

# IMS Version 8 Performance Summary

Author: Mike Gonzales

Contributors: Cedric Chen, Hiram Neal, Bovorit Pibulsonggram, Hsiung Tang

November 2002

## **Abstract**

IBM has announced the eighth version of its premier transaction/database server; IMS Version 8. IMS Version 8 offers features to enhance sysplex and non-sysplex processing to increase the performance and availability of IMS and IMS data. This evaluation demonstrates the impact of new and changed code. As a result of hardware availability, this evaluation also demonstrates the benefit to IMS performance of hardware enhancements in the Enterprise Storage Server model M800 turbo. A final study was performed to examine a high stress Fast Path environment. The high stress study demonstrates the ability of IMS Fast Path to manage ultra-high transaction volumes with updates.

This paper summarizes the performance results obtained at the IBM Silicon Valley Laboratory for purposes of IMS migration and capacity planning.

© IBM 2002  
Software Group  
San Jose

<b>1 Introduction</b>	4
<b>2 Executive Overview</b>	5
<b>3 Base Performance</b>	6
3.1 Introduction	6
3.2 Summary	7
3.3 Base Test Environment	7
3.4 Results	8
<b>4 CSA Virtual Storage Constraint Relief</b>	13
4.1 Introduction	13
4.2 Summary	13
4.3 Test Environment	13
4.4 Results	14
<b>5 Parallel Database Open</b>	16
5.1 Introduction	16
5.2 Summary	16
5.3 Test Environment	16
5.4 Results	17
<b>6 Common Service Layer</b>	22
6.1 Introduction	22
6.2 Summary	22
6.3 Test Environment	22
6.4 Results	23
<b>7 Fast Path Greater-Than-240 Areas Support</b>	24
7.1 Introduction	24
7.2 Summary	24
7.3 Test Environment	24
7.4 Results	25
<b>8 Coupling Facility Structure Duplexing</b>	26
8.1 Introduction	26
8.2 Summary	26
8.3 Test Environment	26
8.4 Results	27
<b>9 IMS Java Message Processing Region</b>	29
9.1 Introduction	29
9.2 Summary	29
9.3 Test Environment	29
9.4 Results	30
<b>10 Fast Path High Stress</b>	32
10.1 Introduction	32
10.2 Summary	32
10.3 Test Environment	32
10.4 Results	33
<b>11 ESS M800 Turbo with IMS Version 8</b>	35

11.1	Introduction .....	35
11.2	Summary .....	35
11.3	Test Environment .....	36
11.4	Results .....	37
<b>12</b>	<b>Test Methodology</b> .....	<b>43</b>
12.1	Disclaimer .....	44
12.2	Trademarks .....	44

# 1 Introduction

IMS Version 8 is the latest addition to the IMS family of products offering performance, availability, and scalability enhancements. These are the significant enhancements evaluated in IMS Version 8:

- CSA Virtual Storage Constraint Relief
- Parallel Database Open
- Common Service Layer
- Fast Path Greater-Than-240 Areas Support
- Coupling Facility Structure Duplexing
- IMS Java Message Processing Region

The evaluation of IMS Version 8 includes a study of the base transaction processing functions as well as studies of enhancements. The results demonstrate that enhancements made in IMS Version 8 improve the performance and availability of IMS processes and data. No variance in device configuration was used to ensure the specific software performance characteristics comparisons were accurate. Where minor variances in processor utilization were observed, results were extrapolated for purposes of comparison.

In addition to the enhancements to IMS Version 8, the ESS M800 Turbo and a 1Gbit FICON channel is included in a few test scenarios to show the potential of the IMS logging performance. This paper includes some additional comparisons with an ESS F20 to demonstrate the extra performance benefits of the ESS M800 Turbo.

This paper discusses the performance characteristics of IMS Version 8 relative to IMS Version 7 with the following products:

- z/OS V1R2
- OS/390 V2R10
- zSeries 900 model 2064

Note: All tests in this paper are executed on zSeries 900 model 2064 using a S390 G6 model 9672 to drive TPNS. With the exception of the ESS M800 and Fast Path High Stress studies, all data sets were on ESS F20.

## 2 Executive Overview

Overall, IMS Version 8 performs better or scales up larger than IMS Version 7 for each of the following areas:

### CSA Virtual Storage Constraint Relief

**IMS Version 8 reduces CSA usage below the 16 MB line by up to 47%** in full function and Fast Path environments when compared to IMS Version 7.

### Parallel Database Open

**IMS Version 8 restarts over 50% faster** in normal restarts and almost **40% faster in emergency restarts** when compared to IMS Version 7. **IMS Version 8 shuts down nearly 20% faster** when compared to IMS Version 7. **IMS Version 8 opens databases 5% faster** when compared to IMS Version 7.

### Fast Path Greater-Than-240 Areas Support

**IMS Version 8 maximum Data Entry Database sizes increase from <1 terabytes to 8 terabytes** when compared to IMS Version 7. **IMS Version 8 increases the maximum number of areas allowed in a single DEDB by 750%, from 240 to 2048** when compared to IMS Version 7.

### Coupling Facility Structure Duplexing

**IMS Version 8 CF Duplexing facilitates uninterrupted IMS transaction processing in the event of the loss of a coupling facility or loss of a duplexed structure.**

### IMS Logging

**IMS Version 8 logging when OLDS are on ESS M800 connected with FICON channels can reach a logging bandwidth of 27 MB per second. IMS Version 8 exhibits improved WADS I/O response times using only a single track on ESS M800. WADS I/O response times with ESS M800 are less than 1 millisecond.**

### IMS Java Message Processing Region

**IMS Version 8 Java Message Processing (JMP) regions allow a Java application program to run under a Java Virtual Machine in IMS using the new Persistent Reusable Java Virtual Machine which replaces the High Performance Java used in IMS Version 7. A Java application program using the IMS Version 8 JMP support was used to reach a steady throughput rate of 2134 transactions per second.**

## 3 Base Performance

### 3.1 Introduction

This section shows the performance characteristics of IMS Version 8 when compared to IMS Version 7 in a variety of workloads used to evaluate the basic transaction processing and database access functions of IMS. IMS Version 8 is rated at 323,000 lines of code which is a considerable increase in total IMS lines of code.

The following workloads were used for the studies described herein

:

The Data Systems Workload HALDB (DSWH) is a mix of Full Function Database with High Availability Large Database (96 volumes) using heavy database update processing with transaction sets performing the following functions: order entry, stock control, hotel reservation, inventory tracking, banking and teller system. This workload uses small messaging for transaction responses.

The Debit-Credit Workload (DCW) uses Fast Path Databases (120 areas) and is designed to produce large volumes of log data (24K per transaction).

The Data Sharing, Full Function (DSFF) workload is a mix of Full Function Database (384 volumes) using medium database update processing with transaction sets performing the following functions: order entry, stock control, hotel reservation, inventory tracking, banking and teller system. This workload uses large messaging for transaction responses and is similar to the workload used to determine the IMS ITR ratings for the IBM LSPR. This workload uses IRLM for the data sharing lock manager. IRLM is used in 'global' mode.

The Banking BMP workload is an imported customer workload which performs heavy update activity in sequential access of a Fast Path database. This workload simulates end-of-day account reconciliation.

The Fast Path Two (FP2) workload simulates a credit card processing system with lean transactions, small messaging and database updates for each transaction. This workload uses an account data entry database (DEDB) which is split into 86, 240, 480, or 960 area data sets.

All IMS performance statistics are gathered during a steady state.

### 3.2 Summary

The following summarizes the performance results of the IMS Version 8 base functions relative to IMS Version 7. Refer to Section 3.4 for detailed results.

IMS Version 8 demonstrates the following performance impact when compared to IMS Version 7. Note, a variance of not more than 5% is generally deemed equivalent between measurements, however, IMS Version 8 targets degradation not to exceed 3% in any base function. The resultant overhead observed in IMS Version 8 is as follows for the following functions:

- Full Function transaction processing (with or without HALDB): 1.6%
- Full Function transaction processing with Shared Message Queues: 2.5%
- Full Function transaction processing from APPC: 2.6%
- Fast Path transaction processing: 2.2%
- Fast Path transaction processing from APPC: 2.6%
- Batch Message Processing : -0.4% (reduced elapsed time)

The CPU cost of migrating a full function transaction processing environment from IMS Version 7 to IMS Version 8 is less than the CPU cost of migrating from IMS Version 6 to IMS Version 7. This is significant in that IMS Version 8 growth in KLOC is about twice the growth of IMS Version 7 in KLOC.

### 3.3 Base Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 20GB storage, 6 CPs 2064-216, 2 x ICF, 2 CPs per ICF (for shared queues test)
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 ESCON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 16GB cache
<b>Operating Systems:</b>	z/OS DFSMS V1R2, OS/390 V2R10, IMS V7.1, IMS V8.1
<b>Other:</b>	3 PAVs per real volume

### 3.4 Results

Figures 3.1 - 3.5 summarize performance statistics of IMS Version 8 when compared to IMS Version 7. The charts show the following statistics:

- Full Function transaction processing rates - HALDB only
- Full Function Shared Message Queues transaction processing rates
- Fast Path transaction processing rates
- APPC transaction processing rates
- BMP elapsed times and EXCP counts

Figure 3.1 shows the IMS transaction rates achieved when the CPU utilization was at a steady state of 70% busy while running the DSWH workload. A similar result was obtained when HALDB was not included in the test. The lack of variance in the HALDB and non-HALDB tests was expected based on prior studies showing HALDB incurs less than 1% overhead when compared to standard full function database access. The statistics show IMS Version 8 Full Function transaction processing capacity is nearly equivalent to IMS Version 7 with a variance of only 1.6%..

**Figure 3.1: Full Function transaction processing - Throughput comparison**

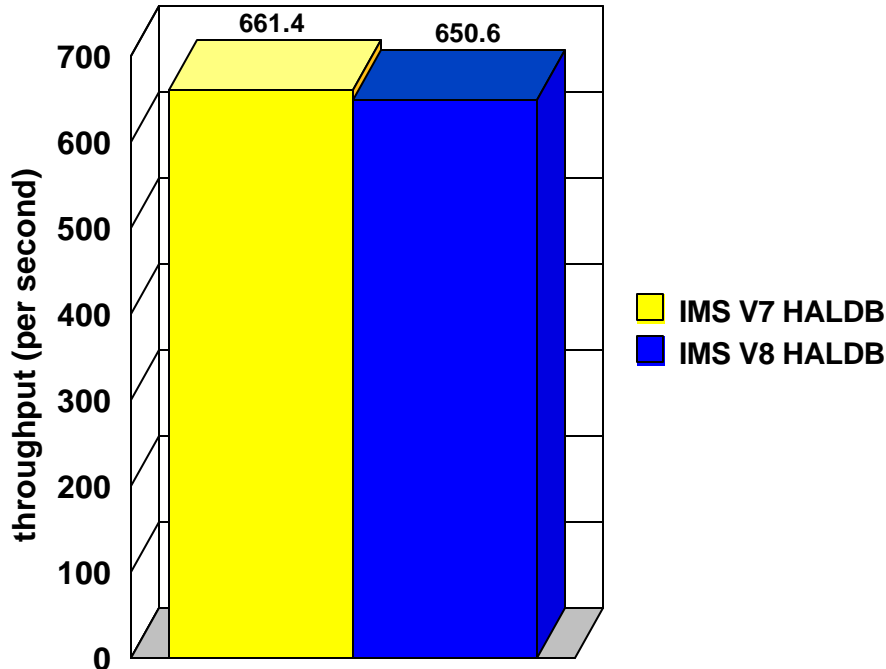






Figure 3.2 shows the IMS transaction rates achieved when the CPU utilization was at a steady state while running the DSFF workload using shared queues. The statistics show the transaction rates achieved in IMS Version 8 relative to IMS Version 7. The statistics show IMS Version 8 Full Function transaction processing capacity is nearly equivalent to IMS Version 7 with a variance of 2.5%.

**Figure 3.2: Full Function transaction processing in SMQ - Throughput comparison**




Figure 3.3 shows the IMS transaction rates achieved when the CPU utilization was at a steady state of 70% busy while running the DCW workload. The statistics show the transaction rates achieved in IMS Version 8 relative to IMS Version 7. The statistics show IMS Version 8 Fast Path transaction processing capacity is nearly equivalent to IMS Version 7 with a variance of only 2.2%.

**Figure 3.3: Fast Path transaction processing - Throughput comparison**




Figure 3.4 shows the IMS transaction rates achieved when the CPU utilization was at a steady state of 70% busy while running the DSFF and FP2 workloads from APPC. The statistics show the transaction rates achieved in IMS Version 8 relative to IMS Version 7. The statistics show IMS Version 8 Full Function transaction processing capacity through APPC is nearly equivalent to IMS Version 7 with a variance of only 2.6%. The statistics show IMS Version 8 Fast Path transaction processing capacity through APPC is nearly equivalent to IMS Version 7 with a variance of only 2.6%.

**Figure 3.4: APPC transaction processing - Throughput comparison**

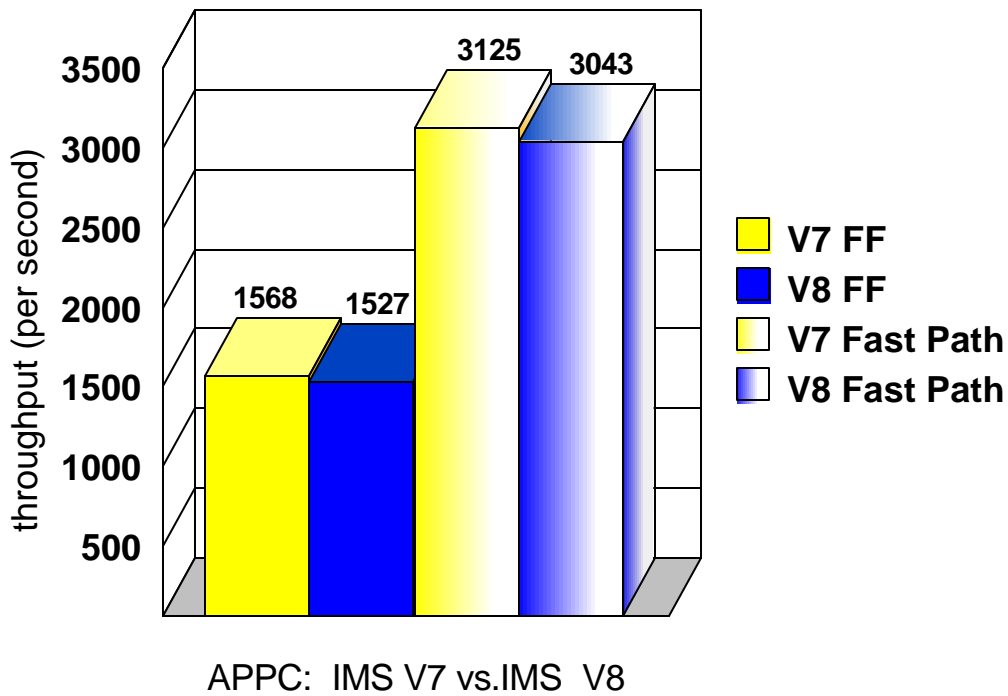
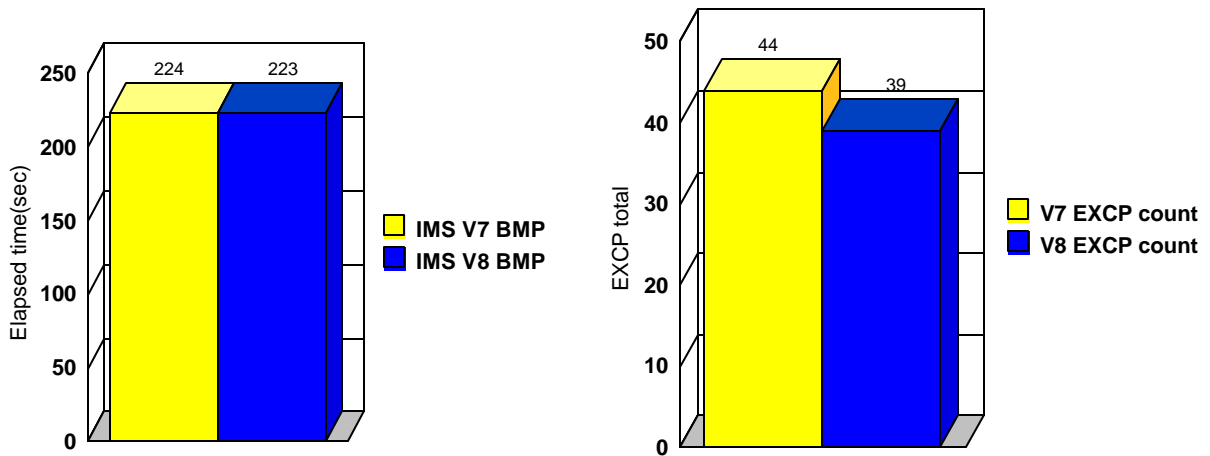


Figure 3.5 shows the results of the BMP banking account reconciliation job. The statistics show the elapsed times and EXCP counts incurred by IMS Version 8 relative to IMS Version 7. The statistics show IMS Version 8 Batch Message Processing elapsed time is reduced and is nearly equivalent to IMS Version 7 with a variance of only 0.4%.

**Figure 3.5: BMP processing - Elapsed time and EXCP count comparisons**



## 4 CSA Virtual Storage Constraint Relief

### 4.1 Introduction

This section shows the reduction in CSA usage below the 16 megabyte line by IMS Version 8 when compared to IMS Version 7.

The following results were obtained during the base testing of IMS Version 8 which is described in Section 3.1 of this document.

### 4.2 Summary

The following summarizes the results of IMS Version 8 CSA usage relative to IMS Version 7. Refer to Section 4.4 for detailed results.

The IMS Version 8 **CSA VSCR** provides the following improvements when compared to IMS Version 7. Note, the amount of CSA used depends on the number of active system tasks and the number of log buffers defined.

- Average CSA reduction improves 36%
- Minimum CSA reduction is 172 Kbytes

### 4.3 Test Environment

The description of the test environment is as follows:

<b>Hardware and Software Environment</b>	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 10GB storage, 3 CPs
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 ESCON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 16GB cache
<b>Operating Systems:</b>	z/OS DFSMS V1R2, OS/390 V2R10, IMS V7.1, IMS V8.1
<b>Other:</b>	3 PAVs per real volume

## 4.4 Results

Figures 4.1 - 4.2 summarize the CSA utilization statistics of IMS Version 8 in a variety of processing configurations. The charts shows the following statistics for the comparison of IMS Version 8 vs IMS Version 7:

- Full function processing
- Fast Path processing
- Batch message processing
- APPC processing

Figure 4.1 shows the reduction in CSA usage below the 16 MB line by IMS Version 8 relative to IMS Version 7. The statistics show IMS Version 8 uses significantly less CSA than IMS Version 7 in all tested configurations.

### **Figure 4.1: CSA VSCR - CSA utilization comparison**







Figure 4.2 shows the reduction in CSA usage below the 16 MB line by IMS Version 8 relative to IMS Version 7 in an APPC environment. The statistics show IMS Version 8 uses less CSA than IMS Version 7 in both Full Function and Fast Path environments when driven through APPC.

**Figure 4.2: CSA VSCR - CSA utilization comparison**



## 5 Parallel Database Open

### 5.1 Introduction

This section shows the performance improvements attributed to the Parallel Database Open feature of IMS Version 8. The Parallel Database Open feature creates 10 threads to process database open/close/EOV tasks.

### 5.2 Summary

The following summarizes the performance results obtained with IMS Version 8 relative to IMS Version 7. Refer to Section 5.4 for detailed results.

The database open times were compared using a batch message processing region to open all databases. The identical BMP was executed under IMS Version 8 and IMS Version 7 with all involved databases restored to their pristine state prior to each test.

IMS Version 8 **Parallel Database Open** provides the following improvements when compared to IMS Version 7:

- Normal restart times reduced by up to 56%
- Emergency restart times reduced by up to 39%
- Shutdown (/CHE FREEZE) times reduced by up to 19%
- DUMPQ shutdown times reduced by up to 13%
- Database open times reduced by up to 5%

Note, these comparisons are based on reaching a 'ready for work' state which IMS Version 8 achieves upon restart compared to IMS Version 7 which requires additional processing to initiate the database open process.

### 5.3 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 10GB storage, 3 CPs
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 ESCON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 16GB cache
<b>Operating</b>	z/OS DFSMS V1R2, IMS V7.1, IMS V8.1

**Systems:**


**Other:**

3 PAVs per real volume

## 5.4 Results

Figure 5.1 shows the improved normal restart time for IMS Version 8 when compared to IMS Version 7. For purposes of accurate comparison, the IMS Version 7 time reported is the aggregate time of the normal restart plus the time required to allocate and open all test databases.

**Figure 5.1: Normal restart to open databases - Elapsed time comparison**



IMS V7 /NRE + DB OPEN vs IMS V8 /NRE

Figure 5.2 shows the improved emergency restart time for IMS Version 8 when compared to IMS Version 7. For purposes of accurate comparison, the IMS Version 7 time reported is the aggregate time of the normal restart plus the time required to allocate and open all test databases.

**Figure 5.2: Emergency restart to databases open - Elapsed time comparison**




Figure 5.3 shows the improved database open time for IMS Version 8 when compared to IMS Version 7.

**Figure 5.3: Database open - Elapsed time comparison**

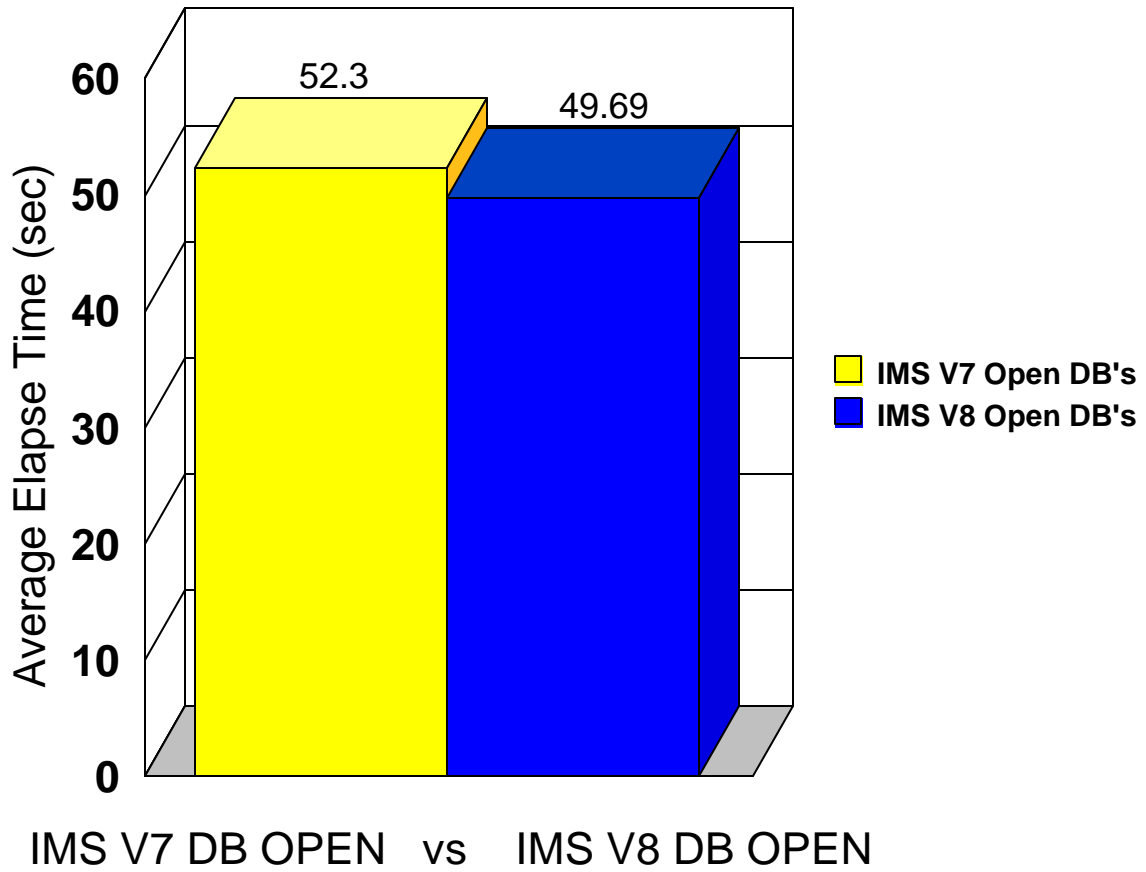


Figure 5.4 shows the improved shutdown time for IMS Version 8 when compared to IMS Version 7 when using the checkpoint freeze command to stop IMS.

**Figure 5.4: Checkpoint freeze command - Elapsed time comparison**





Figure 5.5 shows the improved shutdown time for IMS Version 8 when compared to IMS Version 7 when using the checkpoint dumpq command to stop IMS.

**Figure 5.5: Checkpoint DUMPQ command - Elapsed time comparison**





---

## 6 Common Service Layer

### 6.1 Introduction

This section shows the performance characteristics of the Common Service Layer component of IMS Version 8.

The following measurements were executed using the DSFF workload in a Shared Message Queues environment using the Resource Manager (RM), Operations Manager (OM), and Common Service Layer (CSL) address spaces.

### 6.2 Summary

The following summarizes the performance results observed when the CSL address space was used with IMS Version 8 relative to the CSL address space being absent with IMS Version 8. Refer to Section 6.4 for detailed results.

The IMS Version 8 Common Service Layer adds a single point of control for the IMSplex without incurring any overhead. The Common Service Layer relieves the IMS control region of less than 0.5% transaction pathlength.

### 6.3 Test Environment


The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 20GB storage, 6 CPs
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 ESCON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 16GB cache
<b>Operating Systems:</b>	z/OS DFSMS V1R2, IMS V8.1
<b>Other:</b>	3 PAVs per real volume

## 6.4 Results

Figure 6.1 shows the transaction rates achieved with IMS Version 8 when using the optional Common Service Layer address space compared to IMS Version 8 without the CSL address space. The testing of this environment included artificially high sign-on activity to exercise the CSL activity. The statistics show the Common Service Layer improves transaction performance by less than 0.5% with the benefit of centralized common services for the IMSplex and reduced IMS control region CPU utilization.

**Figure 6.1: Common Service Layer - Throughput comparison**



## 7 Fast Path Greater-Than-240 Areas Support

### 7.1 Introduction

This section shows the performance characteristics of the support added to IMS Version 8 with the Fast Path Greater-Than-240 Areas feature which allows Data Entry Databases (DEDB) to scale up to 8 TB using up to 2048 area data sets. This section focuses on overhead incurred when the number of area data sets is increased above 240 areas. The 240 area baseline was used because that is the DEDB area limit in IMS Version 7.

### 7.2 Summary

The following summarizes the performance results for the Fast Path Greater-Than-240 Areas support study. This study observes the performance impact of increasing the number of areas used for a DEDB with no change to the number or size of each database record. Testing was performed using the FP2 workload. Refer to Section 7.4 for detailed results.

The IMS Version 8 Fast Path Greater-Than-240 Areas support feature:

- incurs no appreciable overhead to DEDB access
- incurs no appreciable overhead when increasing the number of DEDB areas above 240

### 7.3 Test Environment

The description of the test environment is as follows:


<b>Hardware and Software Environment</b>	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 10GB storage, 3 CPs
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 ESCON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 16GB cache
<b>Operating Systems:</b>	OS/390 2.10, IMS V8.1
<b>Other:</b>	3 PAVs per real volume



## 7.4 Results

Figure 7.1 shows the performance impact to IMS transaction rates when doubling the number of area data sets used for a DEDB. The statistics show less than 0.5% variance in the throughput rates such that the results can be deemed equivalent.

**Figure 7.1: Fast Path GT240 Areas - Throughput comparison**



# 8 Coupling Facility Structure Duplexing

## 8.1 Introduction

This section shows the performance impact of Coupling Facility Structure Duplexing with IMS Version 8. This section focuses on structure duplexing for these key IMS CF structures:

- IMS Shared Message Queues
- Global IRLM Lock Structure

## 8.2 Summary

The following summarizes the performance results with the Coupling Facility Structure Duplexing enabled for the IMS Shared Message Queue and global IRLM lock structure. Refer to Section 8.4 for detailed results.

IMS Version 8 Structure Duplexing incurs the expense of duplexing while providing the benefit of uninterrupted availability during a coupling facility structure or connection failure. This is particularly important where high availability outweighs high transaction rates. The study shows IMS Version 8 CF structure duplexing on an internal CF incurs:

- 7.1% pathlength increase for Shared Message Queues structure duplexing
- 9.9% pathlength increase for IRLM structure duplexing
- 12.3% pathlength increase when duplexing both the IRLM lock structure and the Shared Message Queues structure

The study also shows IMS Version 8 CF structure duplexing on a remote CF incurs:

- 14.9% pathlength increase for Shared Message Queues structure duplexing
- 23.1% pathlength increase for IRLM structure duplexing
- 31.9% pathlength increase when duplexing both the IRLM lock structure and the Shared Message Queues structure

## 8.3 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processors:</b>	2064-216 - 20 GB storage , 3 CPs
<b>ESS F20 Disk:</b>	18GB 10K RPM disk drives, 4 x ESCON channels, 8 LSS, 16 ranks, Volume config = 3390-3, 16GB Cache
<b>Operating Systems:</b>	z/OS DFSMS V1R2

**Notes:**

3 PAVs per real volume

## 8.4 Results

The Coupling Facility Structure Duplexing feature of z/OS allows IMS to keep a current copy of the critical coupling facility structures to be used in the event of CF structure unavailability. The process of duplexing the IMS structures bears a cost in terms of IMS transaction rates.

Figure 8.1 shows the transaction rates achieved when IMS Version 8 was active with data sharing and Shared Message Queues. The statistics show the impact to IMS transaction rates when enabling structure duplexing of the shared message queues structure and the IRLM global lock structure on an internal coupling facility. The statistics show that overhead for duplexing multiple structures is less than the aggregate overhead incurred for duplexing the individual structures.

**Figure 8.1: CF duplexing - Throughput comparison**

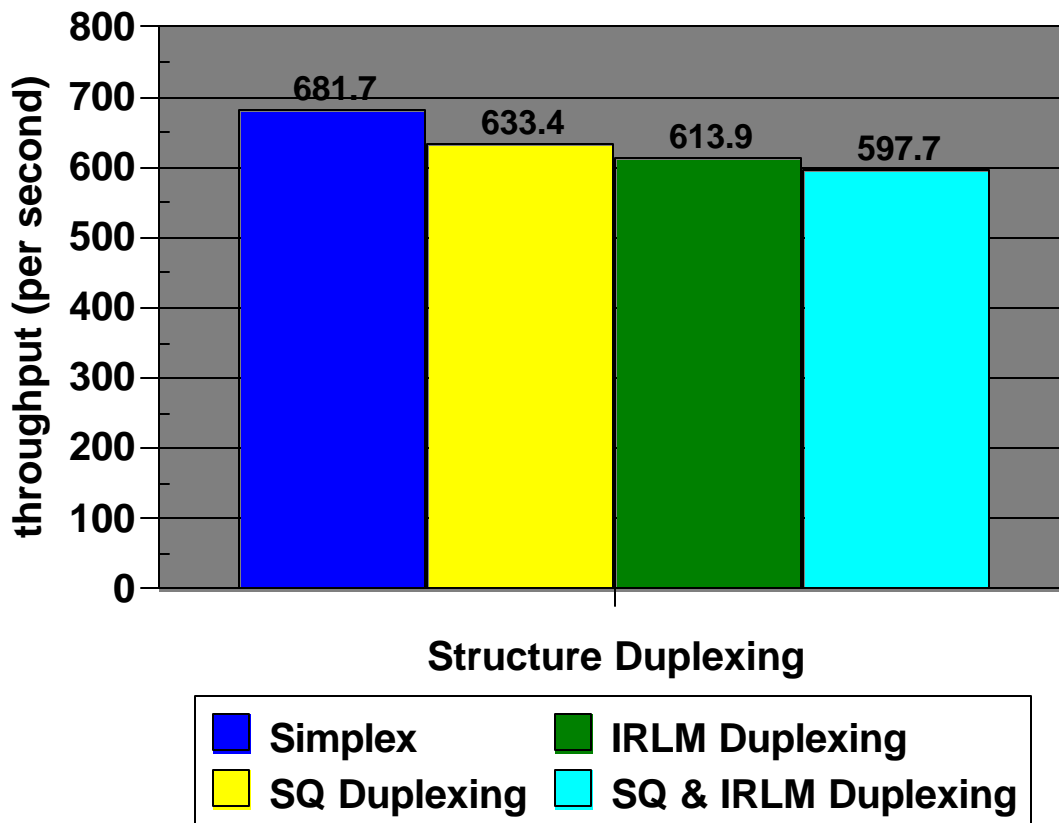
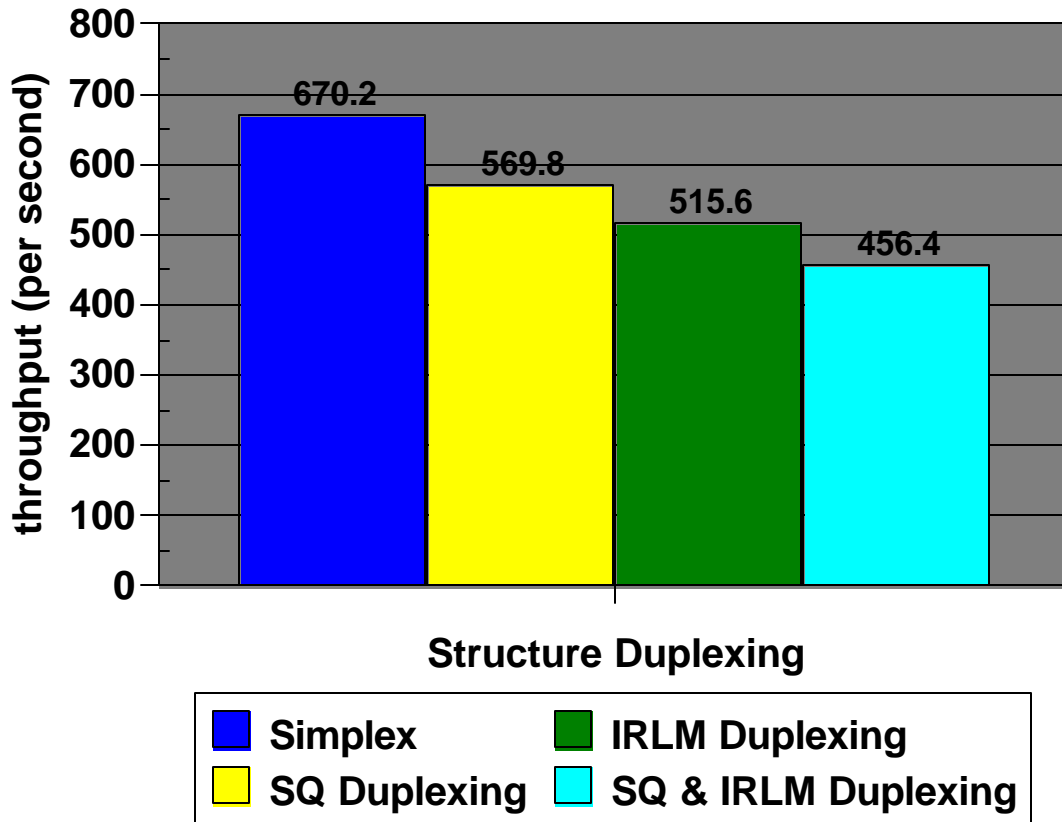




Figure 8.2 shows the transaction rates achieved when IMS Version 8 was active with data sharing and Shared Message Queues. The statistics show the impact to IMS transaction rates when enabling structure duplexing of the shared message queues structure and the IRLM global lock structure on a remote coupling facility. The statistics show that overhead for duplexing multiple structures is less than the aggregate overhead incurred for duplexing the individual structures.

**Figure 8.2: CF duplexing - Throughput comparison**



## 9 IMS Java Message Processing Region

### 9.1 Introduction

This section shows the performance characteristics of the IMS Version 8 Java Message Processing region support when compared to an equivalent IMS application program using COBOL. An evaluation including COBOL was deemed valuable for purposes of application planning and design considerations. At the writing of this paper additional Java JVM (SDK 1.3.1) performance enhancements are in progress so improved performance of the IMS Java Message Processing region support can be expected.

The IMS Java Dependent Region and IMS Java Offerings Performance white paper can be located for additional information at the IMS Family web site, <http://www-3.ibm.com/software/data/ims/> .

### 9.2 Summary

The following summarizes the performance results observed with IMS Version 8 using a Java application in a Java Message Processing region relative to IMS Version 8 using a COBOL application in a Message Processing Program region. Refer to Section 9.4 for detailed results.

The IMS Version 8 Java Message Processing region support using a Java application demonstrates:

- transaction rates up to 2134 transactions per second
- 258% additional pathlength when compared to an equivalent COBOL application

### 9.3 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processor:</b>	2064-116 (zSeries 900 ) - 28GB storage, 14 CPs
<b>Operating Systems:</b>	z/OS DFSMS V1R2, IMS V8.1



## 9.4 Results

Figure 9.1 shows the transaction rates attained with IMS Version 8 executing a Java application program in a JMP region compared to an equivalent COBOL application program executing in an MPP region. Due to variations in CPU utilization the throughput reported for the COBOL application program is an extrapolated value for purposes of comparison.

**Figure 9.1: IMS JMP - Throughput comparison**


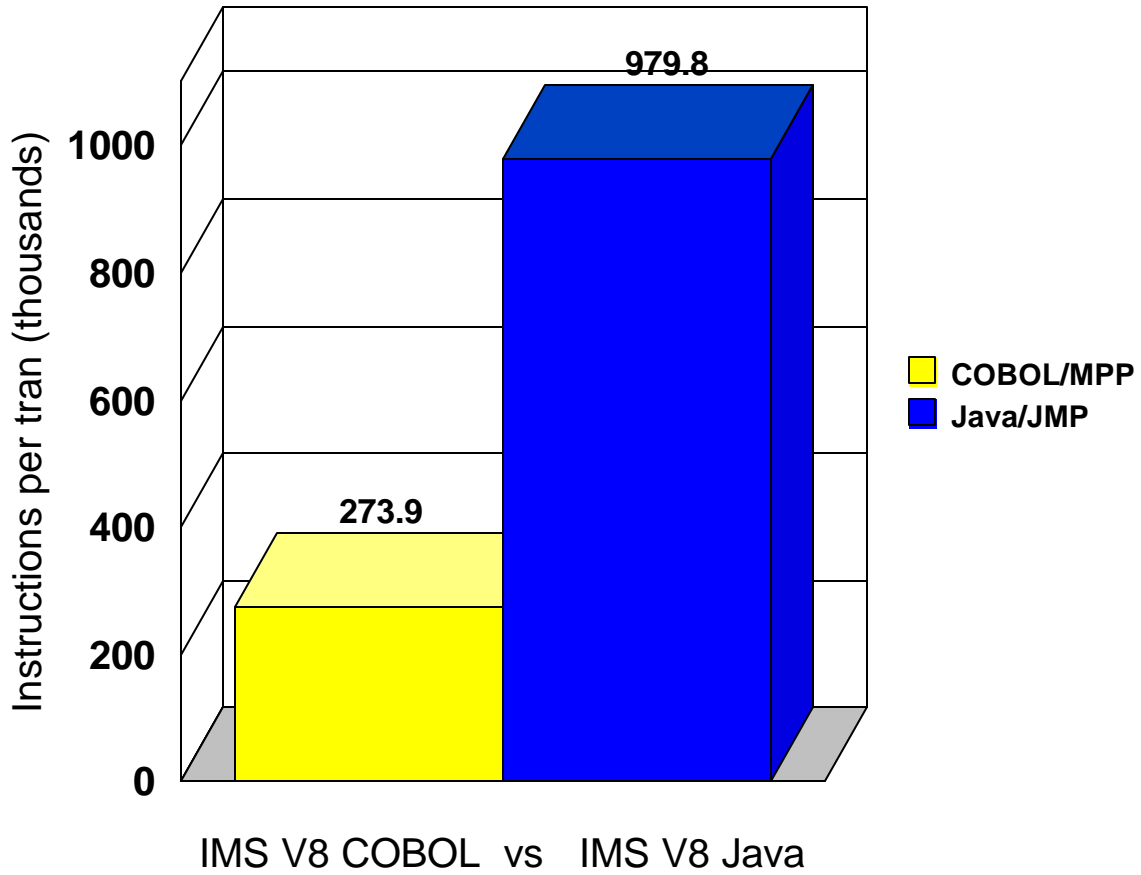


Figure 9.2 shows the pathlengths incurred for the equivalent IMS COBOL and IMS Java application programs under IMS Version 8.

**Figure 9.2: IMS Java - Instruction pathlength comparison**



# 10 Fast Path High Stress

## 10.1 Introduction

This section shows the performance characteristics of IMS Version 8 Fast Path in a high stress Shared EMHQ environment with data sharing in a parallel sysplex.

The following test is executed using the FP2 workload as described in Section 3.1.

## 10.2 Summary

The following summarizes the performance results observed during the IMS Version 8 Fast Path High Stress test. Refer to Section 10.4 for detailed results.

IMS Version 8 configured in a 4-way parallel sysplex using data sharing and the Fast Path shared EMHQ and with all data sets residing on an ESS M800 is observed to achieve:

- throughput up to 11,784 transactions per second
- CPU utilization up to 95%
- DASD I/O rates up to 12,960 per second with an average response time of 1 millisecond

## 10.3 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
<b>Processor:</b>	2064-216 (zSeries 900 ) - 28GB storage, 12 Cps (4 x 3 CPs per LPAR) 2064-216 - 2 x ICF, 2 CPs per ICF - 4 Cps 9672-ZZ7 - 4 CPs
<b>2105-M800 Disk:</b>	36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9, 16GB cache
<b>Operating Systems:</b>	z/OS DFSMS V1R2, IMS V8.1
<b>Other:</b>	7 PAVs per real volume

## 10.4 Results

Figures 10.1 - 10.2 summarize the parallel sysplex configuration used and the results obtained.

Figure 10.1 shows an illustration of the parallel sysplex configuration used in the Fast Path High Stress evaluation. The workload executes on a 2064-216 processor and is driven by a stand-alone 9672-ZZ7 processor. All data sets reside on a single ESS model M800 using 4 native FICON channels. 4-way data sharing through IRLM and shared EMHQ are used.

**Figure 10.1: Fast Path High Stress - Parallel sysplex configuration**

2064-216

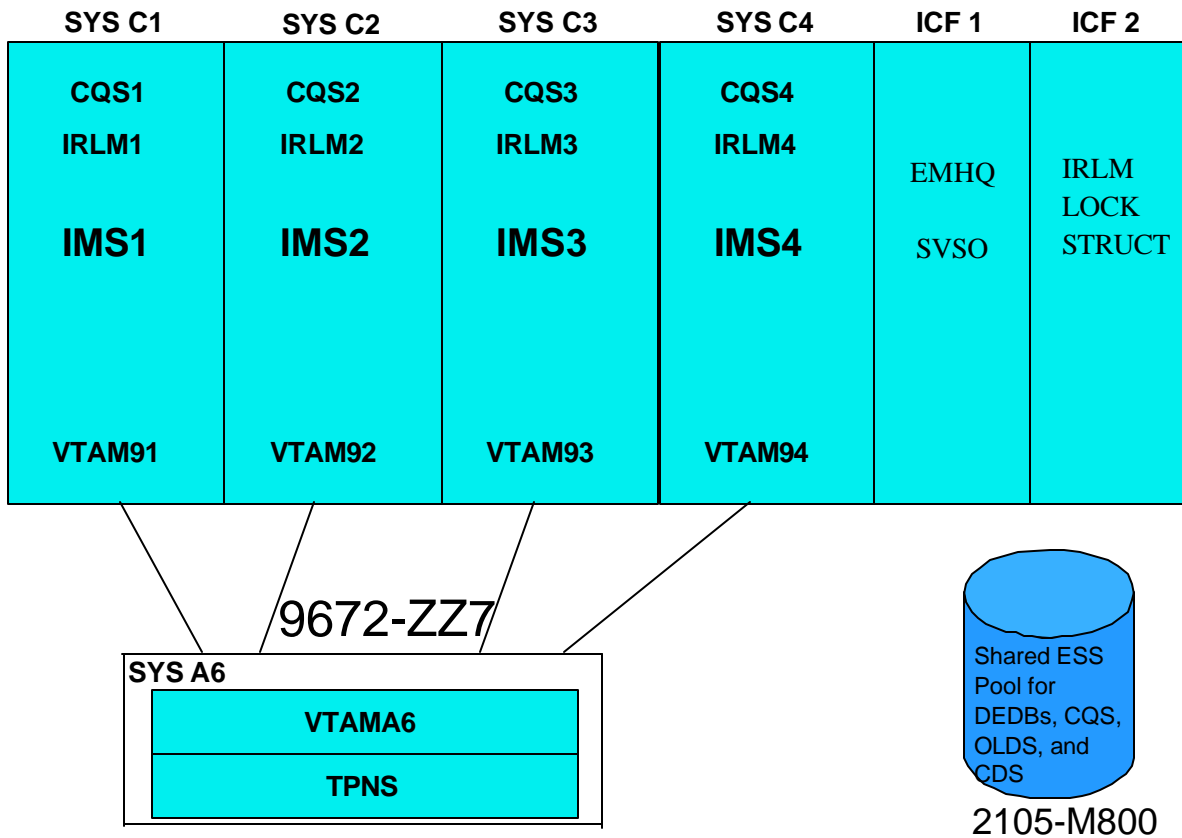






Figure 10.2 summarizes the results observed during the Fast Path High Stress test. This result achieves the benchmark of processing 1 billion transactions per day in a 4-way parallel sysplex.

**Figure 10.2: Fast Path High Stress - Results**

<b>Fast Path High Stress results</b>	
<b>Throughput:</b>	11,784 transactions (aggregate per second)
<b>CF utilization:</b>	52.3% (of 4 CPs)
<b>CPU utilization:</b>	95.4% (of 12 CPs)
<b>DASD I/O rate:</b>	12,960 I/Os (per second)
<b>Ave. I/O resp time:</b>	1 millisecond

## 11 ESS M800 Turbo with IMS Version 8

### 11.1 Introduction

This section shows the performance improvements in IMS Version 8 when using Enterprise Storage Server (ESS) turbo model 800 with 1 Gigabit FICON channels when compared to using ESS model F20 with 1 Gigabit channels. This section focuses on IMS performance for these key areas:

- **ESS 800 versus ESS F20 Performance Results**
  - IMS logging
  - WADS I/O response time
  - IMS Batch Message Processing
  - IMS Database Reorganization

Note, the improvements observed in the MSC environment will also provide improved results in all WADS intensive processes such as IMS Shared Message Queues processing.

### 11.2 Summary

The following summarizes the performance results with the ESS 800 relative to the ESS F20. Refer to Section 11.4 for detailed results.

IMS Version 8 using ESS M800 demonstrates the following performance improvements when compared to the ESS F20 :

- IMS OLDS logging bandwidth increased by 31%
- FICON channel utilization reduced 9.8%
- MSC throughput increased 40%
- WADS I/O response times reduced 65%
- Batch Message Processing elapsed times reduced by 19%
- IMS database reorganization times reduced by 34%

### 11.3 Test Environment

The description of the test environment is as follows:

<b>Hardware and Software Environment</b>	
<b>Processor:</b>	2064-216, 28 GB real storage
<b>2105-800 Disk:</b>	36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
<b>2105-F20 Disk:</b>	18GB 10K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-3, 3 PAVs per real volume
<b>Operating Systems:</b>	z/OS DFSMS V1R2, IMS V8.1

## 11.4 Results

Figures 11.1 - 11.2 show the improved IMS logging performance of the ESS M800 when compared to the ESS F20.

Figure 11.1 shows the IMS OLDS logging bandwidth improvement with the ESS M800 using one 1 Gigabit FICON channel relative to the ESS F20 using either one ESCON channel or one 1 Gigabit FICON channel. The ESS M800 out performs the ESS F20 in maximum IMS logging bandwidth.

**Figure 11.1: OLDS logging - Bandwidth comparison**

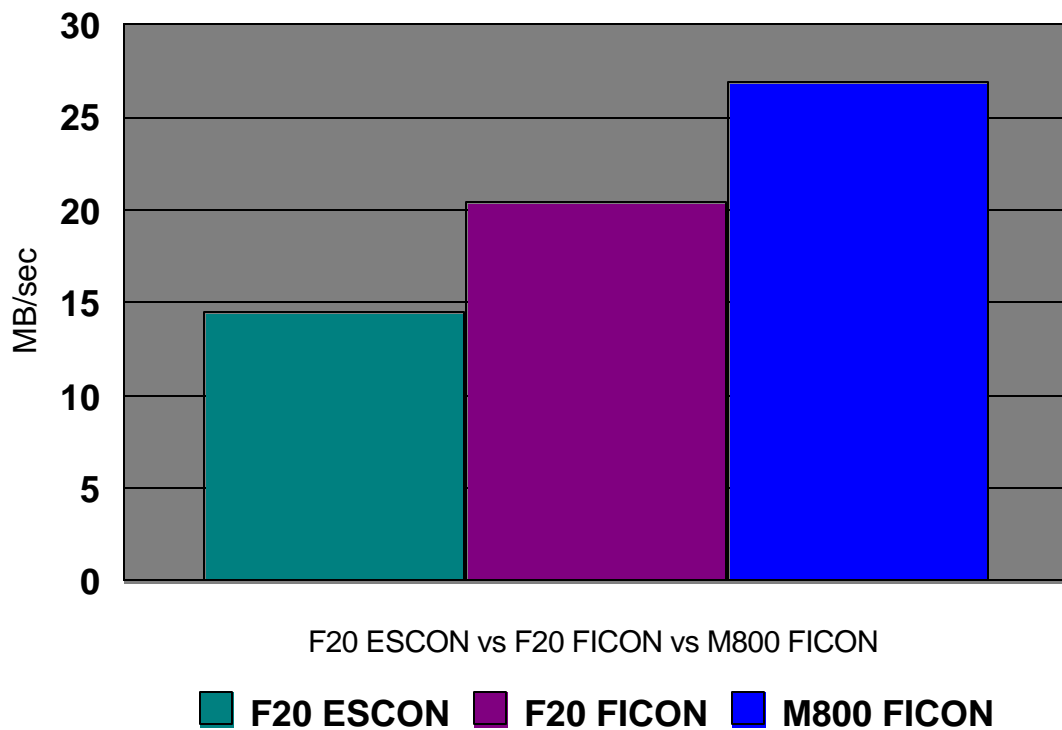



Figure 11.2 shows the improvement in channel utilization by ESS 800 compared to ESS F20 with both ESCON and FICON channels. The statistics show the ESS M800 reduces the channel utilization when compared to ESS F20.

**Figure 11.2: OLDS logging - Channel utilization comparison**



F20 ESCON vs F20 FICON vs M800 FICON

**F20 ESCON** **F20 FICON** **M800 FICON**

All 1 Channel Path

Figures 11.3 - 11.4 show the improved IMS Write Ahead Data Set (WADS) performance of the ESS M800 relative to the ESS F20 in an IMS Multiple Systems Coupling (MSC) environment.

Figure 11.3 shows the improved MSC throughput achieved when the WADS resides on the ESS M800 when compared the WADS residing on the ESS F20.

**Figure 11.3: MSC via VTAM CTC - Throughput comparison**


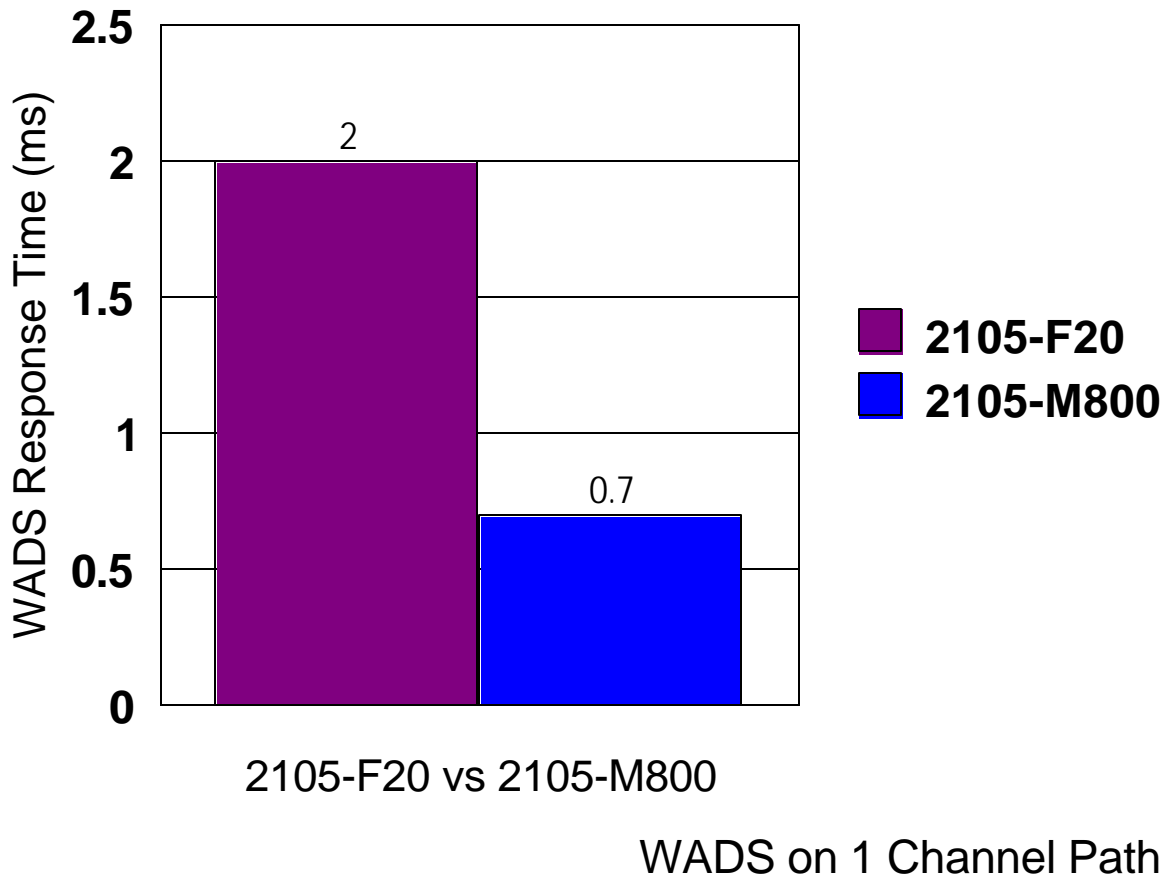




Figure 11.4 shows the improved WADS I/O response time achieved when the WADS resides on the ESS M800 when compared the WADS residing on the ESS F20. This constitutes the first time IMS WADS I/O is measured in sub-millisecond times.



**Figure 11.4: MSC via VTAM CTC - WADS I/O response time comparison**





Figure 11.5 shows the improved Batch Message Processing elapsed times (in seconds) when the target database resides on ESS M800 when compared to the target database residing on ESS F20.

**Figure 11.5: BMP sequential processing - Elapsed time comparison**

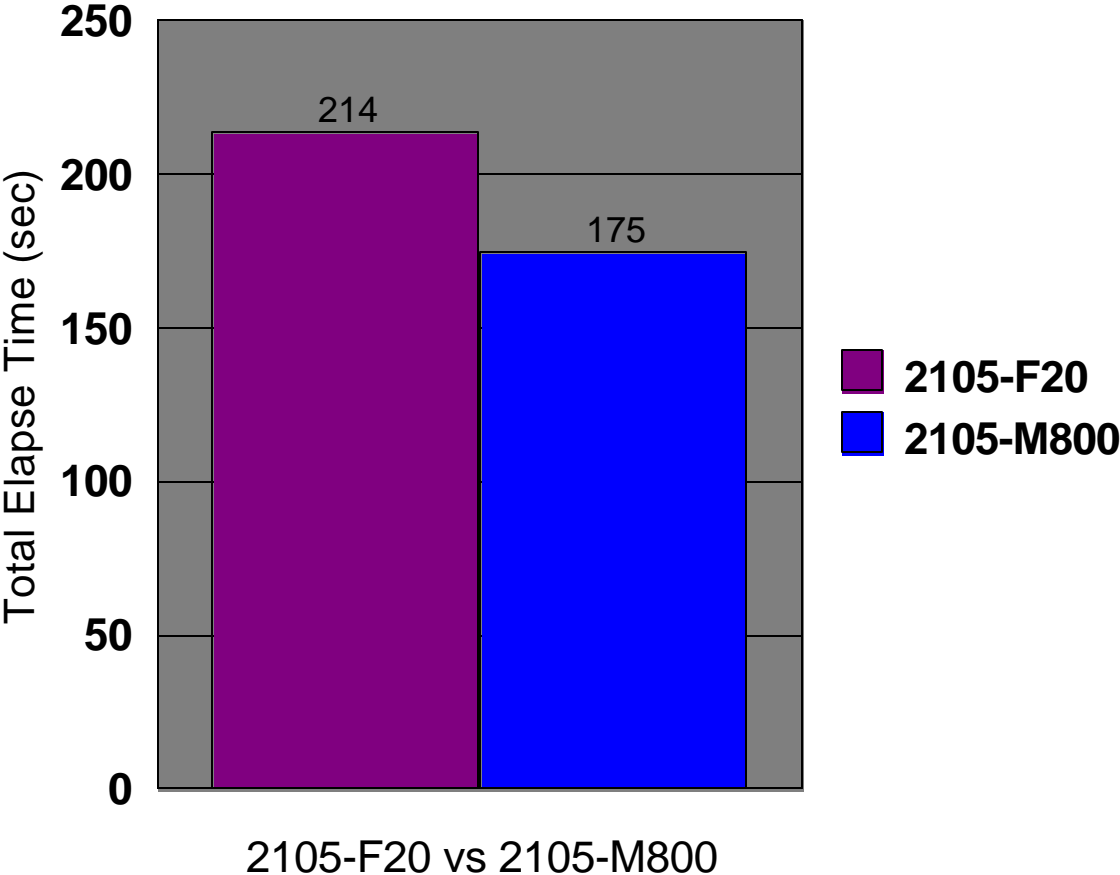
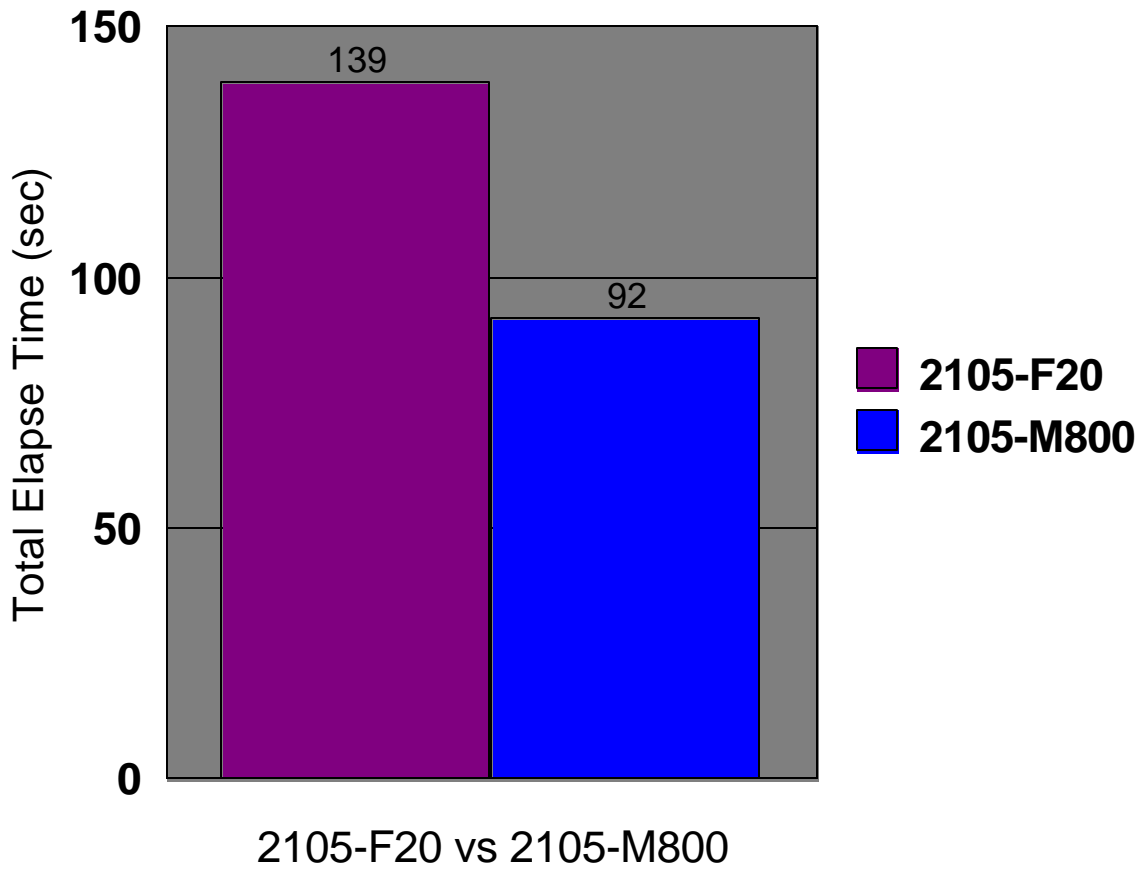


Figure 11.6 shows the reduced batch window requirement for database reorganization when the target database resides on ESS M800 when compared to the target database residing on ESS F20.

**Figure 11.6: Database Reorganization - Elapsed time comparison**



## 12 Test Methodology

The test methodology used in the IMS Version 8 performance study is similar to the methodology described in the IBM Large Systems Performance Reference, document number SC28-1187-08, with the exception of the choice of terminal simulators. This study used the IBM Teleprocessing Network Simulator on a stand-alone processor in place of the proprietary 'internal driver' employed in the LSPR measurements.

The Large System Performance Reference for IBM can be found at:  
<http://www-1.ibm.com/servers/eserver/zseries/lSpr>

The LSPR document can be obtained at:  
<http://www-1.ibm.com/servers/eserver/zseries/lSpr/pdf/SC2811878.pdf>

Measurement data is to be considered equivalent for comparison purposes in this document when it is between +/-3%.

## 12.1 Disclaimer

References in this document to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

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 expressed 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 their 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 environments do so at their own risk.

Any performance data contained in this document was obtained in a controlled environment based on the use of specific data. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data in their specific environment.

The test scenarios (hardware configuration and workloads) used in this document to generate performance data are not considered 'best performance case' scenarios. Performance may be better or worse depending on the hardware configuration, data set types and sizes, and the overall workload on the system.

## 12.2 Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM®

IMS

Enterprise Storage Server

z/OS

FICON

zSeries