TCP/IP for z/OS Performance Tuning Tips and Capacity Planning

Robert Perrone
IBM Corporation
(bperrone@us.ibm.com)

Share 102, Session 3919

February 25, 2004

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.

Presentation Overview

- CS/390 R10, R12, R14 & R15 Performance Summary
- IBM MVS TCP/IP CS/390 R4, R5, R6, R7, R8, R10, R12, R14 & R15 Performance Tuning
- TCP/IP z/OS Unix Performance Tuning
- FTP Performance Tuning / Capacity Planning
- Telnet Performance Tuning / Capacity Planning
- CICS Sockets Tuning Parms
- TCP/IP Tuning Performance Checklist

CS/390 R15 Performance Highlights

- IPv6
  - IPv6 support for TN3270(E) server, CICS Sockets, MPC+
  - Minimal performance impact when running IPV6

    - TN3270E Server:
      Client: IPv6 Server: IPv6 enabled, IPv6 TN3270E Server
      TPUT: + 1.1 to + 1.5 % (vs IPv4)
      CPU/Tran: + 1.8 to + 4.3 % (vs IPv4)

    - CICS Sockets:
      Client: IPv6 Server: IPv6 Listener & Child Server
      TPUT: - 0.43 % (vs IPv4)
      CPU/Tran: + 1.16 % (vs IPv4)

    - FTP Server:
      Client: IPv6 Server: IPv6 enabled, IPv6 FTP Server
      TPUT: - 0.8 to - 3.6 % (Avg: - 2.6 %) (vs IPv4)
      CPU/Tran: + 0.5 to + 3.6 % (Avg: +2.2 %) (vs IPv4)

    - MPC+ (AWM Streams Workload):
      Client: IPv6 Server: IPv6 enabled, IPv6 AWM Server
      TPUT: Equivalent (vs IPv4)
      CPU/Tran: + 2.9 % (vs IPv4)

CS/390 R15 Performance Highlights

- TN3270E Definite Response Performance Improvement :
  - Specify 'TCPCONFIG DELAYACKS' in TCP profile
  - TPUT: + 3.1 % CPU/Tran: - 4.2 % (vs NODELAYACKS)

- Asynchronous I/O :
  - Performance improvement with new async I/O model
  - TPUT: + 0.4 % CPU/Tran: - 12.1 %
    vs old async I/O model (Receive only)

- Enterprise Extender (EE) :

  - Multiple VRN's :
    - TPUT: + 3.2 % CPU/Tran: - 4 to + 2.7 %
      vs single source/destination pair
  - MPC+ HPDT Packing :
    - Supports 1K, 2K, 4K packing
    - TPUT: + 25 to + 26 % CPU/Tran: - 0.3 to + 2.6 %
      vs no packing (AWM Request/Response workload)

- DVIPA Limit enhancements :
  - New limit is 1024 (was 256)

CS/390 R15 Performance Highlights

- New Crypto instructions (available on zSeries 990)
  - IPSEC Crypto Assist Performance Improvements (Hardware vs Software encryption):
    - Request/Response Workload :
      TPUT: + 203 % CPU/Tran: - 67 %
    - Connect Request/Response Workload :
      TPUT: + 228 % CPU/Tran: - 72 to - 75.8 %
    - Streams Workload :
      TPUT: + 741 % CPU/Tran: - 90 %

CS/390 R15 Performance Highlights

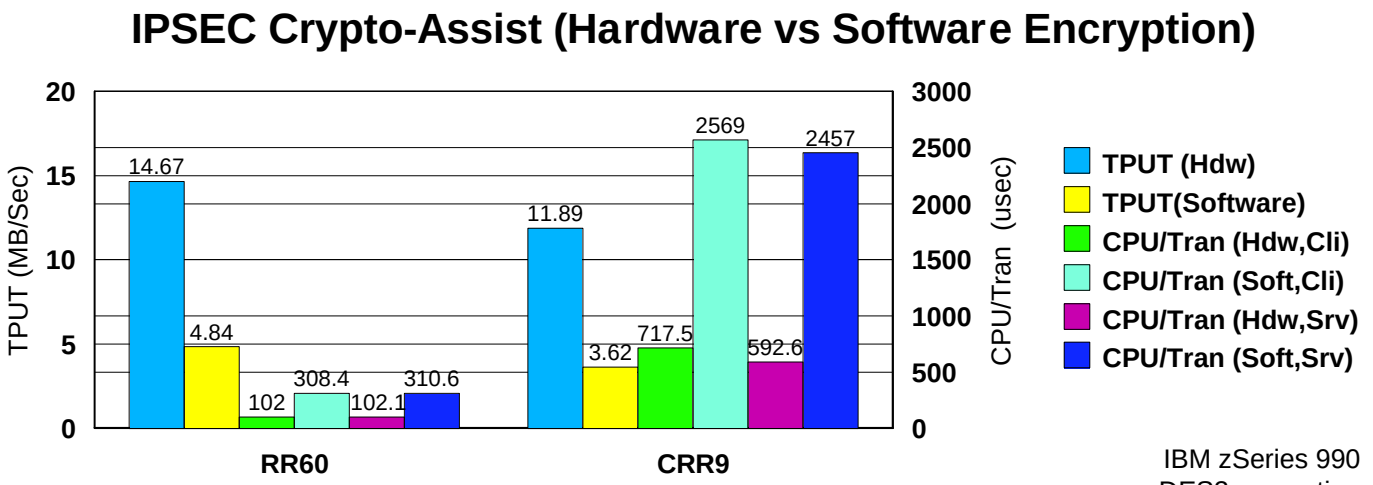

IBM zSeries 990
DES3 encryption
no filter rules, no authentication

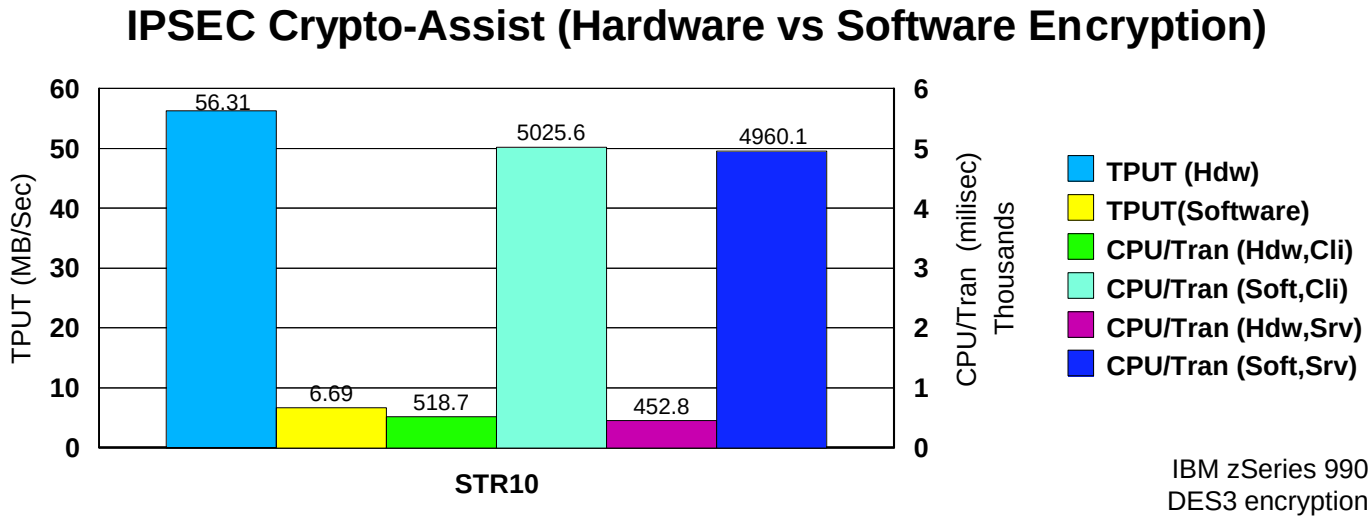

IBM zSeries 990
DES3 encryption
no filter rules, no authentication

CS/390 R15 Performance Highlights

- OSA Express Enhancements :
  - Checksum Offload :
    - TCP/IP stack Checksum function offloaded to OSAE adapter
    - Requires a zSeries 990 system
    - Request/Response Workload :
      TPUT: Equivalent CPU/Tran: - 0.6 to - 1.4 %
    - Connect Request/Response Workload :
      TPUT: - 2.1 % CPU/Tran: - 5.5 to - 6.0 %
    - Streams Workload :
      (TCP) TPUT: - 1.6 % CPU/Tran: - 11.5 to - 13.8 %
      (EE) TPUT: + 145 % CPU/Tran: - 9.6 to - 14.3 %
  - INBPERF :
    - New parm on Link statement (INBPERF xxxxxxx)
    - BALANCED : Default (Recommended)
    - MINCPU : minimizes cpu
      TPUT: - 29.6 to - 0.5 % CPU/Tran: - 5.3 to + 0.7 %
      Reduces cpu/tran for RR & CRR with a loss of tput
    - MINLATENCY: minimizes latency
      TPUT: - 2.1 to + 5.2 % CPU/Tran: - 4.8 to + 2.4 %
      Improves tput for RR and CRR workloads,
      Degrades Streams workload

CS/390 R14 Performance Highlights

- IPV6
  - z/OS V1R4 : First release supporting IPV6
  - Design goal was to incorporate IPV6 protocol into communications stack with minimal impact to IPV4 performance
  - Single TCP/IP stack supporting IPV4 and IPV6 protocols at same time
  - Can be run as V4 or V4/V6 stack
  - To enable IPV6 support: Add AF_INET6 network stmt to BPXPRMxx member
  - Minimal performance impact when running IPV6

    - AWM Primitives (RR60, CRR9, STR10 Put, STR10 Get):
      Client: IPv4 Server: IPv6 enabled, IPv4 application
      TPUT: -0.84 to +0.77 % (Avg: -0.27 %)
      CPU/Tran: -0.38 to +1.36 % (Avg: +0.75 %)
      Client: IPv4 Server: IPv6 enabled, IPv6 application
      TPUT: -0.72 to +0.36 % (Avg: -0.12 %)
      CPU/Tran: -0.9 to +1.47 % (Avg: +0.19 %)
    - FTP Server: (Binary/Ascii Put and Gets, 8 sessions):
      Client: IPv4 Server: IPv6 enabled, IPv6 FTP Server
      TPUT: -0.09 % (Puts) -0.14 % (Gets) (Avg: -0.11 %)
      CPU/Tran: +0.98 % (Puts) +0.34 % (Gets) (Avg: +0.66 %)

    - Overall average z/OS CPU increase is 0.53 % with negligible tput degradation

- DNS bind 9.2
  - Significant performance (CPU & Throughput) and scalability improvements for DNS bind 9.2 compared to DNS bind 9.1 (CS/390 R12)
    - TPUT Improvement: +35 to + 186 % CPU Reduction: 31 to 43 %

CS/390 R14 Performance Highlights

- OSA Express / HiperSockets Storage Enhancement
  - Amount of storage used for read processing for both OSA Express and HiperSockets devices has been increased
  - OSA Express (QDIO): 1 to 4 MB (was 0.5 MB)
  - HiperSockets (IQDIO): 2 to 8 MB (was 2 to 4 MB)
  - New VTAM start options (QDIOSTG, IQDIOSTG), display and modify commands are provided to adjust amount of storage used
  - Defaults settings should be appropriate in most cases
  - Storage adjustment may be necessary when many OSA adapters are used, multiple TCP/IP stacks per LPAR or many VM 2nd level guests
  - Available on CS/390 R12 with APAR OW52291

- OSA Express Microcode Enhancement:
  - Improves inter-packet delay and max block size to optimize performance
  - Interpacket delay: 40 microsecond Max Blk Size: 1 millisecond
  - Requires z/OS V1R4 APAR OW56019 (V1R4 PTF UW94043)
  - Info APAR II11952
  - OSAE Code level :
    - G5/G6: 4.28 or later
    - z/Series 2064: GA2 ==> 2.29 or later, GA3 ==> 3.23 or later

CS/390 R14 Performance Highlights

- SNA Enhancements
  - Open ACB limit expanded from 65K to 1,044,480
  - CSALIMIT start option change:

    VTAM continues to execute beyond the specified CSALIMIT provided sufficient CSA and ECSA storage is available
  - Display command for RTP physical units has been enhanced to provide additional data for HPR/EE performance analysis

    D NET,ID=rtpname,HPRDIAG=YES

CS/390 R12 Performance Highlights

- Hipersockets
  - LPAR to LPAR communication via memory (same CEC)
  - Supported on IBM eServer zSeries 900 processors
  - Speed equivalent to cross addr space memory move
  - Up to 4 hipersocket Lans per CEC
  - MTUs: 8KB, 16KB, 32KB, 56KB
  - Configure with IPCONFIG DYNAMICXCF stmt or Device/Link stmts

- Hipersockets Accelerator
  - Allow a single TCP/IP stack to act as a "router" for entire CEC
  - 'Accelerate' routing by bypassing entire TCP/IP stack
  - CPU/Tran Reduction on Router: 49 % (RR), 75 % (Streams)
  - Reduces number of direct connections to external network
  - IPCONFIG IQDIOROUTING QDIOPRIORITY 3

- Full CLAW Packing
  - Pack multiple datagrams into a single claw channel frame
  - Pack up to 60 KB claw channel buffer
  - Max Read & Write Size: 256 KB (ex. 8 x 32KB = 256KB)
  - TPUT Improvement: 31 to 66 % CPU/KB Reduction: 3 to 24 %
  - R8 & R10: APAR PQ41205 (limited claw packing: 4KB)
  - Recommended for all customers using Cisco 7200/7500 (via claw)
  - MVS: DEVICE CLAW7500 CLAW HOST1 CISCO1
    PACKED 8 8 32K 32K
    Router: CLAW EF73 80 48.1.1.1 HOST1 CISCO1 PACKED PACKED

CS/390 R12 Performance Highlights

- TCP Protocol Configuration Options:
  - TCPTIMESTAMP
    TCPCONFIG TCPTIMESTAMP (default) or
    TCPCONFIG NOTCPTIMESTAMP
  - FINWAIT2
    TCPCONFIG FINWAIT2TIME 60 (60-3600 sec, default = 675)
  - Display current settings:
    D TCPIP,,NETSTAT,CONFIG

- TCP/IP Storage Management:
  - GLOBALCONFIG ECSALIMIT 250M POOLLIMIT 300M
  - Default is to set no storage limits
  - ECSALIMIT does not include ECSA CSM storage
  - POOLLIMIT is authorized private storage in TCP/IP Addr Space
  - D TCPIP,,STOR (Shows current storage limits)
  - 4 TCP/IP Storage Classes :
    Normal: < 80 % of Storage Limit
    Constrained: 80 % <= Constr < 90 %
    Critical: 90 % <= Critical < 98 %
    Exhausted: exhausted >= 98 %
  - Storage warning messages displayed on MVS console

CS/390 R12 Performance Highlights

- TCP/IP Performance Statistics:
  - D TCPIP,,NETSTAT,STATS
  - New MIB objects defined for SNMP
  - Displays IP, TCP, UDP, ICMP performance statistics

- Enterprise Extender Performance Improvements
  - Minimize route lookups
  - EE fastpath (Bypass IP layer, call IF layer)

- 64 bit real addressing support
  - Back most CSM data space above the 2 GB line
  - CSM displays are changed to support 64 bit
  - VTAM Option:
    API64R=YES (64 bit backed storage can be passed to application)
    NO (Forces a copy out of 64 bit backed storage before being given to application)

- IPSEC Performance Improvment
  - Previously, TCP connections using IPSEC were not able to take advantage of Path MTU Discovery.
  - Track negoiated MTU as part of PATH MTU discovery on a tunnel basis
  - Avoid calls to firewall code at IP layer

HiperSockets Performance (Interactive Transactions)

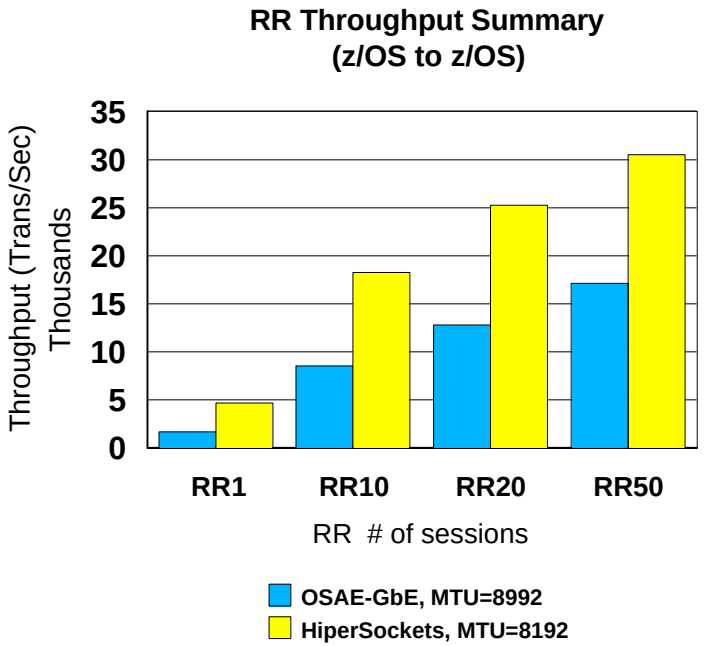

z/OS (4 Ded CP LPAR) to z/OS (4 Ded CP LPAR)
2064-116, Transaction: 200 bytes out, 1000 bytes in
OSAE-GbE (Shared adapter between two LPARs)
z/OS Level: V1R2

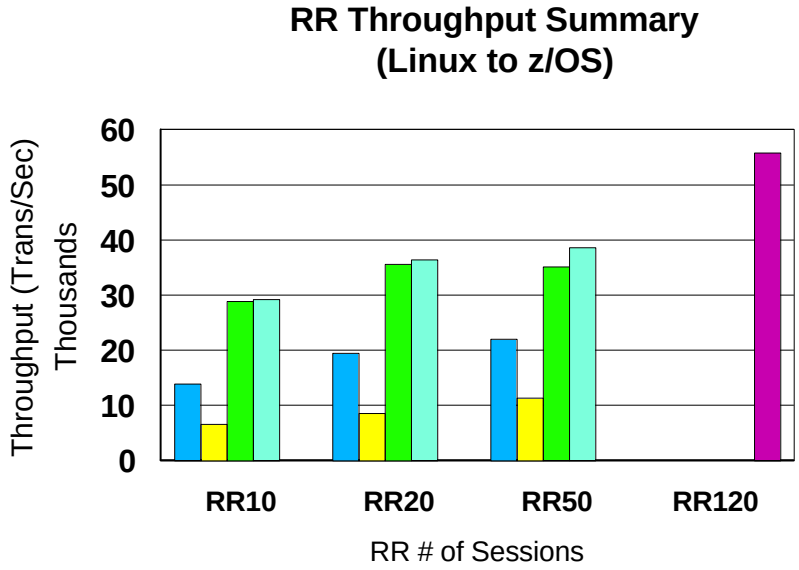

LnNetfin: Pentium 3, 866 Mhz, 2 CPs
2064 -116, Transaction: 200 bytes out, 1000 bytes in
z/OS Level: V1R2, Linux: 2.4.7 kernel

HiperSockets Throughput Summary (RR):

z/OS to z/OS : 1.8 to 2.1 X vs OSAE-GbE (Shared)

Linux/390 to z/OS : 3.4 to 4.5 X vs OSAE-GbE (Shared)

HiperSockets Response Time Reduction (RR):

z/OS to z/OS : 46 to 52 % vs OSAE-GbE (Shared)

Linux/390 to z/OS : 69 to 74 % vs OSAE-GbE (Shared)

HiperSockets Performance (Bulk Data Xfer)

STREAMS (PUT) Throughput Summary (z/OS to z/OS)

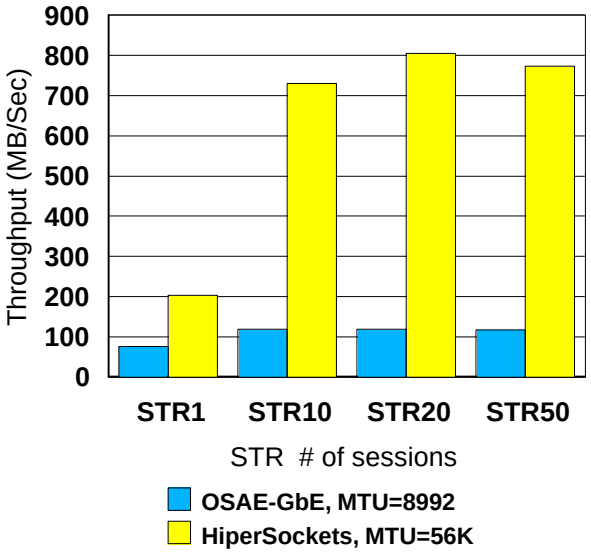

z/OS (4 Ded CP LPAR) to z/OS (4 Ded CP LPAR)
2064-116, Transaction: 20 Mbytes out, 20 bytes in
OSAE-GbE (Shared adapter between two LPARs)

z/OS TCP Send/Receive socket buffers = 262144
z/OS Level: V1R2

STREAMS (PUT) Throughput Summary (Linux to z/OS)

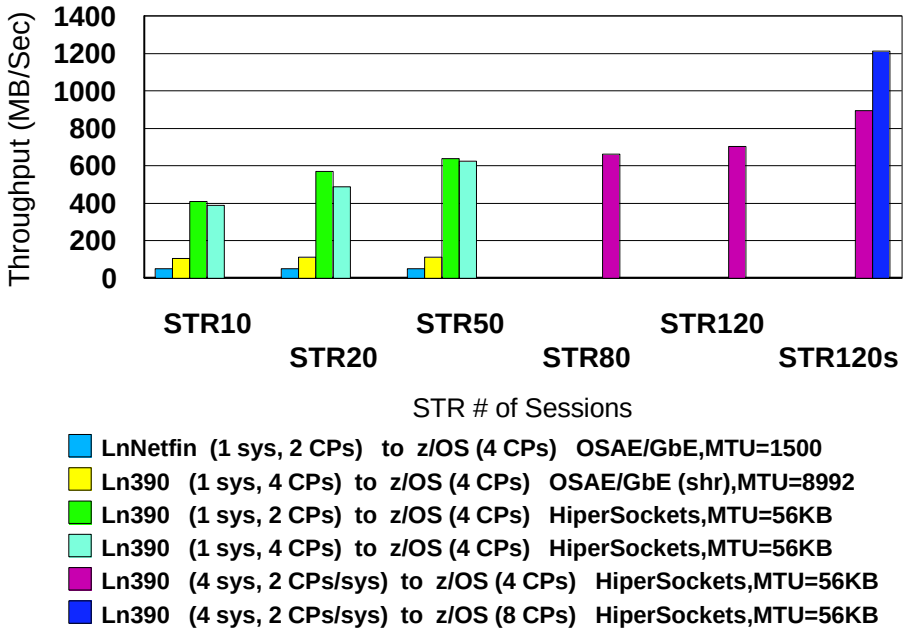

s: Linux TCP Send/Receive Socket Buffers = 131072
LnNetfin: Pentium 3, 866 Mhz, 2 CPs
2064 -116, Transaction: 20 Mbytes out, 20 bytes in

z/OS TCP Send/Receive socket buffers = 262144
Linux TCP Send/Receive socket buffers = 65535
z/OS Level: V1R2, Linux: 2.4.7 kernel

HiperSockets Throughput Summary (STREAMS):

z/OS to z/OS : 6.2 to 6.8 X vs OSAE-GbE (Shared)

Linux/390 to z/OS : 3.7 to 5.5 X vs OSAE-GbE (Shared)

CS/390 R10 Performance Highlights

- SAP Performance
  - 8.1 % ITR (Tran/CPU) Improvement in SAP Sales and Distribution (SD) Benchmark (vs CS/390 R8)
- Fastpath Local Sockets
  - All **non-loopback interfaces** take fast path compared to loopback
  - 2 X increase in transactions/sec (vs CS/390 R8 Local Sockets)
  - 48 % or greater reduction in latency (vs CS/390 R8 Local Sockets)
- Route Lookup Enhancements
  - Up to 50 % CPU reduction when using CS/390 R10 as an intermediate router
- IP Security
  - 25 % reduction in Client CPU (vs CS/390 R8)
  - 10 % improvement in throughput (vs CS/390 R8)
- UNIX Select() Scalability & Performance Improvements
  - Response time of select() with 1000 file descriptors reduced to 1 ms from 60 ms.
  - In OS/390 R10 base, PTFed (OW44754) back to R8.
- FRCA-Webspere responsibility passing
  - Dynamic passing of connection ownership between FRCA and Websphere cache for persistent connections.

IBM MVS TCP/IP CS/390 R4-R15 Performance Tuning

- TCP/IP Buffers :
  - Number of TCP/IP Buffers no longer specified in TCP Profile
  - TCP/IP Buffers are dynamically allocated from **CSM**

- TCP/IP Send / Receive Buffer Sizes:

  Default Size = 16 KB

  - Override Send/Receive buffer size for <u>all</u> applications (TCP/IP Profile):

TCPCONFIG	TCPSENDBFRSIZE 65535
	TCPRCVBUFRSIZE 65535
UDPCONFIG	UDPSENDBFRSIZE 65535
	UDPRCVBUFRSIZE 65535
	NOUDPQUEUELIMIT

  - Override Send Receive buffer size for <u>one</u> application:

    Use setsockopt(SO_SNDBUF) to set TCP Send buffer in application
    Use setsockopt(SO_RCVBUF) to set TCP Receive buffer in application

IBM MVS TCP/IP CS/390 R4-R15
CSM Performance Tuning

- CSM Storage Settings:

SYS1.PARMLIB (IVTPRM00):
Fixed MAX(x M) Recommend 60 M
ECSA MAX(y M) Recommend 40 M

Display cmds: D NET,CSM or
D NET,CSM,ownerid=all

CSM Usage (R12):

Workload	# Users / Clients	TPUT	MAX CSM (ECSA)	MAX CSM (Data Space)	Max CSM (FIXED)
Web Server	80	5425 c/s	2.97 MB	8.0 MB	11.2 MB
CICS Sockets	84	409 c/s	0.736 MB	1.96 MB	3.8 MB
TN3270 (Echo's)	4000 8000 16000 32000 64000	399.1 tr/sec 798.5 tr/sec 1591.5 tr/sec 3115.0 tr/sec 5732.5 tr/sec	0.85 MB 1.40 MB 1.34 MB 8.10 MB 17.90 MB	5.1 MB 11.1 MB 8.1 MB 12.6 MB 27.4 MB	12.5 MB 13.6 MB 17.5 MB 24.2 MB 50.0 MB
FTP Server	16 Inbound 16 Outbound	49.8 MB/S 68.9 MB/S	4.5 MB 0.54 MB	6.8 MB 4.4 MB	9.6 MB 6.4 MB

IBM MVS TCP/IP CS/390 R4-R15
VTAM Buffer Performance Tuning

- VTAM Buffer Settings:
  - VTAM Start Options:

    Set IOBUFF, LFBUF, CRPLBUF, TIBUF and CRA4BUF using application usage below as a guideline.
  - Display cmds: D NET,BFRUSE,BUFFER=SHORT
    D NET,STORUSE

VTAM Buffer Max Usage (R12):

Work load	# Users / Clients	TPUT	VTAM Buffer (IO00)	VTAM Buffer (LF00)	VTAM Buffer (CRPL)	VTAM Buffer (TI00)	VTAM Buffer (CRA4)
Web Server	80	5425 c/s	6	4	2	4	4
CICS Sockets	84	409 c/s	26	5	54	29	6
TN3270 (Echo's)	4000 8000 16000 32000 64000	399.1 tr/sec 798.5 tr/sec 1591.5 tr/sec 3115.0 tr/sec 5732.5 tr/sec	112 112 345 2005 11000	4005 8005 16005 32003 64005	8007 16012 32019 64018 128018	5 5 5 5 5	25 41 49 65 170
FTP Server	16 Inbound 16 Outbound	49.8 MB/Sec 68.9 MB/Sec	5 5	4 4	2 2	3 3	4 4

MVS TCP/IP and z/OS Unix Performance *Tuning*

- Follow the z/OS Unix performance tuning guidelines in the z/OS Unix System Services Planning manual (GA22-7800) or WWW.
  - http://www-1.ibm.com/servers/eserver/zseries/zos/unix/bpxa1tun.html
- Follow IBM MVS TCP/IP Performance checklist.
- Update your MVS TCP/IP Profile, TCPIP.DATA and FTP.DATA files.
- Estimate how many z/OS Unix users, processes, ptys, sockets and threads would be needed for your z/OS Unix installation. Update your BPXPRMxx member in SYS1.PARMLIB.

MVS TCP/IP and z/OS Unix Performance Tuning (con't)

- OS/390 R1-R3 :
  Estimate how many ASCH initiators would be needed for your Unix installation. Update your ASCHPMxx member in SYS1.PARMLIB.

- Spread z/OS Unix user HFS datasets among many DASD volumes for optimal performance.

- Monitor your z/OS Unix resources with RMF and/or system commands (DISPLAY ACTIVE, DISPLAY OMVS, DISPLAY ASCH, DISPLAY APPC, etc.).

- Adjust z/OS Unix system parms to improve performance.

MVS TCP/IP and z/OS Unix Performance Tuning (BPXPRMxx)

BPXPRMxx (SYS1.PARMLIB) Tuning:

- Optimally set Max... parms.
  - Make sure MAXPROCSYS, MAXPROCUSER, MAXUIDS, MAXFILEPROC, MAXPTYS, MAXTHREADTASKS and MAXTHREADS are optimally set.
  - If these parms are not optimally set, your z/OS Unix performance may be degraded. For more information, see the z/OS Unix Services Planning manual (GA22-7800).
- Set MAXSOCKETS(n) to a high number to avoid shortage.
  - Make sure the MAXSOCKETS(n) parm for the AF_INET domain is set high enough to avoid running out of z/OS Unix sockets.
  - As an example, each z/OS Unix telnet session would require 1 z/OS Unix socket and each FTP session would require 1 z/OS Unix socket. Once the MAXSOCKETS limit is reached, no more telnet, FTP sessions or other apps that require z/OS Unix sockets would be allowed to start.

FTP Tuning Summary

- MVS CPU decreases as packet size (MTU) increases
- Set Client TCP/IP send/receive socket buffers optimally.
  - Recommend at least 64 KB
- MVS throughtput increases and MVS CPU decreases as Workstation window size increases
  - Recommended WS Window size = 64 KB
- MVS throughput increases and MVS CPU decreases as MVS TCP/IP TCP Window size increases
  - Recommended MVS Window size = 64 KB

Note: MVS FTP Server and Client sets its TCP send/receive buffers to 180 KB. Thus, the MVS FTP Server and Client will use a 64 KB window size by default.

For other applications that send/receive large streams of data, make sure their TCP or UDP send/receive buffers are set to a minimum of 64 KB.

FTP Tuning Summary (continued)

- For CLAW devices (CISCO Channel Attached Routers), use CLAW packing for better performance.
  Set read and write buffers on CLAW Device statement (MVS TCP/IP Profile):
  V1R2 : Max Read & Write sizes <= 256KB (ex. 8 x 32KB = 256 KB)
  V2R10: Use 8 read & write buffers (ex. 8 x 4 KB = 32 KB)

- MVS throughput greatly increases as MVS file blocksize increases.
  - Recommended file blocksize < = 1/2 DASD track.

- For best performance, keep CHKPTINT parm (in TCPIP.FTP.DATA) = 0.

- The number of ESCON channels to your DASD greatly influences overall throughput. Average throughput over an ESCON channel to your DASD ranges between 10 to 12 MB/Sec.
  - Recommend using multiple ESCON channels to your DASD

- For best FTP performance, use FICON channels to DASD.

- File System characteristics (Caching, file blksize, dasd speed, etc.) can greatly influence FTP performance (CPU & Throughput).

MVS CS/390 V1R5 FTP Server Throughput Performance (MTU=1500)

- V1R5 Single session FTP throughput ranges from 17.9 to 29.1 MB/Sec
- V1R5 Multi session (6) FTP throughput ranges from 59 to 64.9 MB/Sec

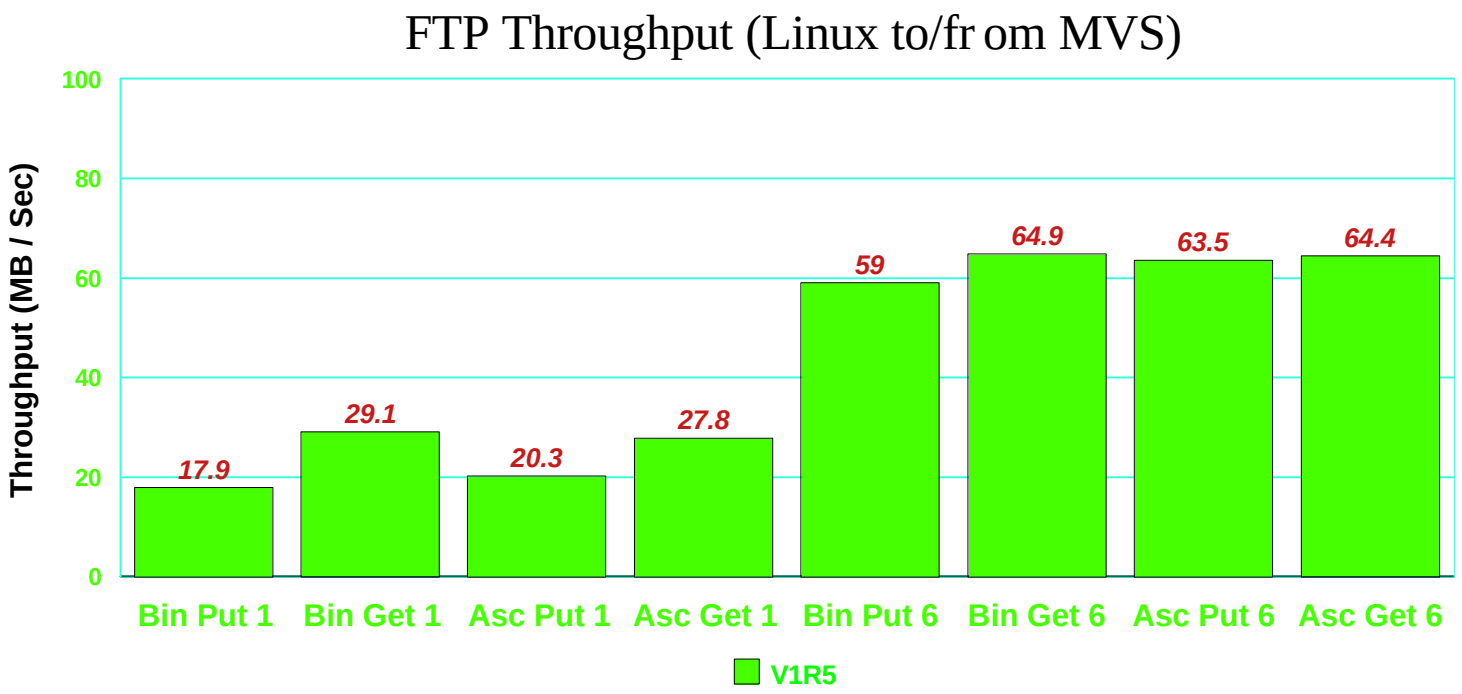

Put: Linux -----> MVS, Get: Linux <------ MVS

Config: Linux WS (1 or 2)---GbE---Switch----OSAE/GbE---2084-B16 (2 CPs)

DASD Config: 2 Banks of Shark DASD (2 FICON Chipids)

MVS CS/390 V1R5 FTP Server CPU Performance (MTU=1500)

- V1R5 Single session FTP CPU/KB ranges from 4.45 to 12.86 microseconds per KB
- V1R5 Multi session (6) FTP CPU/KB ranges from 4.03 to 11.39 microseconds per KB

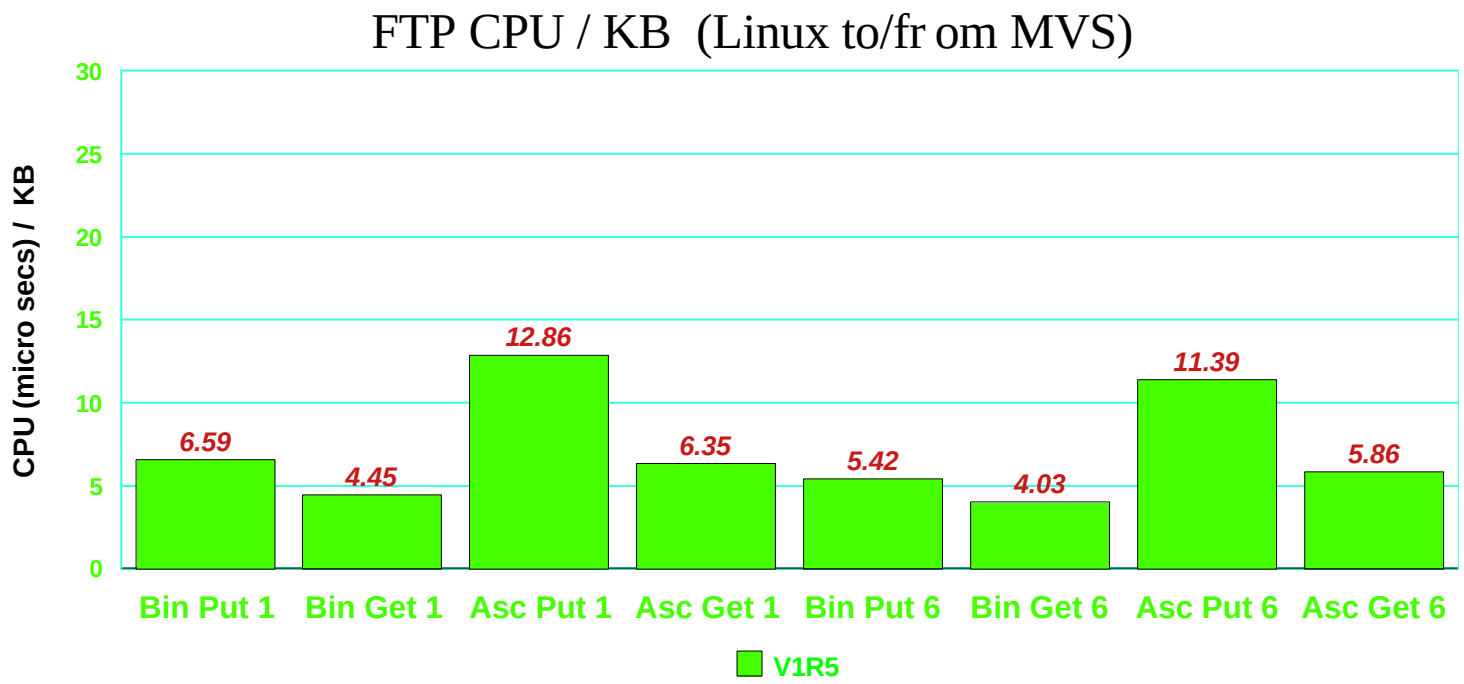

Put: Linux -----> MVS, Get: Linux <------ MVS

Config: Linux WS (1 or 2)----GbE----Switch----OSAE/GbE----2084-B16 (2 CPs)

DASD Config: 2 Banks of Shark DASD (2 FICON Chipids)

FTP Capacity Planning

MVS CPU Requirements:

$$\frac{\text{Max KB}}{\text{Elap secs}} * \frac{\text{CPU secs}}{\text{KB}} = \frac{\text{CPU secs}}{\text{Elap secs}}$$

Example: (59.0 MB/S, WS--> MVS, Bin Put, CS/390 R15, IBM OSAE-GBE)

$$\frac{\text{57617.2 KB}}{\text{Elap secs}} * \frac{\text{.00000542 N1}}{\text{KB}} = \frac{\text{.311 CPU secs}}{\text{Elap secs}}$$

N1: MVS TCP/IP + VTAM + FTP Addr Spaces
(2084-B16 2 CP LPAR)

If the CPU secs/Elap sec ratio is greater than 1, one would need more than one processor (CS/390 R4 - R15).

FTP Capacity Planning con't

MVS CPU Utilization:

$$\frac{\text{CPU secs/Elap Sec}}{\text{\# of processors}} * 100\% = \text{CPU Util \%}$$

\# of processors: Should be equal to the number of number of processors (CS/390 R4 - R15).

Example: (59.0 MB/S, WS--> MVS, Bin Put, CS/390 R15, IBM OSAE-GBE)

$$\frac{0.311 \text{ CPU secs/Elap sec}}{\text{2 processor}} * 100\% = 15.5\%$$

Thus, MVS TCP/IP's + VTAM's + FTP addr spaces CPU requi rement for FTP Binary PUT would require 15.5 % of a two processor 2084-B16 LPAR system. LSPR can be used to adjust for other processors types.

z/OS V1R5 TN3270 Performance Summary

- TN3270

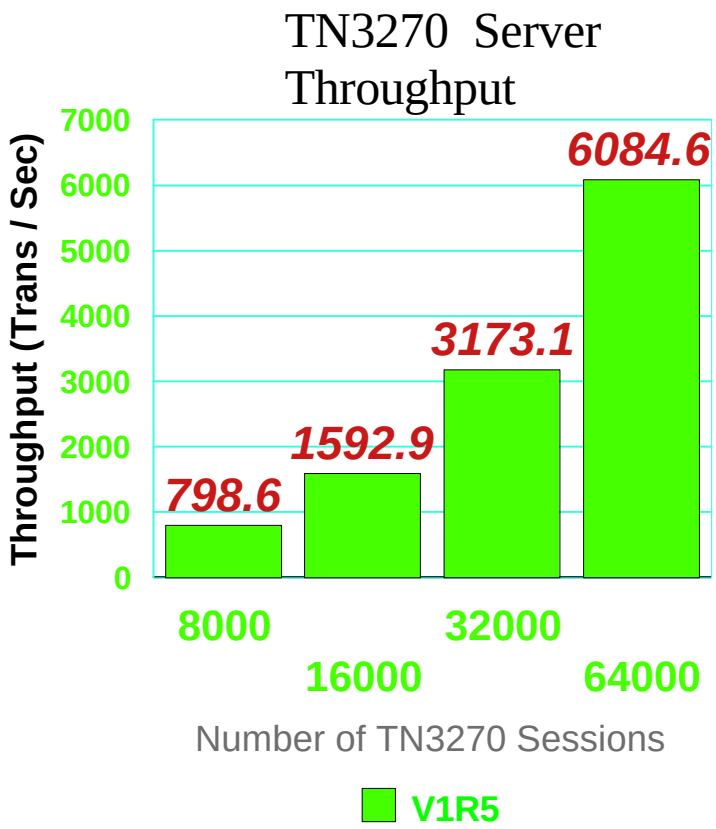

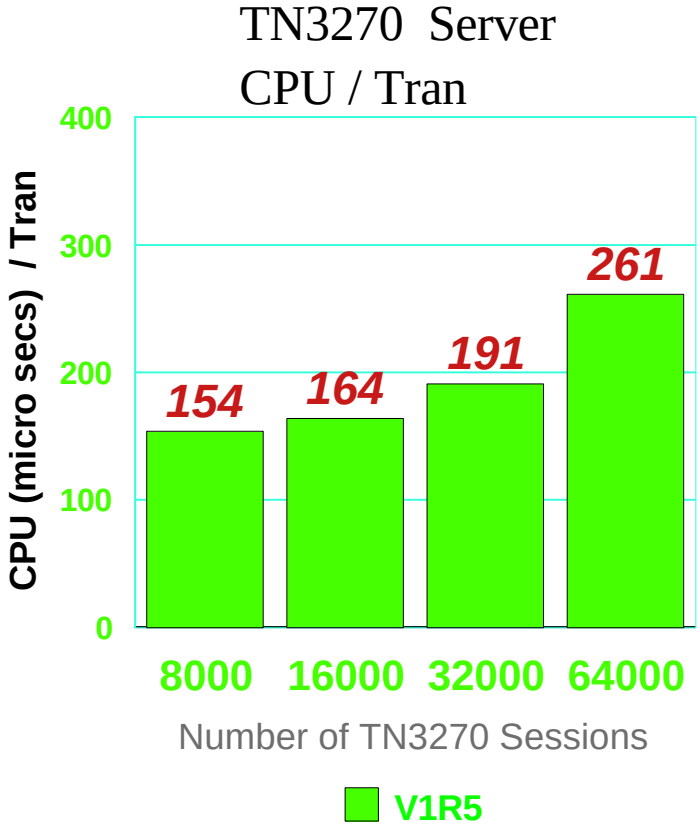

Tran: 100 bytes in / 800 bytes out # clients = 8000 to 64000
Config: 2084-316 (2 CPs, Clients)----OSAE/GbE (1)----2084-316 (2 CPs, Server)
Client: 4 TPNS's simulating TN3270 clients
Server: 4 ITPECHO applications, TN3270 Server
Transaction Rate: 6 / minute / user (10 sec Think Time per user transaction)
CPU : TCP/IP + VTAM + ITPECHO CPU Time

Telnet (TN3270) Storage Utilization (CS/390 V1R5)

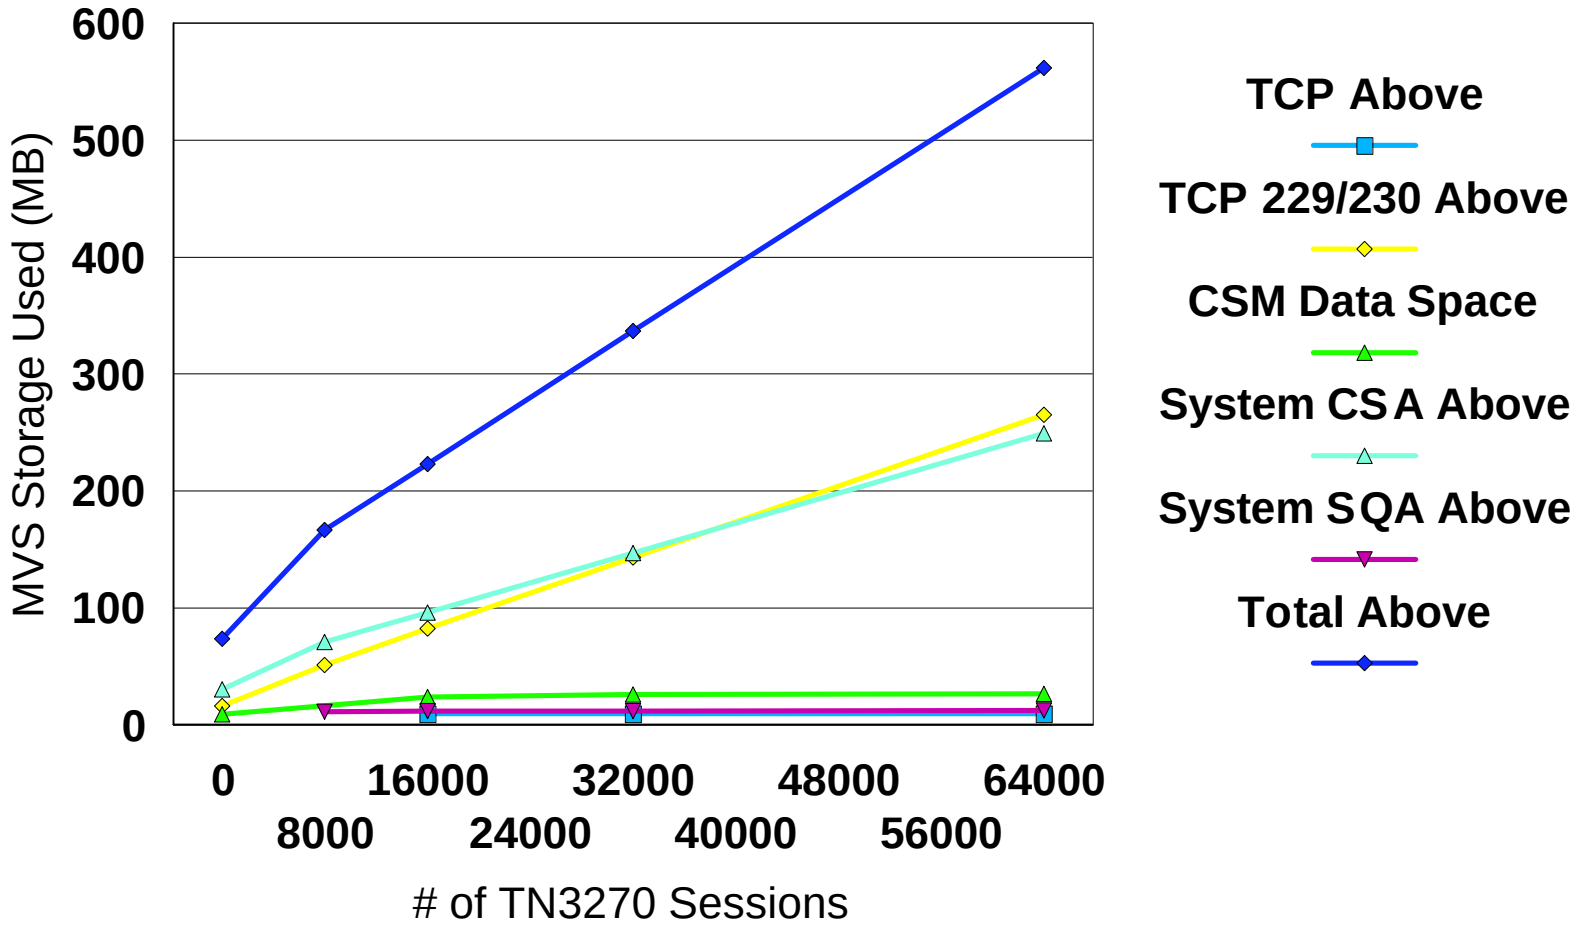

Storage usage (Above 16M line) of TCP/IP Addr Space and MVS System Storage (SQA, CSA) during TN3270 echoes (8000 to 64000 users) when using CS/390 V1R5.
Delta Per User Total: 7.64 to 11.64 KB / user (8000 to 64000 TN3270 sess.)

Telnet (TN3270) Storage Utilization (CS/390 V1R5)

# of TN3270 Sessions	0	8000	16000	32000	64000
TCP/IP Below	0.228 M	0.324 M	0.356 M	0.404 M	0.548 M
TCP/IP Above	8.55 M	9.35 M	9.37 M	9.38 M	9.43 M
TCP/IP LSQA /SWA/**229/230** Below	0.188 M	0.224 M	0.292 M	0.332 M	0.332 M
TCP/IP LSQA /SWA/**229/230** Above	**16.3 M**	**51.0 M**	**82.2 M**	**143 M**	**265 M**
CSM Data Space	8.76 M	23.78 M	23.74 M	26.0 M	26.34 M
System CSA Below	0.320 M	0.320 M	0.320 M	0.320 M	0.320 M
System CSA Above	**30.6 M**	**71.1 M**	**96.1 M**	**147 M**	**249.3 M**
System SQA Below	0.348 M	0.396 M	0.396 M	0.396 M	0.396 M
System SQA Above	9.45 M	11.4 M	11.5 M	11.7 M	12.1 M
Total Below	1.08 M	1.26 M	1.36 M	1.45 M	1.60 M
Total Above	73.66 M	166.63 M	222.9 M	337.08M	561.86M
Total	74.75 M	167.90 M	224.26M	338.54M	563.48M
Delta Per User Total (KB)		11.64 KB	9.34 KB	8.24 KB	7.64 KB

Delta Per User Total: 7.64 to 11.64 KB / user (8000 to 64000 TN3270 sess.)
Storage usage of TCP/IP Addr Space and MVS System Storage (SQA, CSA) during TN3270 echoes (8000 to 64000 users) when using CS/390 V1R5 .

TN3270 Capacity Planning

MVS CPU Requirements:

$$\frac{\text{\# trans/user x \# users x CPU secs/tran}}{\text{\# of Elap secs}} = \frac{\text{CPU secs}}{\text{Elap secs}}$$

Example: CS/390 R15, 8000 users, 6 tr/min/user

$$\frac{\text{6 tr/u x 8000u x 0.000154 CPU secs/tr}}{\text{60 elap. sec}} = 0.123 \ \frac{\text{cpu sec}}{\text{elap sec}}$$

N1: MVS TCP/IP + VTAM + ECHO Application CPU
(2084-B16 2 CP LPAR)

If the CPU secs/Elap sec ratio is greater than 1, one would need more than one processor (CS/390 R4-R15).

TN3270 Capacity Planning con't

MVS CPU Utilization:

$$\frac{\text{CPU secs/Elap Sec}}{\text{\# of processors}} * 100\ \% = \text{CPU Util \%}$$

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

Example: CS/390 R15, 8000 users, 6 tr/min/user

$$\frac{0.123\ \text{CPU secs/Elap sec}}{2\ \text{processors}} * 100\ \% = 6.16\ \%$$

Thus, the MVS TCP/IP + VTAM + Echo Application CPU requirement for 8000 TN3270 users would require 6.16 % of a two processor 2084-B16 LPAR system. LSPR can be used to adjust for other processors types.

MVS Telnet (TN3270) Tuning Parms (V3R1/V3R2/CS390 R5-R15)

- SCANINTERVAL:
  - Used to override the default scan time (120 secs). This time specifies the periodic time that the Telnet server would scan the entire list of TCP/IP connections.
- Timemark:
  - Used to specify how often the Telnet server will send an "are you there" probe to clients that appear to be inactive.
  - Clients who receive three consecutive probes without intervening activity are considered to be inactive.
- INACTIVE:
  - Used to specify how long a terminal can remain unused (no communication with the Telnet server) before it will be deemed inactive and disconnected by the server.
- DISABLESGA:
  - Permits the transmission of GO AHEAD by Telnet. Negotiated by both client and server. Using DISABLESGA increases the overhead for a full duplex terminal using a full duplex connection. Applies only to Linemode, not 3270 connections.
    Default is to suppress transmission of GO AHEAD.
    Recommendation is to use the default (Do not specify DISABLESGA).

MVS CICS Sockets Tuning (V3R2/CS390 R5-R15)

- MVS dispatching priority of VTAM, TCP/IP, and other servers:
  - ►Recommendation: Set following to SYSSTC (WLM).
    VTAM
    TCP/IP
    Routing Deamons
    Other Daemons (FTPD, INETD, etc.)
    Other TCP/IP Servers/Applications

- SOMAXCONN: Maximum # of queued connections on a listening port.
  - ► Recommendation: Set to a large value (ie. 1024 or >)
    in MVS TCP/IP Profile. Default = 10.

- Registration / Deregistration with WLM:
  - ► Modifications to the Listener configuration to allow for up to three group names.
  - ► A listener can be defined to belong to one of the group names. These group names are used to register the CICS listener with the Workload Manager (WLM) so that a BIND-based Domain Name Service (DNS) can be used to balance requests across multiple hosts in a sysplex environment.

- Refer to OS/390 V2R8.0-V2R10.0 SecureWay Communications Server IP CICS Sockets Guide (SC31-8518-01) for more information.

MVS CICS Sockets Tuning Parms (V3R2/CS390 R5-R15)

- CICS Sockets Configuration Parms :
  - EZAC CICS Transaction for CICS:
    - TCPADDR : Set to TCP/IP proc name.
    - NTASKS : Set slightly higher than the max number of concurrent CICS sockets connections. Default = 20.
      NTASKS defines a pool of reuseable subtasks.
      Attached tasks are used for listeners and when the pool of reusable tasks is exhausted.
    - DPRTY : Set to 0 to improve response time for CICS Sockets.
      The difference between the dispatching priority of the subtasks and the attaching CICS task.
  - EZAC CICS Transaction for CICS Listener(s):
    - NUMSOCK : Set slightly higher than NTASKS.
      One less than this number is the maximum number of concurrent GIVESOCKET requests that can be active.
      Default value is 50.
    - BACKLOG : Set to 40. Default = 20.
      The number of unaccepted connections that can be queued to this listener.
    - WLMGN1,2,3 : cicssocgr1 / cicssocgr2 / cicssocgr3

z/OS V1R5 CICS Sockets Performance Summary

- CICS Sockets :

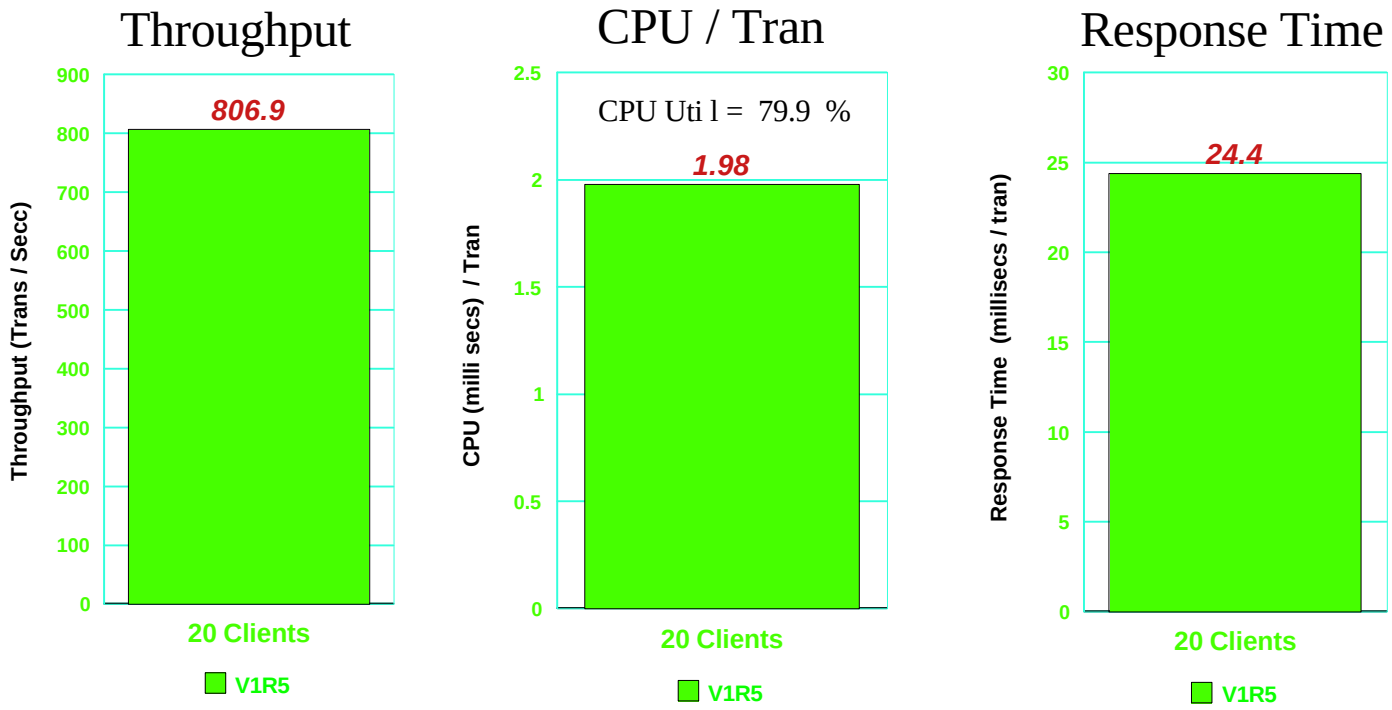

Tran Size : 200 bytes in / 200 bytes out # client sessions = 20

1 CICS Sockets listener, EZACICSC child server.

Client transaction communicates to CICS Sockets listener and then the EZACICSC child server. There is no databas e access. Database access would have additi onal impact on TPUT and CPU.

Config: Linux WS (2)----GbE----Switch----OSAE/GbE----2064-216 (2 CPs)

Enterprise Extender Test Results - Effect of RU Size

- Stream throughput was highest for 8K RU size

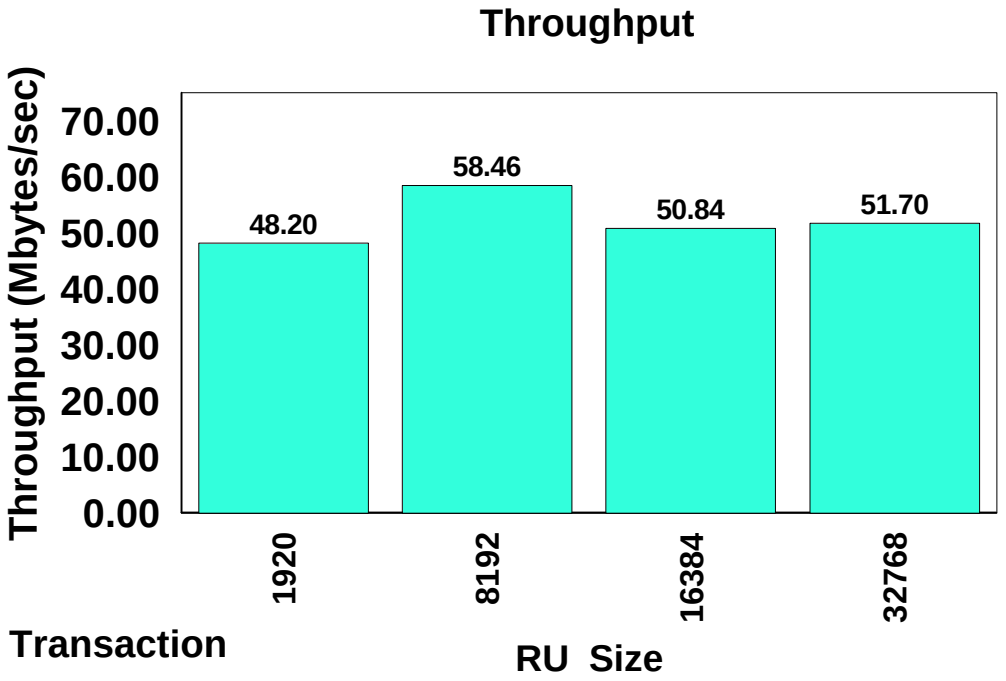

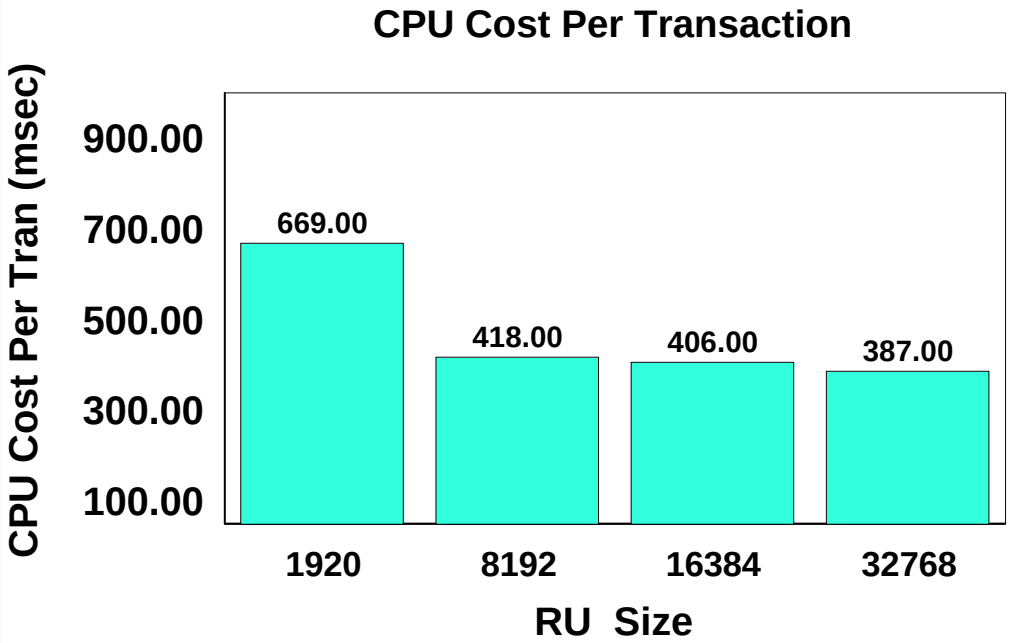

- CPU cost per transaction reduced by 37.5% when RU size increased from 1920 to 8192

Enterprise Extender Performance Recommendations

- Based on analysis of z/OS V1R2 CS Enterprise Extender environment with tests of interactive and stream workloads using 1500 byte MTU size
- Recommended Maintenance:
  - z/OS CS
    - APAR OA02213 (will prereq other VTAM maintenance)
    - APAR PQ69398
  - OSA Microcode
    - G5/G6: 4.28
    - z/Series 2064 GA3: 3.26 or later
  - Monitor EE Info APAR (II11223) for additional information
- When choosing RU size, carefully consider application workload:
  - For streaming workloads, RU size of 8192 seems to provide an optimum balance of improved throughput and reduced CPU cost per transaction
  - As always, results may vary due to network and application specifics

IBM MVS TCP/IP Performance *Checklist*

- MVS dispatching priority of VTAM, TCP/IP, and other servers:
  - Recommendation: Set following to SYSSTC (WLM).
    VTAM
    TCP/IP
    Routing Deamons
    Other Daemons (FTPD, INETD, etc.)
    Other TCP/IP Servers/Applications
- Make sure client and server TCP Window size are equal
  - Recommendation: On MVS, use default window size of 32768 or 65535.
    On Client, set client window size to 32768 or 65535 (if allowed).
- Make sure client and server MTU/packet size are equal
  - Recommendation: For Ethernet lans use 1500,
    Token Ring lans use 1500 or 2000,
    FDDI lans use 4000 or 4352,
    CTC use 65527.

IBM MVS TCP/IP Performance *Checklist*

- Routers: Make sure buffers are set appropriately so that packets are not being dropped.
- 3172: Make sure Delay timer and Max. response length are set correctly for each Lan adapter
  - Recommendation: Delay Timer = 10 ms, Max. response length = 500 bytes
- 2216: Make sure Blk timer and Ack length are set correctly for each LCS or MPC+ definition.
  - Recommendation: Blk Timer = 5 ms, Ack Length = 10 bytes
- RS/6000 ESCON Attachment:
  - Recommendation: Use MPC (instead of CLAW) as Subchannel Type for improved performance.

  By using MPC, FTP throughput improved 61 % (outbound) or 92 % (inbound).

IBM MVS TCP/IP Performance *Checklist*

- FTP: BUFNO, EXTRATASKS and NCP parms are no longer used in CS/390 V2R5 and later releases.
- FTP: Use large dataset blocksizes on MVS
  - Recommendation: DS Blocksize < = 1/2 DASD track (3380: approx. 23424 byt, 3390/9334: approx. 28288 byt)
- FTP: Set Client TCP/IP send/receive socket buffers to at least 64 KB.

- TELNET: Check TIMEMARK, SCANINTERVAL, INACTIVE parms in MVS TCP/IP Profile (INTERNALCLIENTPARMS)
  - Recommendation: TIMEMARK = 10800 (3 hrs)
    SCANINTERVAL = 1800 (30 min)
    INACTIVE = 5400 (90 min)

- SOMAXCONN: Maximum # of queued connections on a listening port.
  - Recommendation: Set to a large value (ie. 1024 or >) in MVS TCP/IP Profile.
    Default = 10.

IBM MVS TCP/IP Performance *Checklist*

- Sockets: Use large msg sizes (> 1 KB) for better performance.

- Gateway Statement (MVS TCP/IP Profile):
  Use a numeric value (ex. 1500) for MVS Packet Size.
  (Do not use DEFAULTSIZE for Packet Size).

- PTF's: Make sure have latest CS/390 (TCP/IP & VTAM) Performance PTF's.

- Traces: Make sure TCP/IP and all other traces are turned off for optimal performance.

  To turn z/OS TCP/IP trace off:
  TRACE CT,OFF,COMP=SYSTCPIP,SUB=(TCPCS12)

  To display z/OS TCP/IP trace status:
  D TRACE,COMP=SYSTCPIP,SUB=(TCPCS12)

Summary

- CS/390 R15
  - IPv6 support for TN3270(E), CICS Sockets, MPC+
  - TN3270E Definite response performance improvement
  - Async I/O performance improvement
  - EE enhancements (Multiple VRNs, MPC+ HPDT Packing)
  - DVIPA limit now 1024
  - New Crypto instructions
  - OSA enhancements (Checksum offload)

- CS/390 R14
  - IPv6
  - DNS bind 9.2
  - OSA Express / HiperSockets Storage Enhancements
  - OSA Express Microcode Enhancement
  - SNA Enhancements

- CS/390 R12
  - Hipersockets: Fast LPAR to LPAR communication (same CEC)
  - Hipersockets Accelerator: Single stack acts as router for entire CEC
  - Full Claw Packing: Up to 60KB claw packing buffers
  - TCP/IP Storage Limits (ECSA & POOLLIMIT)

- Follow MVS TCP/IP Performance Checklist for optimal performance.

Appendix

CS/390 Release Info

- Comm. Server for OS/390 (CS/390):
  Integral part of OS/390 or z/OS.

  - OS/390 R3 , CS/390 R3 (3/97): TCP/IP V3R2
  - OS/390 R4 , CS/390 R4 (9/97): TCP/IP V3R2
  - & R4 (Stack Rewrite)
  - OS/390 R5 , CS/390 R5 (3/98): R5 (Appl Support)
  - OS/390 R6 , CS/390 R6 (9/98): R6
  - OS/390 R7 , CS/390 R7 (3/99): R7
  - OS/390 R8 , CS/390 R8 (9/99): R8
  - OS/390 R9 , CS/390 R8 (3/00): R8
  - OS/390 R10, CS/390 R10 (9/00): R10
  - z/OS V1R1 , CS/390 R10 (3/01): R10
  - z/OS V1R2 , CS/390 R12 (10/01): R12
  - z/OS V1R3 , CS/390 R12 (03/02): R12
  - z/OS V1R4 , CS/390 R14 (10/02): R14 (IPv6)
  - z/OS V1R5 , CS/390 R15 (03/04): R15

Enterprise Class SNA and TCP/IP

Comm. Server for OS/390

Integrated Services

- Comm. Server for OS/390 provides:
  - Reliability of OS/390 system
  - High Availability
  - Capacity
  - Scalability
- Provide common services within the S/390 Communications Server
  - Storage Management
  - Device Attachment (with Data Link Controls)
- TCP/IP and SNA integration (e.g. TN3270)

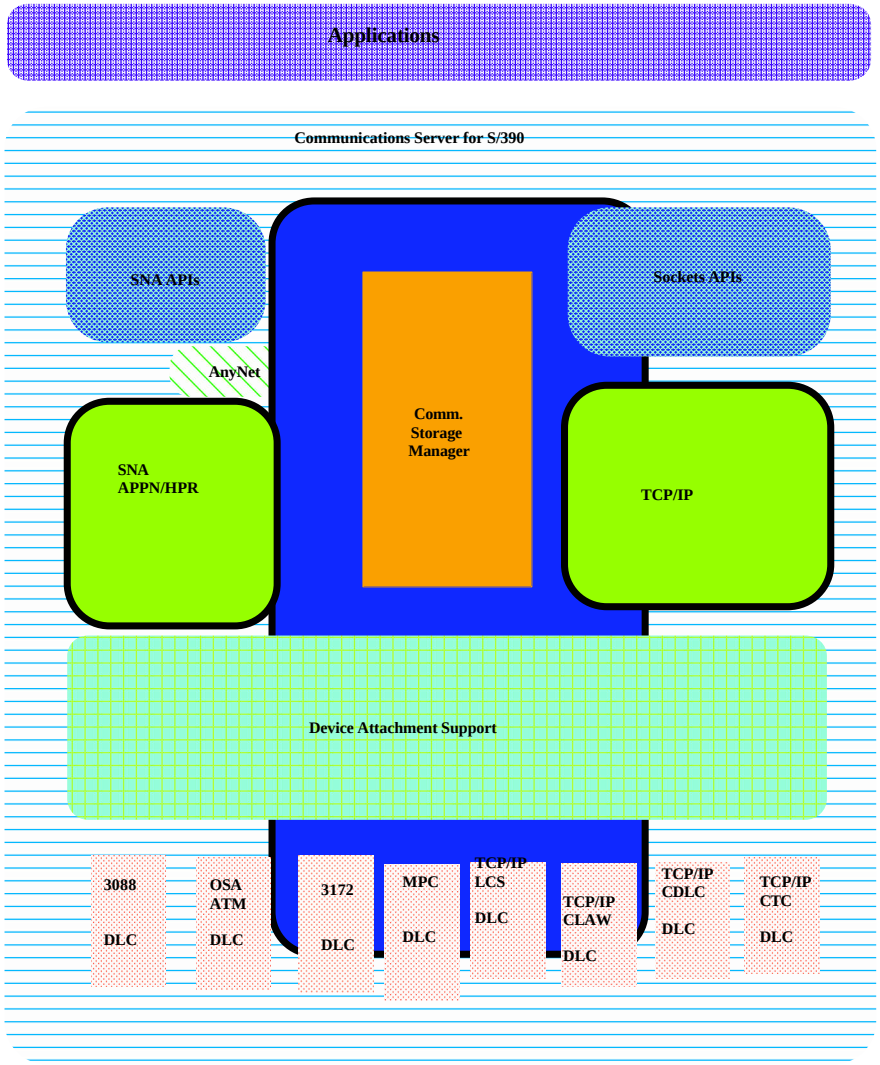

IBM MVS TCP/IP Performance Comparison (CS/390 R12 & TCP/IP V3R2)

CS/390 R12 vs TCP/IP V3R2+ :

Application	TPUT	CPU	TPUT / CPU
Web Server	up to 8.2 X	- 93 %	up to 117 X
CICS Sockets	up to 1.95 X	- 46 %	up to 3.61 X
Telnet(TN3270)	up to +2 % (note 1)	- 48 %	up to 1.96 X
FTP Server	up to 1.27 X	- 48 %	up to 2.46 X
FTP Client	up to 1.31 X	- 63 %	up to 3.52 X

Note 1: Telnet transaction generator is time driven (with a think time between user transactions) so that transactions per second are approx. equal for all releases.

IBM MVS TCP/IP CPU Performance Ratio Summary: (V3R2 GA, V3R2+, CS/390 R5, R6, R7, R8, R10 & R12)

MVS CPU Ratio's:

Application	TCP/IP V3R2 GA	TCP/IP V3R2+	CS/390 R5 +	CS/390 R6 +	CS/390 R7+	CS/390 R8 (GA)	CS/390 R10 (GA)	CS/390 R12 (GA)
Web Server	-----	1.00	0.30	0.26	0.07 (FRCA ON	0.07 (FRCA ON)	0.07 (FRCA ON)	0.07 (FRCA ON)
CICS Sockets	-----	1.00	0.61	0.58	0.57	0.55	0.54	0.54
Telnet (TN3270)	1.00	0.71	0.48	0.44	0.40 LCS	0.39	0.37	0.37
FTP Server	1.00	0.87	0.58	0.52	0.48 LCS	0.47 LCS	0.45	0.45
FTP Client	1.00	0.77	0.31	0.30	0.29 LCS	0.28 LCS	0.28	0.28

+ : means GA code + PTF 's

z/OS V1R2 CLAW Packing Throughput Performance Summary

- CLAW Packing
  - Pack multiple datagrams into a single claw channel frame
  - V1R2 FTP throughput increases from 31 to 66 % vs R10
  - Recommended for all customers using Cisco 7200/7500 routers using CLAW protocol

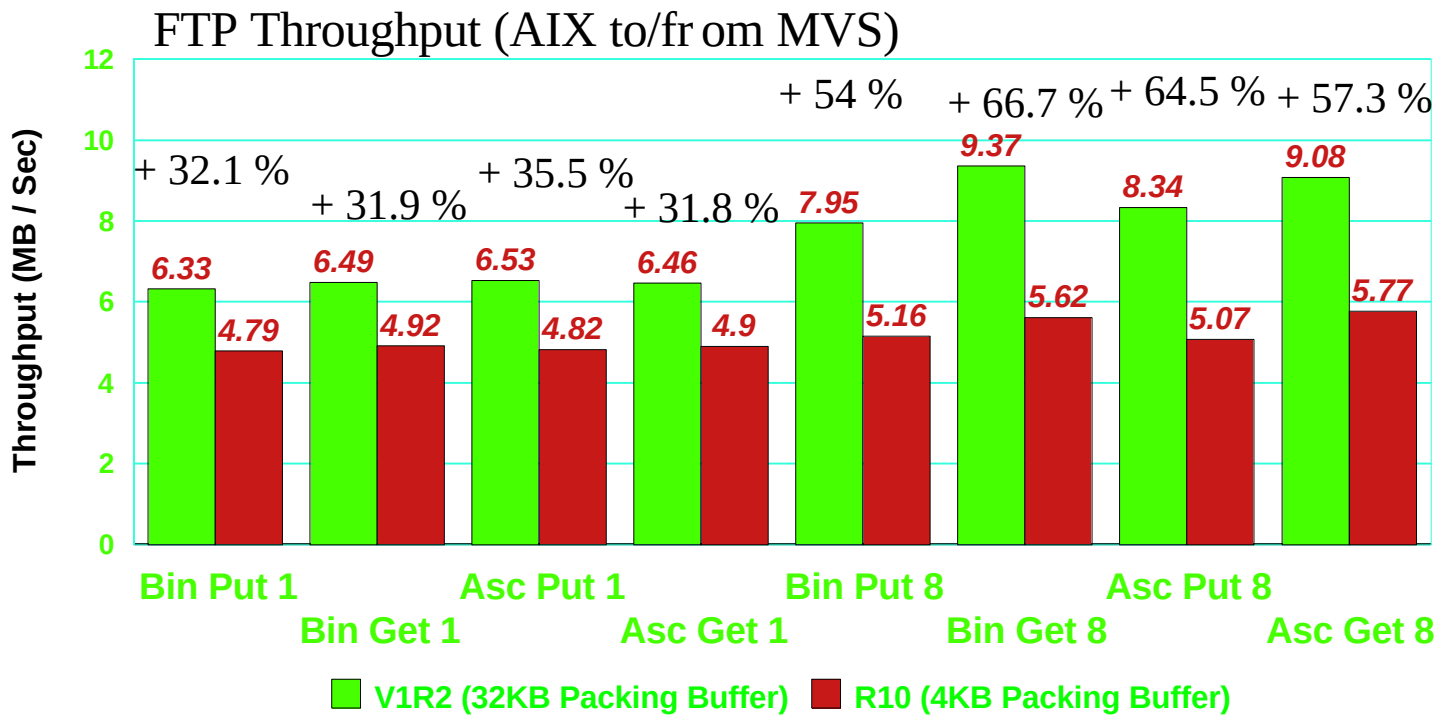

Put: AIX -----> MVS, Get: AIX <------ MVS
Config: AIX WS (1 or 4)---FE---Switch--FE--Cisco Rtr 7505---ESCON---9672-RX6 (4 CPs)
DASD Config: 1 Banks of Shark DASD (4 Chipids)

z/OS V1R2 CLAW Packing
CPU Performance Summary

- CLAW Packing
  - Pack multiple datagrams into a single claw channel frame
  - V1R2 FTP CPU/KB is reduced from 3 to 24 % vs R10
  - Recommended for all customers using Cisco 7200/7500 routers using CLAW protocol

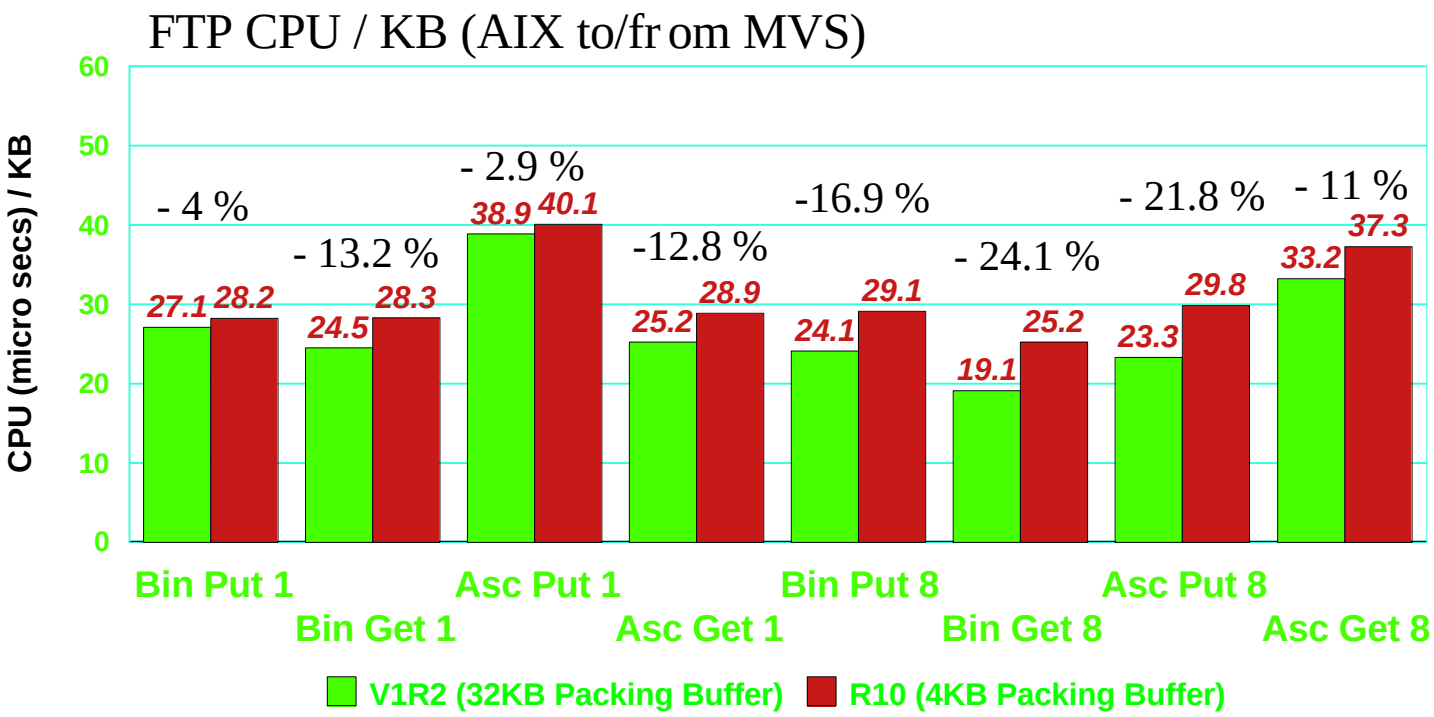

Put: AIX -----> MVS, Get: AIX <------ MVS
Config: AIX WS (1 or 4)---FE---Switch--FE--Cisco Rtr 7505---ESCON---9672-RX6 (4 CPs)
DASD Config: 1 Banks of Shark DASD (4 Chipids)

MVS FRCA Configuration Parms

- Following are the important configuration parms for Fast Response Cache Accelerator (FRCA):
- Parms are specified in the Webserver configuration file (/etc/httpd.conf).

```
EnableFRCA on                          Enables FRCA
FRCACacheSize 8192                     # of 4K Blocks
FRCACacheEntries 1024                  max # of files to be cached
FRCAMaxFileSize 1000000000             max file size (bytes)
FRCAStackName TCPCS12                    TCP/IP Proc name
FRCAWLMParms FRCAHTTP WEBFRCA WEBFRCA
     where   FRCAHTTP  =   subsystem name  (required)
             WEBFRCA   =   application environment name
             WEBFRCA   =   transaction class
FRCACacheOnly   *.gif
FRCACacheOnly   cacheable/*.html
    or   (mutually exclusive)
FRCANoCaching   dontcache/*.html
FRCANoCaching   meeither/*.html
```

z/OS V1R2 Web Serving Performance Summary

- Web Serving :

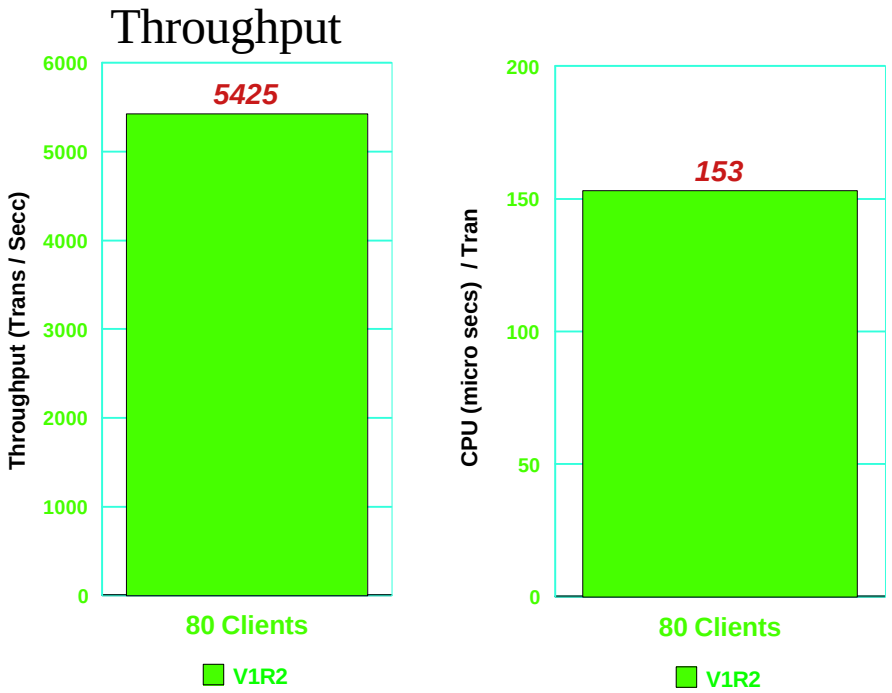

Tran Size (Avg): 73 bytes in / 7.6 KB out # clients = 80
File Sizes: 1 KB to 200 KB
Config: AIX WS (4)----GbE----Switch----OSAE/GbE----2064-108 (4 CPs)
Used FRCA (Fast Response Cache Accelerator), approx. 100% cache hits.