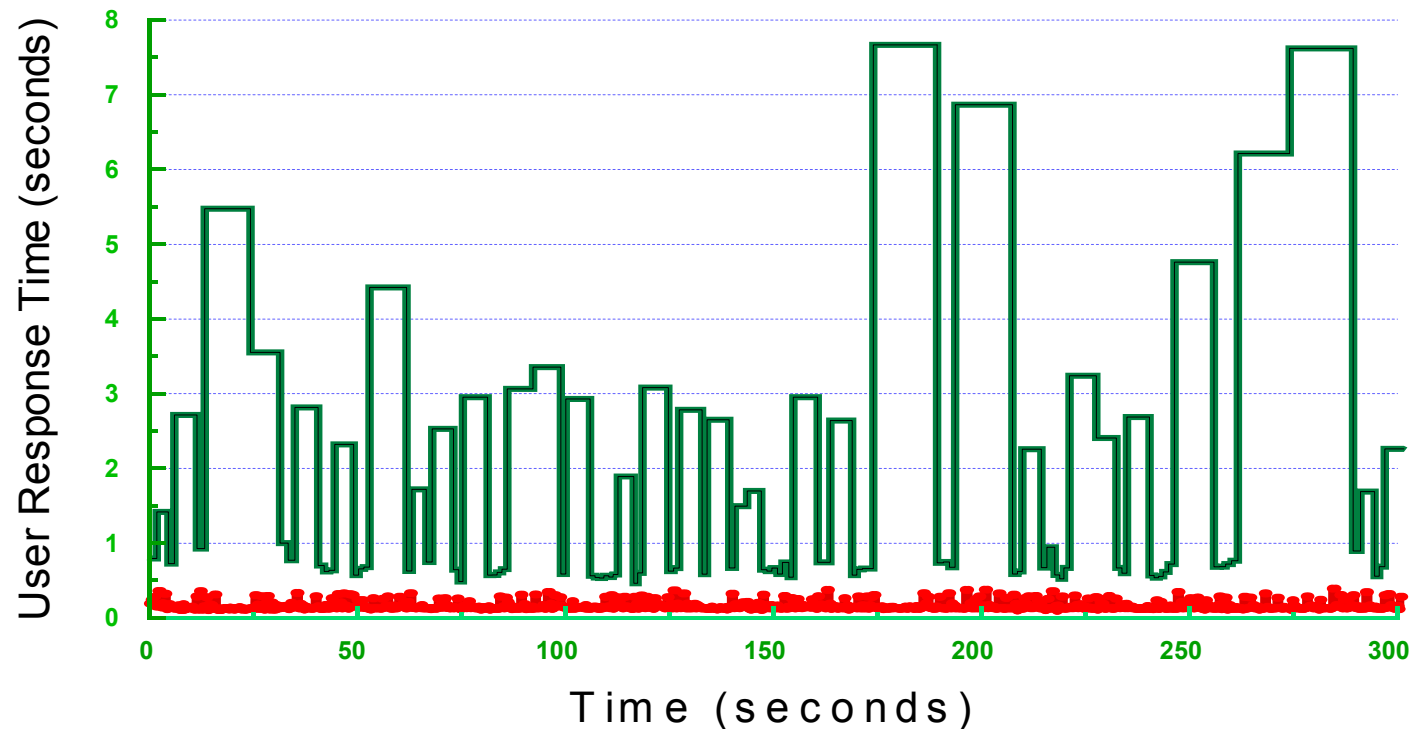
- o Network DiffServ Enabled (CBWFQ)
- o WebSphere PA and QoS Support NOT enabled



QoS WebSphere Measurement Example 2

Transient Behavior of User Response Time (WebSphere + IBM HTTP Server for OS/390)

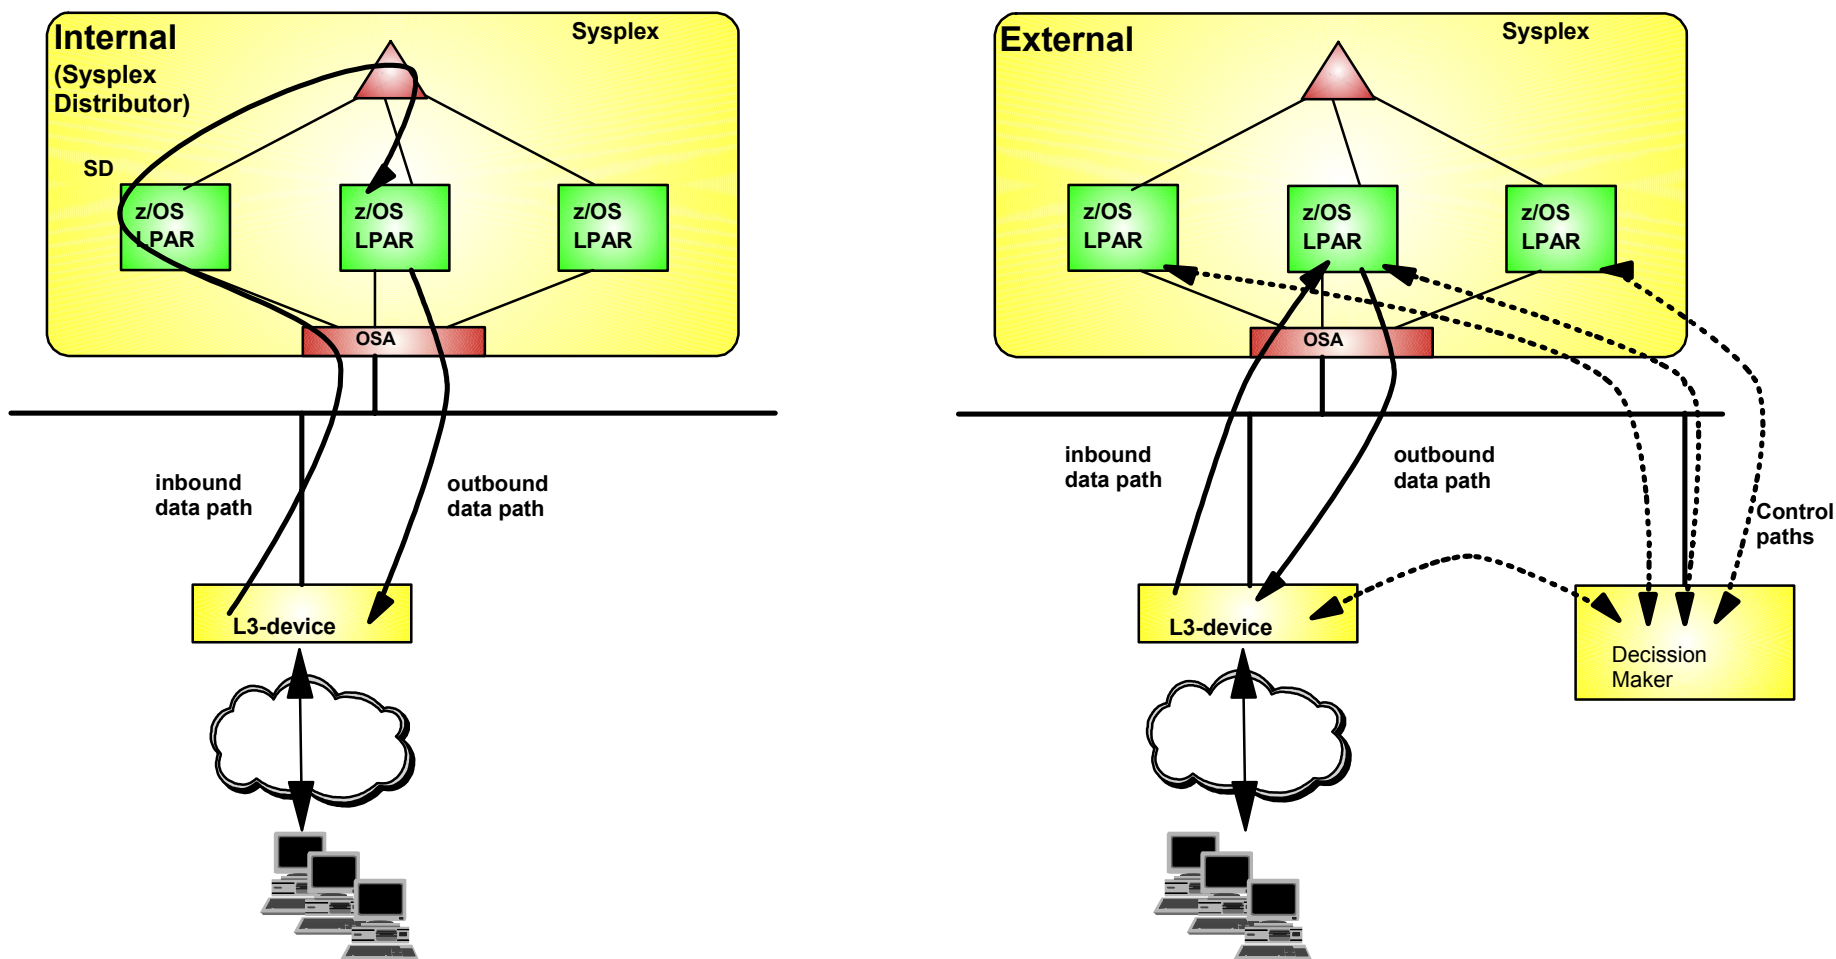
- o Network DiffServ Enabled (CBWFQ)
- o WebSphere PA and QoS Support Enabled



- Higher Priority Access (ToS = 5); Average Response Time = 157 msec (98-393 msec)
- Lower Priority Access (ToS = 0); Average Response Time = 1506 msec (465-7669 msec)

**Application Workload Modeler use in demonstrating
the value of Integrating WebSphere with Network QoS**

Modeling Load Balancing Solutions





- AWM can be used to model and measure the effect of a Load Balancing solution prior to deployment
 - ▶ Capacity planning (How many target servers are needed?)
 - ▶ Network Impact
 - ▶ Effect on end-user response time

AWM usage in z/OS measurements

■ Client/Server Benchmarks (TCP)

- ▶ z/OS V1R4 CS relative to V1R2
- ▶ Comparison - Trans/sec, CPU Cost Per Transaction



- V1R4 provides Transaction rates within -0.66% to 17.85 % of V1R2
- V1R4 provides lower CPU cost per transaction than V1R2 and the percentages are within (-0.62 to -6.43)

■ AWM used extensively internally for performance measurements

- ▶ For example, refer to the "z/OS V1R4 Communication Server Performance Summary", at:
 - <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS582>

IBM Services for AWM

- Many services offerings are available
 - ▶ Try and buy
 - Evaluate the AWM product in your environment while benefiting from IBM consultant's expertise
 - ▶ Capacity projection and analysis
 - IBM will use the AWM product and IBM systems to perform complex modeling. We will provide a report documenting the results along with tuning recommendations and capacity projections.
 - ▶ Installation, training and configuration workshop
 - A workshop to enable rapid deployment, and to help your staff get the most out of AWM.
 - ▶ Custom offerings
 - IBM is able to build custom offerings tailored to your specific needs.
 - Examples could include
 - 3745/6 to APPN Enterprise Extender migration
 - Projecting the impact of data center or server consolidation
 - Estimating network and system requirements for a new application

For more information...

- See your IBM client representative or call IBM DIRECT at 1-800-IBM-CALL in the US and Canada.
- Application Workload Modeler web site:
<http://www.ibm.com/software/network/awm/index.html>
 - Additional Product information
 - Announcement letter
 - Product Documentation
 - Links to other relevant sites
- Or send an e-mail with any questions to ***awm@us.ibm.com***