

# VM/ESA Performance Update

VM & VSE Tech Conference  
May 2000 - Orlando  
Session M70

Bill Bitner

VM Performance

607-752-6022

[bitner@vnet.ibm.com](mailto:bitner@vnet.ibm.com)

Last Updated: April 10, 2000

[RETURN TO INDEX](#)

# 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

The following are trademarks of Sun Microsystems:

Java, JDK

# Introduction

- VM/ESA 2.4.0 Available
  - ▶ Hardware Support
  - ▶ Scheduler Enhancements
  - ▶ TCP/IP
  - ▶ Incremental monitor improvements
- Other points of interest:
  - ▶ ADSM Version 3
  - ▶ Java
  - ▶ Connectivity Performance

# VM/ESA Regression

- CMS Regression V2.3.0 to V2.4.0
  - ▶ ITR decreased 0.4 to 0.6%.
  - ▶ Response time was equivalent.
- VSE
  - ▶ Equivalent performance
- TCP/IP
  - ▶ FTP equivalent to FL 310
    - VM FTP client "get" throughput improved 2%
  - ▶ Telnet equivalent to FL 310
  - ▶ NFS big improvements (see later charts)

# Support for FICON and Friends

- Capacity gains:
  - ▶ Bandwidth of 100 MB/Sec
- Additional Exploitation:
  - ▶ Synchronize Control extends current prefetching
    - Paging, Spooling, and Guest
  - ▶ Avoidance of nullification window
    - Requires Enterprise Storage Server (ESS) as well as FICON for support
- Should use Enhanced CPMF for channel utilization, not STCPS method.

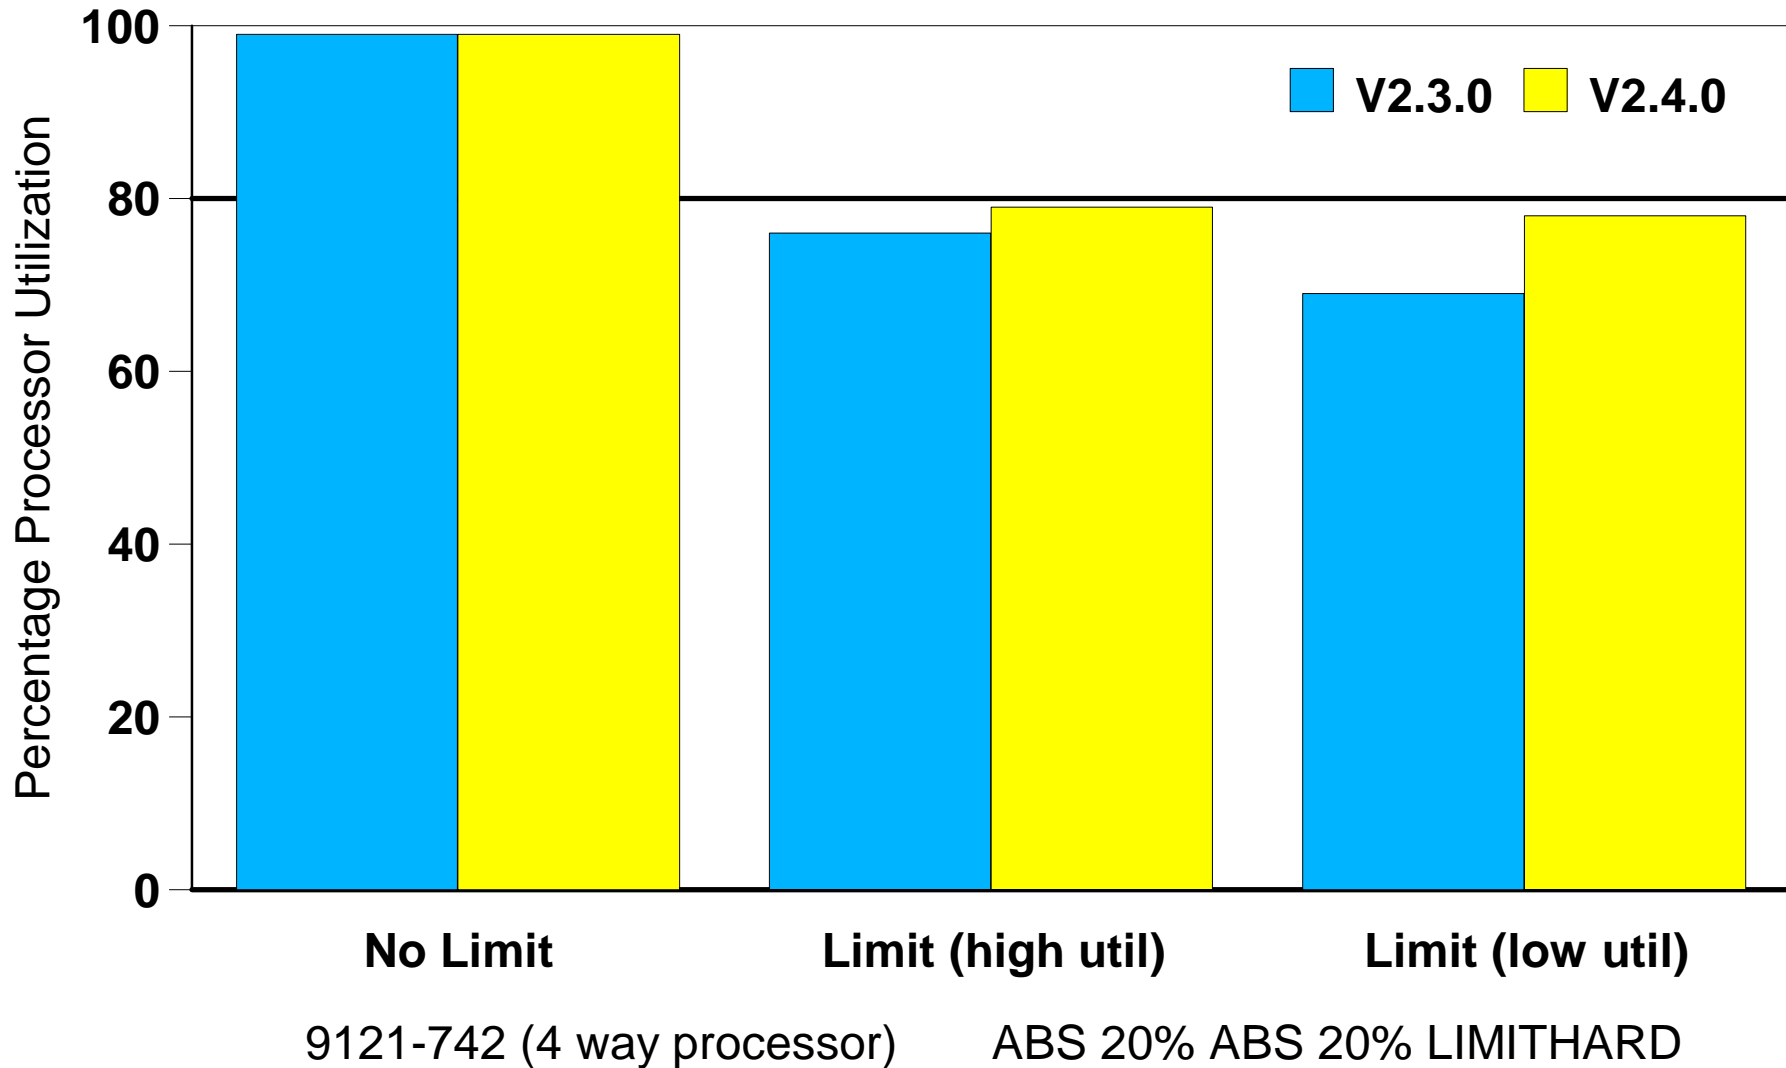
## 2.4.0 CP Monitor Changes

- Enhanced Channel Path Measurement
  - ▶ New System Domain Record (per channel)
- Indication of source of device active time (HW or SW)
- Synchronization of SCM block statistics
- CP return free storage requests now accurate
- Support for Parallel Access Volumes on ESS DASD.
- New I/O record for state change events such as PAV in ESS.

# Improved Limit Shares

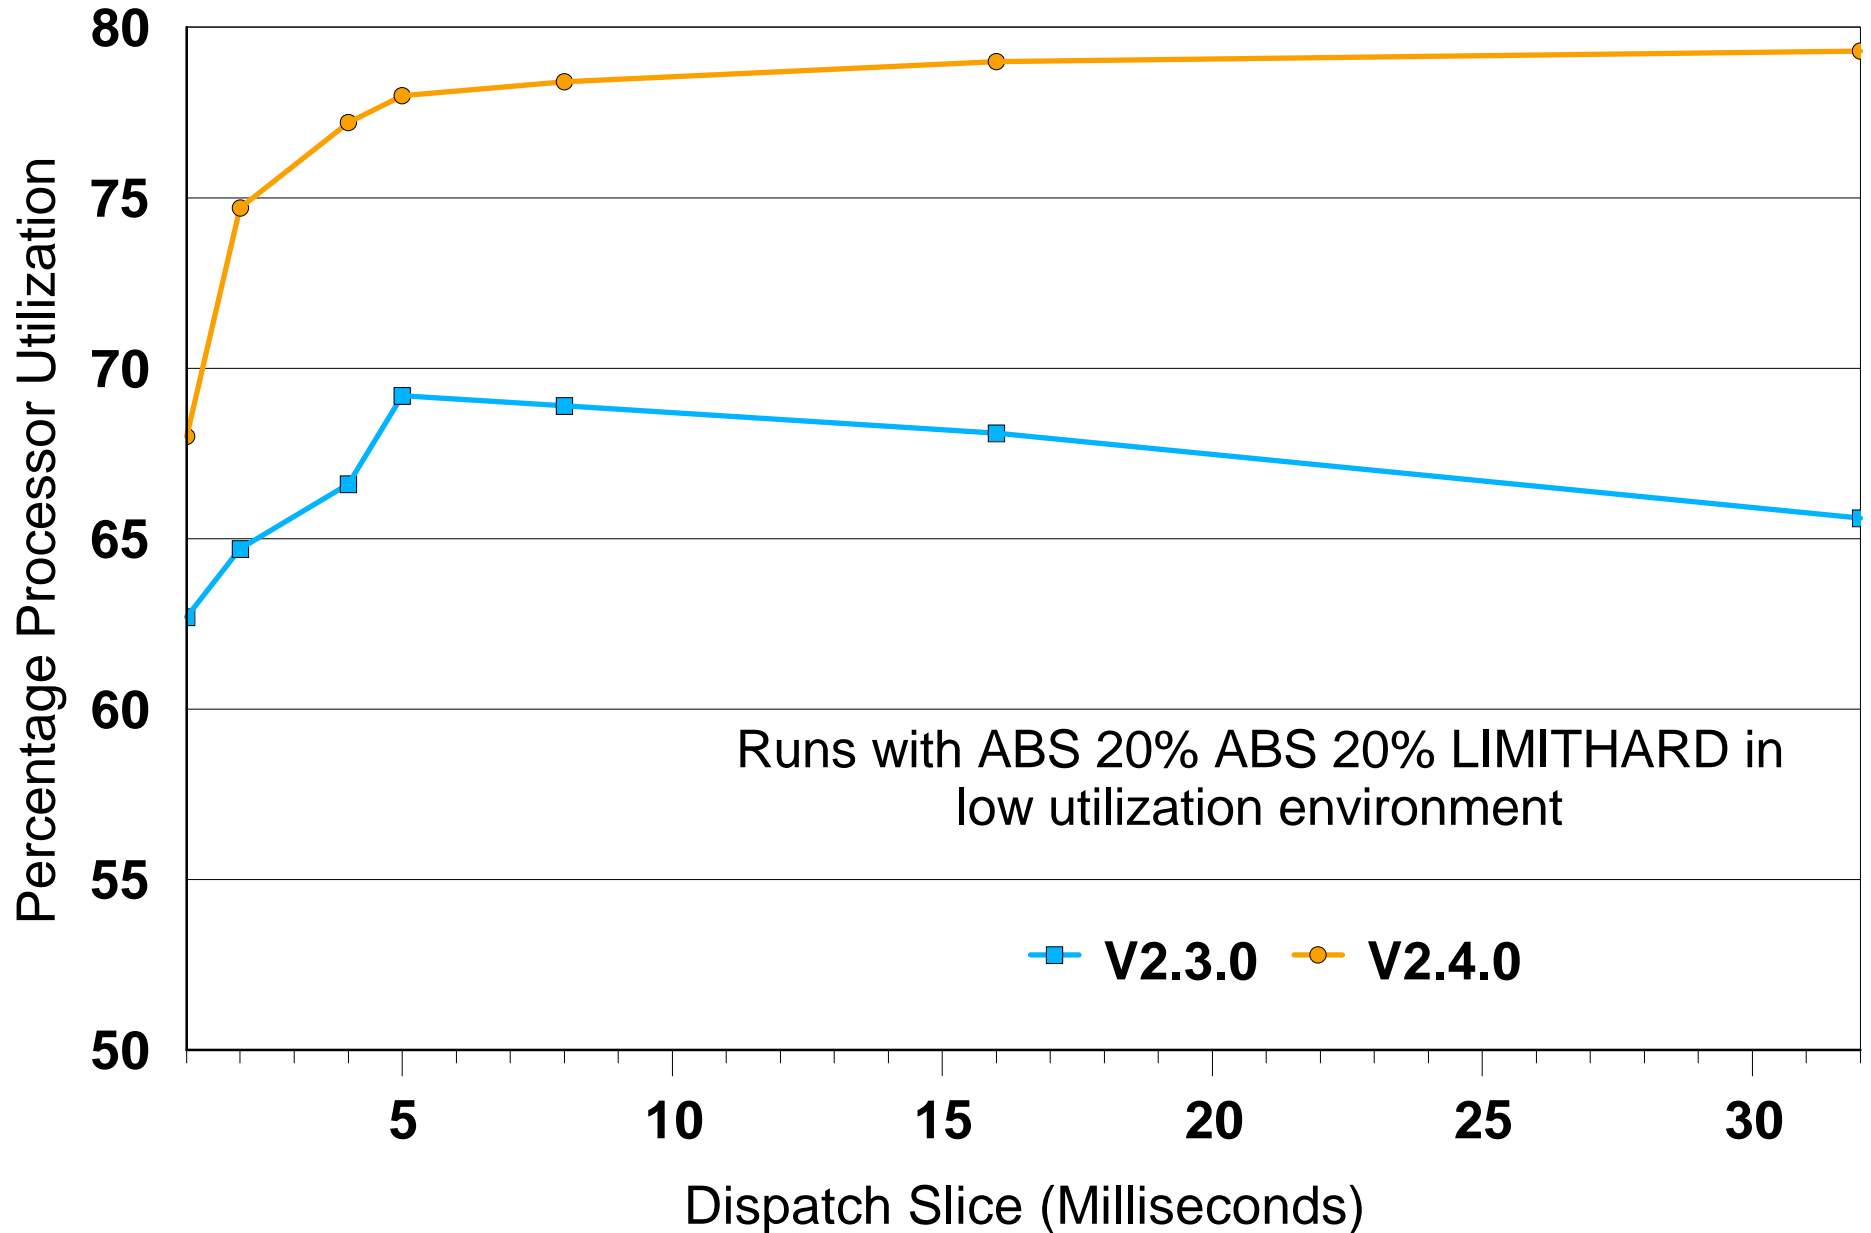
- VM/ESA 1.2.2 Introduced Limit Shares
- Two flavors:
  - ▶ LIMITHARD - limit regardless of capacity
  - ▶ LIMITSOFT - limit unless extra capacity exists
- Worked great... except in
  - ▶ Virtual MP environments
  - ▶ Low system utilization
- Some minor improvements through service stream
- FIN APAR VM61527 now in VM/ESA 2.4.0

# Problem Scenario

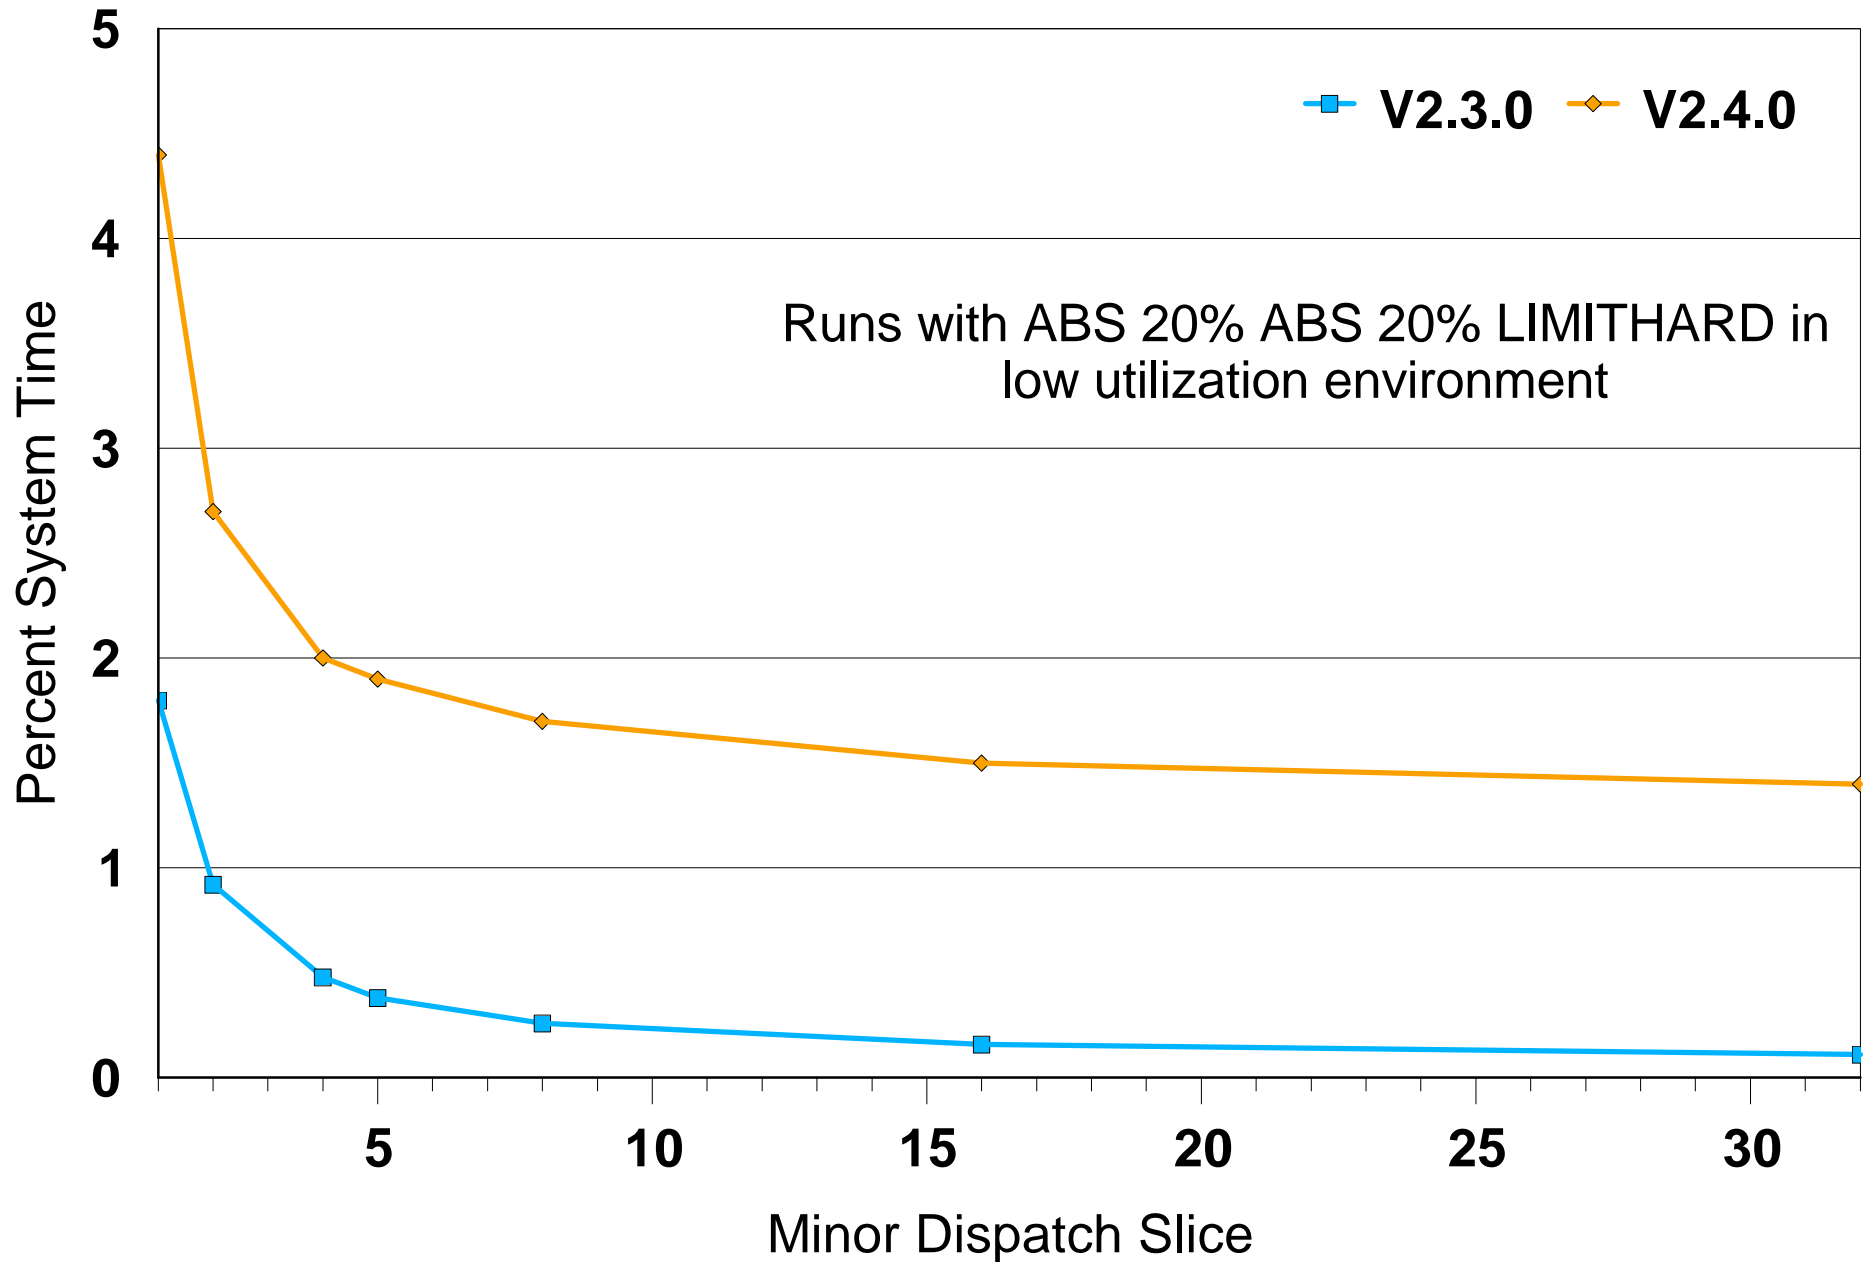




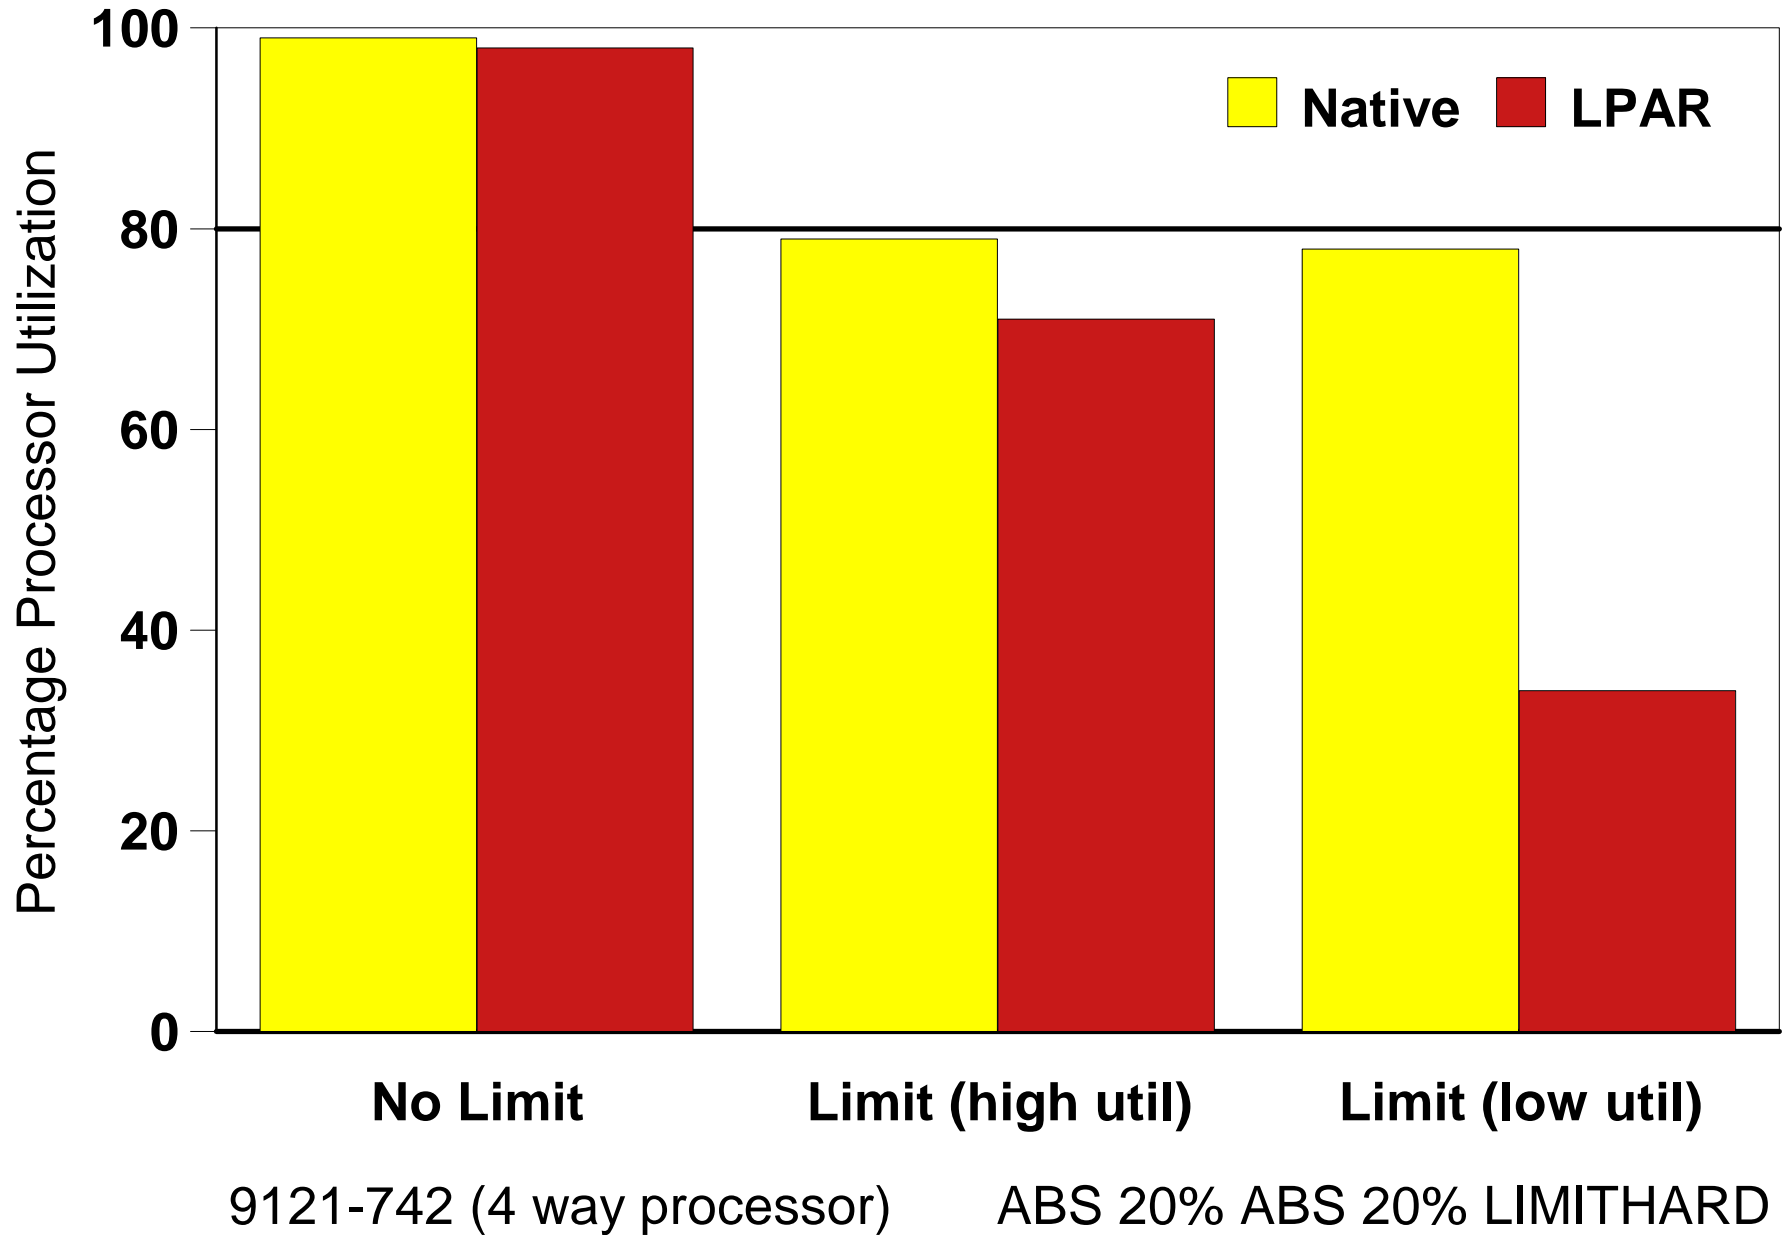
# Dispatch Slice and Limit Share



# Dispatch Slice and System Overhead



# VM/ESA 2.4.0 Native vs. LPAR Results



# Share Capping Summary

- Less restrictive while holding a LIMITHARD in native environments.
- IBM tests show LPAR environment tends to hold user below the LIMITHARD setting.
- One ESP customer LPAR environment shows user getting more than the LIMITHARD.
- Use LIMITHARD with care in an LPAR environment.

# SFS Performance Improvements

- Recent performance APARs rolled into VM/ESA 2.4.0
  - ▶ VM61547 - mitigate "lock out" scenario when deleting very large files (>512KB)
  - ▶ VM62008 - follow-on to VM61547
  - ▶ VM62086 - mitigate "lock out" scenario for long open-write-close nocommit sequences

# TCP/IP Improvements

- Feedback on TCP FL 310 with RFC 1323
  - ▶ OSA 2 Fast Ethernet saw a factor of 3 improvement
  - ▶ Unlike VM/ESA, some stacks default RFC1323 off.
- APAR PQ18391 - extends TCP Maximum Segment Size (MSS)
- FL 320: TCP Header prediction
  - ▶ Lower pathlengths for inbound processing

# TCP/IP Monitor Improvements


- APAR PQ16942 rolled into FL 320
  - ▶ Allow for recording of larger amounts of data on TCB and UCB close records
- FL 320 Changes:
  - ▶ count of packets discarded for LAND attack
  - ▶ count segment headers predicted correctly
  - ▶ TCP close record now includes window scaling factors and local IP address
  - ▶ UDP open/close records now created for sessions initiated through sockets interface

# NFS Improvements

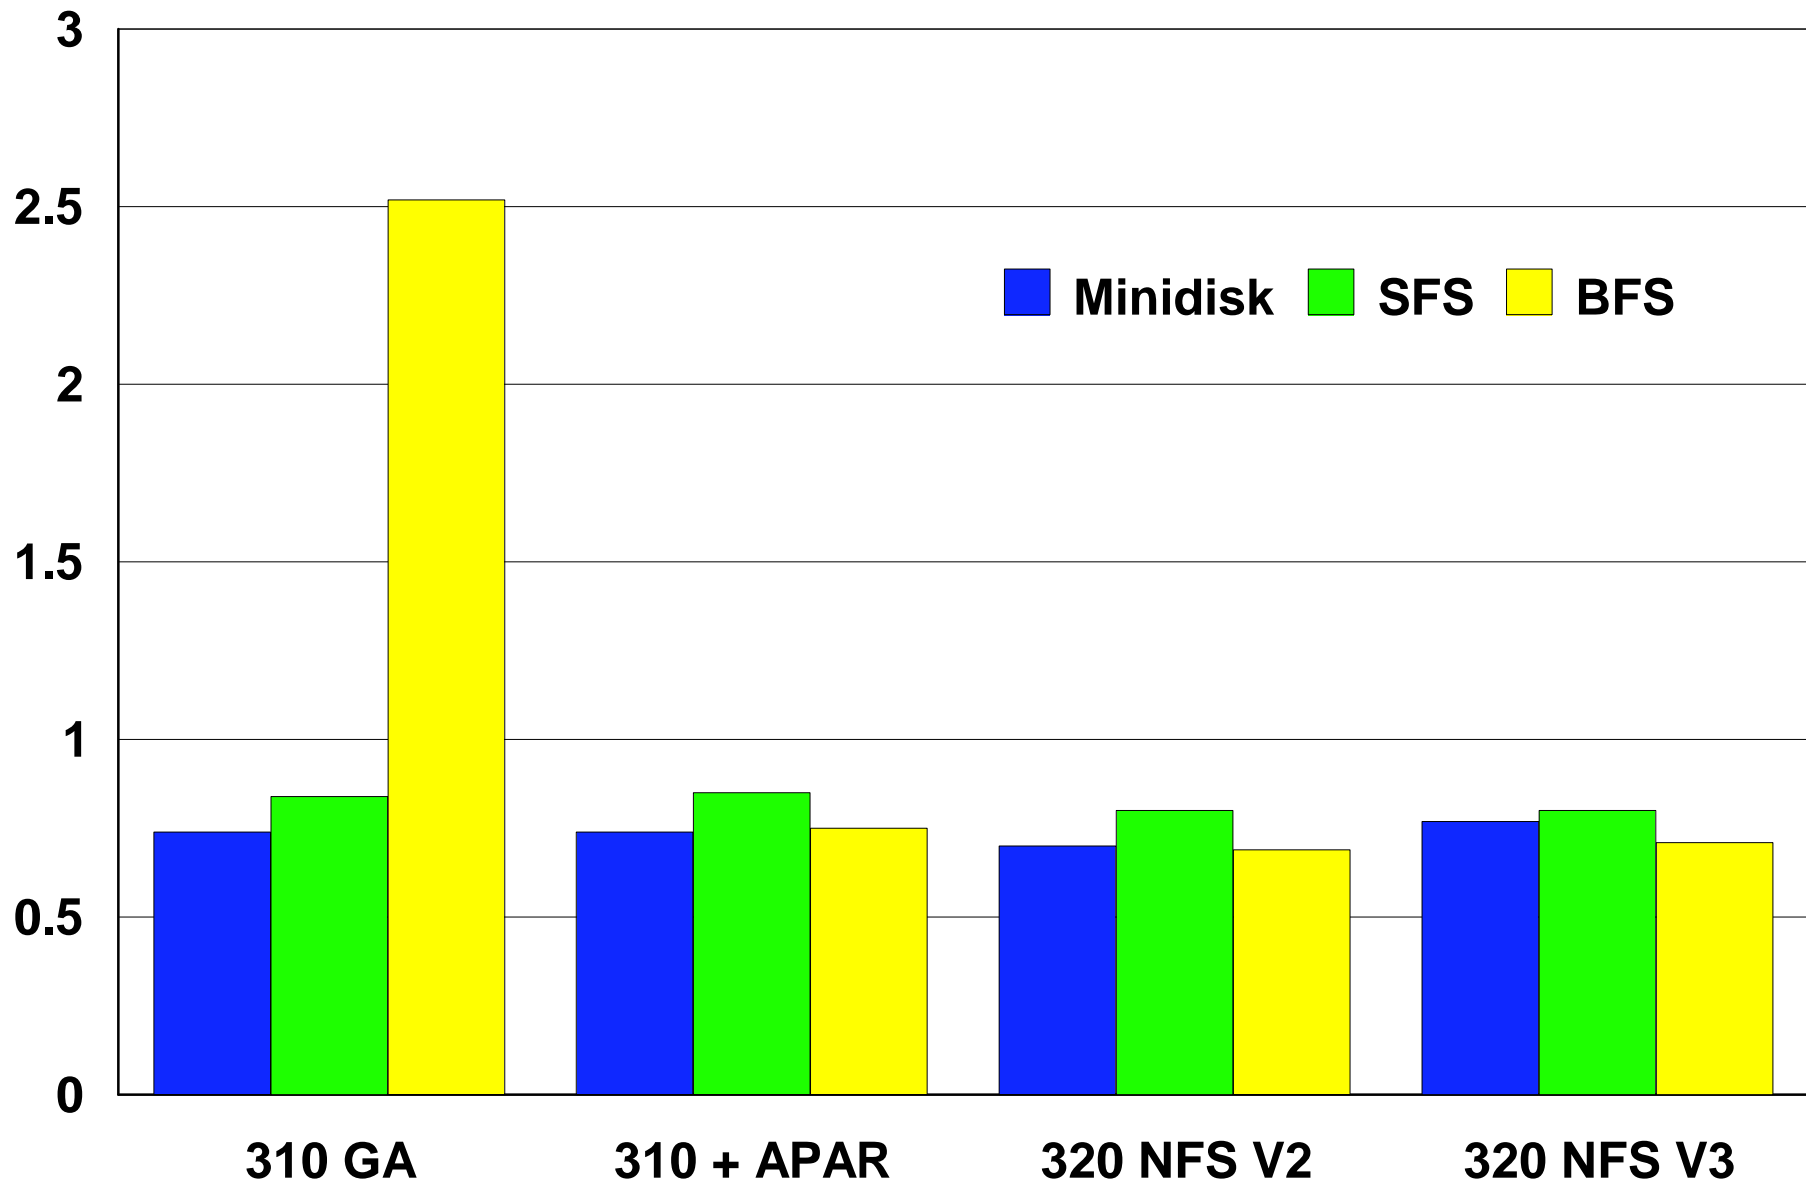
- TCP/IP FL 310:
  - ▶ APAR PQ16183 (helps BFS only)
    - improves reading large files
- TCP/IP FL 320:
  - ▶ NFS Version 3 Protocol
    - larger block sizes helps large file processing
    - REaddirPlus helps directory displays
  - ▶ allow TCP connections
  - ▶ Improvements to BFS interface




# NFS- BFS: Processor Time Breakdown



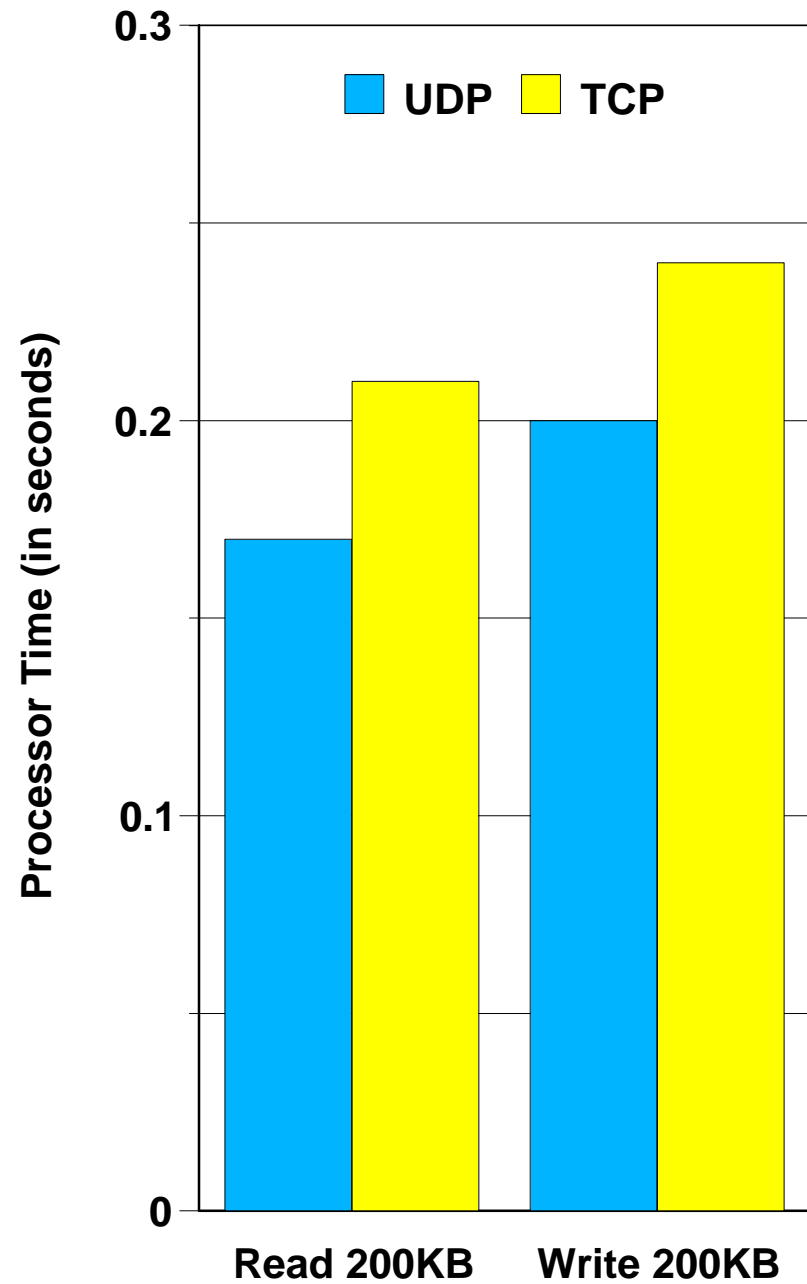

# NFS: Elapsed Time to read 200KB file



# NFS: CPU Time to read 200KB file



# NFS Support for TCP



# JAVA

- VM/ESA JDK 1.1.4 level performance challenges:
  - ▶ slow execution (no compiler)
  - ▶ large cost for initialization
  - ▶ while multithreading, it is not multiprocessing
- JIT (Just-In-Time) compiler
  - ▶ Greatly speeds up execution
    - up to 2.5 times improvement compared to no JIT in portable BOB workload (1 thread only)
    - Kernel benchmarks: 1 to >50 X faster
  - ▶ Does add a hit to initialization

# JAVA

- Java Initialization
  - ▶ Currently 5 seconds on a 19 MIPS/engine box
  - ▶ Improved in JDK 1.1.6 with new CMS 15.
- RAWT (Remote Abstract Windowing Toolkit)
  - ▶ Is available with JDK 1.1.6
  - ▶ RAWT not recommended for performance sensitive applications.
- CMS 15 Improves Java by eliminating:
  - ▶ Large number of ThreadCreates
  - ▶ Redundant BFS QueueOpen calls
  - ▶ Thousands of CMSSTOR calls

# ADSM Version 3


- Much better performance than version 2
- Backup throughput improvements
  - ▶ lower processor and DASD I/O requirements
  - ▶ smaller files saw larger improvement due server file aggregation item
  - ▶ measured throughput improvements of 9% to 126%
- Restore throughput showed little change
  - ▶ lower processor and DASD I/O requirements
  - ▶ restore throughput much lower than backup

# ADSM Measurement Config

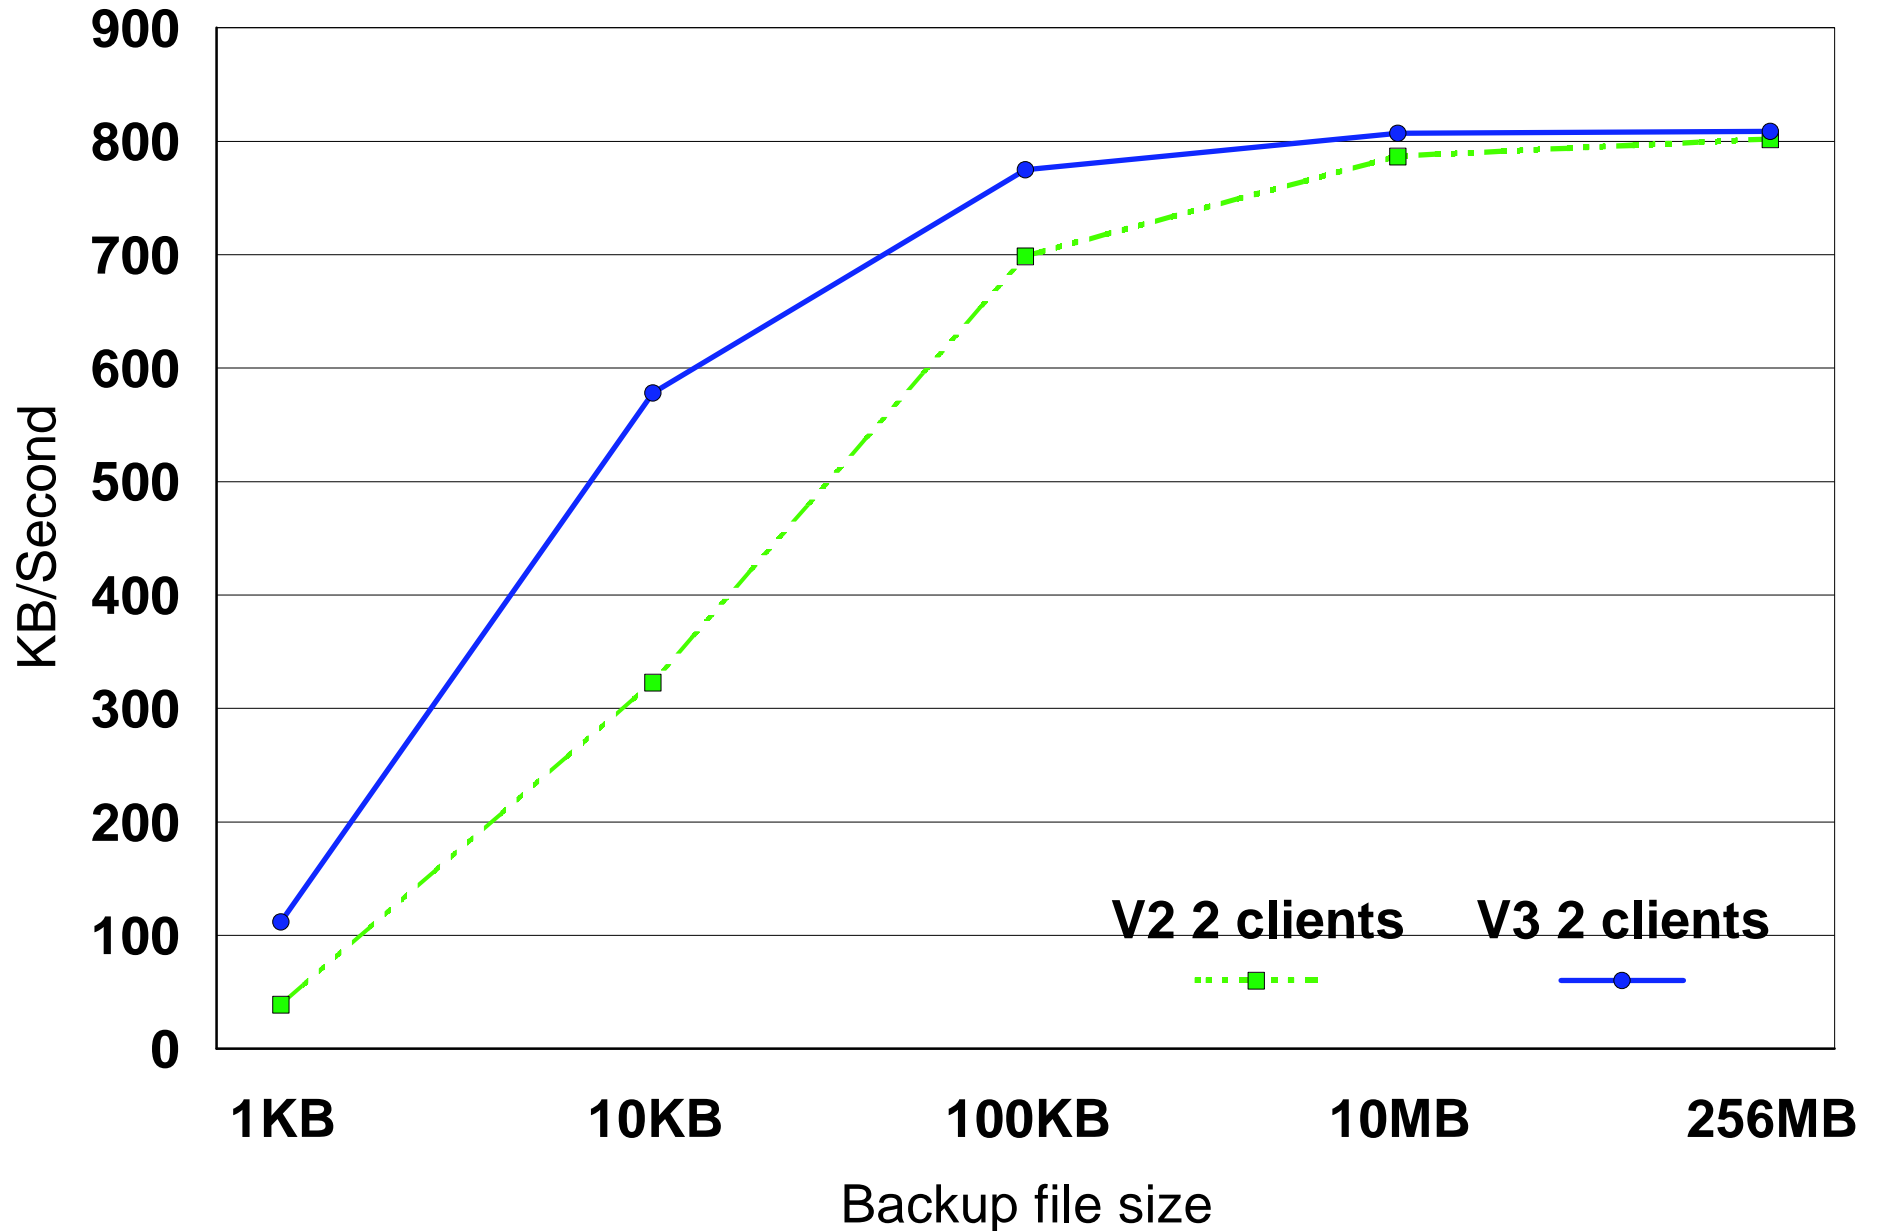
- Server:
  - VM/ESA 2.3.0 on 9121-480
- Clients:
  - AIX 4.1.4 on RS/6000 model 250
  - ADSM clients: version matched server
- Connection:
  - 16 Mbit IBM Token Ring
  - VM connected via 3172-3
  - TCP/IP FL310




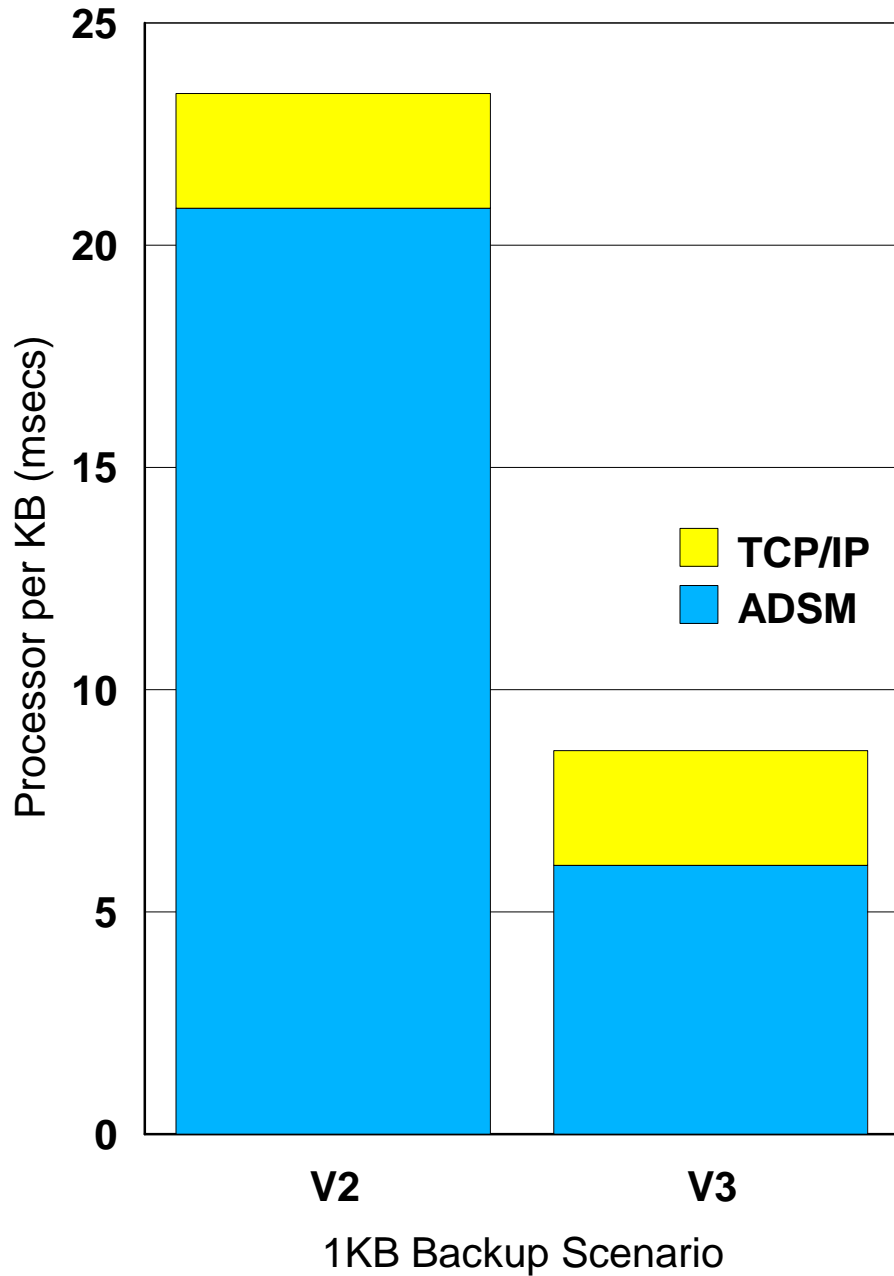
# ADSM Backup Throughput




# ADSM Backup Throughput



# ADSM Single Client Backup




# Single Client Backups DASD I/O



	1KB	10KB	100KB	10MB	256MB
V2	1.14	0.12	0.05	0.03	0.03
V3	0.16	0.02	0.01	0.01	0.01

# ADSM RestoreThroughput



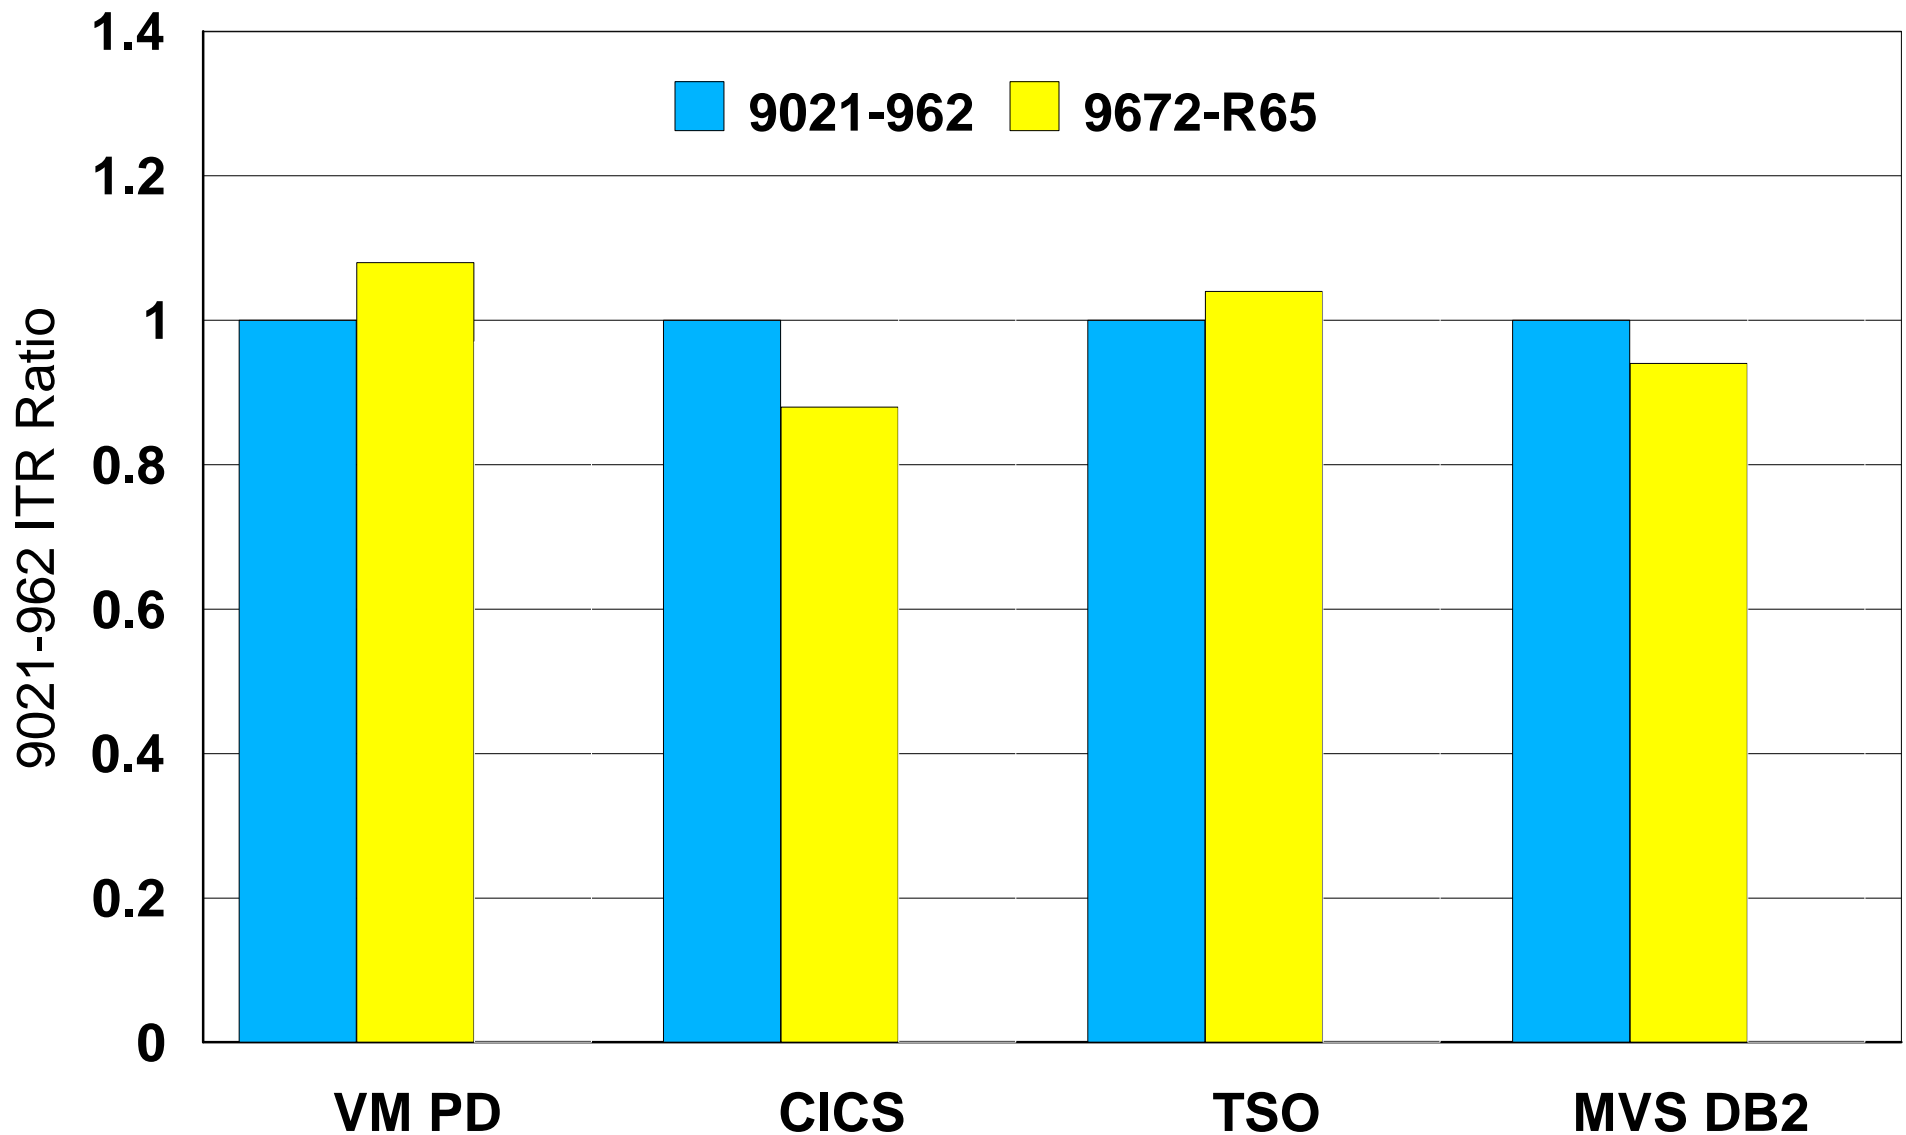
# Dirmaint APAR VM62262

- Performance Improvement APAR
- Avoids need to reopen/reclose files
- Range of change for key indicators:
  - ▶ Elapsed time 0 to -46%
  - ▶ Virtual I/O 0 to -97%
  - ▶ Virtual processor time 1 to -24%

# G5/G6 Performance Improvements

- First  $>1$  BIPS machine
  - ▶ BIPS = billion instructions per second
  - ▶ BIP = Baffling Indicator of Performance
- Improvements
  - ▶ decrease memory access costs
  - ▶ improved processor caching
- Much better performance for workloads with poor locality of reference and/or very short transactions


# Prior to G5 Sizing Advice



VM PD may be optimistic. Check MVS workloads.



# New 9672 Sizing Advice

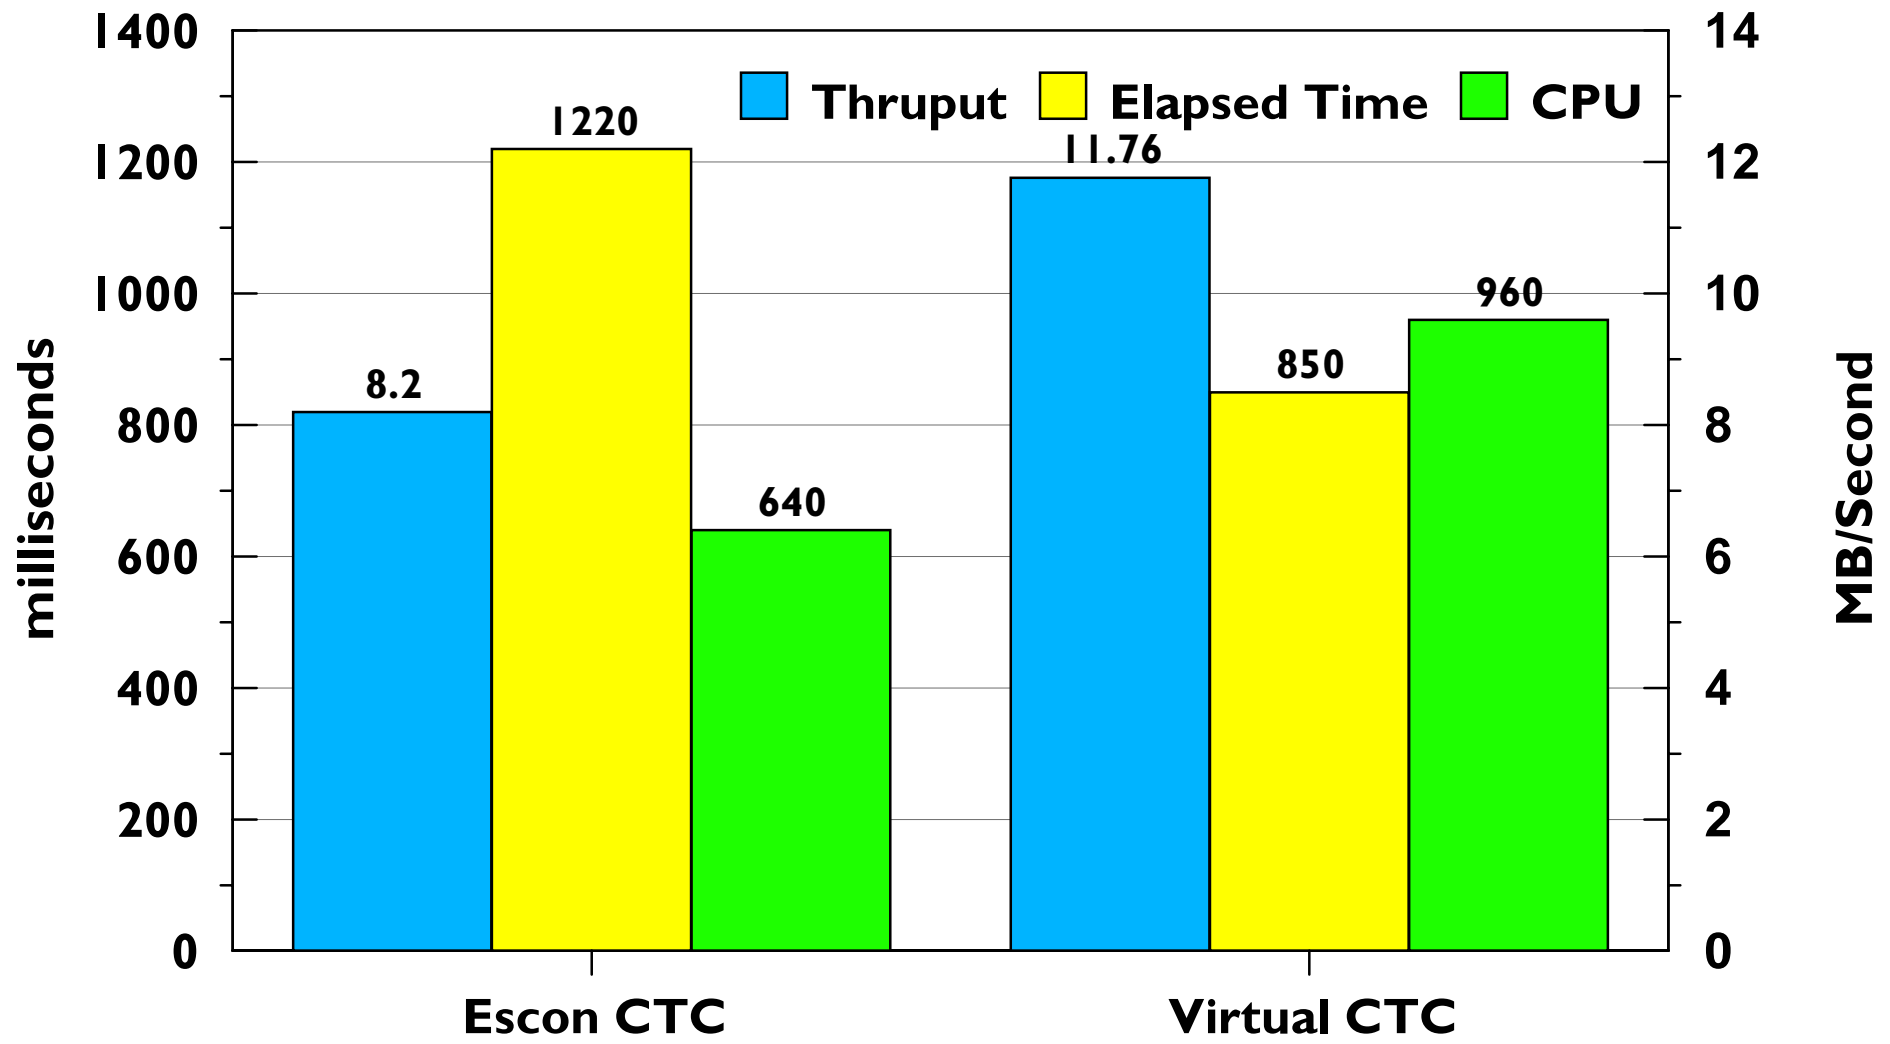


**Check all workloads to validate worse case scenario.**

# LSPR Workload Changes


- Large System Performance Reference
  - <http://www.ibm.com/s390/lSpr/>
- Helps for capacity planning by providing a series of workloads across operating systems for various processors
- Workload Enhancements
  - ▶ TPNS replaces FSID
  - ▶ Allow for measuring larger systems
  - ▶ Make more realistic
    - Server virtual machines
    - Virtual disk in storage
    - Program products

# Connectivity Measurements



FTP 10MB files between 2 V=V Guests on 9672-R55.

# Improved Virtual CTC



FTP 10MB files between 2 V=V Guests on 9672-R55.

# Summary

- VM/ESA Development team continues to keep an eye on performance
- Full VM/ESA 2.4.0 Performance Report  
<http://www.ibm.com/s390/vm/perf/docs/>
- Wider scope than traditional "regression" CMS performance
- For news, keep checking:  
<http://www.ibm.com/s390/vm/perf/>