z/OS V1R4 Communication Server Performance Summary

Dan Patel
(danpatel@us.ibm.com
Kevin Stalnaker
(kevins@us.ibm.com)

October 15, 2002

Trademarks

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

BatchPipes*
CICS*
DB2*
DB2 Connect
DB2 Universal Database
e-business logo
ESCON*
FICON
HiperSockets
IBM*
IBM logo*
Multiprise
NetView*
OS/390*
Parallel Sysplex*
PR/SM
QMF
RMF
S/390*
S/390 Parallel Enterprise Server
Tivoli
VTAM
xSeries
WebSphere
z/Architecture
z/OS
z/VM
zSeries

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Cisco is a registered trademark of Cisco Systems, Inc.
Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

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

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. 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.

Performance Disclaimer

- The performance data discussed in this presentation was collected in dedicated system environments. Therefore, the results obtained in other configurations or operating system environments may vary.

e-business

Presentation Overview

- z/OS V1R4 CS Performance Summary
  - Performance Highlights
    - Comm Server Performance Summary
    - Netmarks Primitive Workload
    - Hardware/Software configurations
    - Release to Release Comparison (z/OS V1R4 CS vs. z/OS V1R2 CS)
      - Netmarks Primitive Benchmarks (TCP): RR,CRR,STR (Throughput and CPU cost per Transaction)
      - Netmarks Primitive Benchmarks (UDP): RR (Throughput and CPU cost per Transaction)
      - DNS Bind v9.2 vs. Bind v9.1 - Throughput and CPU Cost per transaction
      - FTP Server Throughput Comparison
      - FTP Server CPU Cost per Transaction
      - FTP Server Throughput and CPU Cost Comparison Relative to z/OS V1R2 CS
      - TN3270 Server Throughput Comparison
      - TN3270 Server CPU Cost Comparison
      - TN3270 Server Throughput and CPU Cost Comparison Relative to z/OS V1R2 CS
    - Telnet (TN3270) Storage Utilization (z/OS V1R4 CS)
    - Telnet TN3270 Capacity Planning (z/OS V1R4 CS)
    - z/OS V1R4 CS Performance comparisons
      - Effect on performance of enabling IPv6 - Comparisons of Throughput and CPU cost per transaction
      - Effect of enabling IPv6 - FTP Server
    - Summary

z/OS V1R4 Comm Server Performance Summary

z/OS V1R4 CS Performance Highlights

- *z/OS V1R4 CS: First release supporting IPv6*
- *Design goal was to incorporate IPv6 protocol into the communication stack with minimal impact to IPv4 performance*
- *Single TCP/IP stack supporting IPv4 and IPv6 protocols at the same time*
- *Can be configured as v4 or v4/v6 stack*
- *To configure IPv6 support: Add AF_INET6 network statement to the BPXPRMxx member of System Parameters*
- *DNS Bind 9.2*
  - ***Significant performance (CPU and Throughput) and scalability improvements for DNS bind 9.2 (z/OS V1R4 CS) compared to DNS bind 9.1 (z/OS V1R2 CS)***

- **Performance Comparisons:**
  - *Release to Release performance comparisons (V1R4 vs. V1R2)*
    - ***Netmarks Primitive Benchmarks (Throughput and CPU Cost)***
    - ***Applications: FTP Server and TN3270 (Throughput and CPU cost comparisons)***
  - *z/OS V1R4 CS Performance Comparisons:*
    - ***IPv6 Enabled vs. Disabled - Throughput and CPU comparisons***

z/OS V1R4 Comm Server Performance Summary

Hardware/Software Configurations

- **Netmarks Primitives (RR, CRR, Streams) and TN3270**

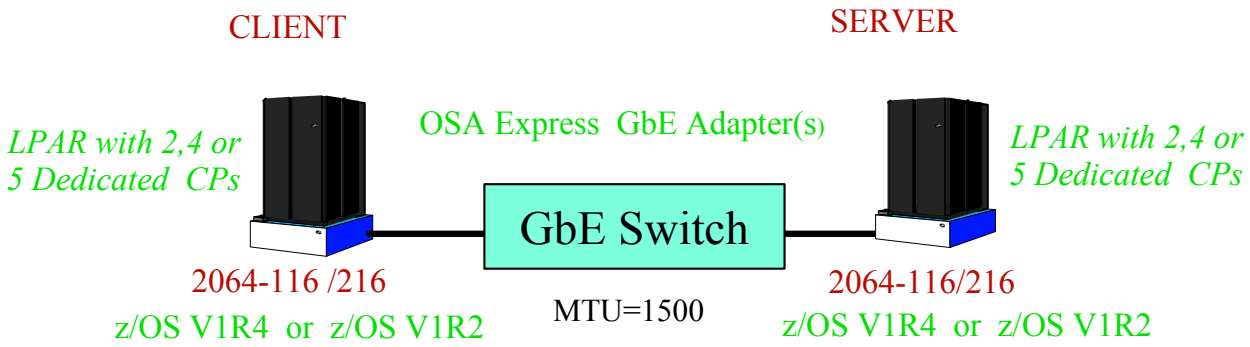

- **FTP Server, DNS :**

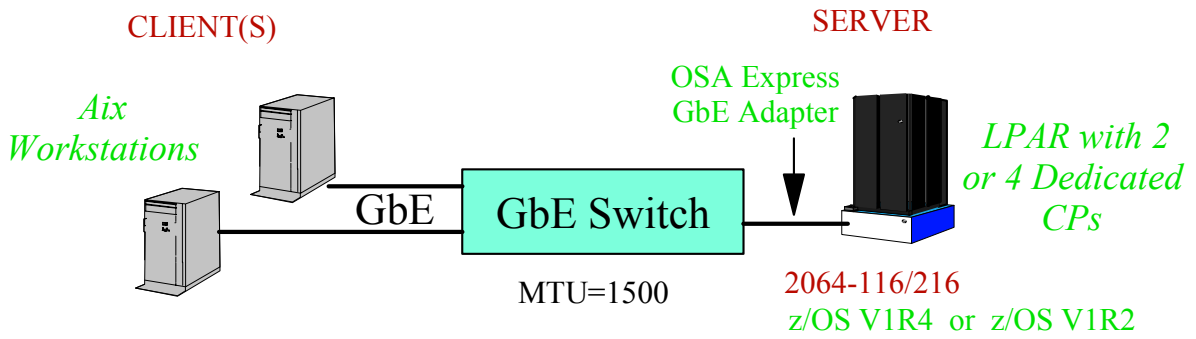

z/OS V1R4 Comm Server Performance Summary

PRIMITIVE WORKLOADS FOR TCP/IP

RR Workload

- Simulate TN3270
- Interactive workloads

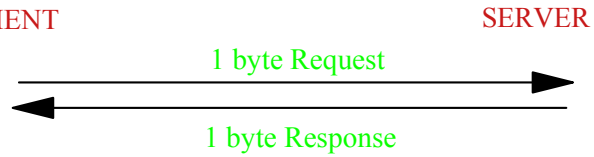

- 60 connections

Connect-Request-Response (CRR) Workload

- Static Web Serving

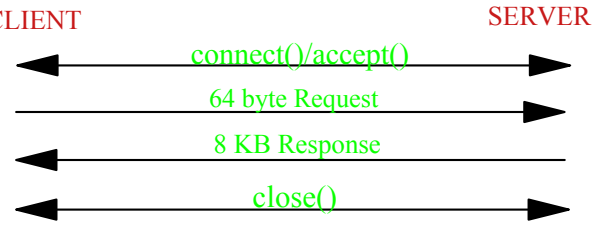

- 9 Connections

Streams Workload

- Simulate FTP
- Bulk data transfers applications such as ADSM, DB/2
- Memory to Memory (no DASD)
- STRG used for outbound
- STRP used for Inbound

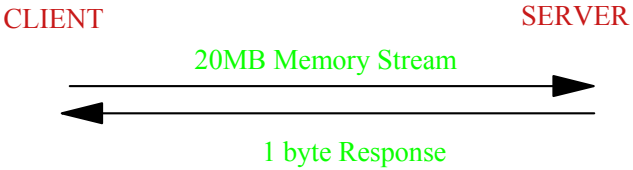

- 10 Connections

z/OS V1R4/V1R2 Comm Server Performance Scenarios

Three Scenarios considered for comparison

- **1- Do not specify AF_INET6 in the BPXPRMxx members (both server and client). Configure the netmarks client to use a v4 address, the Netmarks server to use a v4 address, and the Netmarks server to use an AF_INET socket (IPv6: N selected by Netmarks default). This is the pure v4-v4 scenario and the V1R4 stack is expected to behave similar to V1R2**
  - ✓ **Compare V1R2 (v4-v4) to V1R4 (v4-v4, no IPv6 configured)**

- **2- Specify AF_INET6 in the BPXPRMxx member (only in the server, client BPXPRMxx will use AF_INET and not AF_INET6). Configure the Netmarks client to use a v4 address, the Netmarks server to use a v4 address, and the Netmarks server to use an AF_INET6 socket (IPv6: Y specified in Netmarks server host file). This is an AF_INET6 socket communicating to an AF_INET socket, but the TCP/IP stack where the server resides will convert these IPv4 addresses into mapped addresses before delivering the packets to the server. This scenario will show customers the impact that will occur if they enable IPv6 on their stack and modify their application to be an IPv6 application (eg. FTP server does this in V1R4).**
  - ✓ **Compare V1R4 (v4-v4, IPv6 configured and used) to V1R4 (v4-v4, no IPv6 configured as in case 1)**

z/OS V1R4/V1R2 Comm Server Performance Scenarios ...

Three Scenarios considered for comparison ...

- **3- Specify AF_INET6 in the BPXPRMxx member (only in the server, client BPXPRMxx uses AF_INET and not AF_INET6). Configure the Netmarks client to use a v4 address, the Netmarks server to use a v4 address, and the Netmarks server to use an AF_INET socket (IPv6: N, selected by Netmarks default). This is a v4-v4 scenario, similar to case 1 above. This scenario will show customers the impact that will occur to enable IPv6 on their stack and continue to run unmodified IPv4 applications.**
  - ✓ **Compare V1R4 (v4-v4, no IPv6 configured as in case 1) to V1R4 (v4-v4, IPv6 enabled but not used)**

Release to Release Comparison

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

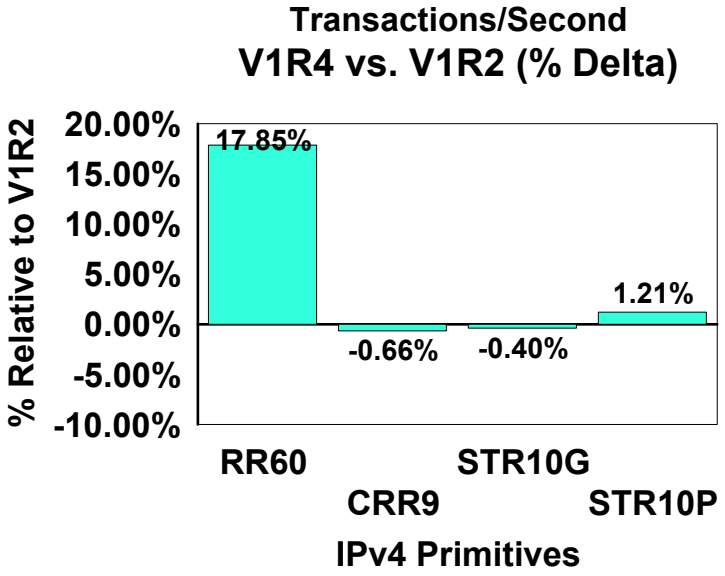

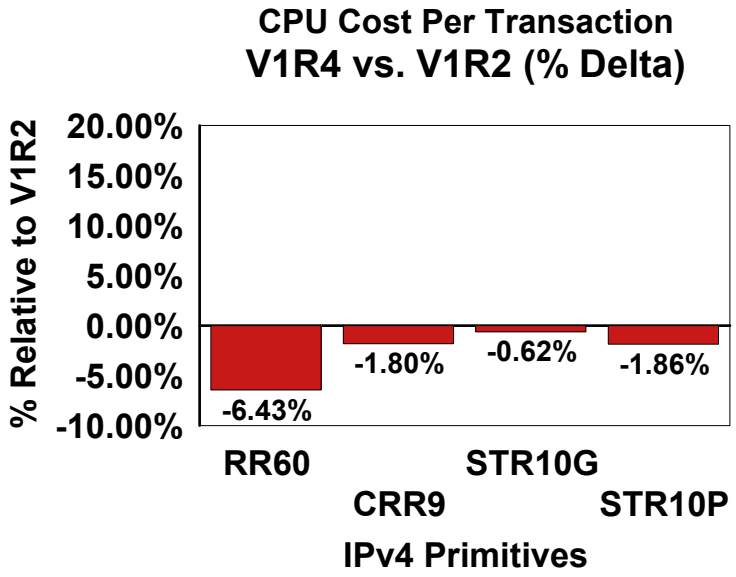

- V1R4 provides Transaction rates within -0.66% to 17.85 % of V1R2
- V1R4 provides lower CPU cost per transaction than V1R2 and the percentages are within (-0.62 to -6.43)

Release to Release Comparison

- Netmarks Primitive Benchmarks (UDP)
  - z/OS V1R4 CS relative to V1R2
    - RR Workload (60 Connections)
    - Comparison - Trans/sec, CPU Cost Per Transaction

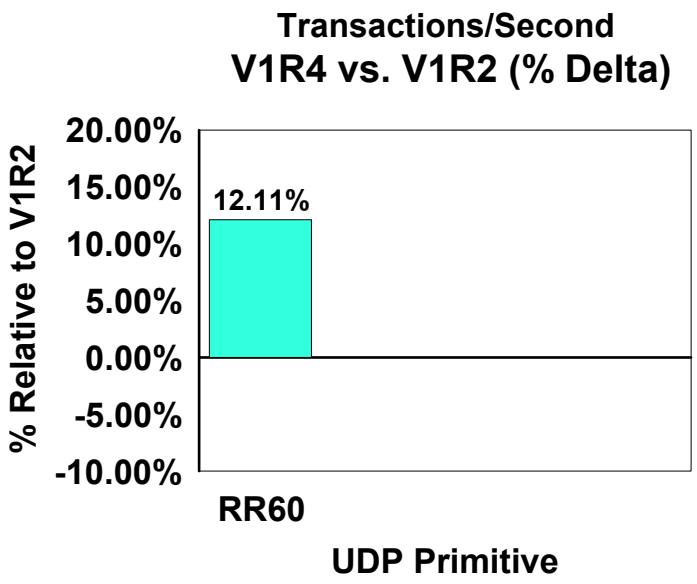

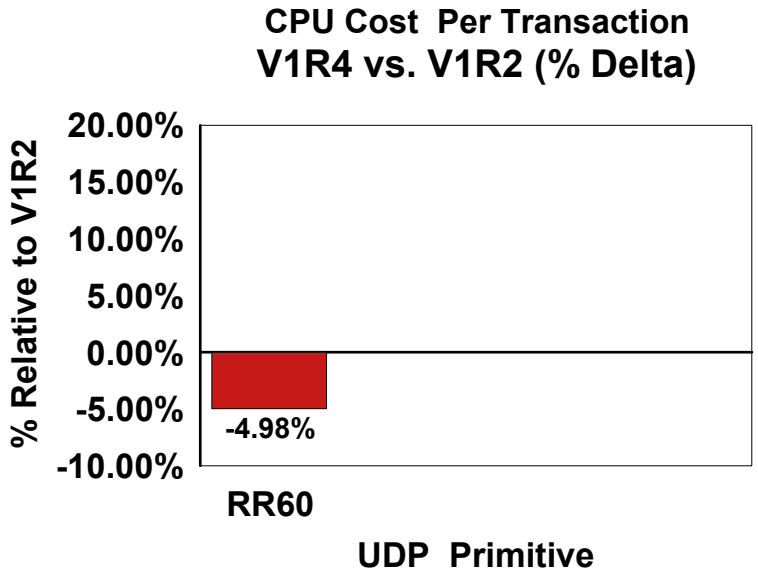

- V1R4 UDP-RR60 provides 12.11% higher transaction rate than V1R2

- V1R4 UDP-RR60 transaction costs is 4.98% less than V1R2

Release to Release Comparison

- Netmarks Primitive Benchmarks
  - Bind V9.2 vs. Bind V9.1 (Domain Name System (DNS) name server)
  - Comparison - Trans/Second, CPU Cost Per Transaction

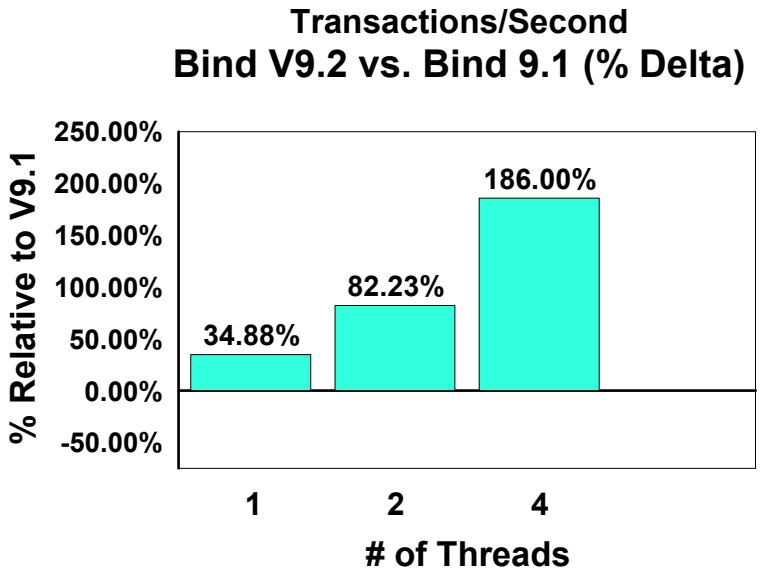

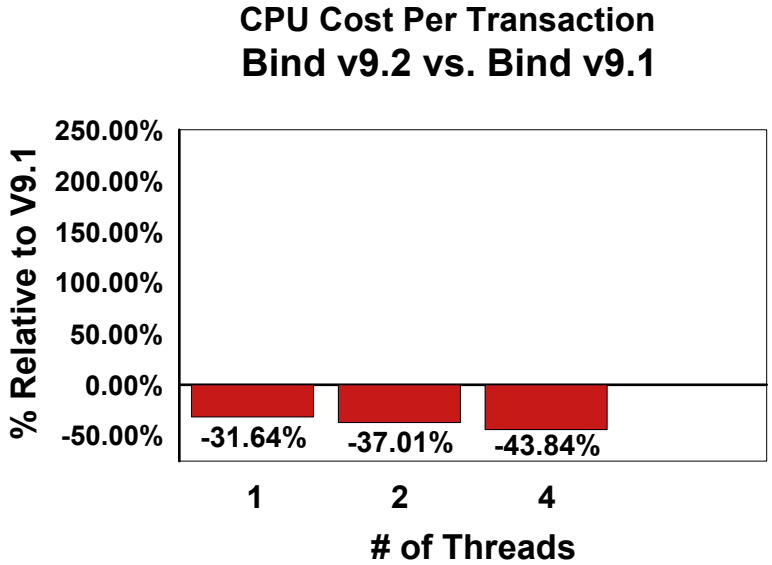

9672-RX6 (4 CPs)---OSAE/GbE--- GbE---Switch----AIX Workstation (Client)
- V1R4: Bind **v9.2** and V1R2: Bind **v 9.1**

- Bind v 9.2 significantly improved transaction rates compared to v9.1 and bind 9.2 scales better than bind 9.1 as the number of threads increases from 1 to 4

- Bind v9.2 significantly reduces CPU cost per transaction

Release to Release Comparison

- **FTP Server : Throughput Comparison (V1R4 vs. V1R2)**

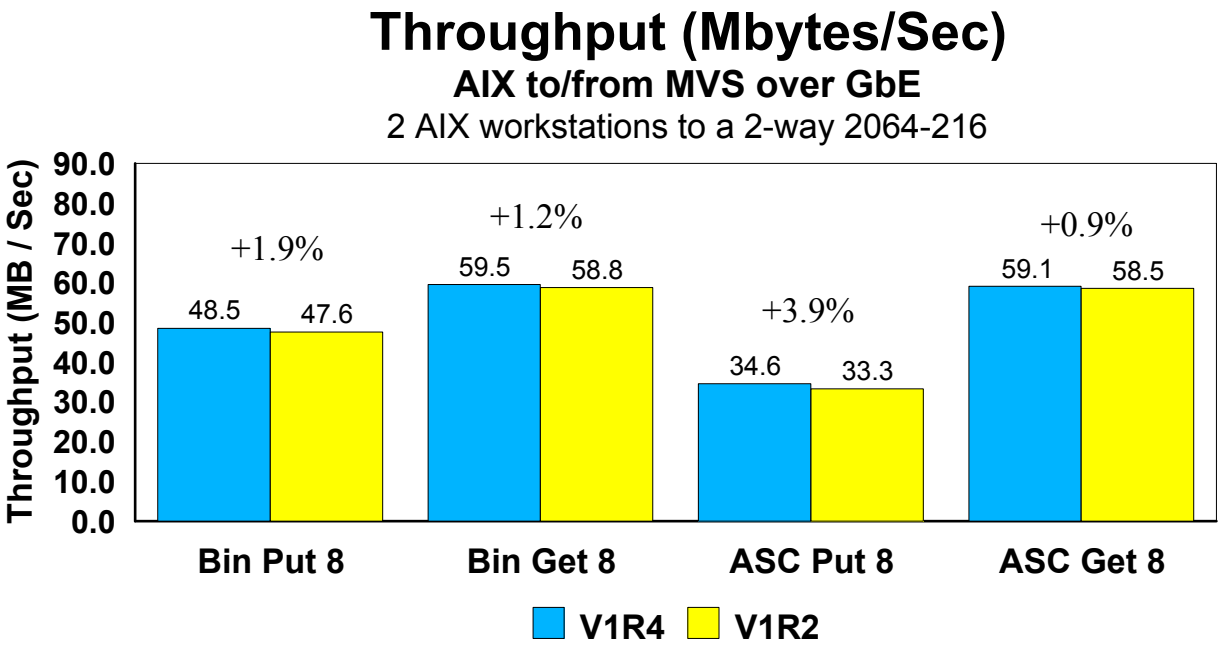

PUT: AIX -----> MVS (8 FTP Sessions, binary/ASCII PUT initiated from AIX clients)
GET: AIX <----- MVS (8 FTP Sessions, binary/ASCII GET initiated from AIX clients)

- Generally, throughput differences between V1R4 and V1R2 are negligible

Release to Release Comparison

- **FTP Server : CPU Cost Comparison (V1R4 vs. V1R2)**

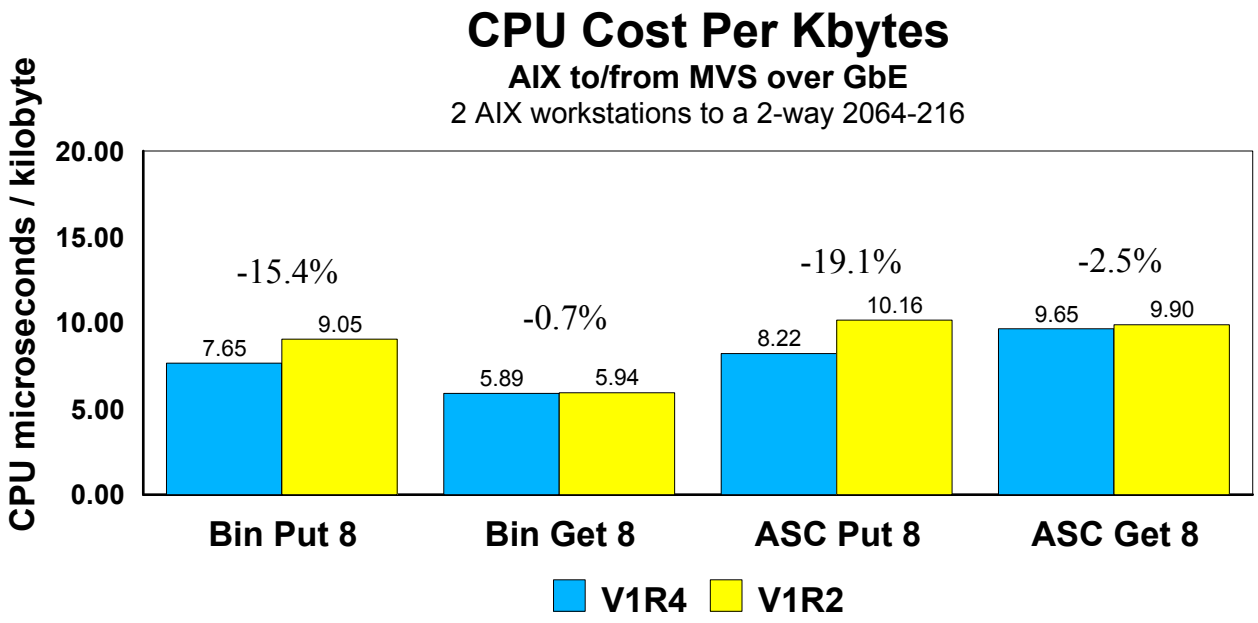

PUT: AIX -----> MVS (8 FTP Sessions, binary/ASCII PUT initiated from AIX clients)
GET: AIX <----- MVS (8 FTP Sessions, binary/ASCII GET initiated from AIX clients)

- For inbound data, V1R4 significantly reduces CPU cost when compared to V1R2

Release to Release Comparison

- **FTP Server: V1R4 relative to V1R2**
  - **Comparison - Throughput, CPU Cost Per Transaction**

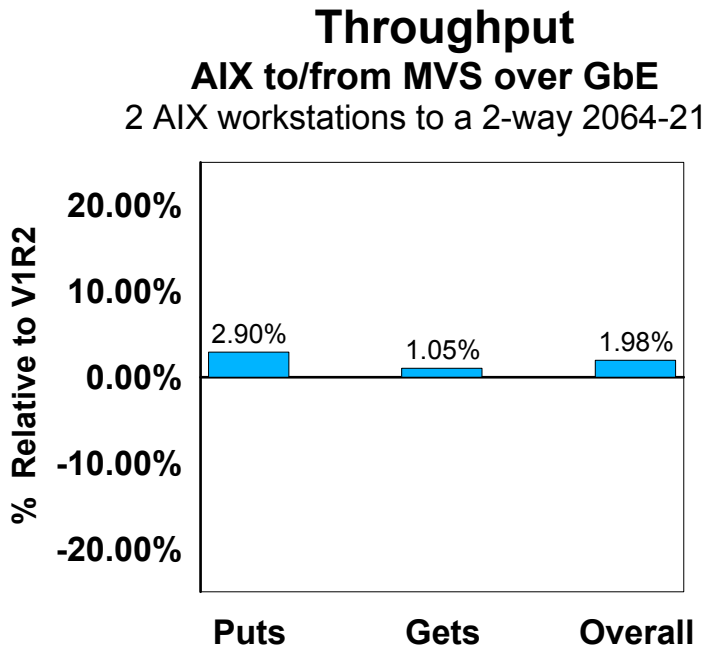

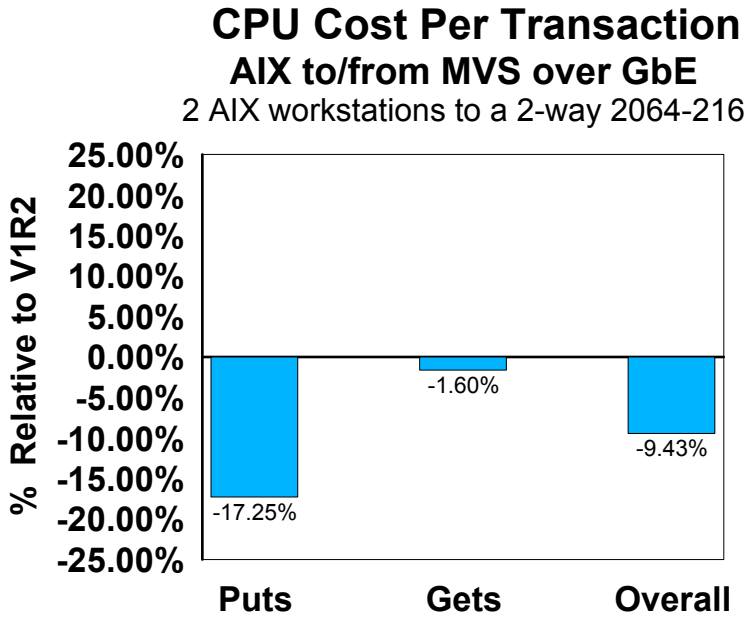

PUT: AIX -----> MVS (8 FTP Sessions, binary/ASCII PUT initiated from AIX clients)
GET: AIX <----- MVS (8 FTP Sessions, binary/ASCII GET initiated from AIX clients)

- V1R4 FTP Server provides better throughput compared to V1R2
- V1R4 FTP Server provides significantly lower CPU cost per transaction compared to V1R2

Release to Release Comparison

- **TN3270 Server: Transaction Rate Comparison (V1R4 vs. V1R2):**

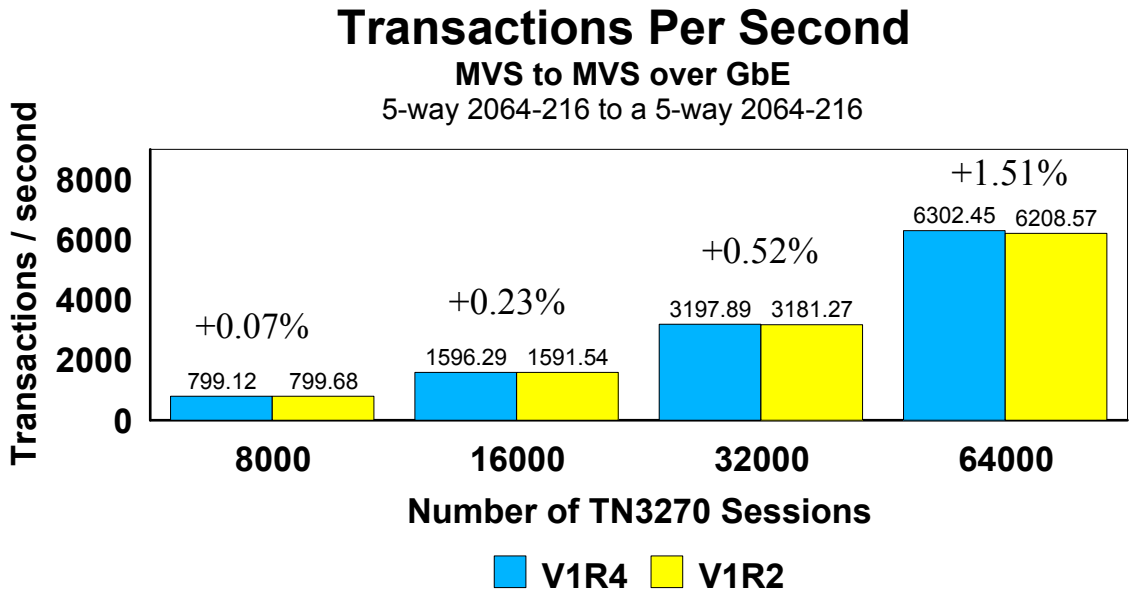

TN3270 Configuration:

Tran: 100 bytes in / 800 bytes out # clients = 8000 to 64000
Config: 2064-216 (5 CPs, Clients)----OSAE/GbE (2)----2064-216 (5 CPs, Server)
Client: 4 TPNS's simulating TN3270 clients
Server: 4 ITPECHO applications, TN3270 Server
Transaction Rate: Six transactions per minute for each user (10 second think time)
Effective transaction rate is shown in the graph

Release to Release Comparison

- **TN3270 Server : CPU Cost Comparison (V1R4 vs. V1R2)**

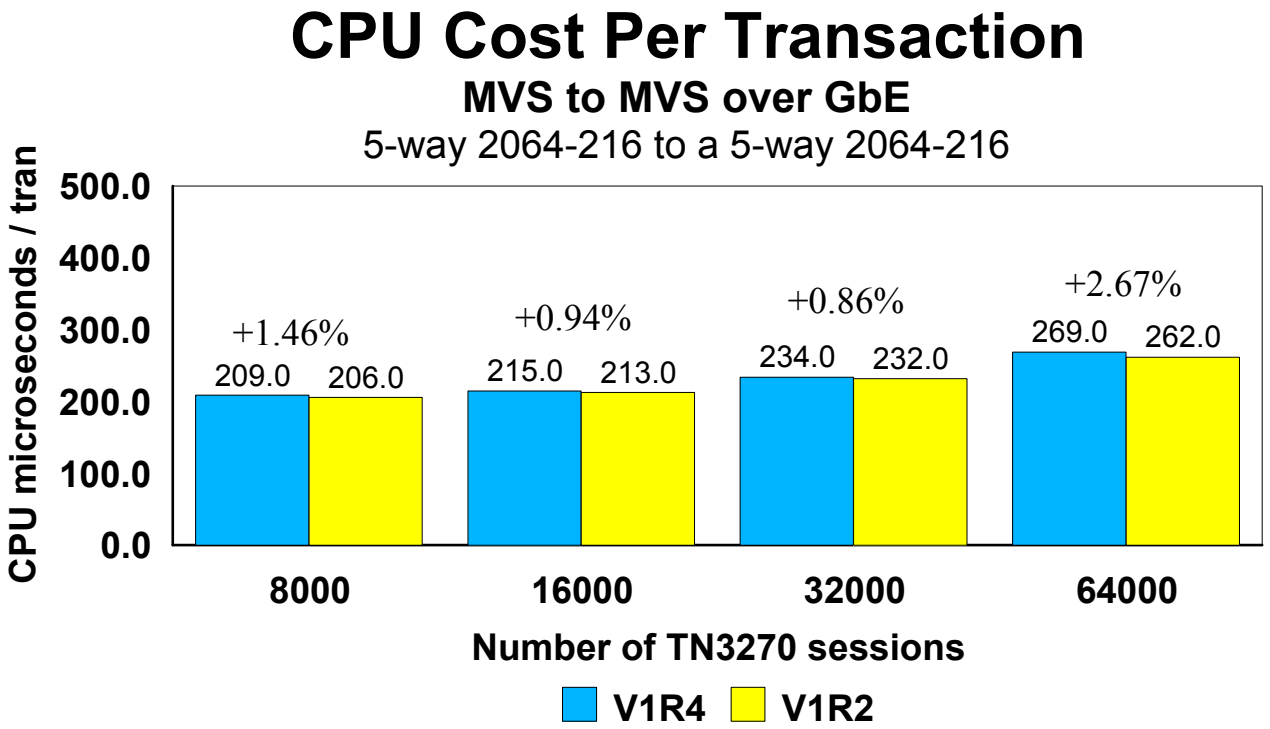

- V1R4 Telnet TN3270 CPU cost per transaction is within (0.86 to 2.67)% of V1R2

Release to Release Comparison

- **TN3270 Server: V1R4 relative to V1R2**

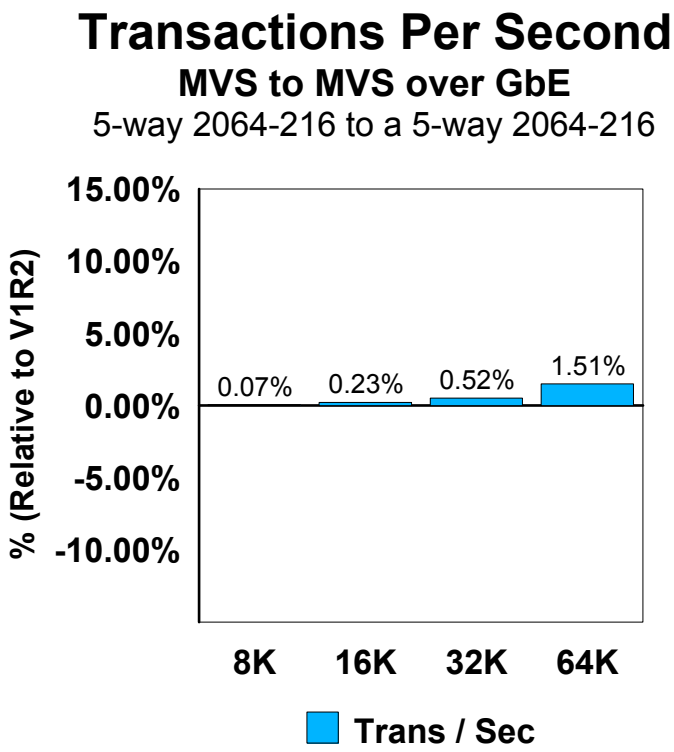

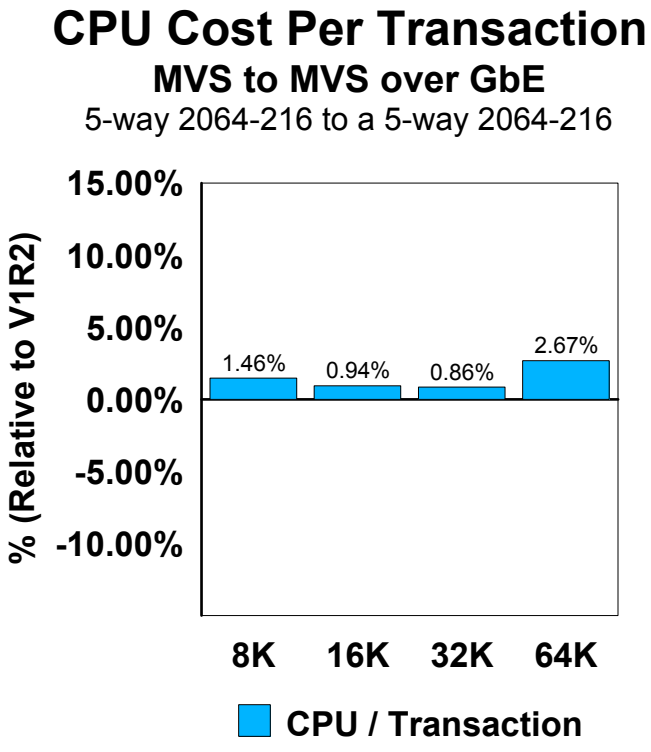

- V1R4 Telnet TN3270 provides transaction rates within (0.07 to 1.51)% and CPU cost per transaction within (0.86 to 2.67)% of V1R2

Telnet (TN3270) Storage Utilization (z/OS V1R4 CS)

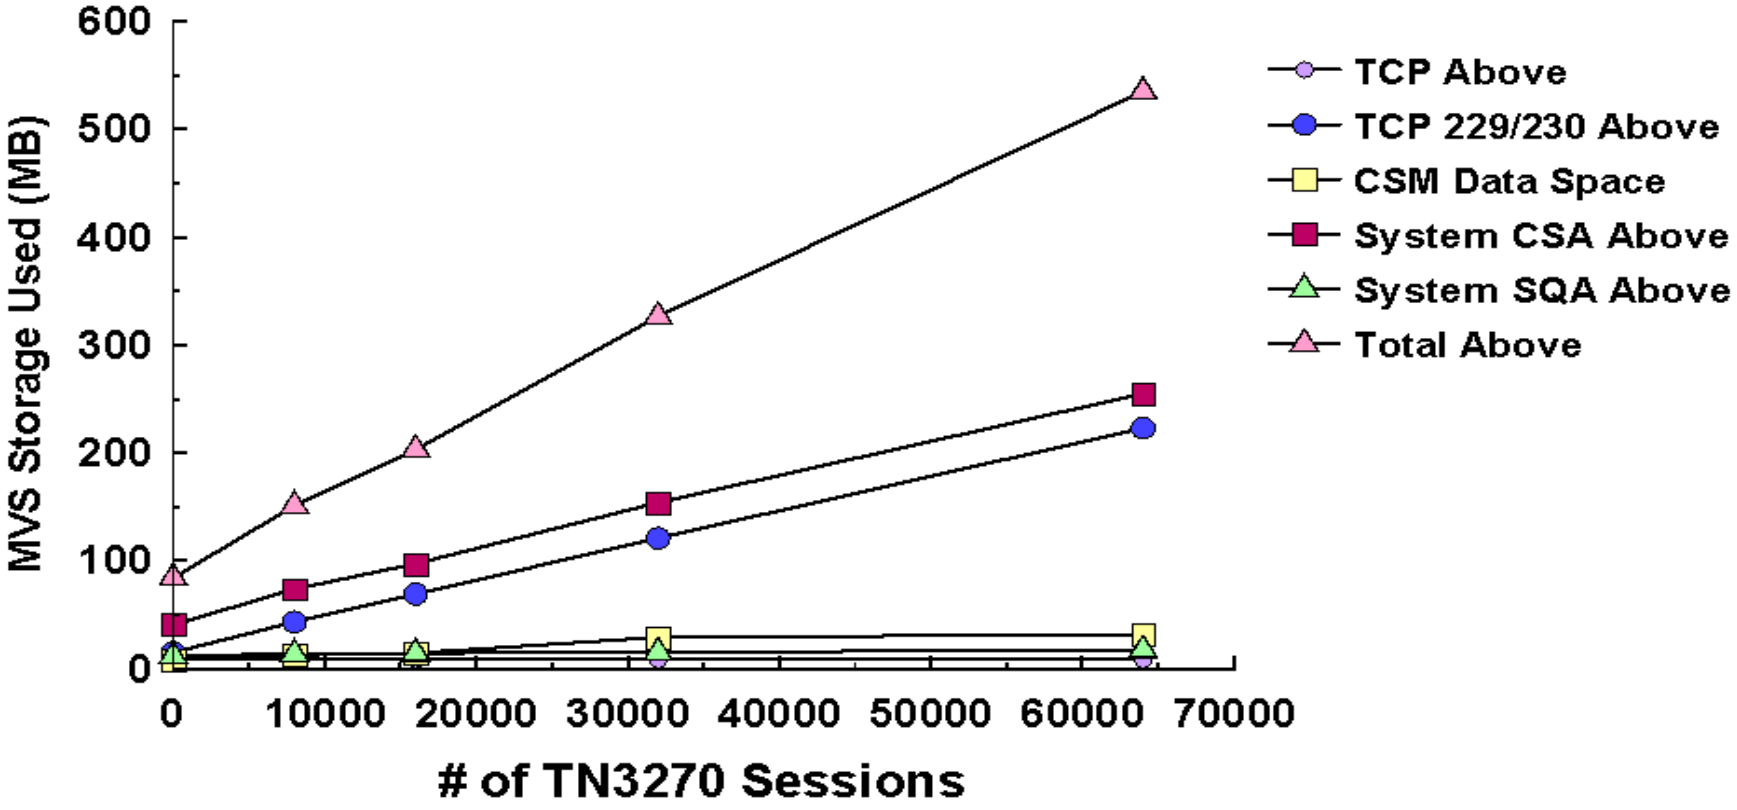

Storage usage (Above 16M line) of the TCP/IP Address Space and MVS System Storage (SQA and CSA) during TN3270 echoes (0 to 64000 users) when using z/OS V1R4 CS.

Telnet (TN3270) Storage Utilization (z/OS V1R4 CS)

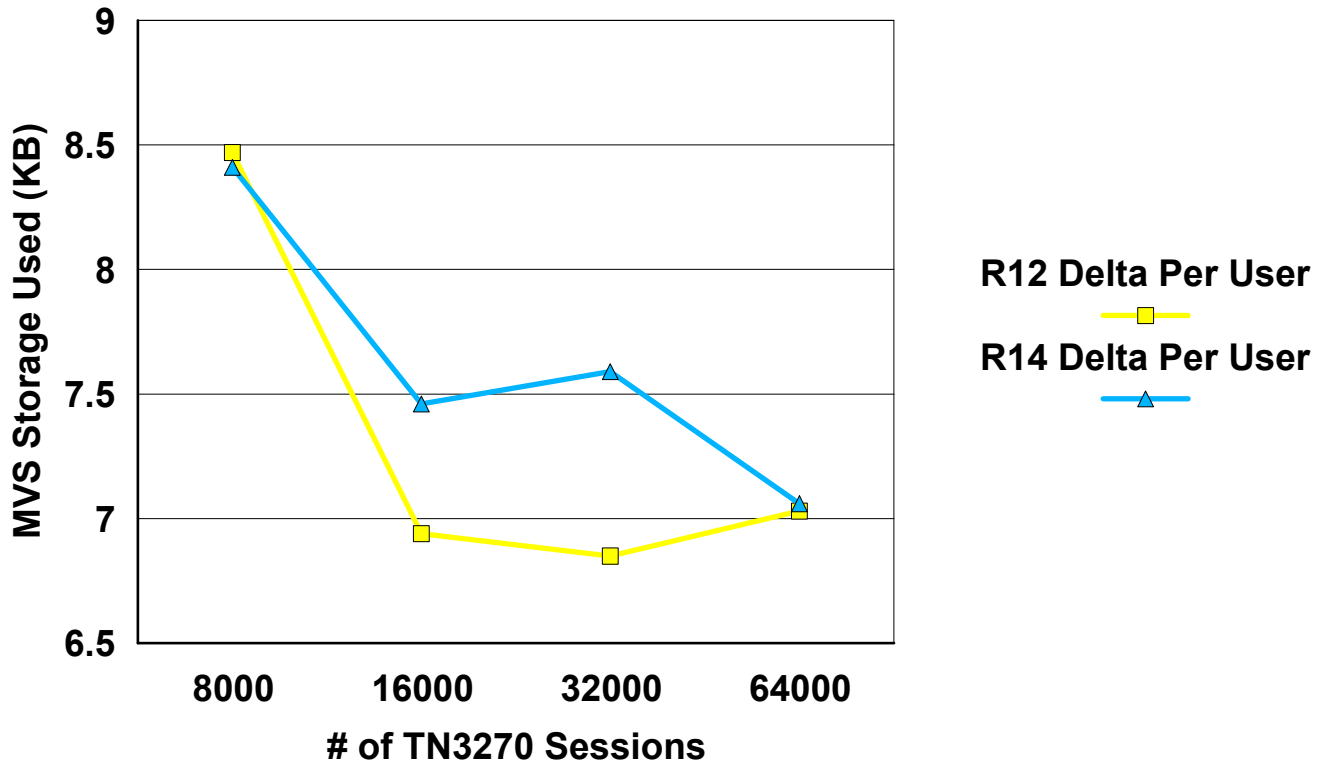

This chart shows z/OS V1R4 CS storage usage delta per user compared to V1R2 for TN3270 echoes (8000 to 64000 users).

Delta Per User Total: 7.06 to 8.41 KB / user

Telnet (TN3270) Storage Utilization (z/OS V1R4 CS)

# of TN3270 Sessions	0	8000	16000	32000	64000
TCP/IP Below	.180 M	.248 M	.264 M	.312 M	.440 M
TCP/IP Above	8.08 M	8.08 M	8.08 M	8.08 M	8.12 M
TCP/IP LSQA /SWA/**229/230** Below	.212 M	.224 M	.284 M	.284 M	.284 M
TCP/IP LSQA /SWA/**229/230** Above	**15.1 M**	**43.3 M**	**69.1 M**	**121 M**	**223 M**
CSM Data Space	8.71 M	12.54 M	14.47 M	288.3 M	31.84 M
System CSA Below	.516 M	.512 M	.516 M	.516 M	.516 M
System CSA Above	**40.600 M**	**73.6 M**	**97.3 M**	**154 M**	**255 M**
System SQA Below	.704 M	.776 M	.776 M	.776 M	.776 M
System SQA Above	11.6 M	13.7 M	14.3 M	15.3 M	17.3 M
Total Below	1.612 M	1.760 M	1.840 M	1.89 M	2.02 M
Total Above	84.09 M	151.22 M	203.25 M	326.68 M	535.26 M
Total	85.70 M	152.98 M	205.09 M	328.57 M	537.28 M
Delta Per User (KB)		8.41 KB	7.46 KB	7.59 KB	7.06 KB

Storage usage of the TCP/IP Address Space and MVS System Storage (SQA and CSA) during TN3270 echoes (0 to 64000 users) when using z/OS V1R4 CS.

TN3270 Capacity Planning (z/OS V1R4 CS)

MVS CPU Requirements:

$$\frac{\text{\# trans/user} \times \text{\# users} \times \text{CPU secs/tran}}{\text{\# of elapsed secs}} = \frac{\text{CPU secs}}{\text{elap secs}}$$

Example: z/OS V1R4, 4000 users, 6 tr/min/user

$$\frac{6 \text{ tr/u} \times 4000\text{u} \times 0.000209 \text{ CPU secs/tr}}{60 \text{ elap secs}} = 0.0836 \frac{\text{CPU secs}}{\text{elap secs}}$$

N1: MVS TCP/IP + VTAM + ECHO Application CPU
(2064-216 5 CP LPAR)

If the CPU secs/elapsed secs ratio is greater than 1, more than one processor would be required (z/OS V1R4 CS).

TN3270 Capacity Planning (z/OS V1R4 CS)...

MVS CPU Utilization:

$$\frac{\text{CPU secs/elapsed sec}}{\text{\# of processors}} \times 100\ \% = \text{CPU Util \%}$$

\# of processors: Should be equal to the number of 390 processors.

Example: z/OS V1R4, 4000 users, 6 tr/min/user

$$\frac{0.0836\ \text{CPU secs/elapsed sec}}{5\ \text{processors}} \times 100\ \% = 1.672\ \%$$

Thus, the MVS TCP/IP + VTAM + ECHO Application CPU requirement for 4000 TN3270 users would require 1.672% of a five processor LPAR 2064-216 system. LSPR can be used to adjust for other processors types.

z/OS V1R4 CS Performance - Configuring IPv6

- Netmarks Primitive Benchmarks (TCP)
  - V1R4 Stack and Application IPv6 Enabled vs. V1R4 IPv6 Disabled
  - Comparison - Trans/Second, CPU Cost Per Transaction

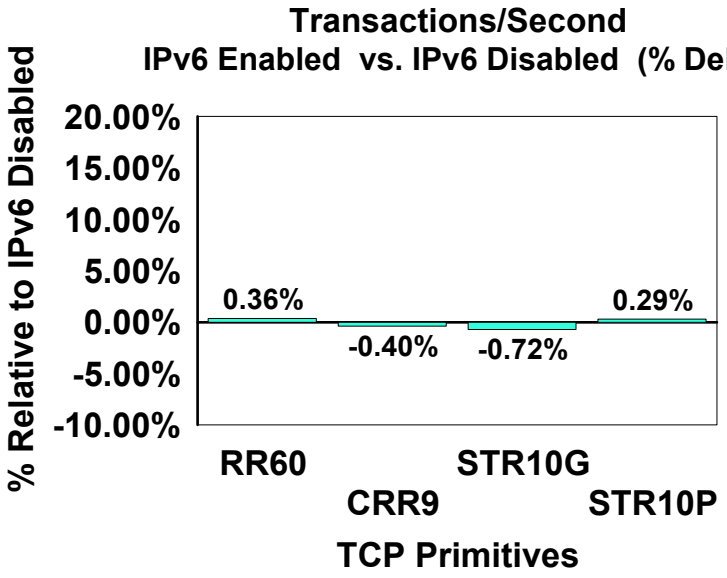

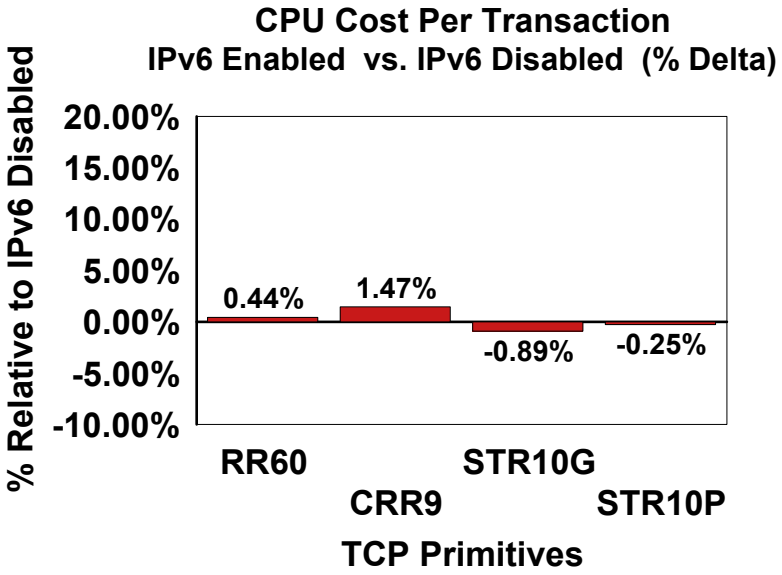

- V1R4 TCP/IP stack with IPv6 enabled and server application modified to be an IPv6 application; Client is an IPv4 application running on an IPv4 stack

- Differences in Transaction rate and CPU cost per transaction are minimal between an IPv6 enabled stack and application, and an IPv4 stack and application

z/OS V1R4 CS Performance - Configuring IPv6

- Netmarks Primitive Benchmarks (UDP)
  - V1R4 Stack and Application IPv6 Enabled vs. V1R4 IPv6 Disabled
  - Comparison - Trans/Second, CPU Cost Per Transaction

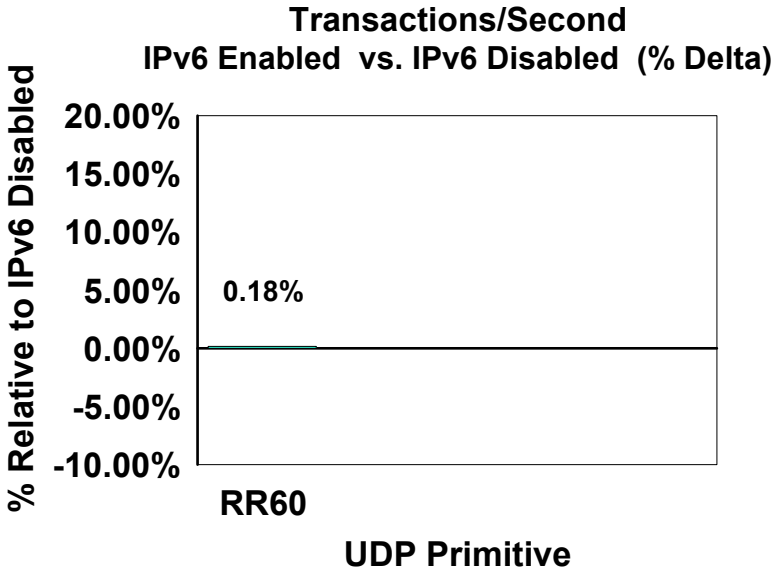

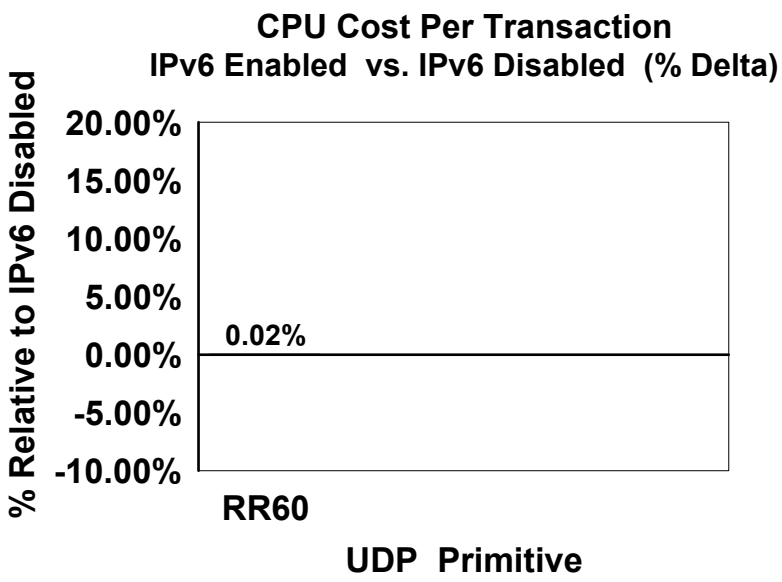

- V1R4 TCP/IP stack with IPv6 enabled and server application modified to be an IPv6 application; Client is an IPv4 application running on an IPv4 stack

- Differences in Transaction rate and CPU cost per transaction are minimal between an IPv6 enabled stack and application, and an IPv4 stack and application

z/OS V1R4 CS Performance - Configuring IPv6

- Netmarks Primitive Benchmarks (TCP)
  - V1R4 Stack IPv6 Enabled vs. V1R4 stack IPv6 disabled
  - Comparison - Trans/Second, CPU Cost Per Transaction

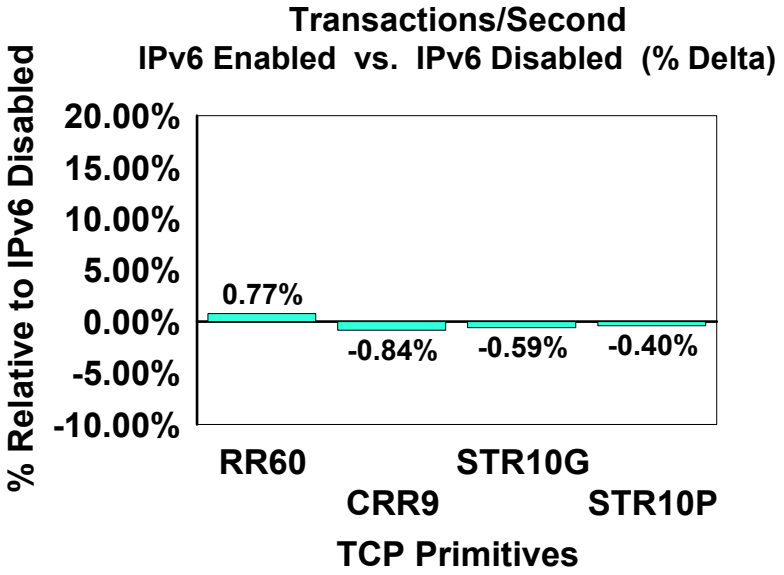

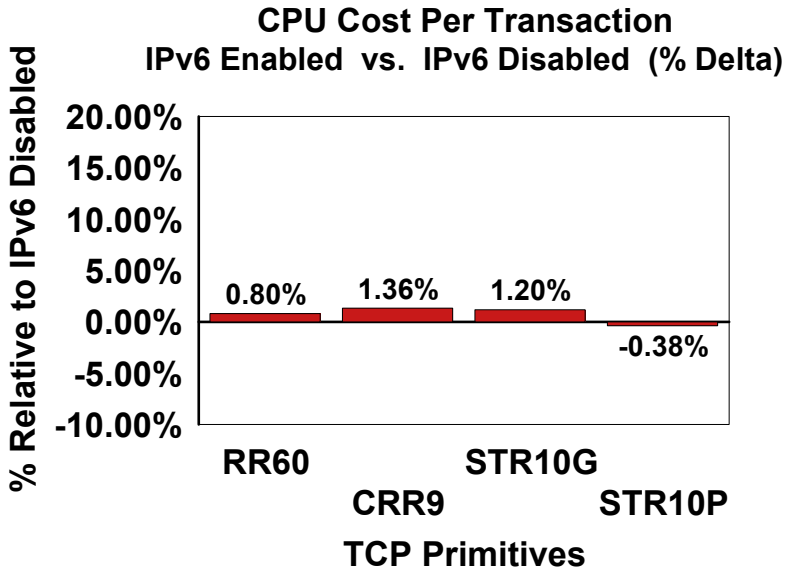

- Server is a V1R4 TCP/IP stack with IPv6 enabled running an unmodified IPv4 application; Client is an IPv4 application running on an IPv4 stack

- Differences in Transaction rate and CPU cost per transaction are minimal between running IPv4 applications on an IPv6 enabled versus disabled stack

z/OS V1R4 CS Performance - Configuring IPv6

- Netmarks Primitive Benchmarks (UDP)
  - V1R4 Stack IPv6 Enabled vs. V1R4 stack IPv6 disabled
  - Comparison - Trans/Second, CPU Cost Per Transaction

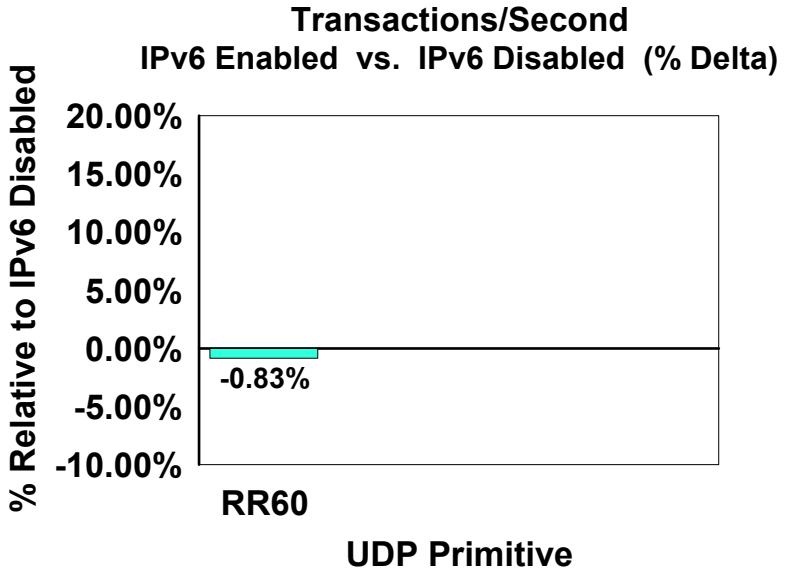

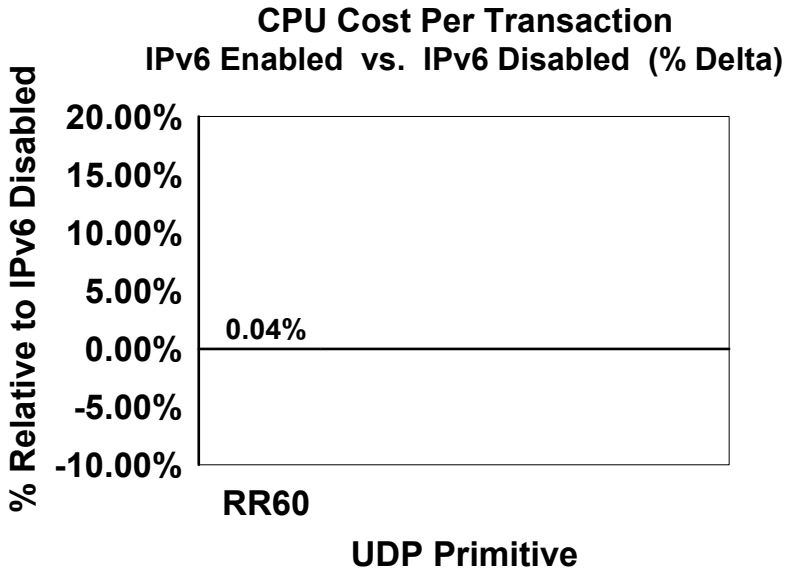

- Server is a V1R4 TCP/IP stack with IPv6 enabled running an unmodified IPv4 application; Client is an IPv4 application running on an IPv4 stack

- Differences in Transaction rate and CPU cost per transaction are minimal between running IPv4 applications on an IPv6 enabled versus disabled stack

Release to Release Comparison

- FTP Server
  - Effect on Throughput and CPU Cost Per Transaction of Enabling IPv6

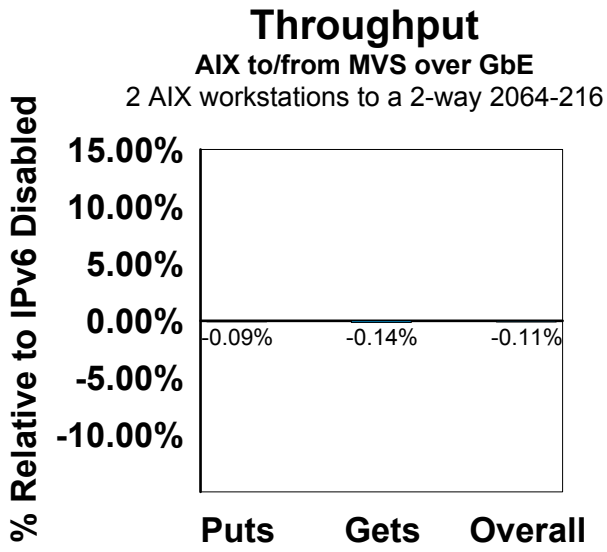

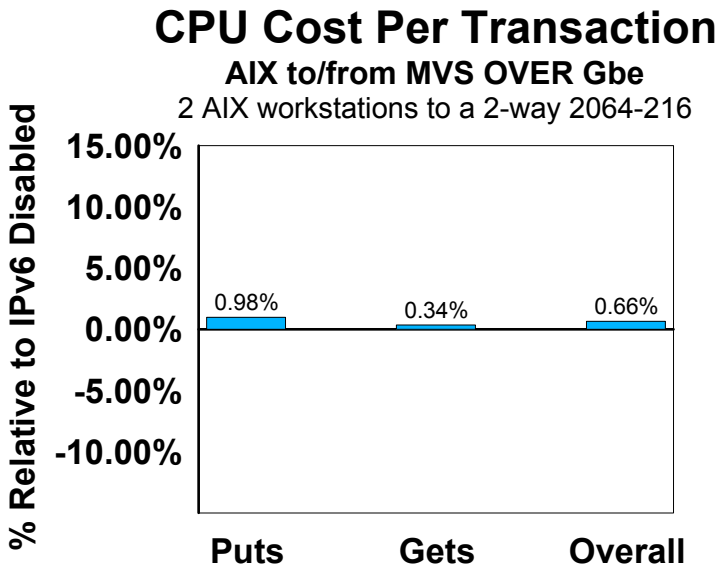

PUT: AIX -----> MVS (8 FTP Sessions, binary/ASCII PUT initiated from AIX clients)
GET: AIX <----- MVS (8 FTP Sessions, binary/ASCII GET initiated from AIX clients)

- There is negligible performance degradation when the V1R4 stack is enabled for IPv6 and running FTP server which has been modified to be an IPv6 application

z/OS V1R4 Comm Server Performance Summary

- Performance Tuning Considerations
  - General Recommendations (Refer session 3916 on this page for detail)
    - Turning traces off
    - Verify network and packet size
  - FTP Server and Client Tuning
    - Dataset block size (1/2 track as the size of dataset block size), FB with record length of 80 bytes
    - Client (use TCP send/recv buffer size of 128K)
  - Recommended PTFs - z/OS V1R4
    - APAR OW56019
    - Info APAR II11952
    - OSA Code level :

      G5/G6 ==>4.28, z/Series 2064: GA2==> 2.29, GA3 ==> 3.23
- z/OS V1R2 CS Performance References
  - http://w3-1.ibm.com/support/techdocs/atsmastr.nsf (Document: TD100541
    - z/OS V1R2 CS Performance Summary
    - IBM MVS TCP/IP Performance Tuning Tips and Capacity Planning (session 3916, 7/26/01)
  - http://www.share.org/proceedings/sh98/share.htm
    - IBM MVS TCP/IP Performance Tuning Tips and Capacity Planning (session 3916, 3/07/02)

z/OS V1R4 Comm Server Performance Summary

- Performance Summary
  - Release to Release Performance
    - Overall z/OS V1R4 Communication Server provides equivalent or better performance compared to z/OS V1R2
      - V1R4 provides better performance:
        - ✓ Bind 9.2 significantly improved throughput, scalability and cost per transaction compared to bind 9.1 used with V1R2
        - ✓ FTP server reduced CPU cost for inbound data transfer significantly compared to V1R2 FTP server
  - Effect of Enabling IPv6
    - Overall marginal differences were observed in performance between an IPv6 enabled stack/application communicating with IPv4 partners versus an IPv6 disabled stack communicating with IPv4 partners

z/OS V1R4 Performance References

- **z/OS V1R4 Performance References**
  - **S/390 e-business Performance (USS, LE, Java, Websphere, XML, net commerce)**
    - **http://www.ibm.com/servers/eserver/zseries/ebusiness/perform.html**
  - **Software capacity planning tool (SOFTCAP - to plan an upgrade to a new OS/390 or z/OS release)**
    - **The SOFTCAP tool can be downloaded from the following web sites:**
      - **External version: http://www.ibm.com/support/techdocs**
      - **IBM Business partners: http://www-1.ibm.com/partnerworld/sales/systems**