VM Performance Update

Bill Bitner
IBM Endicott
607-752-6022
bitnerb@us.ibm.com
Last Updated: June 21, 2001

Legal Stuff

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly.

Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country or not yet announced by IBM. Such references or information should not be construed to mean that IBM intends to announce such IBM products, programming, or services.

Should the speaker start getting too silly, IBM will deny any knowledge of his association with the corporation.

Trademarks

The following are trademarks of the IBM Corporation:

IBM, VM/ESA, z/VM

LINUX is a registered trademark of Linus Torvalds

Penguin (Tux) complements of Larry Ewing

Overview

- Regression performance
- Large Storage Considerations
  - MDC
  - Greater than 2GB
  - V=R Area
- TCP/IP - QDIO - Gigabit Ethernet
- Monitor & Performance Products

z/VM 3.1.0 Background

- CP will run in 31-bit or 64-bit
- Much of the code is still common
- RIO370 dropped
- V=R area still must reside below 2GB
- Storage above 2GB used for DPA & MDC

VSE Guest V=R Regression

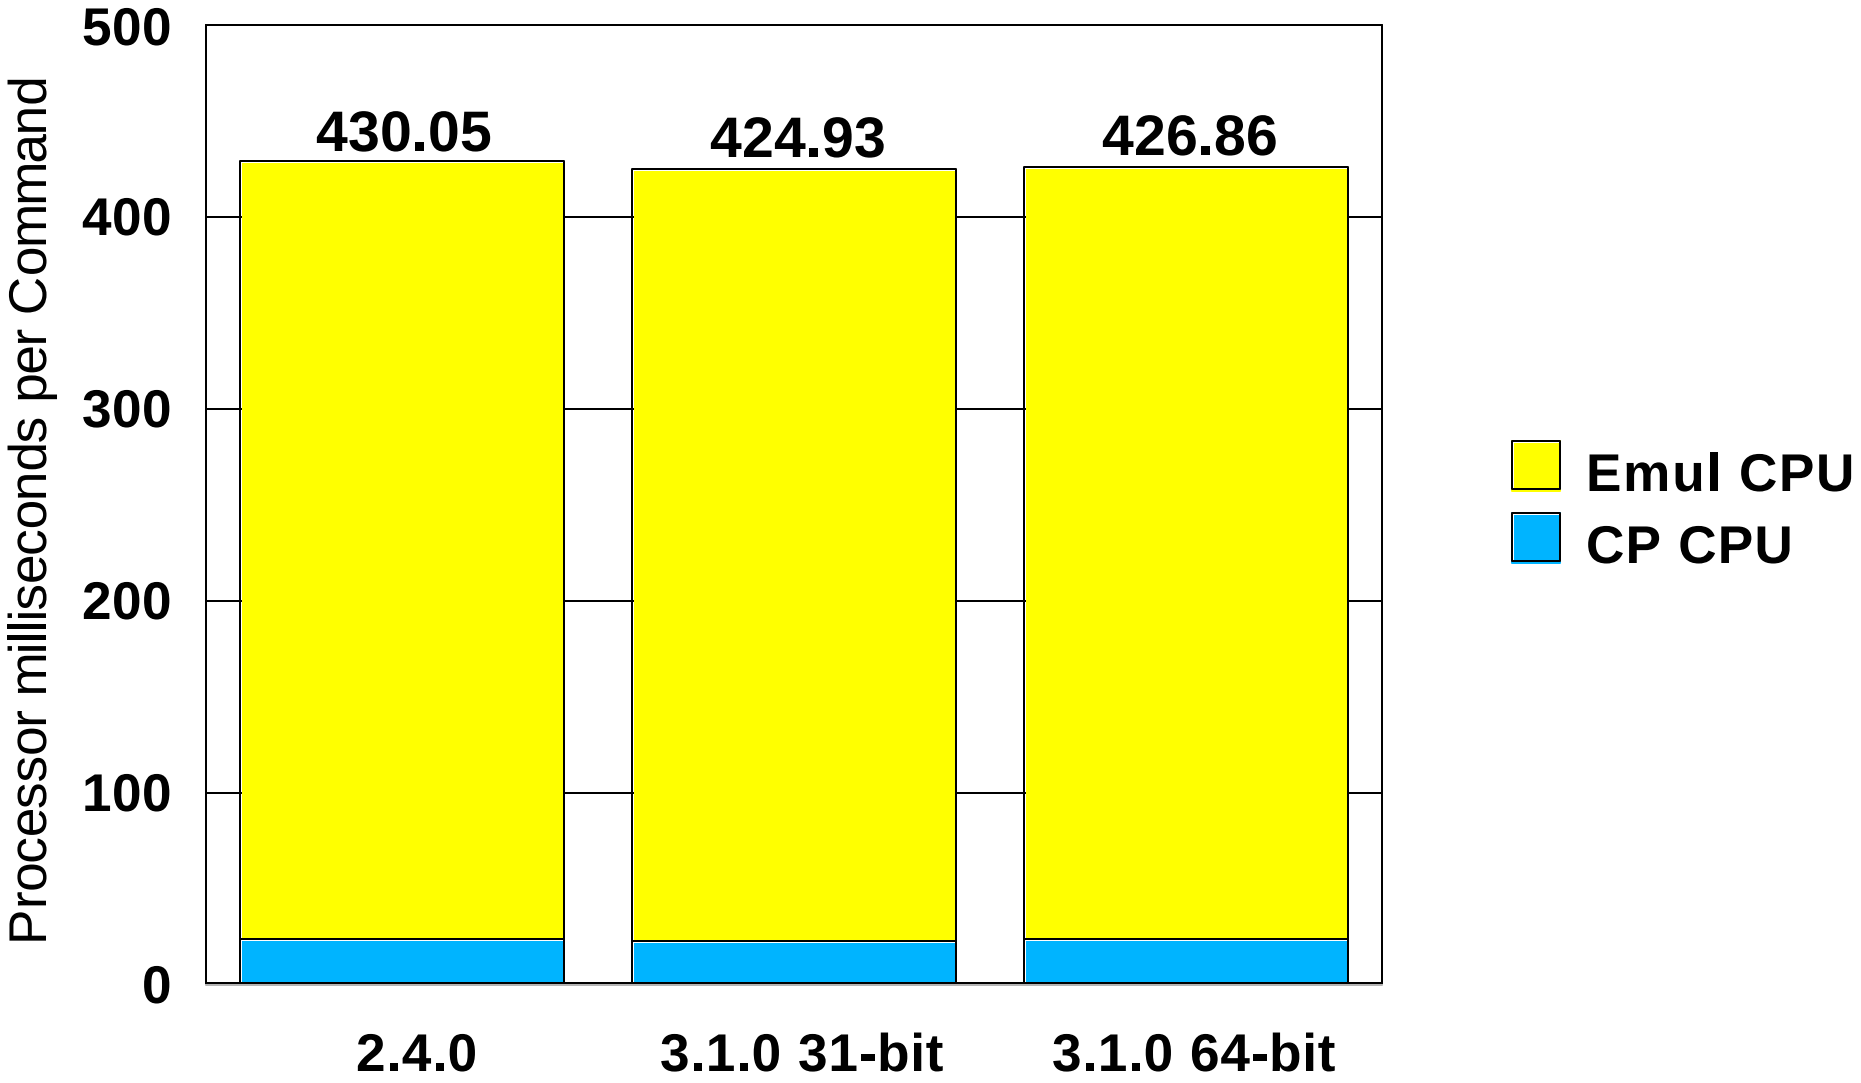

2064-109 2-proceesors online; 2G/2G; V=R VSE Dynapace

VSE Guest V=V Regression

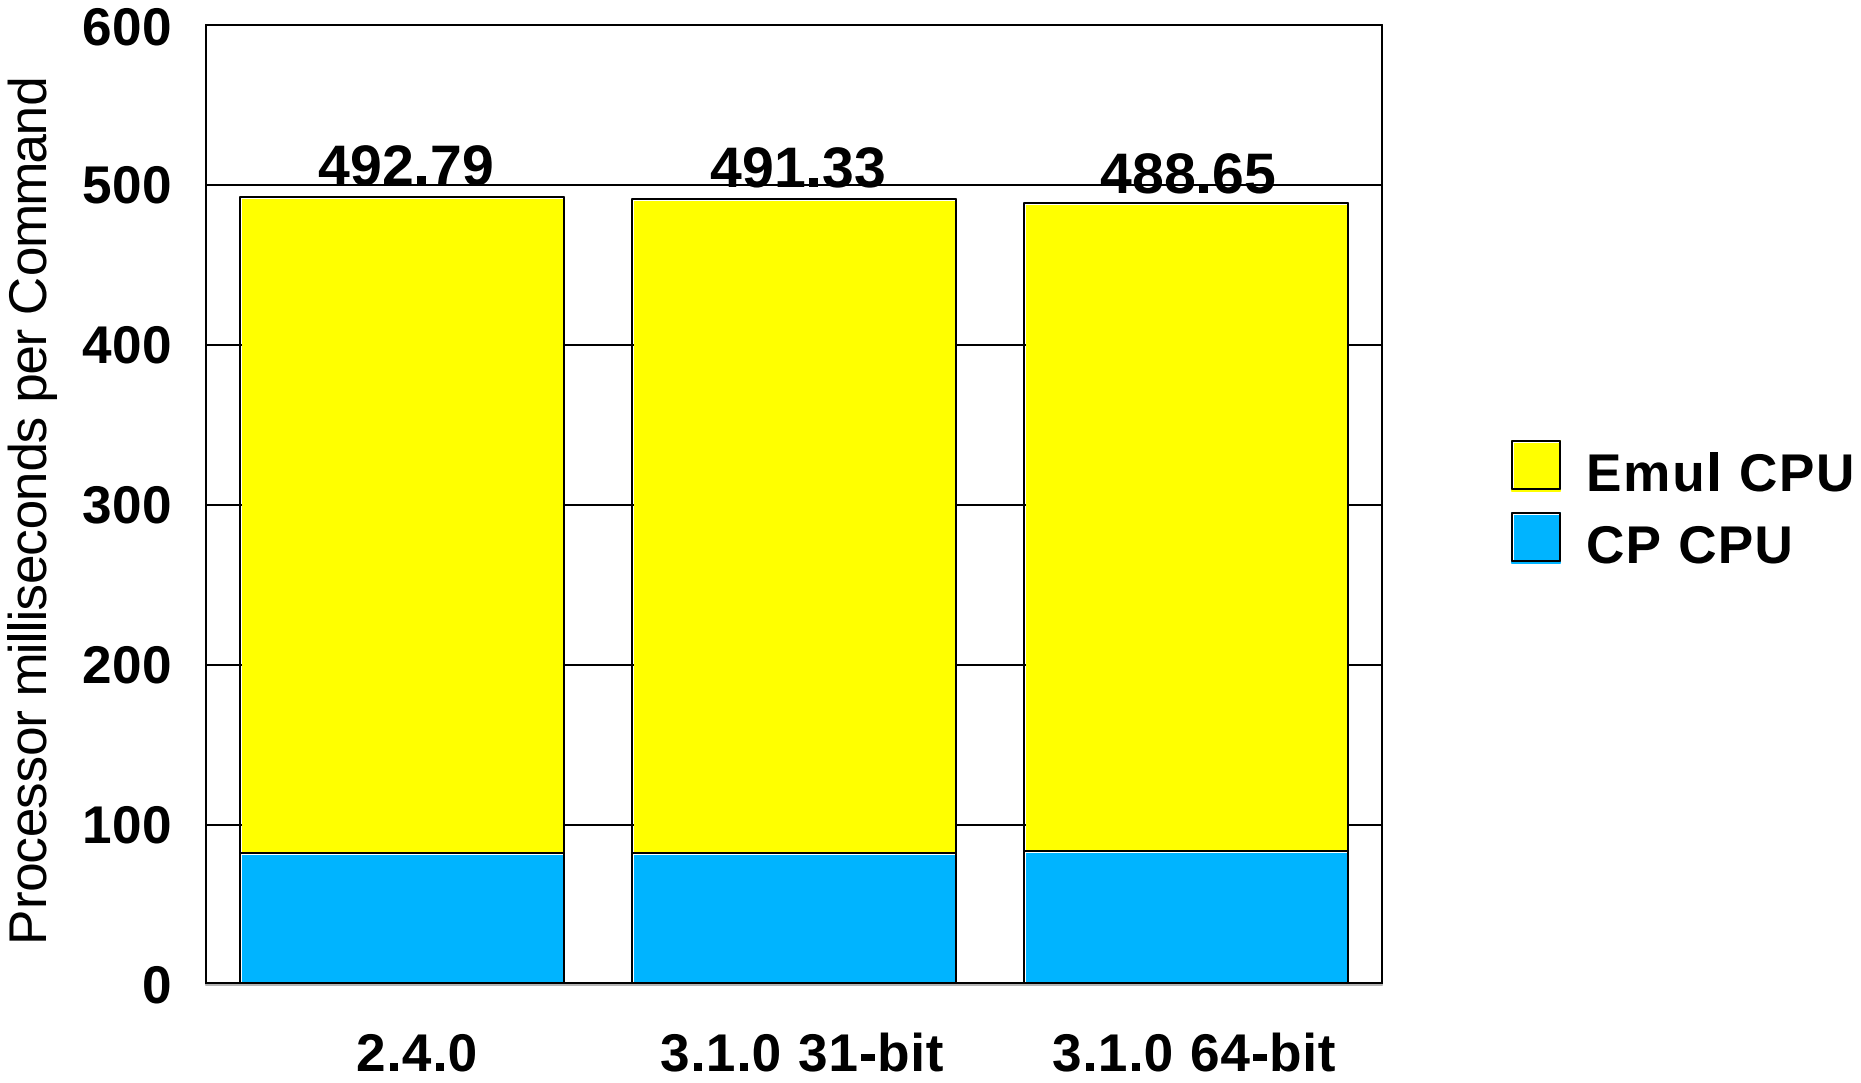

2064-109 2-proceesors online; 2G/2G; V=V VSE Dynapace

CMS Regression

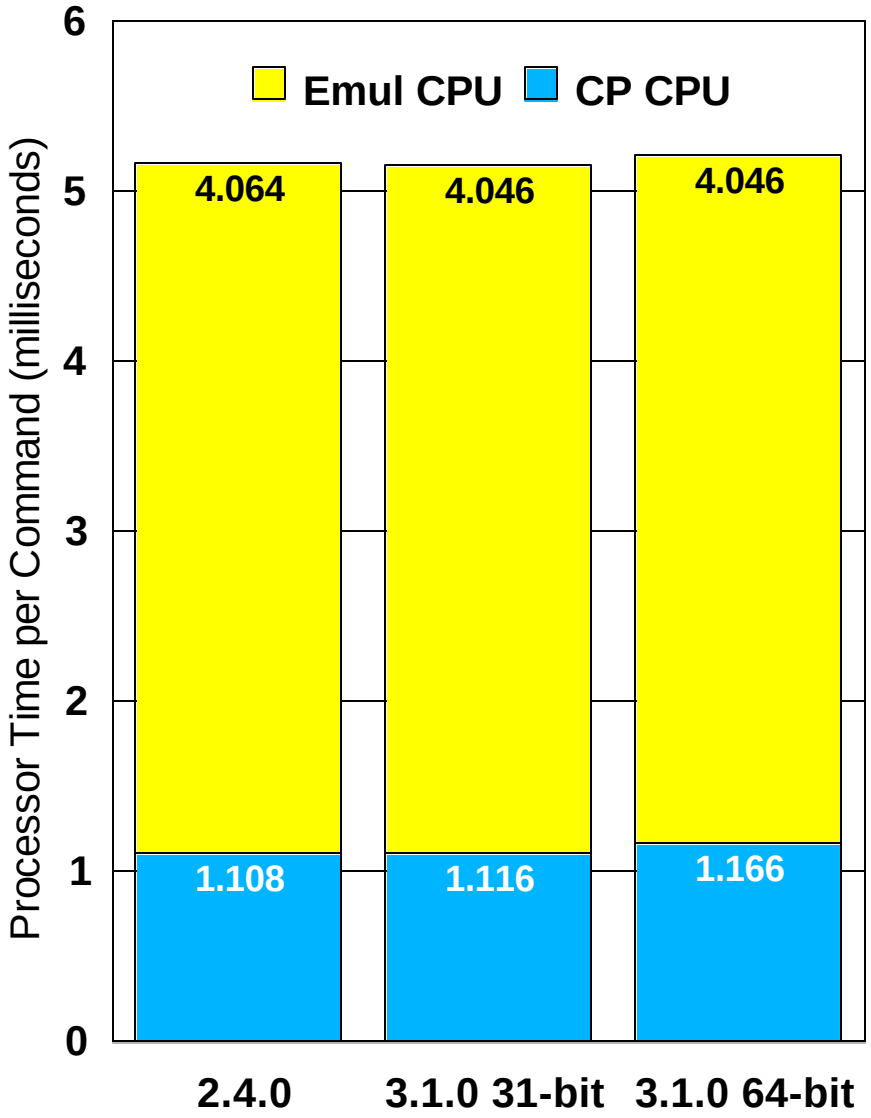

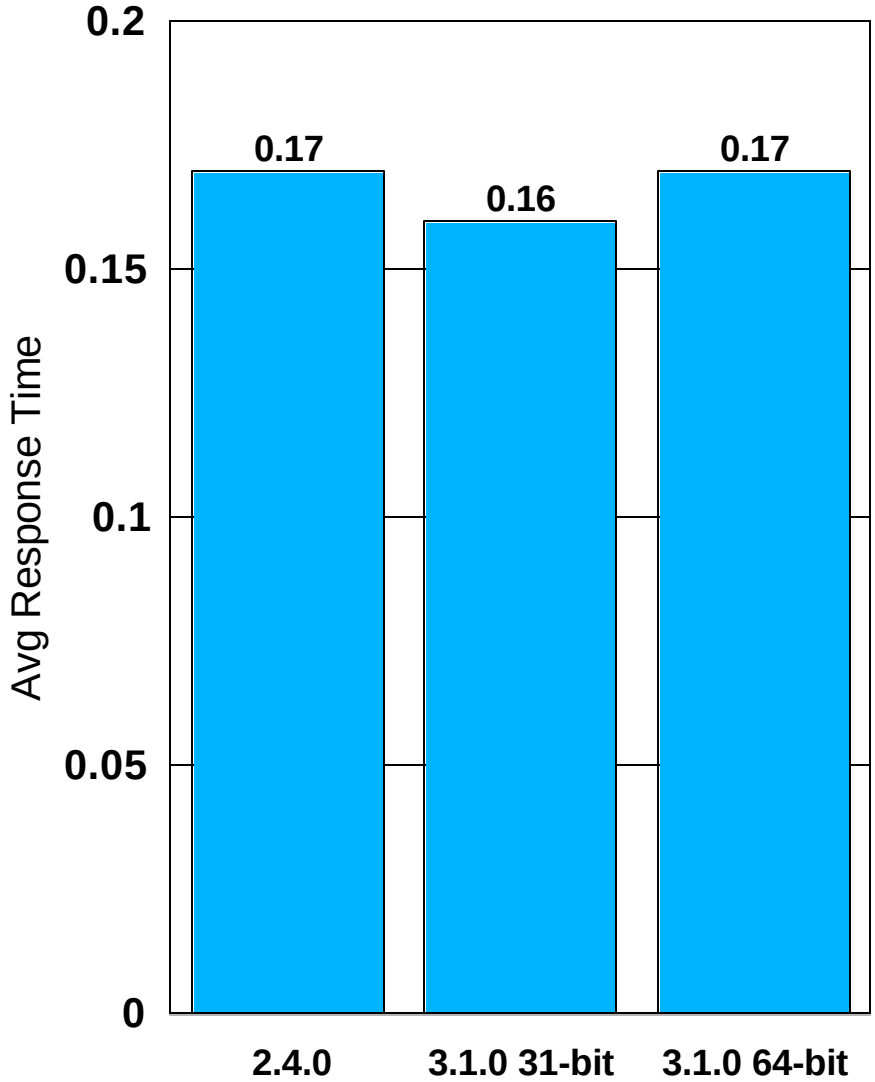

2064-109 LPAR 2-way; 1G/2G; CMS1 External TPNS

CMS Regression

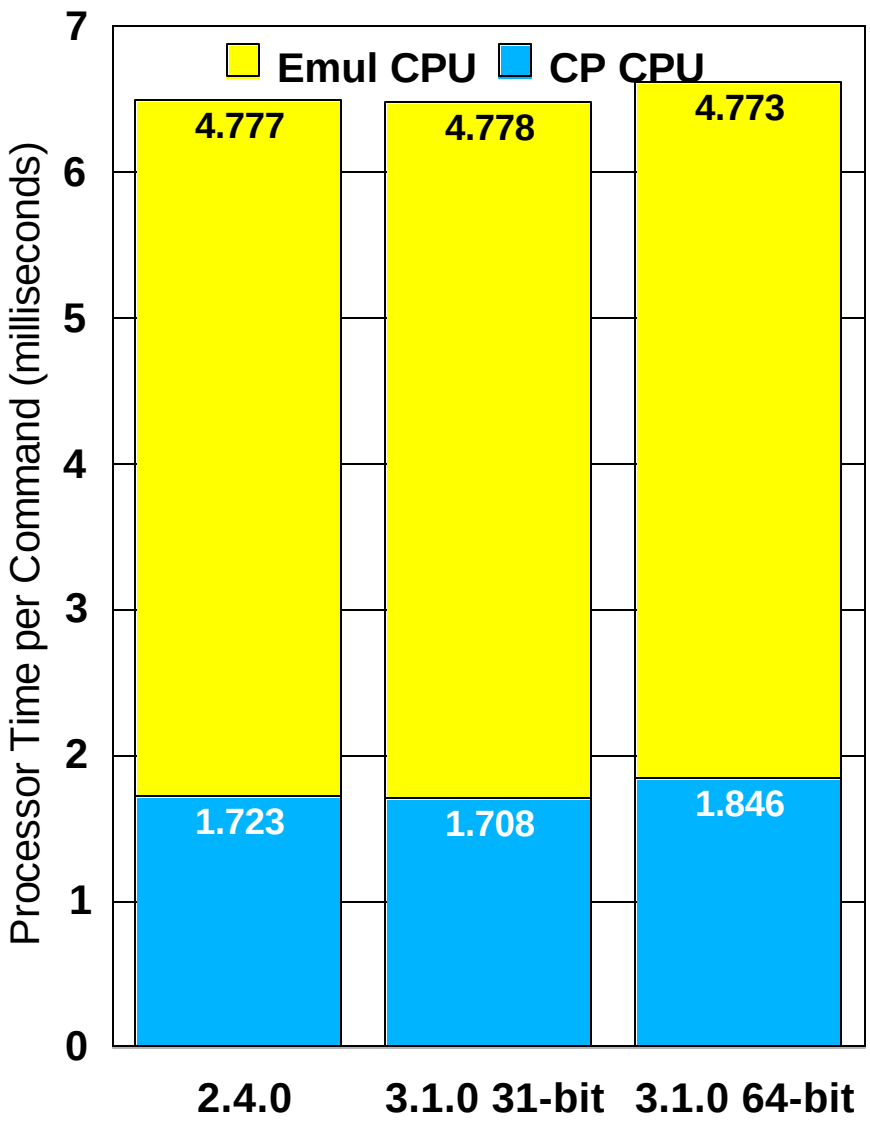

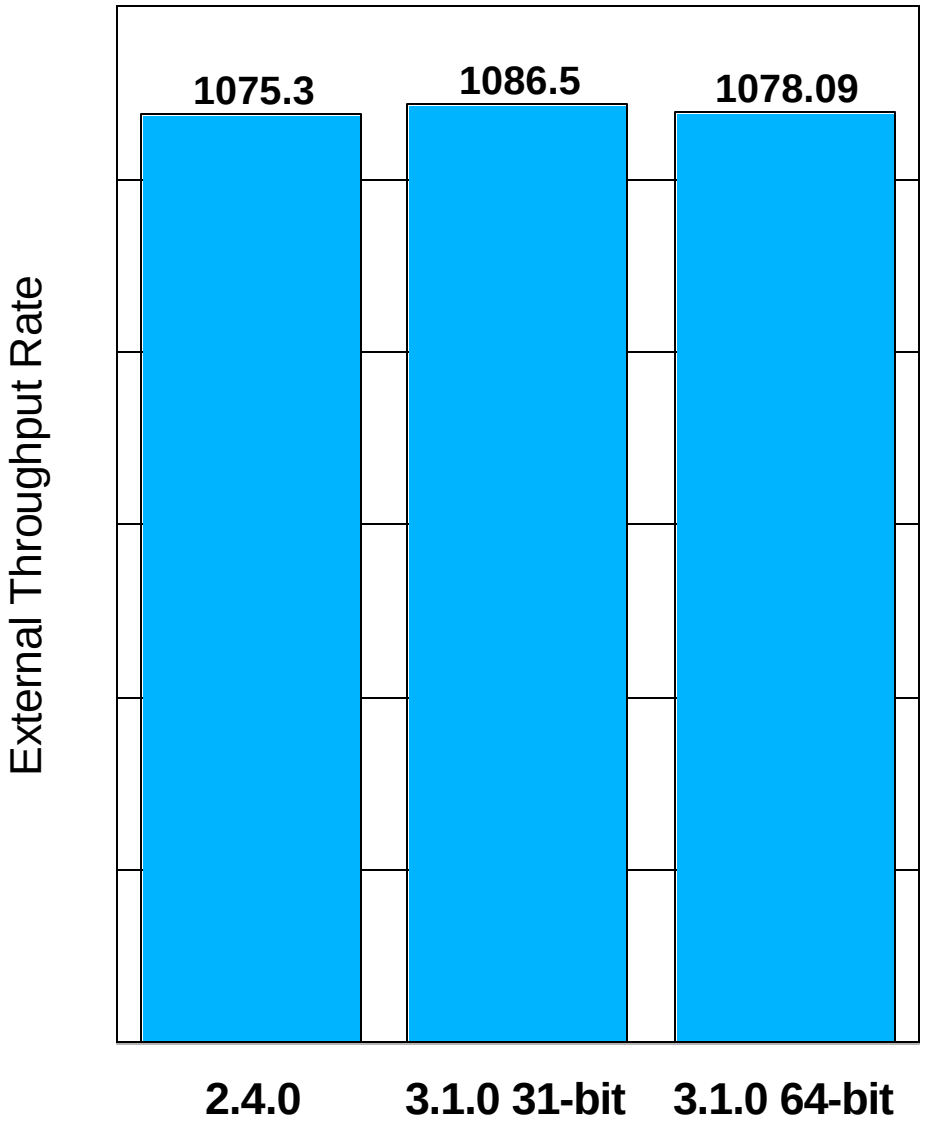

2064-1C8 8-way; 2G/6G; CMS1 Internal TPNS

Storage Allocation Considerations

- Can now run VM with greater than 2GB of real storage on 2064 processors.
  - Should there be any expanded storage?
  - How should storage be used for MDC?
    - Real only?
    - Expanded only?
  - How much can I use for the V=R area?

Storage Allocation - 8 GB

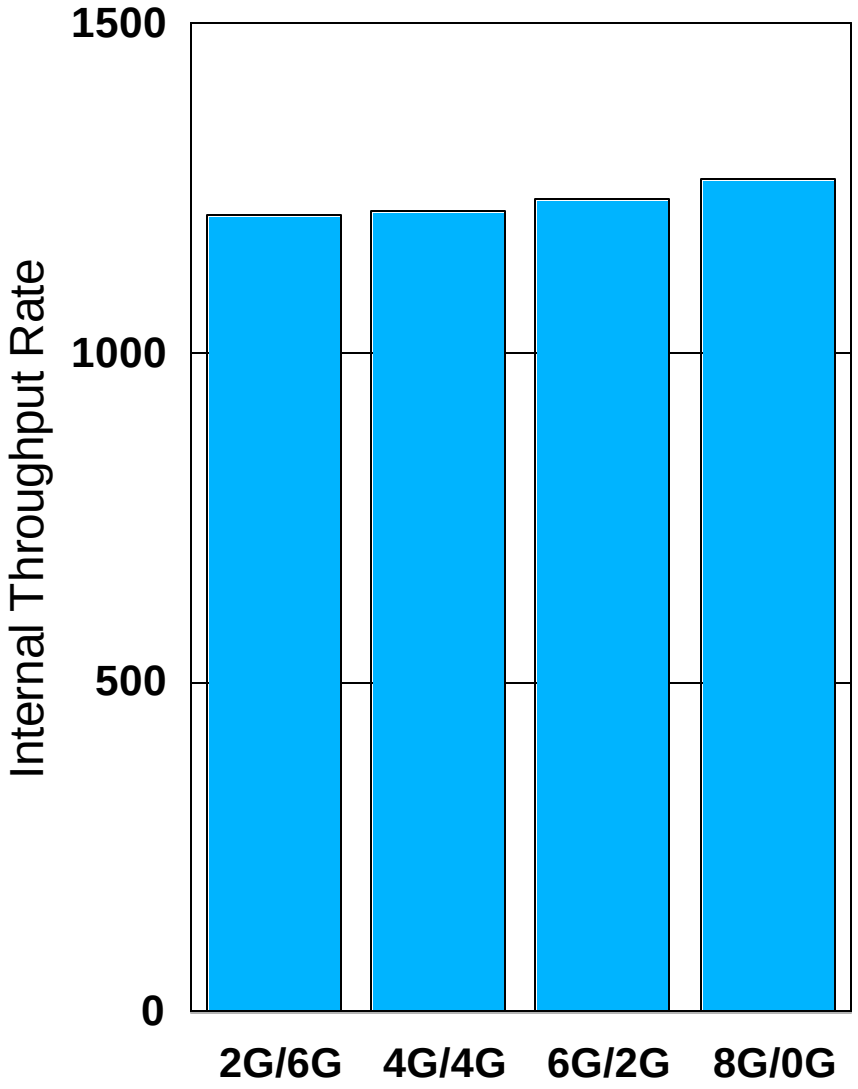

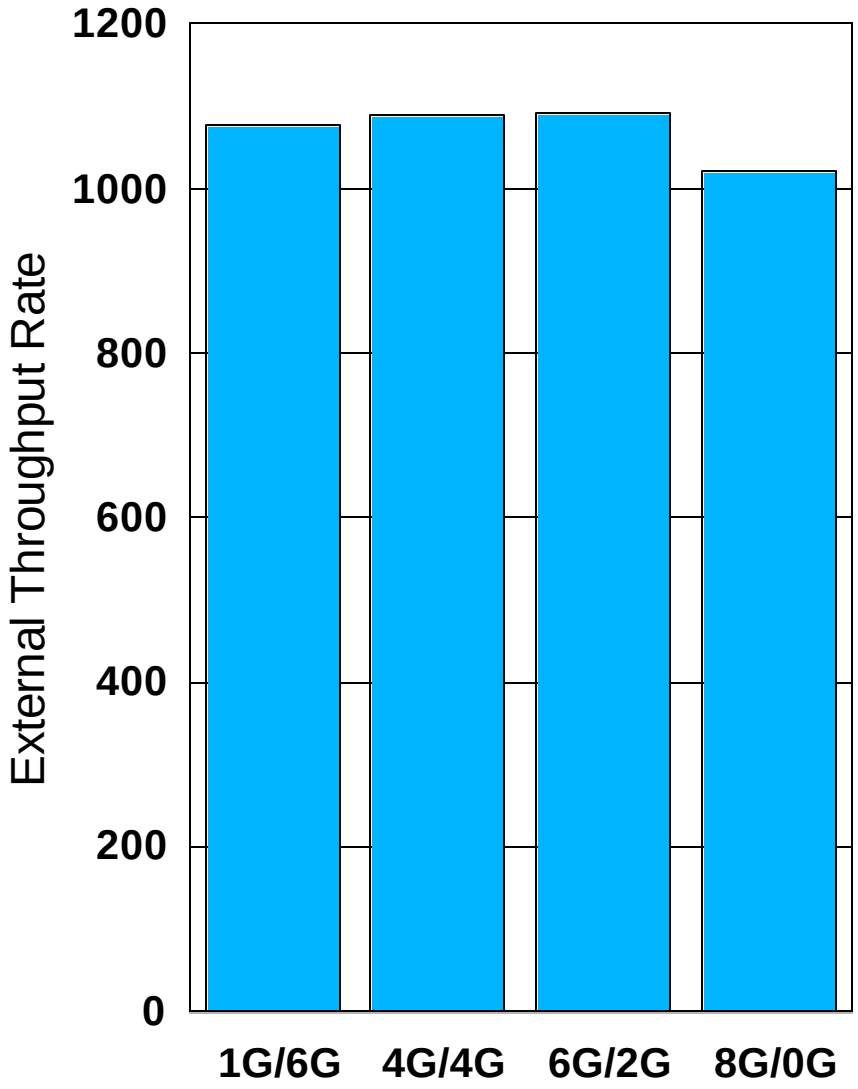

2064-1C8, 8-way; 10800 users; CMS1 Internal TPNS

Storage Allocation - 12 GB

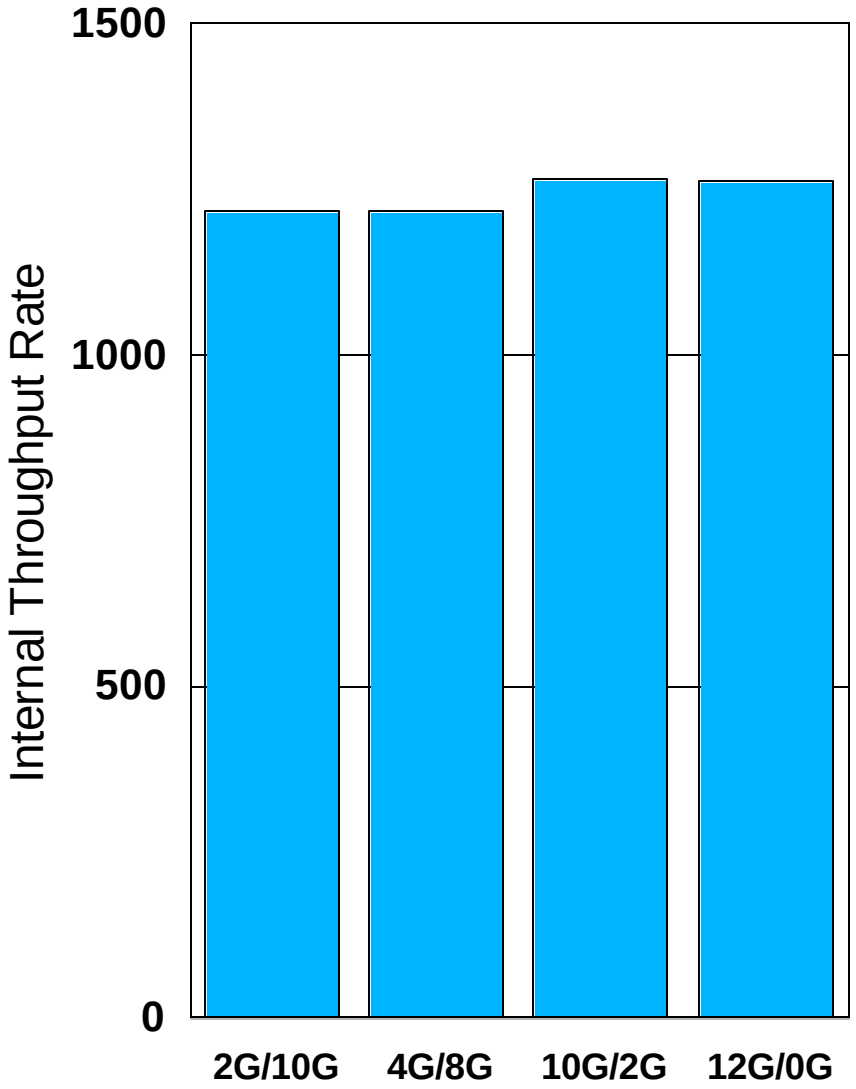

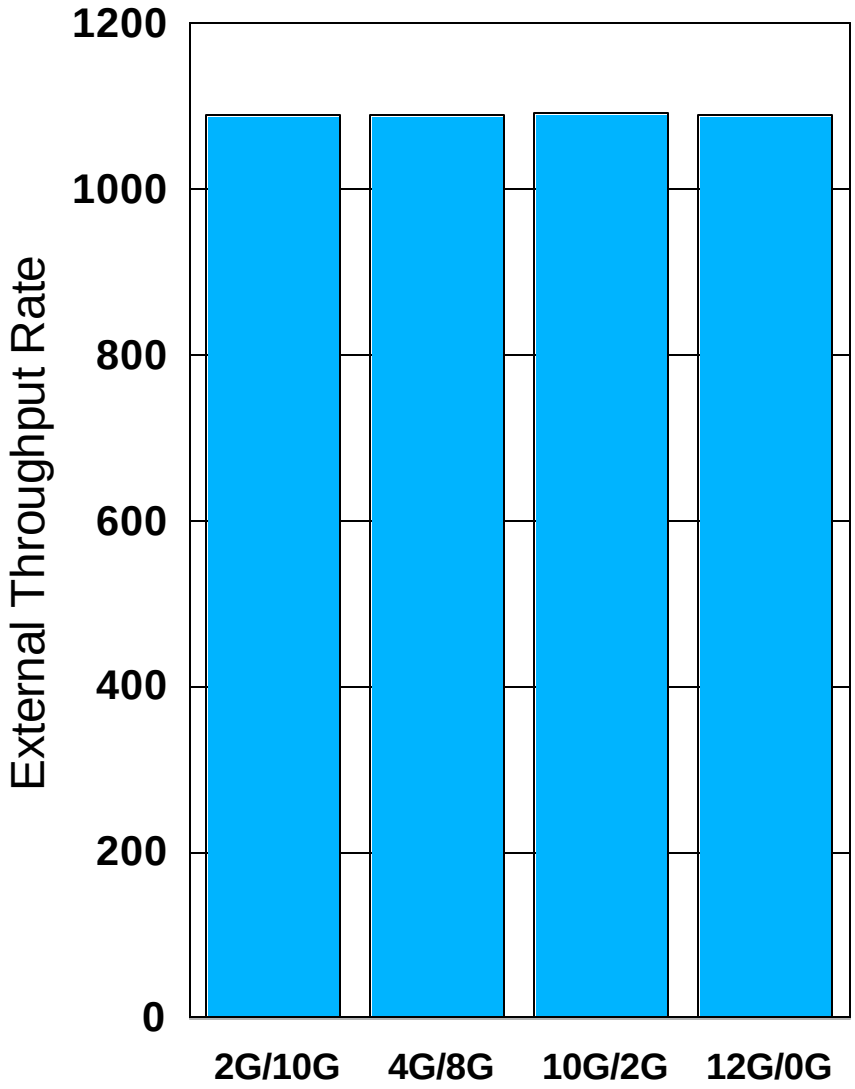

2064-1C8, 8-way; 10800 users; CMS1 Internal TPNS

MDC Tuning 8GB

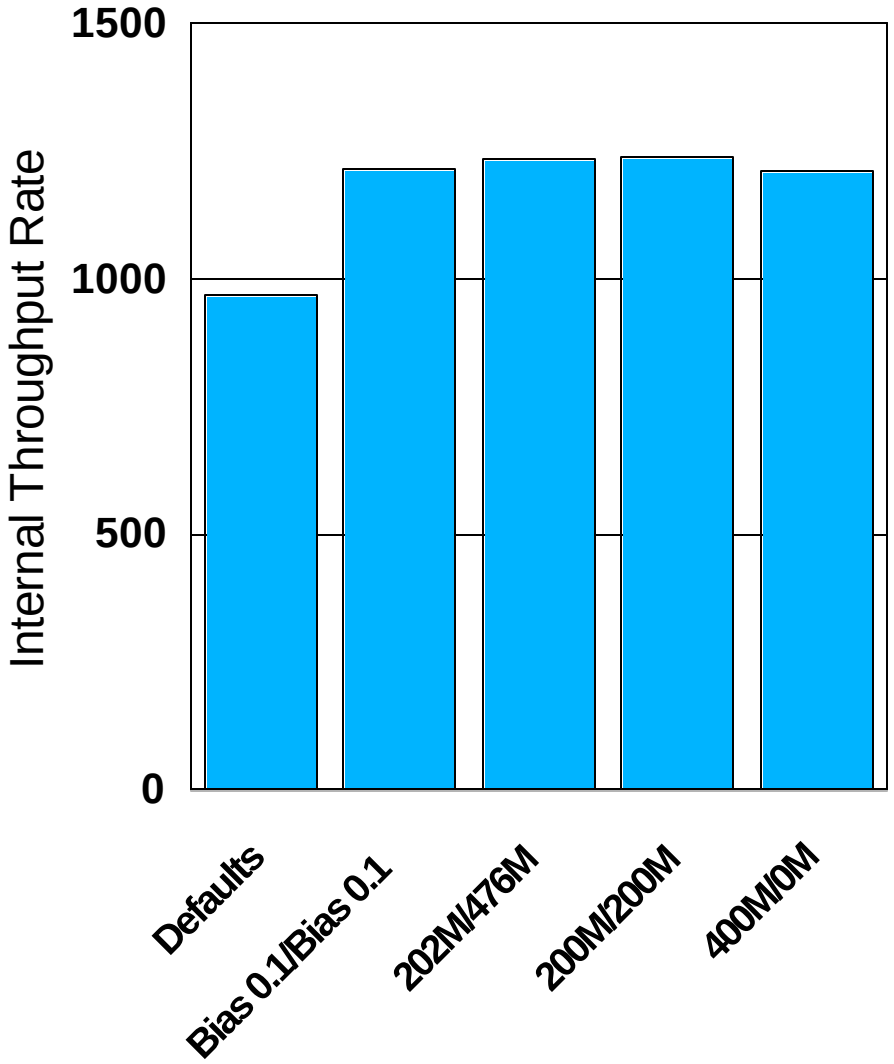

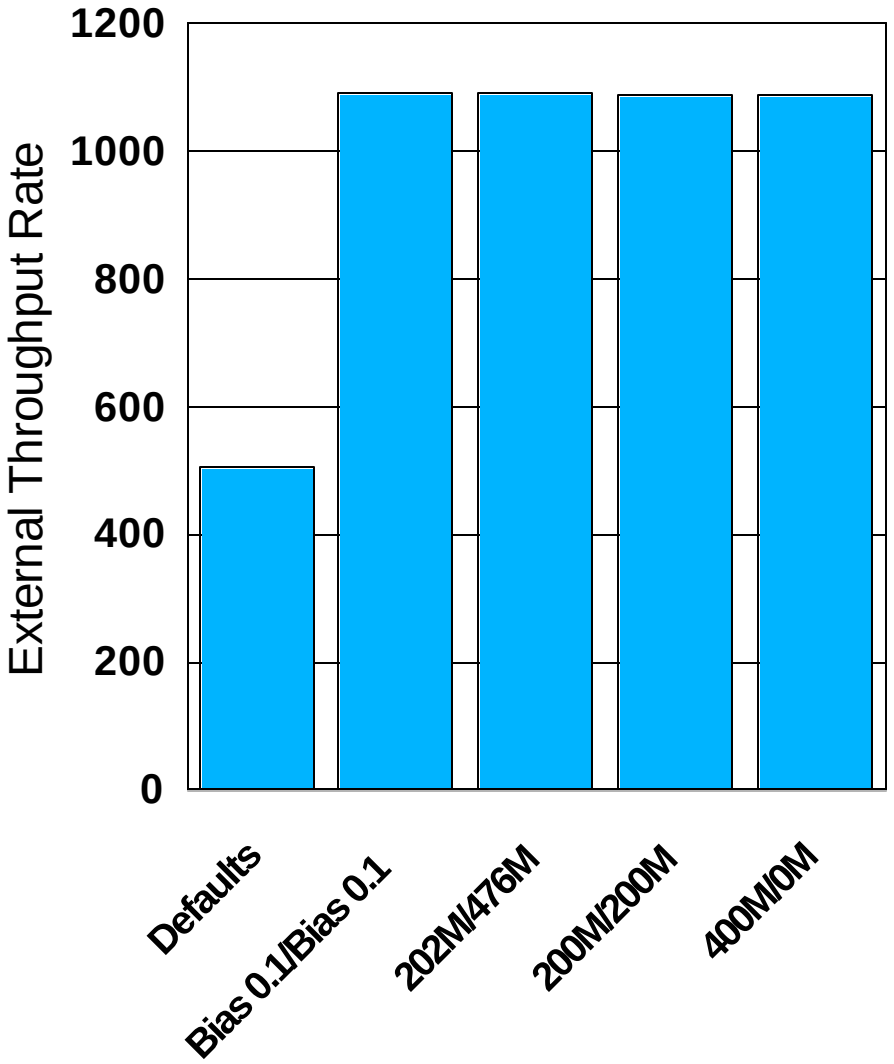

2064-1C8, 8-way; 6G real/2G xstore; 10800 users; CMS1 Internal TPNS

MDC 8GB - I/O per Command

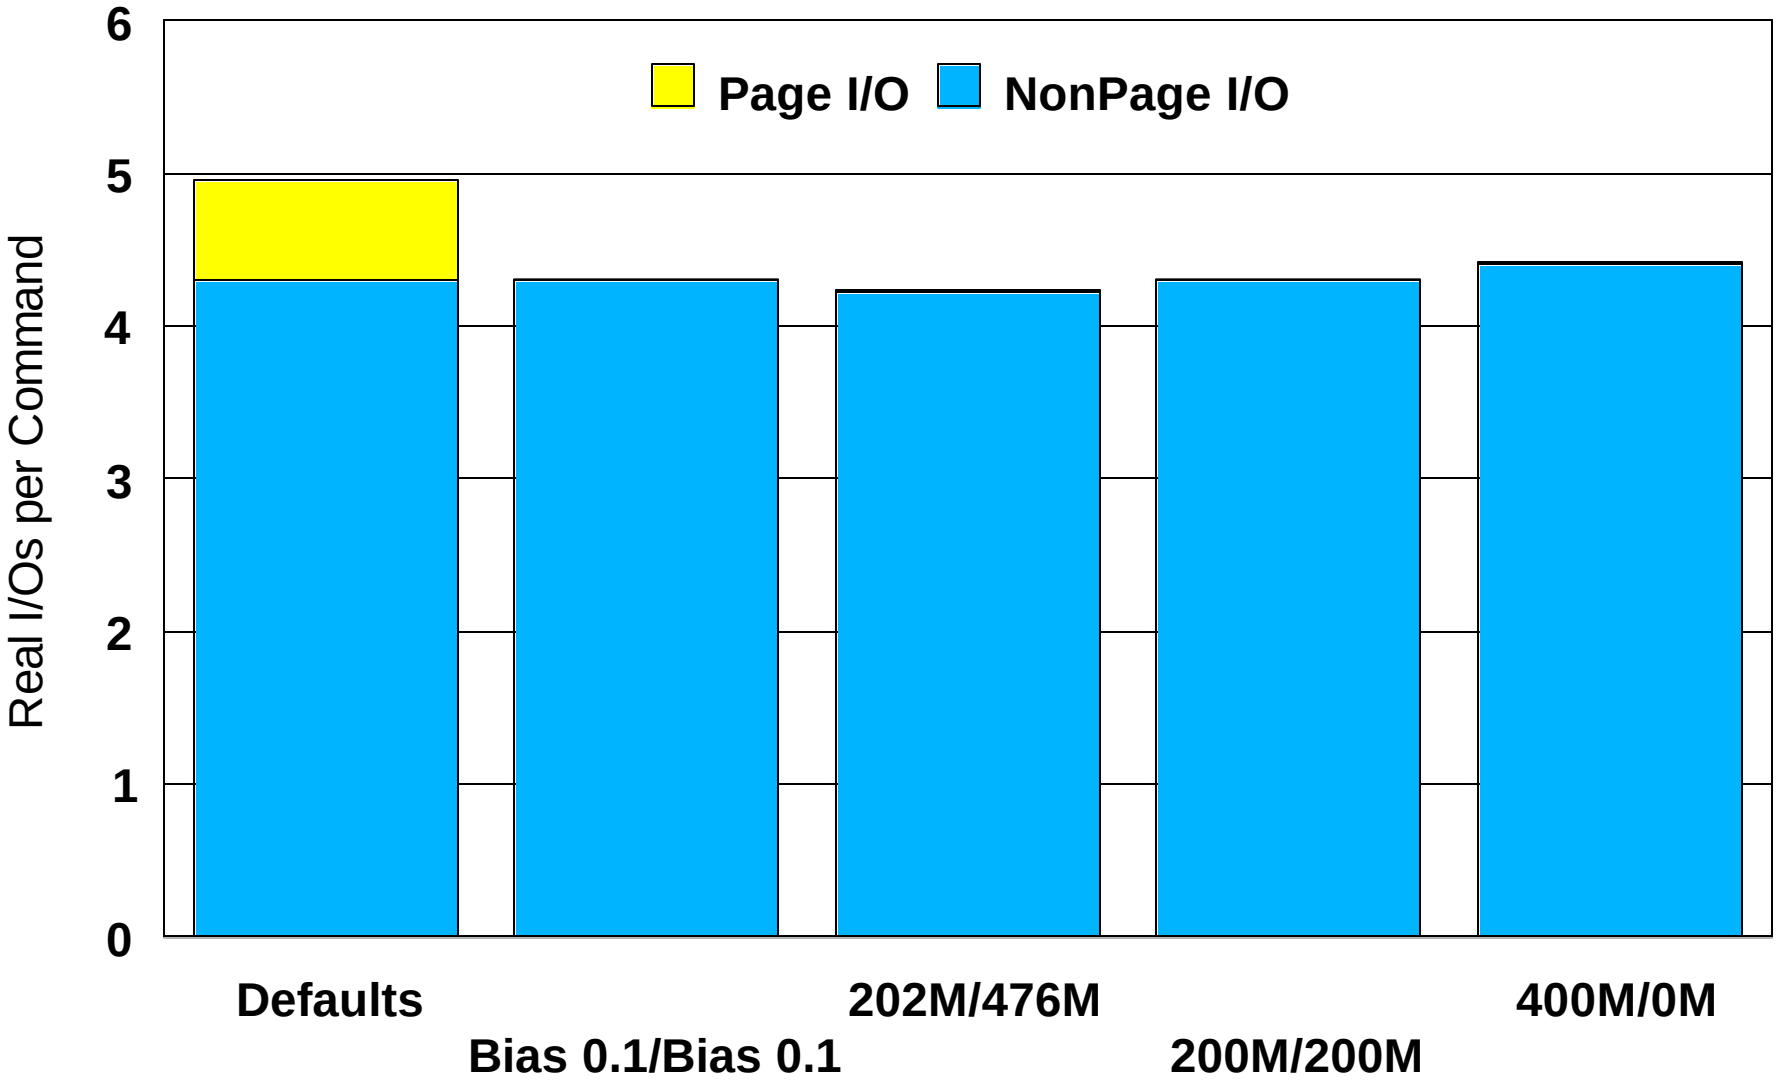

2064-1C8, 8-way; 6G real/2G xstore; 10800 users; CMS1 Internal TPNS

MDC Tuning 12 GB

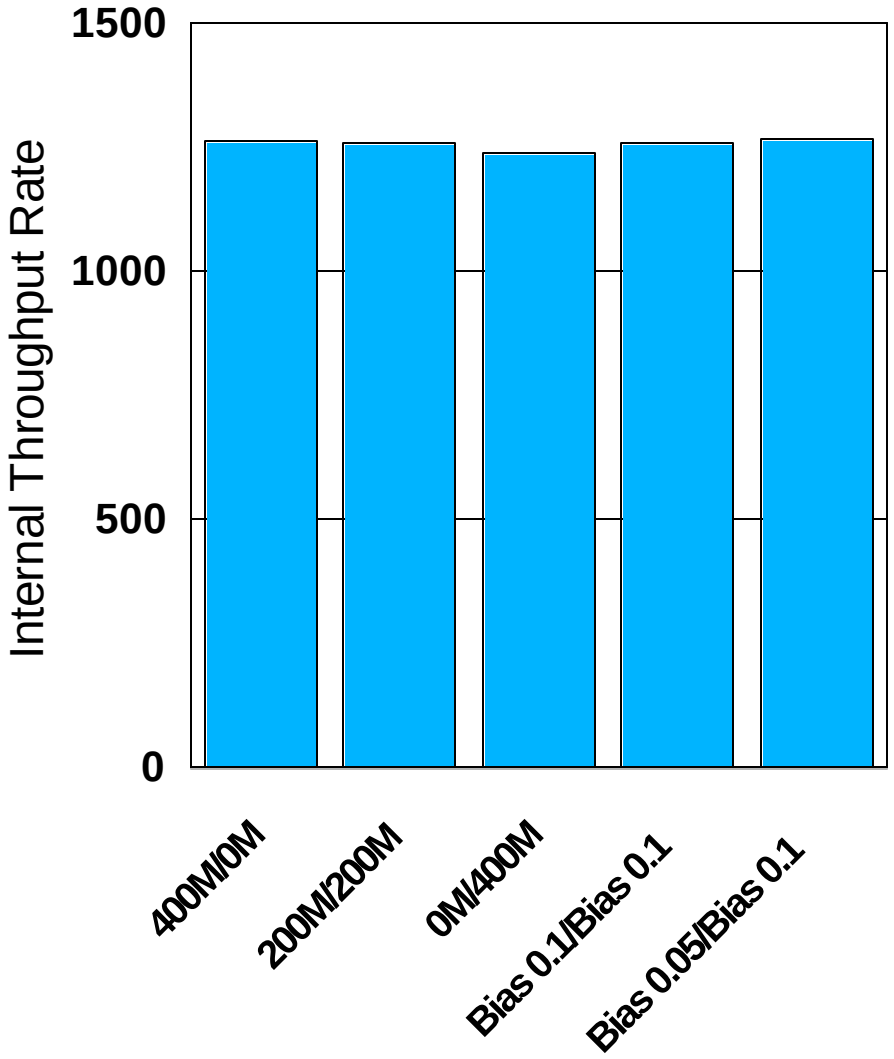

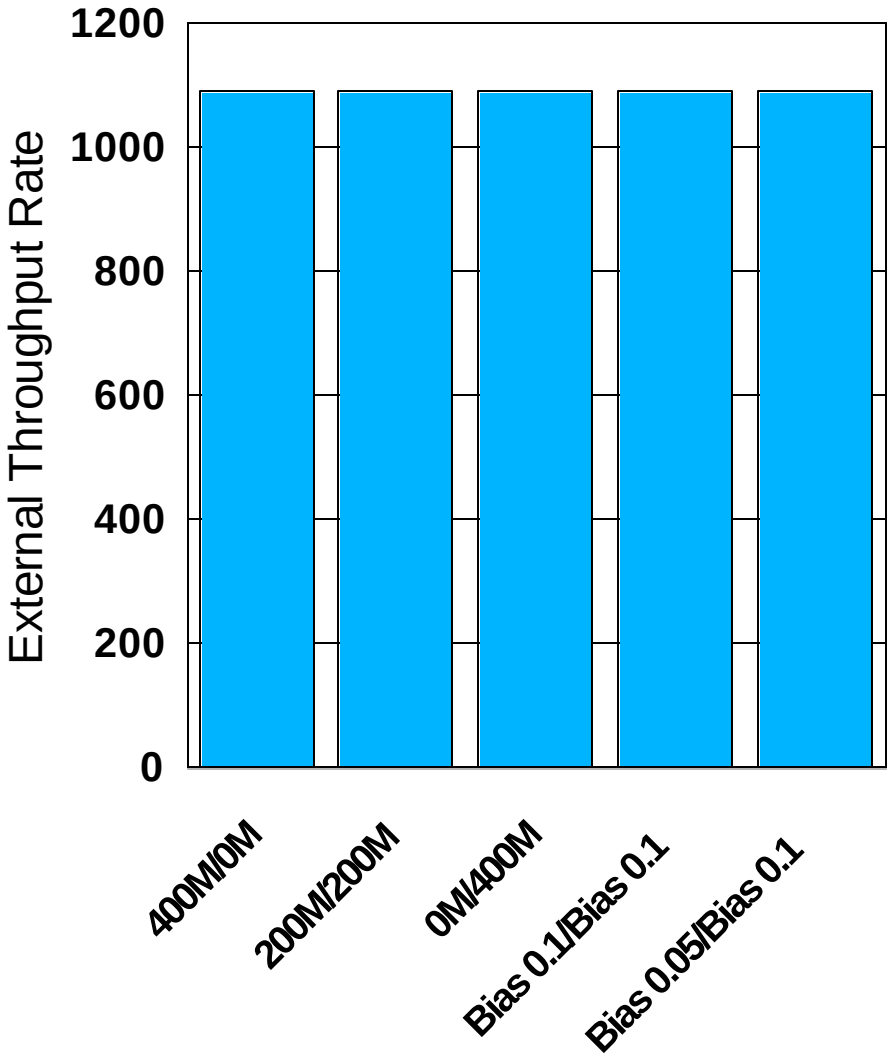

2064-1C8, 8-way; 10G real/2G xstore; 10800 users; CMS1 Internal TPNS

MDC 12GB - I/O per Command

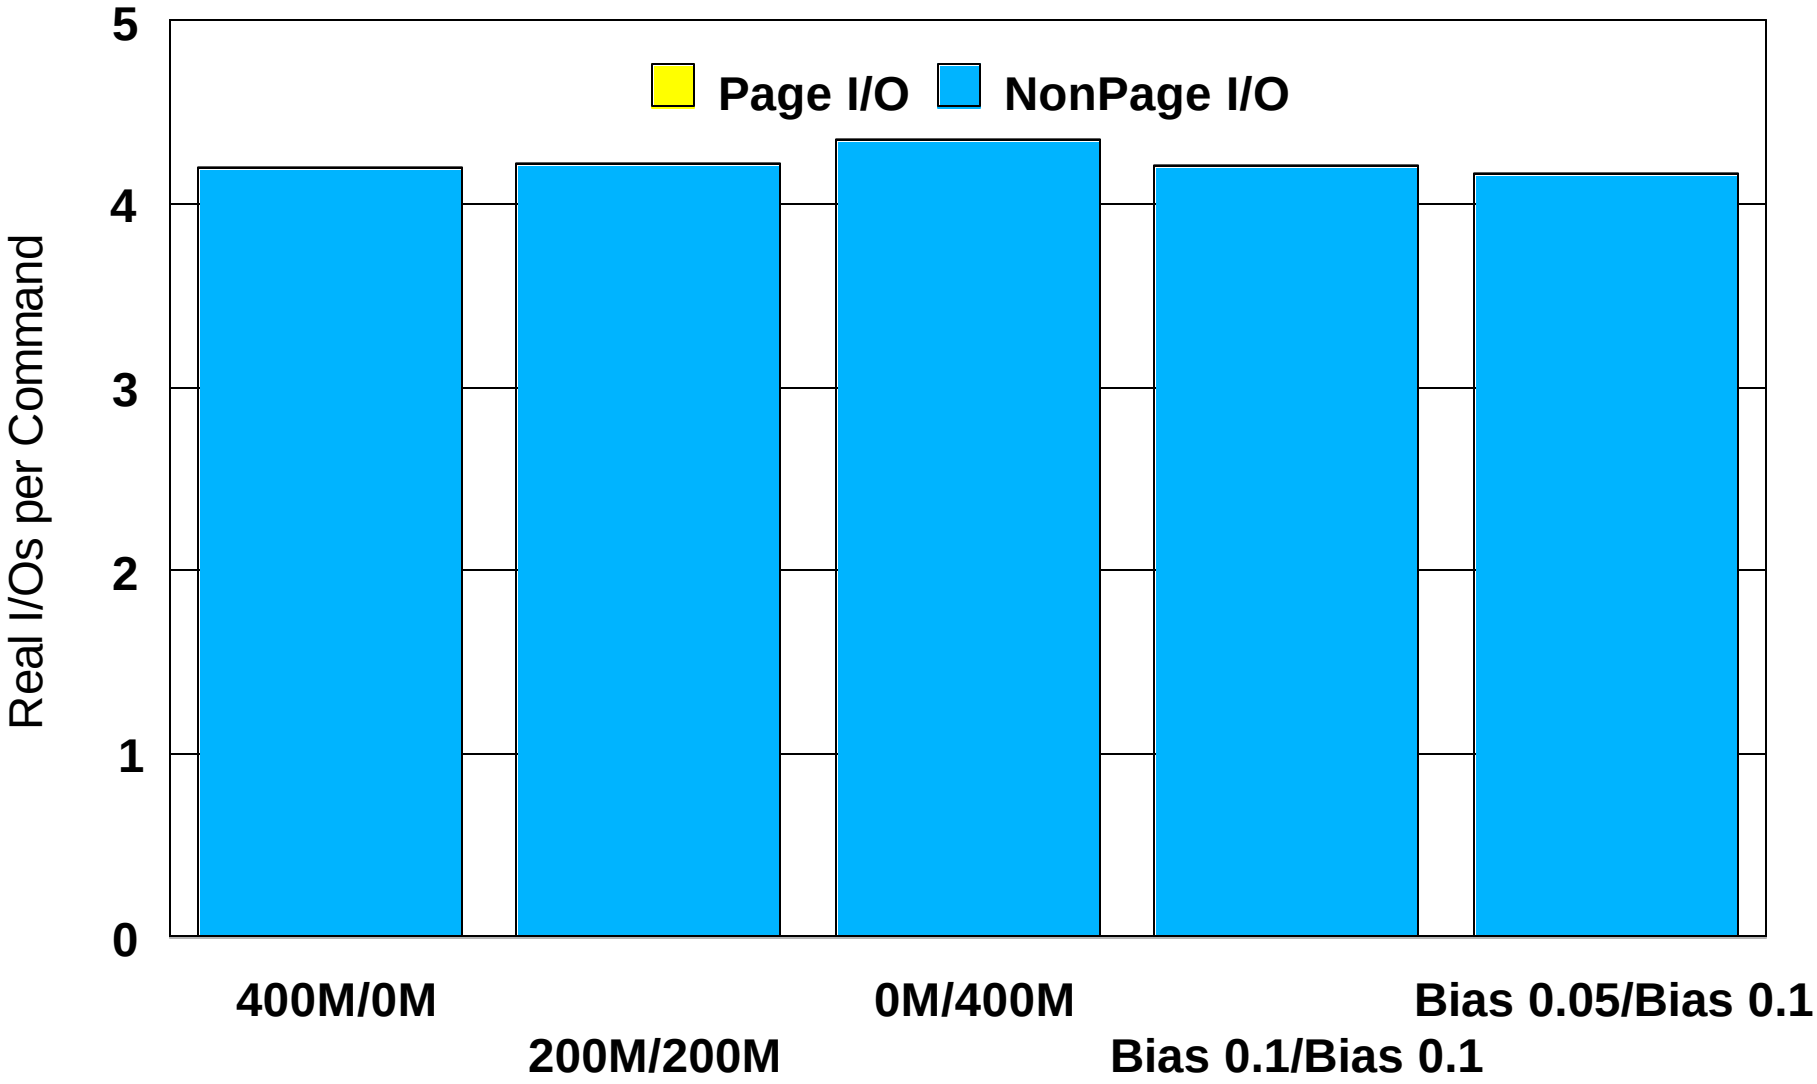

2064-1C8, 8-way; 6G real/2G xstore; 10800 users; CMS1 Internal TPNS

Contention Below 2G

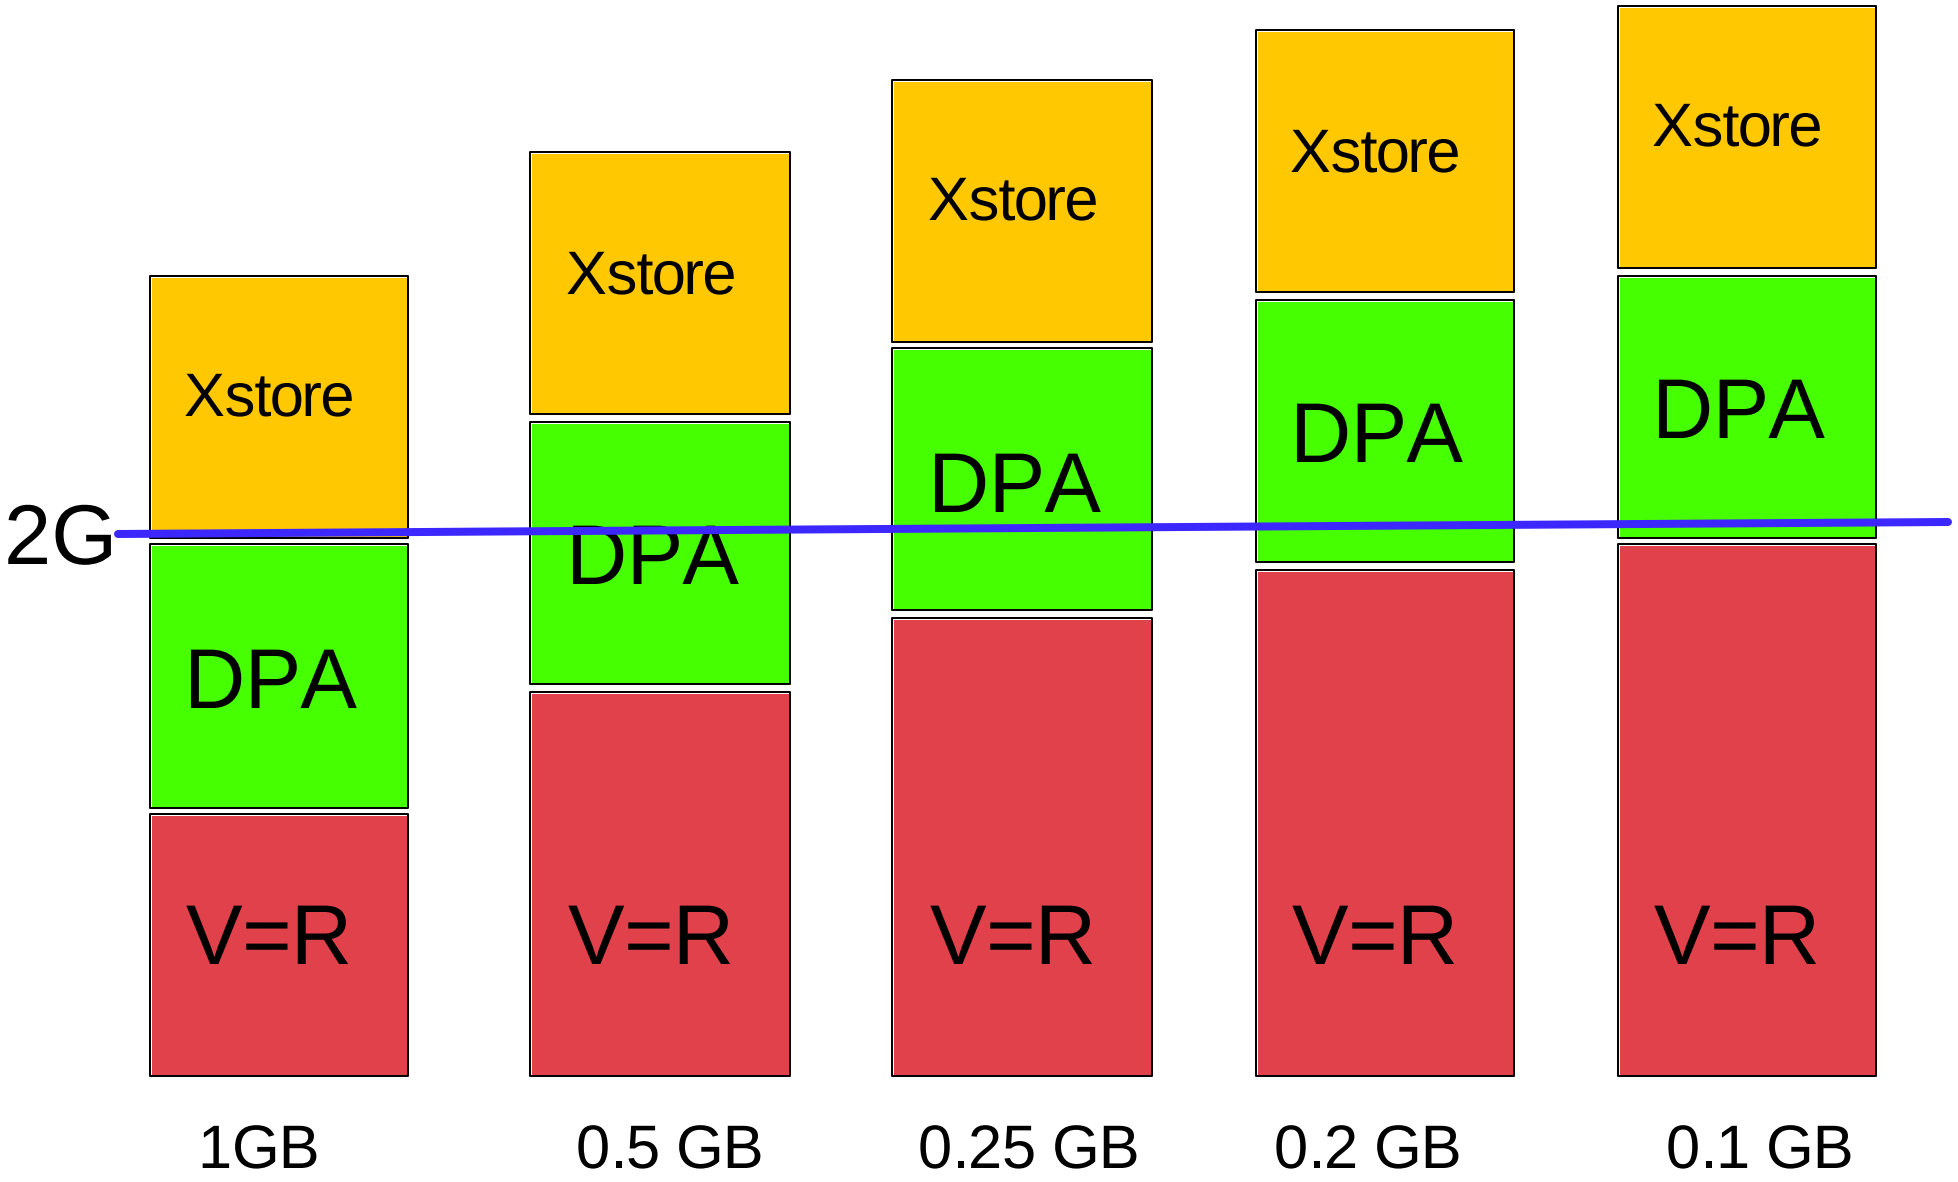

Contention Below 2G

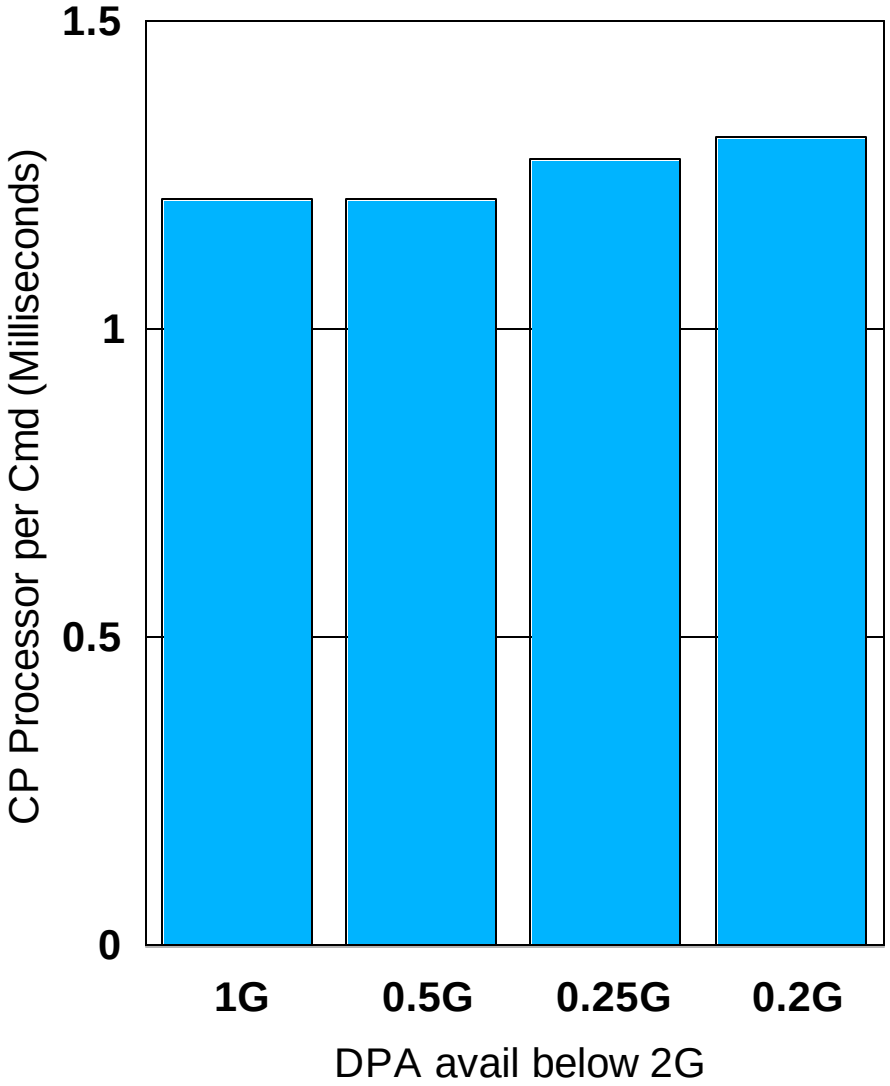

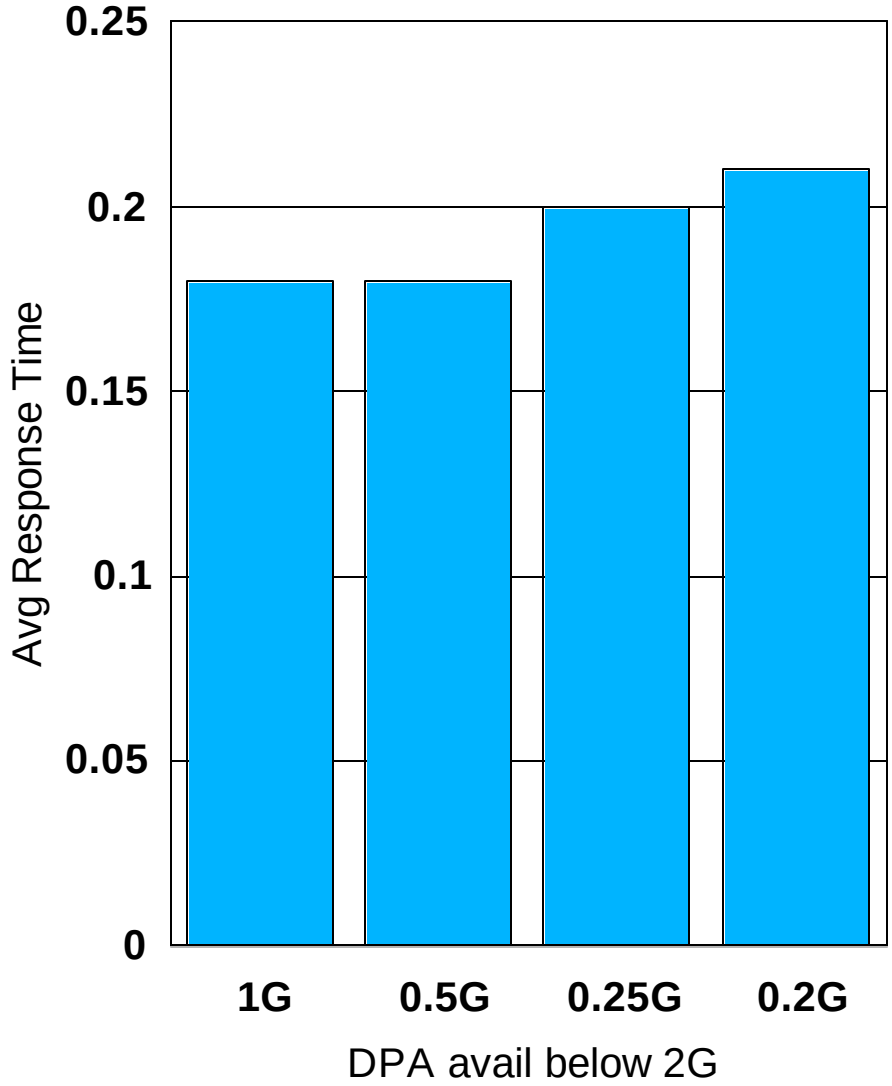

2064-109; LPAR 2-way; 3420 users; CMS1 External TPNS

Storage Recommendations

- Configure some Expanded Storage
- MDC
  - With larger real storage, limit MDC with either maximum or bias settings
  - Allow real and expanded storage
- Need to save some storage below 2G
- APAR VM62827 - corrects reorder frequency being too high.

Storage Exploitation

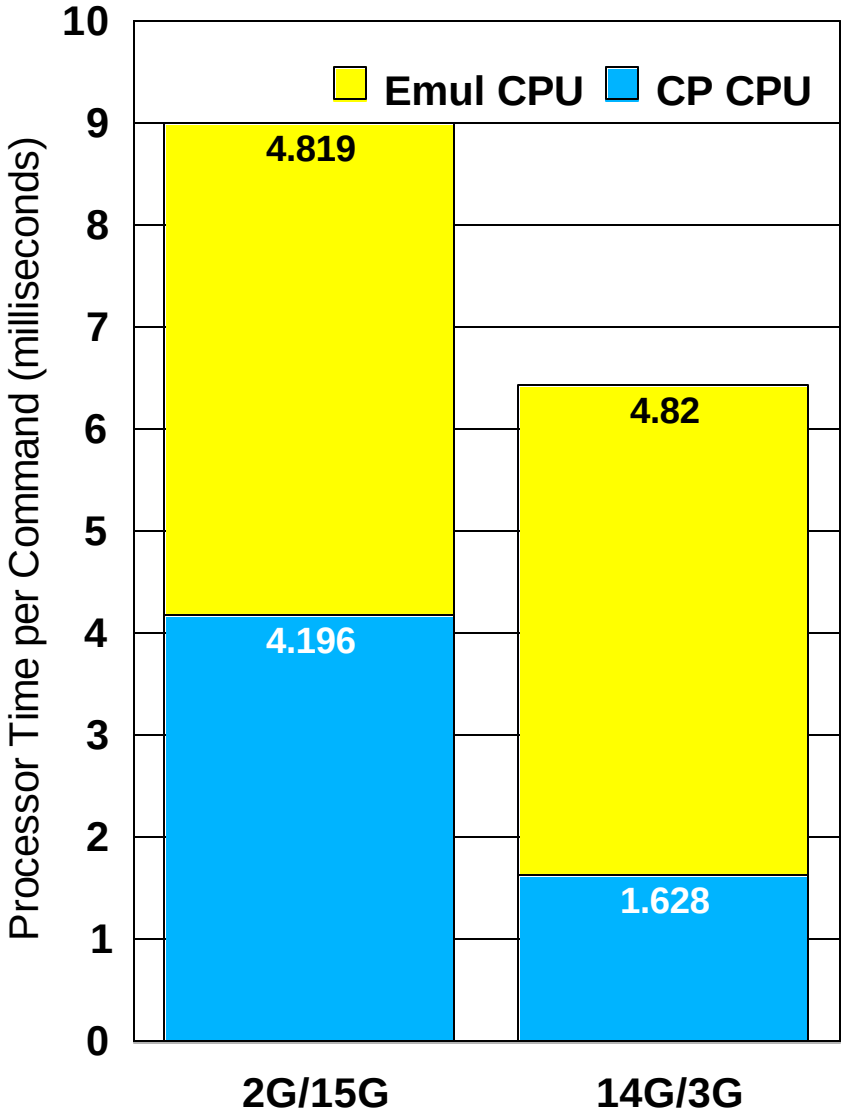

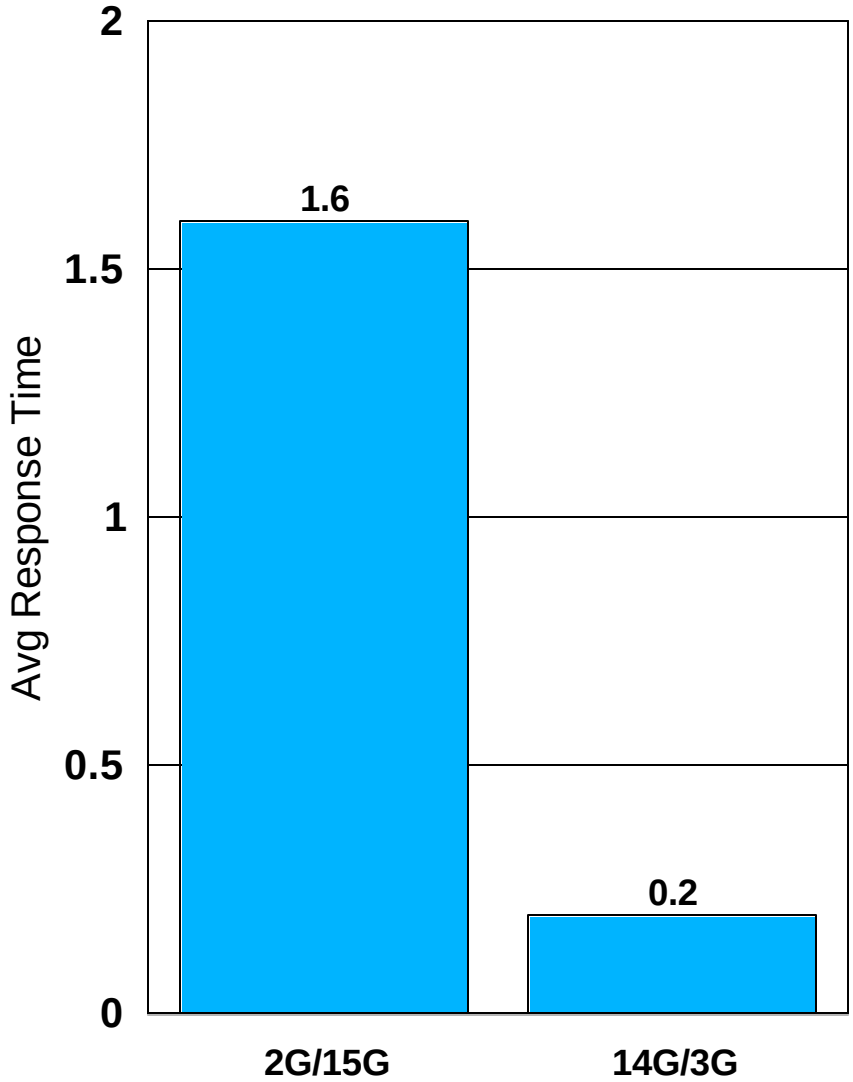

2064-110; 17G total; CMS1 Internal TPNS; zVM 3.1.0

Queued Direct I/O Support

- Previously QDIO available to guests
- TCP/IP Level 3A0 uses for Gigabit Ethernet
- Available on G5, G6, and zSeries processors
- Data transfer via data queues instead of SSCH
- Controlled via state-change-signaling protocol
- Also supports ATM and Fast Ethernet

QDIO Datastream Results

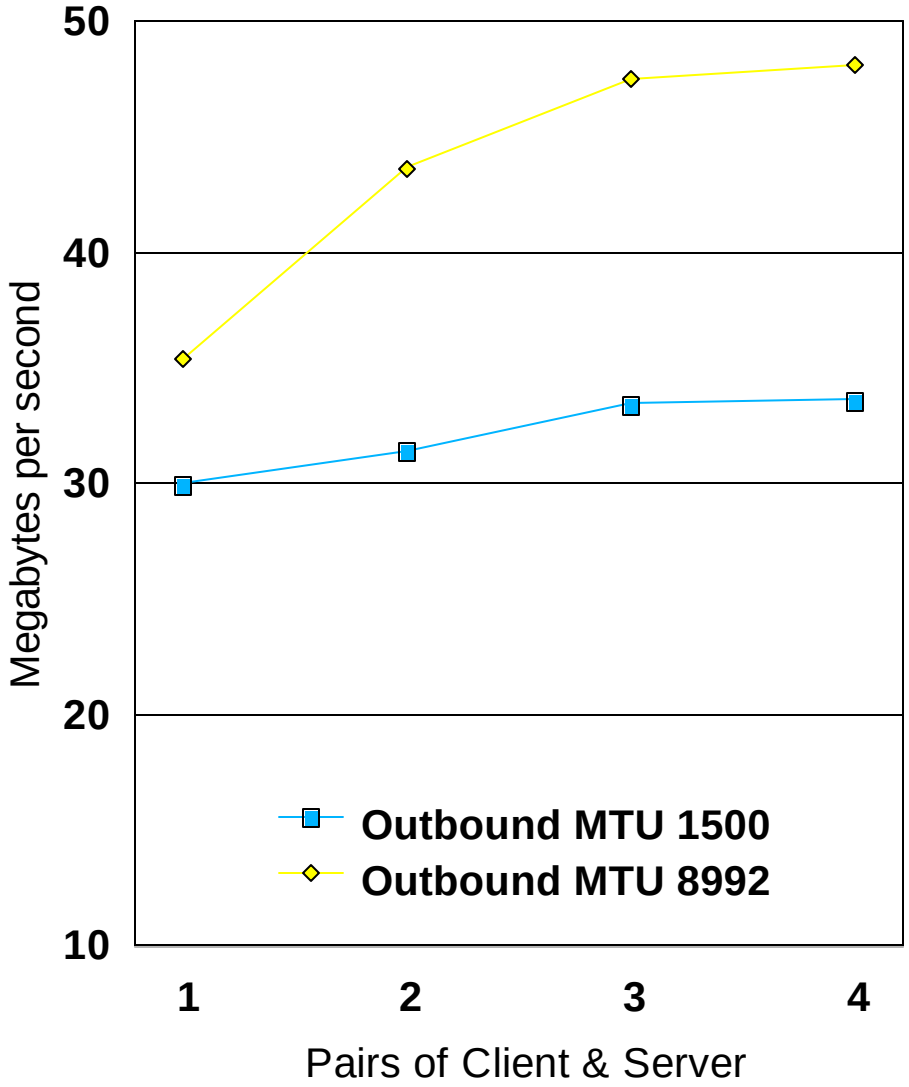

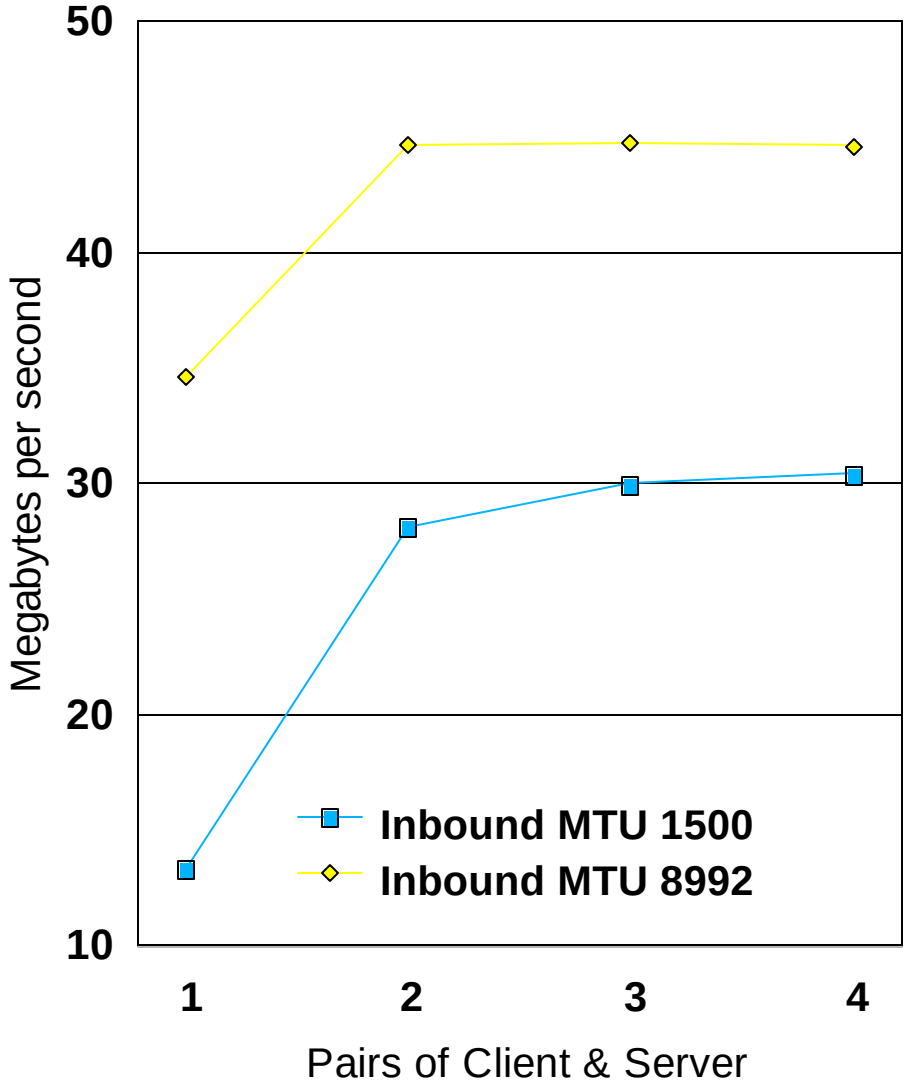

9672-ZZ7 LPAR; z/VM 3.1.0 TCP/IP 3A0

Secure Socket Layer Support

- Provided by new SSL server virtual machine
- Additional processing for secure connections
  - Handshaking at connect time
    - determine cryptographic parameters
    - some data can be cached
  - Encrypt/decrypt overhead while transferring data

SSL Environment

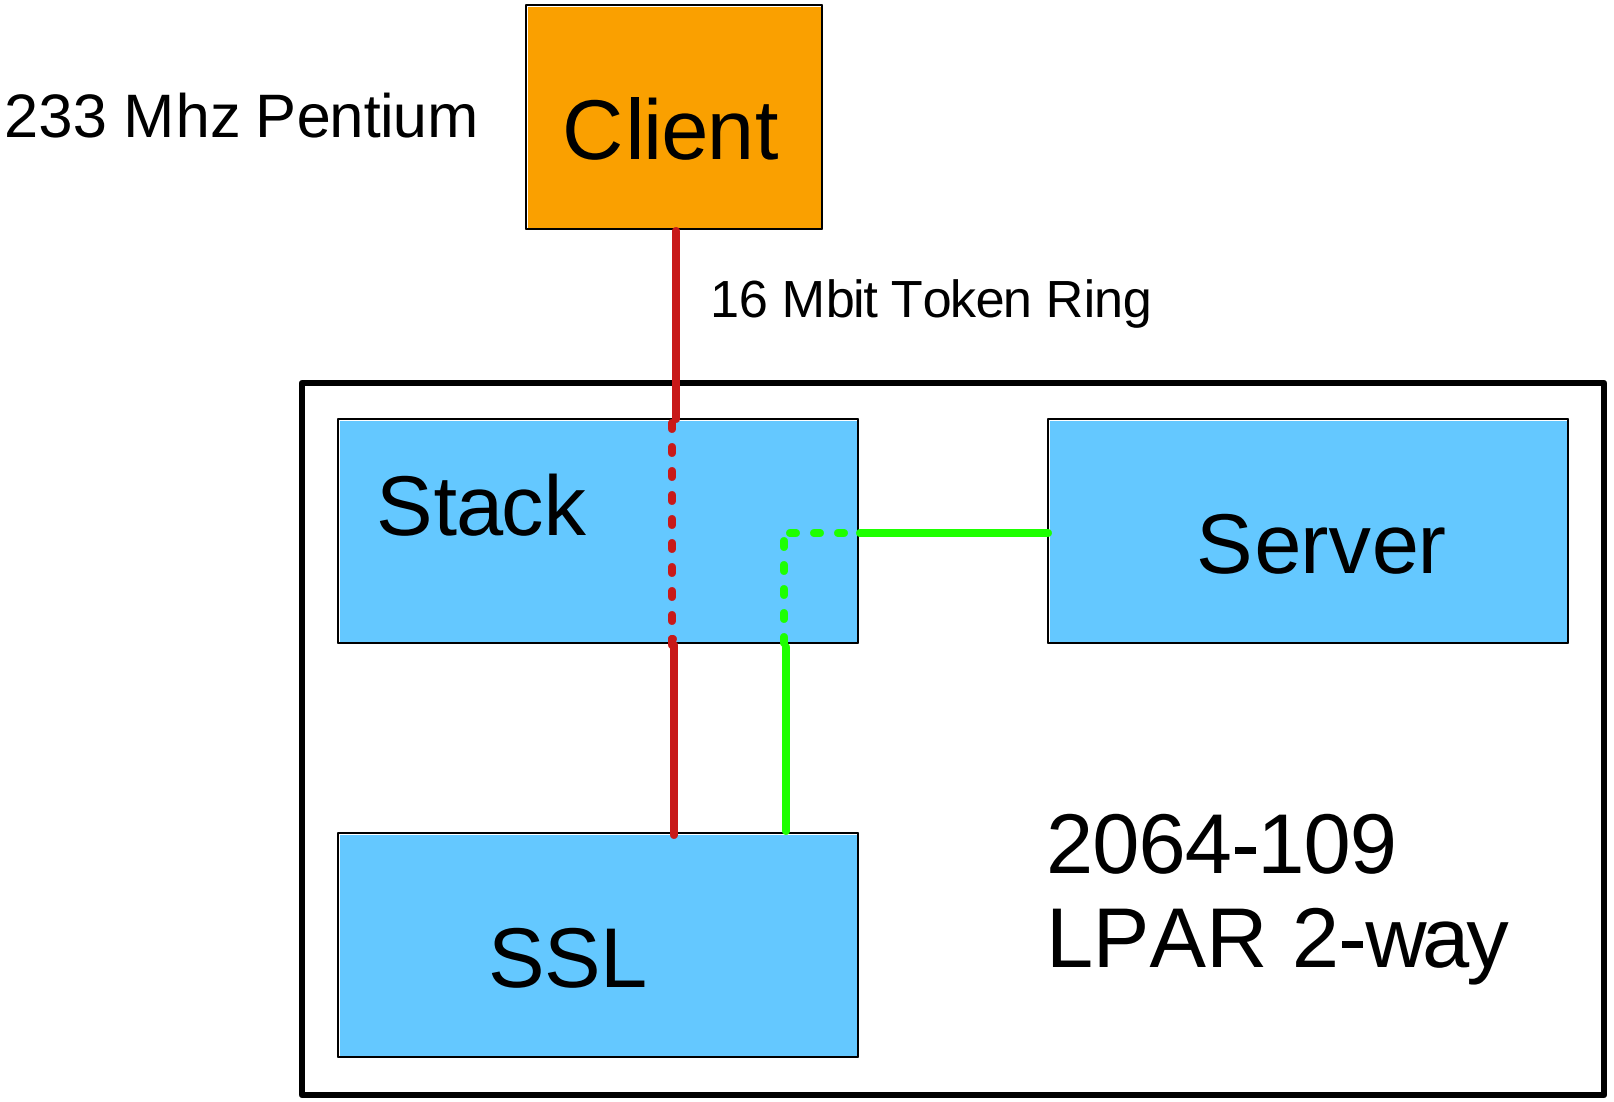

SSL Connect - New Session

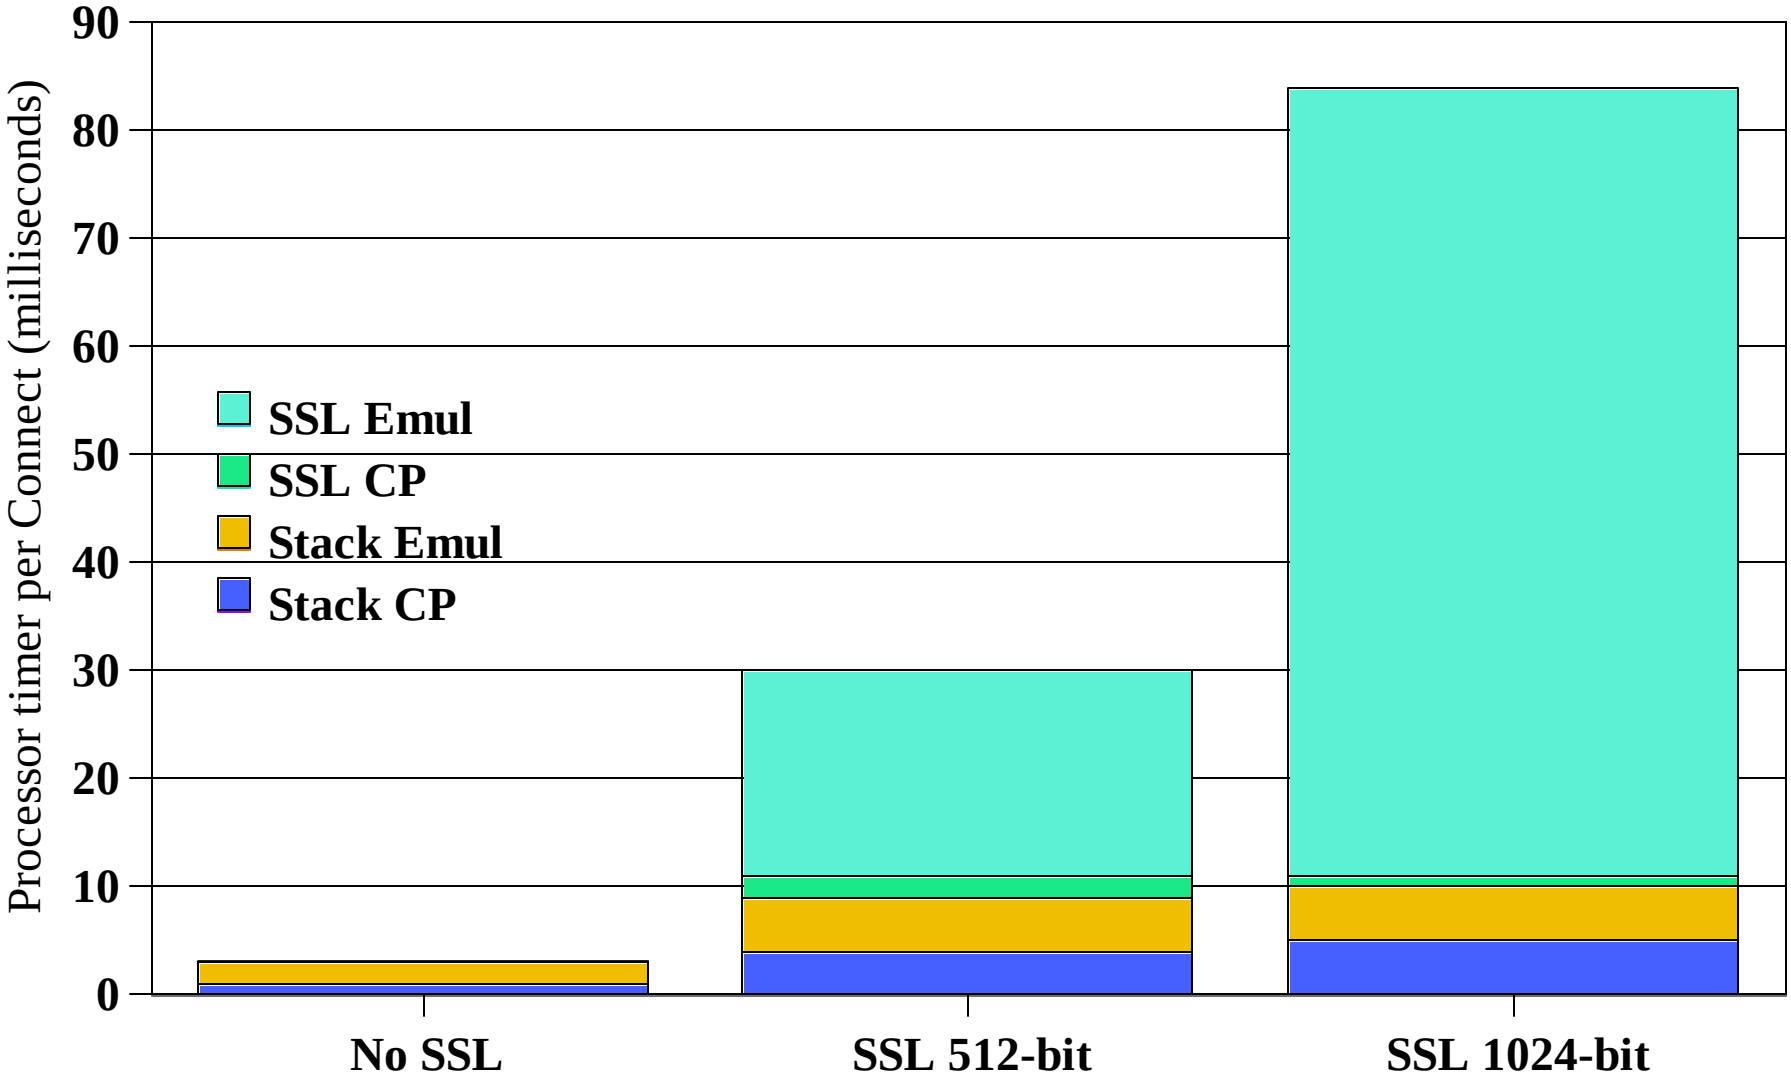

SSL Connect - Resume Session

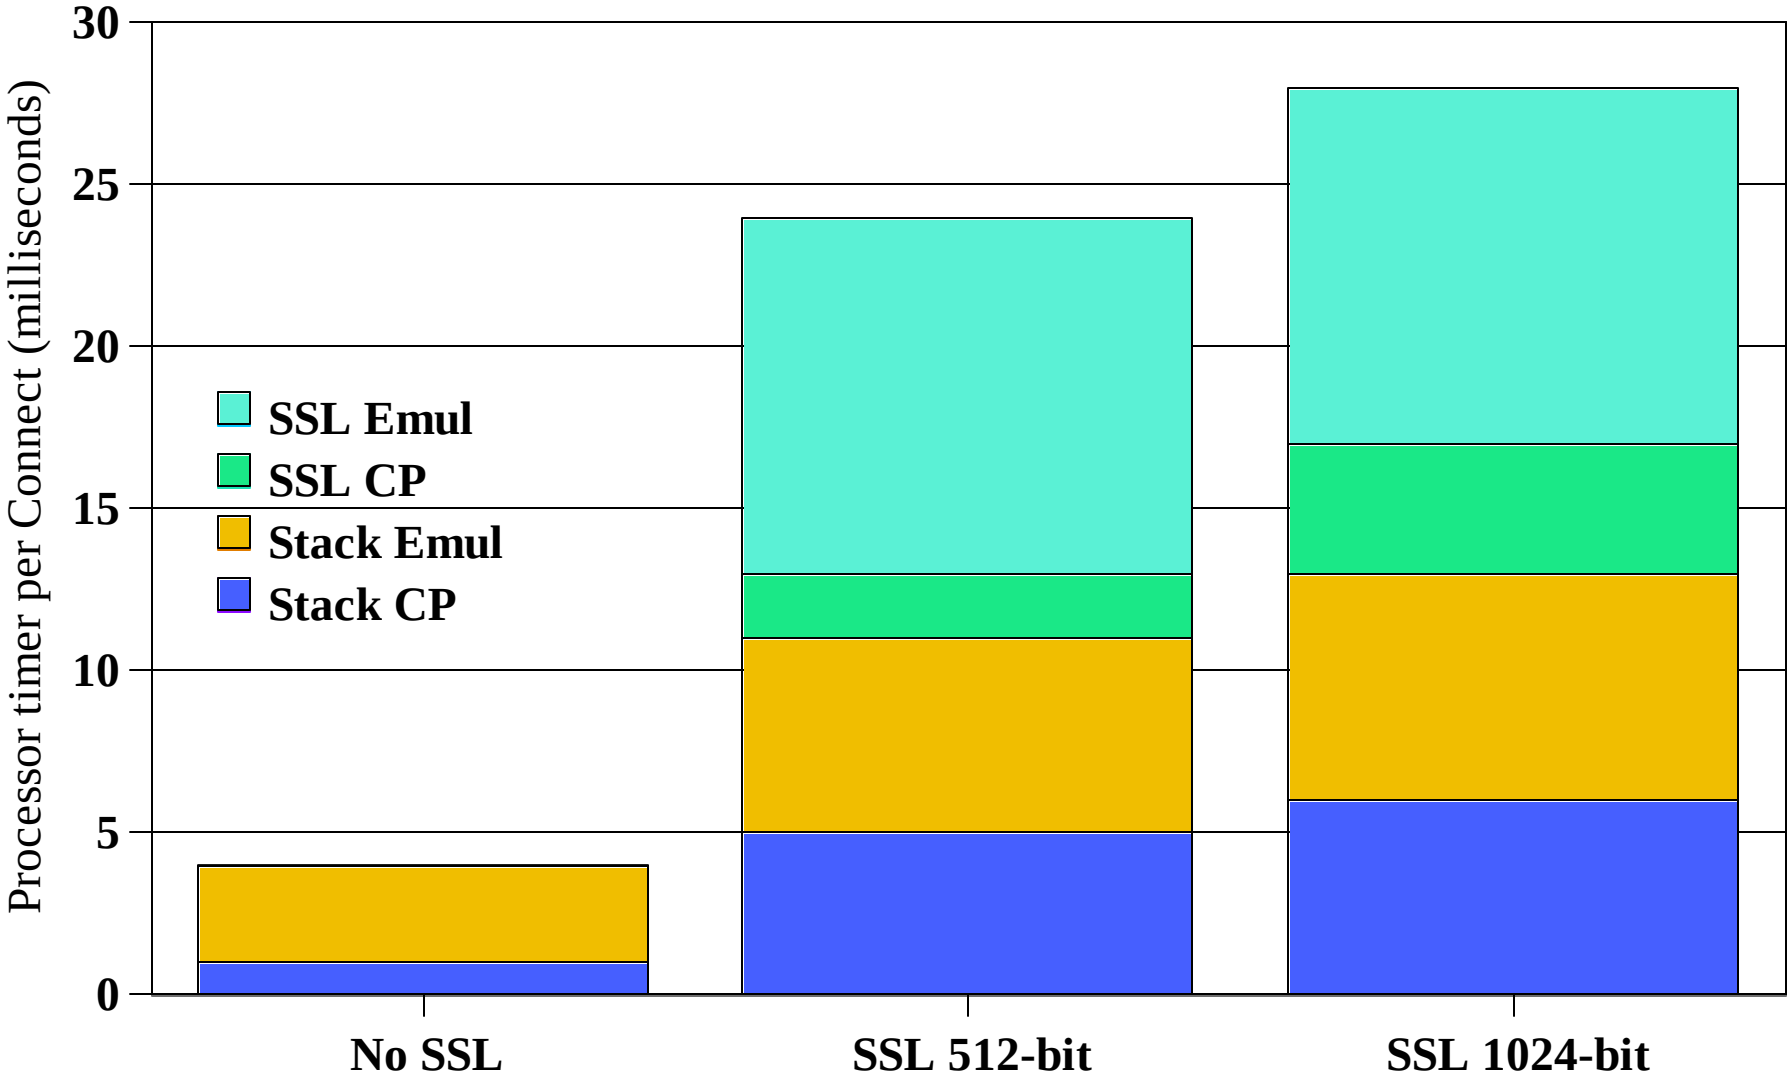

SSL - FTP Binary Get 10M

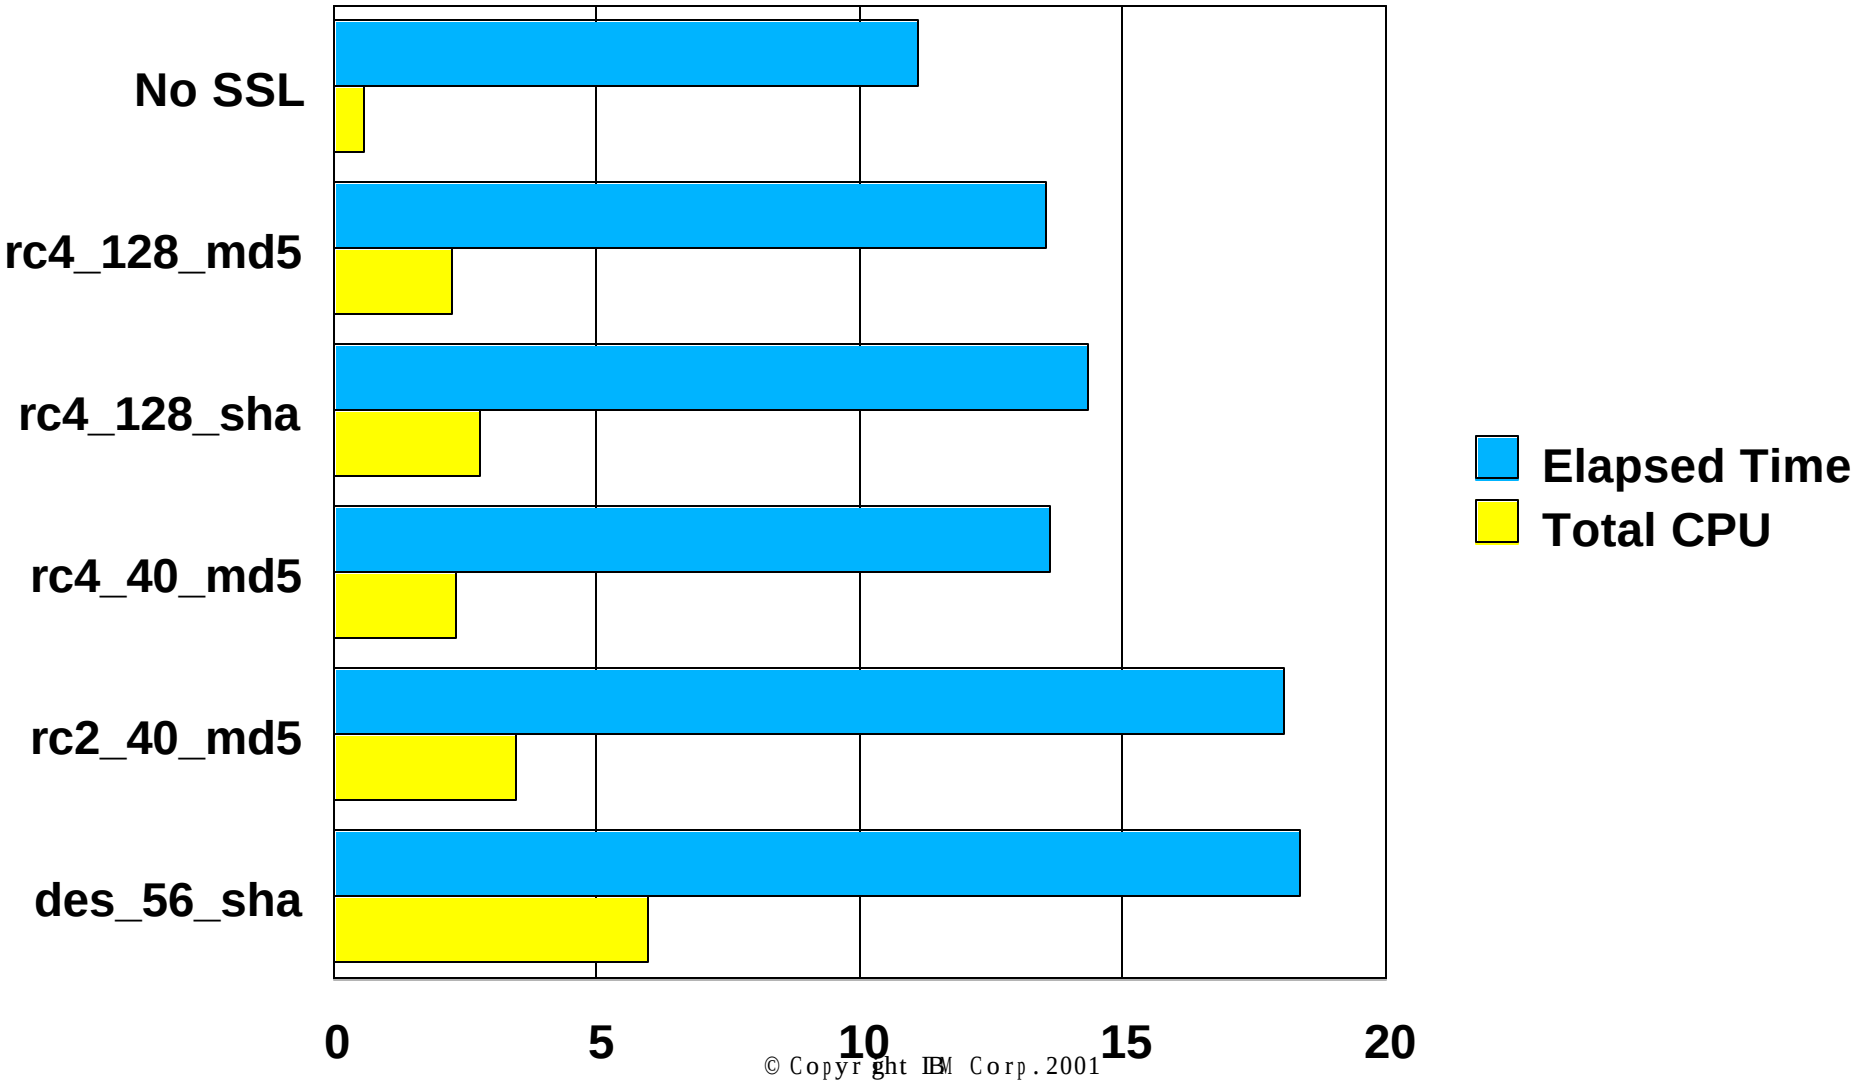

Monitor Enhancements

- Most monitor reduction programs should work without change for regression environments
- Larger fields to record virtual and real storage sizes
- Indication of virtual machines in 64-bit
- Record use of storage above/below 2G
- APAR VM62794 - correct shared segment numbers
- Stack records enhanced for QDIO support

IBM Performance Products

- VMPRF 1.2.2
  - 64-bit support
  - New reports
    - SYSTEM_SUMMARY2_BY_TIME
    - AUXSTORE_BY_TIME
    - NONDASD_BY_ACTIVITY or _BY CONFIG
- RTM for z/VM 3.1.0
  - 64-bit support
  - No longer requires 370 Accommodation
  - Configuration file avoids some mods
- FCON/ESA Version 3.2.02
  - 64-bit support
  - TCP/IP Level 3A0 support
- VM/PAF 1.1.3
  - Runs on z/VM 3.1.0

Performance Management

- CP logic and control blocks drastically changed
  - Review CP mods
  - Review tools that pull data from CP control blocks
- CP Trace Table Changes
  - Some entries double in size
- QUERY FRAMES

```
SYSGEN   REAL     USABLE   OFFLINE
524287   524287   524287   000000
V=R      RESNUC   PAGING   TRACE    RIO370
000000   000667   523070   000550   000000
AVAIL    PAGNUC   LOCKRS   LOCKCP   SAVE     FREE     LOCKRIO
506751   009916   000300   000000   000061   006042   000000
Storage >= 2G:
   Online          = 786432      Available List = 58941
   Not init        = 0           Offline        = 0
```

Linux Virtual Connectivity

- CMS Driver
  - Really synchronous APPC/VM
  - Very little application/protocol overhead
- Linux using IUCV
  - Internal Tool to drive networks
  - Application and protocol overhead included
- Linux using Virtual CTC
  - Internal Tool to drive networks
  - Application and protocol overhead included
- Test Environment
  - 9672-XZ7, Two processor LPAR with 2G/2G
  - z/VM 3.1.0 running 128MB Linux guests

Linux Virtual Communication

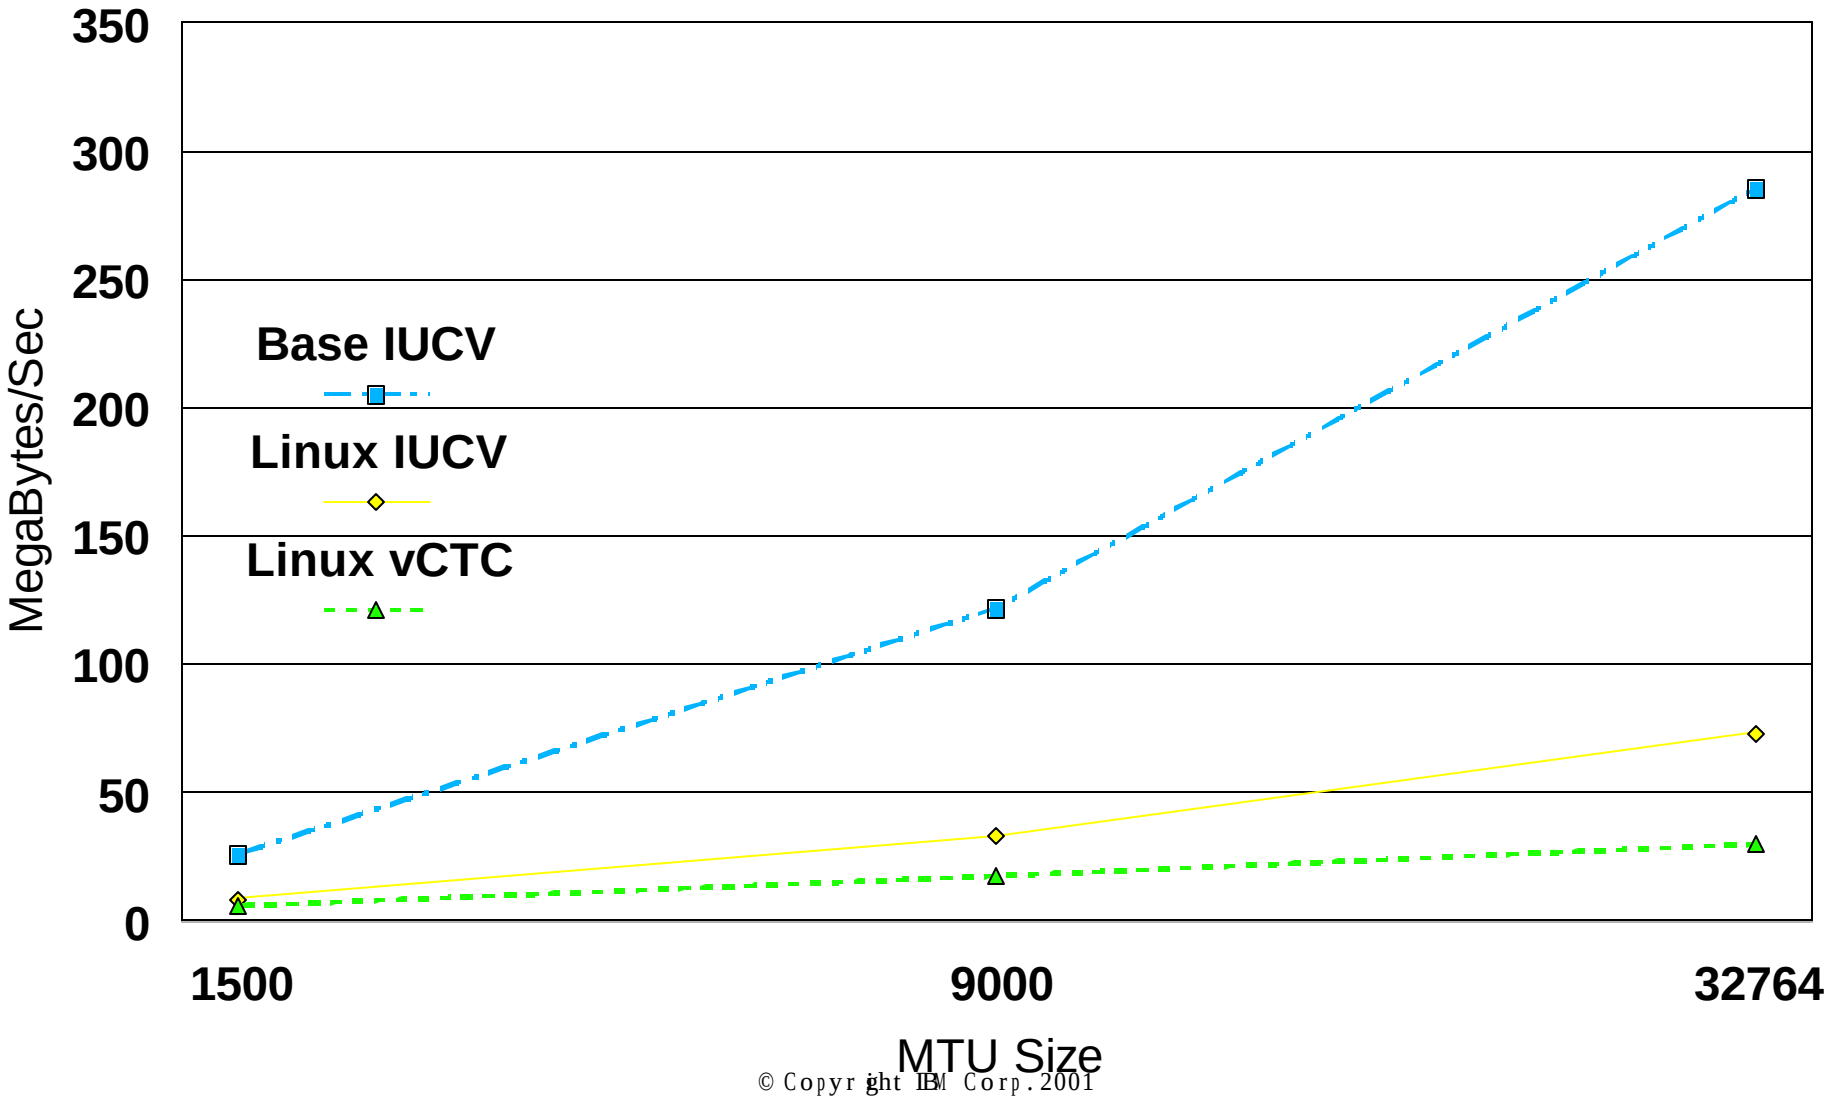

Communication Processor Time

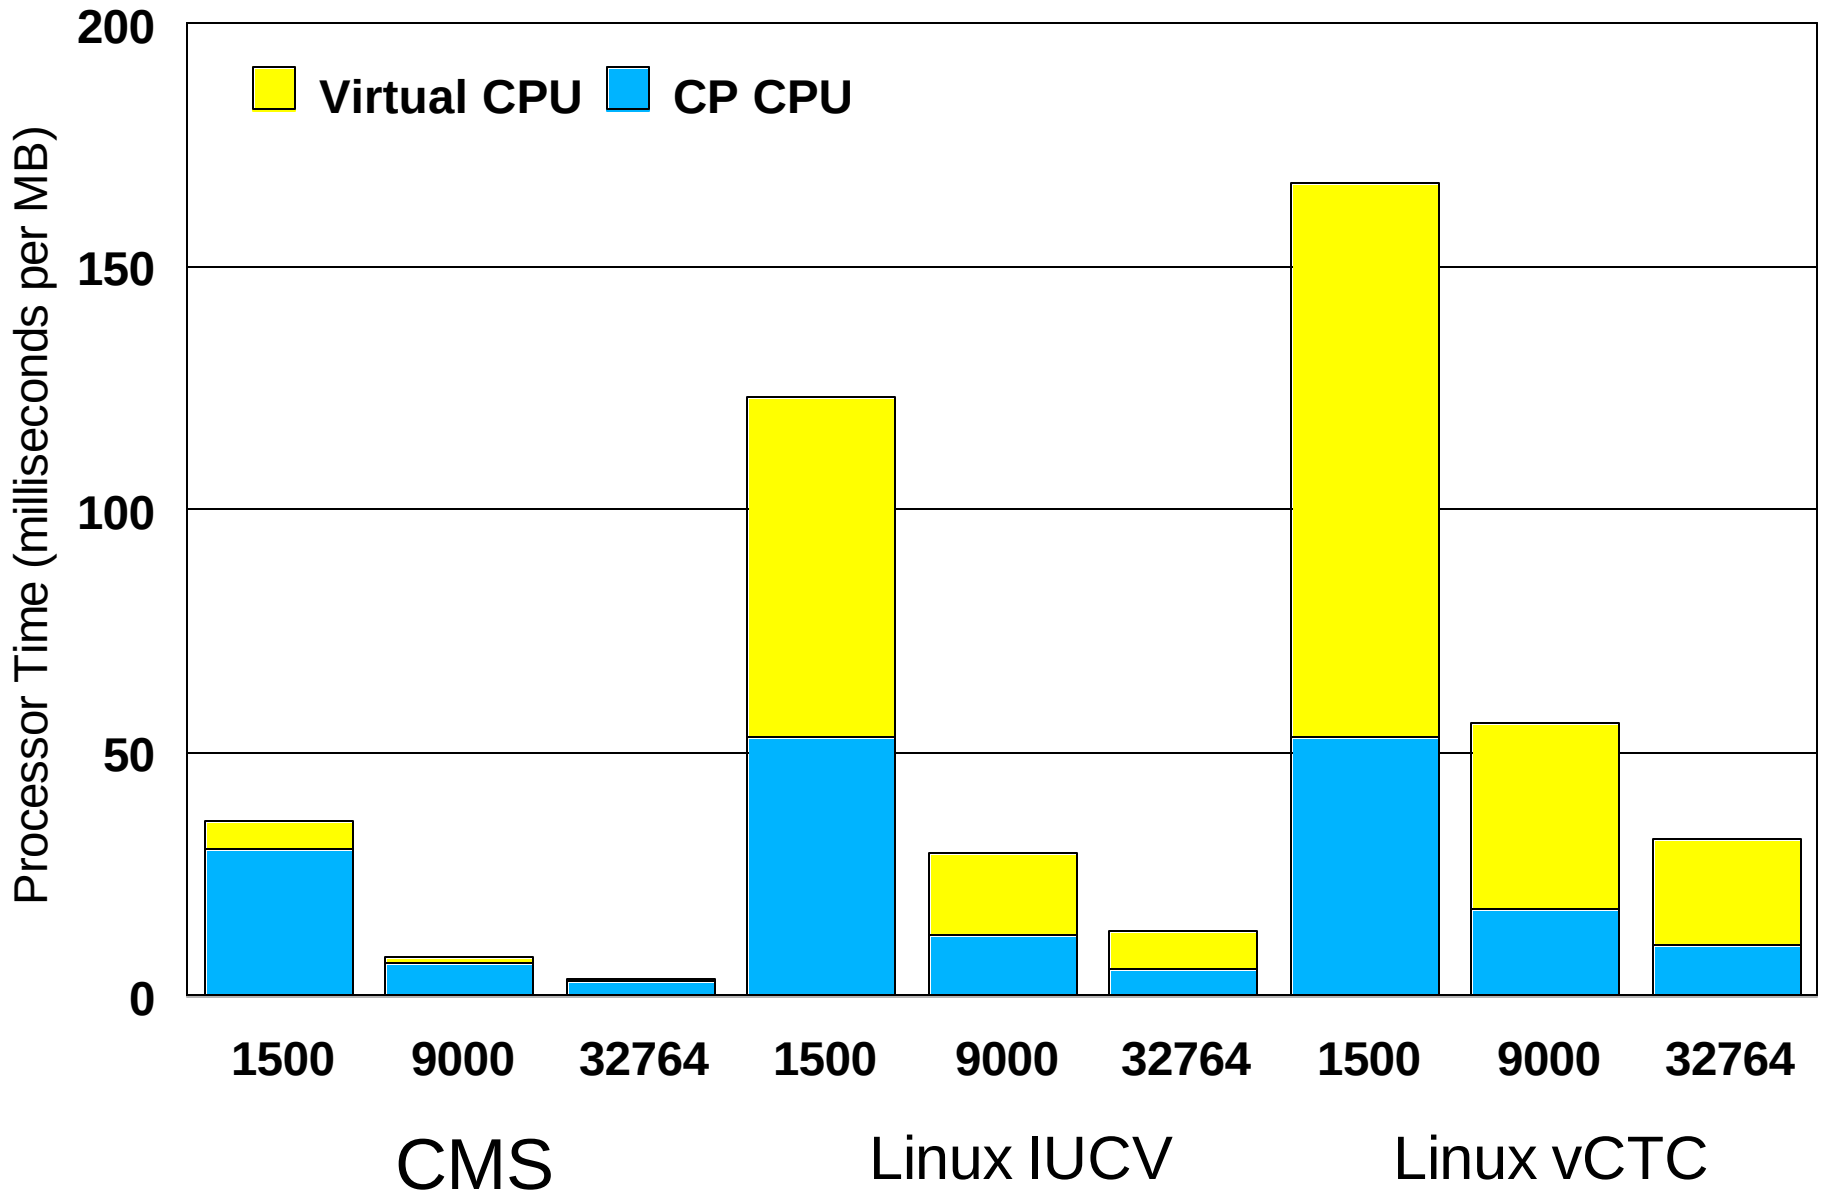

Some APARs of Interest

- VM62869 - corrects pages used for QDIO staying locked.
- VM62827 - corrects reorder frequency being too high.
- VM62794 - correct shared segment numbers

Summary

- Good regression story
- 64-bit
  - Lifts many aspects of the 2G limit
  - Still should have some expanded storage
  - Do not make V=R area too large
- Networking improvements
  - QDIO Gigabit Ethernet
  - SSL
- Full details in Performance Report on web
  - http://www.ibm.com/s390/vm/perf/