IMS Version 9 Performance Summary

October 2004

Author: B. Pibulsonggram
Contributors: C. Chen, M. Gonzales, H. Neal, H. Tang

Abstract

IMS Version 9 offers the features to enhance the availability, performance, integration, manageability and scalability of IMS and IMS data. This paper illustrates the performance characteristics of new enhanced IMS Version 9.

The information provided in this paper obtained at the IBM Silicon Valley Laboratory. It is for migration and capacity planning purposes.

Software Group
San Jose

1 Introduction

IMS Version 9 offers availability, performance, integration, manageability and scalability enhancements to the customers. This paper illustrates the performance characteristics of the new enhanced IMS Version 9:

- Base Functions:
  - Full Function
  - Fast Path
  - BMP
  - APPC
  - Full Function Shared Queues
- CSA VSCR
- HALDB Online Reorganization
- Fast Path Area Open/Close Enhancements
- Fast Path DEDB Shared VSO Multi-Area Structure Support
- DBRC Enhancements

In addition, IBM eServer zSeries 2084 Model 316 was used to demonstrate the ability of IMS to manage extremely high transaction volumes - IMS Fast Path High Stress Study.

The following products were used in the study:

- IBM eServer zSeries 2064 Model 216 and IBM eServer zSeries 2084 Model 316
- IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800
- z/OS 1.4 and 1.5
- TPNS 3.5

2 Executive Overview

According to the evaluation, the performance of IMS Version 9 surpasses IMS Version 8 in the following areas:

HALDB Online Reorganization

Overcomes the limitation of HALDB databases unavailability during databases reorganization by providing a nondisruptive online reorganization**.**

CSA Virtual Storage Constraint Relief

Reduces CSA usage below the 16 MB when compared to IMS Version 8.

Fast Path DEDB Shared VSO Multi-Area Structure Support

Enhances the usability of Shared VSO Structure Support by allowing multiple areas to share a single Coupling Facility Structure.

Fast Path Area Open/Close Enhancements

Enable DEDB areas to be automatically opened after normal restart (warm) or emergency restart or IRLM reconnect, and use up to 10 Task Control Blocks (TCBs) to open, preopen, or close DEDB areas, allowing parallelism for these tasks.

DBRC Enhancements

Provide support for tape block sizes greater than 32 KB for the output from the Database Image Copy Utility and Online Database Image Copy Utility.

Base Functions

IMS Version 9 ’s base functions continue to perform within the guideline.

3 Base Performance

3.1 Introduction

The performance characteristics of the base functions of IMS Version 9 and IMS Version 8 were compared using the basic transaction processing and database access functions of IMS.

The following workloads were selected for the evaluation: DSWH, DCW; DSFF; Banking BMP; FP2; and OLR

DSWH : The Data Systems Workload HALDB, has a mixture of Full Function databases and a High Availability Large database, performs extensive database updates, and uses small messages for transaction responses.

DCW : The Debit-Credit Workload, has Fast Path databases and is designed to produce large volumes of log data (24K per transaction).

DSFF : The Data Sharing Full Function Workload, has a mixture of Full Function databases, performs medium database updates, and uses large messages for transaction responses. This workload uses IRLM for the datasharing lock manager.

The Banking BMP Workload has the characteristics of a customer workload, performs extensive sequential Fast Path database updates. This workload simulates end-of-day account reconciliation.

FP2: The Fast Path Two Workload, has the characteristics of a credit card processing system with DEDB Fast Path databases and performs light database updates.

OLR : HALDB Online Reorganization Workload, has a mixture of PHIDAM and PHDAM databases, performs medium database updates concurrently with online reorganization.

All IMS performance results were gathered during a steady state.

3.2 Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2064 Model 216 (zSeries 900) - 12 GB storage, 3 CPs; IBM eServer zSeries 2064 Model 216 (zSeries 900) - 2 x ICFs, 2 CPs and 8 GB storage per ICF (for shared queues test)
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4, IMS V8.1, IMS V9.1

3.3 Results

Figures 3.1 - 3.5 illustrate the performance characteristics of IMS Version 9 when compared to IMS Version 8:

- Full Function transaction processing rates
- Full Function Shared Message Queues transaction processing rates
- Fast Path transaction processing rates
- APPC transaction processing rates
- BMP elapsed times and CPU busy

Figure 3.1 Illustrates the ITR achieved while running the DSWH workload. The results show IMS Version 9 Full Function transaction processing capacity is within 1.0 percent variance of IMS Version 8.

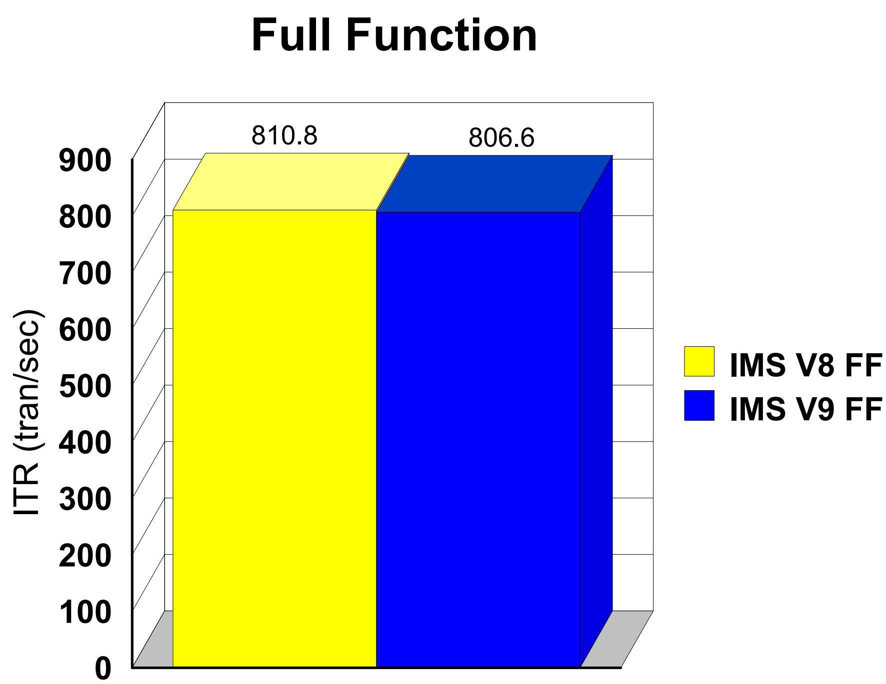

Figure 3.1: Full Function transaction processing - Transaction rates comparison

Figure 3.2 illustrates the ITR achieved while running the DSFF workload using shared queues. IMS Version 9 Full Function Shared Queues transaction processing capacity is within 3.0 percent variance of IMS Version 8.

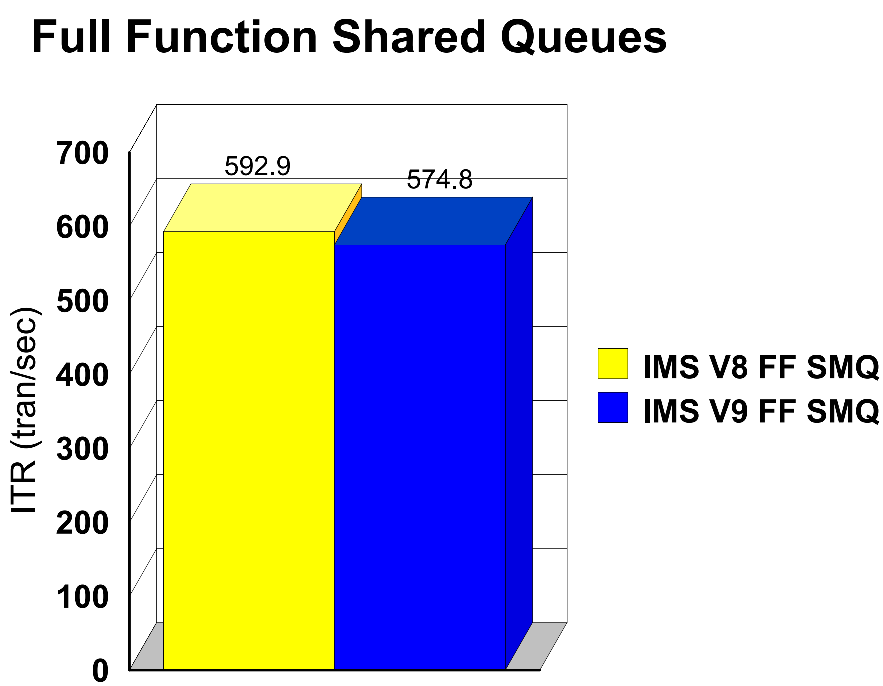

Figure 3.2: Full Function transaction processing in SMQ - Transaction rates comparison

Figure 3.3 illustrates the ITR achieved while running the DCW workload. IMS Version 9 Fast Path transaction processing capacity is within 2.3 percent variance of IMS Version 8.

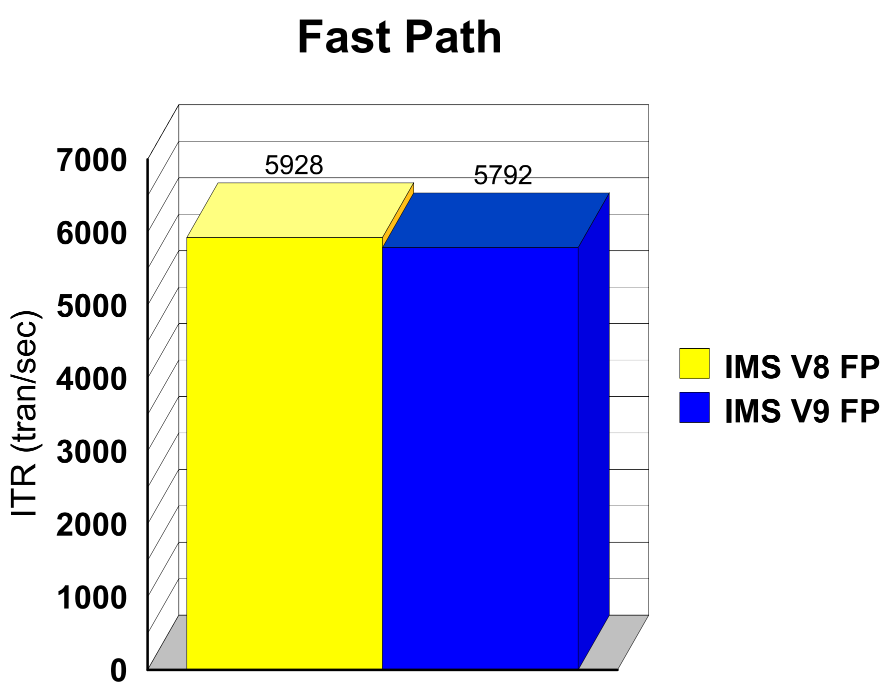

Figure 3.3: Fast Path transaction processing - Transaction rates comparison

Figure 3.4 illustrates the ITR achieved while running the DSFF workload. IMS Version 9 Full Function transaction processing capacity through APPC is within 3.0 percent variance of IMS Version 8.

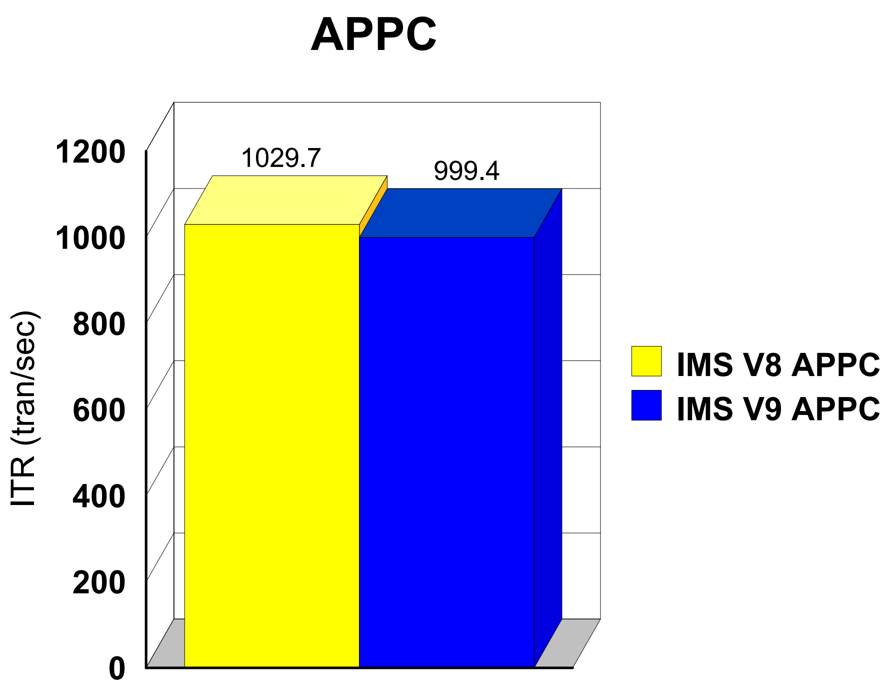

Figure 3.4: APPC transaction processing - Transaction rates comparison

Figures 3.5 - 3.6 illustrate the elapsed times and CPU busy while running the BMP banking workload. The elapsed times and CPU busy incurred by IMS Version 9 are within 0.6 percent and 2.3 percent variance of IMS Version 8 respectively.

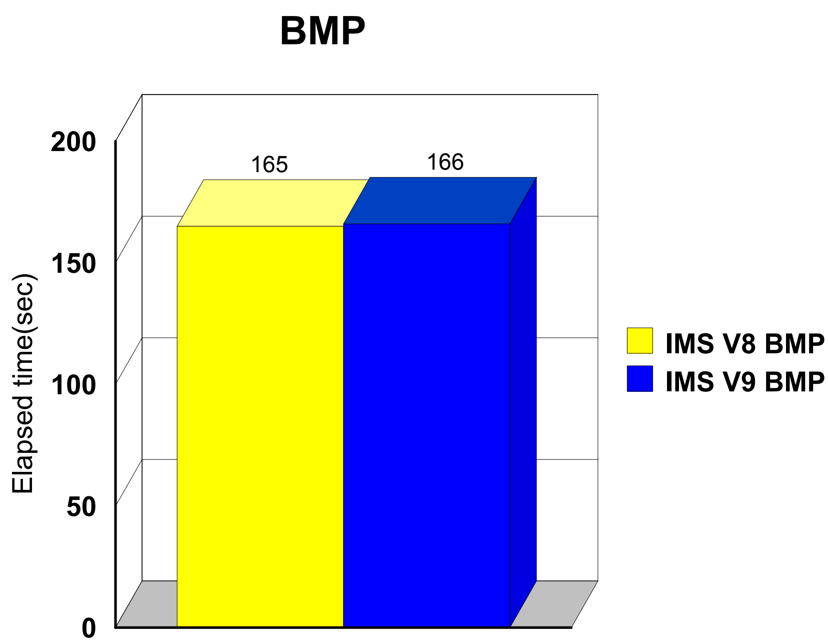

Figure 3.5: BMP elapsed times

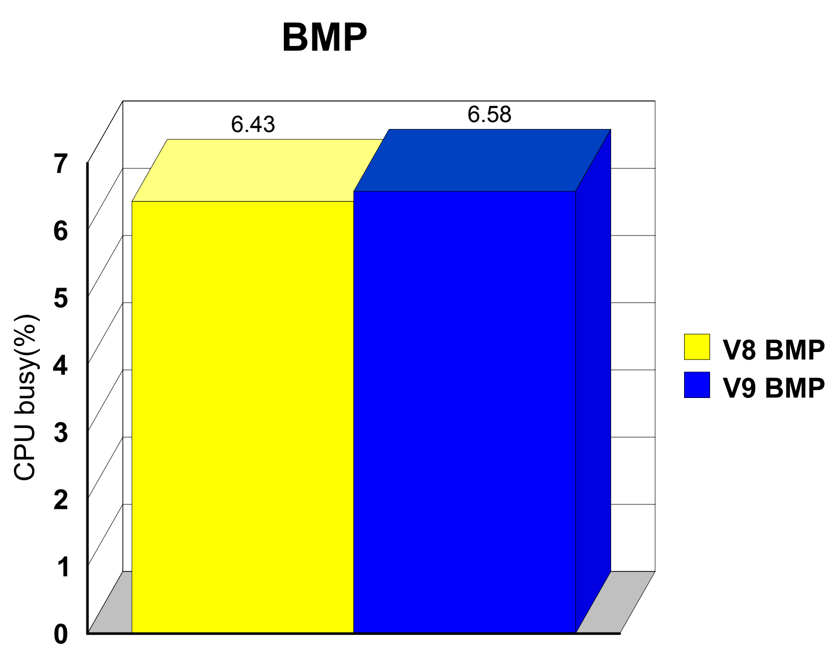

Figure 3.6: BMP CPU busy

3.4 Summary

IMS Version 9 demonstrates the following performance characteristics when compared to IMS Version 8.

- Full Function transaction processing : ITR degradation is within 1.0 percent
- Full Function transaction processing with Shared Message Queues: ITR degradation is within 3.0 percent
- Full Function transaction processing from APPC: ITR degradation is within 3.0 percent
- Fast Path transaction processing: ITR degradation is within 2.3 percent
- Batch Message Processing : Elapsed time increased by 0.6 percent and CPU busy increased by 2.3 percent

Generally, a variance below 5 percent is considered equivalent between measurements, however, the degradation for IMS Version 9 is targeted at or below 3 percent in any base function.

4 CSA Virtual Storage Constraint Relief

4.1 Introduction

This section illustrates the CSA usage below the 16 MB line by IMS Version 9 when compared to IMS Version 8.
.

4.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2084 Model 316 (zSeries 990) - 12 GB storage, 3 CPs; IBM eServer zSeries 2084 Model 316 (zSeries 990) - 2 x ICFs, 2 CPs and 8 GB storage per ICF (for shared queues test)
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R5, IMS V8.1, IMS V9.1

4.3 Results

Figures 4.1 - 4.2 illustrate the CSA utilization of IMS Version 9 when compared to IMS Version 8. The results show their characteristics in a variety of configurations:

- Full function processing
- Full Function Shared Queues
- Fast Path processing
- Batch message processing

The CSA usage below the 16 MB line by IMS Version 9 when compared to IMS Version 8 is illustrated in Figure 4.1. Full Function and BMP gain 8 KB and 4 KB respectively.

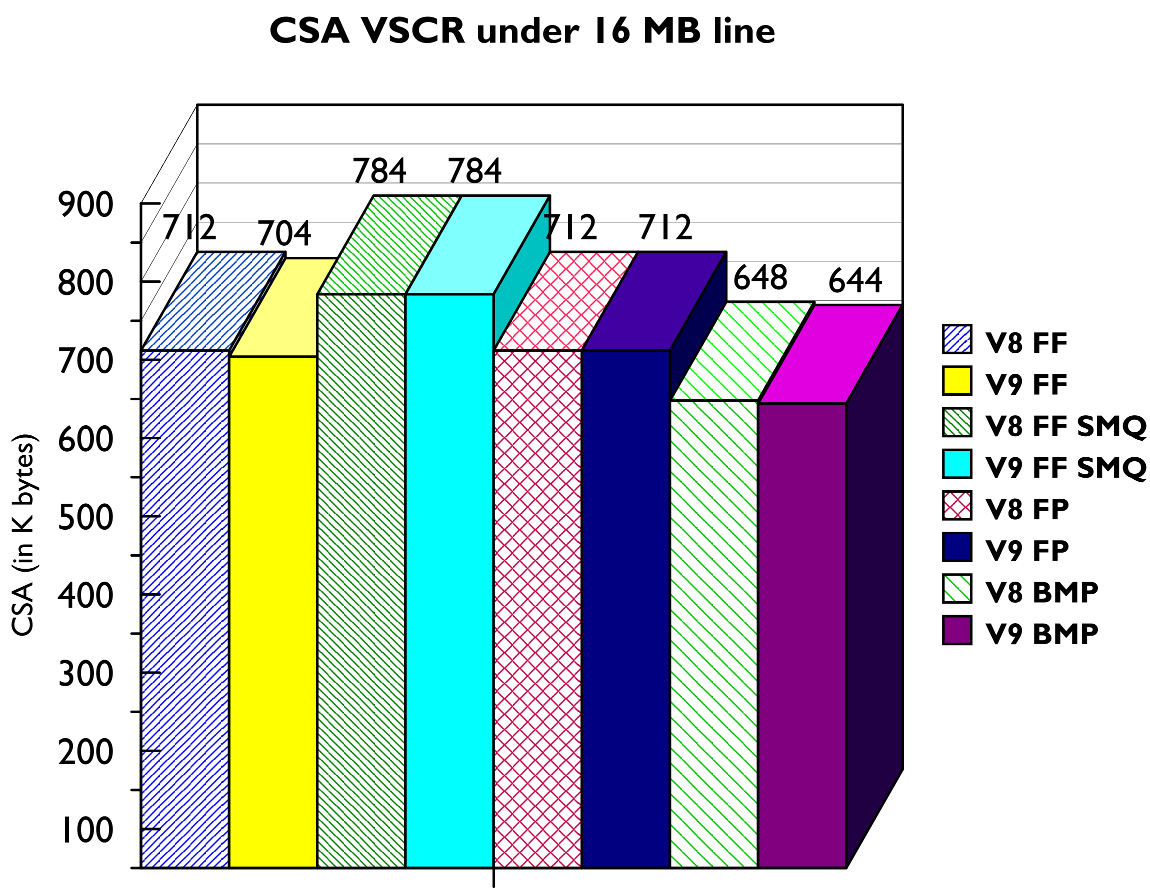

Figure 4.1: CSA VSCR - CSA utilization comparison

4.4 Summary

The IMS Version 9 **CSA VSCR** provides the following improvements when compared to IMS Version 8.

- Average CSA reduction is 3 KB
- Maximum CSA reduction is 8 KB

5 HALDB Online Reorganization

5.1 Introduction

HALDB Online Reorganization (OLR) overcomes the limitation of HALDB databases unavailability during databases reorganization by providing a nondisruptive integrated online reorganization.

5.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2064 Model 216 (zSeries 900) - 12 GB storage, 3 CPs; IBM eServer zSeries 2064 Model 216 (zSeries 900) - 2 x ICFs, 2 CPs and 8 GB storage per ICF (for sysplex test) IBM eServer zSeries 2084 Model 316 (zSeries 990) - 24 GB storage, 6 Cps; IBM eServer zSeries 2084 Model 316 (zSeries 990) - 2 x ICFs, 2 CPs and 8 GB per ICF (for sysplex test)
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4 and V1R5, IMS V8.1, IMS V9.1

5.3 Results

The initial study was performed using IBM eServer zSeries 2064 Model 216 (2064-216). After the test environment was migrated to the IBM eServer zSeries 2084 Model 316 (2084-316), the study was continued in the new environment.

Tables 5.1 - 5.8 and Figures 5.1 - 5.20 illustrate the evaluation results: The impact to online response time involving and not involving a HALDB partition under OLR in a N-way data sharing environment; The impact of different RATE values (100/50/25); The impact to update- intensive BMP against a partition under OLR; And the impact of the add-on OLRs OLDS logging rate.

The results of OLR PHIDAM/OSAM studies with various partition sizes - .5 GB, 1 GB, 2 GB, and 4 GB, in 2064-216 environment, are illustrated in table 5.1

Part. Size (GB)	OLR Elap. Time (mm:ss)	CPU Busy(%)	Log Volume (#Cyls/GB)	OLDS LOGG. Rate (MB/sec)	IRLM Tbl Acc./ SynSvt(ms)
0.5	4:44	12.0	1,056/.8	2.5	2,040/9.4
1	8:56	11.8	2,096/1.5	2.6	1,991/9.3
2	17:48	11.9	4,180/3	2.5	2002/9.2
4	35:30	11.8	8,355/6	2.6	1,971/9.2

Table 5.1: OLR PHIDAM/OSAM studies with .5 GB, 1 GB, 2 GB and 4 GB partition sizes

Figure 5.1 illustrates the OLR PHIDAM/OSAM elapsed times for various partition sizes - .5 GB, 1 GB, 2 GB, and 4 GB, in 2064-216 environment.

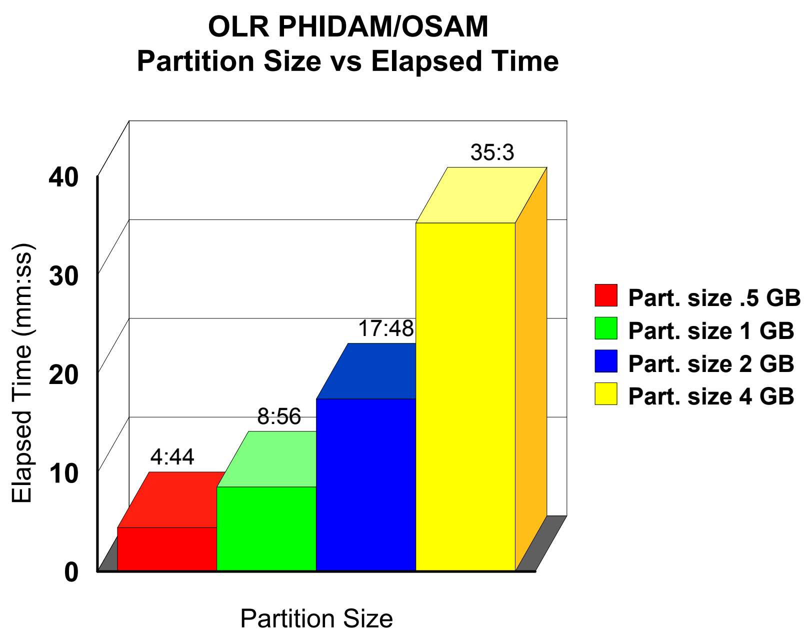

Figure 5.1: Partition sizes vs Elapsed times

Figure 5.2 illustrates the OLR PHIDAM/OSAM logging rate for various partition sizes - .5 GB, 1 GB, 2 GB, and 4 GB, in 2064-216 environment.

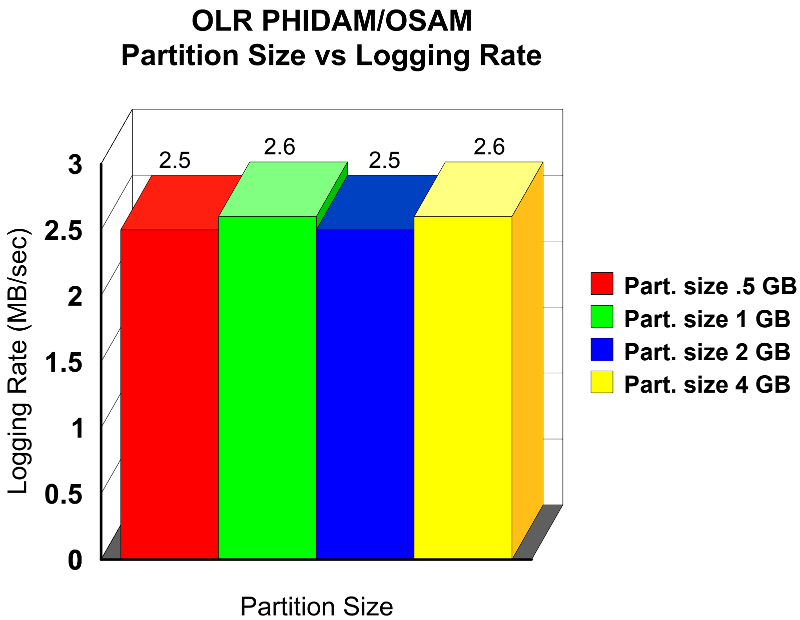

Figure 5.2: Partition sizes vs Logging rates

Figure 5.3 illustrates the OLR PHIDAM/OSAM CPU Busy for various partition sizes - .5 GB, 1 GB, 2 GB, and 4 GB, in 2064-216 environment.

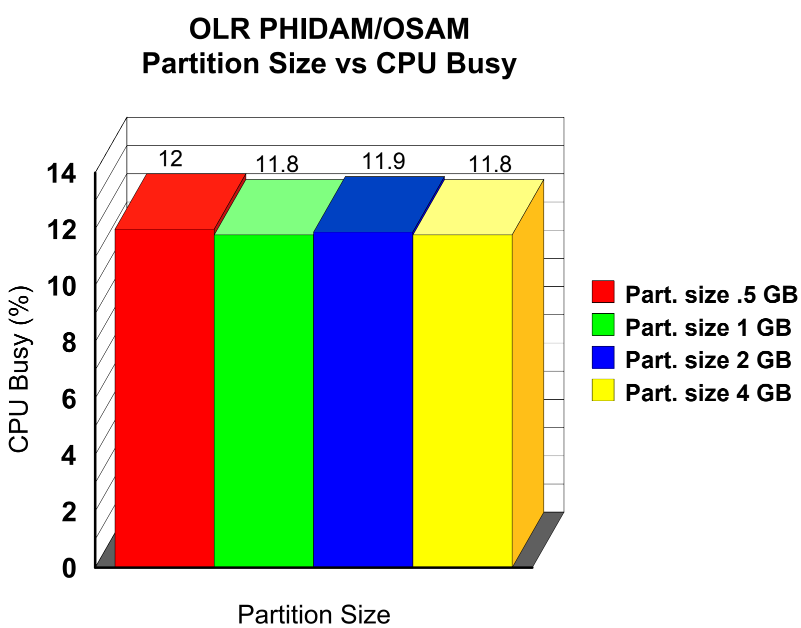

Figure 5.3: Partition sizes vs CPU busy

Figure 5.4 illustrates the OLR PHIDAM/OSAM IRLM Table access for various partition sizes - .5 GB, 1 GB, 2 GB, and 4 GB, in 2064-216 environment.

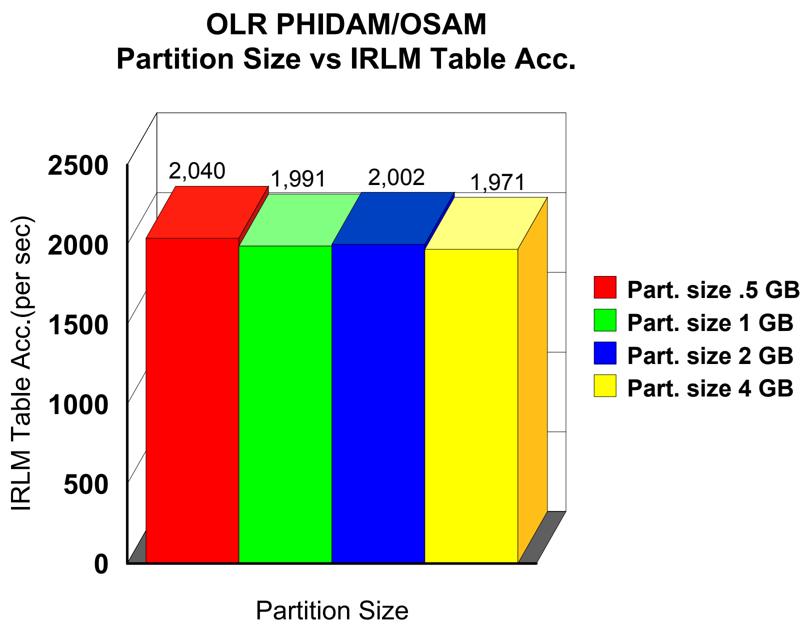

Figure 5.4: Partition sizes vs IRLM table Access

The results of OLR PHIDAM/OSAM (2 GB partition) studies, with various RATES- 100, 50, 25, in 2064-216 environment, are illustrated in table 5.2 and figures 5.5 - 5.8.

Rate	OLR Elap. Time (mm:ss)	CPU Busy(%)	Log Volume (#Cyls/GB)	OLDS LOGG. Rate (MB/sec)	IRLM Tbl Acc./ SynSvt(ms)
100	17:48	11.9	4,180/3	2.5	2002/9.2
50	34:28	6.8	4,180/3	1.4	1,079/9.2
25	68:08	4.2	4,180/3	0.7	578/9.3

Table 5.2: OLR PHIDAM/OSAM studies with RATES 100/50/25

Figure 5.5 illustrates the OLR PHIDAM/OSAM(2 GB partition) elapsed times for various RATEs - 100/50/25, in 2064-216 environment.

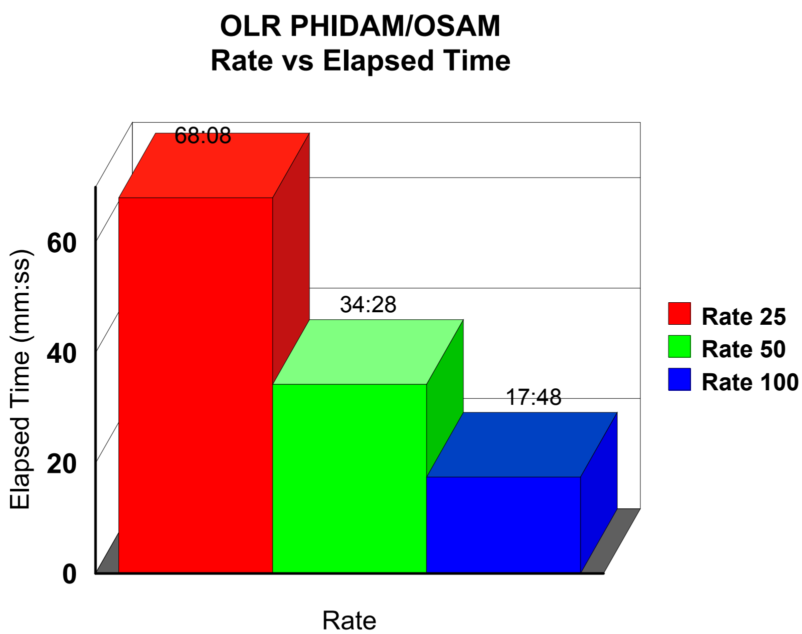

Figure 5.5: RATES vs Elapsed times

Figure 5.6 illustrates the OLR PHIDAM/OSAM(2 GB partition) logging rates for various RATEs - 100/50/25, in 2064-216 environment

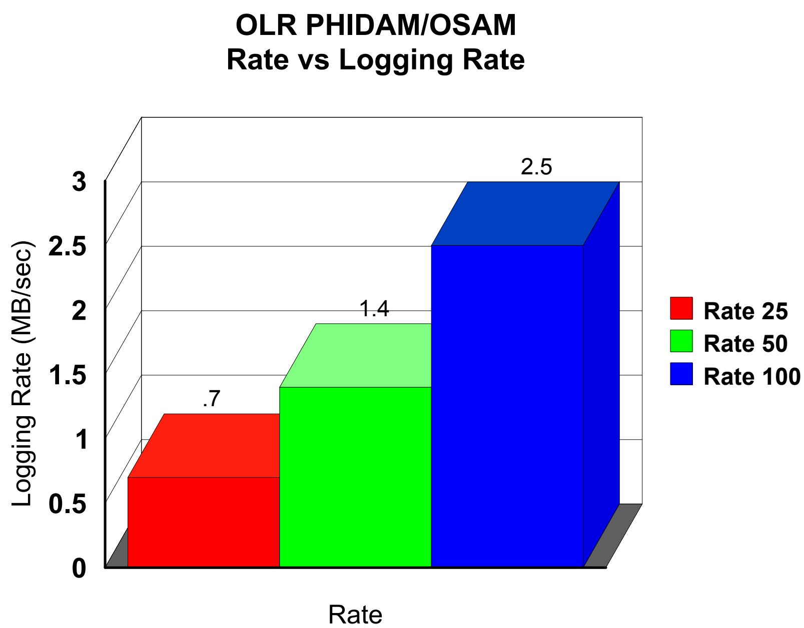

Figure 5.6: RATES vs Logging rates

Figure 5.7 illustrates the OLR PHIDAM/OSAM(2 GB partition) CPU busy for various RATEs - 100/50/25, in 2064-216 environment

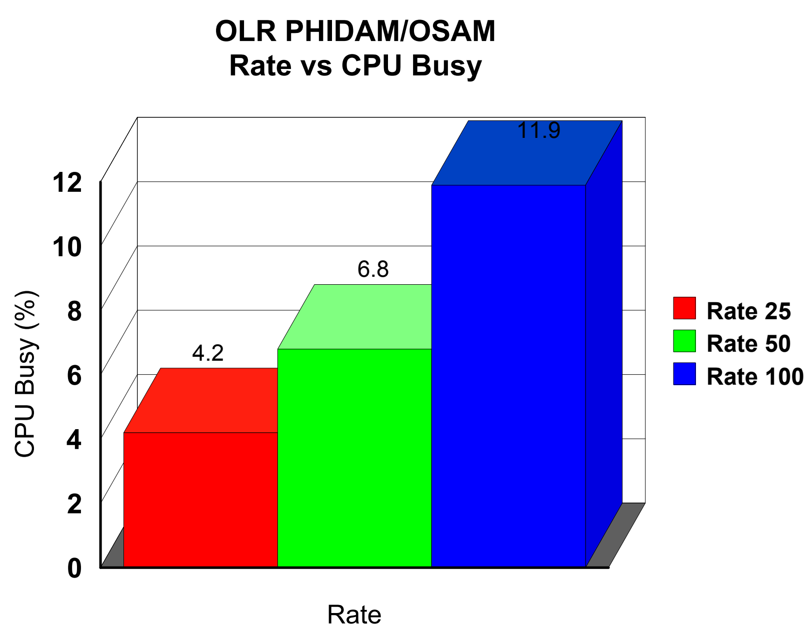

Figure 5.7: RATES vs CPU busy

Figure 5.8 illustrates the OLR PHIDAM/OSAM(2 GB partition) IRLM table access for various RATEs - 100/50/25, in 2064-216 environment**.**

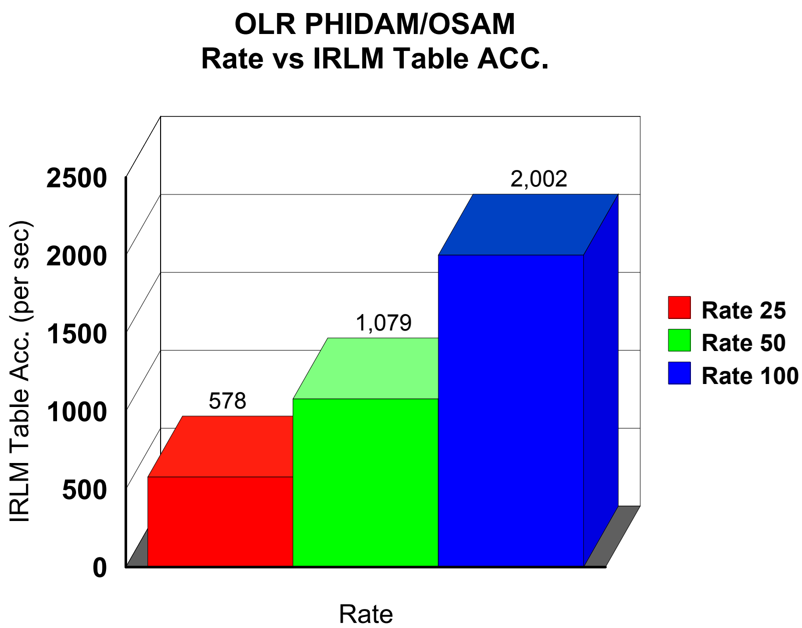

Figure 5.8:   RATES vs IRLM table Access

The results of OLR PHDAM/VSAM (2 GB partition) studies with various RATES- 100, 50, 25, in 2064-216 environment, are illustrated in table 5.3 and figures 5.9 - 5.12.

Rate	OLR Elap. Time (mm:ss)	CPU Busy(%)	Log Volume (#Cyls/GB)	OLDS LOGG. Rate (MB/sec)	IRLM Tbl Acc./ SynSvt(ms)
100 (1st A-M)	22:02	10.5	3,505/2.5	1.9	2,160/10.9
100 (1st M-A)	19:38	10.1	2,597/1.9	1.6	1,518/10.9
50	34:02	6	2,597/1.9	0.9	1,320/11.5
25	63:23	3.8	2,597/1.9	0.5	664/10.9

Table 5.3 OLR PHDAM/VSAM studies (2 GB Partition) with RATEs 100, 50, 25

Figure 5.9 illustrates the OLR PHDAM/VSAM(2 GB partition) elapsed times for various RATEs - 100/50/25, in 2064-216 environment

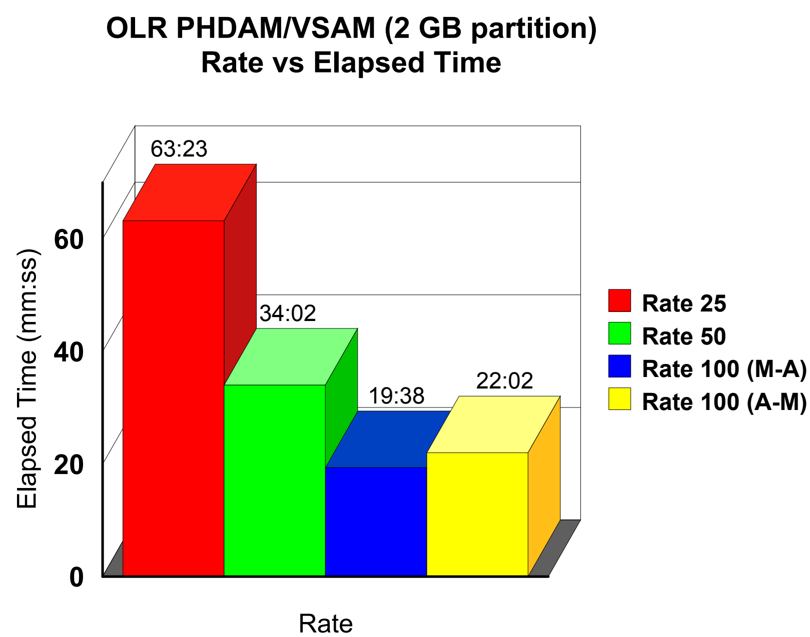

Figure 5.9: RATES vs Elapsed times

Figure 5.10 illustrates the OLR PHDAM/VSAM(2 GB partition) logging rates for various RATEs - 100/50/25, in 2064-216 environment

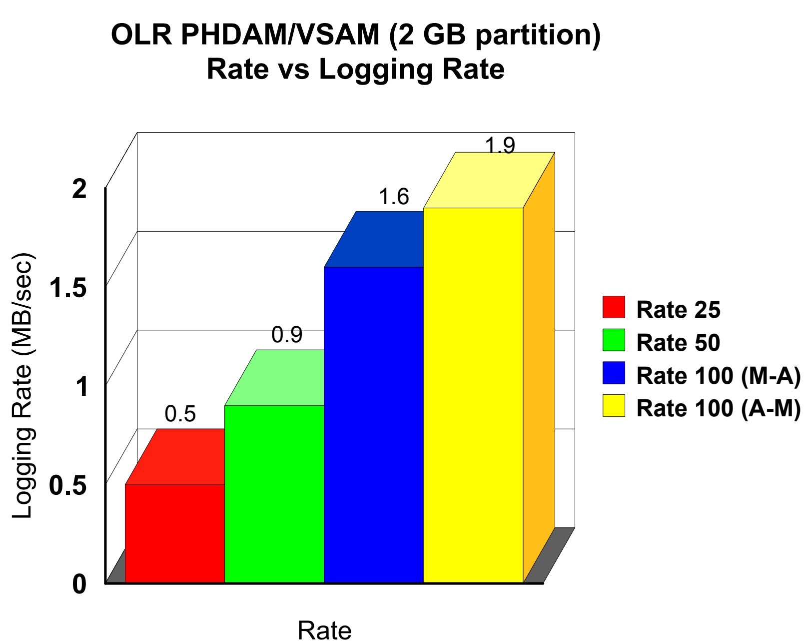

Figure 5.10: RATES vs Logging rates

Figure 5.11 illustrates the OLR PHDAM/VSAM(2 GB partition) CPU busy for various RATEs - 100/50/25, in 2064-216 environment

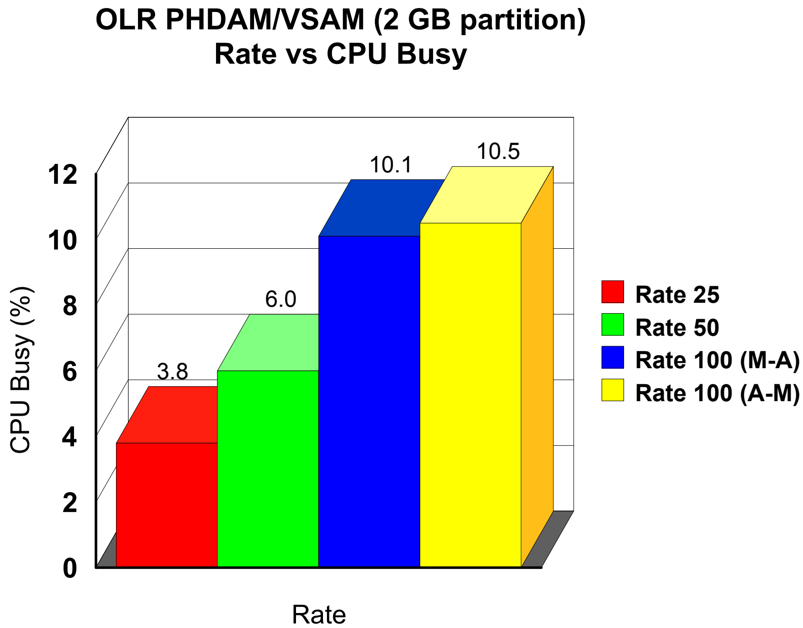

Figure 5.11: RATES vs CPU busy

Figure 5.12 illustrates the OLR PHDAM/VSAM(2 GB partition) IRLM table access for various RATEs - 100/50/25, in 2064-216 environment

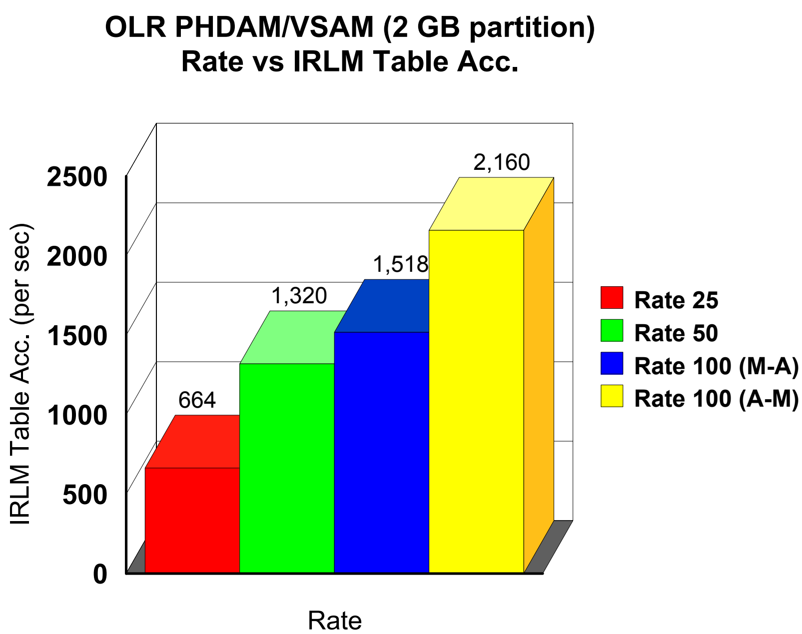

Figure 5.12: RATES vs IRLM table access

The results of OLR PHIDAM/OSAM (1 GB partition) studies with concurrent OLRs - 1/2/4/8, in 2064-216 environment, are illustrated in table 5.4 and figure 5.13 - 5.16.

No. Of OLRs	Avg OLR Elap. Time (mm:ss)	CPU Busy(%)	OLDS LOGG. Rate (MB/sec)	LOGL Latch (MB/sec)	DBBP Latch (count/sec)	IRLM Tbl Acc./ SynSvt(ms)
1	8:56	11.9	2.7	0	0	1,991/9.3
2	9:30	22.8	2.5	142.9	35	3,761/10
4	10:30	41.1	8.8	958.6	302	7,019/11.2
8 1 subpool	14:19	60.6	13.2	2,402	921.3	9,084/12.3
8 2 subpools	13:50	64.7	13.5	2,595	365.7	10,501/12. 5

Table 5.4: OLRs PHIDAM/OSAM studies with 1/2/4/8 Concurrent OLRs

Figure 5.13 illustrates the OLR PHIDAM/OSAM (1 GB partition) elapsed times for various concurrent OLRs - 1/2/4/8, in 2064-216 environment

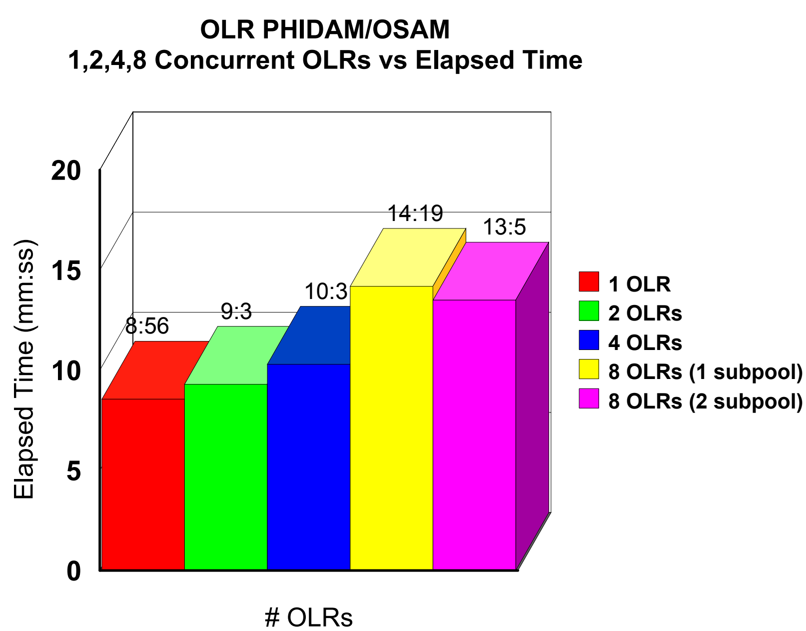

Figure 5.13: Concurrent OLRs vs Elapsed times

Figure 5.14 illustrates the OLR PHIDAM/OSAM(1 GB partition) logging rates for various concurrent OLRs - 1/2/4/8, in 2064-216 environment

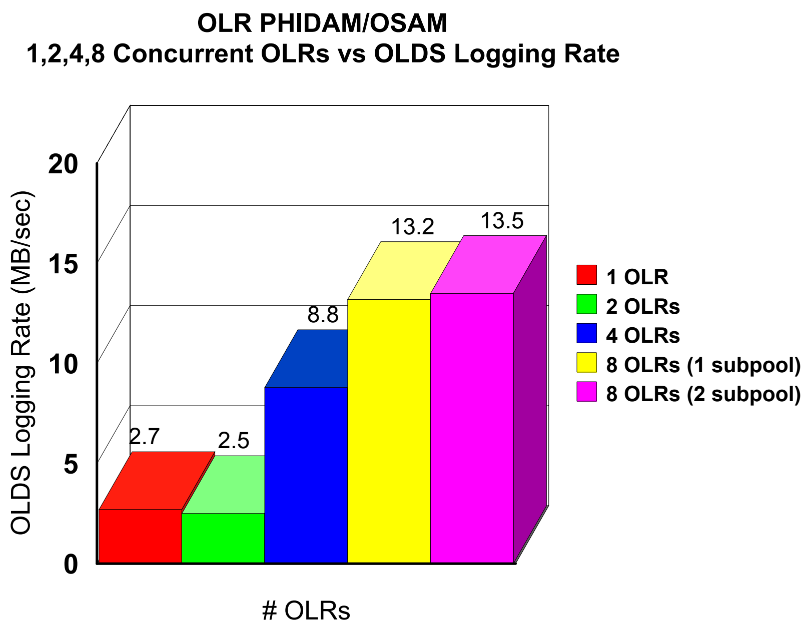

Figure 5.14: Concurrent OLRs vs Logging rates

Figure 5.15 illustrates the OLR PHIDAM/OSAM(1 GB partition) CPU busy for various concurrent OLRs - 1/2/4/8, in 2064-216 environment

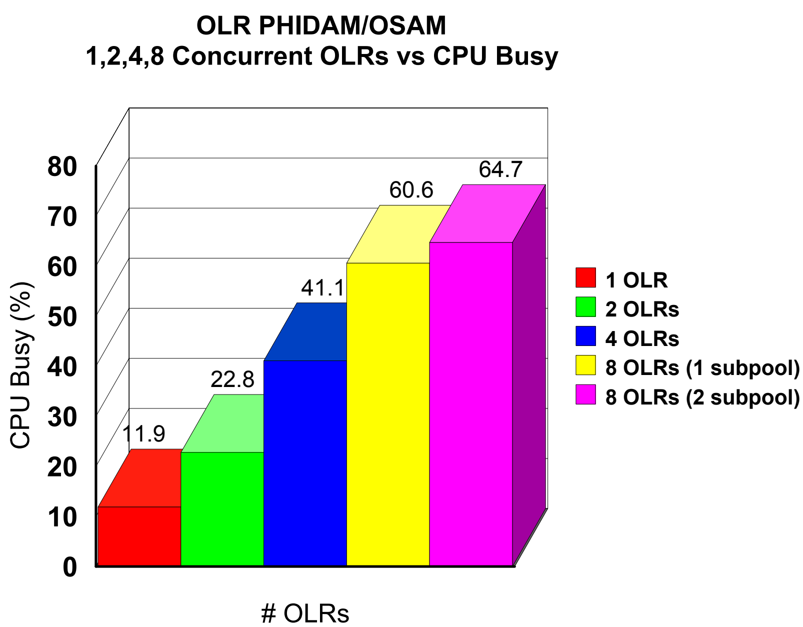

Figure 5.15: Concurrent OLRs vs CPU busy

Figure 5.16 illustrates the OLR PHIDAM/OSAM(1 GB partition) LOGL Latch for various concurrent OLRs - 1/2/4/8, in 2064-216 environment

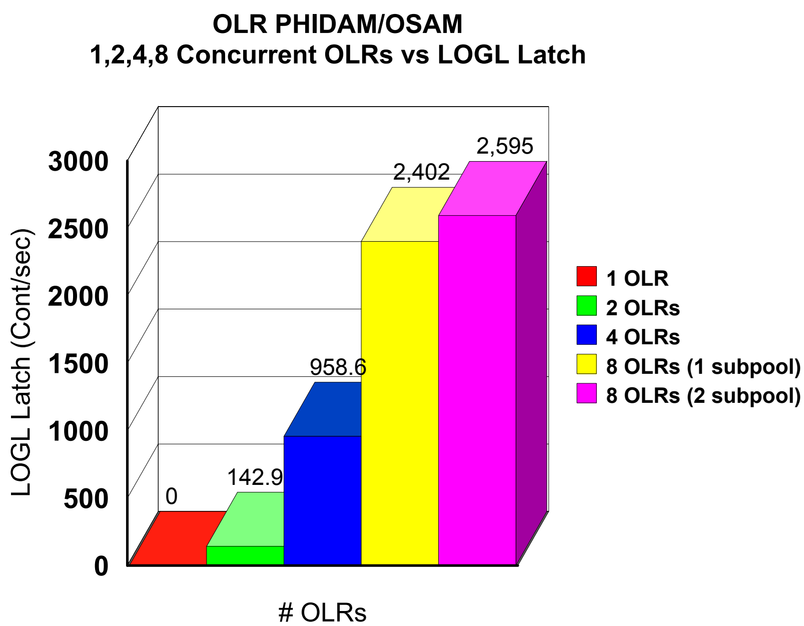

Figure 5.16: Concurrent OLRs vs LOGL latch

The results of OLR PHIDAM/OSAM (1 GB partition) studies with concurrent OLRs - 1/2/4/6/8, in 2084-316 environment, are illustrated in table 5.5 and figure 5.17 - 5.20.

No. of OLRs	OLR Elap. Time (mm:ss)	CPU Busy(%)	OLDS LOGG. Rate (MB/sec)	LOGL Latch (MB/sec)	DBBP Latch (count/sec)
1	7.59	10.09	3.65	0	0
2	8:15	19.56	5.92	149.51	27
4	8:43	38.51	11.23	1,011.15	239.17
6	10:14	50.82	14.34	1,990	500.06
8	12:26	55.92	15.74	2,626	738.06

Table 5.5 OLRs PHIDAM/OSAM studies with 1/2/4/6/8/ Concurrent OLRs

Figure 5.17 illustrates the OLR PHIDAM/OSAM(1 GB partition) elapsed times for various concurrent OLRs - 1/2/4/6/8, in 2084-316 environment

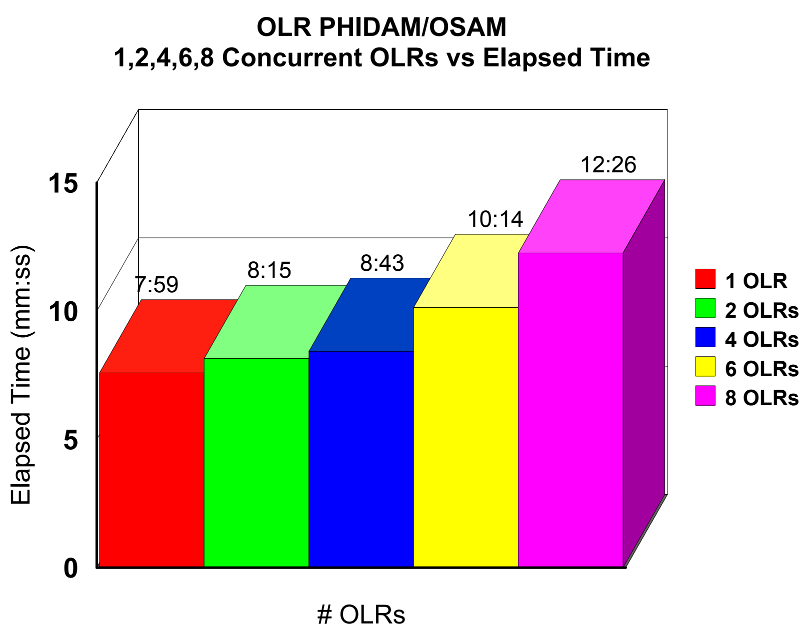

Figure 5.17: Concurrent OLRs vs Elapsed times

Figure 5.18 illustrates the OLR PHIDAM/OSAM(1 GB partition) logging rates for various concurrent OLRs - 1/2/4/6/8, in 2084-316 environment

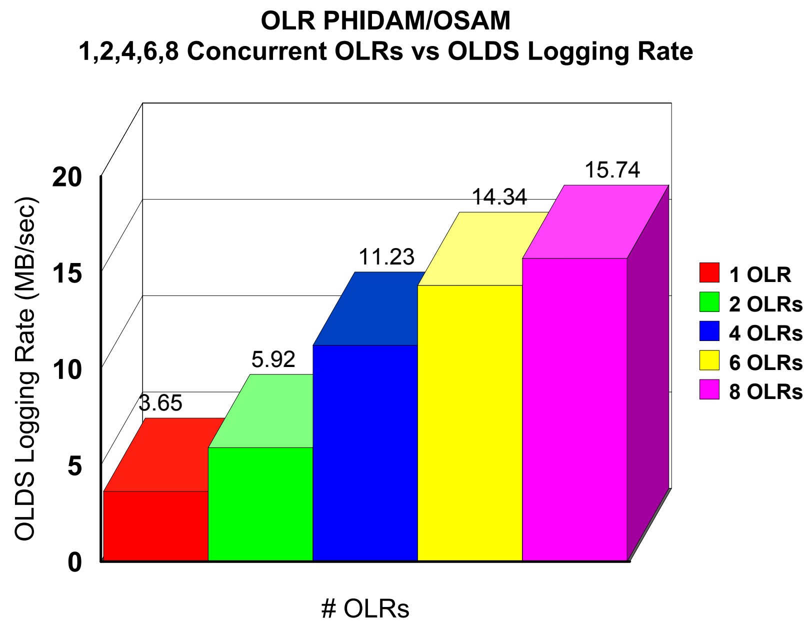

Figure 5.18: Concurrent OLRs vs Logging rates

Figure 5.19 illustrates the OLR PHIDAM/OSAM(1 GB partition) CPU busy for various concurrent OLRs - 1/2/4/6/8, in 2084-316 environment

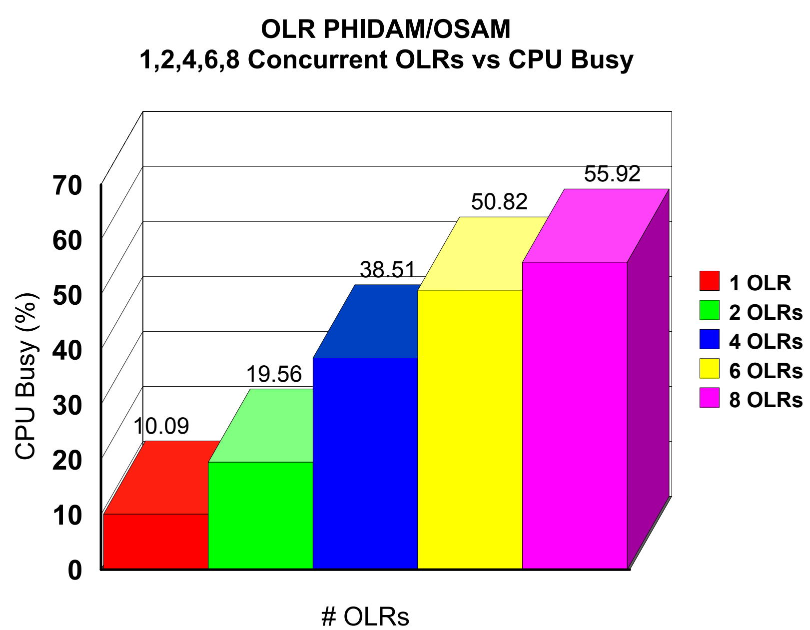

Figure 5.19: Concurrent OLRs vs CPU busy

Figure 5.20 illustrates the OLR PHIDAM/OSAM(1 GB partition) LOGL latch for various concurrent OLRs - 1/2/4/6/8, in 2084-316 environment

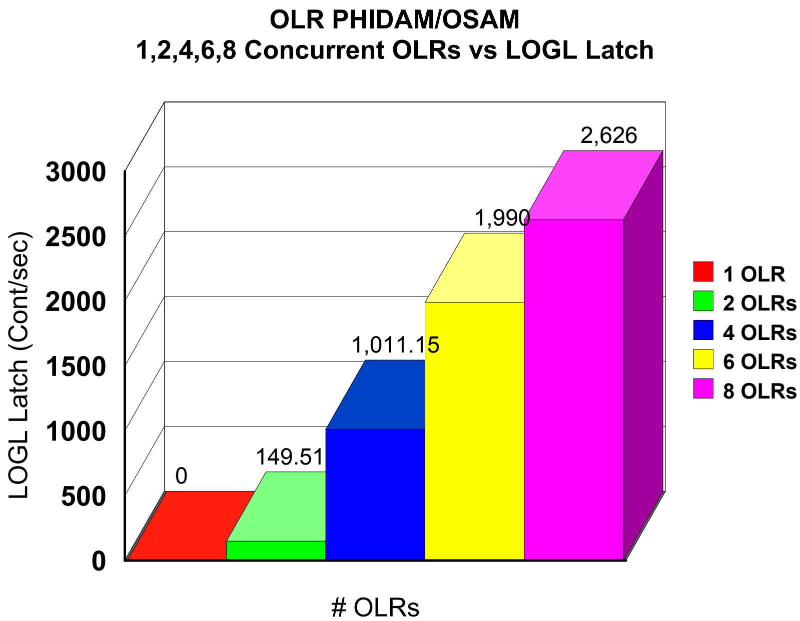

Figure 5.20: Concurrent OLRs vs LOGL latch

The results of OLR PHIDAM/OSAM online impact studies, in 2064-216 and 2084-316 environment, are illustrated in tables 5.6 - 5.7.

	Base 1-w w/o OLR	1-w Online with 2 xOLRs	1-w Online with 3 xOLRs
CPU Busy (%)	65.1	83.3	85.4
Tran rate -- Overall (tx/sec)	199.6	195.3	196.1
Tran rate -- against OLR part. (tx/sec)	8	7.6	7.8
Transit time Overall (msec)	310	393	410
Transit time Against OLR part.(msec)	195	252	255
Avg OLR elap. Time (mm:ss)	-	11:55	15:34
Logging rate(MB/sec)	1.8	5.9	6.6
LOGL (cont/sec)	23	211	251
DBBP (cont/sec)	1	21	25

Table 5.6: OLRs PHIDAM/OSAM online impact studies in 2064-216 environment.

	Base 2-w w/o OLR	2-w Online with 4 xOLRs	2-w Online with 6 x OLRs	2-w Online with 7 x OLRs
CPU Busy (%)	66.25	88.18	90.91	92.22
Tran rate -- Overall (tx/sec)	1,041.23	1,047.64	1,041.23	1,026.13
Tran rate -- against OLR part. (tx/sec)	6.72	6.55	6.58	6.57
Transit time Against OLR part.(msec)	291	286	303	333
Avg OLR elap. Time (mm:ss)	-	15:08	20:03	22:21
Logging rate(MB/sec)	8.64	15.04	15.96	16.45
LOGL (cont/sec)	494.56	1,317.67	1,702.15	1,950.79
DBBP (cont/sec)	75.78	197.05	286.27	337.23

Table 5.7: OLRs PHIDAM/OSAM online impact studies in 2084-316 environment

The results of OLR BMP impact studies, in 2084-316 environment are illustrated in table 5.8 and Figure 5.21

	2 LPAR w 1 BMP w/o OLR	2 LPAR w 1 BMP w OLR
Elap. time (mm:ss)	34:03	35:36
Impact on Elap. time(%)		-4.55

Table 5.8: OLR BMP Impact

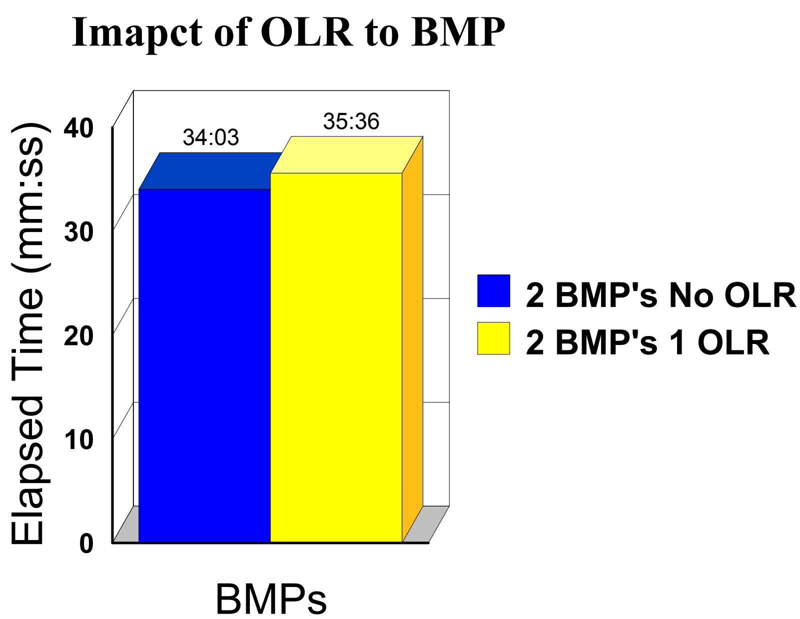

Figure 5.21:
BMP Elapsed time comparison

5.4 Summary

OLR improves the availability of databases. Without OLR, database information is unavailable during database reorganization. OLR overcomes this limitation by providing a nondisruptive integrated online reorganization for HALDB databases.

Rate parameter provides an ability (a command) to dynamically slow down, or speed up, an OLR in execution. Thus it can reduce the OLR resource requirement in CPU busy, OLDS logging bandwidth, Coupling Facility structures and buffer pool accesses. RATE(50) uses approximately 50 percent of resources when compared to RATE(100), but the elapsed time would be approximately doubled.

Executing multiple OLRs concurrently in an IMS region affect OLR total elapsed time, however, normal workloads are sustainable under 4 concurrent OLRs. Second subpool should be considered to relieve DB buffer contention when executing more than 4 concurrent OLRs.

6 Fast Path Area Open /Close Enhancements

6.1 Introduction

To increase parallelism during DEDB area open, preopen, or close IMS Version 9 now uses up to Ten TCBs during those Fast Path functions. Adding to that ability, IMS Version 9 allows the decision to reopen DEDB areas automatically after IMS warm or emergency restart or IRLM reconnect.

6.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2064 Model 216 (zSeries 900) - 12 GB storage, 3 CPs
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4, IMS V8.1, IMS V9.1

6.3 Results

The IMS restart times were compared by using Fast Path area preopen and reopen functions and the IMS normal termination (/CHE FREEZE) was used to compare the Shutdown times. These comparisons were based on reaching a ‘ready for work’ state which IMS Version 9 achieved upon restarts when compared to IMS Version 8.

Table 6.1 and Figure 6.1-6.3 illustrate the results of normal and emergency restart times for IMS Version 9 in a single image environment. The multiple TCBs and reengineering of the area open/close code contributed to the improvements.

	IMS V8 Elapsed time (mm:ss)	IMS V9 Elapsed time (mm:ss)	Improvement
/NRE (Cold)	10:20	7:18	29.35 percent
/NRE (Warm)	10:47	7:24	31.38 percent
/ERE	11:32	9:34	17.05 percent
/CHE FREEZE	3:37	3:36	< 1 percent

Table 6.1: Fast Path Area Open/Close Enhancements improvements

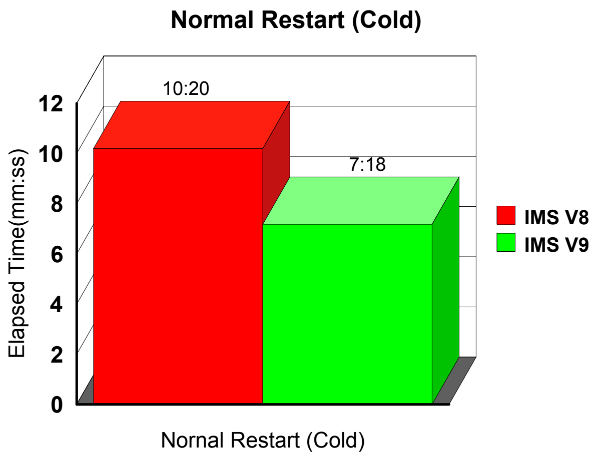

Figure 6.1: Normal restart (cold) Elapsed time comparison

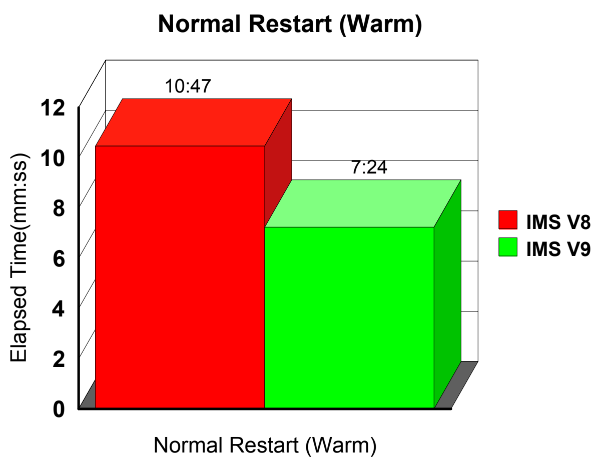

Figure 6.2: Normal restart (warm) Elapsed time comparison

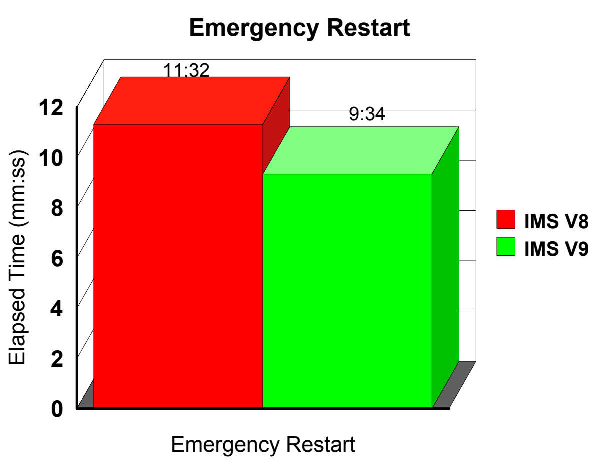

Figure 6.3: Emergency Restart - Elapsed time comparison

Figure 6.4 shows some improvement to normal shutdown time (/CHE FREEZE) for IMS Version 9 when compared to IMS Version 8.

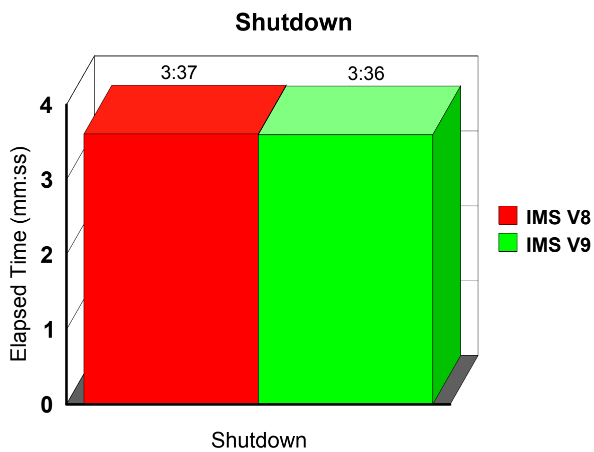

Figure 6.4: Checkpoint freeze command - Elapsed time comparison

6.4 Summary

IMS Version 9 Fast Path Area Open/Close Enhancements provide the following improvements when compared to IMS Version 8:

- Elapsed time for normal restart time (cold) has reduced by 29.35 percent
- Elapsed time for normal restart time (warm) has reduced by 31.38 percent
- Elapsed time for emergency has reduced by 17.05 percent
- Elapsed time for normal shutdown (/CHE FREEZE) has reduced less than 1 percent

7 Fast Path DEDB Shared VSO Multi-Area Structure (MAS) Support

7.1 Introduction

Fast Path DEDB areas can reside in Coupling Facility structures by using the base Shared VSO Structure support. Multiple structures are allowed with only one area per structure. In an environment where there are many Shared VSO structures it becomes a usability issue to manage all the structures. In IMS V9, Shared VSO MAS, enhances the usability of Shared VSO Structure Support, and overcomes this limitation by allowing multiple areas to share a single Coupling Facility Structure.

7.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2064 Model 216 (zSeries 900) - 24 GB storage, 6 CPs; IBM eServer zSeries 2064 Model 216 (zSeries 900) - 2 x ICFs, 2 CPs and 8 GB per ICF (for shared vso test)
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4, IMS V8.1, IMS V9.1

7.3 Results

The study was performed by comparing the performance characteristics of Fast Path DEDB Shared VSO Multi-Area Structure Support which had one Coupling Facility structure (four areas per structure), to Fast Path DEDB base Shared VSO Structure Support which had four Coupling Facility structures (one area per structure).

The performance characteristics of Fast Path DEDB Shared VSO Multi-Area Structure Support and Fast Path DEDB base Shared VSO Structure Support are illustrated in figure 7.1.

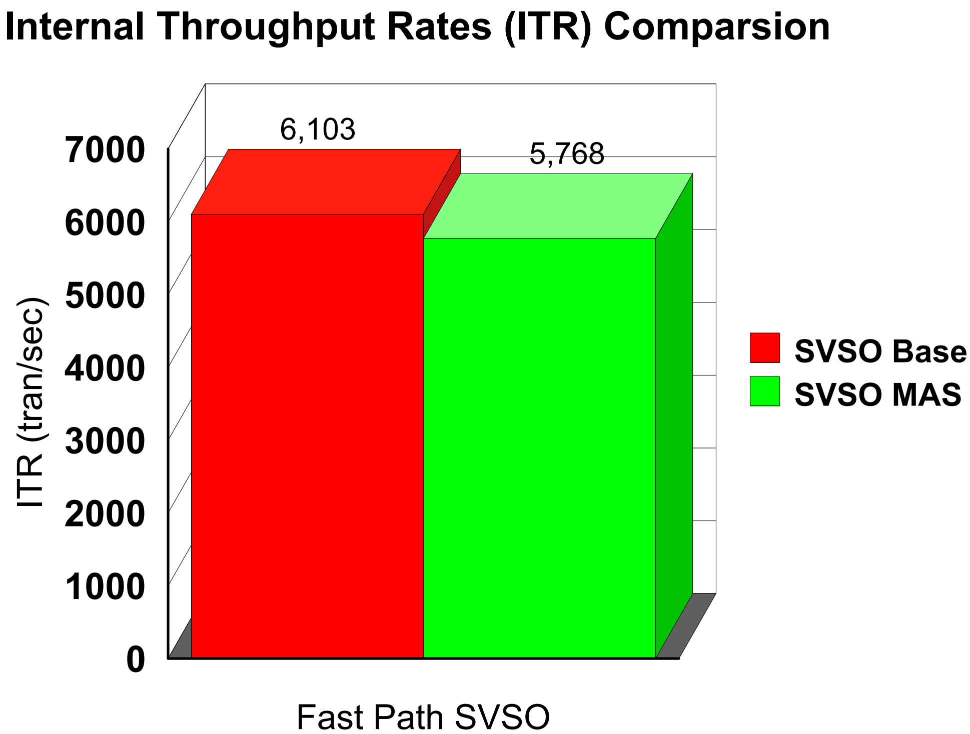

Figure 7.1: ITRs comparison

7.4 Summary

While enhancing the usability for Fast Path Shared VSO Structure Support, the cost incurred by Fast Path DEDB Shared VSO MAS in terms of the ITR is at 5.4 percent.

8 DBRC Enhancements

8.1 Introduction

The performance characteristics of the DBRC enhancements using Large Blocksize support - greater than 32 K - are identified.

8.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2064 Model 216 (zSeries 900) - 12 GB storage, 3 CPs;
2105-M800 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4, IMS V8.1, IMS V9.1

8.3 Results

The elapsed time improvements for Image Copy and Recovery Utility are illustrated in table 8.1 and Figures 8.1 - 8-3. The DBRC Enhancements (Large Blocksize support - greater than 32 K) contributed to the improvements.

	Image Copy Elapsed time (mm:ss)	Recover (DFSURDB0) Elapsed time (mm:ss)
IMS V8 TAPEBLKSZLIM=32,760	22:44	37:27
IMS V9 TAPEBLKSZLIM=32,760	22:46	36:10
IMS V9 TAPEBLKSZLIM=2G BLKSZLIM=2G on IC Job	19:53	32:45
Improvement	12.53 percent	12.55 percent

Table 8.1: Elapsed time improvements

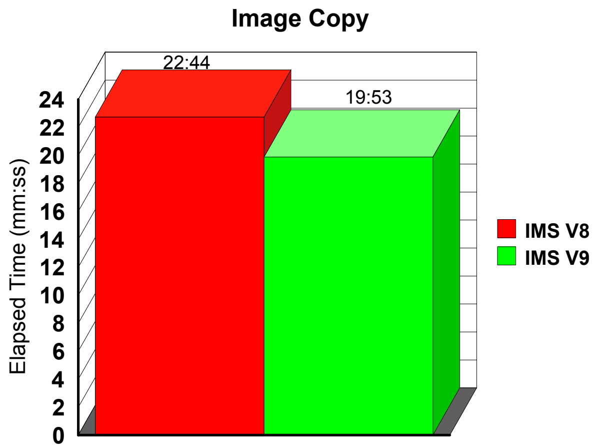

Figure 8.1: Elapsed times comparison

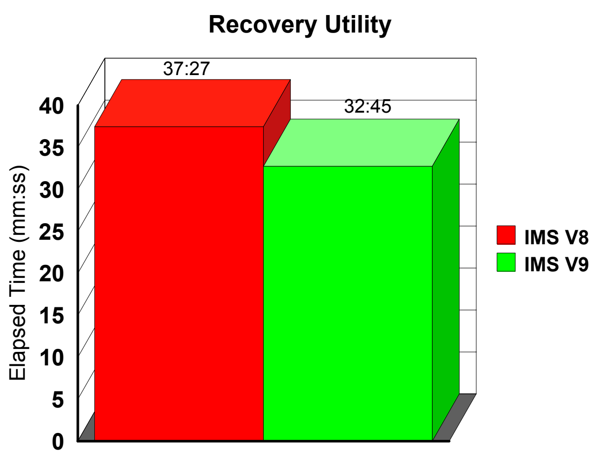

Figure 8.2: Elapsed Times comparison

8.4 Summary

IMS Version 9 with DBRC Enhancements has improved the elapsed times for Image Copy and Recovery Utility by 12.53 and 12.55 percent respectively when compared to IMS Version 8.

9 Fast Path High Stress

9.1 Introduction

This section shows the performance characteristics of IMS Version 9 Fast Path in the parallel sysplex, Shared EMHQ, 4-way datasharing environment.

9.2 Test Environment

The description of the test environment is as follows:

Hardware and Software Environment	
Processor:	IBM eServer zSeries 2084 Model 316 (zSeries 990) - 48 GB storage, 12 CPs ; IBM eServer zSeries 2084 Model 316 (zSeries 990) - 2 x ICFs, 2 CPs and 8 GB per ICF (for sysplex test)
2105-M00 Disk:	IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800: 36GB 15K RPM disk drives, 4 FICON channels, 8 LSS, 16 Ranks, Volume config = 3390-9 , 7 PAVs per real volume
Operating Systems:	z/OS DFSMS V1R4, IMS V8.1, IMS V9.1

9.3 Results

The Fast Path High Stress was performed using FP2 workload in the parallel sysplex environment: IBM eServer zSeries 2084 Model 316 with IBM TotalStorage Enterprise Storage Server (ESS) 2105 Model 800; 4-way datasharing; And Shared EMHQ. The parallel sysplex configuration used is illustrated in figure 9.1.

Figure 9.1: Fast Path High Stress - Parallel sysplex configuration

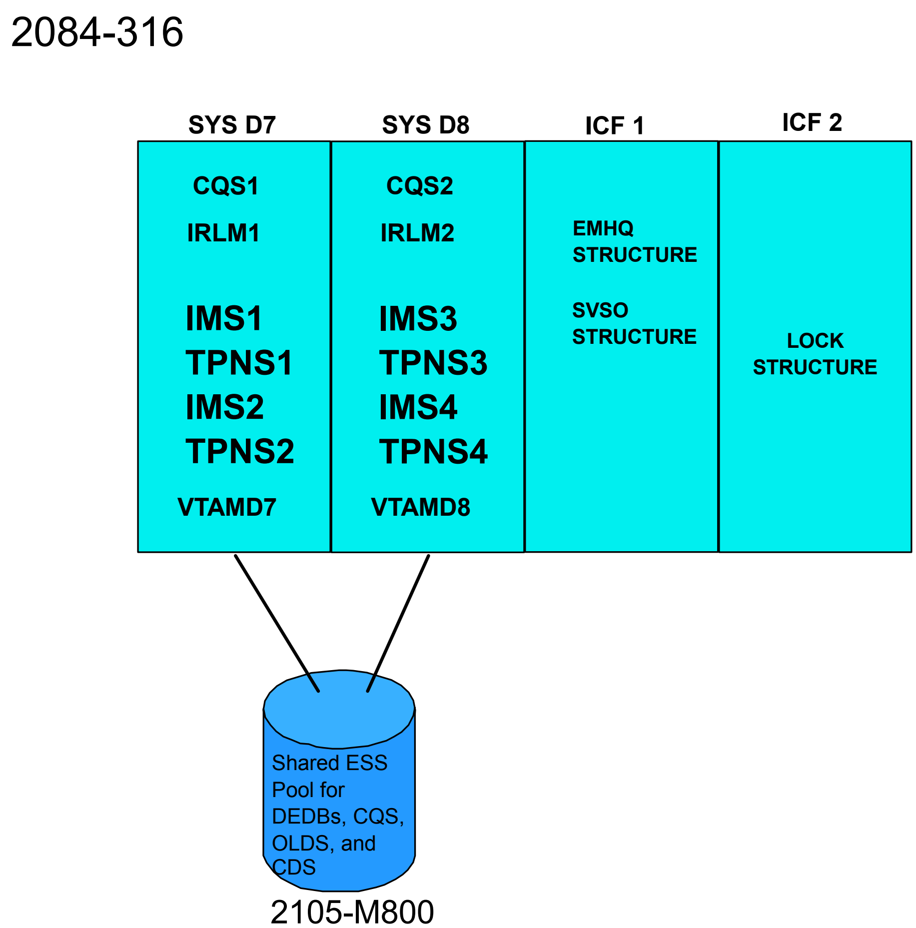

All performance results were gathered during a steady state and the details are illustrated in figure 9.2.

Tran Rate -

21,396 per second (nearly 2 billion per day)

Total DASD I/O rate -

27,448 I/Os per sec

28.8 MB/sec Logging bandwidth

6.55 ms average response time

Total CF utilization -

28.0 % of 4 cps

Total CPU utilization -

99.65 % of 12 cps

Figure 9.2: Fast Path High Stress results

9.4 Summary

The transaction rate achieved by IMS Version 9 Fast Path, executing in the parallel sysplex environment, displayed a significant improvement from the benchmark recorded in the previous version of IMS - It represents 81.56 percent improvement in the transaction rate.

10 Test Methodology

The test methodology used in the IMS Version 9 performance study is similar to the methodology described in the IBM Large Systems Performance Reference, document number SC28-1187-09, 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/SC2811879.pdf

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

10.1 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.

10.2 Trademarks

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

IBM®	Enterprise Storage Server	FICON
IMS	z/OS	zSeries