IBM Software Group

Application Workload Modeler for z/OS and Linux on zSeries
Technical Overview

e business on demand software

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 addresse 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 e-Business 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.)?
    - Service Level Agreements for new and changed workloads must be met
      - Degraded performance is often perceived as "down time"
- Uncertainty often inhibits change
  - For example, considering SSL to protect the communications of some key applications
    - What will be the impact of this change
      - On end user response time?
      - On system resources (CPU, etc.)?
      - On network traffic?

Network Modeling and Simulation Tools

- Network modeling and simulation tools provide an inexpensive solution
  - Changes can be evaluated before they are deployed
  - Detailed performance data is available for accurate capacity planning
  - "What if" questions can be analyzed while still in the planning phase
- Several approaches to modeling network performance
  - Simulation modeling
    - Computer program simulates network flows to predict future behavior
  - Analytic Modeling
    - Computer program uses mathematical models to predict future network behavior
  - Modeling through load testing and performance measurements
    - Generate **real** network traffic that represents existing/new workload conditions
      - Measure results and provide detailed performance data

Application Workload Modeler (AWM): Overview

- An application workload simulation tool for measuring and modeling the performance of the network infrastructure, end-to-end
  - Including the network, systems, and applications
  - Generates real network traffic and provides detailed performance measurements

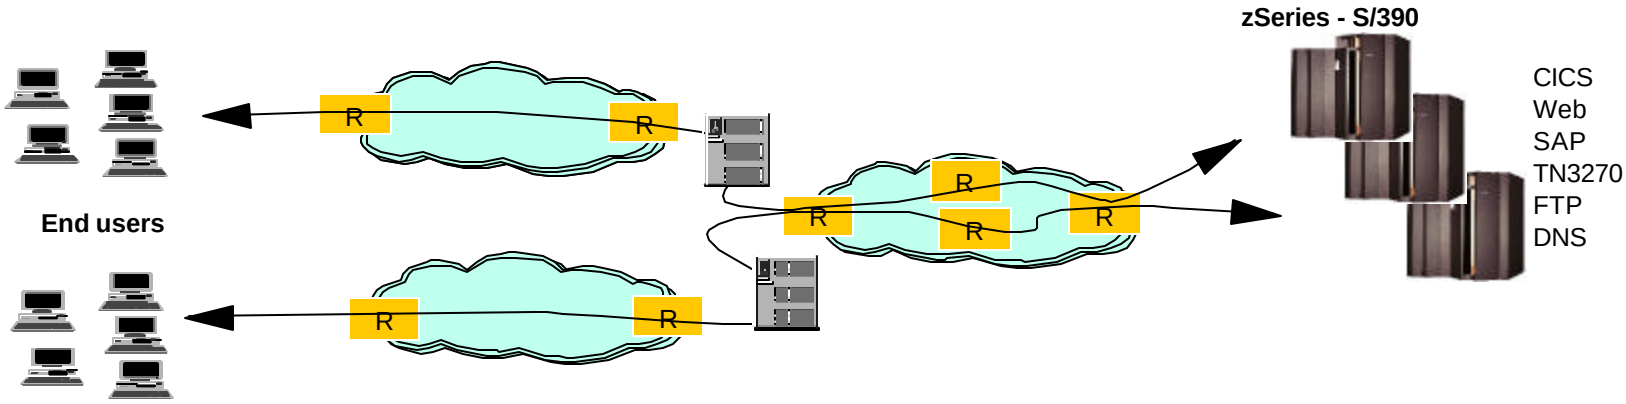

- Allows enterprise or service providers to model the impact of various types of workloads on servers and networks
- Users can evaluate the impact of a change before the change is deployed in a production environment
  - Do you need to upgrade network components to meet performance objectives?
  - Is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

Questions, Questions, Questions

- What is the impact on performance when making changes?
  - Network infrastructure changes
    - Enterprise Extender, Virtual Private Networks (VPN), Quality of Service (QoS)
    - 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 and software changes
    - Upgrading hardware or software
    - Considering a load balancing solution
    - Server consolidation
    - Deploying a TN3270E solution

Key Features

- Generates real network traffic
  - Representative of real application network patterns
  - Tests the network end-to-end
  - 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
- Provides detailed statistics
  - Per "session" as well as cumulative averages for Throughput, Response time, Transaction rate
  - Transient behavior of workloads
  - Helps evaluate the performance of workloads

Client/Server Mode Benchmarks

- Generates ***real*** network traffic that simulates communications for both client and server 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 "network behavior" of new application
    - ***Prior*** to application development/deployment
  - Application-specific bottlenecks eliminated
  - SNA and TCP/IP support

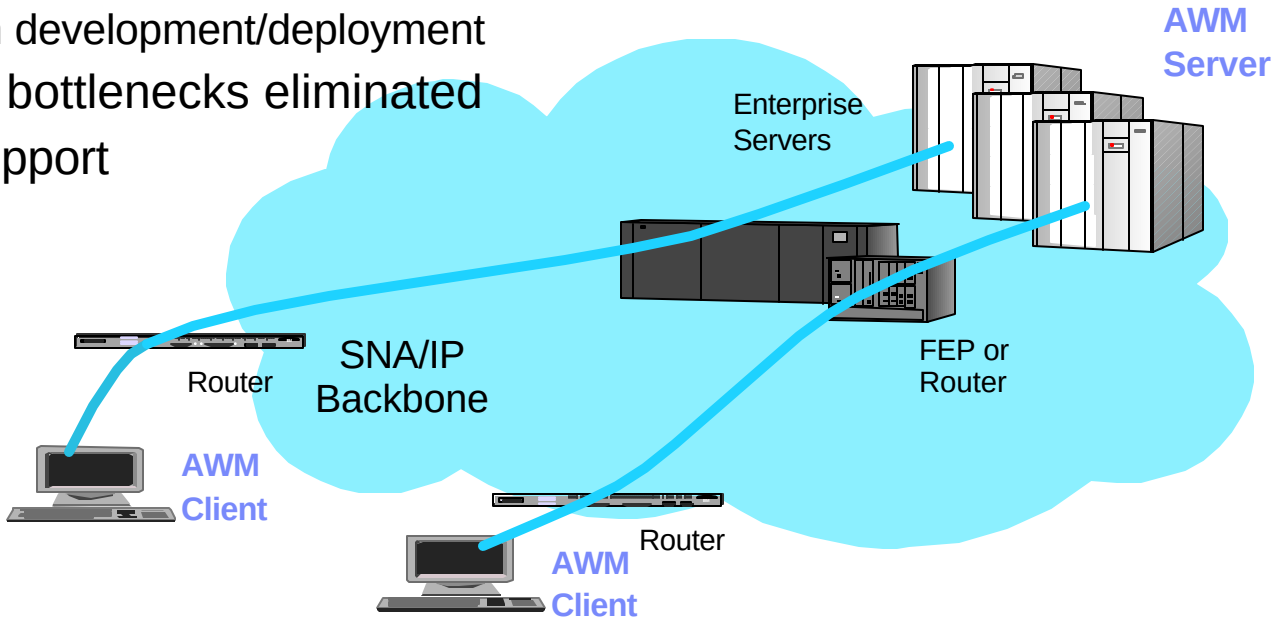

Application Client Mode Benchmarks

- Generates ***real*** network traffic that simulates client communications for well known, standard TCP/IP server applications
  - Focuses 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 ICLI, 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?

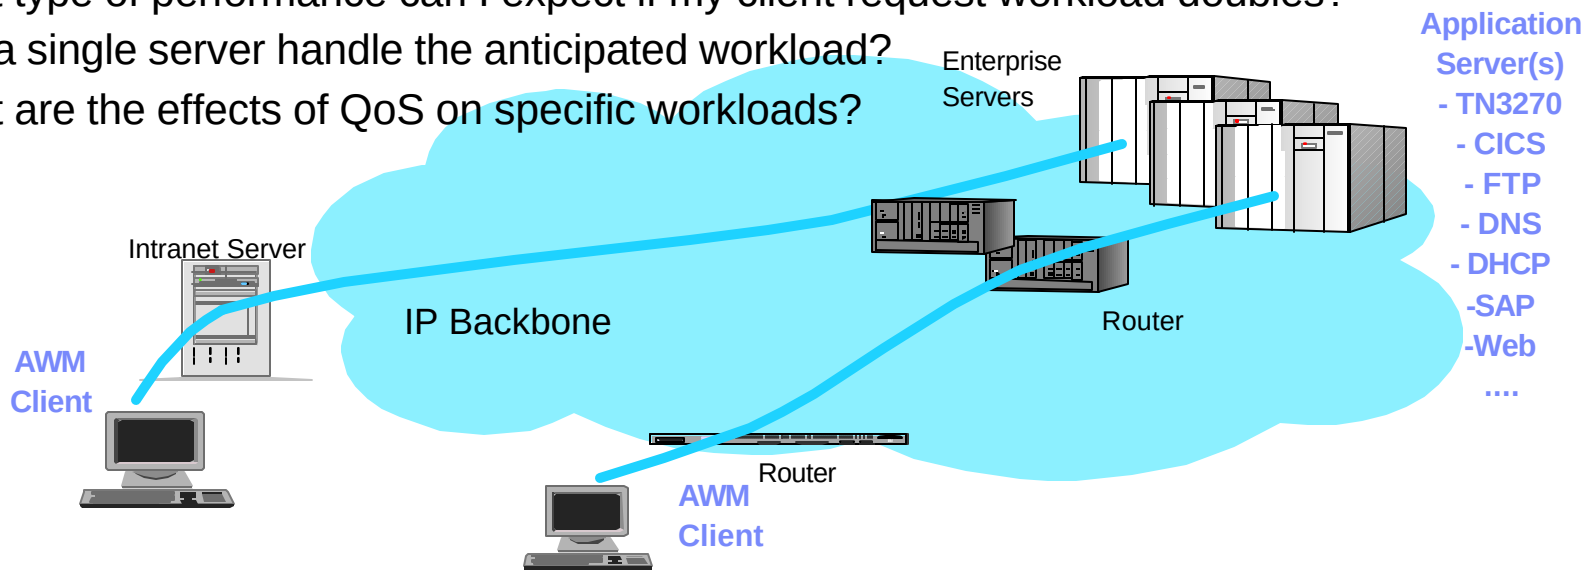

Application Workload Modeler Statistics

- 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.

Multiplatform Support

- AWM for z/OS
  - Supports current z/OS platforms and OS/390 V2R10
- AWM for Linux on zSeries supports
  - Red Hat Linux for S/390, or later
  - SuSE Linux Enterprise Server 7 for S/390 and zSeries, or later
- AWM for Linux supports the following Intel distributions
  - Red Hat 7.1, 7.2, 7.3, SuSE 7.3

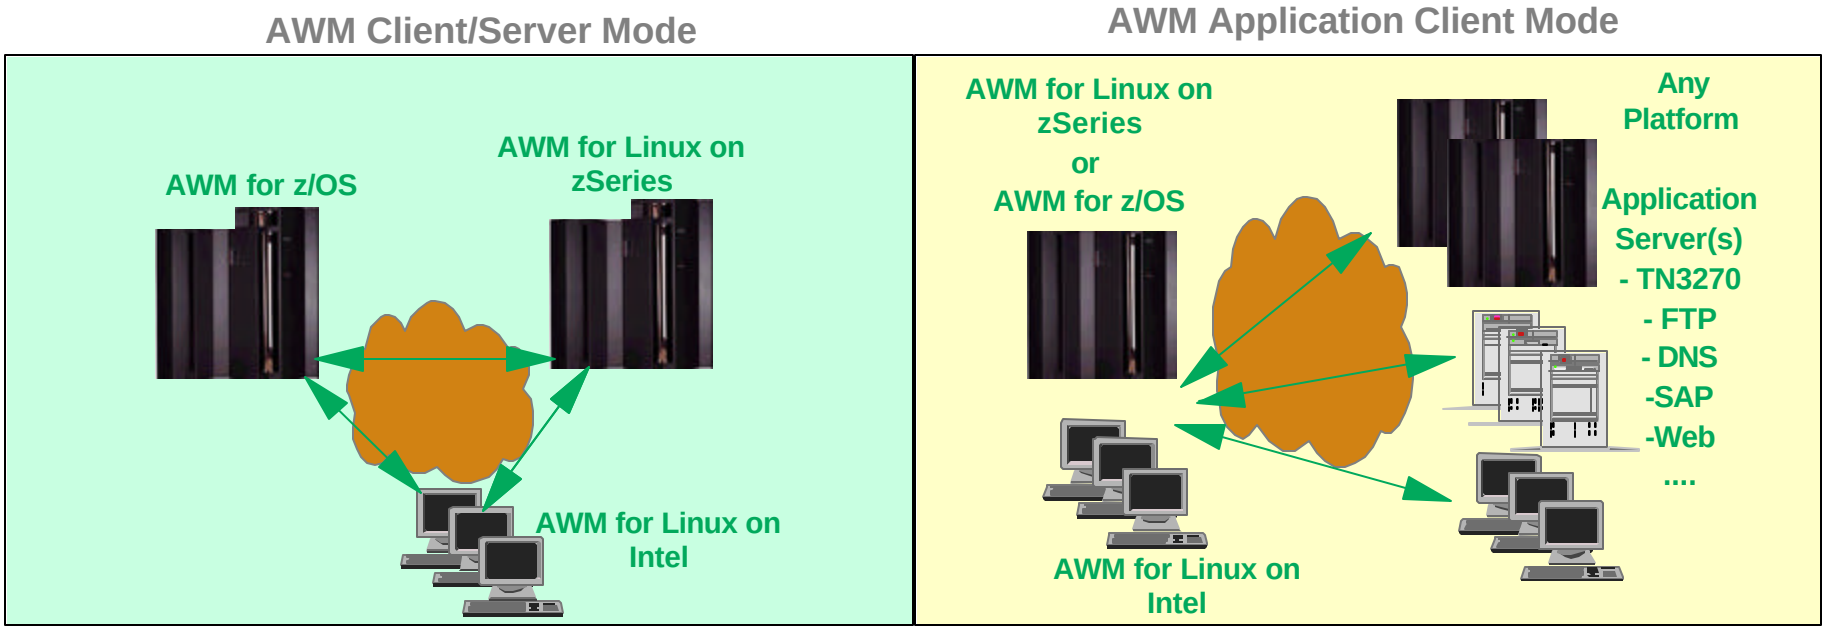

Multi-instance Simulation Support

- A single AWM client instance can generate the network traffic equivalent of hundreds/thousands of 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
    - Configures, executes, monitors the test
    - Obtains performance statistics
    - Provides individual and aggregated reports

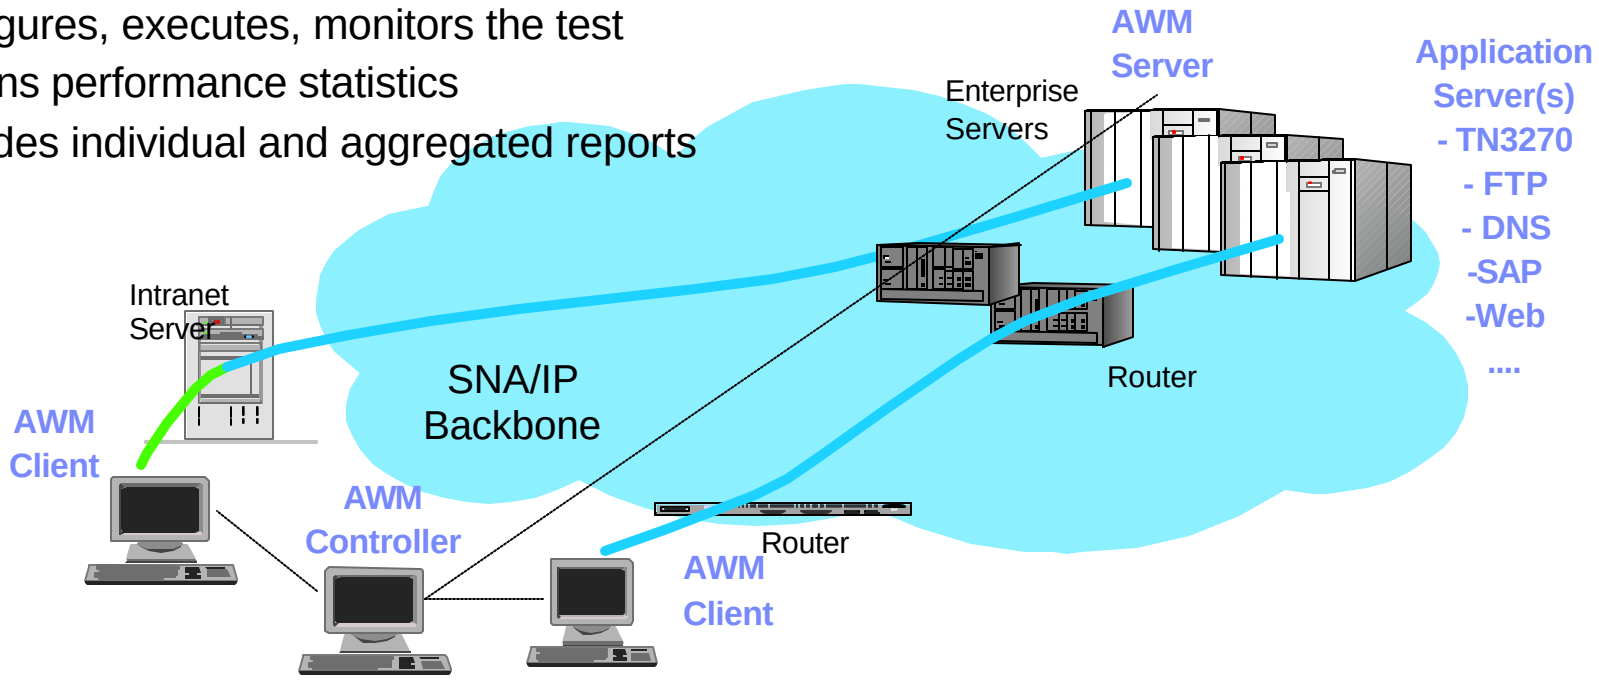

Measuring Network Performance End-to-End

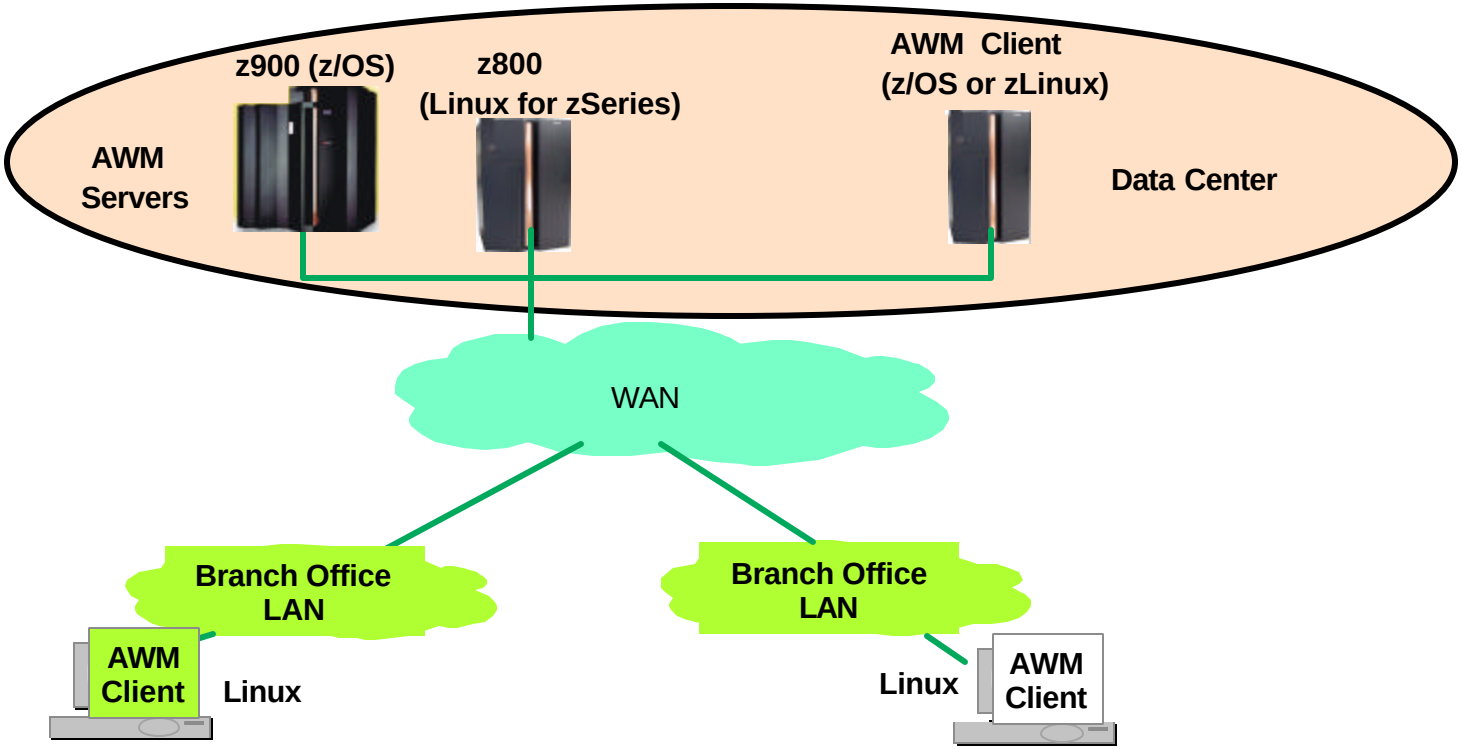

- Flexible Client/Server placement
  - Within data center
  - Between data centers
  - Between remote users and data center
- Allows testing and measuring of distinct network paths
  - Comparison of network performance for different sets of users
    - Local and remote end users
    - End users in different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

Client/Server Mode

Details and Examples

AWM Client/Server Mode

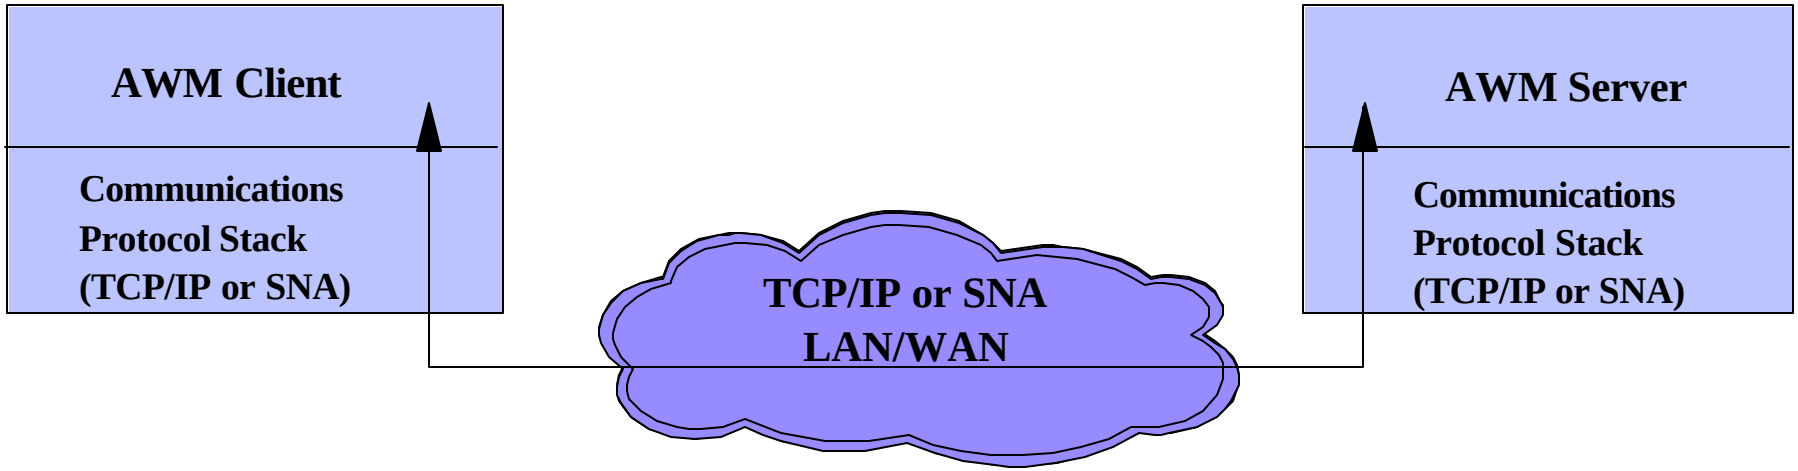

- Application-specific bottlenecks eliminated
  - Allows true measurement of end-to-end network infrastructure performance
- Simulates both client and server application using standard APIs
  - TCP/IP sockets
    - TCP, UDP sockets, including SSL
    - IPv4 and IPv6 support
    - Multicast support
  - 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/receive data length

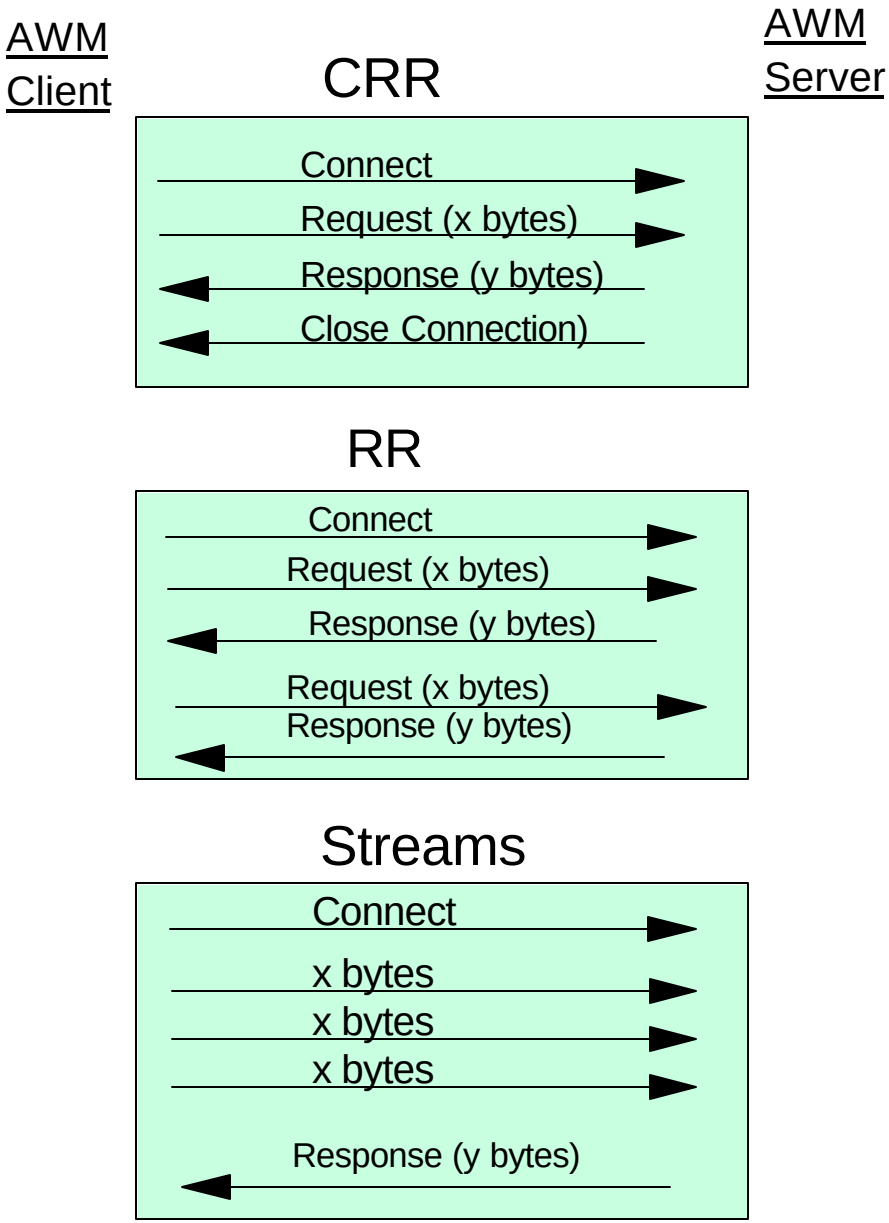

Client/Server Mode...

- 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

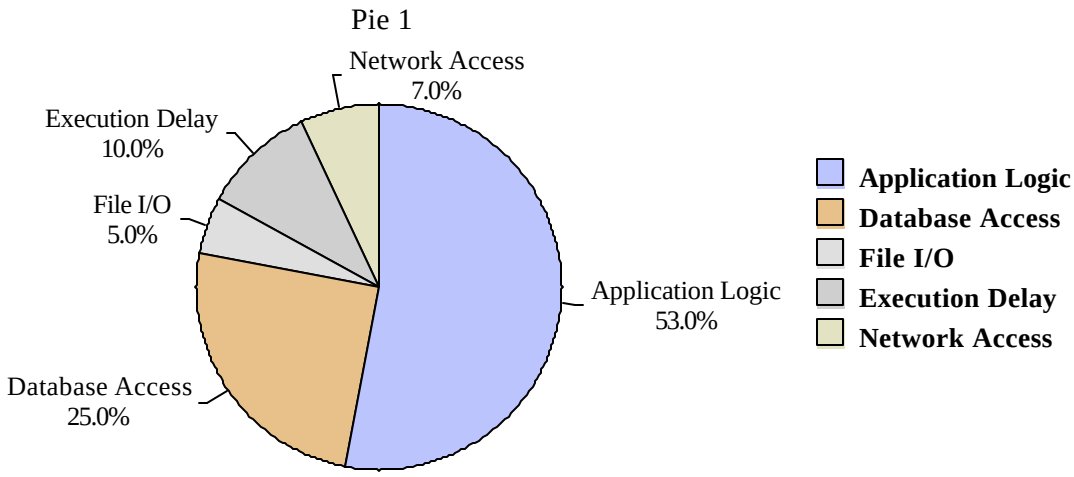

AWM Client/Server Benchmarks
Response Time Decomposition

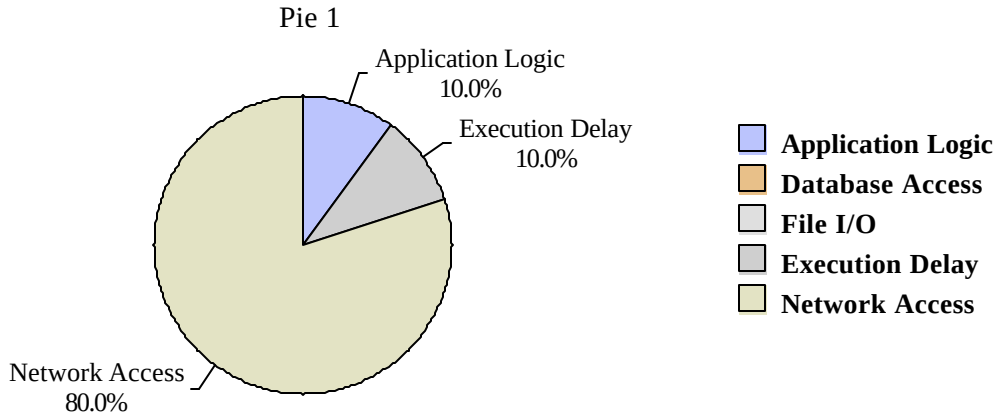

AWM Usage in IBM z/OS Measurements

- Client/Server Benchmarks (TCP)
  - z/OS V1R4 CS relative to V1R2
  - Comparison - Trans/sec, CPU Cost Per Transaction

Transactions/Second
V1R4 vs.. V1R2 (% Delta)

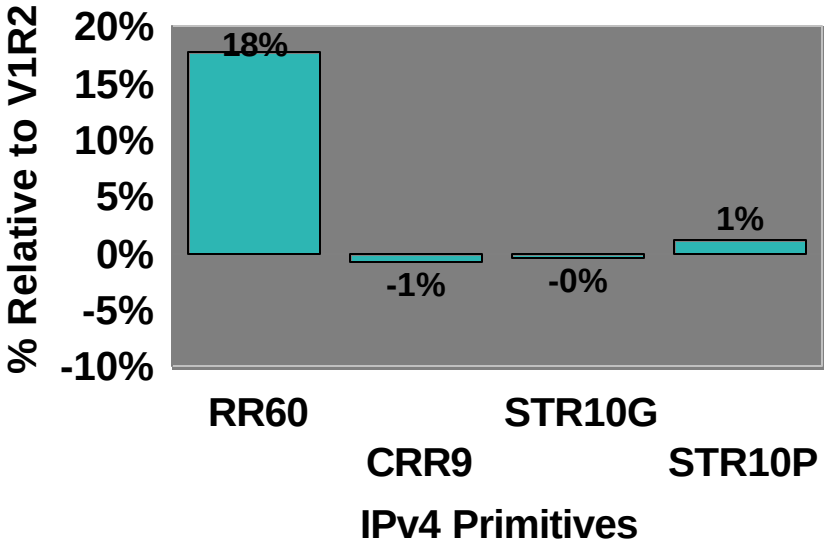

CPU Cost Per Transaction
V1R4 vs.. V1R2 (% Delta)

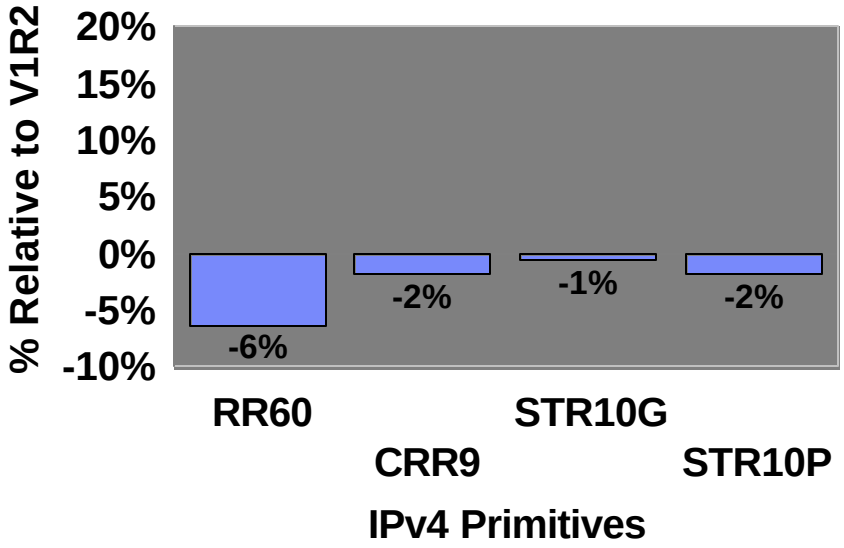

- V1R4 provides Transaction rates within -0.66% to 17.85 % of V1R2
- V1R4 provides lower CPU cost per transaction than V1R2 and the percentages are within (-0.62 to -6.43)

- AWM used extensively internally for performance measurements
  - Example in performance report on AWM Web site

Example 1: Enterprise Extender Modeling

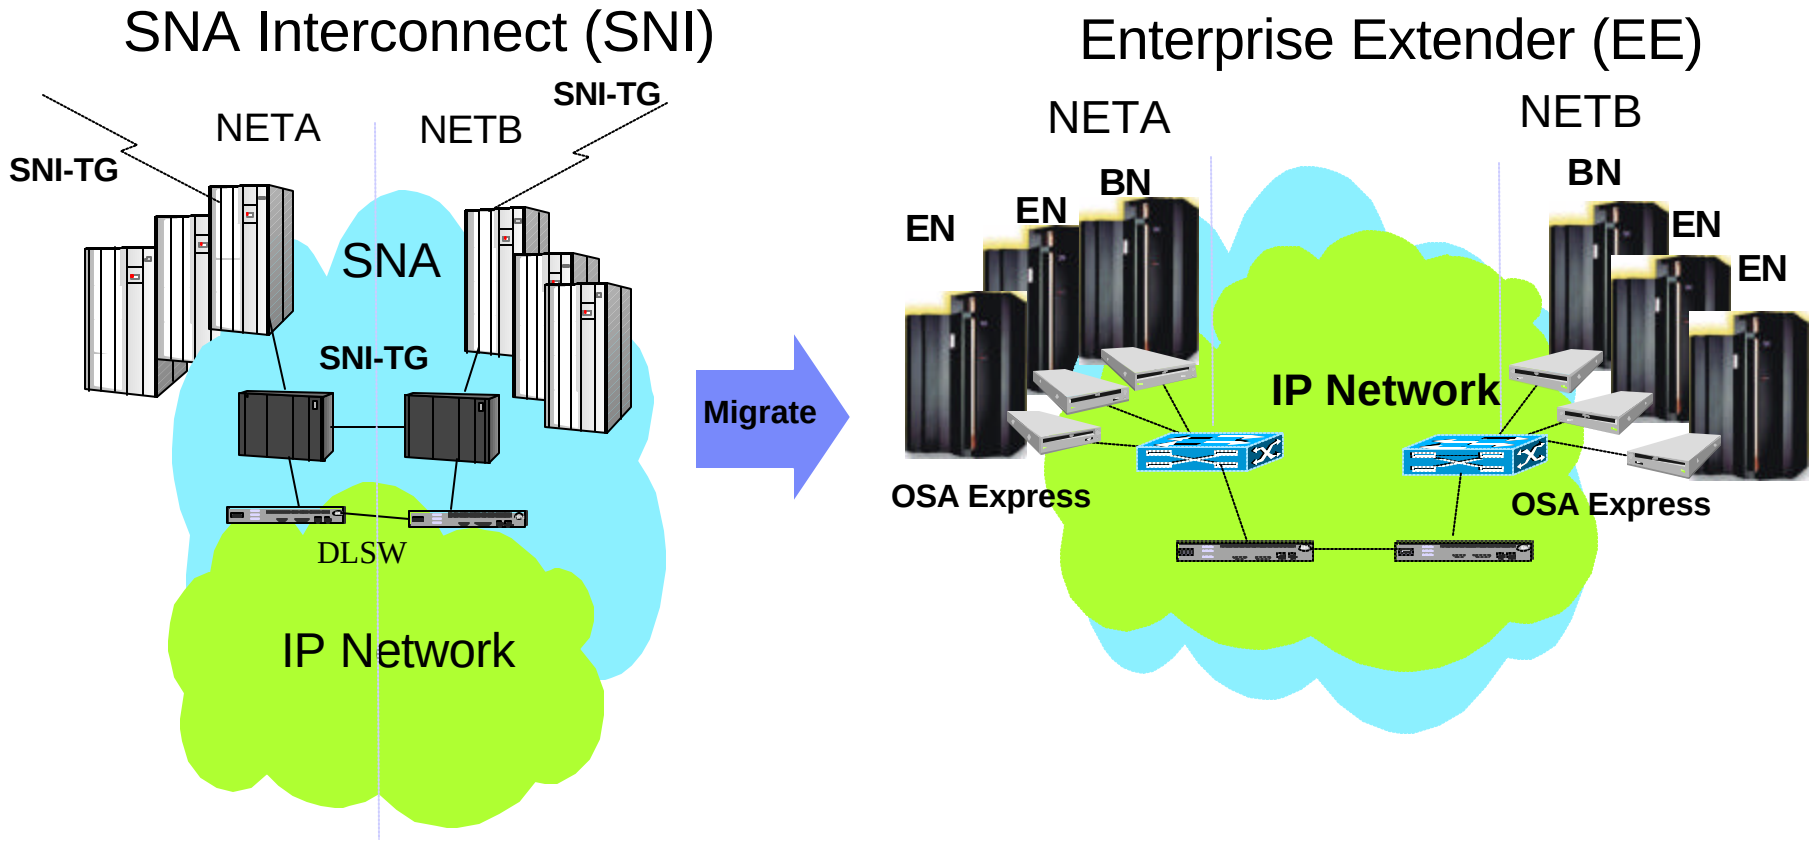

- Model EE connectivity between data centers or business partners
  - Migration to OSA Express and EE attractive
    - Withdrawal of 3745/6
    - High speed and reduced networking costs
  - AWM assists in network tuning, capacity planning, ensuring Quality of Service
    - SNA traffic across IP network

Example 2: Server Consolidation with HiperSockets

- 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

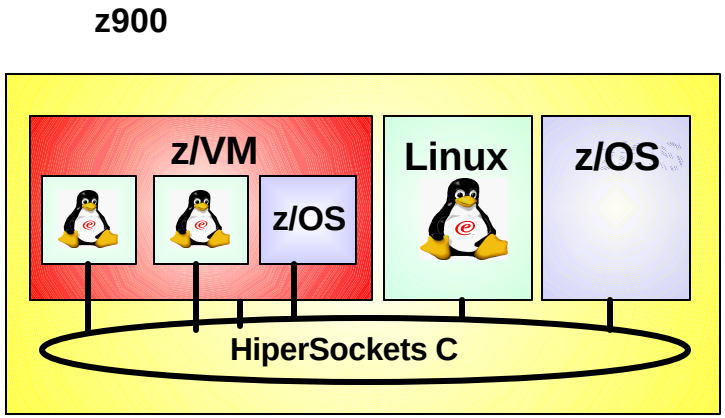

- How will my applications perform over HiperSockets vs.. existing external network?
  - Use Client/Server mode to simulate existing application workloads
  - Run on traditional external network and on z900 processor with HiperSockets
  - Determine cost, performance, and scalability characteristics of each solution

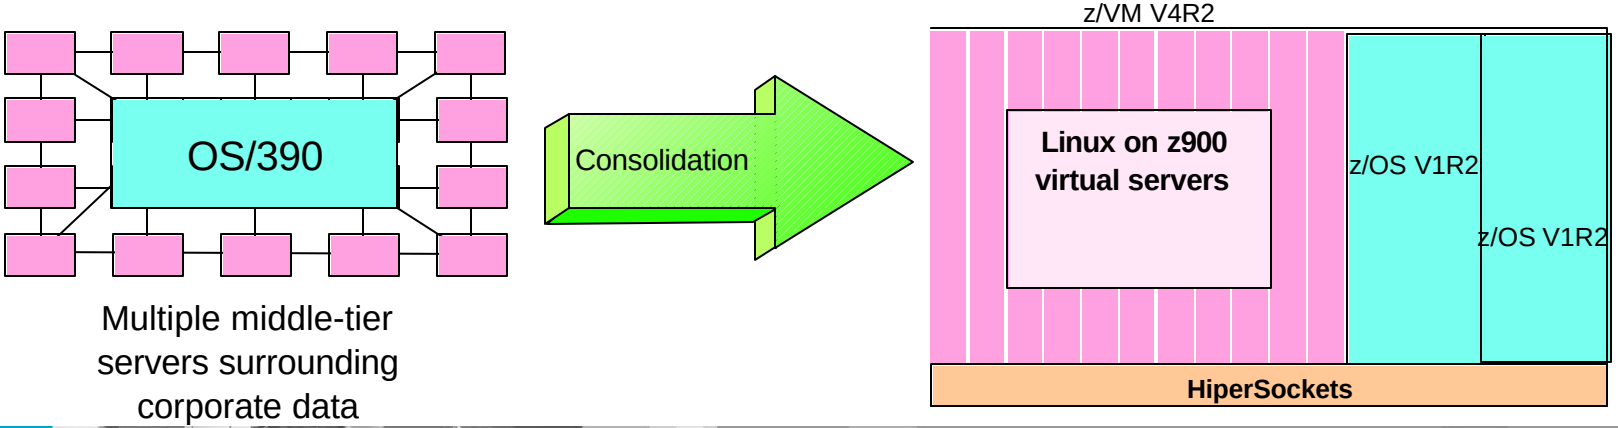

Multiple middle-tier servers surrounding corporate data

Example 3: Performance Evaluation of IPSec/SSL

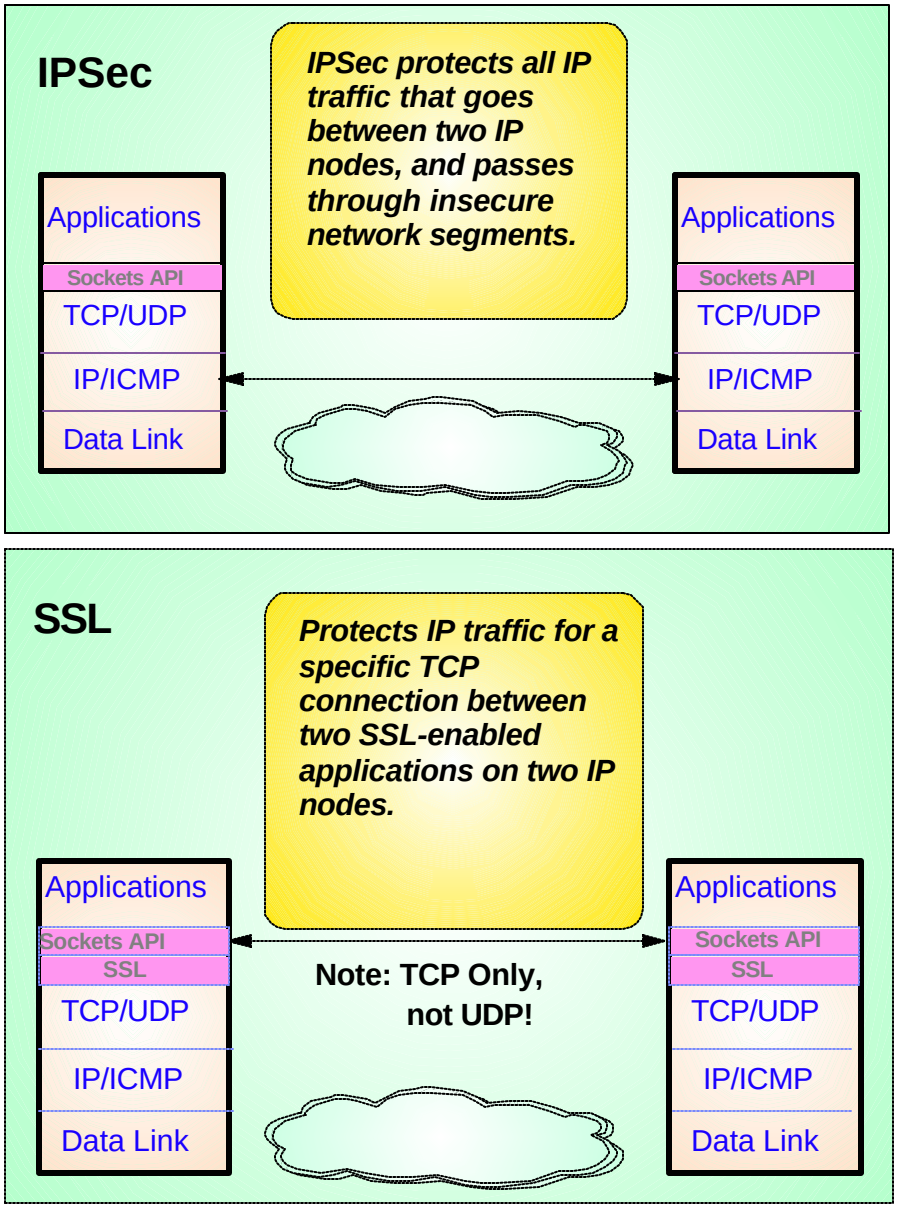

- Model and measure the effect of IPSec/SSL prior to production deployment
  - Generate predictable workload
    - Through a VPN (Tunnel or Transport mode)
      - Client/Server or Application Client mode
    - Using SSL
      - Client/Server mode
      - GSKIT and OPENSSL support (Linux)
      - System 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

Example 4: IPv4 to IPv6 Internet Evolution

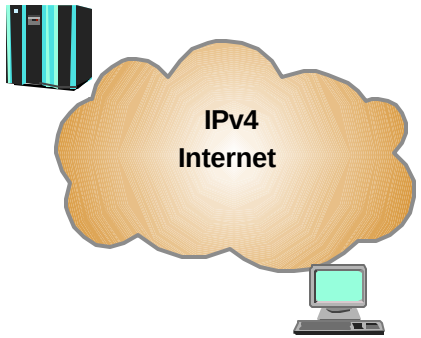

Today

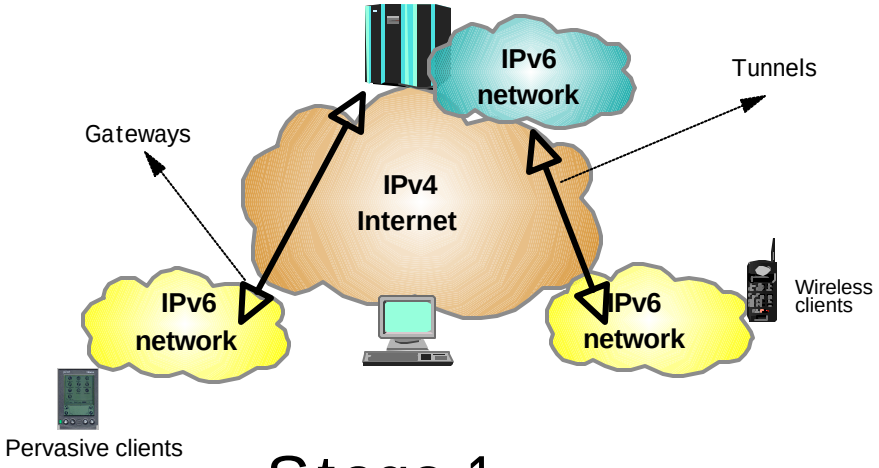

Stage 1

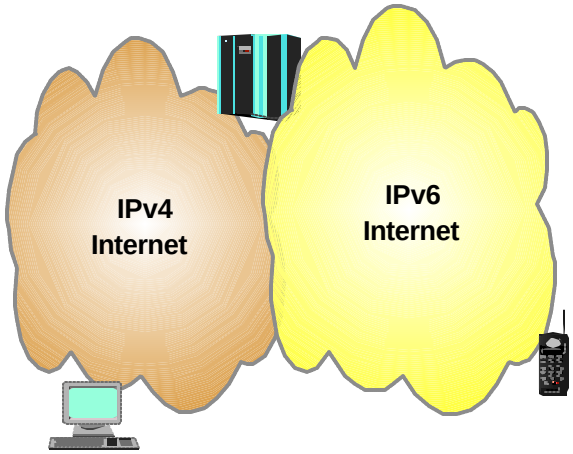

Stage 2

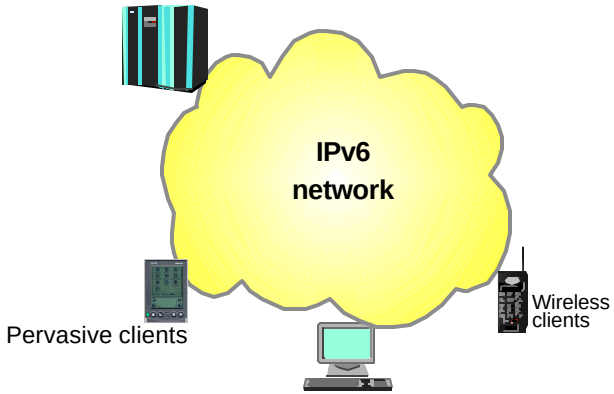

Stage 3

IPv6 - Modeling the Performance Impact

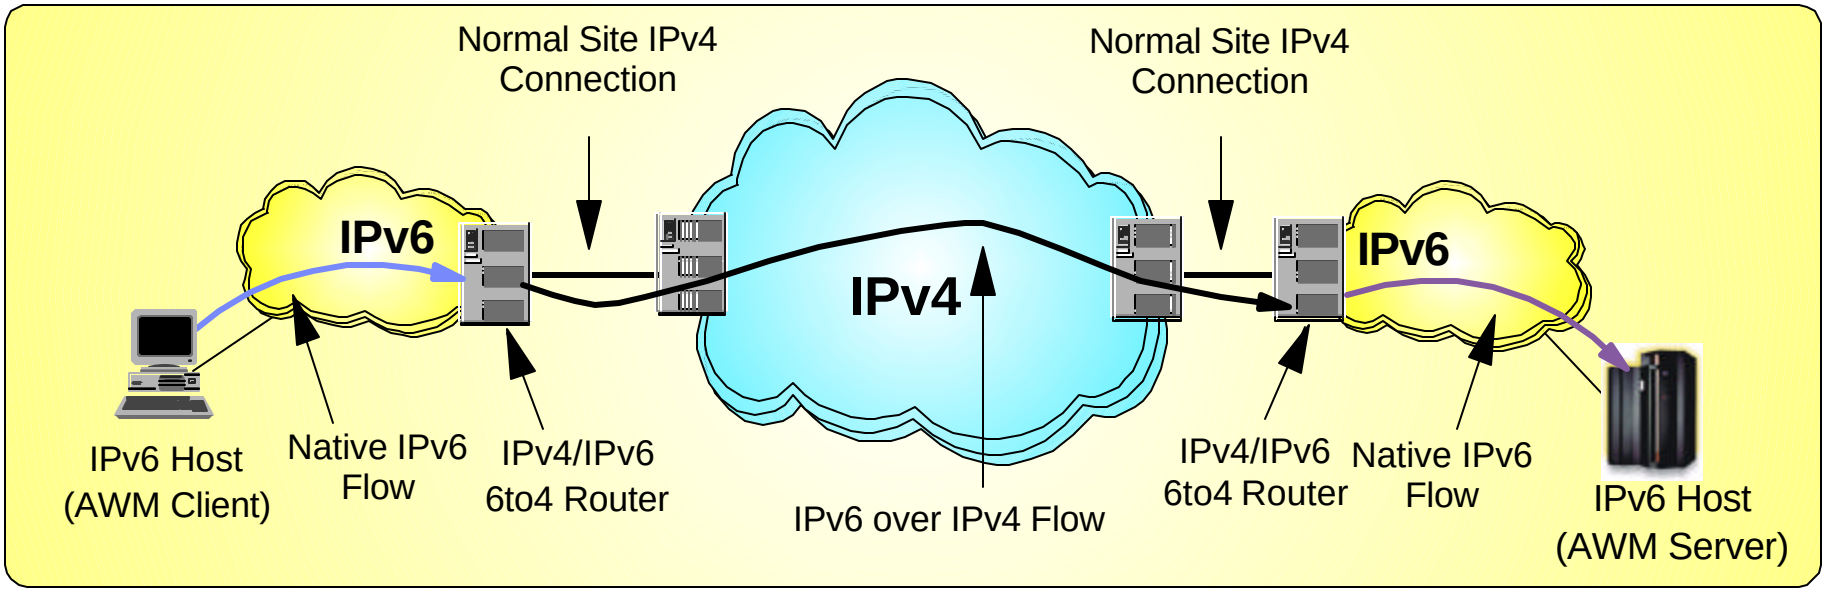

- Tunnel: encapsulate an IPv6 packet in an IPv4 and send to tunnel IPv4 endpoint address
  - Tunnel endpoint placement depends on connectivity needs
    - Placing endpoints in routers allows entire sites to be connected over an IPv4 network
    - Placing endpoints in hosts allows access to remote IPv6 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 IPv6/IPv6 tunnels
  - Effect on IPv6 networks and existing IPv4 applications

Application Client Mode

Details and Examples

Example 1: Evaluating Load Balancing Solutions

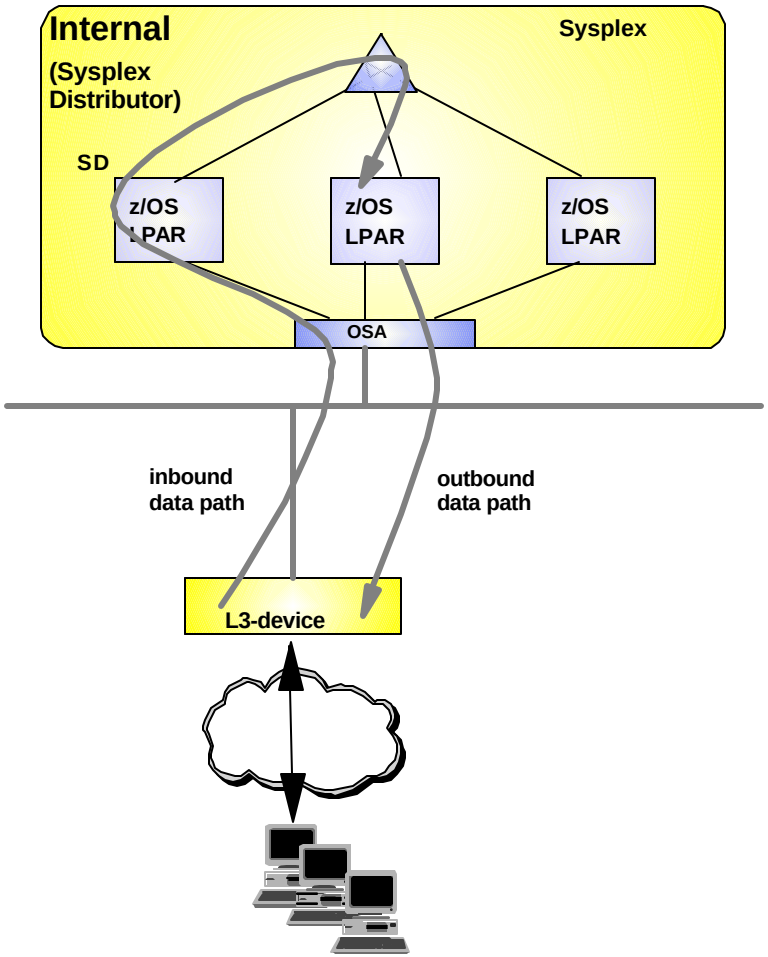

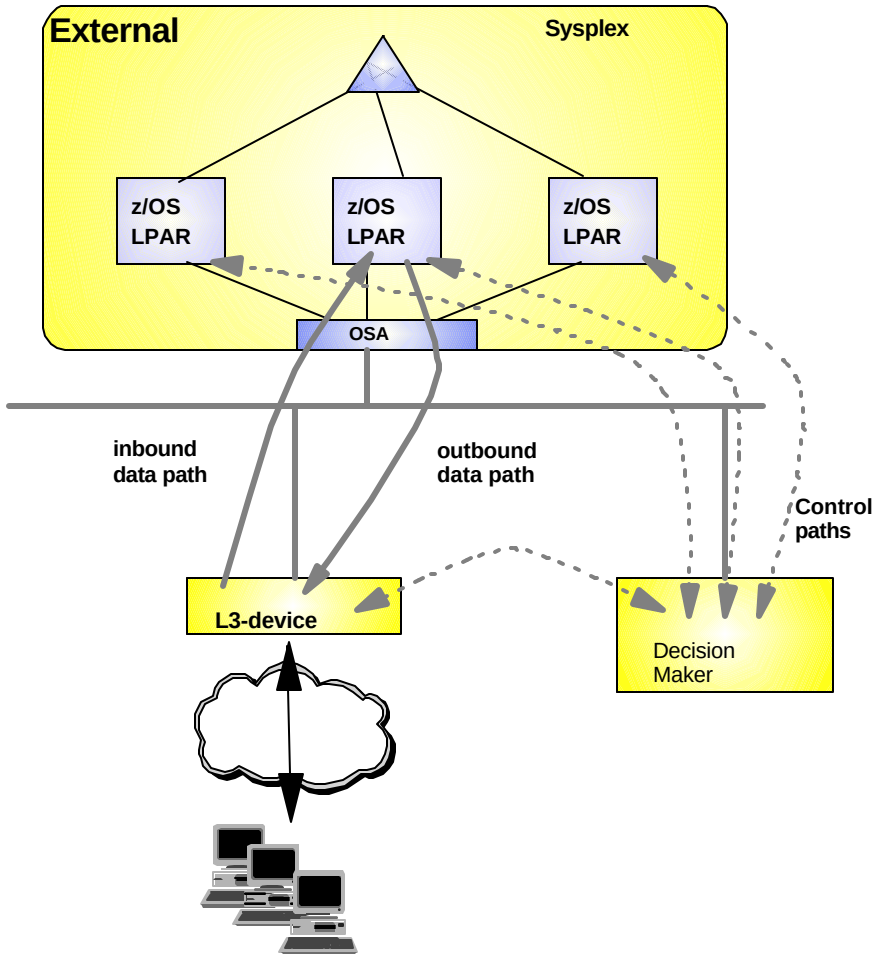

- Model and measure the effect of a load balancing solution prior to deployment
  - Capacity planning - How many target servers are needed?
  - Network impact
  - Effect on end-user response time

Example 2: TN3270(E) Server Placement

- Which is better: Inboard or Outboard TN3270(E) Server?

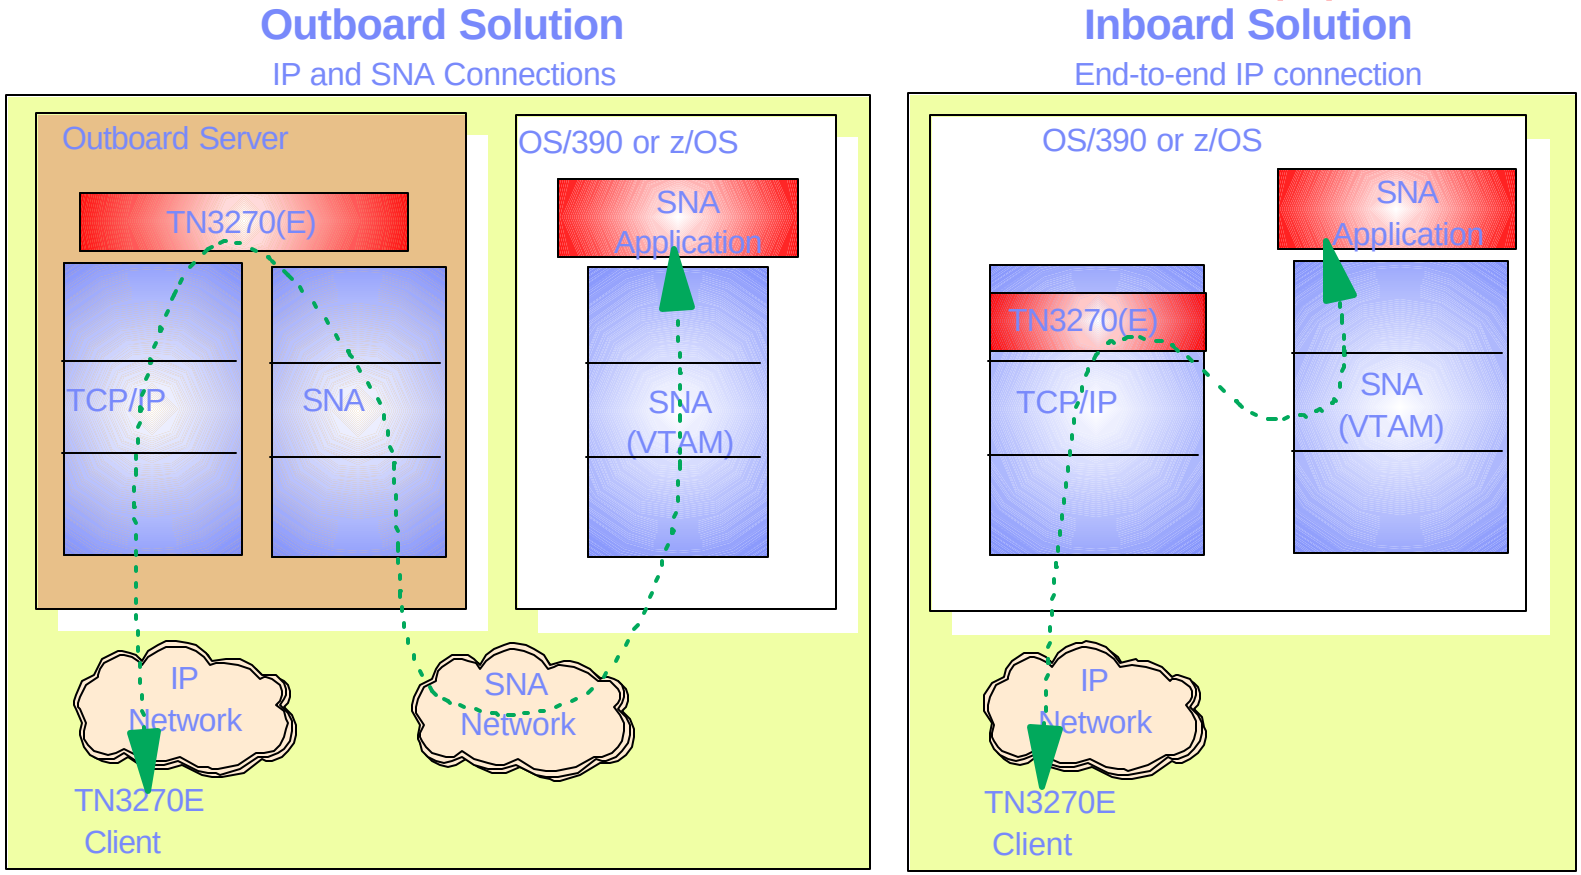

- Use Application Client mode to simulate a large number of TN3270(E) clients/activity
- Run against inboard and outboard server
- Determine cost/performance/scalability characteristics of each solution

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**

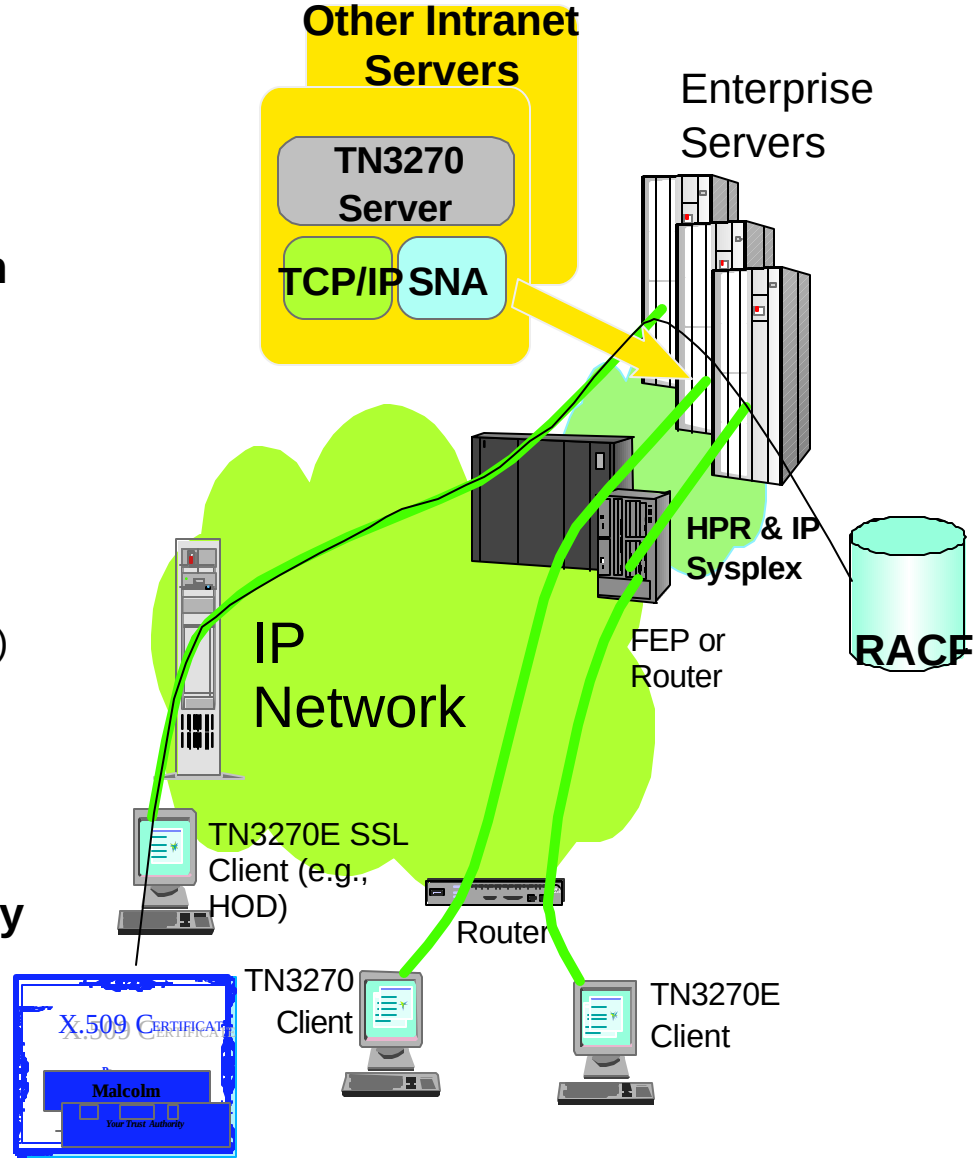

Example 3: SAP R/3 Consolidation on zSeries

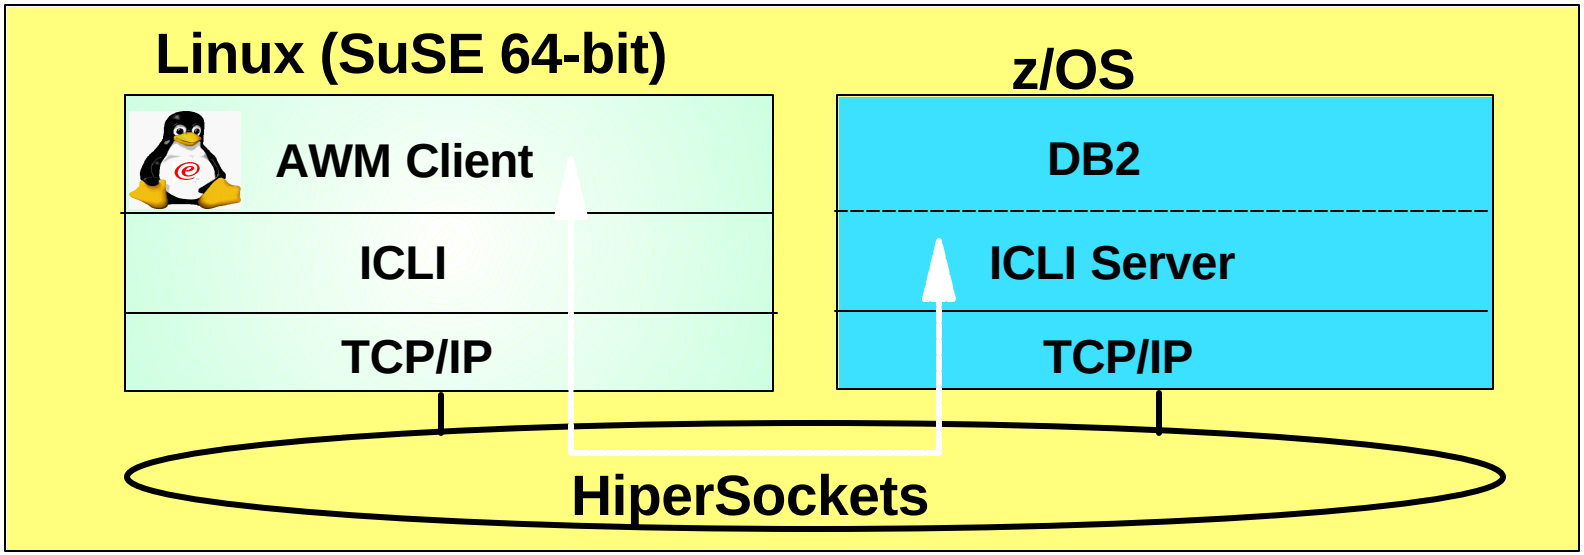

- SAP R/3 communications modeling
  - Model application server to database server communications
    - Using SAP R/3 Integrated Call Level Interface
    - Drives the entire communications path
  - Allows simulation of multiple application servers
    - Each servicing hundreds/thousands of client requests
  - Useful in evaluating SAP R/3 consolidation on zSeries
    - Detailed response time metrics, throughput rates

Example 4: 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

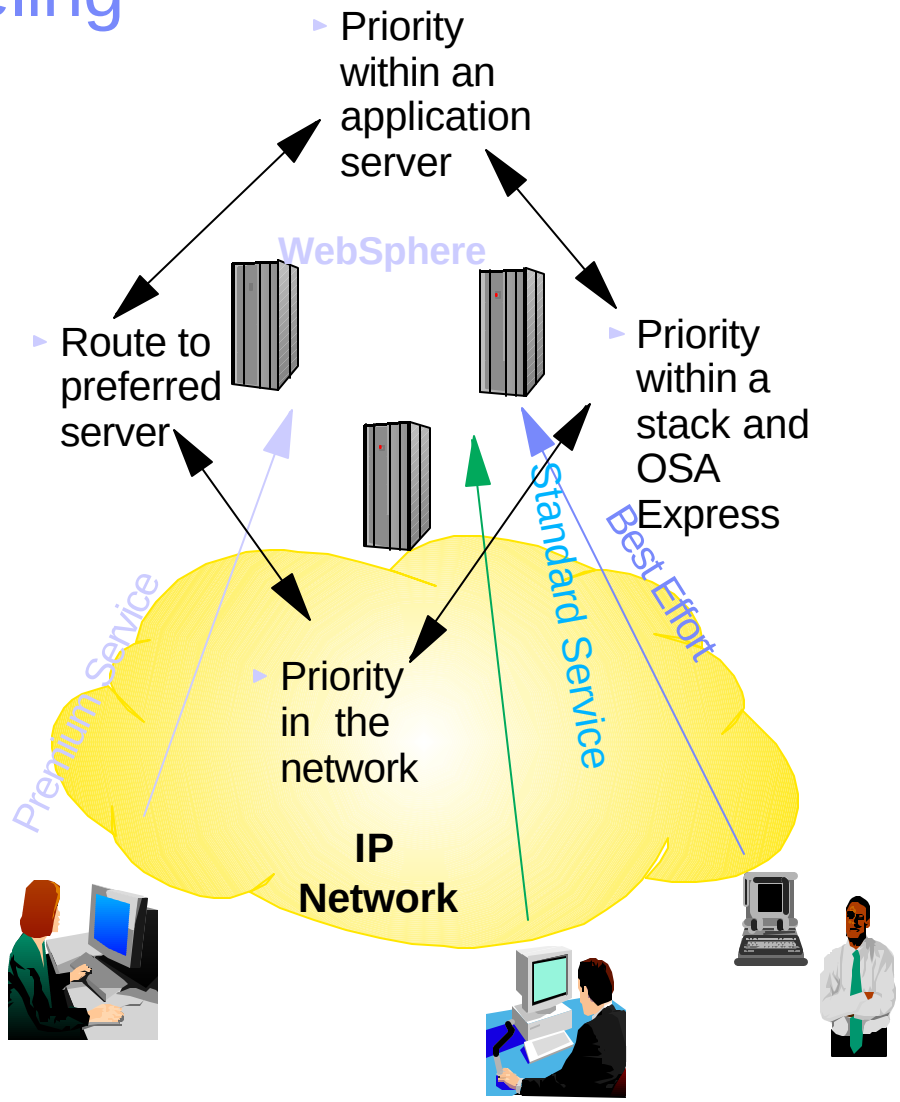

WebSphere Measurement without QoS

Transient Behavior of User Response Time
(WebSphere)

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

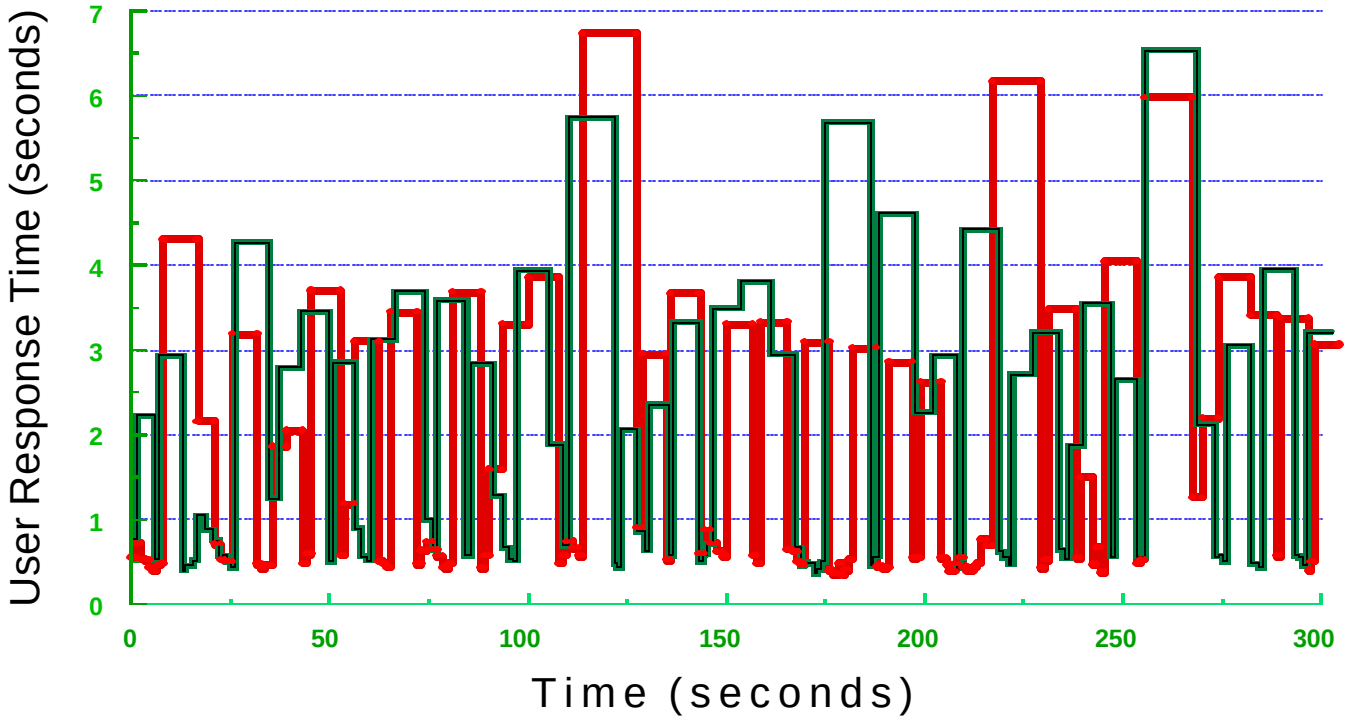

Higher Priority Access (ToS = 5); **Average Response Time = 1269 msec (341-6739 msec)**
Lower Priority Access (ToS = 0); **Average Response Time = 1560 msec (367-6537 msec)**

WebSphere Measurement with QoS

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

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

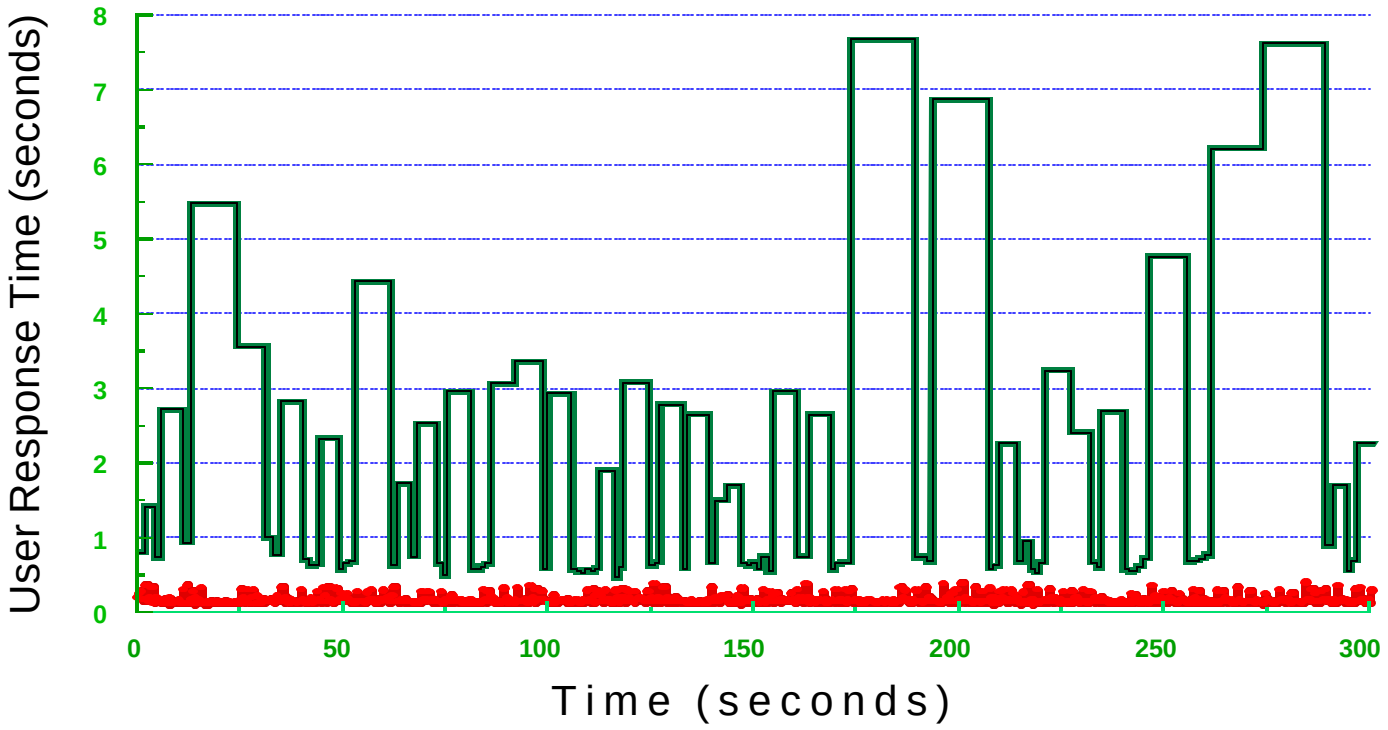

Higher Priority Access (ToS = 5); **Average Response Time = 157 msec (98-393 msec)**
Lower Priority Access (ToS = 0); **Average Response Time = 1506 msec (465-7669 msec)**

Additional Application Client Mode Functions

- CICS Sockets
  - Simulate client traffic to TCP/IP CICS sockets app
- 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

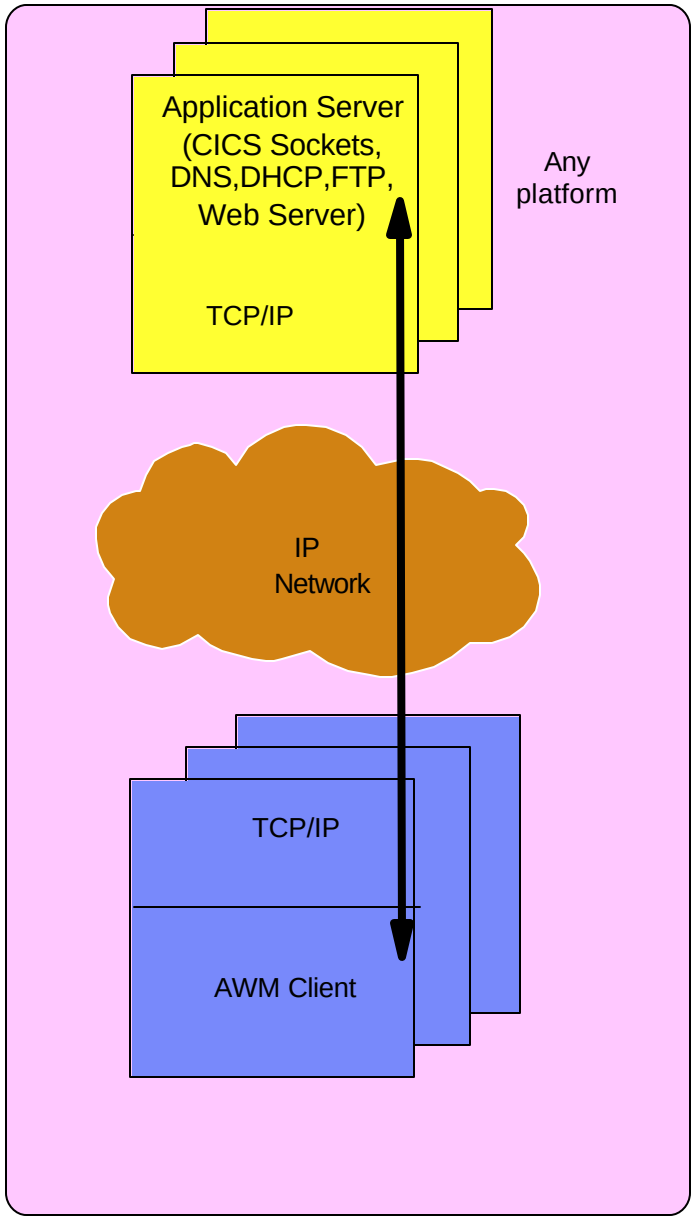

Client/Server Mode Run Example

Configuration, Run, Reports and Graphs

Client/Server Mode Benchmark Run Configuration

Client Host file

Client Suite file

Client Workload file

Client Input file

Server Host file

Server Suite file

Server Workload file

Client/Server Mode Benchmark Run Status

```
awm>show status-run
AWM8888I: Summary Status Information for CLIENT host shah1
AWM8888I: Number of Clients Started: 10  Active: 10  Failed: 0
AWM8888I: State update at: 0.034522  State: AWM Main task clients created
AWM8888I: Last statistics update time: 8.740137
AWM8888I: TPS: 10963.371173  TRT: 0.000912
AWM8888I: TPUT(in): 10963371.173194  TPUT(out): 2192674.234639
AWM8888I: Transactions completed: 80000.000000
AWM8888I: Bytes(in): 80000000.000000  Bytes(out): 16000000.000000
AWM8888I: - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
AWM8888I: Summary Status Information for SERVER host shah3
AWM8888I: Number of Servers Started: 10  Active: 10  Failed: 0
AWM8888I: State update at: 0.017694  State: AWM Main task servers created
AWM8888I: Last statstics update time: 23.749921
AWM8888I: TPS: 10971.566243  TRT: 0.000911
AWM8888I: TPUT(in): 2194313.248565  TPUT(out): 10971566.242825
AWM8888I: Transactions completed: 80000.000000
AWM8888I: Bytes(in): 16000000.000000  Bytes(out): 80000000.000000
AWM8888I: - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
```

Client/Server Mode Benchmark Graphs

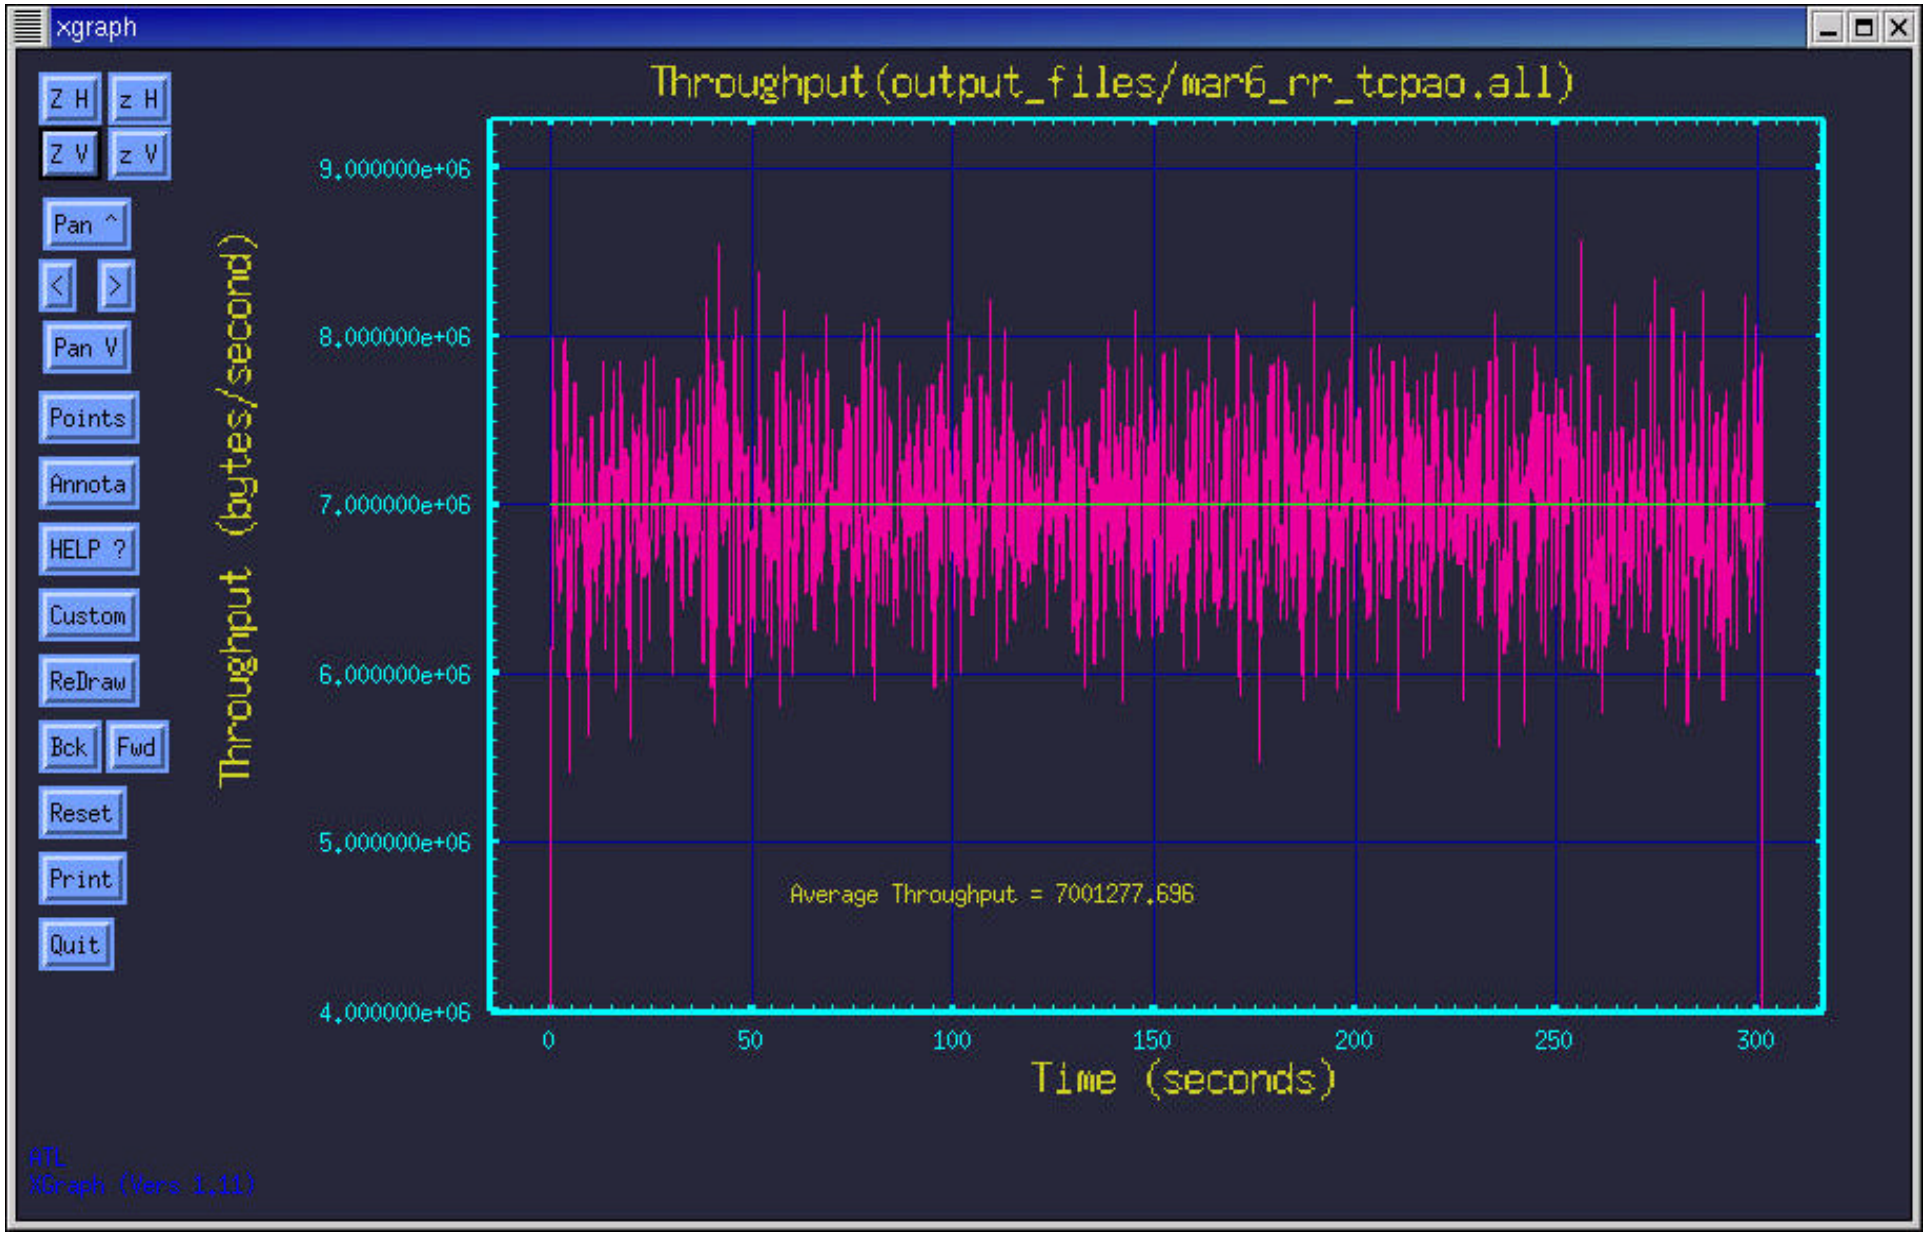

Client/Server Mode Benchmark Graphs

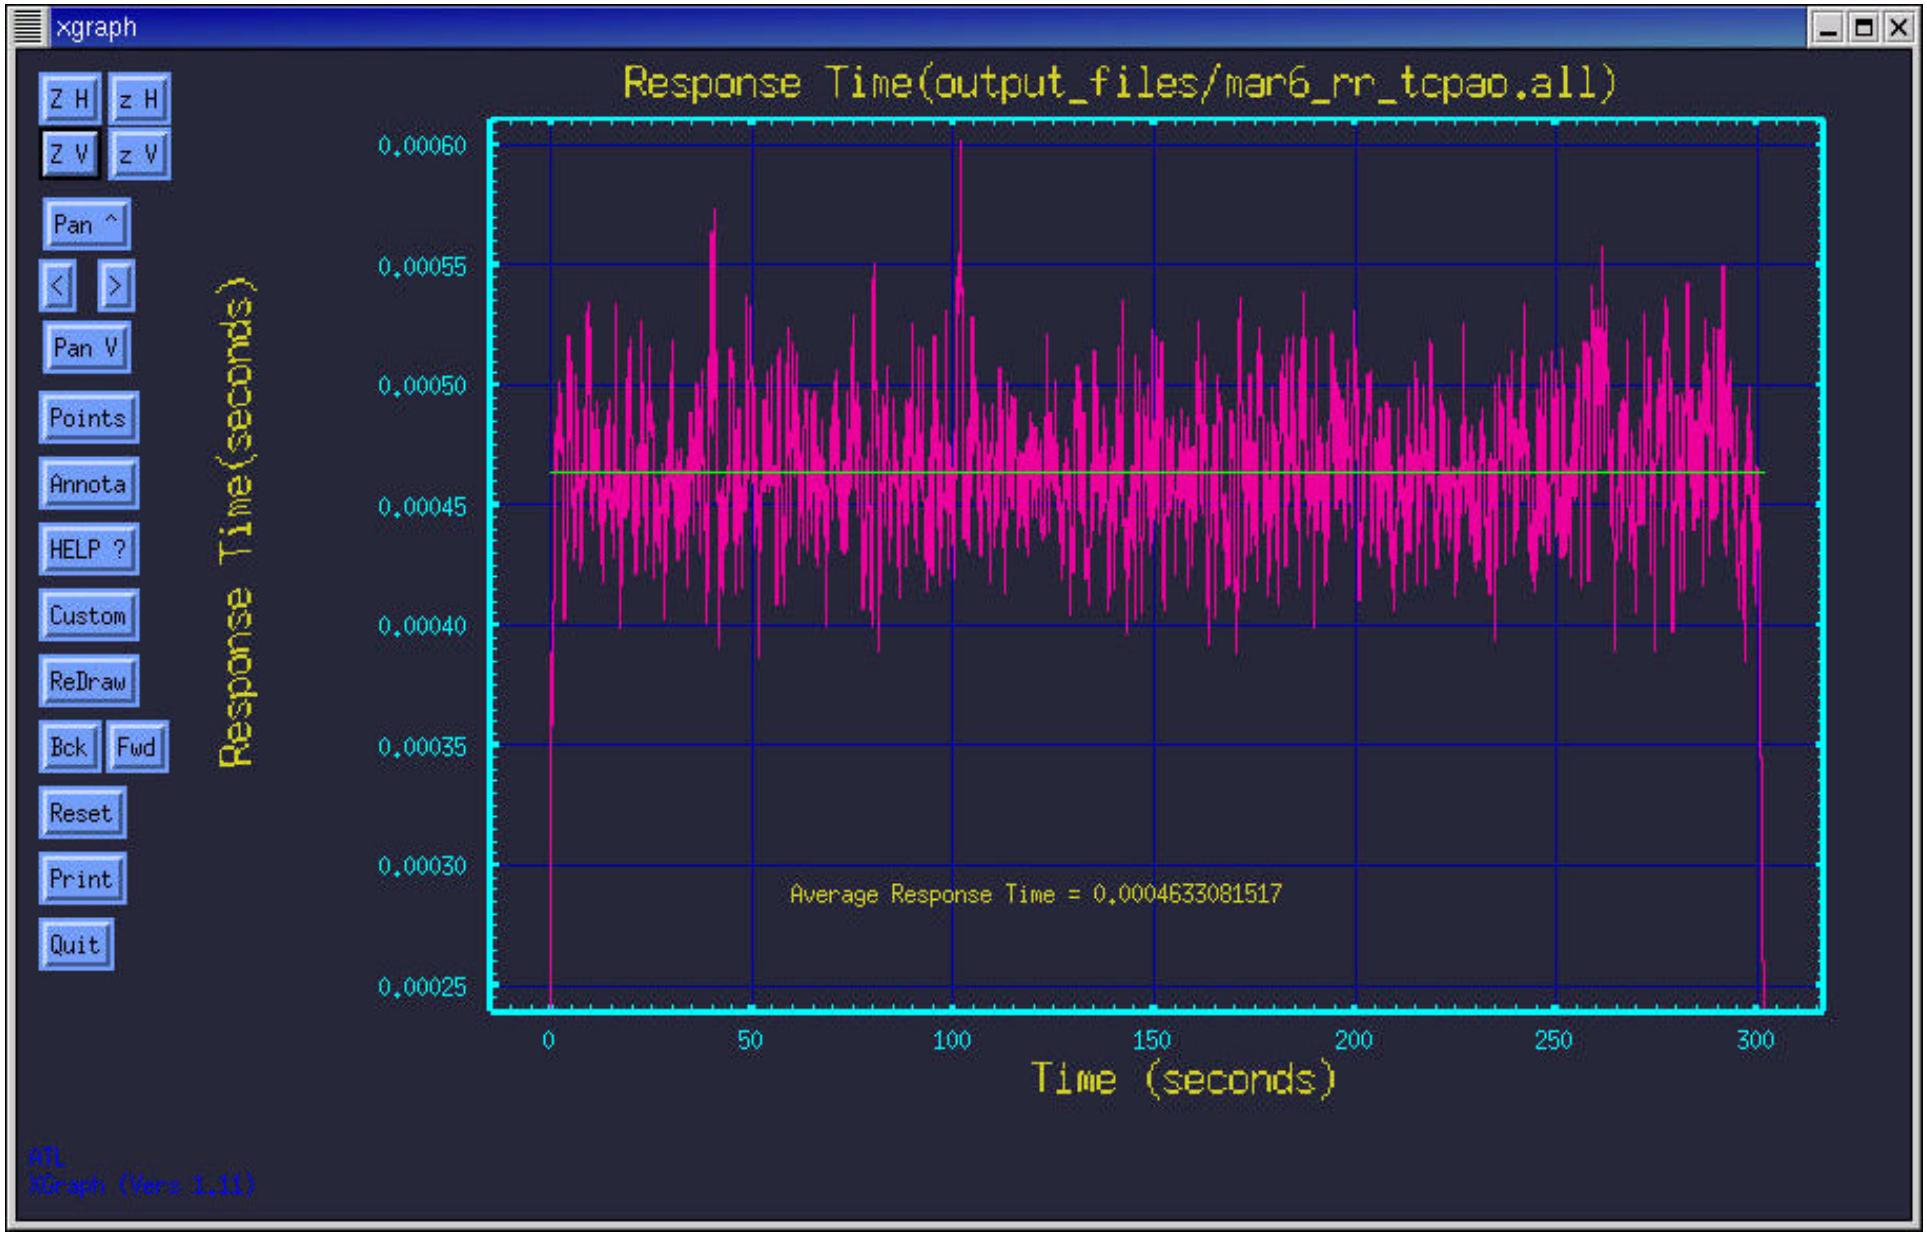

Client/Server Mode Benchmark Graphs

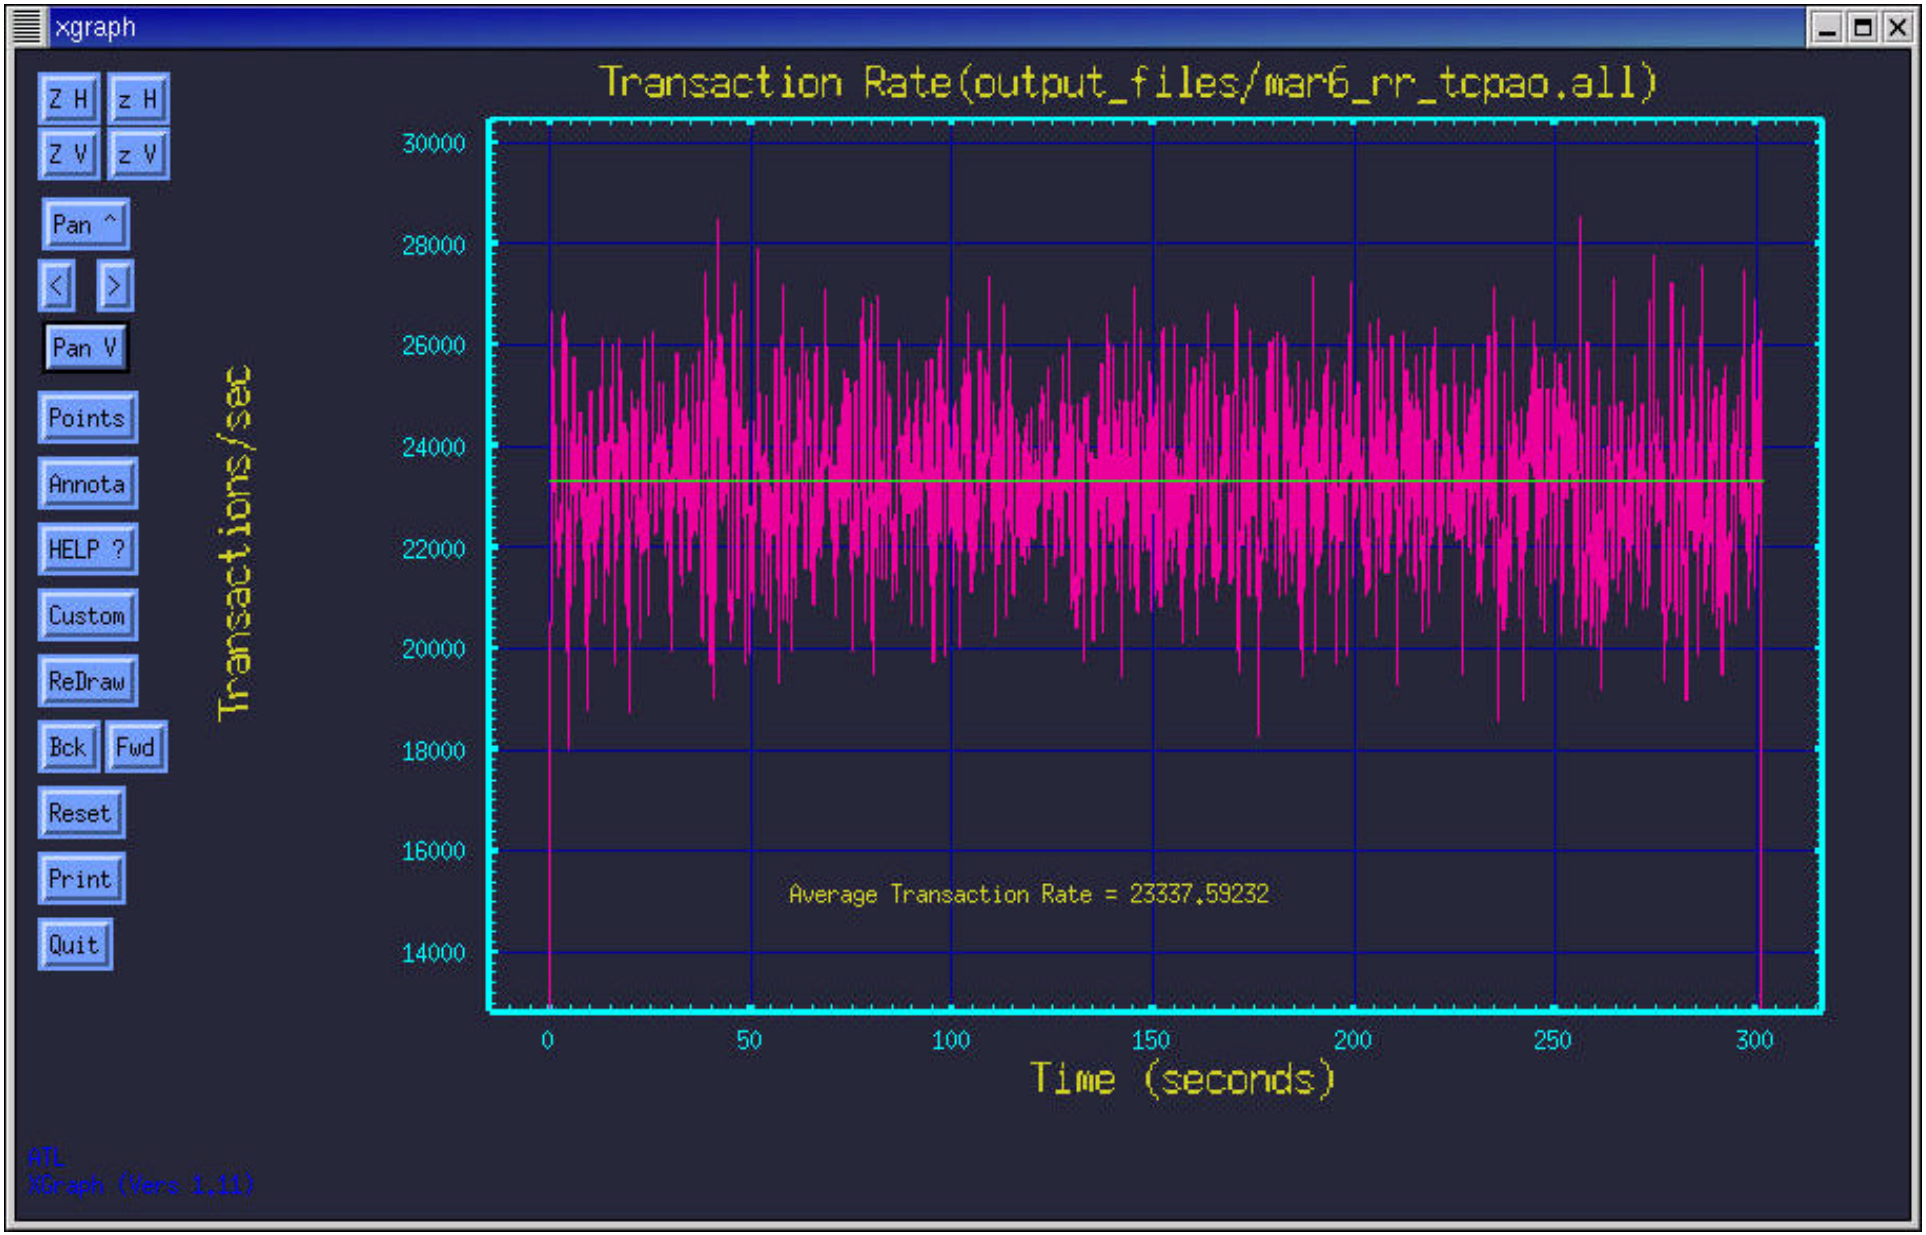

Client/Server Mode Benchmark Reports

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

Summary

- Performance and capacity planning tool for your network and networked applications
  - Generates real network traffic
  - Verify network and server can handle additional application or increased workload from existing application before deployment
- Available as two products
  - Application Workload Modeler for z/OS R1
    - Includes z/OS, Linux on zSeries, and Linux/Intel versions of the product
    - Worldwide GA December 20, 2002
    - IBM product available as 5655-J62
  - Application Workload Modeler for Linux on zSeries
    - Includes Linux on zSeries and Linux/Intel versions of the product
    - Worldwide GA January 17, 2003
    - IBM product available via Passport Advantage

For More Information

- 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
- Send an e-mail with any questions to ***awm@us.ibm.com***