IMS 15 Performance Evaluation Summary

IMS Performance Evaluation Team (IBM Silicon Valley Laboratory)
October 2018

Contributors: Richard Antikoll, Kevin Hite, Bruce Naylor, Jasdeep Singh, Justin Soe, and Jane Yang

Abstract

This paper presents the performance characteristics of IMS™ Version 15 (IMS 15), as evaluated by the IMS Performance Evaluation Team at IBM Silicon Valley Laboratory. It provides comparisons of IMS 15 performance with the prior IMS release's (Version 13 and 14) performance under the same hardware and workload configurations. It also examines performance characteristics for specific IMS 14 and IMS 15 enhancements.

IBM Z
San Jose, CA

Note: Before you use this information and the product it supports, read the information in "Notices" on page 145.

Note: Performance is based on measurements and projections using IMS benchmarks in a controlled environment. The results that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, the amount of zIIP capacity available during processing, and the workload processed. Therefore, results may vary significantly, and no assurance can be given that an individual user will achieve results similar to those stated here. Results should be used for reference purposes only.

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.

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.

The information in this paper provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Purpose Engines/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

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

1 Introduction

This paper documents the performance evaluation of IMS 13, 14, and 15 conducted by the IMS Performance Evaluation Team at the IBM Silicon Valley Laboratory. IMS 15 is compared with the prior version (IMS 14) using different workload types under the same hardware configurations. Although IMS 13 performance results are not directly compared with IMS 15, the performance data is included in this paper since the workloads are run against the latest IBM Z-family processor, z14. Additionally, several specific IMS 14 and IMS 15 enhancements are evaluated from a performance perspective.

As with all IMS releases, IMS 15 contains functional enhancements that affect all areas of the product: IMS Database Manager, IMS Transaction Manager, and IMS Systems. Some of these changes that may affect performance include:

- Support for 64-bit Java Virtual Machine (JVM) for Java dependent regions
- Write-Ahead Data Sets (WADS) encryption
- WADS and Online Data Sets (OLDS) I/O processing through High Performance FICON® for z Systems™ (zHPF) and zHyperWrite
- Management of runtime Application Control Blocks (ACBs) for database and program views.
- Network Security Credential Propagation

The performance results, due to these and other changes in the IMS software, depends on the installation's specific application, transaction, and database characteristics.

2 Executive Overview

The IMS 15 performance workload evaluation demonstrates that the latest release of IMS has similar overall performance compared to IMS 14 while providing additional functionality and usability features. However, some IMS 15 enhancements provide significant improvements and expands IMS features. Section 7 IMS 15 Enhancements describes the benefits of these enhancements.

Customers migrating to IMS 15 from a prior IMS release should expect a similar performance outcome when operating on the same hardware. An increase in performance can be gained when upgrading to the latest IBM Z processor as explained in the “IMS 14 Performance Benchmark on z14” white paper published in February 2018. (See IMS Resources for a link to the z14 white paper.)

While performance in specific production environments will vary, results of IBM internal testing in a controlled laboratory environment revealed that, depending on the specific workload and other factors, IMS 15 is capable of performing as described below:

- IMS 15 demonstrated a reduction in IMS Connect service time of up to 18% for some performance workloads.

- Enhancements in IMS Logger functions decreased WADS and OLDS I/O response times of greater than 60% when utilizing both zHPF and zHyperWrite.

- The encryption of WADS, OLDS, and DEDBs with the Fast Path workload had a CPU cost of about 2% compared to without encryption.

3 IMS Performance Environment Information

This section explains the Performance evaluation environment at Silicon Valley Lab including hardware and software specifications.

3.1 Hardware Environment

All measurements were conducted on the z14 as shown in Figure 1: IBM z14 Model 3906 (M05).

3.1.1 z14 Processor

The z14 processor is the newest member of the IBM Z family that provides higher capacity and processing power, efficient pervasive encryption capability, and micro architecture improvements compared to the z13 processor.

The main features of the z14 technology include:

- 14nm 17 metal layers technology node
- 5.2 GHz system frequency
- 10-core processor chip
- 6-processor chips + 1-SC chips per drawer
- 32 TB max memory capacity
- 4-drawer 170 cores system max configuration
- 192 GB Hardware System Area (HSA)
- Levels of cache:
  - First-level cache (L1) private: 128 KB for instructions, 128 KB for data
  - Second-level cache (L2): 2 MB for instructions, 4 MB for data
  - Third-level cache (L3): 128 MB
  - Fourth-level cache (L4): 672 MB

10% more capacity than z13
32 TB max memory capacity
5.2 GHz core processing speed

Figure 1: IBM z14 Model 3906 (M05)

3.1.2 Storage

IBM System Storage® DS8000® series latest model DS8886 (16 FICON® channel paths) were used with 64 real volumes and 128 alias volumes per LCU using dynamic Hyper Parallel Access Volumes (HyperPAVs). The DASD volumes and paths for each measurement evaluation remained consistent for comparison purposes.

3.1.3 Coupling Facilities

All measurements were performed with Internal Coupling Facilities (ICFs) with Coupling Facility Control Code (CFCC) level 22.

3.2 Software Environment

z/OS: The performance evaluations were performed on z/OS® Version 2 Release 2 (5650-ZOS).

IMS: IMS™ Version 13, 14, and 15

IRLM: IRLM Version 2.3 or later (5635-A04)

CICS subsystems: Customer Information Control System (CICS®) Transaction Server for z/OS® Version 5 Release 2 connecting to IMS Database Manager (DB).

Workload Driver: Teleprocessing Networks Simulator (TPNS), Version 3 Release 5.0, Service Level 9711, was used as the workload driver simulating TCP/IP clients.

Java: Java 8 Service Release 5.

Java Workload Driver: Java-based workload driver simulating client connections and requests.

Db2: Db2® Version 12

4 IMS Workloads

IMS supports multiple types of enterprise databases and communication access methods, so that customers can exploit the technology best suited for their requirements. This section describes the various types of workloads used in the IMS 15 performance evaluations to exercise specific IMS code paths listed below:

- Fast Path (FP)
- Full Function (FF)
- Shared Message Queues (SMQ)
- Open Transaction Manager Access (OTMA)
- CICS 5.2 – IMS Database Control (DBCTL)
- Batch Message Processing (BMP)
- z/OS Connect Enterprise Edition V3.0.10
- IMS TM-Db2
- Open Database Manager (ODBM)
- Java Message Processing (JMP)

The workloads are used for comparison sets to maintain a consistent environment suitable for IMS version to version comparisons. These workloads were used to simulate typical banking, purchase, warehouse, hotel, and inventory customer-like workloads.

4.1 Full Function (FF) with High Availability Large Database (HALDB) Workload

Full Function databases are accessed through the Data Language I (DL/I) call interface and can be processed by application programs running in IMS dependent regions, IMS batch regions, CICS, z/OS WebSphere Application Server (WAS), Db2® Stored Procedures, and through the IMS Open Database Access (ODBA) interface. IMS dependent region types include:

- IMS Fast Path Program (IFP)
- Message Processing Program (MPP)
- Batch Message Processing Region (BMP)
- Java Message Processing Region (JMP)
- Java Batch Processing Region (JBP)

Full function databases can store data using Virtual Storage Access Method (VSAM), a native z/OS access method, or Overflow Sequential Access Method (OSAM), an IMS-specific access method that optimizes the I/O channel program for IMS access patterns.

A High Availability Large Database (HALDB) is a partitioned Full Function database that allows the grouping of Full Function database records into sets of partitions or replicates that are treated as a single database.

This section details the database description, application transactions, and the workload distribution for the Full Function with HALDB workload.

4.1.1 Database Description

The Full Function with HALDB workload contains databases that are a mix of OSAM and VSAM with Hierarchical Direct Access Method (HDAM), Hierarchical Indexed Direct Access Method (HIDAM), Partitioned HDAM (PHDAM), and Partitioned Indexed HDAM (PHIDAM) databases. This workload has 32 replicates. Each replicate contains eight Full Function databases. Half of the replicates are OSAM while the other half are VSAM as described below:

Table 1: Full Function with HALDB Workload Database Description

<table>
<tr><th>Database Name</th><th>Database Overview</th></tr>
<tr><td>CUSTDA</td><td>CUSTDA is a HALDB customer database containing information for customers. A CUSTDA record is composed of the following segments:
• AMFROOT (root segment) – Contains information for customer
• Sixteen direct dependent segments containing dummy data
The following is a graphical view of the database layout:
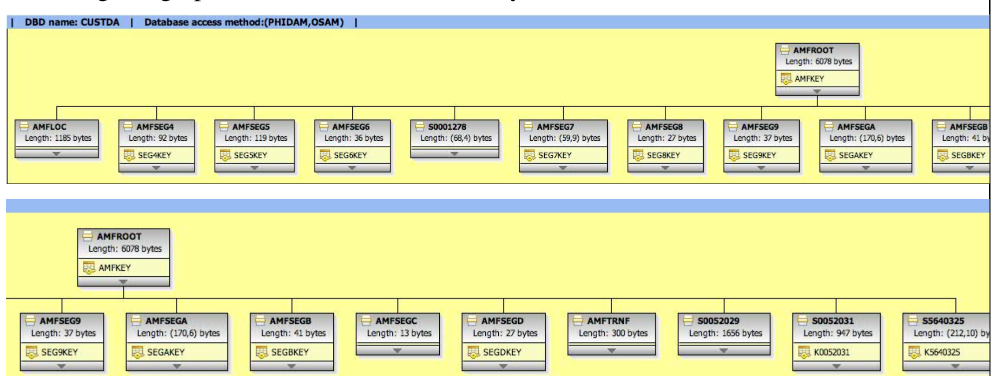
</td></tr>
<tr><td>CUSTOMRA</td><td>CUSTOMRA is a customer database containing information for the customer directory. A CUSTOMRA record is composed of the following segments:
• CUST001 (root segment key) - Contains information for customer
• ADDR002 (direct dependent segment key) - Contains information for address
• ORCOND003 (direct dependent segment key) - Contains information for order conditions
• ORDER004 (direct dependent segment key) - Contains information for order number
• BORD0005 (dependent segment key) - Contains information for back orders
• INV00006 (dependent segment key) - Contains information for invoice parts</td></tr>
</table>

<table>
<tr><td></td><td>

- OFLW0007 (dependent segment key) - Contains information for overflow parts

The following is a graphical view of the database layout:

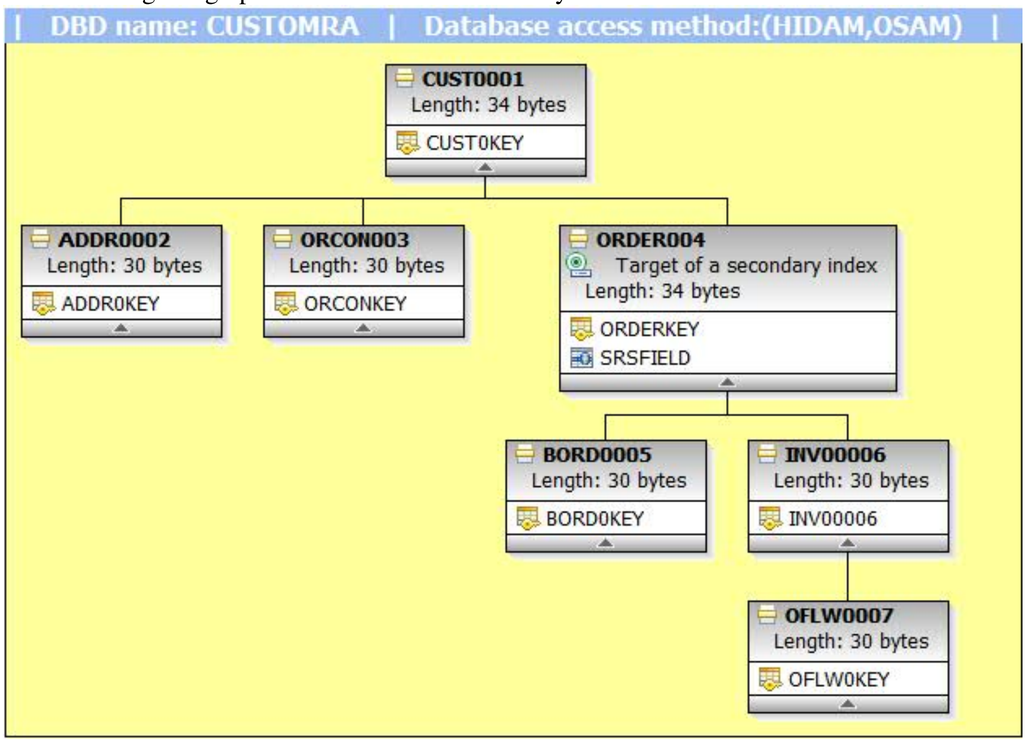

</td></tr>
<tr><td>INVENTRA</td><td>

INVENTRA is an inventory database containing information for all the orders and parts for the Stock Control application. A INVENTRA record is composed of the following segments:

- IN010PAR (root segment key) - Contains information for parts
- IN020REO (direct dependent segment key) - Contains information for reorders
- IN030ENG (direct dependent segment key) - Contains information for engineering
- IN040SLQ (direct dependent segment key) - Contains information for secondary location quantity
- IN050EOI (direct dependent segment key) - Contains information for EOI
- IN070BO (direct dependent segment key) - Contains information for back orders
- IN060SUP (dependent segment key) - Contains information for suppliers

The following is a graphical view of the database layout:
</td></tr>
</table>

<table>
<tr>
<td></td>
<td>
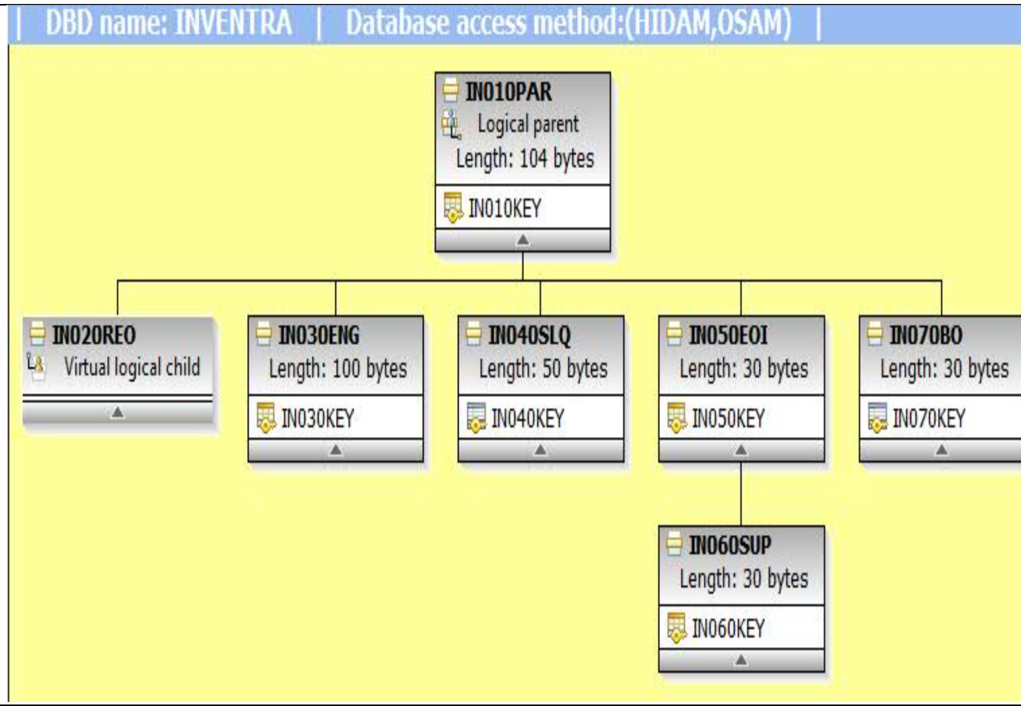

</td>
</tr>
<tr>
<td>HOTELDBA</td>
<td>
HOTELDBA is hotel reservation database pertaining information for the hotel reservation for specific requests like accommodation type, dates and location. A HOTELDBA record is composed of the following segments:

HOTEL001 (root segment key) - Contains information for hotel
ACCOM002 (direct dependent segment key) - Contains information for accommodation type
AHOTL006 (direct dependent segment key) - Contains information for alternate hotel
RESRV003 (dependent segment key) - Contains information for reservations
DATE0004 (dependent segment key) - Contains information for date
ATYPE005 (dependent segment key) - Contains information for alternate type

The following is a graphical view of the database layout:
</td>
</tr>
</table>

	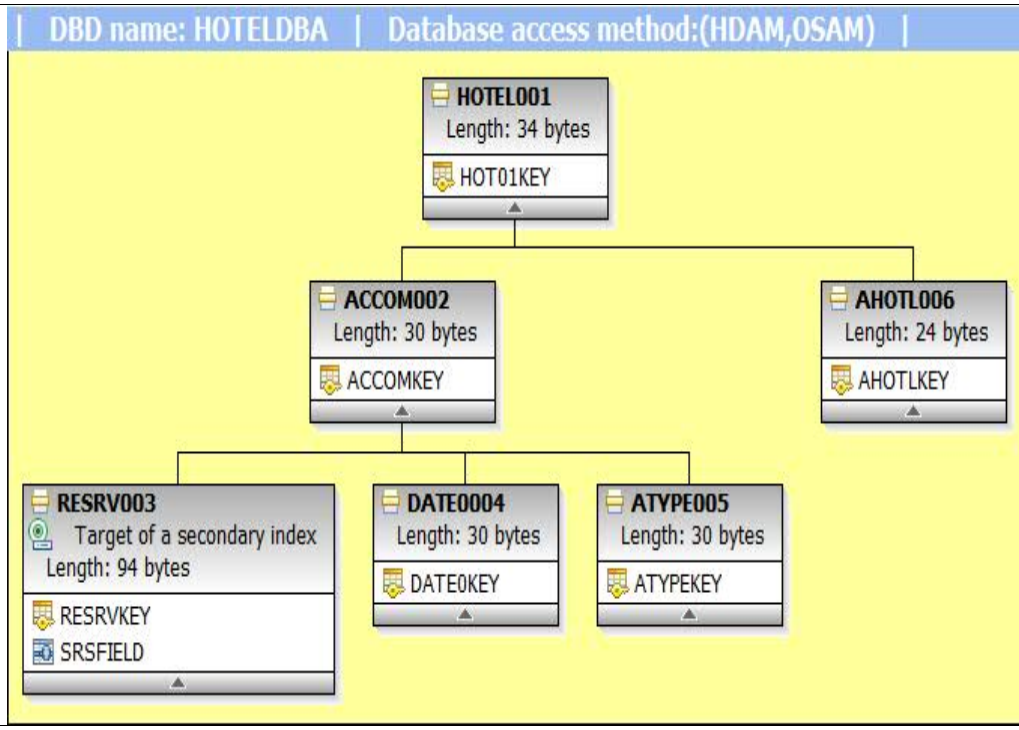
ITEMACTA	ITEMACTA is an item activity database containing information on receiving purchase order items, storage control and tracking material, as it moves through production operation units. A ITEMACTA record is composed of the following segments: • IA010ITM (root segment key) - Contains information for activity item number • IA020DTE (direct dependent segment key) - Contains information for activity date • IA060LOC (direct dependent segment key) - Contains information for location detail • IA080LNK (direct dependent segment key) - Contains information for purchase orders • IA070MOV (direct dependent segment key) - Contains information for move orders • IA040TRK (direct dependent segment key) - Contains information for track operations • IA030LOG (dependent segment key) - Contains information for activity log • IA050DIS (dependent segment key) - Contains information for disc material The following is a graphical view of the database layout:

	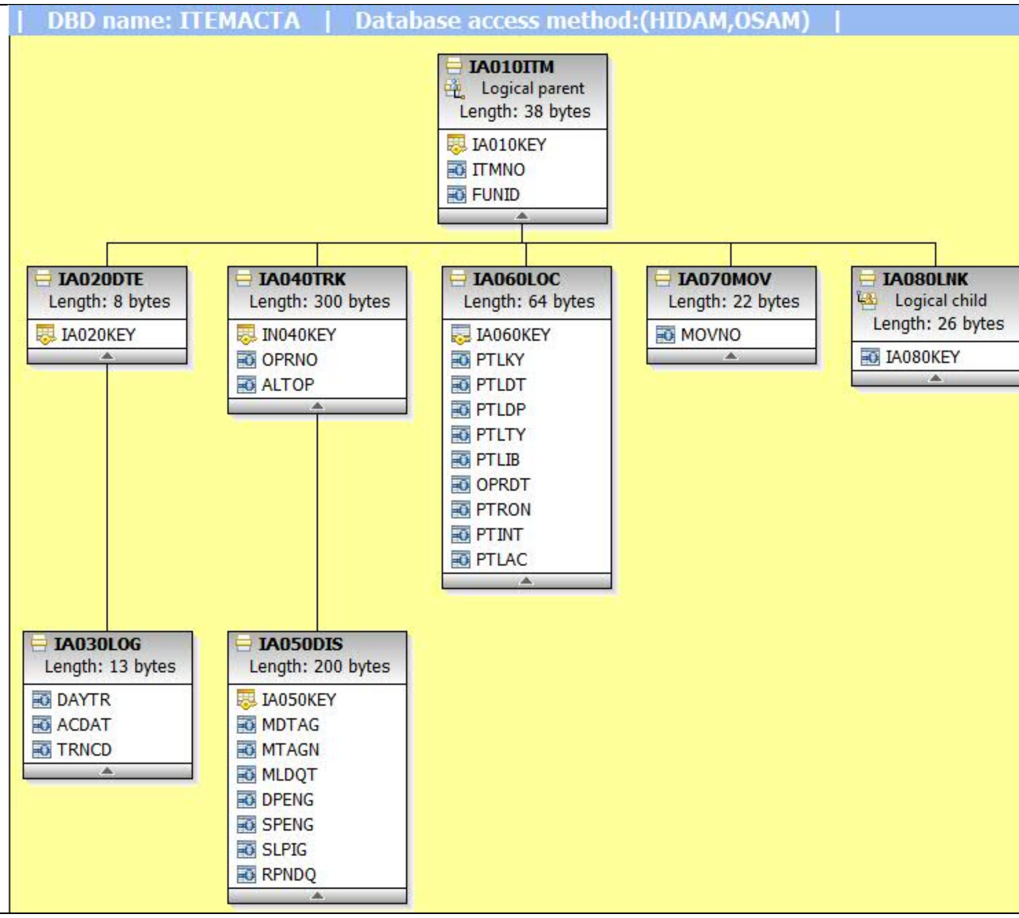
EMPACTDA	EMPACTDA is an employee activity database containing labor operations information for the employees, as the item is moved from one work station to another. A EMPACTDA record is composed of the following segments: • EMP00001 (root segment key) - Contains information for employee number • LAB00002 (direct dependent segment key) - Contains information for labor status The following is a graphical view of the database layout:

<table>
<tr><td></td><td>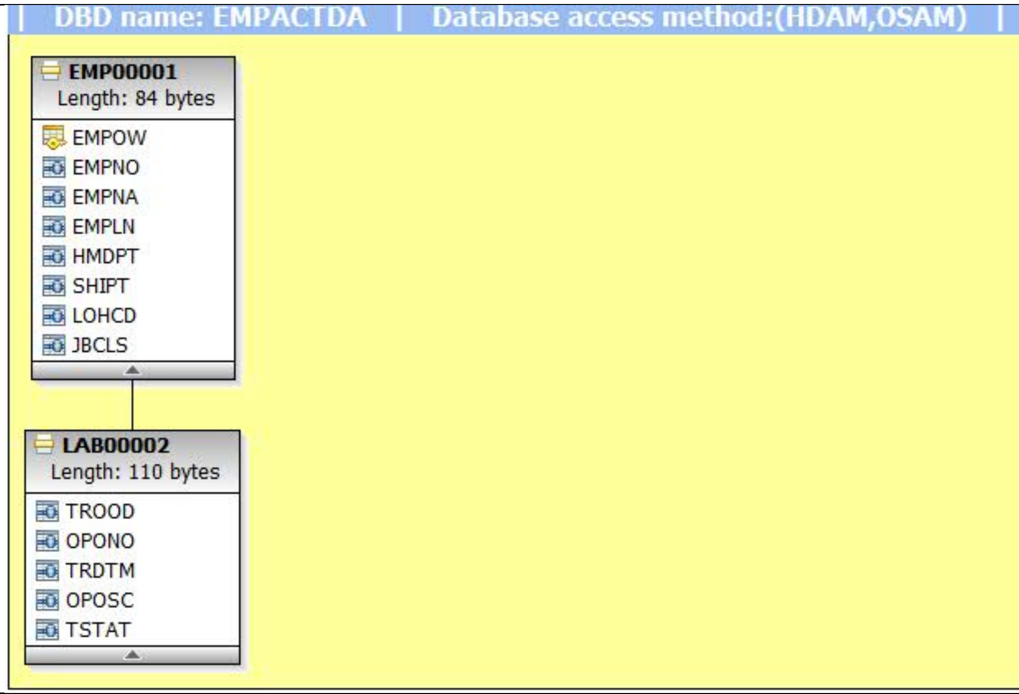
</td></tr>
<tr><td>ITEMMASA</td><td>PRODMSTR is a product item master database containing information for the product specifics. A PRODMSTR record is composed of the following segments:
• PRODMSTR (root segment key) - Contains information for product master
• PRODSTRC (direct dependent segment key) - Contains information for production structure
• PRODSTWC (direct dependent segment key) - Contains information for production schedule
• PRODLDTL (direct dependent segment key) - Contains information for product unit detail
• LABOROPS (direct dependent segment key) - Contains information for labor operations
• MATSPECS (direct dependent segment key) - Contains information for material specifications
• GENNOTES (direct dependent segment key) - Contains information for general notes
• IPYKOAL (direct dependent segment key) - Contains information for control
• HISTORY (direct dependent segment key) - Contains information for history
The following is a graphical view of the database layout divided into two graphics shown below:</td></tr>
</table>

	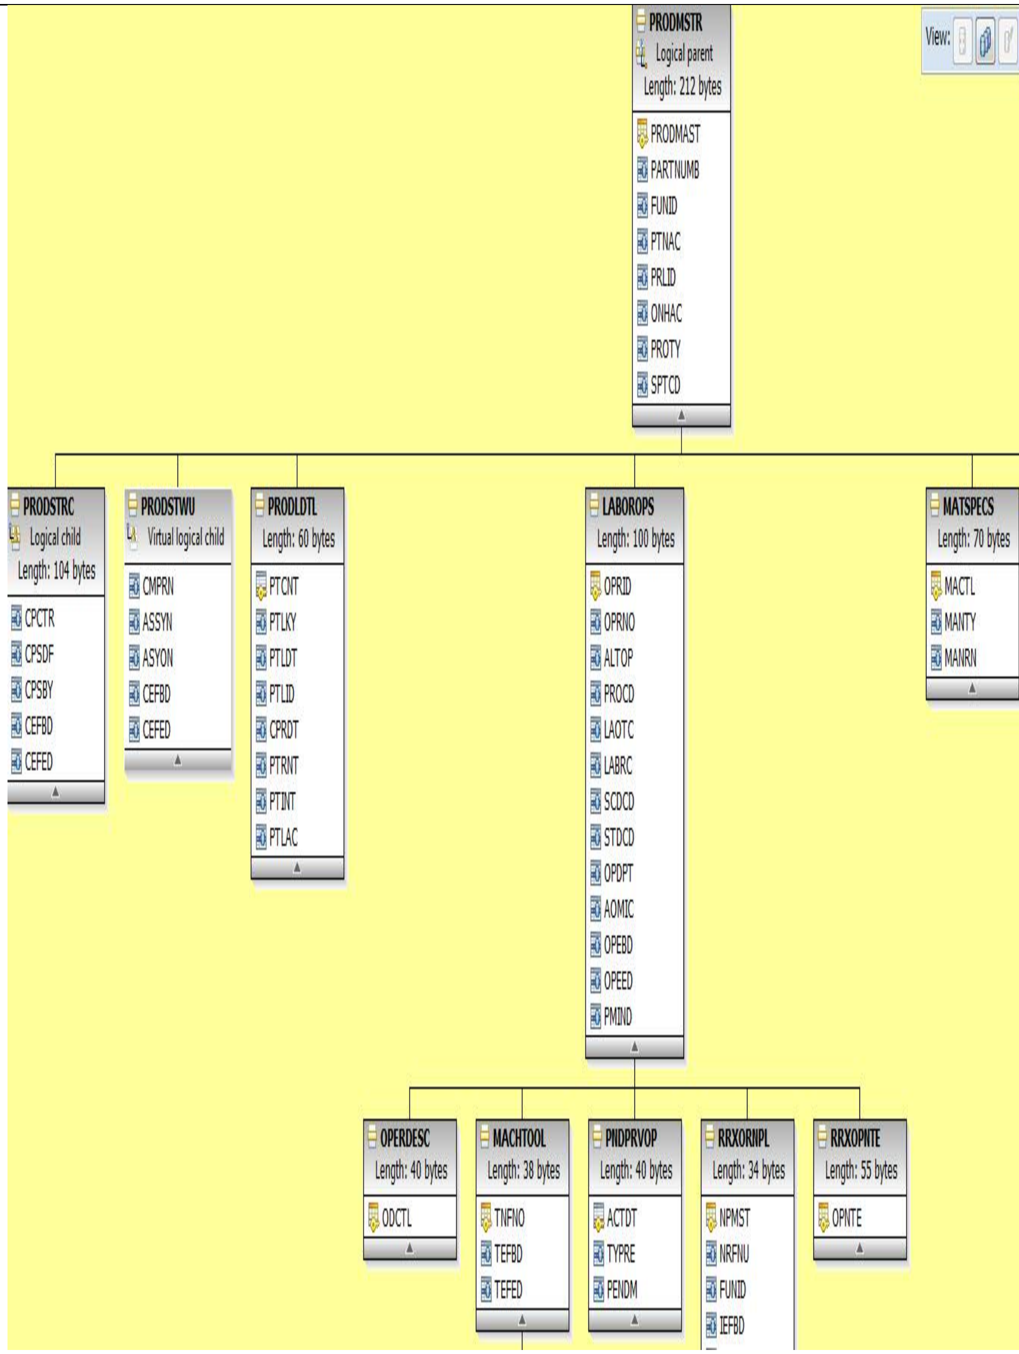View:
PRODMSTR
Logical parent
Length: 212 bytes
PRODMAST
PARTNUMB
FUNID
PTNAC
PRLID
ONHAC
PROTY
SPTCD
PRODSTRC
Logical child
Length: 104 bytes
CPCTR
CPSDF
CPSBY
CEFBD
CEFED
PRODSTWU
Virtual logical child
CMPRN
ASSYN
ASYON
CEFBD
CEFED
PRODLDTL
Length: 60 bytes
PTCNT
PTLKY
PTLDT
PTLID
CPRDT
PTRNT
PTINT
PTLAC
LABOROPS
Length: 100 bytes
OPRID
OPRNO
ALTOP
PROCD
LAOTC
LABRC
SCDCD
STDCD
OPDPT
AOMIC
OPEBD
OPEED
PMIND
MATSPECS
Length: 70 bytes
MACTL
MANTY
MANRN
OPERDESC
Length: 40 bytes
ODCTL
MACHTOOL
Length: 38 bytes
TNFNO
TEFBD
TEFED
PNDPRVOP
Length: 40 bytes
ACTDT
TYPRE
PENDM
RRXORNPL
Length: 34 bytes
NPMST
NRFNU
FUNID
IEFBD
RRXOPNTE
Length: 55 bytes
OPNTE

	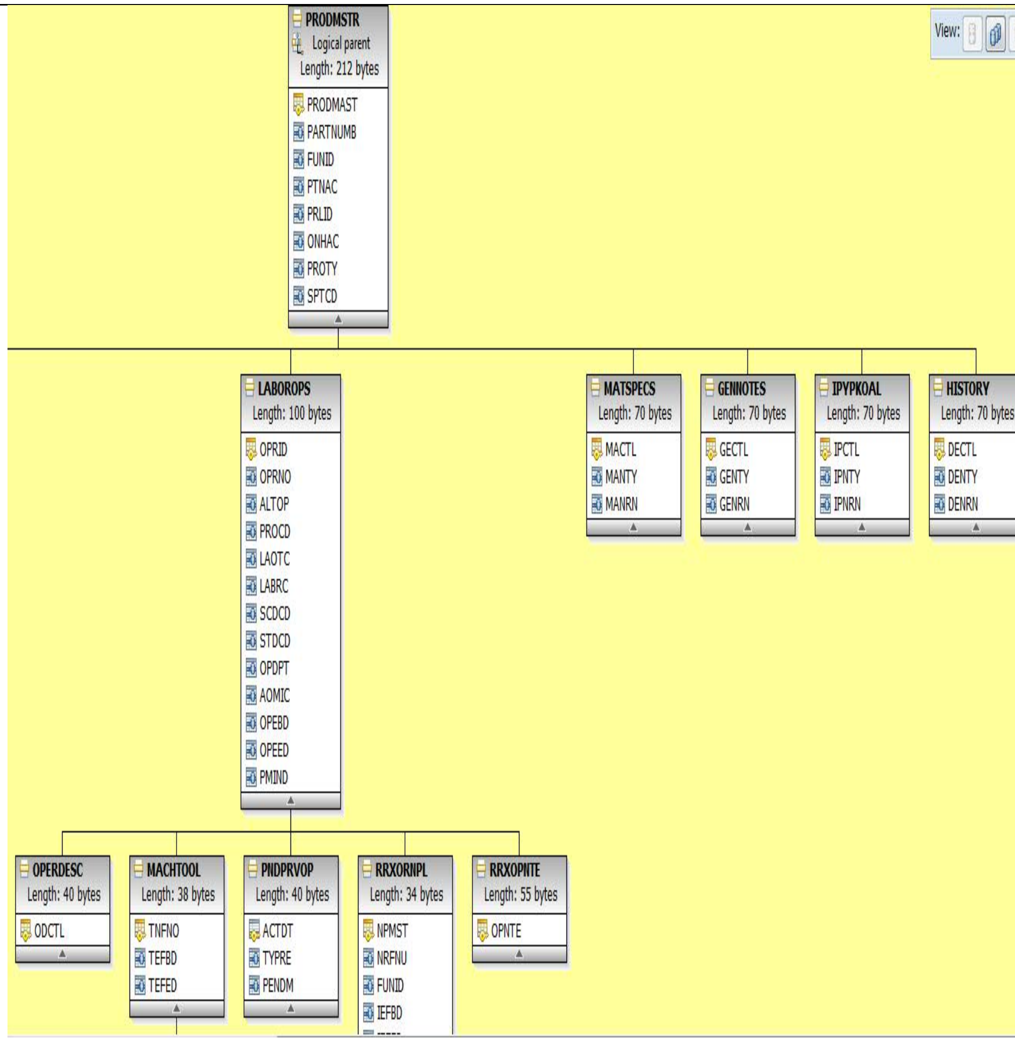
COMPOSDA	COMPOSDA is component position database containing information for the part numbers and their position data. A COMPOSDA record is composed of the following segments: • PAR00001 (root segment key) - Contains information for part number • POS00002 (direct dependent segment key) - Contains information for position data The following is a graphical view of the database layout:

	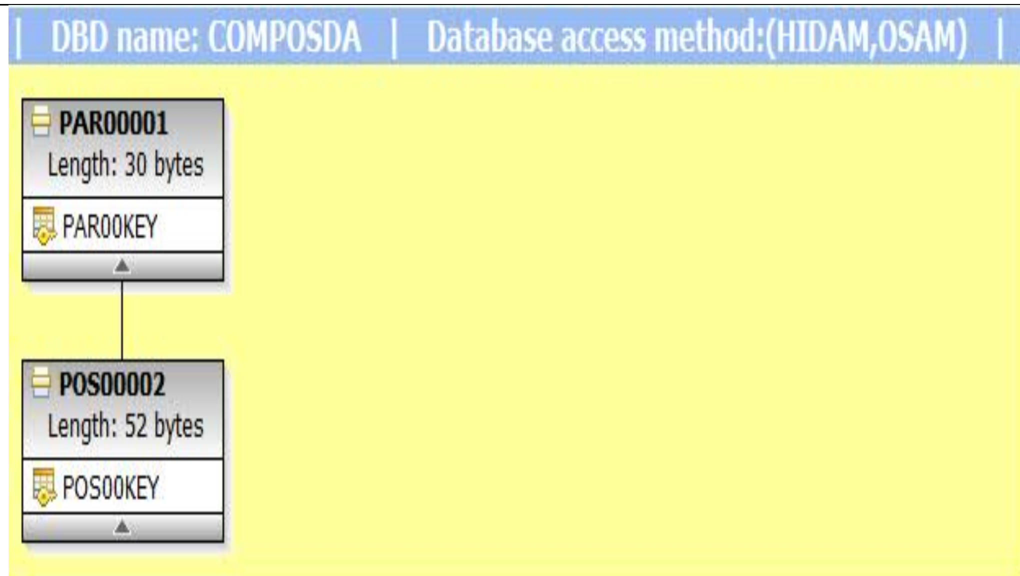
DEPSUMDA	DEPSUMDA is a department summary database containing information for the home department and work center. A DEPSUMDA record is composed of the following segments: • DEPT0001 (root segment key) - Contains information for home department • LINK0002 (direct dependent segment key) - Contains information for work center The following is a graphical view of the database layout: 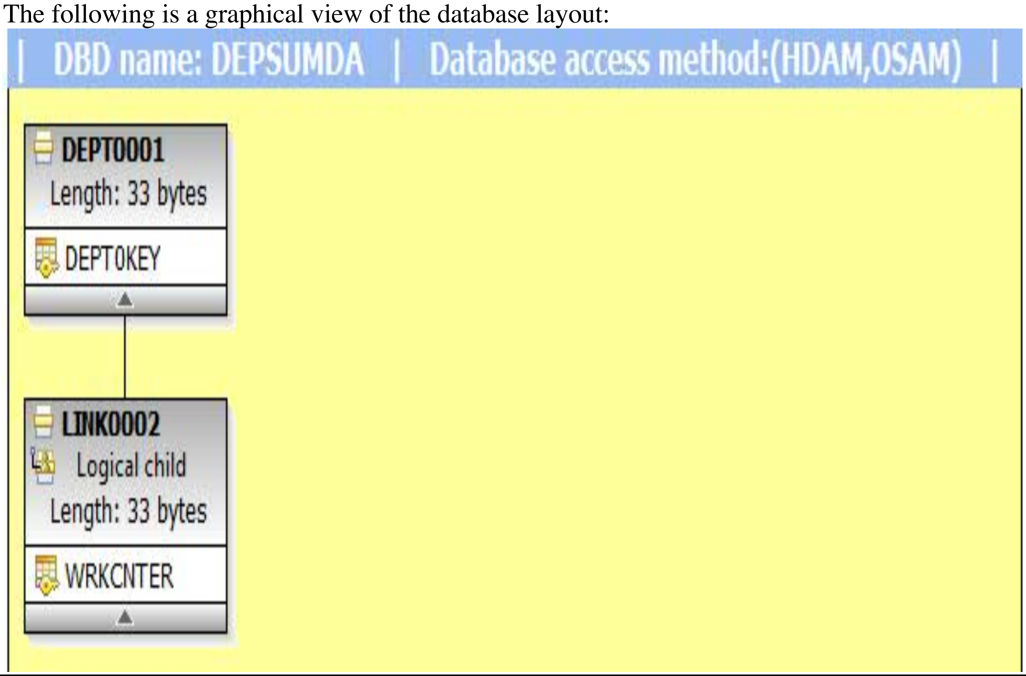
VENDORDA	VENDORDA is a vendor database containing information for all the vendors. VENDORDA record is composed of the following segments: • VN010VDR (root segment key) - Contains information for vendor key • VN020REO (logical parent segment key) - Contains information for reorder segment of vendor database via secondary index • VN030PAR (virtual logical child segment key) – Contains information for replacing a reorder segment

	The following is a graphical view of the database layout: 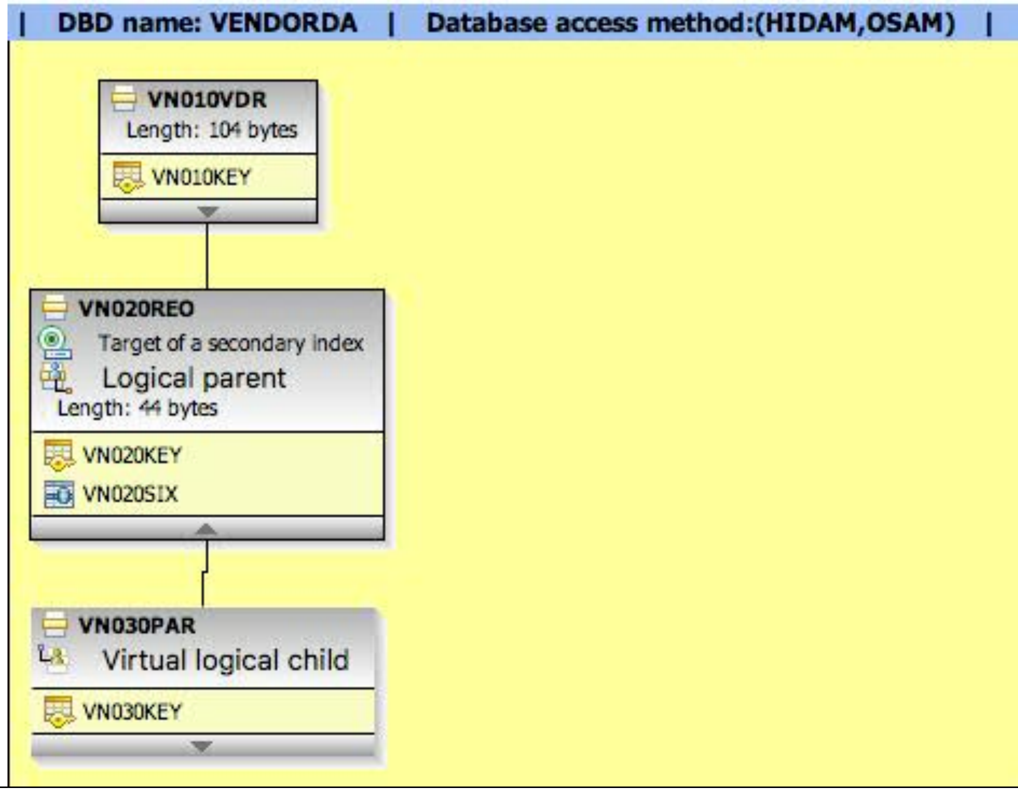

4.1.2 Application - Transaction Description

The Full Function with HALDB workload runs a mix of transactions using teller system, inventory, hotel, and warehouse type transactions doing read, replace, delete and insert calls as described in Table 2 below. The application is written in COBOL and compiled with COBOL V4R2.

Below is the complete list of transactions run by the Full Function with HALDB and CICS Full Function workloads as described in Section 4.5.

Table 2: Full Function with HALDB Workload Transaction Description

Transaction Name	Transaction Overview
HR1	Hotel Reservation Application - Hotel Reservation Transaction: • Processes reservation requests for specific accommodations at specific locations for specific dates • If accommodation type is unavailable, alternate arrangements are recommended • Hotel Reservation Database accessed

HR2	Hotel Reservation Application – Reservation Segment Create Transaction: • Creates a third level reservation segment for the already held reservation • Includes new information in the dependent segment including arrival date, departure date, room type, bed type, number of persons, rate and personal information • Hotel Reservation Database accessed
IT2	Inventory Tracking Application - Part Location and Inventory Inquiry Transaction: • Reports part location and available inventory information • Item Activity Database accessed
IT8	Inventory Tracking Application - Update in process Inventory Transaction: • Records the completion of labor operations on manufactured items and updates employee activity records • Item is tracked as it moves from work station to work station on manufacturing shop floor • Records two labor operations by a single employee on a single in-process item • Employee Activity Database, Item Activity, Product Item Master and Department Summary Database accessed
OE1	Order Entry Application - Customer Inquiry Transaction: • Inquires the customer database indexed by customer number • The customer number, if not directly known, is looked up in a customer directory • The output consists of customer name, ship to and bill to address, possible discounts and current credit • Conversational transaction. Data is stored in the SPA for further inquiry • Customer Database (CUSTOMRA) accessed
OE2	Order Entry Transaction - Customer Change: • Makes changes to the customer database • Mostly affects the root and second level address segment • Input consists of altered customer information from the customer inquiry step • Conversational transaction • Customer Database (CUSTOMRA) accessed
OE4	Order Entry Transaction - Parts Inquiry: • Reads a list of part numbers and quantities requested from inputting terminal • Makes inquiry into the status of each of these • Conversational transaction • Inventory database accessed
OE5	Order Entry Transaction - Parts Processing: • Makes actual updates to the inventory database when filling an order and producing an invoice • Updates the quantity on hand for each of the parts and creates a back-order entry if required • Also affects customer database, as order must be produced under appropriate customer • Final step for Order Entry conversational transactions • Output consists of terminal response and an invoice • Inventory and customer database (CUSTOMRA) accessed
PS2	Production Specification Application - Bill of Materials Transaction: • Bills of Materials for design level components • Product Item Master and Component Position Database accessed

PS3	Production Specification Application - Labor Operation/Incentive Standards for Parts Transaction: • Inquires the labor operations and incentive standards for a given part • Product Item Master Database accessed
SC2	Stock Control Application – Receiving and Processing Transaction: • Handles the receipt of ordered goods • Vendor and Inventory database accessed
SC6	Stock Control Application - Supplier Delete Transaction: • Delete obsolete suppliers from the Inventory database • Input is supplier's number and part number. Record may be deleted if there are no longer any orders pending • Database maintenance function • Inventory Database accessed
TS1	Teller System Application - Data Entry Transaction: • This is a HALDB transaction that does inserts and deletes • Customer Database (CUSTDA) accessed

4.1.3 Application - Workload Distribution

The Full Function with HALDB workload uses transactions with the following execution distributions as shown in Table 3 below.

Table 3: Full Function with HALDB Workload Transaction Distribution

Transaction	Distribution
HR1	8.89%
IT2	8.89%
IT8	8.89%
OE1	8.88%
PS2	24.44%
PS3	22.23%
SC6	8.89%
TS1	8.89%

Order Entry (OE) transactions are conversational, Program to Program switch transactions. OE1 calls OE2 which initiates OE4 which finally calls OE5.

4.2 Data Sharing Full Function (DSFF) with HALDB Workload using Shared Queues

IMS provides the ability for multiple IMS systems in a parallel sysplex environment to share a single set of message queues and the messages are held in structures within a Coupling Facility (CF).

The IMS Shared Message Queue (SMQ) uses the Common Queue Server (CQS), a generalized server that manages objects on CF structures, such as list or message queue structures. CQS receives, maintains, and distributes data objects from shared queues on behalf of its clients.

The DSFF with HALDB workload using Shared Queues contains the same Full Function databases and transactions described in Section 4.1; however, it uses IMS Shared Queues (vs. local queues) for its message processing.

The workload performs updates to a variety of OSAM and VSAM Full Function databases and uses the IRLM address space as the data sharing lock manager.

The databases were defined to DBRC using block level data sharing (SHARELVL (3)) to support data sharing for the Shared Message Queues (SMQ) workload.

4.2.1 Database Description

The DSFF with HALDB workload using Shared Queues consist of the same databases explained in Section 4.1.1

4.2.2 Application - Transaction Description

The Data Sharing Full Function workload with HALDB using Shared Queues consist of the same transactions described in Section 4.1.2

4.2.3 Application - Workload Distribution

The DSFF with HALDB workload using Shared Queues uses transactions with the same execution distributions as shown in Table 3 in Section 4.1.3.

4.3 Fast Path (FP) Banking Workload

Fast Path database and message processing provide high performance processing in a highly-available data environment for IMS applications. The FP banking workload represents application functions that might comprise the high-volume online workload in a credit card environment. Its capabilities include credit card validation, credit limit check, debit/credit account, and lost/stolen card reporting.

This section provides details about the database description, application transactions, and workload distributions.

4.3.1 Database Description

The FP banking workload consists of three Fast Path Data Entry Databases (DEDB) as described in Table 4 below.

Table 4: Fast Path Workload Database Description

Database Name	Database Overview
ACCTA	The Account Database (DEDB) is a customer account database containing general information. An ACCTA record is composed of the following segments: • ACCTINFO (root segment) - Contains information for account information. Major fields include account number, credit limit, and used credit • ACCTCUST (direct dependent segment) - Contains information for account customer. Major fields include customer name and address • ACCTACTV (direct dependent segment) - Contains information for account activity. One ACCTACTV segment instance exists for every debit/credit against an account. Major fields include the amount of the debit/credit, store code number, and date of transaction This database consists of 480 DEDB Area Data sets. The following is a graphical view of the database layout: 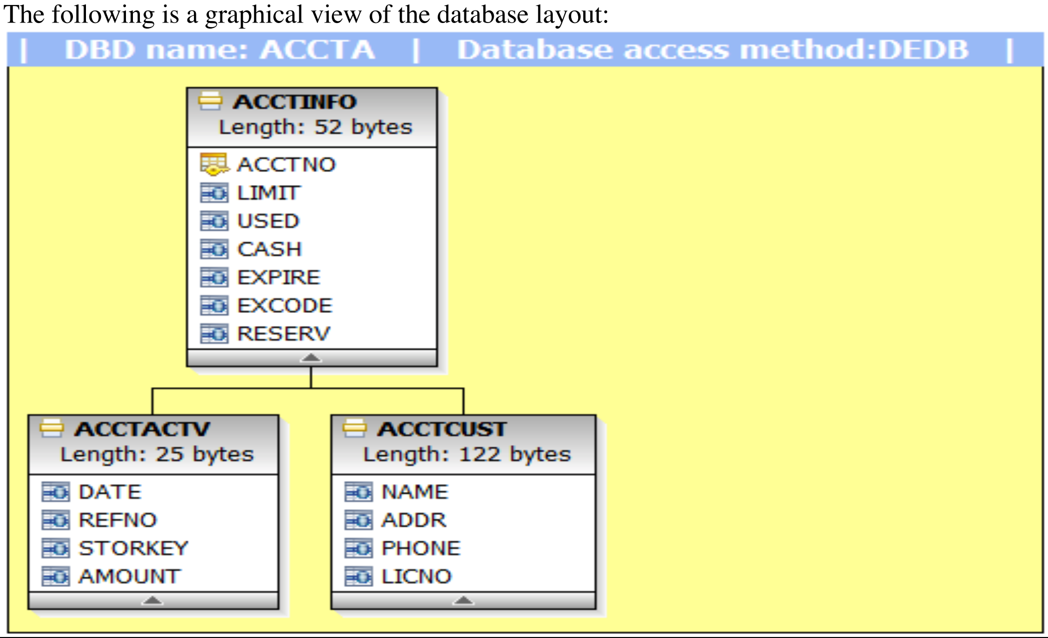

EXCEPTA	The Exception Card Database (DEDB) is an exception card database containing records for lost, stolen, or otherwise invalid cards. This database is periodically maintained (off-line) with EXCADD cards plus other updates. An EXCEPTA record consists of the following segment: • EXCECARD (root segment) - Contains information for exception card. Major fields include account number, customer name, and number of attempted uses This database consists of one DEDB Area Data set. The following is a graphical view of the database layout: 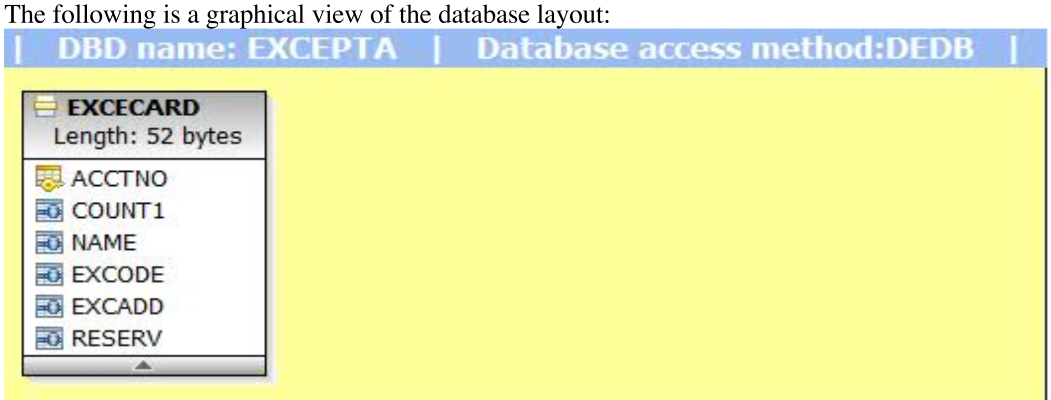
STOREDBA	Store Database (DEDB) is a database containing a record for each "establishment/retailer" subscribing to the credit card service. A STOREDB record consists of the following segment: • STORE (root segment) - Contains information for store and about the establishment. Major fields include store code number, store name, and number (by type) of each transaction invoked This database consists of 16 DEDB area data sets. The following is a graphical view of the database layout: 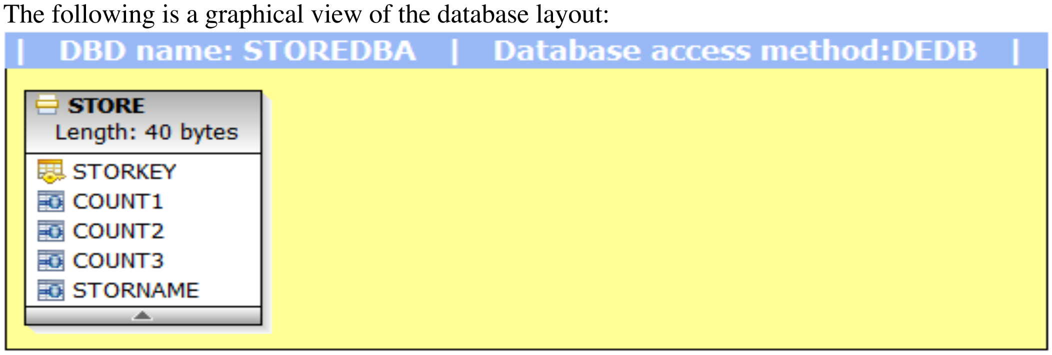

4.3.2 Application - Transaction Description

The Fast Path banking workload runs a mix of transactions using an on-line credit card environment executing read, replace, and insert calls.

The Fast Path banking workload includes four unique transactions with varying levels of processing. All transactions run in the IMS Fast Path (IFP) regions and are Fast Path only; that is, they are Fast Path Expedited Message Handler (EMH) messages and they access only Fast Path Data Entry Databases (DEDBs). They do not access any Full Function DL/I databases or issue calls to the IMS message queue. The transactions provide an online update capability with full integrity and recovery facilities for both the Fast Path databases and the Fast Path output response messages. Each transaction starts with an 80-byte input message and replies with an 83-byte output message. The application is written in COBOL and compiled with COBOL V5R2.

The IMS transactions CCCK, CLCK, DEBIT, and CREDIT would normally be issued by establishments subscribing to the credit card service. Counters in the STOREDBA database's store record keeps a tally of the number of DEBIT and CREDIT transactions issued. Table 5 provides an overview of the processing done by each of these transactions.

Table 5: IMS Fast Path workload transaction overview

Transaction Name	Transaction Overview
CCCK	Credit Card Authorization Check: • Search the Exception Card database to see if this card has been reported as lost or stolen • If necessary, also search the Added Exception Card database • If the card was found, update the "attempted uses" field in the corresponding record • Increment the transaction count in the Store Database • Return a message indicating the outcome of the search ("authorization ok" or "authorization denied") to the requesting client *Note:* Usually there is no I/O since most cards are in good standing and the Added Exception Card database will not need to be searched. The vast majority of transactions processed will stop at first bullet point with "authorization ok" returned in last bullet point.
CLCK	Credit Limit Check: • Fetch the Account Database root and check the transaction amount against the available credit • Increment the transaction count in the Store Database • Return a message authorizing or denying the purchase to the requesting client
DEBIT/CREDIT	Debit/Credit: - Fetch the Account Database root and update the balance - Insert a direct dependent under this root to journal the account activity - Increment the transaction count in the Store Database - Return "transaction complete" message to the requesting client *Note:* Each transaction requires DEDB read and an DEDB write.

4.3.3 Application - Workload Distribution

The transactions in the IMS FP banking workload have the following execution distributions as shown in Table 6 below.

Table 6: IMS Fast Path Workload Transaction Distribution

Transaction	Distribution
CCCK	33.33%
CLCK	33.33%
DEBIT	16.66%
CREDIT	16.66%

4.4 Batch Message Processing (BMP) Banking Workload

The BMP banking workload performs extensive sequential updates to Fast Path databases simulating end-of-day bank account batch reconciliation and reconsolidation to a DEDB account database with 85,142 segments.

4.5 Customer Information Control System (CICS) – IMS DBCTL Workload

The Customer Information Control System (CICS) transaction server workload was tested using the following IMS databases:

- CICS FF (retail/warehouse) workload as described in Section 4.1
- CICS FP (banking) workload as described in Section 4.3

The CICS FF workload uses transactions with the following execution distribution as shown in Table 7 below. This workload averages about 17 DL/I call per transaction.

Table 7: CICS Full Function Workload Transaction Distribution

Transaction	Distribution
HR1	10%
HR2	10%
IT2	10%
IT8	10%
PS2	25%
PS3	25%
SC2	5%

SC6	5%

The CICS FP workload uses transactions with the following execution distribution as shown in Table 8 below.

Table 8: CICS Fast Path Workload Transaction Distribution

Transaction	Distribution
CCCK	25%
CLCK	25%
DEBIT	25%
CREDIT	25%

4.6 z/OS Connect Enterprise Edition IMS Service Provider Workload

The IBM z/OS Connect Enterprise Edition (z/OS Connect EE) is an integrated solution that enables developers to merge business applications into today's growing mobile, cloud and hybrid cloud application ecosystems. z/OS Connect EE combines IBM and industry state of the art technologies to deliver a performant, intuitive solution for defining services and APIs to access your IMS assets using industry standard REST technology.

The following three components in z/OS Connect EE provide API solutions for IMS:

- The z/OS Connect EE Server
- The IMS Service Provider (IMS SP)
- z/OS Connect EE API toolkit

Each of these components integrates seamlessly to provide a fast and reliable experience for developers as they build applications for mobile and cloud use cases where speed to market is critical.

This workload focuses on stressing the IMS SP code and consists of message only transactions without any database activity.

4.6.1 Application - Transaction Description

The workload consists of 25 message-only transactions with transaction codes VARTX001 to VARTX025. All transactions run in the MPP regions. Figure 2 below shows the transaction attributes from the IMS command "/DIS TRAN ALL".

TRAN	CLS	ENQCT	QCT	LCT	PLCT	CP	NP	LP	SEGSZ	SEGNO	PARLM	RC
VARTX001	1	0	0	65535	65535	1	1	1	0	0	0	0
VARTX002	2	0	0	65535	65535	1	1	1	0	0	0	0
VARTX003	3	0	0	65535	65535	1	1	1	0	0	0	0
VARTX004	4	0	0	65535	65535	1	1	1	0	0	0	0
VARTX005	5	0	0	65535	65535	1	1	1	0	0	0	0
VARTX006	6	0	0	65535	65535	1	1	1	0	0	0	0
VARTX007	7	0	0	65535	65535	1	1	1	0	0	0	0
VARTX008	8	0	0	65535	65535	1	1	1	0	0	0	0
VARTX009	9	0	0	65535	65535	1	1	1	0	0	0	0
VARTX010	10	0	0	65535	65535	1	1	1	0	0	0	0
VARTX011	11	0	0	65535	65535	1	1	1	0	0	0	0
VARTX012	12	0	0	65535	65535	1	1	1	0	0	0	0
VARTX013	13	0	0	65535	65535	1	1	1	0	0	0	0
VARTX014	14	0	0	65535	65535	1	1	1	0	0	0	0
VARTX015	15	0	0	65535	65535	1	1	1	0	0	0	0
VARTX016	16	0	0	65535	65535	1	1	1	0	0	0	0
VARTX017	17	0	0	65535	65535	1	1	1	0	0	0	0
VARTX018	18	0	0	65535	65535	1	1	1	0	0	0	0
VARTX019	19	0	0	65535	65535	1	1	1	0	0	0	0
VARTX020	20	0	0	65535	65535	1	1	1	0	0	0	0
VARTX021	21	0	0	65535	65535	1	1	1	0	0	0	0
VARTX022	22	0	0	65535	65535	1	1	1	0	0	0	0
VARTX023	23	0	0	65535	65535	1	1	1	0	0	0	0
VARTX024	24	0	0	65535	65535	1	1	1	0	0	0	0
VARTX025	25	0	0	65535	65535	1	1	1	0	0	0	0

Figure 2: IMS Service Provider Workload Transaction Attributes

4.6.2 Application – Message Structures

The application was designed to receive and return messages with variable lengths based on the incoming message's field values. The evaluations in this paper focused on an I/O message size of 1KB only. The I/O messages contain simple character fields.

The IMS SP provides data transformation between Java Script Object Notation (JSON) and the binary format that the IMS transaction expects. This is done using message metadata imported from COBOL copybooks or PL/I includes. At runtime, the data and message metadata structures are dynamically inspected and converted.

4.6.3 Application – Workload

A stand-alone Java application was used to drive the workload by sending concurrent REST API requests across several threads to simulate multiple clients. Each request calls a REST API using HTTPS GET and randomly invokes one of the 25 services deployed on the z/OS Connect EE server.

4.7 IMS TM-Db2 with IBM Relational Warehouse Workload (IRWW) - Full PL/I

The IMS-Db2 IRWW relational database is based on a retail type of environment. The IRWW database resides on a Db2® for z/OS® database. The IMS transaction executes in 68 dependent regions (66 IFPs and 2 MPPs) accessing the Db2 warehouse database through the External Subsystem Attach Facility (ESAF).

4.7.1 Database Description

The workload consists of seven transactions with each transaction consisting of one to many SQL statements and containing inventory stock warehouses and sales districts. This workload has nine tables with five of the nine tables being in the partitioned tablespace as described below:

Table 9: IMS TM-Db2 Database Description

Table Name	Table Overview
CUSTOMER	CUSTOMER is 10 parts with index/part + one additional index and has 21 columns per row
DISTRICT	DISTRICT is partition-by-growth with one unique index and has 11 columns per row
HISTORY	HISTORY is partition-by-growth without index and has eight columns per row
ITEM	ITEM is partition-by-growth with one unique index and has four columns per row
NEWORDERS	NEWORDERS is 10 parts with index/part and has three columns per row
ORDERLINE	ORDERLINE is 10 parts with index/part and has 10 columns per row
ORDERS	ORDERS is 10 parts with index/part and has eight columns per row
STOCK	STOCK is 10 parts with index/part and has 17 columns per row
WAREHOUSE	WAREHOUSE is partition-by-growth with one unique index and has nine columns per row

4.7.2 Application - Transaction Description

The IRWW workload runs a mix of transactions using retail warehouse type transactions executing select, fetch, update, insert, and delete calls.

Table 10: TM-Db2 Application Transaction Description

Transaction Name	Transaction Description
Delivery	Performs various SELECTs, UPDATEs, and DELETEs in support of the delivery of a group of orders
New Order	Performs various SELECTs, FETCHs, UPDATEs, and INSERTs in support of the new customer
Order Status	Performs various SELECTs, FETCHs in support of providing the status of an order
Payment	Performs various SELECTs, FETCHs, UPDATEs, and INSERTs in support of the received customer payment
Price Change	Performs an UPDATE in support of changing the price of an item
Price Quote	Performs a SELECT in support of providing the price of a set of items
Stock Inquiry	Performs a JOIN and various SELECTs in support of providing the current stock level of an item

4.7.3 Application - Workload Distribution

The transactions in the IRWW have the following execution distribution as shown in Table 11 below.

Table 11: TM-Db2 Workload Distribution

Transaction	Distribution
Delivery	2%
New Order	22%
Order Status	24%
Payment	22%
Price Change	1%
Price Quote	25%
Stock Inquiry	4%

4.8 Java Message Processing (JMP) Workload

JMP regions are similar to MPP regions except they schedule Java programs rather than native applications such as COBOL or PL/I. In the PSB source associated with the Java program, the option LANG=JAVA must be specified. A JMP application is started when there is a message in the queue for the JMP application and IMS schedules the message to be processed. JMP applications, like MPP applications, are executed through transaction codes submitted by users at terminals and from other applications. Each transaction code represents a transaction that the JMP application processes.

This section provides details about the database description, application transactions, and workload distributions.

4.8.1 Database Description

The JMP workload consists of six Full Function OSAM databases described below in Table 12.

Table 12: JMP Workload Database Description

Database Name	Database Overview
ACCOUNT	The Account Database is a customer account database containing general information. It contains one root segment: • ACCOUNT (root segment) - Contains information for all customer accounts. Major fields include Account ID, Account type, Balance, and Last transaction ID The following is a graphical view of the database layout: 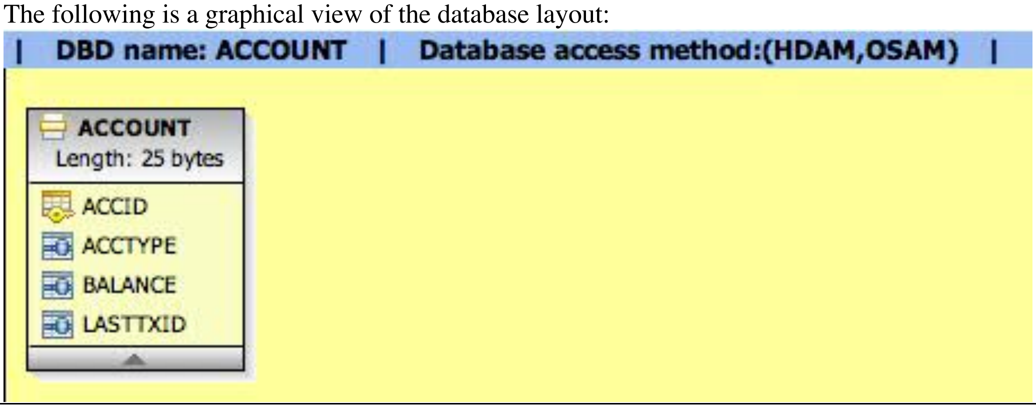
ACCTYPE	The Account Type Database records all types of accounts in the account database. It contains one root segment: • ACCTYPE (root segment) - Contains information for the type of customer account. Major fields include Code and Description of the account The following is a graphical view of the database layout:

	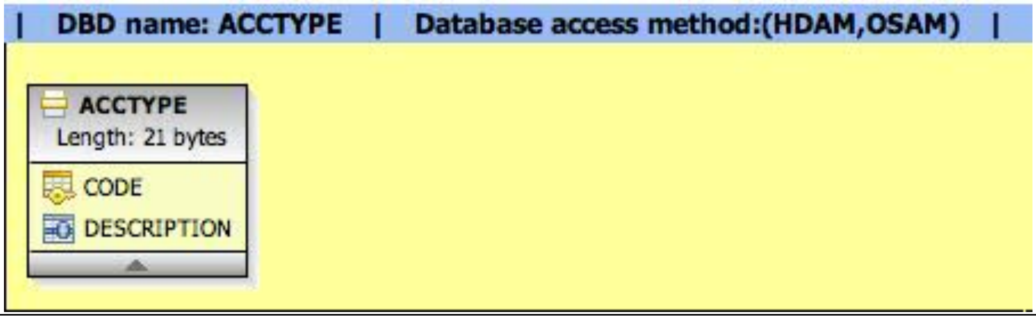
CUSTOMER	Customer Database is a database containing a record for all customers. It contains one root segment: • CUSTOMER (root segment) - Contains information for all customers. Major fields include customer ID, Last name, First name, Address, City, State, Zip code, Phone, Status, Password, Customer type, and Last login The following is a graphical view of the database layout: 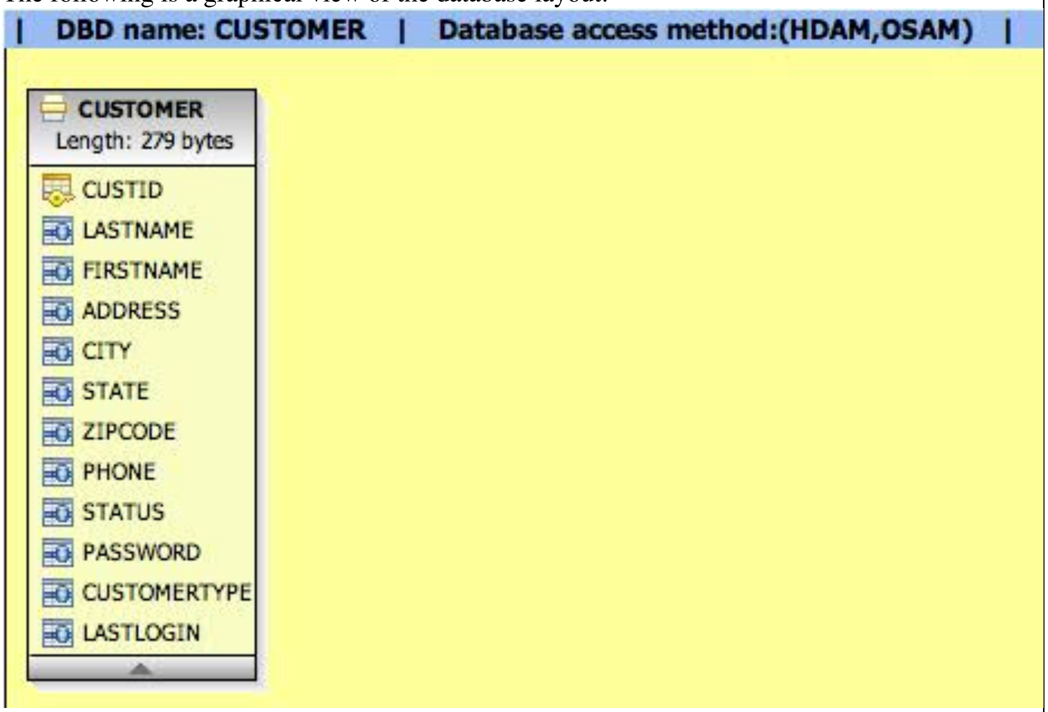
CUSTTYPE	The Customer Type Database records all types of customers in the customer database. It contains one root segment: • CUSTTYPE (root segment) - Contains information for all customer types. Major fields include Customer code and Description The following is a graphical view of the database layout:

<table>
<tr><td></td><td>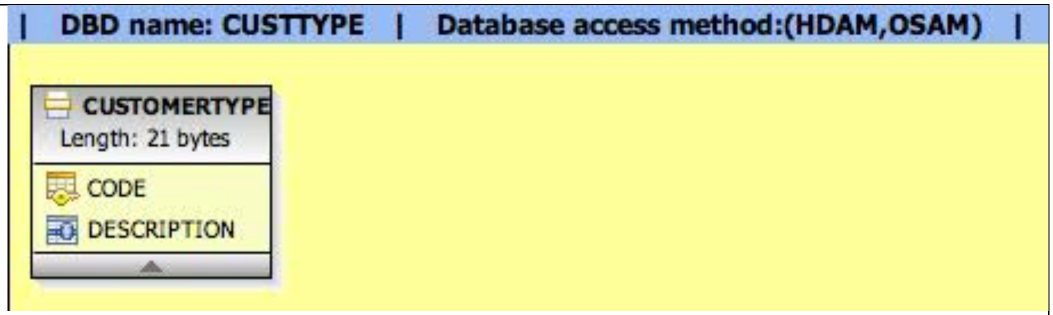
</td></tr>
<tr><td>CUSTACCS</td><td>The Customer Account Segment Database records all types of customers in the customer account database. It contains one root segment:

- CUSTACCS (root segment) - Contains information for all customer accounts. Major fields include Customer ID, Customer account ID, and Account number

The following is a graphical view of the database layout:
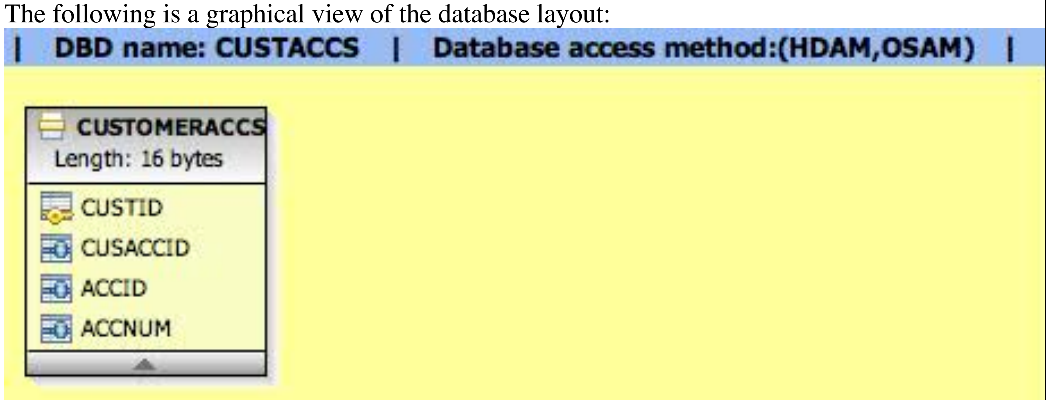
</td></tr>
<tr><td>HISTORY</td><td>The History Database records all types of customers transactions, time, Transaction types, Amount and Reference transaction IDs in History database. It contains one root segment:

- HISTORY (root segment) - Contains information for all transactions. Major fields include Transaction ID, Time, Transaction type, amount and Account ID

The following is a graphical view of the database layout:
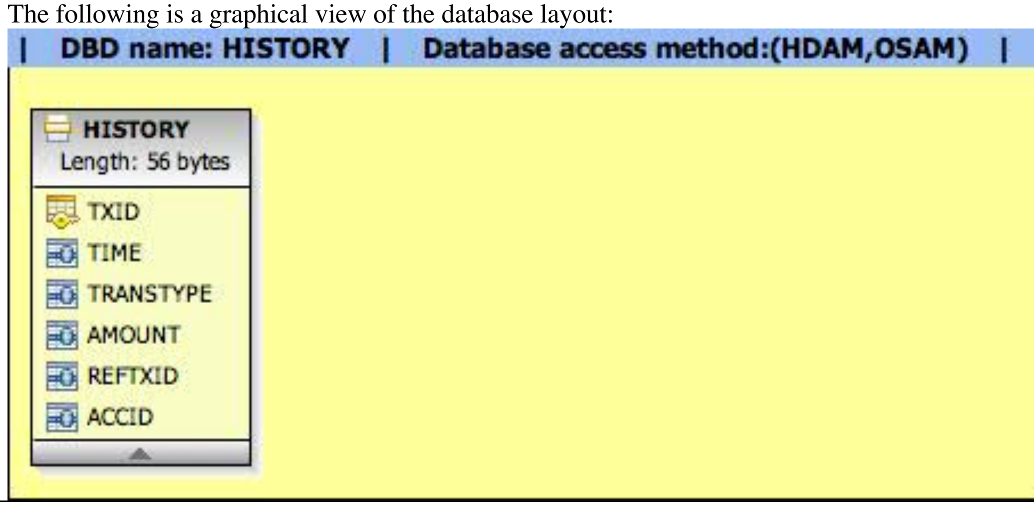
</td></tr>
</table>

4.8.2 Application - Transaction Description

The JMP workload runs a mix of transactions in an online credit card environment executing read, replace, and insert calls.

The IMS transactions FBTRAN, FBLOGIN, FBLOGOUT, FBACSUM, FBGCUDAT and FBSCUDAT are issued by establishments doing credit, account inquiries, logging in and logging out a customer. Table 13 provides an overview of the processing done by each of these transactions.

Table 13: JMP Workload Transaction Description

Transaction Name	Transaction Overview
FBTRAN	Friendly Bank Credit or Debit Check: • Fetch the Account Database root and check the transaction amount against the available credit • Insert a record and increment the transaction count in the History Database • Update Account Summary and balance
FBLOGIN	Friendly Bank Account Login: • Perform login for customer ID into the account, verify password match and check if already logged on • Update last login information
FBLOGOUT	Friendly Bank Account Logout: • Perform logout for customer ID from the account • Update customer status
FBACSUM	Friendly Bank Account Summary: • Fetch the customer account information summary like balance, account type and account ID
FBGCUDAT	Friendly Bank Get Customer Data: • Fetch the customer information like first name, last name and address

4.8.3 Application - Workload Distribution

The transactions in the JMP workload have the following execution distribution as shown in Table 14 below.

Table 14: JMP Workload Database Distribution

Transaction	Distribution
FBTRAN	33.33%

FBLOGIN	16.67%
FBLOGOUT	16.67%
FBACSUM	16.66%
FBGCUDAT	16.67%

4.9 Open Database (ODBM) Workload

The IMS Open Database solution provides distributed access to IMS database resources driving industry standard and open technology into IMS. The distributed nature is two-fold:

- At an IMSPlex level, allows cross LPAR access to any IMS database in the IMSPlex
- At a pure distributed level, allows open platforms access to IMS resources directly using industry standard interfaces

This is accomplished by three main components:

- Client-side libraries implementing the industry standard interfaces and protocols
- IMS Connect to process the distributed requests
- Open Database Manager (ODBM) address space to process database access requests

4.9.1 Database Description

IMS Open Database consists of six Full Function OSAM databases described in Section 4.8.1 and presented in Table 12: JMP Workload Database Description.

4.9.2 Application - Workload Distribution

The Friendly Bank workload runs a mix of transactions in an online credit card environment executing read, replace, and insert calls as described in Section 4.8.2.

5 Testing Methodology

The IMS performance evaluation cycle, as shown in Figure 3 below, is a repetitive process where a test environment is created and/or customized for a specific measurement, performance tests are run, and data is analyzed. Depending on the analyzed results, changes are made in the environment, if applicable, and the whole test process is repeated. All testing is done using an isolated and stable environment to produce consistent and repeatable performance measurement results.

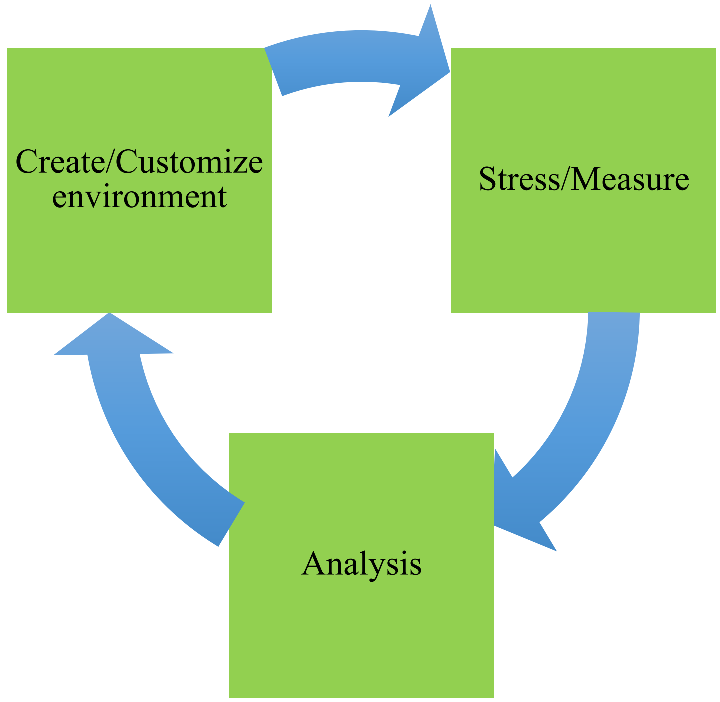

Figure 3: IMS Performance Evaluation Cycle

The performance test environment for each measurement evaluation described in this paper involves a specific system configuration as to the number of LPARs, GPs and Specialty Engines (e.g. zIIPs) that were active.

The system configuration is documented in the introduction of each measurement evaluation and kept constant when comparing IMS versions.

A series of scaling tests were run to compare IMS versions at various CPU percent busy values ranging from 10% to 80%. For each workload, the 80% CPU busy measurements were used for IMS version to version comparisons.

5.1 Pre-Measurement Procedure

The following generic procedure was used to setup the measurement environment prior to the start of the measurement procedure. As mentioned above, the purpose of the procedure is to ensure repeatability from run to run.

1. Restore the IMS database datasets

2. Allocate IMS system datasets including Online Log Datasets (OLDS), Write Ahead Datasets (WADS) and Recovery Control Datasets (RECON)
3. Initialize IMS RECON for database recoverability
4. Start Structured Call Interface (SCI), Operations Manager (OM) and Internal Resource Lock Manager (IRLM)
5. Initialize and 'Cold Start' IMS
6. Start the dependent regions (e.g. IFP, MPP, JMP, BMP) as required
7. Start other workload-specific address spaces (e.g., IMS Connect, CICS, Db2, z/OS Connect EE server) as required

5.2 Measurement Procedure

After the databases are restored and the environment initialized, the following generic steps were taken to measure the performance of a specific workload. The measurement procedure captures key performance data about the overall z/OS system as well as data specific to IMS. The following procedure was used to capture the measurement:

1. Start all TPNS address spaces or Java-based workload drivers
2. Initialize and start all TPNS networks to begin to drive transaction requests through IMS Connect over TCP/IP socket connections or directly to IMS using SNA
3. Ramp up the workload by adjusting the number of users for Java drivers or decreasing the 'think time' (think time specifies the time a TPNS simulated "client" waits between each transaction invocation) for TPNS networks until a target of between 80% and 85% of CPU utilization is reached
4. Start Resource Management Facility (RMF) Monitor I and III. Two-minute RMF intervals were used
5. Issue the /SWITCH OLDS command to force an IMS OLDS switch
6. Issue the /CHECKPOINT STATISTICS command to request that IMS performance records be created and written to the IMS log
7. Wait for the two-minute measurement interval

8. Issue the /CHECKPOINT STATISTICS command to request that IMS performance records be created and written to the IMS log

9. Issue the /SWITCH OLDS command to force an IMS OLDS switch

10. Quiesce the TPNS networks

11. Stop the TPNS address spaces or Java workload drivers

5.3 Post Measurement Procedure

After completing the measurement, the following generic steps were used to capture the performance data:

1. Run the SMF logstream dump utility (IFASMFDL) to allow for post processing

2. Run RMF post processing against the dumped SMF data to produce various RMF reports detailing z/OS system activity

3. Run IMS Performance Analyzer (IMSPA) against all of the IMS OLDS processed between step 5-9 of Section 5.2 to produce various reports detailing IMS activity.

5.4 Measurement Metrics and Analysis

The results of any performance evaluation run include many data points from different measurement sources. RMF provides information about z/OS and hardware resources such as CPU utilization, memory consumption, and I/O rates. IMSPA provides IMS internal statistics such as transaction rate, logging rate, and latch contention rates. All of this data is captured and saved for future research and analysis. However, there are a few basic metrics that apply to almost all measurements as shown in Table 15 below:

Table 15: Performance Metrics for IBM Z Processor Comparison

Metric	Description
CPU % Busy	The average percent busy across all general CPs on an LPAR during the measurement interval.
zIIP % Busy	The average percent busy across all zIIPs, if applicable, on an LPAR during the measurement interval.

Total LPAR % Busy	The average percent busy across all general CPs and zIIPs on an LPAR during the measurement interval.
ETR (Tran/Sec)	External Transaction Rate is the observed average transaction rate in transactions per second (TPS) over the measurement interval captured from the RMF report.
ITR	Internal Transaction Rate is a projection of the observed transaction rate (ETR) to what the transaction rate would be if the processors were running at 100% CPU busy, assuming linear scaling. ITR was calculated by dividing ETR by CPU % Busy
Total IMS Response Time (ms)	The total IMS transaction response time from the IMSPA report.
Total General CPU µs/Tran	The total CPU microseconds spent per transaction was calculated using the following formula: (Number of CPs * CPU % Busy * 1000000)/ETR
Total zIIP µs/Tran	The total CPU microseconds spent per transaction on zIIP was calculated using the following formula: (Number of zIIPs * zIIP % Busy * 1000000)/ETR (Number of zIIPs * zIIP % Busy * AVG TD * 1000000)/ETR
Total IMS CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for all IMS address spaces captured from the RMF report. IMS address spaces include Control region, DL/I region, DBRC region, SCI, OM and IRLM. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
IMS Connect CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for the IMS Connect address space captured from the RMF report. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
Dependent Regions CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for IMS dependent region address space captured from the RMF report. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
z/OS Connect EE CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for the z/OS Connect EE address space captured from the RMF report. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
z/OS Connect EE or JMP zIIP Service Time (s) per transaction	The amount of zIIP time, if applicable, consumed by the processors on an LPAR captured from the RMF report. It is reported on a per transaction basis by dividing the total zIIP service time consumed during an interval by the number of transactions processed in that interval. (IIP APPL% * 10000)/ETR

	IIP APPL% (or zIIP service time) is an estimated value trying to project what zIIP time would have been without SMT active. Its accuracy is dependent on sufficient samples of data when 1 and 2 threads are active. For low utilized zIIPs and/or for low AVG TD this value may not be accurate.
CQS CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for the CQS address space, if applicable, captured from the RMF report. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
IXLOGR CPU Service Time (s) per transaction	The amount of CPU time consumed by the processors for the z/OS Logger (IXGLOGR) address space, if applicable, captured from the RMF report. It is reported on a per transaction basis by dividing the total CPU service time consumed during an interval by the number of transactions processed in that interval.
Total Number of Offloads	The number of z/OS Logger offloads, if applicable, captured from the RMF report.
Avg. CSA Below 16MB Key 7	The average usage of Key 7 common storage below 16MB from the RMF report.
Avg. CSA Above 16MB Key 7	The average usage of Key 7 common storage above 16MB from the RMF report.
Avg. LSQA Private	The average usage of private LSQA storage below 16MB from the RMF report.
Avg. LSQA EPrivate	The average usage of extended private LSQA storage above 16MB from the RMF report.
Avg. USER Private	The average usage of private USER storage below 16MB from the RMF report.
Avg. User EPrivate	The average usage of extended private USER storage above 16MB from the RMF report.

The IMS Shared Queues metrics were calculated using both LPARs in a data sharing environment shown in Table 16 below:

Table 16: Performance Metrics for Shared Message Queues Comparison

Combined Metric	Description
Avg. CPU % Busy	The average percent busy across two LPARs during the measurement interval.
Total ETR (Tran/Sec)	The sum of the average transaction rates in transactions per second (TPS) during the measurement interval captured from the RMF report on both LPARs.
Combined ITR	The combined ITR during measurement interval calculated by dividing Total ETR with Avg. CPU % Busy.
Average Total IMS Response Time (ms)	The average of total IMS transaction response time from the IMSPA reports on both LPARs.

Combined General CPU µs/Tran	The total CPU microseconds spent per transaction calculated using the following formula: (Sum of total number of CPs on both LPARs * Avg. CPU % Busy * 1000000)/Total ETR
IMS CPU Service Time (s) per transaction	The sum of total CPU time consumed by the processors for all IMS address spaces captured from the RMF report on both LPARs. IMS address spaces include Control region, DL/I region, DBRC region, SCI, OM and IRLM. It is reported on a per transaction basis by dividing the sum of total CPU service time consumed during an interval by the sum of total number of transactions processed in that interval on both LPARs.
IMS Connect CPU Service Time (s) per transaction	The sum of total CPU time consumed by the processors for the IMS Connect address spaces captured from the RMF report on both LPARs. It is reported on a per transaction basis by dividing the sum of total CPU service time consumed during an interval by the sum of total number of transactions processed in that interval on both LPARs.
Dependent Regions CPU Service Time (s) per transaction	The sum of total CPU time consumed by the processors for IMS dependent region address space captured from the RMF report on both LPARs. It is reported on a per transaction basis by dividing the sum of total CPU service time consumed during an interval by the sum of total number of transactions processed in that interval on both LPARs.
CQS CPU Service Time (s) per transaction	The sum of total CPU time consumed by the processors for the CQS address space captured from the RMF report on both LPARs. It is reported on a per transaction basis by dividing the sum of total CPU service time consumed during an interval by the sum of total number of transactions processed in that interval on both LPARs.
IXGLOGR CPU Service Time (s) per transaction	The sum of total CPU time consumed by the processors for the z/OS Logger (IXGLOGR) address space captured from the RMF report on both LPARs. It is reported on a per transaction basis by dividing the sum of total CPU service time consumed during an interval by the sum of total number of transactions processed in that interval on both LPARs.
Total Number of Offloads	The total number of z/OS Logger offloads on both LPARs.

ITR is one way of comparing the processor efficiency of two IBM Z processors using the same software environment. ITR normalizes the observed transaction rate to the engine capacity of the machine. It answers the question "assuming the transaction rate scales linearly with CPU usage, what is the maximum transaction rate possible on this particular hardware configuration (i.e. when CPU percent busy is 100%)".

Note that ITR does not take into considerations other possible bottlenecks besides CPU (for instance, I/O, latch contention, lock contention) that could further limit the theoretical maximum transaction rate.

6 IMS 15 Base Performance

6.1 Introduction

The base performance workloads were processed on the newest IBM Z platform, z14, for the three latest IMS releases (IMS 13, 14, and 15). The performance data of approximately 80% CPU utilization for all three IMSs will be presented in the results tables, however, only IMS 14 will have the delta comparisons against IMS 15.

The base performance workloads are listed below, and the details of these workloads were described above in Section 4.

- Full Function with HALDB
- Data Sharing Full Function with HALDB and Shared Message Queues
- Fast Path Banking
- Batch Message Processing Banking
- CICS IMS DBCTL
- IMS TM-Db2 IRWW
- z/OS Connect EE IMS Service Provider
- Java Message Processing
- Open Database Management

IMS 15 was executed in a Managed ACB environment for these base performance workloads as the future direction of IMS. The performance comparison between an IMS 15 running in a Managed ACB environment and a non-Managed ACB environment is detailed in Section 7.2: IMS 15 Managed Application Control Blocks (ACB).

The z14 hardware configuration (e.g. the number of processors, I/O channels, DASD, LPAR memory) was kept constant for any given workload measurement for each IMS release level.

The software configuration was also kept constant for each IMS release level and each performance workload. As with any new release, new functionality can increase the execution path length, resulting in some normally acceptable increased processing cost. The ITR values for each of the specific types of IMS workloads, except BMP, were used to evaluate the general performance and CPU efficiency. Comparisons between IMS releases running in the same environmental configuration are considered equivalent in this document where the delta percentage difference is within +/- 1% in ITR.

Table 17 below shows the hardware and software testing environment.

Table 17: Performance Evaluation Environment

Hardware and Software Environment	
Processor	IBM z14 Model 3964-7E7 (M05)
DASD	IBM System Storage DS8886
IBM z/OS Operating System	z/OS Version 2 Release 2
IBM Enterprise COBOL for z/OS	Version 4 Release 2
IMS	IMS 13, IMS 14, IMS 15
z/OS Resource Measurement Facility (RMF)	Version 2 Release 2
IMS Performance Analyzer for z/OS (IMSPA)	Version 4 Release 4
z/OS Resource Access Control Facility (RACF)	Version 2 Release 2
Java	Java 8 Service Release 5
Db2	Db2 Version 12

In addition to the performance comparisons at approximately 80% CPU utilization, a series of scaling runs were performed ranging from 10% to 80% CPU utilization for each IMS release. Charts comparing the ITR vs Transaction Rate, IMS Response Time vs Transaction Rate, and CPU % Busy vs Transaction Rate are presented.

6.2 Fast Path Performance Evaluation

Fast Path is capable of performing transaction and database processing at high rates. If your system requirements include a high transaction volume with relatively non-complicated database structure and message processing (e.g. no logical relationships; no message switching), Fast Path can be advantageous over Full Function processing. The Fast Path evaluation measured the following scenario:

- TCP/IP input through IMS Connect with Fast Path 64-bit buffer manager enabled

The Fast Path 64-bit buffer manager was introduced in IMS V11 and it autonomically controls the number and size of Fast Path buffer pools for Data Entry Databases (DEDBs). This autonomic control eliminates the need for users to manually define Fast Path buffer pools during system definition.

This evaluation used the Fast Path workload including credit card type transactions such as CCCK (Credit Card Check), CLCK (Credit Card Limit Check), CREDIT, DEBIT, and LOST as described in Section 4.3.

The objective of the Fast Path base evaluation was to compare the ITR between IMS 14 and IMS 15 Managed ACB in the same hardware environment with TCP/IP message protocols.

6.2.1 System Configuration

The Fast Path base evaluation using TCP/IP was executed on a z14 configured in a two LPAR environment as shown in Figure 4.

- LPAR 1 hosts IMS, 375 IFP regions, and 5 IMS Connects with seven General Purpose Engines
- LPAR 2 hosts TPNS driving a total of 8,000 IMS Connect clients via TCP/IP with 10 General Purpose Engines

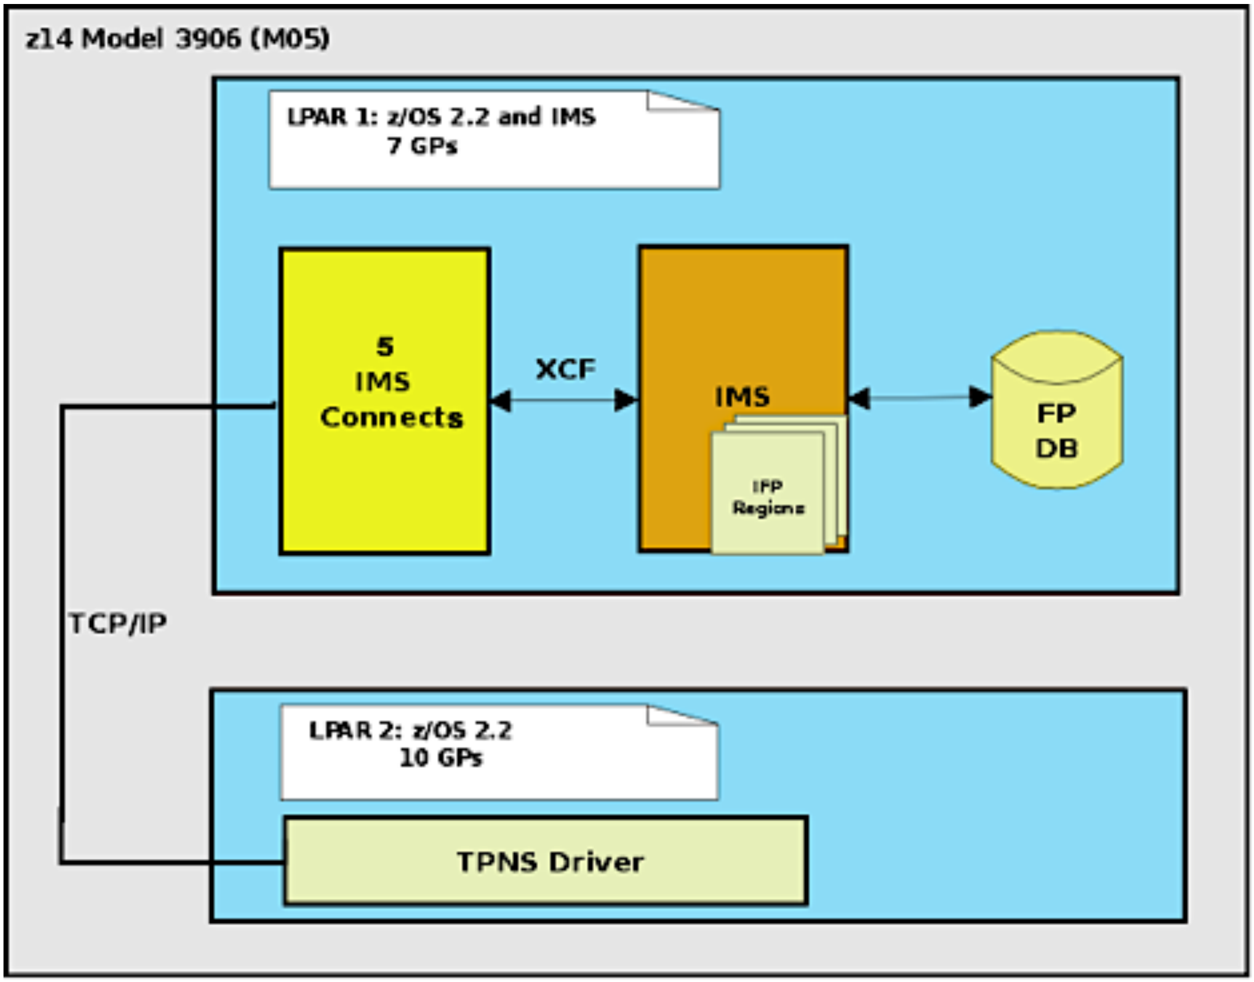

Figure 4: Fast Path Base Evaluation using TCP/IP Configuration

6.2.2 Evaluation Results

The ITR for IMS 15 Managed ACB showed a 1.6% increase over IMS 14 for the FP banking workload. Table 18 shows the results of the Fast Path evaluation for IMS 13, IMS 14 and IMS 15.

Table 18: Fast Path Evaluation Results

Fast Path Banking Evaluation					
	IMS 13	**IMS 14**	**IMS 15**	**Delta**	**Delta %**

CPU % Busy	85.59%	83.06%	81.50%	-1.56%	-1.88%
ETR (Tran/Sec)	98972.20	99031.82	98733.87	-297.95	-0.30%
ITR	115635.24	119229.26	121145.85	1916.59	1.61%
Total IMS Response Time (ms)	0.608	0.653	0.642	-0.011	-1.68%
Total General CPU µs/Tran	60.535	58.710	57.782	-0.929	-1.58%
IMS CPU Service Time/Tran (µs)	11.252	10.678	11.196	0.518	4.85%
ICON CPU Service Time/Tran (µs)	19.302	18.748	17.464	-1.284	-6.85%
IFP CPU Service Time/Tran (µs)	18.308	17.957	18.024	0.067	0.37%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	388	368	368	0	0.00%
Avg. CSA Above 16M Key 7 (M)	715	714	715	1	0.14%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	1275	1276	1280	4	0.31%
Avg. LSQA EPrivate (M)	14.7	14.6	15.1	0.5	3.42%
Avg. USER Private (K)	2652	2276	2272	-4	-0.18%
Avg. USER EPrivate (M)	105.0	93.5	96.2	2.7	2.89%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	352	352	352	0	0.00%
Avg. LSQA EPrivate (M)	8.48	8.46	8.89	0.43	5.13%
Avg. USER Private (K)	568	624	628	4	0.64%

Avg. USER EPrivate (M)	264	264	265	1	0.38%
	Private Storage IMS Connect Regions				
Avg. LSQA Private (K)	360.8	360.0	360.8	0.8	0.22%
Avg. LSQA EPrivate (M)	15.24	15.26	15.44	0.18	1.18%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	204	204	206	2	0.98%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 5 shows the ITR vs Transaction Rate, Figure 6 shows the IMS Response Time vs Transaction Rate, and Figure 7 shows the CPU % Busy vs Transaction Rate.

The Fast Path banking workload on IMS 15 Managed ACB showed slight improvement in ITR and Total IMS response time compared to IMS 14.

- ITR improved by 1.61%
- Total IMS response time improved by 1.68%

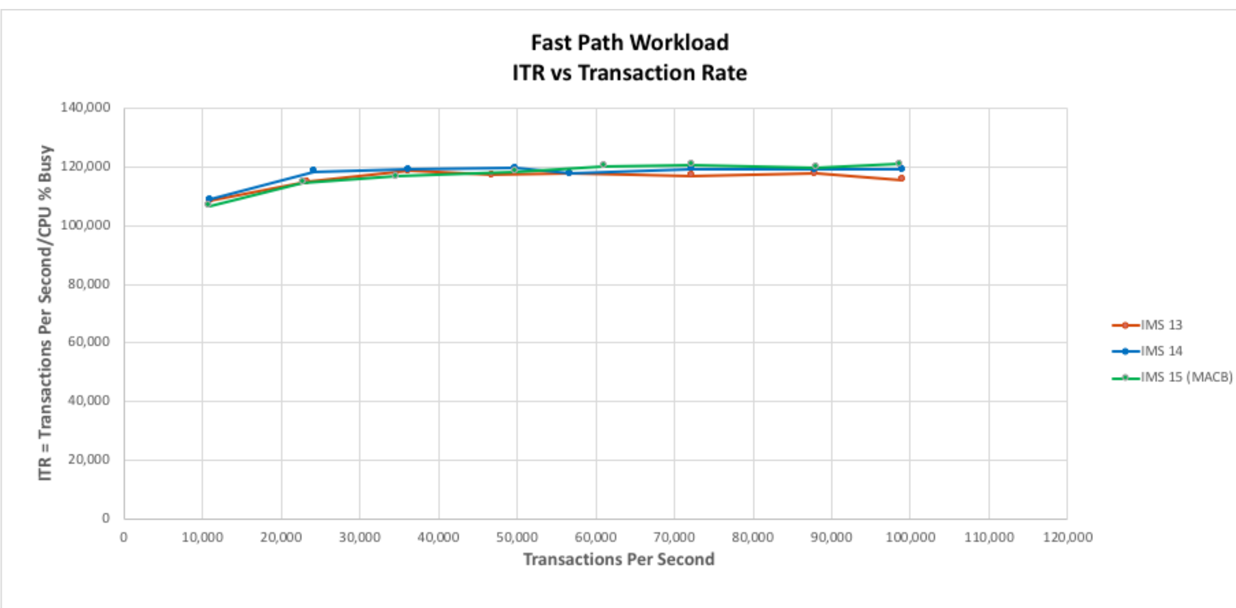

Figure 5: Fast Path Banking ITR vs Transaction Rate Comparison Results

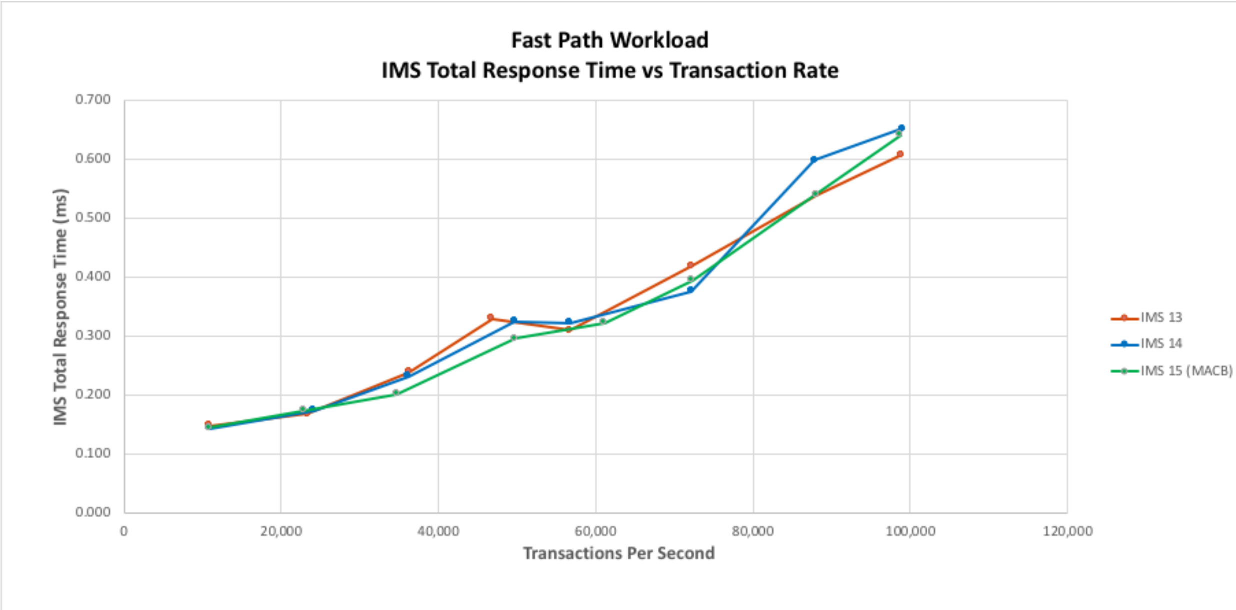

Figure 6: Fast Path Banking IMS Total Response Time vs Transaction Rate Comparison Results

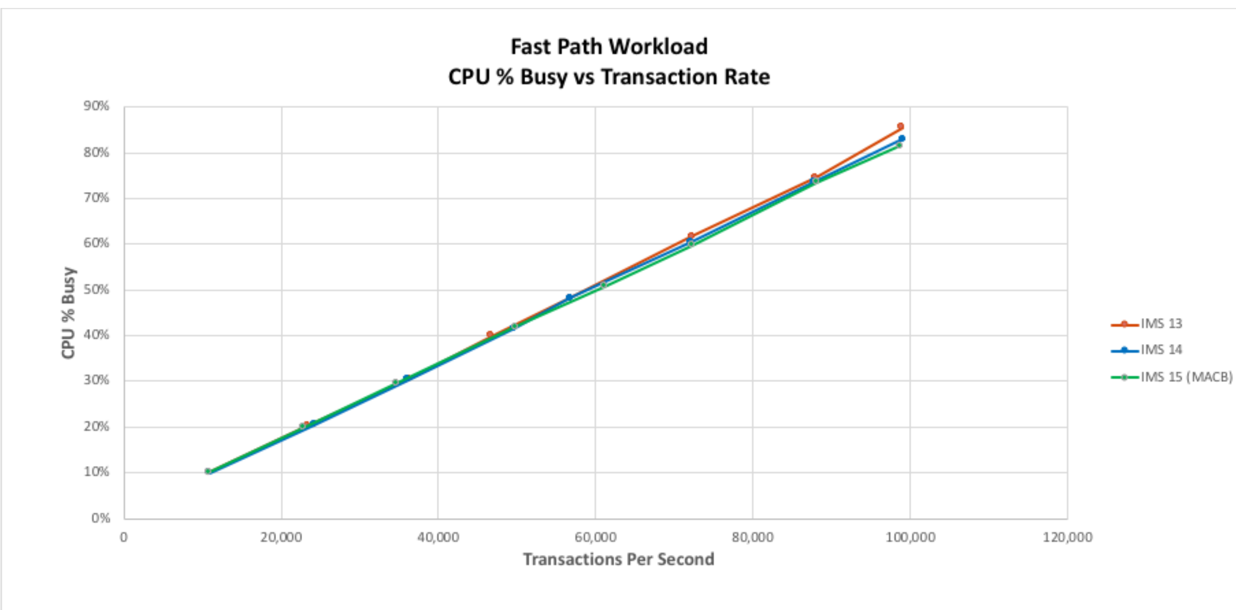

Figure 7: Fast Path Banking CPU % Busy vs Transaction Rate Comparison Results

6.3 Full Function with High Availability Large Database (HALDB) Performance Evaluation

Full Function databases support the full set of IMS database functions and can be used in a wide variety of IMS applications.

This evaluation uses the Full Function with HALDB workload as described in Section 4.1. The workload consists a mix of OSAM and VSAM with HDAM, HIDAM, PHDAM, and PHIDAM databases using inventory, hotel, and warehouse-type transactions that perform read, replace, and insert database calls.

The objective of the Full Function with HALDB evaluation was to compare the ITR between IMS 13, IMS 14, and IMS 15 in the same hardware configuration with TCP/IP message protocols.

6.3.1 System Configuration

The Full Function with HALDB workload was executed on a z14 configured in a two LPAR environment as shown in Figure 8:

- LPAR 1 hosts IMS, 128 MPP regions, and IMS Connect with three General Purpose Engines
- LPAR 2 hosts TPNS driving a total of 4,000 IMS Connect clients via TCP/IP with 10 General Purpose Engines

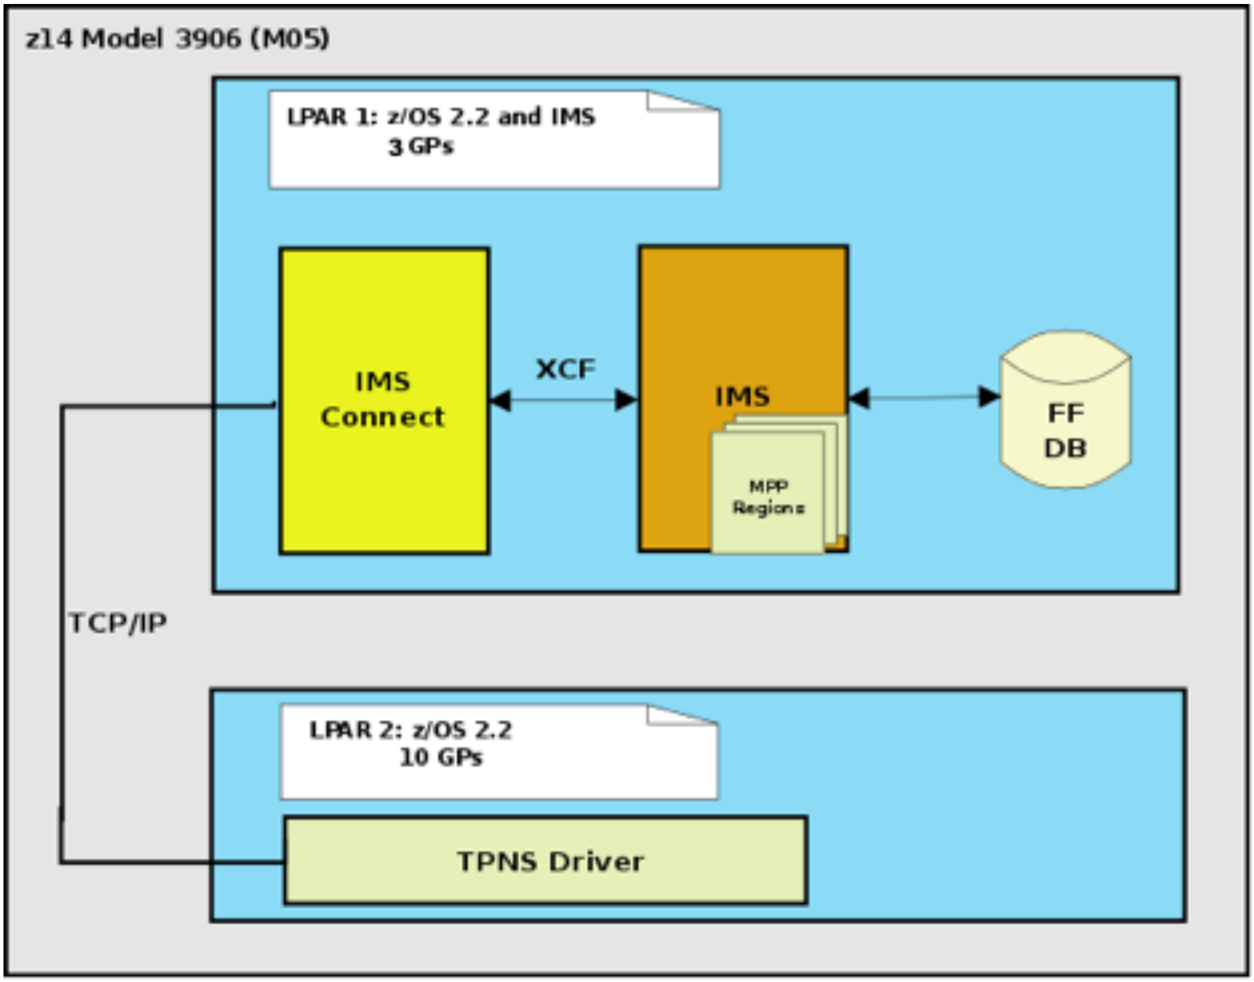

Figure 8: Full Function with HALDB Base using TCP/IP Configuration

6.3.2 Evaluation Results

The IMS 15 Managed ACB showed the ITR to be within the nominal variability range over IMS 14 for the Full Function with HALDB workload. Table 19 shows the comparisons between IMS 13, IMS 14, and IMS 15.

Table 19: Full Function with HALDB Evaluation Results

Full Function with HALDB Evaluation					
	IMS 13	**IMS 14**	**IMS 15**	**Delta**	**Delta %**
CPU % Busy	79.38%	79.20%	81.81%	2.61%	3.30%
ETR (Tran/Sec)	4834.19	4770.88	4951.39	180.51	3.78%

ITR	6089.93	6023.84	6052.30	28.46	0.47%
Total IMS Response Time (ms)	18.384	18.555	18.387	-0.168	-0.91%
Total General CPU µs/Tran	492.616	498.021	495.679	-2.342	-0.47%
IMS CPU Service Time/Tran (µs)	35.632	35.945	36.185	0.240	0.67%
ICON CPU Service Time/Tran (µs)	29.190	29.546	26.429	-3.117	-10.55%
MPP CPU Service Time/Tran (µs)	354.514	358.936	360.642	1.706	0.48%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	280	260	260	0	0.00%
Avg. CSA Above 16M Key 7 (M)	19.4	19.3	19.2	-0.1	-0.52%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	604	612	617	5	0.82%
Avg. LSQA EPrivate (M)	13.1	13.3	14.3	1.0	7.52%
Avg. USER Private (K)	2644	2264	2288	24	1.06%
Avg. USER EPrivate (M)	70.6	73.1	72.0	-1.1	-1.50%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	1912	1944	1952	8	0.41%
Avg. LSQA EPrivate (M)	11.9	11.9	12.6	0.7	5.88%
Avg. USER Private (K)	652	704	704	0	0.00%
Avg. USER EPrivate (M)	266	266	266	0	0.00%
	Private Storage IMS Connect Region				

Avg. LSQA Private (K)	364	360	360	0	0.00%
Avg. LSQA EPrivate (M)	14.2	15.0	14.2	-0.8	-5.33%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	247	248	249	1	0.40%

A series of scaling tests were run to compare IMS 13, IMS 14 and IMS 15 at various CPU percent busy values. Figure 9 shows the ITR vs Transaction Rate, Figure 10 shows the IMS Response Time vs Transaction Rate, and Figure 11 shows the CPU % Busy vs Transaction Rate.

The Full Function with HALDB workload on IMS 15 Managed ACB showed no significant change in both ITR and in Total IMS response time compared to IMS 14:

- ITR improved by 0.47%
- Total IMS response time degraded by 0.91%

Figure 9: Full Function with HALDB ITR vs Transaction Rate Comparison Results

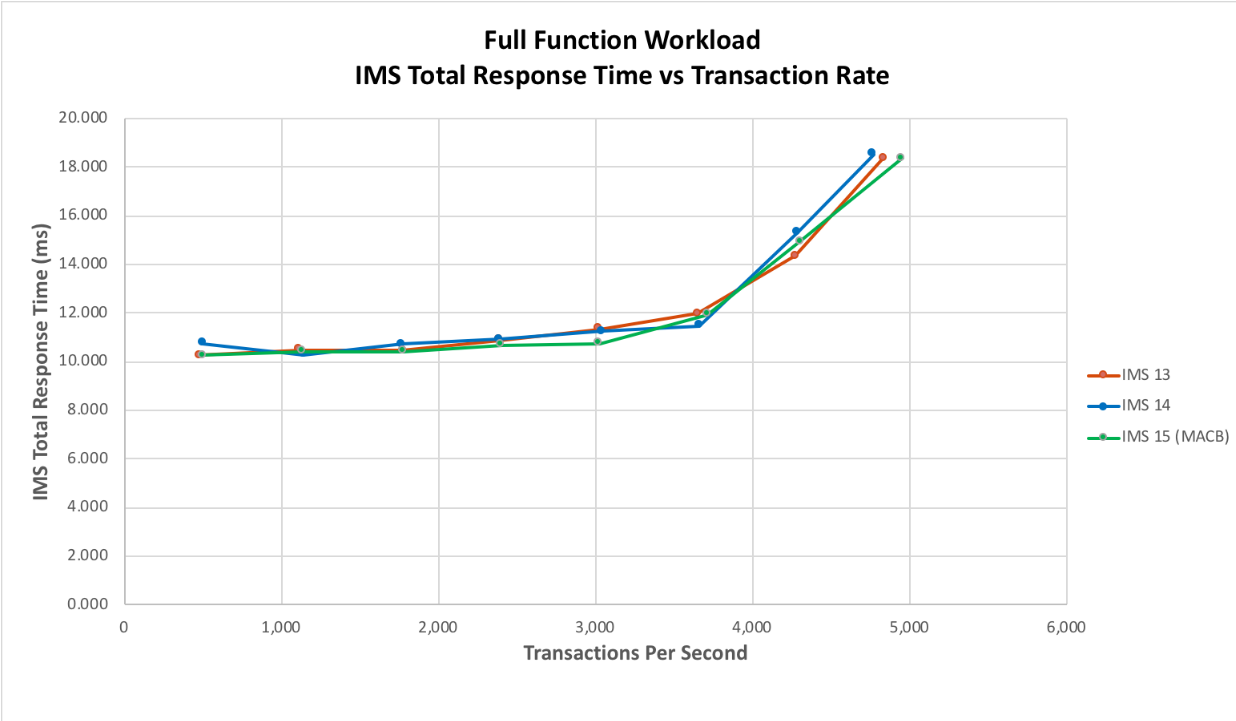

Figure 10: Full Function with HALDB IMS Total Response Time vs Transaction Rate Comparison Results

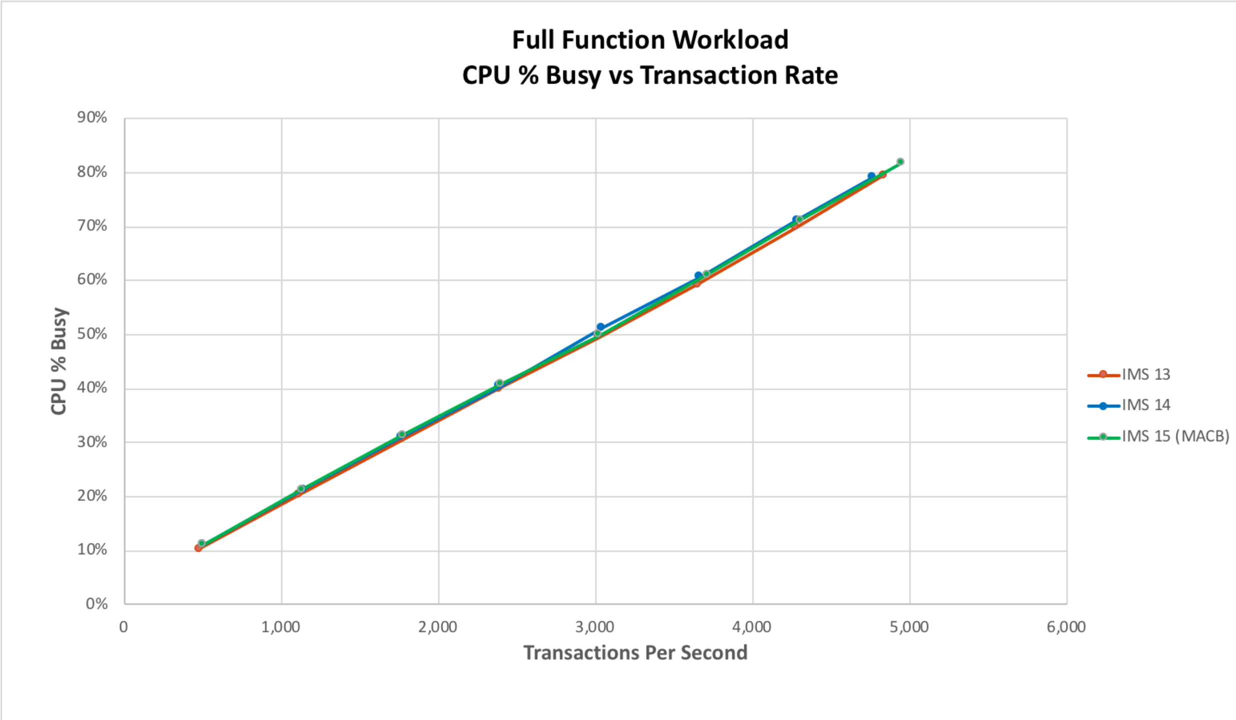

Figure 11: Full Function with HALDB Total CPU % Busy vs Transaction Rate Comparison Results

6.4 Data Sharing Full Function (DSFF) with HALDB using Shared Message Queues (SMQ) Performance Evaluation

The IMS DSFF with HALDB using SMQ uses the Common Queue Server (CQS) which manages a shared queue of messages residing on a Coupling Facility list structure in a two-way shared-queue data sharing environment. This evaluation uses the Full Function workload described in Section 4.2. Note that with IMS 12 and subsequent releases, CQS request response processing runs under enclave Service Request Blocks (SRBs) and are eligible to run on System z Integrated Information Processor (zIIP) specialty engines.

The objective of the DSFF with HALDB using SMQ base evaluation was to compare the ITR between IMS 13, IMS 14 and IMS 15 Managed ACB in the same hardware configuration with TCP/IP message protocols.

6.4.1 System Configuration

The DSFF with HALDB using SMQ workload using TCP/IP was executed on a z14 configured in a three LPAR environment as shown in Figure 12:

- LPAR 1 hosts IMS, 128 MPP regions, and IMS Connect with three General Purpose Engines
- LPAR 2 hosts IMS, 128 MPP regions, and IMS Connect with three General Purpose Engines
- LPAR 3 hosts TPNS driving 4,000 IMS Connect clients into each IMS via TCP/IP with 10 General Purpose Engines

Each IMS was started with CQS and utilized a shared message queue structure and a z/OS log stream structure for the CQS logging. All structures resided in an Internal Coupling Facility (ICF).

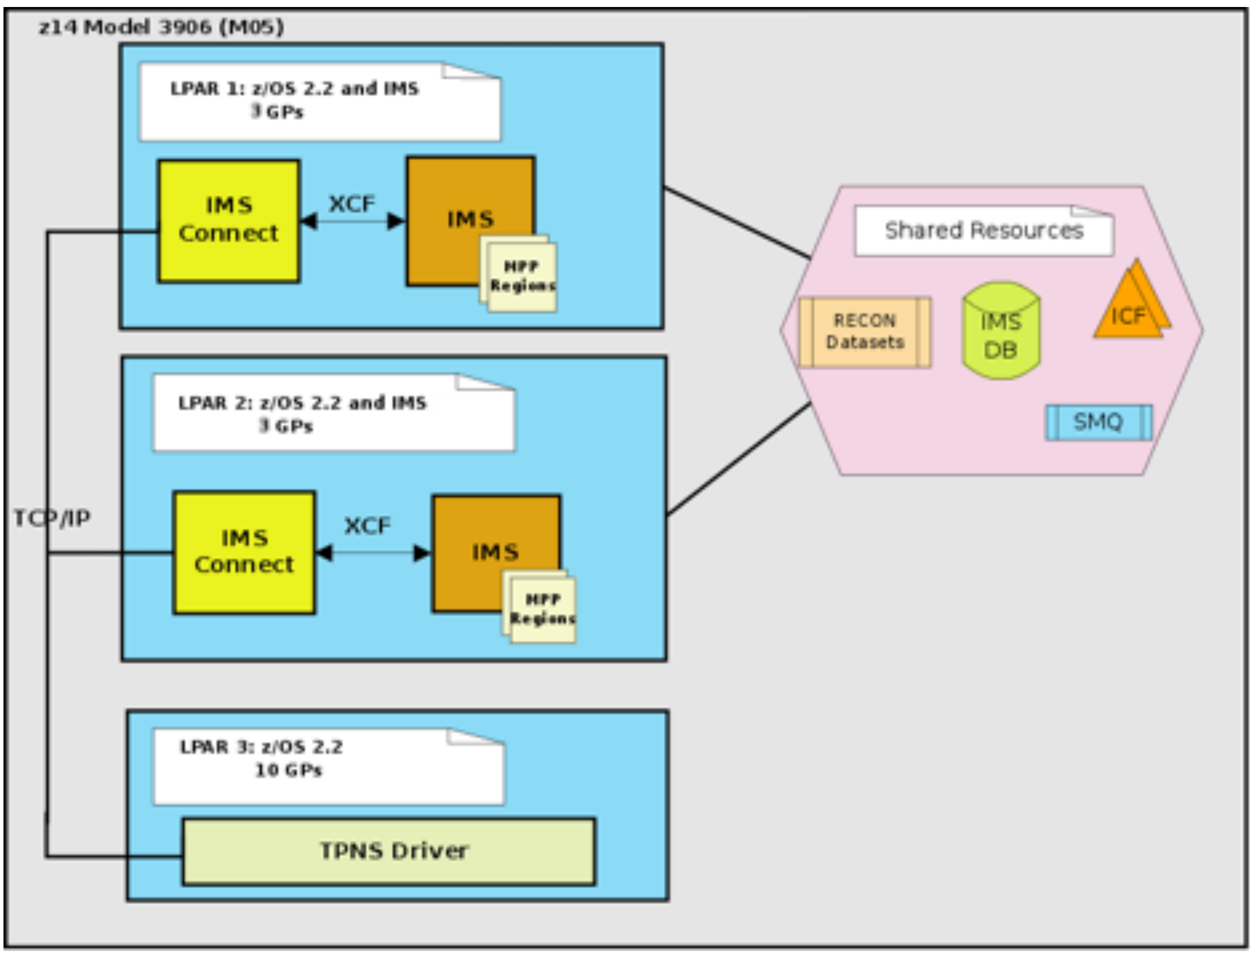

Figure 12: DSFF with HALDB using SMQ Environment Configuration

6.4.2 Evaluation Results

IMS 15 Managed ACB showed a slight increase in ITR compared to IMS 14 for the DSFF with HALDB using SMQ workload. Table 20 shows the comparisons for the combined IMSs in the Shared Queue environment calculated using formulas from Table 16 in Section 5.4. Table 21 and Table 22 show the comparisons for IMS1 and IMS2 for each IMS version.

Table 20: DSFF with HALDB using SMQ Evaluation Combined Results

DSFF with HALDB using SMQ Evaluation Combined					
	IMS 13	IMS 14	IMS 15	Delta	Delta %
Average CPU % Busy	77.09%	79.27%	78.28%	-0.99%	-1.25%
Sum of ETRs (Tran/Sec)	5487.60	5548.70	5529.30	-19.40	-0.35%
Combined ITR	7118.43	7000.19	7063.94	63.75	0.91%
Average Total IMS Response Time (ms)	15.918	17.689	16.067	-1.622	-9.17%
Combined General CPU µs/Tran	842.882	857.120	849.384	-7.735	-0.90%
Sum Total IMS CPU Service Time/tran (µs)	97.343	110.838	112.101	1.263	1.14%
Sum ICON CPU Service Time/Tran (µs)	61.996	60.711	54.611	-6.100	-10.05%
Sum MPP CPU Service Time/Tran (µs)	1149.898	1136.905	1127.102	-9.803	-0.86%
Sum CQS CPU Service Time/Tran (µs)	145.655	154.043	151.869	-2.174	-1.41%
Sum IXGLOGR CPU Service Time/Tran (µs)	13.703	17.517	13.042	-4.475	-25.55%
Sum Total Number of offloads	1.00	1.00	1.00	0.00	0.00%

Table 21: DSFF with HALDB using SMQ Evaluation Results

DSFF with HALDB using SMQ Evaluation			
	IMS 13	IMS 14	IMS 15

	IMS1	IMS2	IMS1	IMS2	IMS1	IMS2
CPU % Busy	79.58%	74.60%	80.85%	77.68%	79.40%	77.15%
ETR (Tran/Sec)	2743.26	2744.34	2762.85	2785.85	2761.39	2767.91
ITR	3447.24	3680.27	3417.22	3587.76	3477.90	3628.28
Total IMS Response Time (ms)	16.474	15.362	17.781	17.596	15.997	16.137
Total General CPU µs/Tran	870.263	815.158	877.907	836.177	862.590	826.837
Total IMS CPU Service Time/tran (µs)	51.110	46.233	57.444	53.394	57.946	54.155
ICON CPU Service Time/Tran (µs)	32.832	29.164	31.753	28.958	28.750	25.861
MPP CPU Service Time/Tran (µs)	591.725	558.173	582.154	554.751	574.418	552.684
CQS CPU Service Time/Tran (µs)	74.242	71.413	78.377	75.666	77.235	74.634
IXGLOGR CPU Service Time/Tran (µs)	9.235	4.468	9.034	8.483	7.433	5.609
Number of offloads	1	0	0	1	0	1

Table 22: DSFF with HALDB using SMQ Evaluation Storage Usage

	IMS1			IMS2		
	IMS 13	IMS 14	IMS 15	IMS 13	IMS 14	IMS 15
	Common Storage Below and Above 16MB for Avg. Key 7					
Avg. CSA Below 16M Key 7 (K)	280	260	260`	280	260	260
Avg. CSA Above 16M Key 7 (M)	23.8	24.2	24.0	23.7	24.0	23.9
	Private Storage IMS Control Region					
Avg. LSQA Private (K)	612	620	624	612	624	628

Avg. LSQA EPrivate (M)	13.4	13.9	14.6	13.4	13.9	14.9
Avg. USER Private (K)	2656	2272	2284	2666	2280	2280
Avg. USER EPrivate (M)	71.0	75.0	72.7	69.4	73.8	71.9
	Private Storage IMS DL/I Region					
Avg. LSQA Private (K)	1900	1952	1948	1900	1940	1944
Avg. LSQA EPrivate (M)	12.5	12.6	13.3	12.5	12.6	13.4
Avg. USER Private (K)	652	704	704	652	704	704
Avg. USER EPrivate (M)	267	267	267	267	267	267
	Private Storage IMS Connect Region					
Avg. LSQA Private (K)	360	360	360	360	360	360
Avg. LSQA EPrivate (M)	14.3	14.3	14.2	14.1	14.3	14.1
Avg. USER Private (K)	52	56	56	52	56	56
Avg. USER EPrivate (M)	247	248	250	248	248	250
	Private Storage CQS Region					
Avg. LSQA Private (K)	344	348	344	348	348	352
Avg. LSQA EPrivate (M)	13.1	13.4	13.3	13.2	13.1	13.3
Avg. USER Private (K)	16	16	16	16	16	16
Avg. USER EPrivate (M)	11.0	12.0	12.0	11.0	11.9	12.0

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 13 shows the ITR vs Transaction Rate comparison and Figure 14 shows the IMS Response Time vs Transaction Rate comparison. Figure 15 shows the CPU % Busy vs Transaction Rate.

The DSFF with HALDB using SMQ workload on IMS 15 Managed ACB showed improvements in both ITR and Total IMS response time compared to IMS 14.

- ITR improved by 0.91%
- Total IMS response time improved by 9.17%

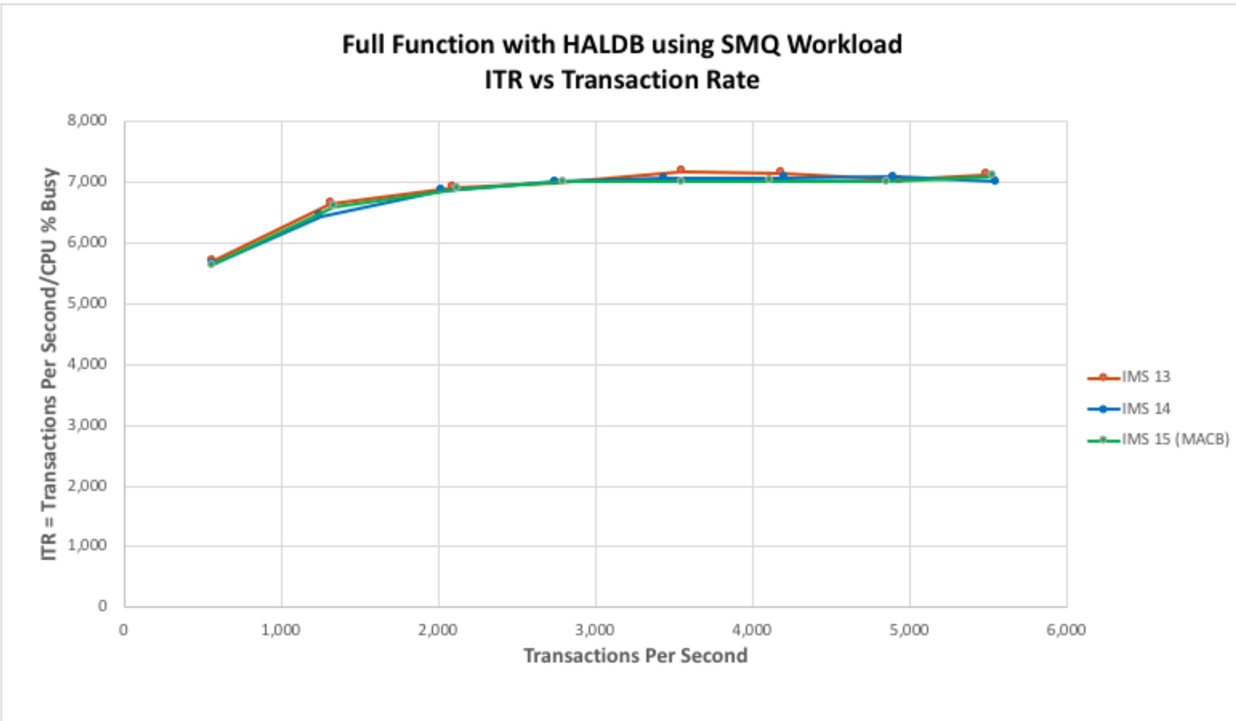

Figure 13: DSFF with HALDB using SMQ ITR vs Transaction Rate Comparison

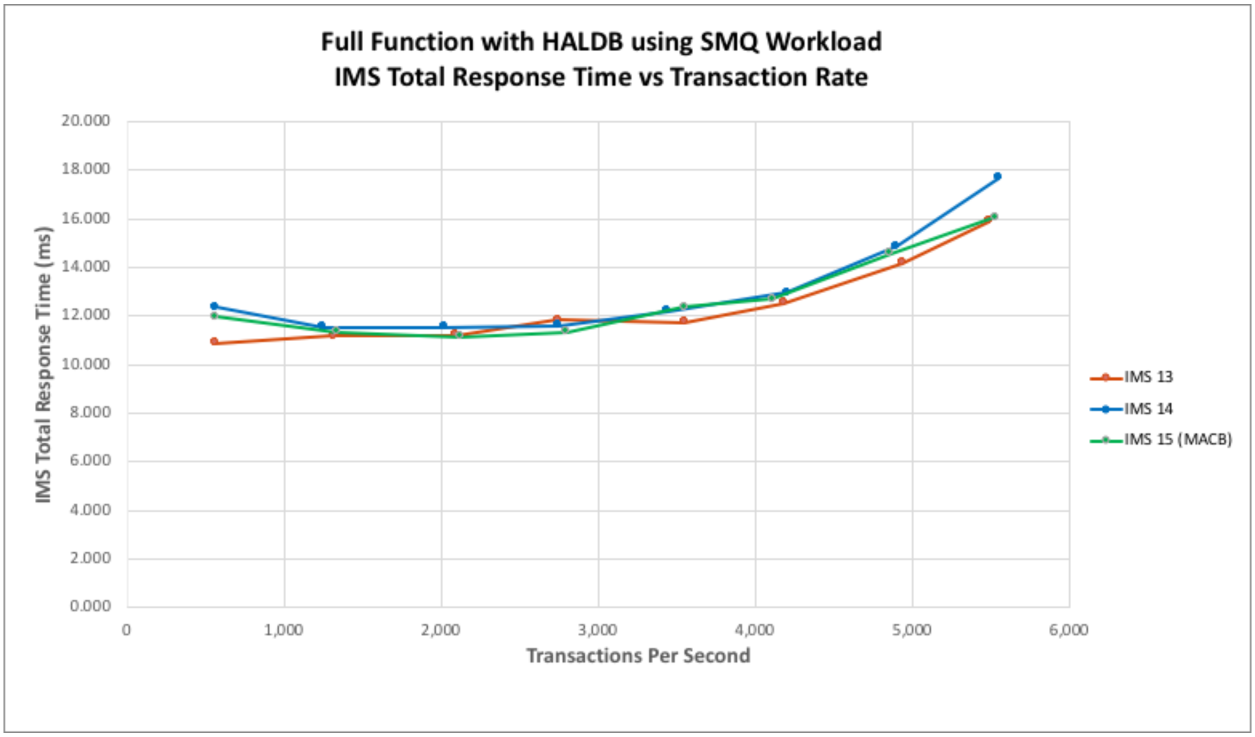

Figure 14: DSFF with HALDB using SMQ IMS Total Response Time vs Transaction Rate Comparison

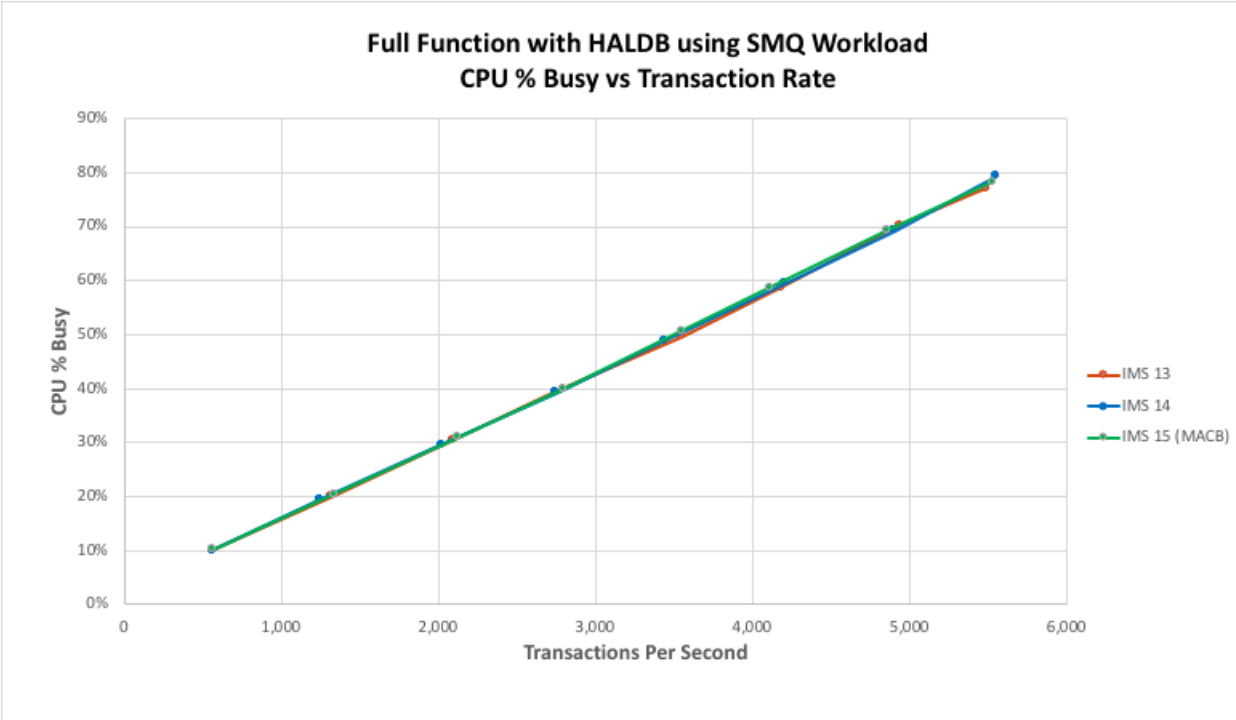

Figure 15: DSFF with HALDB using SMQ CPU % Busy vs Transaction Rate Comparison

6.5 Batch Message Processing (BMP) Banking Performance Evaluation

The objective of the Batch Message Processing (BMP) banking base evaluation was to compare the BMP elapsed time and CPU task time required to execute a set of banking-like BMPs using IMS 13, IMS 14, and IMS 15 Managed ACB in the same hardware configuration.

6.5.1 System Configuration

The Banking BMP base evaluation was executed on a the z14 in a single LPAR configuration as shown in Figure 16.

- LPAR 1 host a single IMS running a single BMP against a FP database with three General Purpose Engines

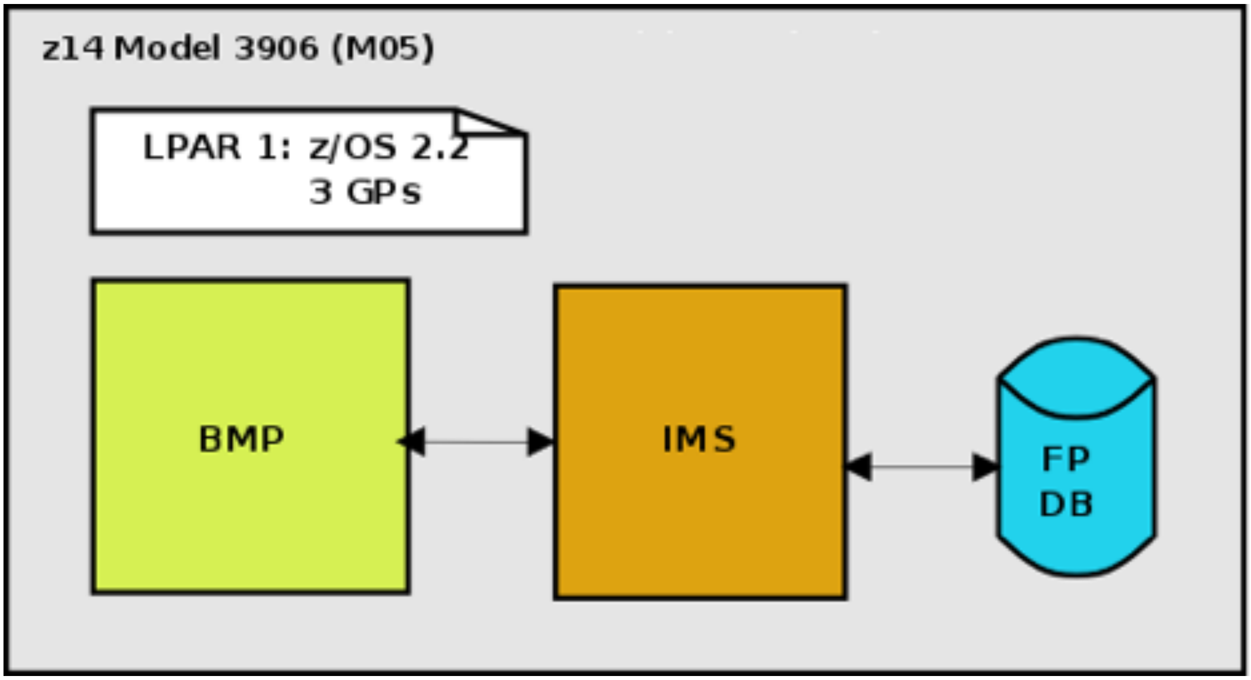

Figure 16: BMP Banking Environment Configuration

6.5.2 Evaluation Results

IMS 15 showed negligible changes in elapsed and CPU task time over IMS 14 for the BMP banking workload. Table 23 shows the comparison between IMS 13, IMS 14, and IMS 15.

Table 23: BMP Banking Evaluation Results

	BMP Banking Workload				
	IMS 13	IMS 14	IMS 15	Delta	Delta %
BMP Execution Time (sec)	309	310	308	-2	-0.65%
Task CPU Task Time (sec)	48.70	49.75	49.09	-0.66	-1.33%
SRB CPU Time (sec)	5.31	5.37	5.43	0.06	1.12%
CPU % Busy	8.34%	8.50%	8.54%	0.04%	0.47%
	Common Storage Below and Above 16MB for Avg. Key 7				

Avg. CSA Below 16MB Key 7	244	224	224	0	0.00%
Avg. CSA Above 16MB Key 7	15.6	15.6	15.6	0.0	0.00%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	619	620	615	-5	-0.81%
Avg. LSQA EPrivate (M)	13.0	13.1	13.4	0.3	2.29%
Avg. USER Private (K)	2644	2272	2268	-4	-0.18%
Avg. USER EPrivate (M)	53.1	55.7	55.0	-0.7	-1.26%
	Private Storage DL/I Region				
Avg. LSQA Private (K)	352	352	348	-4	-1.14%
Avg. LSQA EPrivate (K)	8.6	8.6	8.9	0.3	3.69%
Avg. USER Private (K)	568	624	624	0	0.00%
Avg. USER EPrivate (M)	263	263	263	0	0.00%

The BMP banking workload had minimal changes in both the elapsed and CPU task times between IMS 13, IMS 14, and IMS 15 as shown in Figure 17 and Figure 18.

The BMP banking workload on IMS 15 Managed ACB showed an improvement compared to IMS 14:

- Elapsed time improved by 0.65%
- Total CPU task time improved by 1.33%

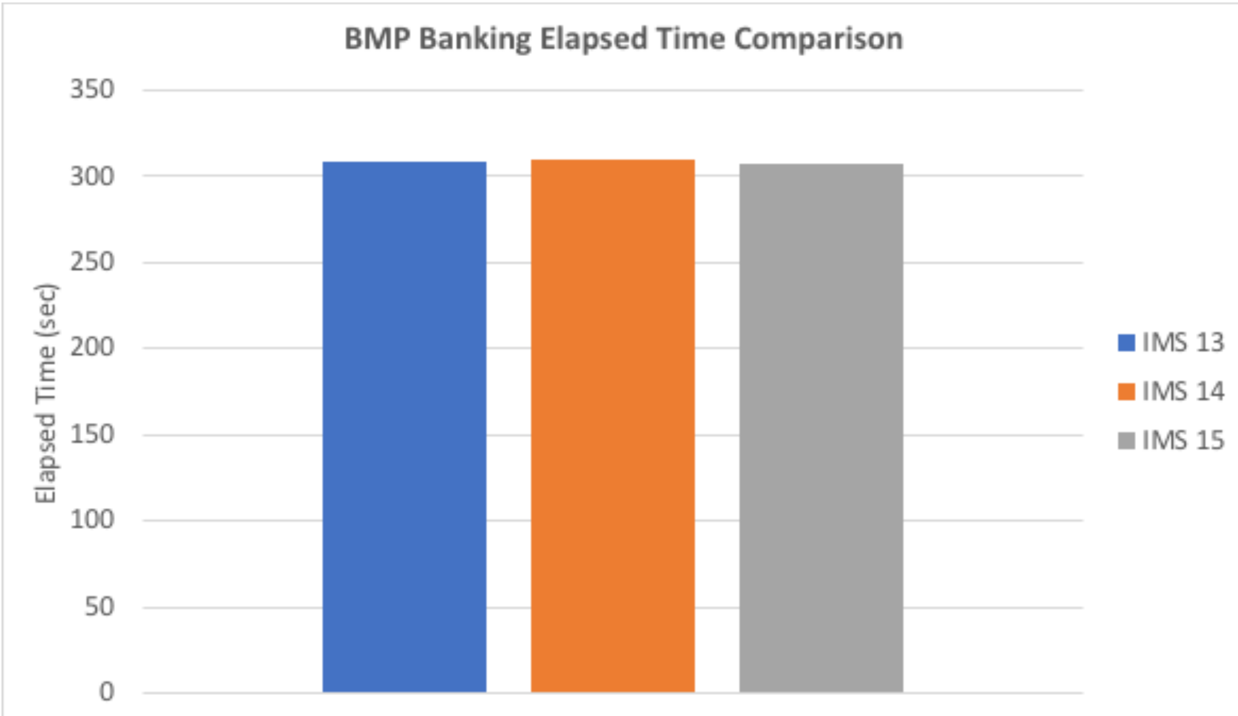

Figure 17: BMP Banking Elapsed Time Comparison

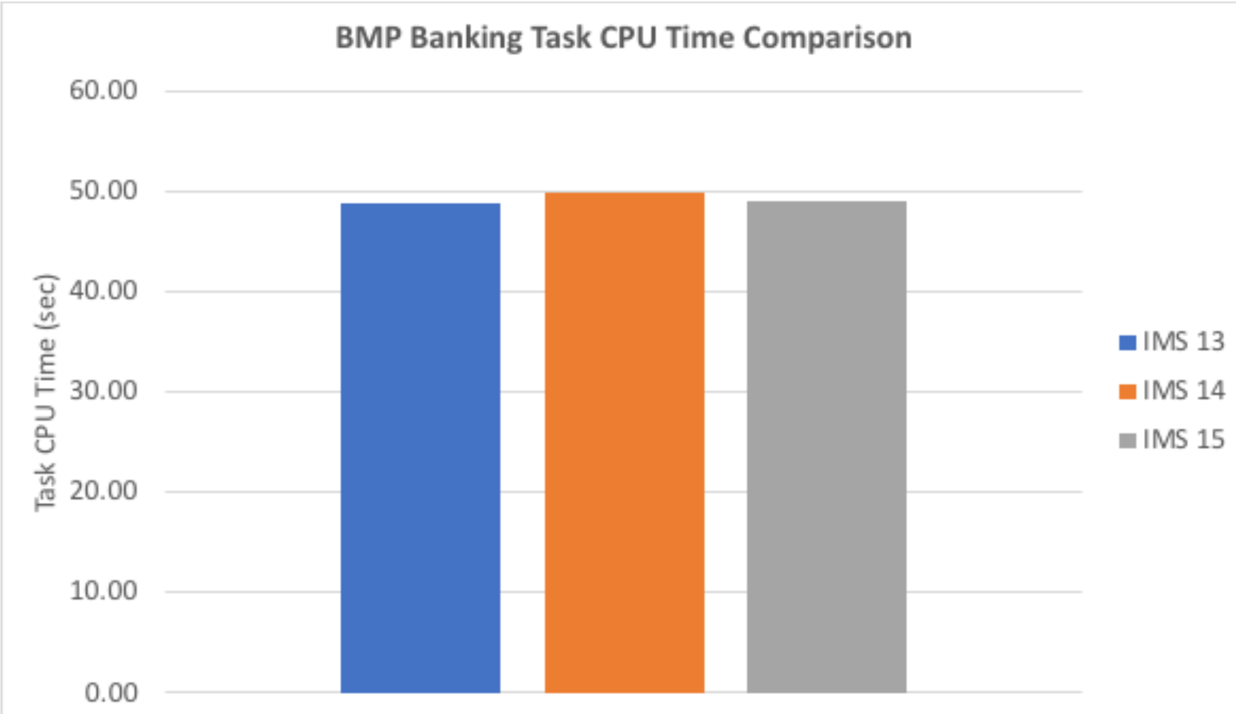

Figure 18: BMP Banking Task CPU Time Comparison

6.6 Customer Information Control System (CICS) - IMS Database Control Performance Evaluation

The CICS Version 5 Release 2 was used with IMS enabled support for Open Transaction Environment (OTE). The CICS programs accessing IMS DBCTL were coded to be reentrant and deployed in a mixture of CONCURRENCY(THREADSAFE) and CONCURRENCY(QUASIRENT) configurations.

The CICS-IMS DBCTL base evaluation was performed for two workloads, IMS Full Function and IMS Fast Path, as described in Section 4.5. The objective of the base evaluation was to compare the ITR between IMS 13, IMS 14, and IMS 15 Managed ACB in the same hardware configuration using both Full Function and Fast Path workloads.

6.6.1 System Configuration

The CICS-IMS DBCTL with Full Function database environment included a two LPAR configuration as shown in Figure 19.

- LPAR 1 hosts CICS and IMS using four General Purpose Engines
- LPAR 2 hosts TPNS driving a total of 2,000 clients via SNA with 10 General Purpose Engines

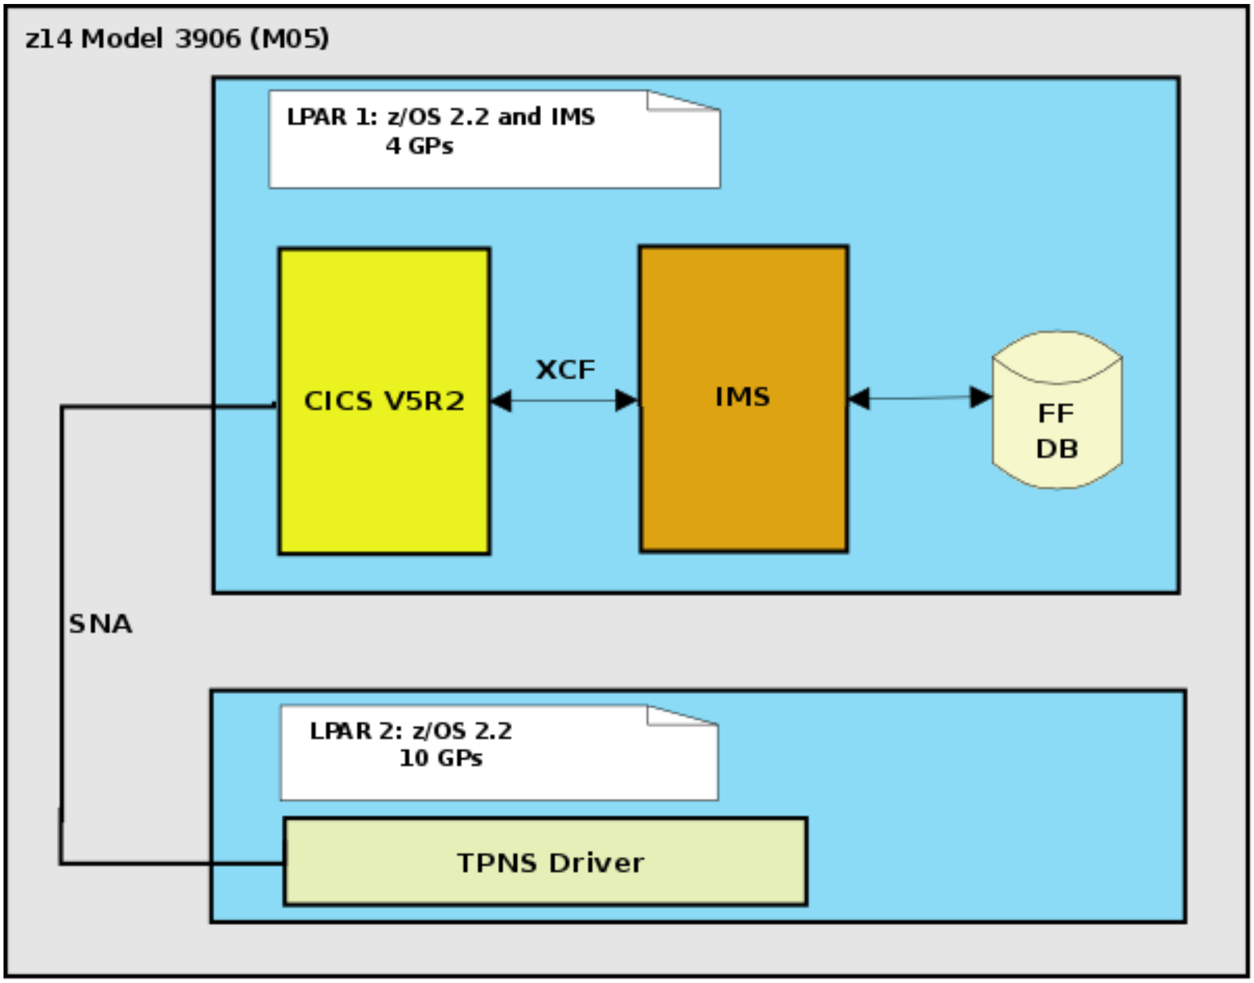

Figure 19: CICS-IMS DBCTL with Full Function DB Environment Configuration

6.6.2 Evaluation Results

IMS 15 Managed ACB showed a less than 1% difference in ITR over IMS 14 for the CICS-IMS DBCTL Full Function workload. Table 24 shows the comparisons between IMS 13, IMS 14, and IMS 15 Managed ACB.

Table 24: CICS-IMS DBCTL with Full Function Performance Results

CICS IMS DBCTL Full Function Workload Evaluation					
	IMS 13	IMS 14	IMS 15	Delta	Delta %
CPU % Busy	79.76%	81.67%	84.52%	2.85%	3.49%
ETR (Tran/Sec)	8480.89	8770.72	9085.75	315.03	3.59%
ITR	10633.01	10739.22	10749.82	10.60	0.10%
Total IMS Response Time (ms)	2.735	2.764	2.985	0.221	8.00%
Total General CPU µs/Tran	376.187	372.467	372.099	-0.367	-0.10%
IMS CPU Service Time/Tran (µs)	13.932	14.043	14.151	0.107	0.76%
CICS CPU Service Time/Tran (µs)	270.524	267.211	265.708	-1.503	-0.56%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	276	256	260	4	1.56%
Avg. CSA Above 16M Key 7 (M)	17.8	17.8	17.8	0.0	0.00%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	604	608	608	0	0.00%
Avg. LSQA EPrivate (M)	12.8	12.9	13.8	0.9	6.98%
Avg. USER Private (K)	2644	2264	2272	8	0.35%
Avg. USER EPrivate (M)	53.3	55.9	55.2	-0.7	-1.25%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	1520	1512	1512	0	0.00%
Avg. LSQA EPrivate (M)	10.9	11.1	11.9	0.8	7.21%
Avg. USER Private (K)	636	692	692	0	0.00%

Avg. USER EPrivate (M)	266	266	266	0	0.00%
	Private Storage CICS Region				
Avg. LSQA Private (K)	611	605	615	10	1.65%
Avg. LSQA EPrivate (M)	14.7	14.7	14.7	0.0	0.00%
Avg. USER Private (K)	4450	4464	4462	-2	-0.04%
Avg. USER EPrivate (M)	98.5	98.1	98.4	0.3	0.31%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 Managed ACB at various CPU percent busy values. Figure 20 shows the ITR vs Transaction Rate comparison and Figure 21 shows the IMS Response Time vs Transaction Rate comparison. Figure 22 shows the CPU % Busy vs Transaction Rate.

The CICS-IMS DBCTL Full Function workload on IMS 15 Managed ACB showed improvement in ITR and degradation in Total IMS response time compared to IMS 14.

- ITR improved by 0.10%
- Total IMS response time increased by 0.221 milliseconds which caused a degradation of 8%.

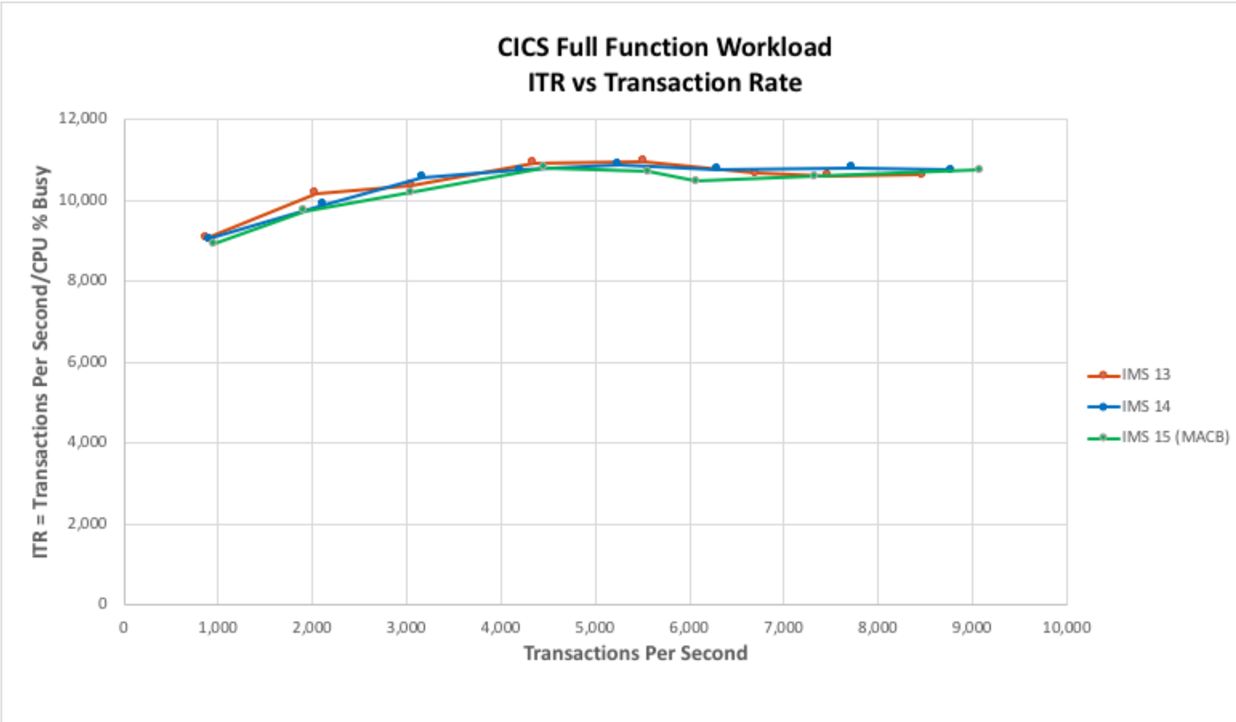

Figure 20: CICS-IMS DBCTL with Full Function ITR vs Transaction Rate Comparison

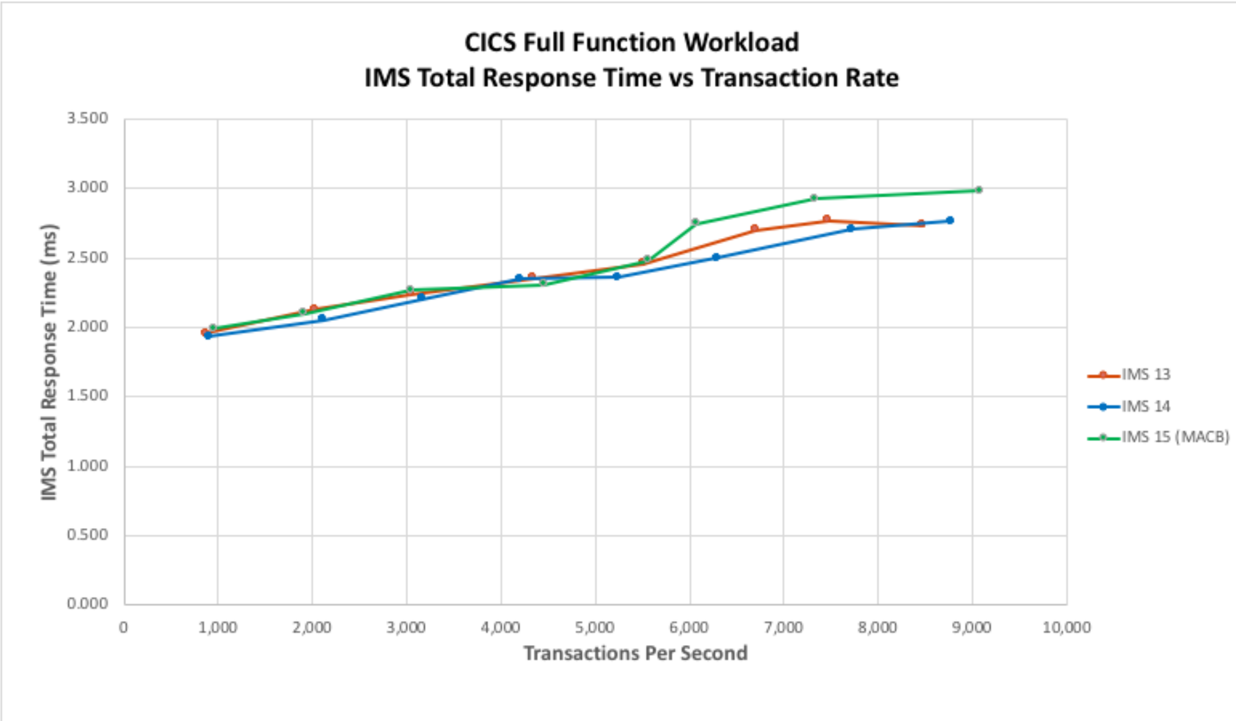

Figure 21: CICS-IMS DBCTL with Full Function Total Response Rate vs Transaction Rate Comparison

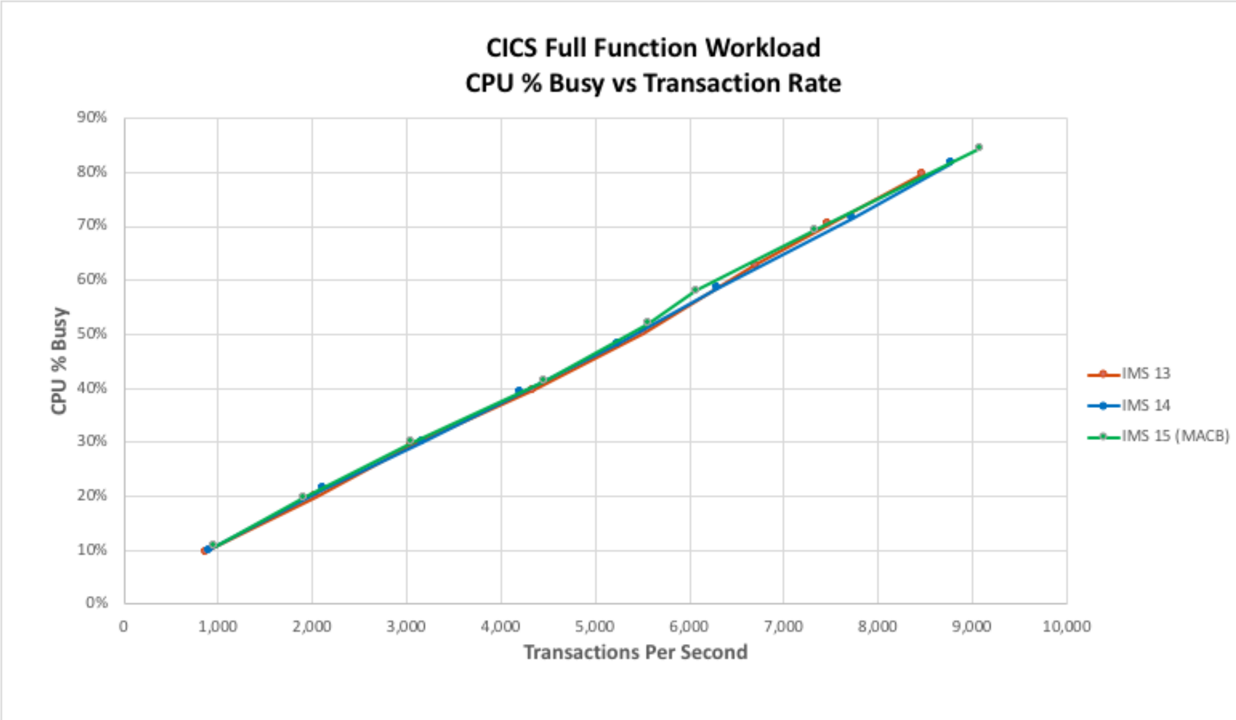

Figure 22: CICS-IMS DBCTL with Full Function CPU % Busy vs Transaction Rate Comparison

6.6.3 System Configuration

The CICS-IMS DBCTL with Fast Path database environment was executed in a two LPAR configuration as shown in Figure 23:

- LPAR 1 hosts CICS and IMS with three General Purpose Engine
- LPAR 2 hosts TPNS driving a total of 2,000 clients via SNA with 10 General Purpose Engines

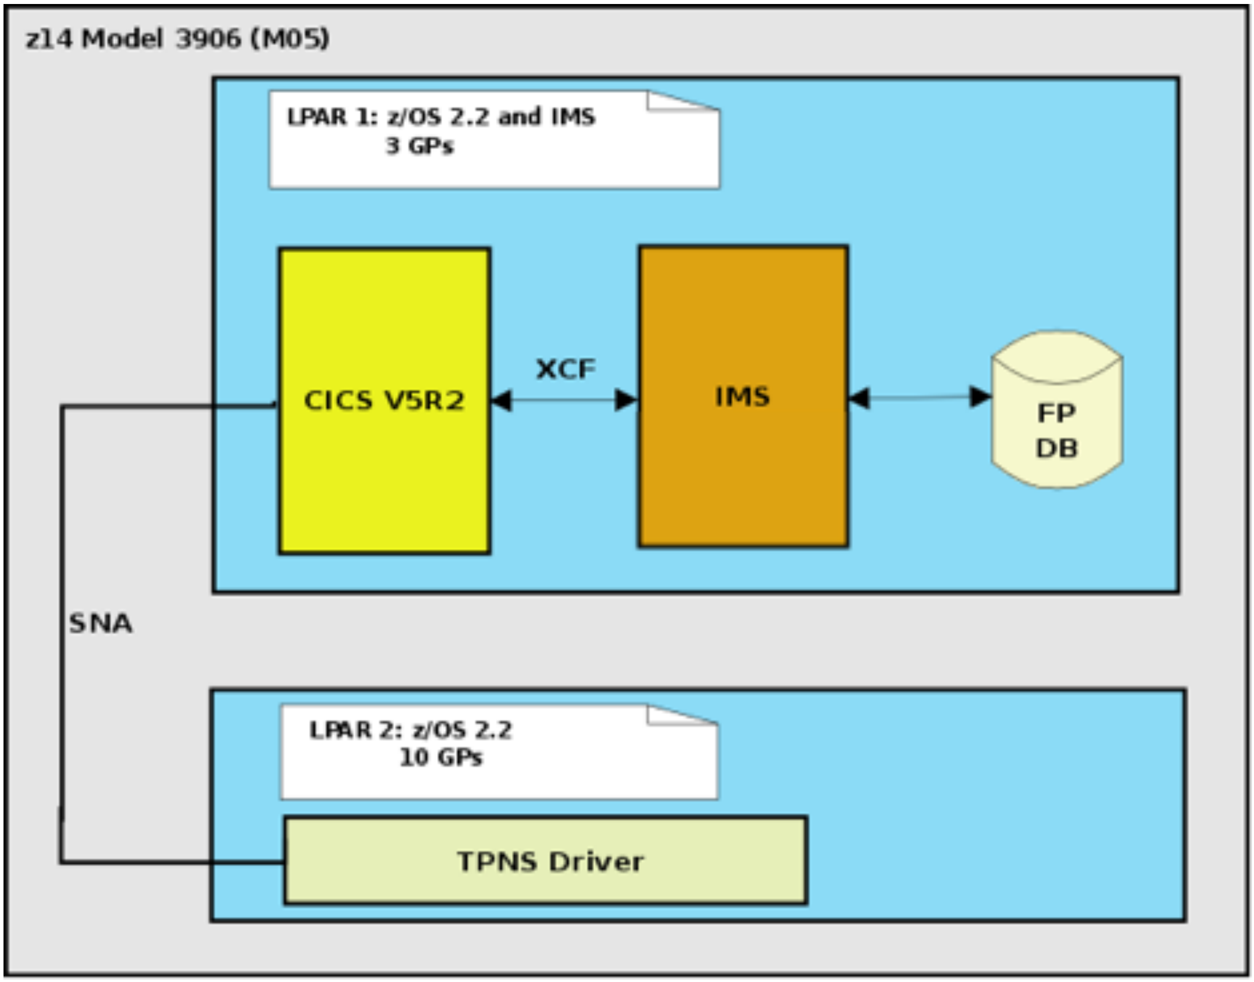

Figure 23: CICS-IMS DBCTL with Fast Path Environment Configuration

6.6.4 Evaluation Results

The delta between IMS 15 Managed ACB and IMS 14 showed a difference of less than 1% in ITR and about a 2% improvement in the total IMS response time for the CICS-IMS DBCTL Fast Path workload. Table 25 shows the comparisons between IMS 13, IMS 14, and IMS 15.

Table 25: CICS-IMS DBCTL Fast Path Performances Results

CICS IMS DBCTL Fast Path Workload Evaluation					
	IMS 13	IMS 14	IMS 15	Delta	Delta %

CPU % Busy	80.48%	80.64%	80.58%	-0.06%	-0.07%
ETR (Tran/Sec)	10890.03	10854.56	10907.57	53.01	0.49%
ITR	13531.35	13460.52	13536.32	75.80	0.56%
Total IMS Response Time (ms)	0.982	0.980	0.960	-0.020	-2.04%
Total General CPU µs/Tran	221.707	222.874	221.626	-1.248	-0.56%
IMS CPU Service Time/Tran (µs)	22.604	22.558	23.117	0.559	2.48%
CICS CPU Service Time/Tran (µs)	117.983	118.667	117.718	-0.949	-0.80%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	284	264	264	0	0.00%
Avg. CSA Above 16M Key 7 (M)	32.0	31.9	32.0	0.1	0.31%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	1224	1224	1220	-4	-0.33%
Avg. LSQA EPrivate (M)	13.5	13.6	13.9	0.3	2.21%
Avg. USER Private (K)	2648	2268	2272	4	0.18%
Avg. USER EPrivate (M)	63.8	66.4	65.7	-0.7	-1.05%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	352	352	348	-4	-1.14%
Avg. LSQA EPrivate (M)	8.5	8.5	8.9	0.4	4.74%
Avg. USER Private (K)	568	624	628	4	0.64%
Avg. USER EPrivate (M)	264	264	265	1	0.38%
	Private Storage CICS Region				

Avg. LSQA Private (K)	611	607	604	-3	-0.49%
Avg. LSQA EPrivate (M)	14.5	14.5	14.5	0.0	0.00%
Avg. USER Private (K)	4430	4441	4428	-13	-0.29%
Avg. USER EPrivate (M)	98.3	98.1	98.1	0.0	0.00%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 24 shows the ITR vs Transaction Rate comparison and Figure 25 shows the IMS Response Time vs Transaction Rate comparison. Figure 26 shows the CPU % Busy vs Transaction Rate.

The CICS-IMS DBCTL Fast Path workload on IMS 15 Managed ACB showed improvement in both ITR and Total IMS response time compared to IMS 14.

- ITR improved by 0.56%
- Total IMS response time improved by 2.04%

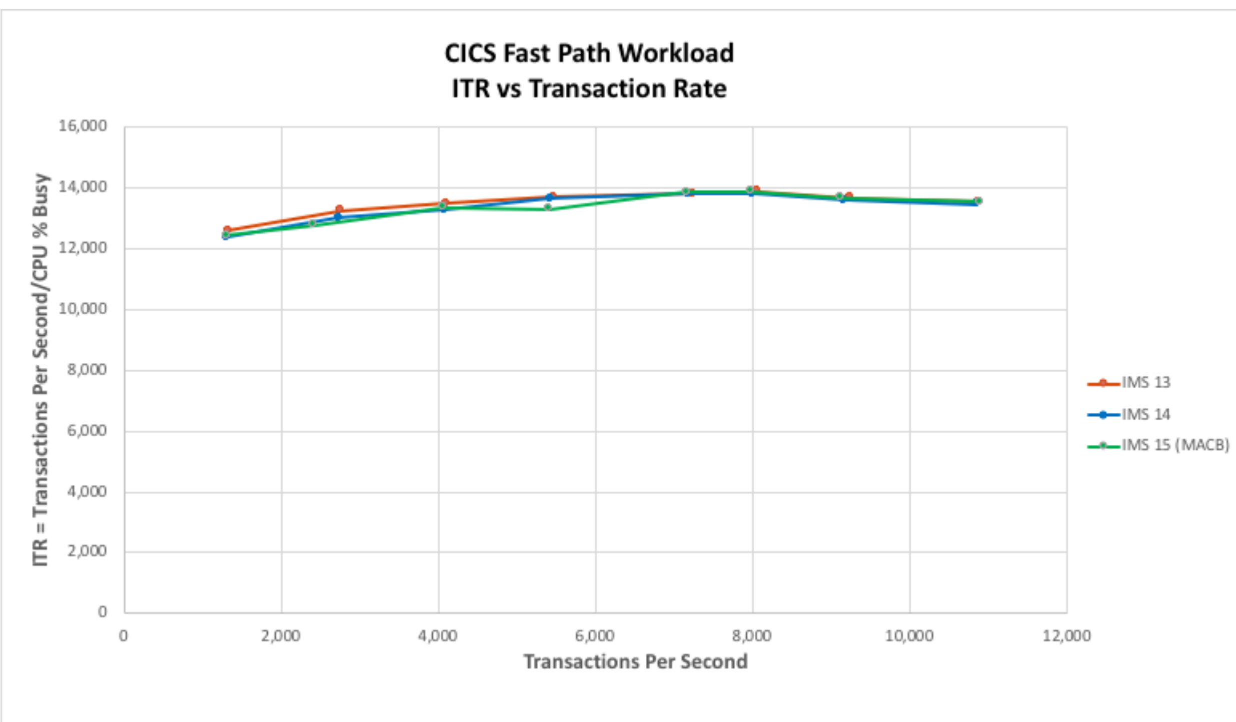

Figure 24: CICS-IMS DBCTL Fast Path ITR vs Transaction Rate Comparison

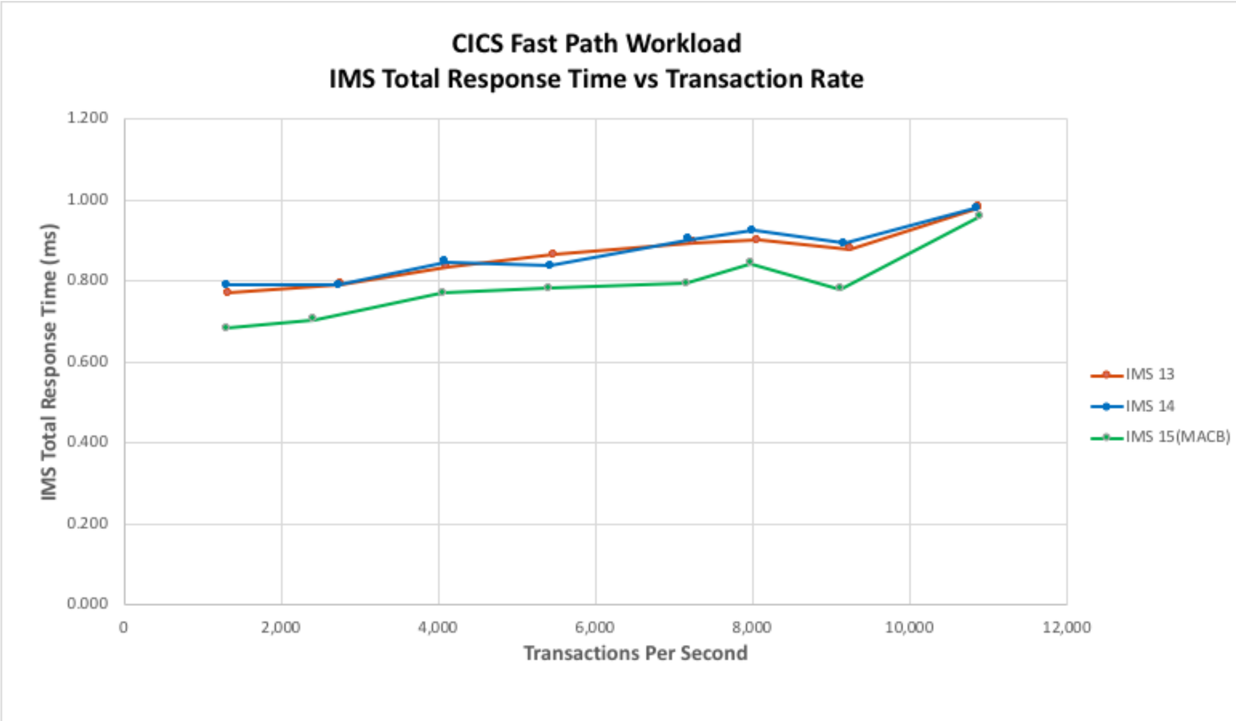

Figure 25: CICS-IMS DBCTL Fast Path IMS Total Response Time vs Transaction Rate Comparison

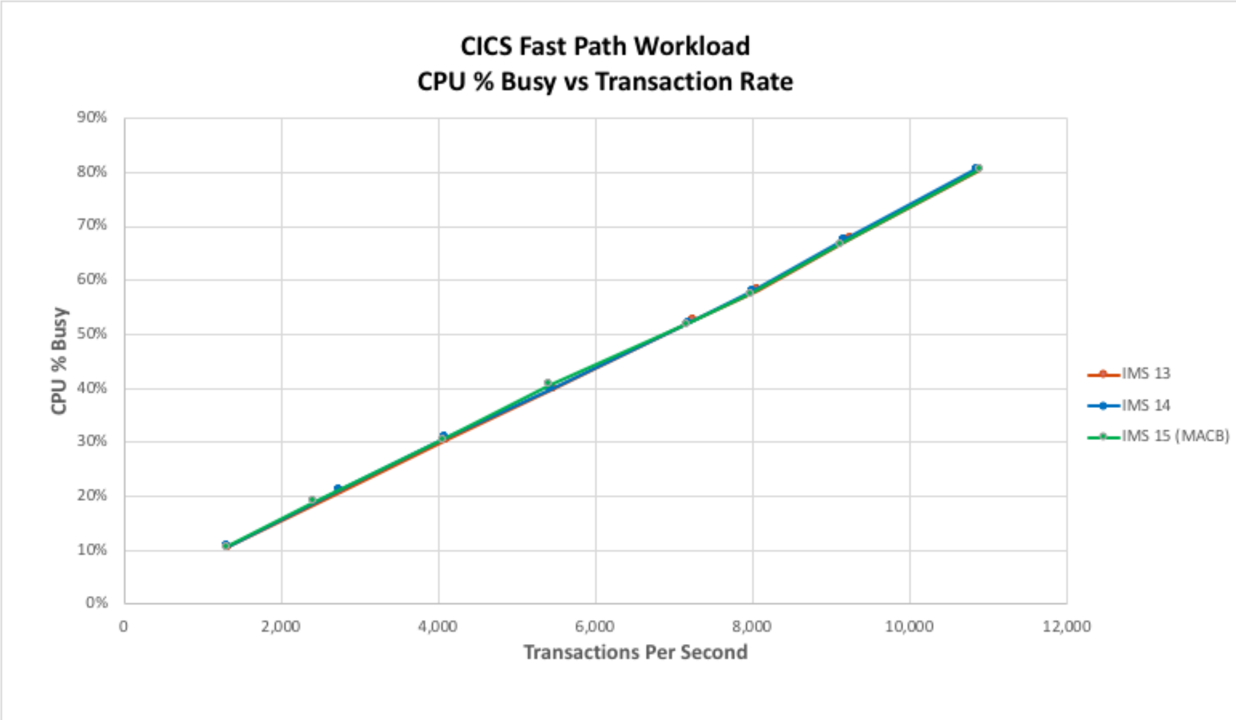

Figure 26: CICS-IMS DBCTL Fast Path CPU % Busy vs Transaction Rate Comparison

6.7 IMS TM-Db2 IRWW (Full PL/I) Performance Evaluation

IMS provides access to external Db2® for z/OS® databases through the External Subsystem Attach Facility (ESAF) by acting as the transaction manager (TM). The IMS TM-Db2 IBM Retail Warehouse Workload (IRWW) performance testing used ESAF to access a Db2® for z/OS® relational database which is based on a retail type of environment.

6.7.1 System Configuration

The IMS TM-Db2 IRWW base evaluation processes the IMS transactions in 66 IFP and two MPP regions accessing the Db2® for z/OS® retail warehouse database through ESAF.
The workload was executed on a z14 configured in a two LPAR environment as shown in Figure 27:

- LPAR 1 hosts IMS, Db2, 66 IFPs and 2 MPP regions, and ESAF with seven General Purpose Engines

- LPAR 2 hosts a Java-based workload driver driving 38 clients via TCPIP and two General Purpose Engines

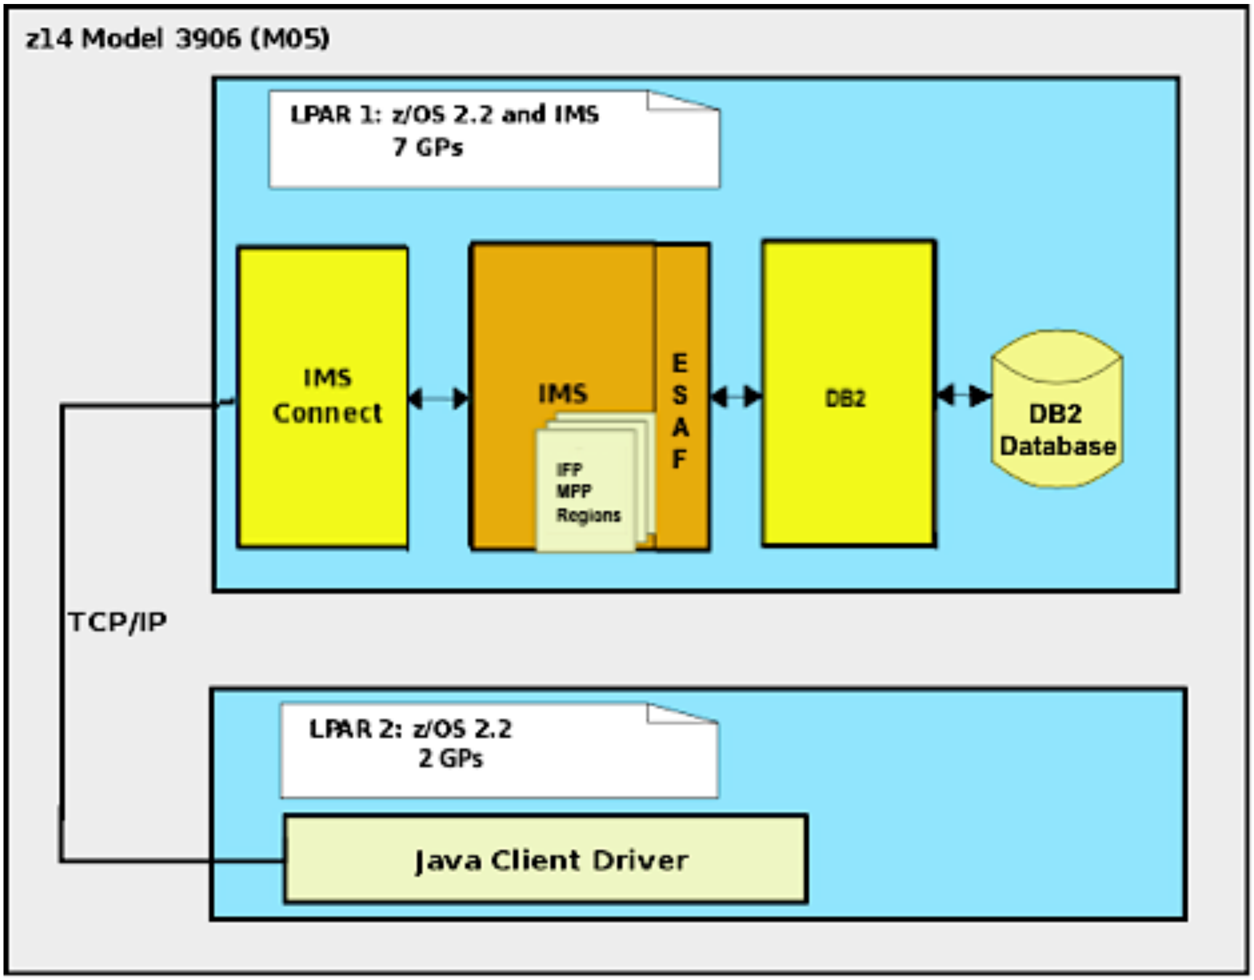

Figure 27: IMS TM-Db2 Configuration

6.7.2 Evaluation Results

The ITR comparison between IMS 15 Managed ACB and IMS 14 performance showed an improvement of 4% for the IMS TM-Db2 IRWW workload. Table 26 shows the results of the workload for IMS 13, IMS 14 and IMS 15.

Table 26: IMS TM-Db2 Evaluation Results

IMS TM-Db2 Evaluation					
	IMS 13	**IMS 14**	**IMS 15**	**Delta**	**Delta**
CPU % Busy	79.57%	80.72%	80.07%	-0.65%	-0.81%
ETR (Tran/Sec)	6799.87	6857.50	7076.77	219.27	3.20%
ITR	8545.77	8495.42	8838.23	342.81	4.04%
Total IMS Response Time (ms)	4.903	4.842	4.783	-0.059	-1.22%
Total General CPU µs/Tran	819.119	823.974	792.014	-31.960	-3.88%
IMS CPU Service Time/Tran (µs)	14.272	14.389	13.967	-0.422	-2.93%
ICON CPU Service Time/Tran (µs)	27.221	24.402	22.771	-1.631	-6.68%
IFP/MPP CPU Service Time/Tran (µs)	553.765	554.970	540.225	-14.745	-2.66%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	284	264	264	0	0.00%
Avg. CSA Above 16M Key 7 (M)	33.4	33.3	33.3	0.0	0.00%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	648	652	652	0	0.00%
Avg. LSQA EPrivate (M)	14.1	14.1	14.5	0.4	2.84%
Avg. USER Private (K)	2736	2364	2372	8	0.34%
Avg. USER EPrivate (M)	55.5	58.1	56.5	-1.6	-2.75%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	352	352	348	-4	-1.14%

Avg. LSQA EPrivate (M)	8.6	8.5	9.0	0.5	5.64%
Avg. USER Private (K)	568	624	624	0	0.00%
Avg. USER EPrivate (M)	263	263	263	0	0.00%
	Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.7	13.5	13.5	0.0	0.00%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	174	174	176	2	1.15%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 28 shows the ITR vs Transaction Rate, Figure 29 shows the IMS Response Time vs Transaction Rate, and Figure 30 shows the CPU % Busy vs Transaction Rate.

The IMS TM-Db2 IRWW workload on IMS 15 Managed ACB showed improvement in ITR and Total IMS response time compared to IMS 14:

- ITR improved by 4.04%
- Total IMS response time improved by 1.22%

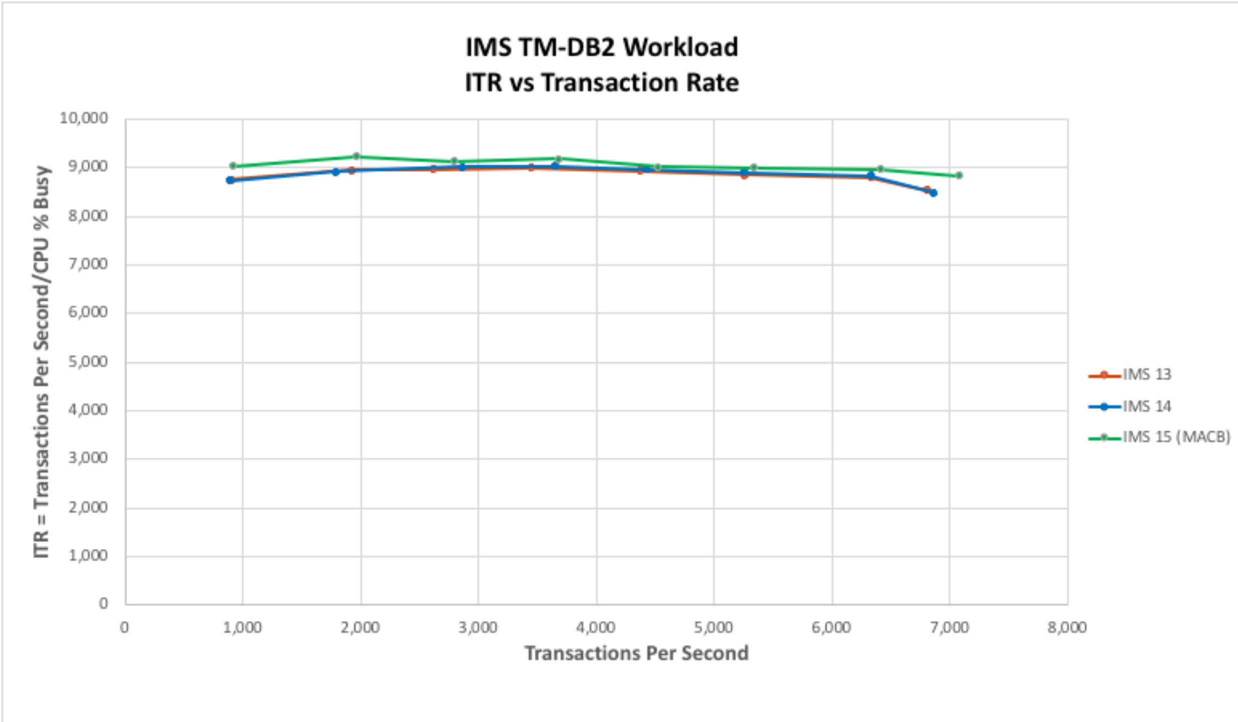

Figure 28: IMS TM-Db2 ITR vs Transaction Rate Comparison Results

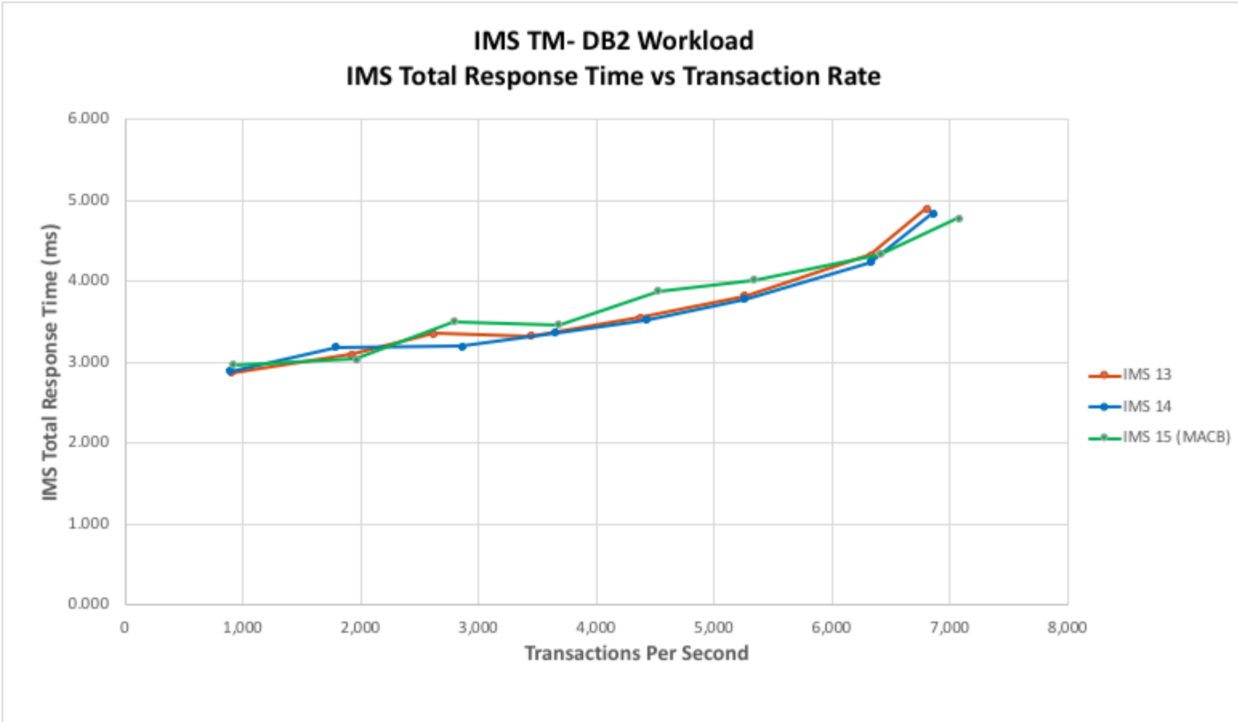

Figure 29: IMS TM-Db2 IMS Total Response Time vs Transaction Rate Comparison Results

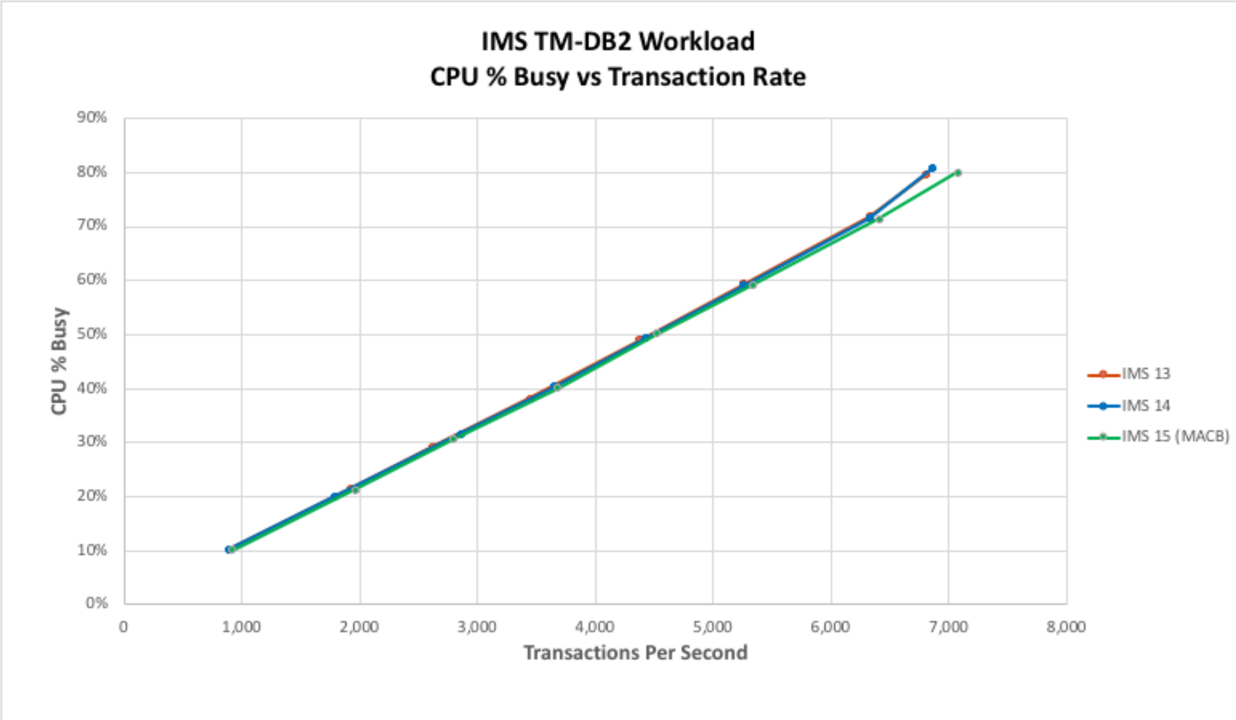

Figure 30: IMS TM-Db2 CPU % Busy vs Transaction Rate Comparison Results

6.8 Java Message Processing (JMP) 31-bit Performance Evaluation

The objective of the JMP 31-bit evaluation was to compare the ITR between IMS 13, IMS 14, and IMS 15 Managed ACB in the same hardware configuration. The comparison between JMP 31-bit and 64-bit performance is found in Section 8.3.

6.8.1 System Configuration

The JMP workload evaluation was executed on the z14 in a two LPAR configuration as shown in Figure 31:

- LPAR 1 hosts IMS, IMS Connect, and 200 JMP regions with six General Purpose Engines
- LPAR 2 hosts a Java-based workload driver driving 38 clients via TCPIP and 10 General Purpose Engines

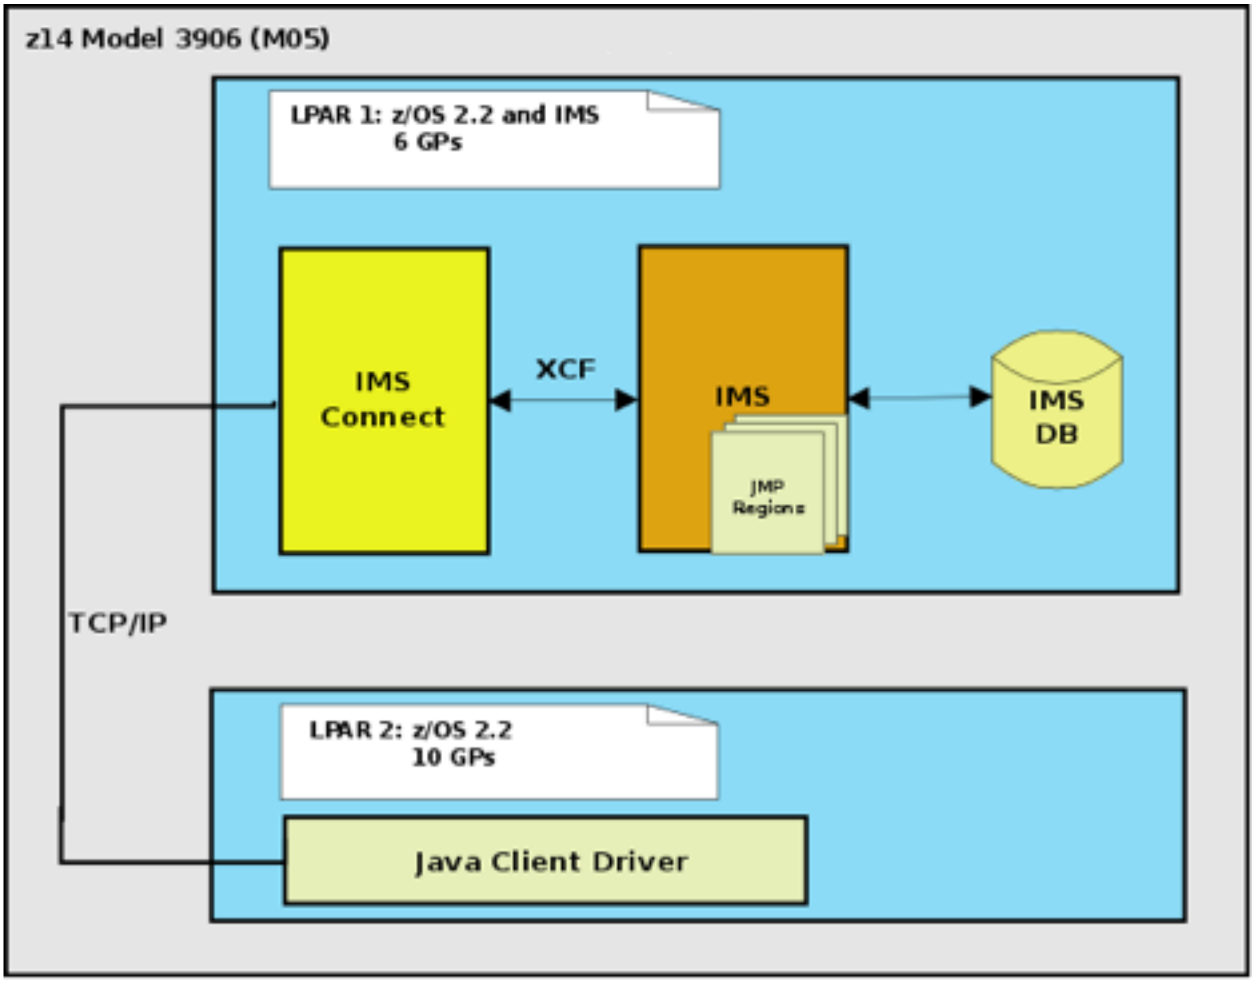

Figure 31: Java Message Processing Configuration

6.8.2 Evaluation Results

The JMP performance in IMS 15 Managed ACB had about a 3.5% increase in ITR over IMS 14 for the JMP workload. Table 27 shows the comparison between IMS 13, IMS 14, and IMS 15 Managed ACB.

Table 27: Java Message Processing Performance Result

Java Message Processing with GPs					
	IMS 13	IMS 14	IMS 15	Delta	Delta %

CPU % Busy	81.76%	78.92%	78.88%	-0.04%	-0.05%
ETR (Tran/Sec)	27210.01	28054.28	29031.51	977.23	3.48%
ITR	33280.34	35547.74	36804.65	1256.91	3.54%
Total IMS Response Time (ms)	0.795	0.759	0.757	-0.002	-0.26%
Total General CPU µs/Tran	180.29	168.79	163.02	-5.77	-3.42%
IMS CPU Service Time/Tran (µs)	11.54	11.64	12.12	0.48	4.07%
ICON CPU Service Time/Tran (µs)	18.008	17.727	16.088	-1.639	-9.25%
JMP CPU Service Time/Tran (µs)	137.126	125.703	121.827	-3.876	-3.08%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	304	284	284	0	0.00%
Avg. CSA Above 16M Key 7 (M)	20.6	20.8	20.7	-0.1	-0.48%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	604	608	608	0	0.00%
Avg. LSQA EPrivate (M)	13.2	13.3	13.8	0.5	3.76%
Avg. USER Private (K)	2636	2264	2272	8	0.35%
Avg. USER EPrivate (M)	55.9	58.7	57.0	-1.6	-2.73%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	384	388	384	-4	-1.03%
Avg. LSQA EPrivate (M)	8.7	8.8	9.1	0.3	4.14%
Avg. USER Private (K)	572	628	628	0	0.00%

Avg. USER EPrivate (M)	264	264	264	0	0.00%
	Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.9	14	13.8	-0.2	-1.43%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	174	174	176	2	1.15%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 32 shows the ITR vs Transaction Rate comparison and Figure 33 shows the IMS Response Time vs Transaction Rate comparison. Figure 34 shows the CPU % Busy vs Transaction Rate comparison.

The JMP workload on IMS 15 Managed ACB showed improvements in ITR and Total IMS response time compared to IMS 14.

- ITR improved by 3.54%
- Total IMS response time improved by 0.26%

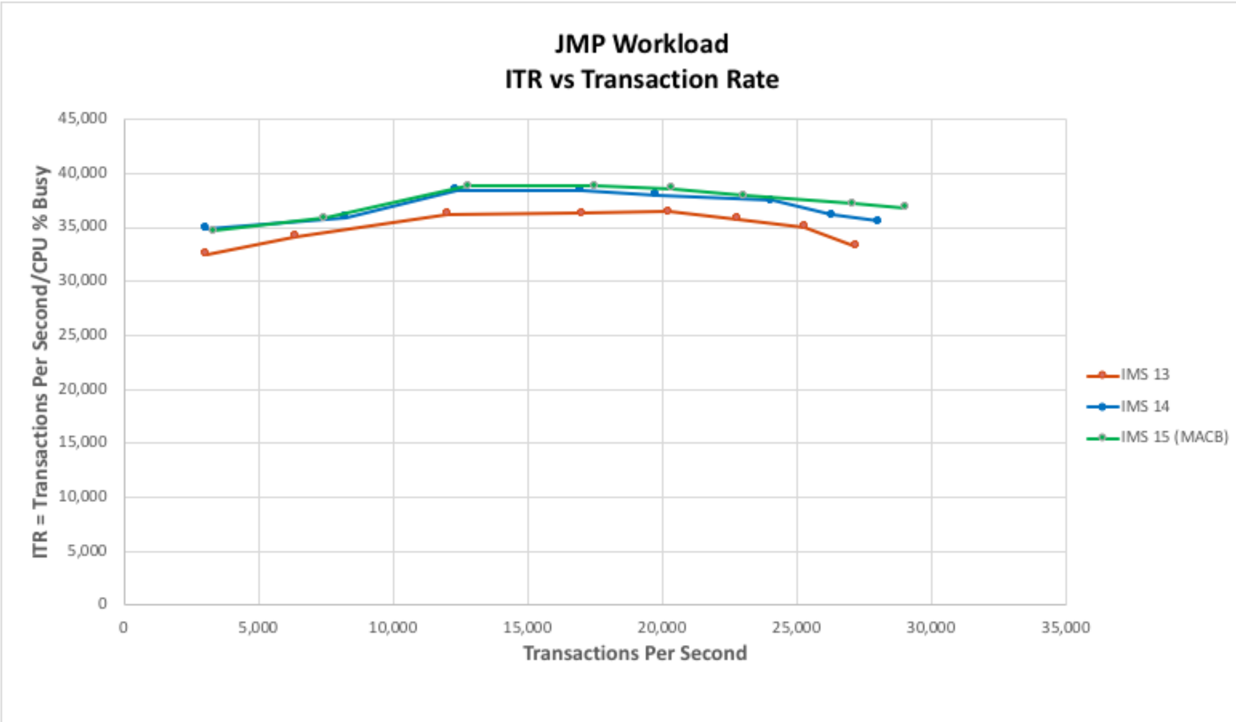

Figure 32: JMP ITR vs Transaction Rate Comparison

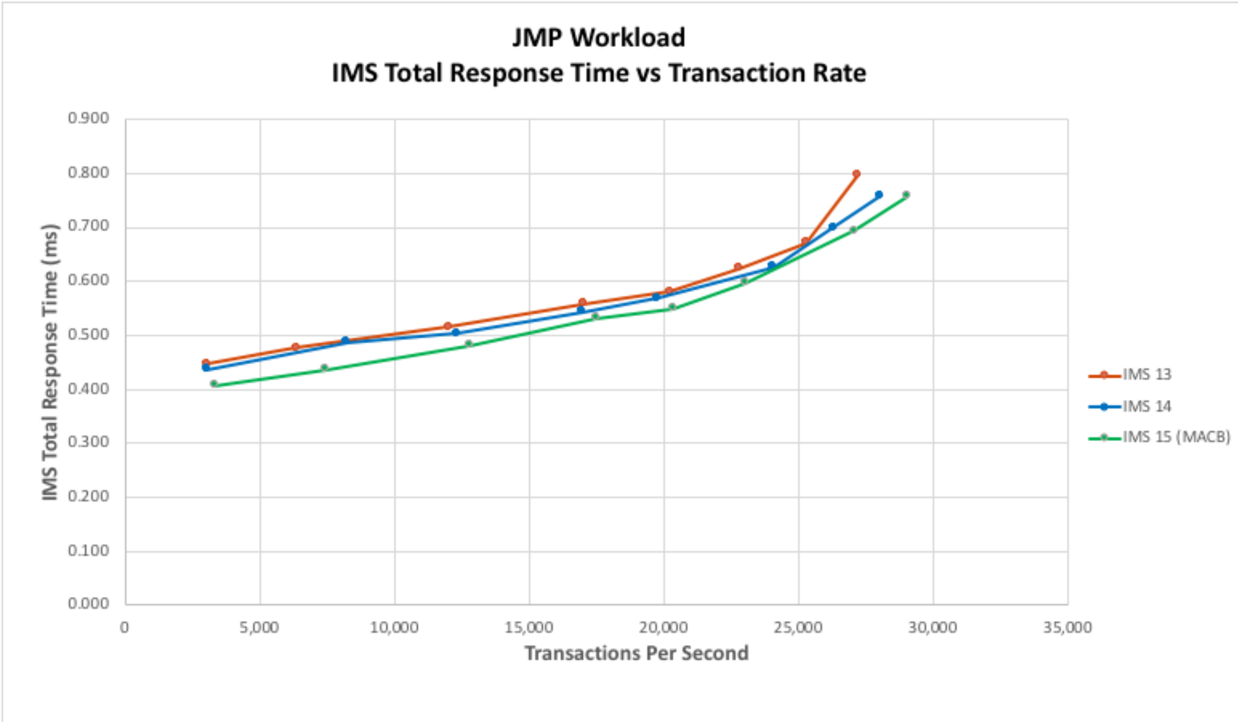

Figure 33: JMP IMS Total Response Time vs Transaction Rate Comparison

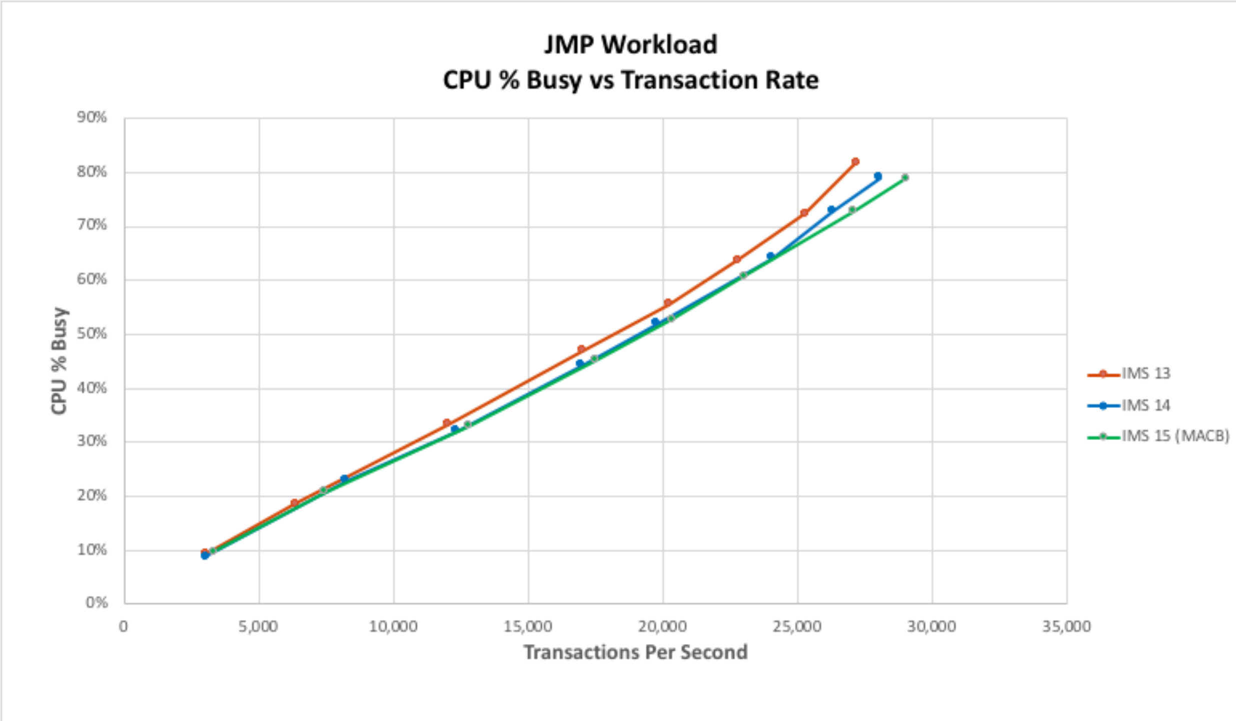

Figure 34: JMP CPU % Busy vs Transaction Rate Comparison

6.9 z/OS Connect Enterprise Edition (z/OS Connect EE) IMS Service Provider Performance Evaluation

The objective of the z/OS Connect EE IMS Service Provider (SP) evaluation was to compare the ITR between IMS 15 Managed ACB and IMS 14 in the same hardware configuration using z/OS Connect EE V3.0.10.

6.9.1 System Configuration

The z/OS Connect EE IMS SP evaluation was executed on the z14 in a two LPAR configuration as shown in Figure 35:

- LPAR 1 hosts z/OS Connect EE with IMS SP, IMS, 350 MPP regions and IMS Connect with eight General Purpose Engines
- LPAR 2 hosts a Java-based workload driver driving 30 clients via TCPIP and 10 General Purpose Engines

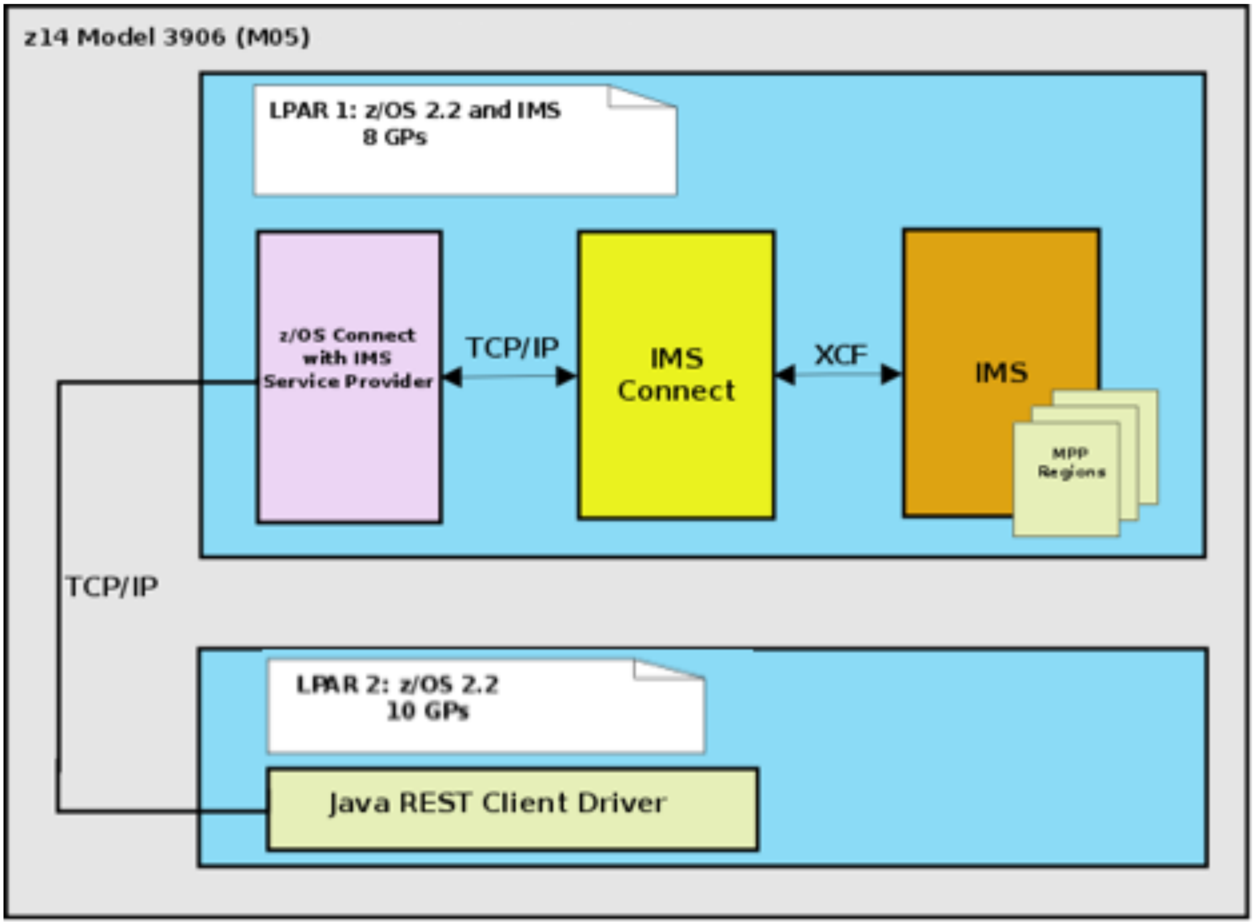

Figure 35: z/OS Connect EE IMS SP Workload Environment Configuration

6.9.2 Evaluation Results

The z/OS Connect EE IMS SP performance in IMS 15 Managed ACB showed less than a 1% difference in ITR over IMS 14. Table 28 shows the comparisons between IMS 13, IMS 14, and IMS 15.

Table 28: z/OS Connect EE IMS SP Comparison Results

z/OS Connect IMS Service Provider Workload Evaluation GPs only					
	IMS 13	IMS 14	IMS 15	Delta	Delta %
CPU % Busy	80.86%	80.47%	81.05%	0.58%	0.72%

ETR (Tran/Sec)	18684.29	18698.80	18808.05	109.25	0.58%
ITR	23106.96	23236.98	23205.49	-31.49	-0.14%
Total IMS Response Time (ms)	0.147	0.139	0.145	0.006	4.32%
Total General CPU µs/Tran	346.22	344.28	344.75	0.47	0.14%
IMS CPU Service Time/Tran (µs)	12.72	12.74	13.22	0.47	3.72%
ICON CPU Service Time/Tran (µs)	22.031	21.590	19.851	-1.739	-8.05%
MPP CPU Service Time/Tran (µs)	25.958	25.295	25.307	0.012	0.05%
z/OS Connect EE **CPU Service Time/Tran (µs)**	264.744	264.572	266.398	1.826	0.69%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	352	332	332	0	0.00%
Avg. CSA Above 16M Key 7 (M)	23.7	23.7	23.7	0.0	0.00%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	604	608	608	0	0.00%
Avg. LSQA EPrivate (M)	13.2	13.3	13.8	0.5	3.76%
Avg. USER Private (K)	2640	2264	2268	4	0.18%
Avg. USER EPrivate (M)	55.5	58.2	56.7	-1.5	-2.58%
	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	352	352	348	-4	-1.14%
Avg. LSQA EPrivate (M)	8.6	8.6	8.9	0.3	3.85%

Avg. USER Private (K)	568	624	628	4	0.64%
Avg. USER EPrivate (M)	263	264	264	0	0.00%
	Private Storage ICON Region				
Avg. LSQA Private (K)	360	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.7	13.7	13.7	0.0	0.00%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	174	174	176	2	1.15%

A series of scaling tests were run to compare IMS 13, IMS 14, and IMS 15 at various CPU percent busy values. Figure 36 shows the ITR vs Transaction Rate comparison and Figure 37 shows the IMS Response Time vs Transaction Rate comparison. Figure 38 shows the CPU % Busy vs Transaction Rate comparison.

The z/OS Connect EE IMS SP workload showed a degradation in ITR and Total IMS response time when comparing IMS 15 Managed ACB to IMS 14.

- ITR degraded by 0.14%
- Total IMS response time increased by 0.006 milliseconds which caused a degradation of 4.32%

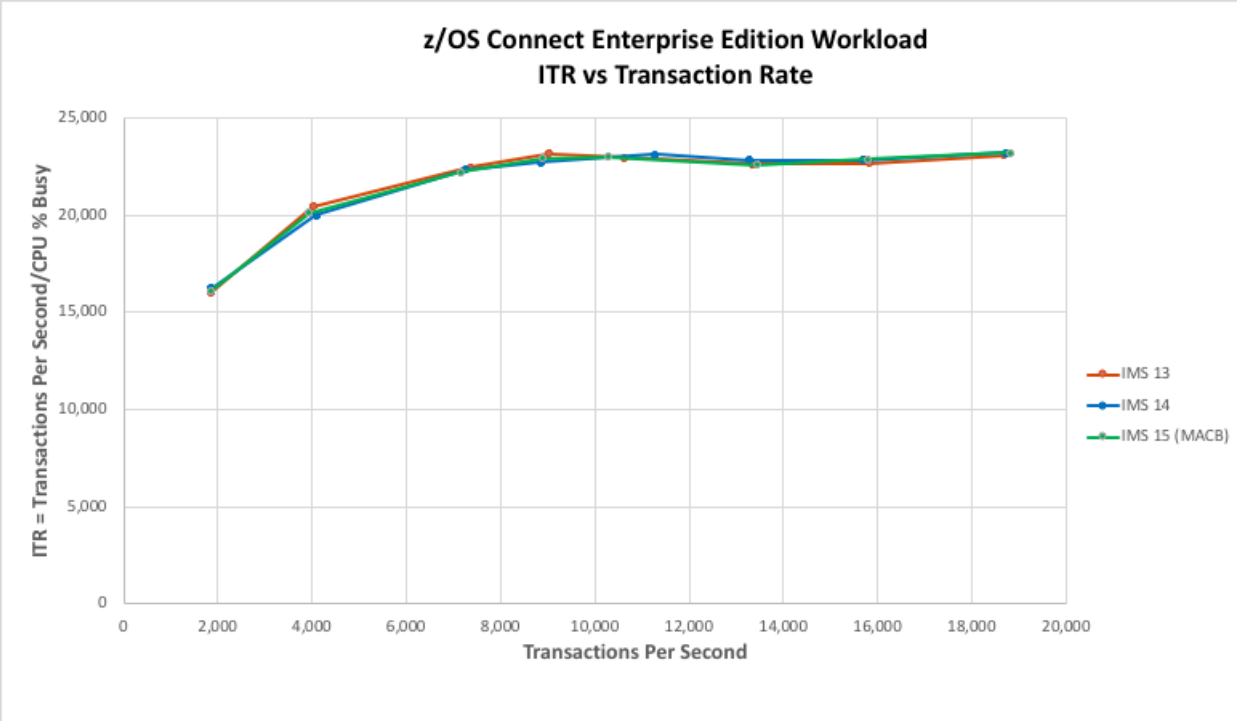

Figure 36: z/OS Connect EE SP ITR vs Transaction Rate Comparison

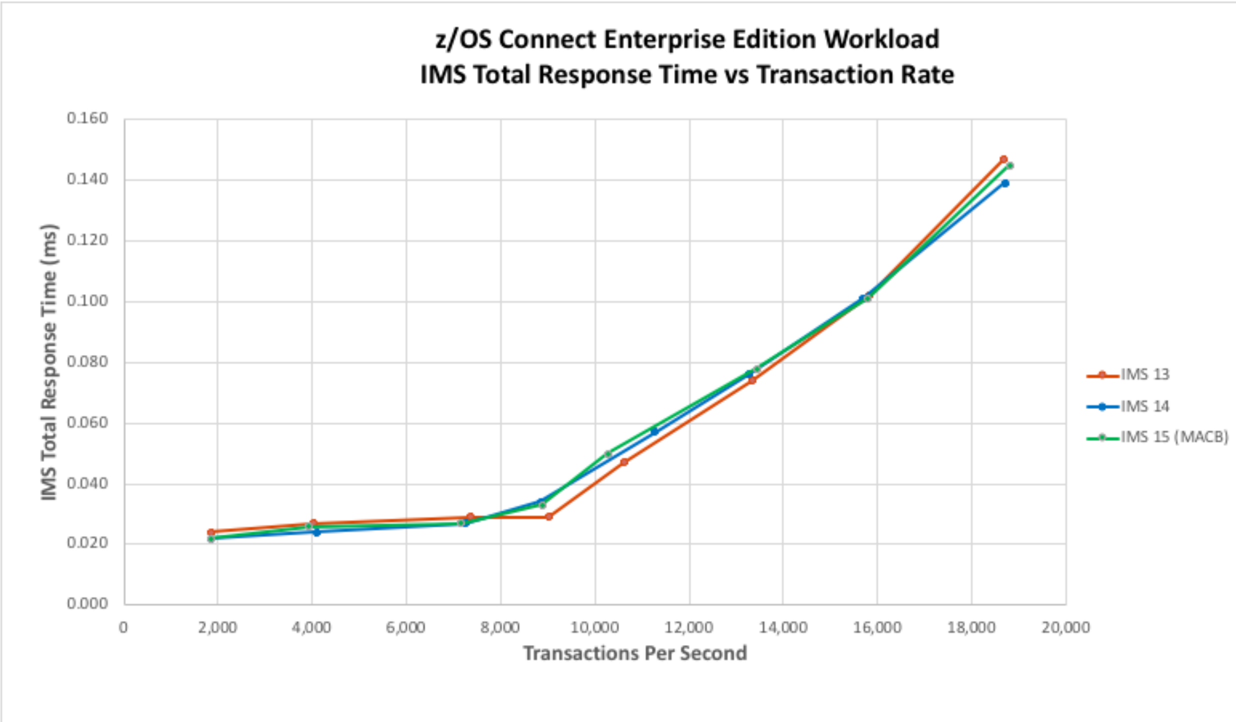

Figure 37: z/OS Connect EE SP IMS Total Response Time vs Transaction Rate Comparison

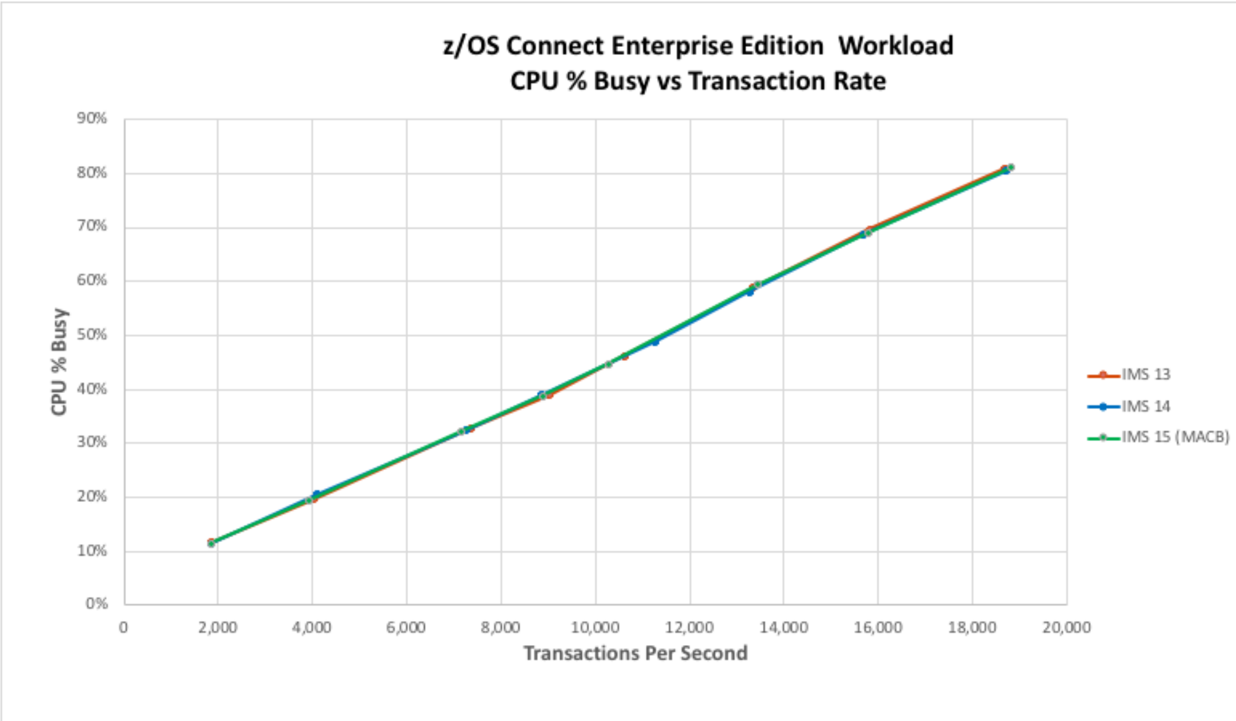

Figure 38: z/OS Connect EE SP CPU % Busy vs Transaction Rate Comparison

6.10 Open Database Management (ODBM) Performance Evaluation

ODBM is a CSL region that manages database connections and access requests from application programs that use the following resource adapters and APIs:

- IMS Universal Database resource adapter
- IMS Universal JDBC driver
- IMS Universal DL/I driver
- Open Database Access interface (ODBA)
- ODBM CSLDMI interface

For the ODBM workload performance measurement, the IMS Universal JDBC driver was used to communicate with ODBM through IMS Connect, using the open standard Distributed Relational Database Architecture (DRDA) as the low-level communication protocol using the Distributed Data Management (DDM) architecture. ODBM translates the DDM into DL/I calls and packages the IMS output as DDM to be returned back to the client. Also, ODBM can run with or without z/OS Resource Recovery Services (RRS). By default, ODBM runs with RRS.

6.10.1 System Configuration

The ODBM Friendly Bank workload with RRS was executed on a z14 configured in a two LPAR environment as shown in Figure 39:

- LPAR 1 hosts IMS, ODBM and IMS Connect with four General Purpose Engines
- LPAR 2 hosts a Java-based workload driver driving 38 clients via TCPIP with 10 General Purpose Engines

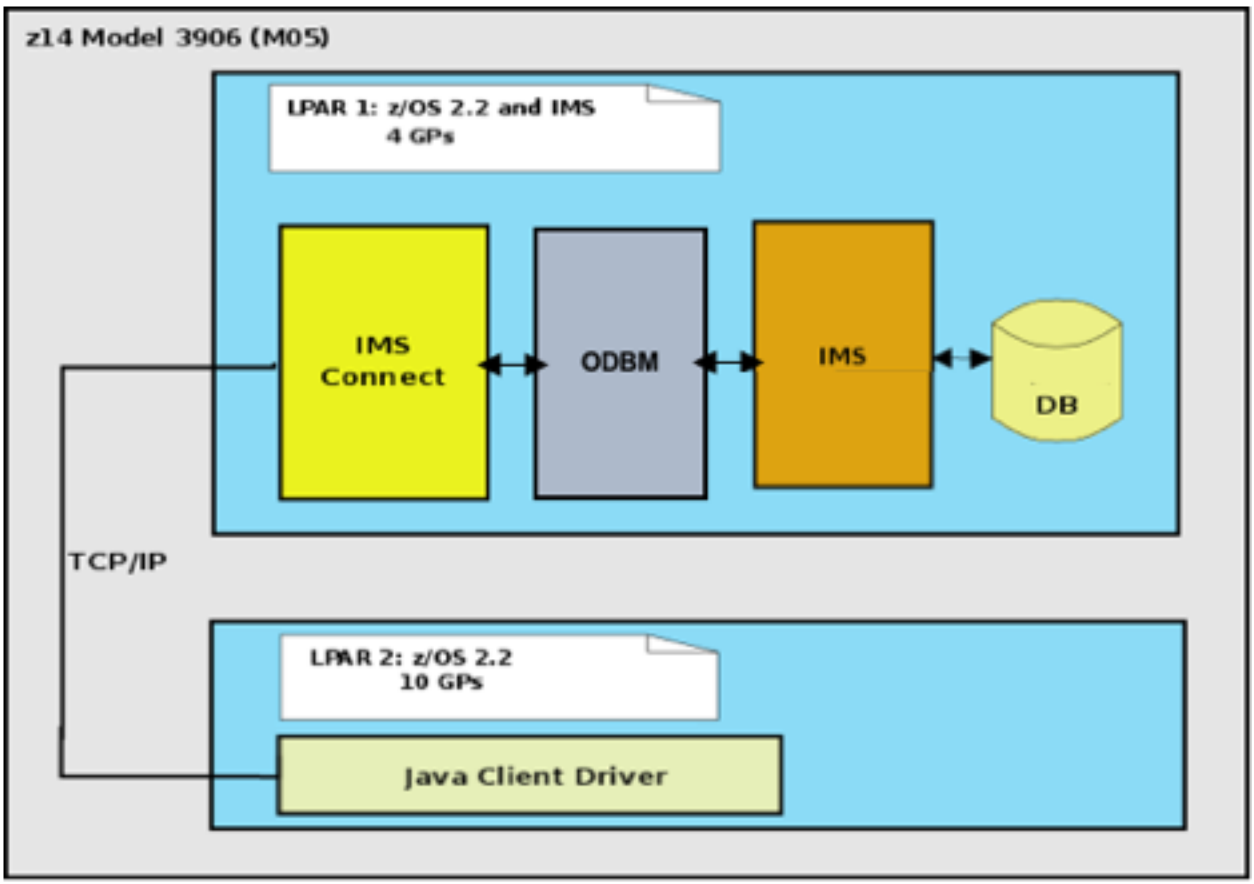

Figure 39: ODBM Processing Configuration

6.10.2 Evaluation Results

The ODBM Friendly Bank workload had a less than 1% difference in ITR for IMS 15 Managed ACB over IMS 14. Table 29 shows the comparison between IMS 13, IMS 14, and IMS 15.

Table 29: ODBM Comparison Results

Open Database Management Processing					
	IMS 13	**IMS 14**	**IMS 15**	**Delta**	**Delta %**
CPU % Busy	80.12%	81.31%	81.85%	0.54%	0.66%
ETR (Tran/Sec)	4469.08	4442.55	4487.49	44.94	1.01%
ITR	5577.98	5463.72	5482.58	18.86	0.35%
Total IMS Response Time (ms)	7.89	7.964	7.863	-0.101	-1.27%
Total General CPU µs/Tran	717.105	732.102	729.584	-2.518	-0.34%
IMS CPU Service Time/Tran (µs)	13.060	12.885	13.943	1.058	8.21%
ICON CPU Service Time/Tran (µs)	228.056	239.277	240.663	1.386	0.58%
ODBM CPU Service Time/Tran (µs)	338.301	339.474	335.073	-4.401	-1.30%
	Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	268	248	252	4	1.61%
Avg. CSA Above 16M Key 7 (M)	17.5	17.7	17.7	0.0	0.00%
	Private Storage IMS Control Region				
Avg. LSQA Private (K)	608	612	612	0	0.00%
Avg. LSQA EPrivate (M)	12.9	12.9	13.5	0.6	4.65%
Avg. USER Private (K)	2640	2268	2272	4	0.18%
Avg. USER EPrivate (M)	54.7	57.2	55.6	-1.6	-2.80%

	Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	384	388	384	-4	-1.03%
Avg. LSQA EPrivate (M)	8.7	8.7	9.1	0.4	5.00%
Avg. USER Private (K)	572	628	628	0	0.00%
Avg. USER EPrivate (M)	263	264	264	0	0.00%
	Private Storage IMS Connect Region				
Avg. LSQA Private (K)	364	364	364	0	0.00%
Avg. LSQA EPrivate (M)	13.5	13.5	13.6	0.1	0.74%
Avg. USER Private (K)	52	56	56	0	0.00%
Avg. USER EPrivate (M)	178	179	181	2	1.12%

A series of scaling tests were run to compare IMS 13, IMS 14 and IMS 15 at various CPU percent busy values. Figure 40 shows the ITR vs Transaction Rate, Figure 41 shows the IMS Response Time vs Transaction Rate, and Figure 42 shows the CPU % Busy vs Transaction Rate.

The ODBM Friendly Bank workload on IMS 15 Managed ACB showed improvements in ITR and Total IMS response time compared to IMS 14.

- ITR improved by 0.35%
- Total IMS response time improved by 1.27%

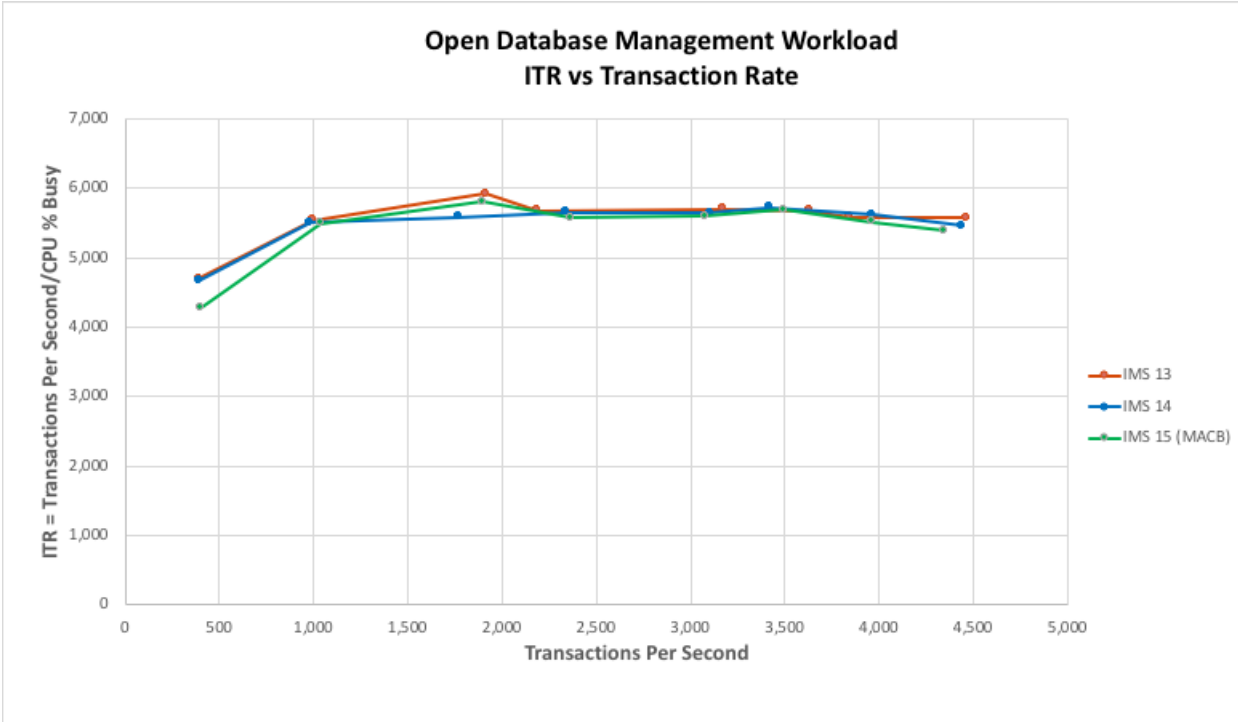

Figure 40: ODBM ITR vs Transaction Rate Comparison Results

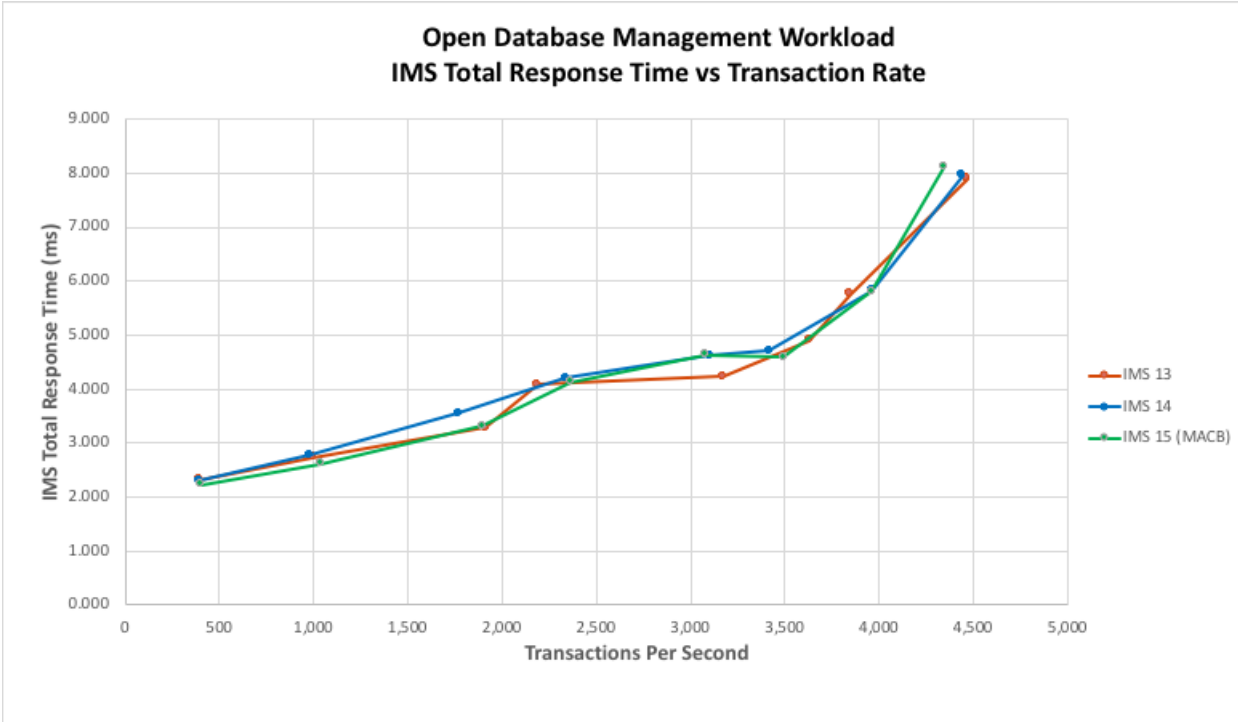

Figure 41: ODBM IMS Total Response Time vs Transaction Rate Comparison Results

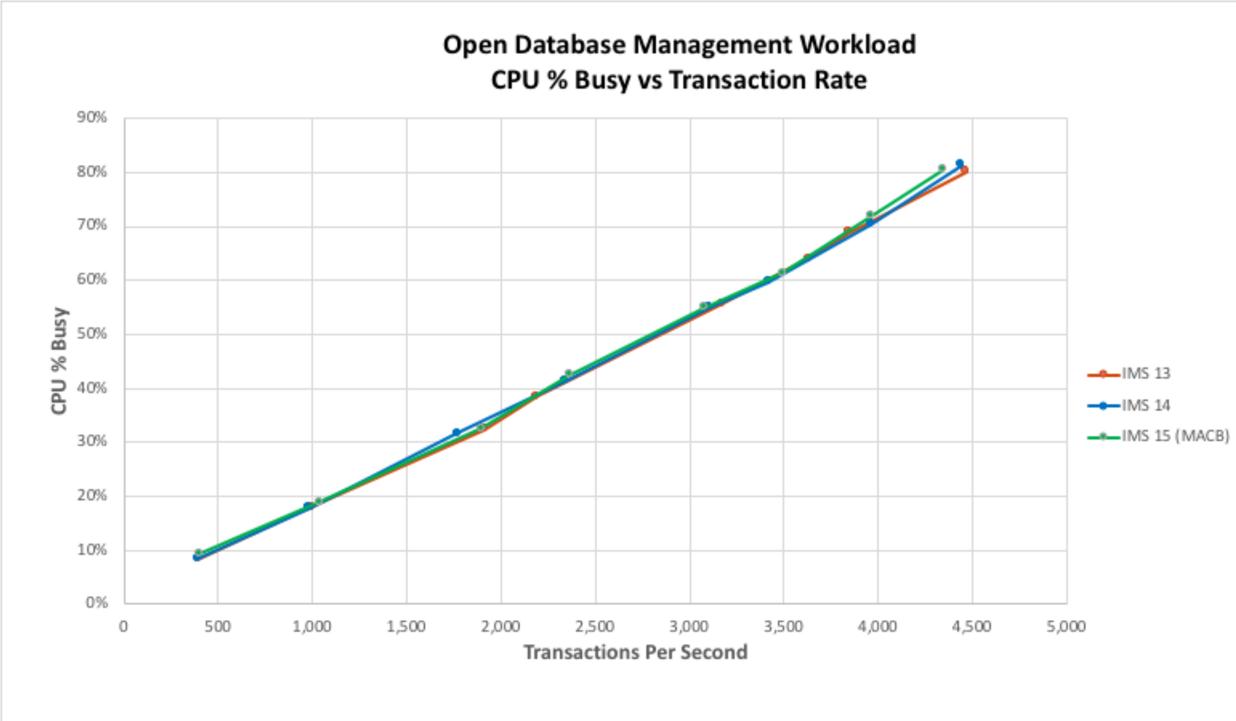

Figure 42: ODBM CPU % Busy vs Transaction Rate Comparison Results

6.11 Base Workload Performance Summary

The results of the base performance workloads in our in-house environment yielded no significant changes in ITR, however a few of the workloads showed improvements with IMS 15 Managed ACB over IMS 14. Figure 43 shows the summary of the base workload ITR delta percentage for each of the base performance workloads.

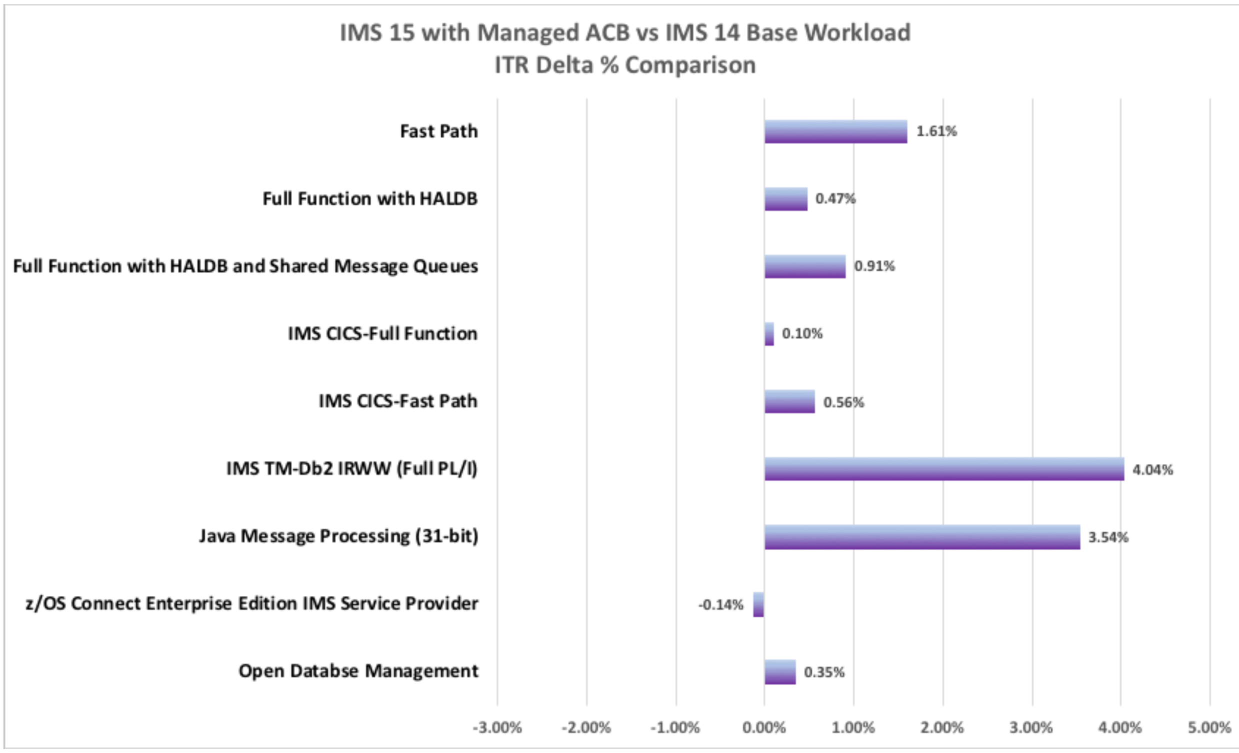

Figure 43: IMS 15 Managed ACB vs IMS 14 Base Workloads ITR Delta % Comparison

7 IMS 15 Enhancements

7.1 Introduction

The latest release of IMS consists of a variety of software enhancements to encourage new application development and modernizing legacy IMS applications as well as taking advantage of new zSystems hardware and operating system improvements and features. Please refer to the IMS 15 release planning guide for the complete list of enhancements in IMS Database Manager, Transaction Manager, and Systems components.

In the following sections we will highlight the performance results of four noteworthy and beneficial IMS 15 enhancements running on the z14 system.

- IMS 15 Managed Application Control Blocks
- IMS 15 Logger Media Manager Support for zHPF and zHyperWrite
- IMS 15 Fast Path Encryption
- IMS 15 Network Security Credential Propagation

7.2 IMS 15 Managed Application Control Blocks

IMS 15 can manage the runtime Application Control Blocks (ACBs) for databases and program views. When IMS manages ACBs, IMS does not use DBD, PSB, nor ACB libraries, and databases and program views can be defined by using SQL DDL statements instead of using generation utilities.

ACBs are the runtime blocks that represent the active databases and program views in online and batch IMS environments. They are created from the databases and program views that are defined to the IMS system.

Most ACBs are pre-built, stored in a data set in binary format, and loaded into memory by the online IMS system or by batch application programs. Some application programs and utilities, such as those that run in an offline DL/I batch region, build ACBs dynamically during run time.

When IMS manages ACBs, as indicated by ACBMGMT=CATALOG in the <SECTION= CATALOG> section of the DFSDF*xxx* member, IMS can build, activate, and load ACBs into memory dynamically when database and program view definitions are submitted to IMS by using SQL DDL statements. Upon receiving the SQL statements, IMS automatically updates the IMS catalog and can, if directed to do so, activate certain changes to database or program view definitions automatically. Changes that are not activated automatically or that are saved in IMS for later activation, can be activated by issuing the IMPORT DEFN SOURCE(CATALOG) command.

Defining databases and program views with SQL DDL statements is an alternative to the process of coding DBD and PSB source, generating DBDs, PSBs, and ACBs with utilities, and performing an Online Change process to activate the resulting ACBs. The DDL statements can be submitted to IMS through a separate product, such as the IMS™ Enterprise Suite Explorer for Development.

When IMS manages ACBs, you can still define databases and program views by using the DBD and PSB generation utilities and build ACBs by using either the ACB Maintenance utility or the ACB Generation and Catalog Populate utility (DFS3UACB).

If you use the DFS3UACB utility, in addition to building the ACBs, the utility can also update the IMS catalog, flag the active resources in the IMS catalog, and activate the ACBs by loading them into the IMS directory, a set of IMS-managed system data sets that are an extension of the IMS catalog. If you do not use the DFS3UACB utility, you can achieve the same results by using the ACB Maintenance utility and the IMS Catalog Populate utility (DFS3PU00).

The IMS Catalog Directory Recovery utility (DFS3RU00) can also be used to rebuild the IMS directory and write the online resources into the IMS directory data sets.

Both the DFS3UACB and the DFS3PU00 utilities require exclusive access when they update the IMS directory directly. So, if the utilities are run in the UPDATE mode, the IMS systems that use the IMS catalog that is being updated must be shut down. To avoid shutting down the IMS systems, run the utilities in STAGE mode and add them to the IMS directory by using the IMPORT DEFN SOURCE(CATALOG) command.

In many cases, whether you use DDL or the generation utilities, you must finalize the activation of the ACBs by issuing the IMPORT DEFN SOURCE(CATALOG) command. The Online Change process is not required. In other cases, such as when you delete a resource, the changes to the ACBs are activated automatically.

When you enable IMS management of ACBs, IMS performance is similar to what it is when an ACB library is used. IMS uses the same amount of I/O to access IMS-managed ACBs as it does to access ACBs in an ACB library.

Before IMS can manage ACBs, IMS and the IMS catalog must be set up to support ACB management by running the DFS3PU00 utility with the MANAGEDACBS=SETUP control statement that you provide as input. After IMS is set up for ACB management, you specify ACBMGMT=CATALOG in the <SECTION=CATALOG> section of the DFSDF*xxx* PROCLIB member.

7.2.1 System Configuration

The environment and system configurations for the IMS 15 Managed ACBs workloads are the same as those workloads described in Section 6.

7.2.2 Evaluation Results

The workloads described in Section 6 were run in both IMS 15 Managed ACB and Non-Managed ACB environments. The result comparison between the IMS 15 Managed ACB and Non-Managed ACB showed that the ITR had a mix of improvements and degradations. However, the degradations were less than 1% while the improvements were as high as 2.5%. A chart of the ITR delta % between IMS 15 Managed ACB and Non-Managed ACB for each workload is shown below in Figure 44.

The LSQA storage above 16M saw an increase of up to 1.0 MB for all workloads in the IMS control and DL/I regions under the IMS 15 Managed ACB environment. The increase is the result of accessing and utilizing the catalog information in building the Managed ACB libraries. The size of storage increase will vary depending on the size of the IMS DB Catalog. Figure 45 shows the delta in megabytes and the delta percentage increase of the LSQA storage for each workload.

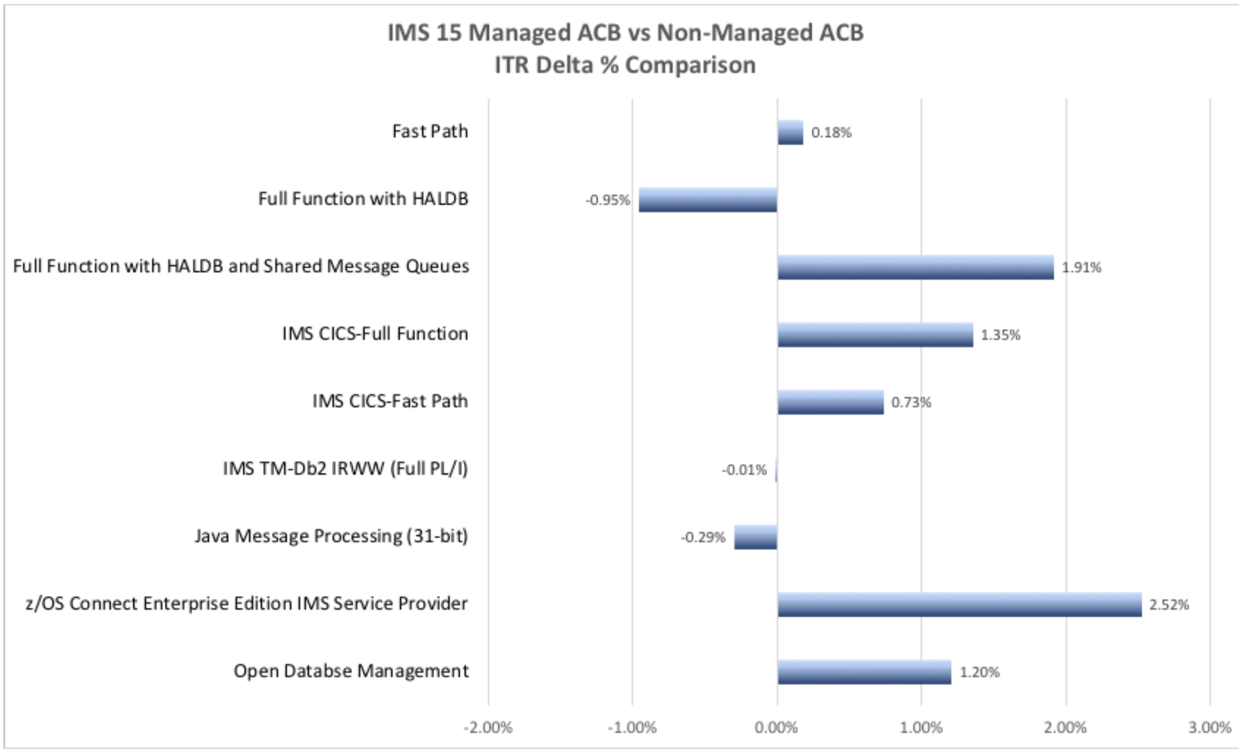

Figure 44: IMS 15 Managed ACB vs Non-Managed ACB ITR Delta % Comparison

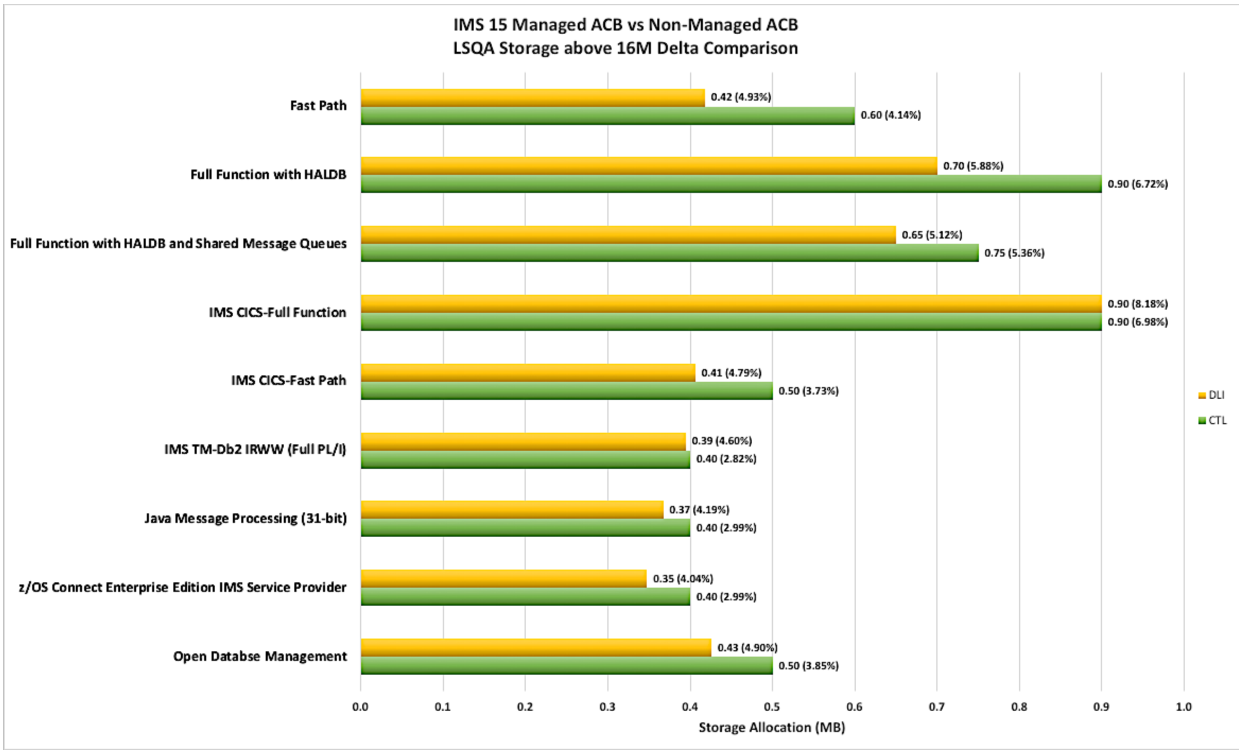

Figure 45: IMS 15 Managed ACB vs Non-Managed ACB LSQA Storage Above 16M Comparison

7.3 IMS 15 Logger Media Manager Support for High Performance FICON® (zHPF) and zHyperWrite

IMS 15 enhances the IMS Logger function to use the Data Facility Storage Management Subsystem (DFSMS) Media Manager for the IMS Write-Ahead Log Data Set (WADS). The DFSMS Media Manager exploits new hardware features like zHPF and zHyperWrite. Additionally, IMS 15 uses the new function added in z/OS APAR OA51385 to allow the IMS Online Log Data Set (OLDS) to also exploit zHyperWrite.

The zHPF feature provides improved channel and control unit efficiency, lower latency, and improved reliability/availability as compared to FICON which increases I/O throughput. The zHyperWrite feature is able to reduce the synchronous replication delays that occur when using Metro Mirror™ technology for continuous availability and disaster recovery reducing latency time for synchronous replication products from direct-access storage devices (DASD) in VSAM files. The Media Manager is a no charge component of DFSMS.

7.3.1 System Configuration

In IMS 15, the logger configuration has been changed. The WADS must be defined as a linear VSAM data set with a control interval (CI) size of 4 KB (4096-bytes), secondary space allocation of 0, and include the SHAREOPTIONS(3 3) parameter. Also, the logger parameter definitions have been moved from the DFSVSMxx to the DFSDFxxx PROCLIB member and supports a new optional parameter to enable zHyperWrite for WADS and OLDS. For more information, refer to the IMS 15 Release Planning Guide and the IMS 15 IMS System Administration Guide.

The Full Function with HALDB workload was used to evaluate the performance of the IMS Logger Media Manager. The workload was executed on a z14 Model 3906 (M04) hardware system located at the IBM Poughkeepsie Laboratory running with z/OS 2 Release 2, IBM® DS8880 Hybrid Storage system, and Metro Mirror. The IBM Poughkeepsie z14 is not an isolated environment and not all the data metrics were available, but the key ones are provided.

Two LPARs within a sysplex were configured with General Purpose Engines to conduct the Logger Media Manager enhancement evaluation:

- LPAR 1 hosts IMS and 130 MPP regions using Full Function database and Metro Mirror™ with four General Purpose Engines
- LPAR 2 hosts TPNS driving over 1,000 transactions per second.

7.3.2 Evaluation Results

IMS 15 showed a significant reduction in the WADS and OLDS response times when the Full Function workload was executed on the IBM Poughkeepsie system with a Metro Mirror distance of zero kilometer as shown in Table 30 and graphed in Figure 46. Note: The measurement data used as the base for comparison had both zHPF and zHyperWrite disabled.

- When only zHPF was enabled there was more than a 10% improvement in both WADS and OLDS response times.
- When only zHyperWrite was enabled, there was a 68% and 64% improvement in WADS and OLDS response times respectively.
- When both zHPF and zHyperWrite were enabled, there was more than a 75% improvement in WADS and OLDS response times.

Table 30: IMS 15 Logger Manager with zHPF and zHyperWrite at 0 km

IMS 15 Logger Media Manager with zHPF and zHyperWrite at 0 km

	Avg Base	Avg zHPF enabled	Delta % from Base	Avg zHyperWrite enabled	Delta % from Base	Avg zHPF zHyperWrite enabled	Delta % from Base
CPU % Busy	49.59%	50.09%	1.00%	49.70%	0.22%	48.90%	-1.39%
ETR (Tran/Sec)	985.14	997.87	1.29%	980.85	-0.44%	966.76	-1.87%
ITR	1986.40	1992.46	0.30%	1973.44	-0.65%	1976.72	-0.49%
Total IMS Response Time (ms)	1012.54	1024.97	1.23%	1006.07	-0.64%	995.91	-1.64%
WADS Response Time	1.167	1.047	-10.28%	0.371	-68.21%	0.254	-78.23%
OLDS Response Time	1.110	0.991	-10.72%	0.399	-64.05%	0.274	-75.32%

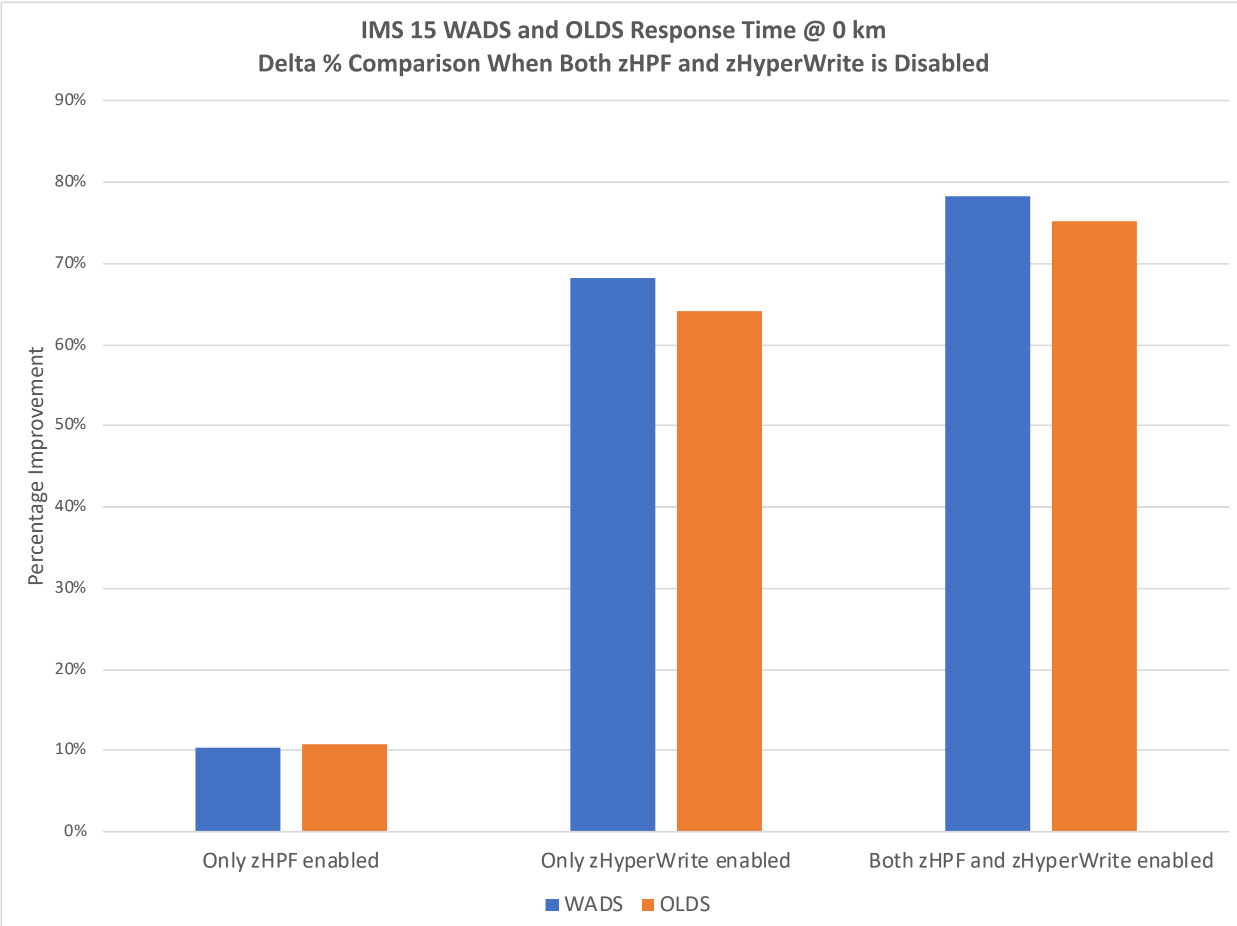

Figure 46: IMS 15 WADS and OLDS Response Time at 0km

When the same workload was executed with the Metro Mirror™ configuration at a distance of 100 km, we saw about the same improvement percentage in the WADS and OLDS response times as shown in Table 31 and graphed in Figure 47.

- When only zHPF was enabled, there was more than a 2% and 16% improvement in the WADS and OLDS response times respectively.
- When only zHyperWrite was enabled, there was more than a 75% improvement in both WADS and OLDS response times.
- When both zHPF and zHyperWrite were enabled, there was more than 80% improvement in WADS and OLDS response times

Table 31: IMS 15 Logger Manager with zHPF and zHyperWrite at 100 km

IMS 15 Logger Media Manager with zHPF and zHyperWrite at 100 km							
	Avg Base	Avg zHPF enabled	Delta % from Base	Avg zHyperWrite enabled	Delta % from Base	Avg zHPF zHyperWrite enabled	Delta % from Base
CPU % Busy	45.18%	47.62%	5.40%	45.65%	1.04%	45.77%	1.31%
ETR (Tran/Sec)	934.92	969.14	3.66%	1062.66	13.66%	970.88	3.85%
ITR	2069.55	2035.85	-1.63%	2327.86	12.48%	2121.48	2.51%
Total IMS Response Time (ms)	961.27	993.23	3.32%	1002.77	4.32%	997.23	3.74%
WADS Response Time	1.677	1.633	-2.62%	0.392	-76.62%	0.273	-83.72%
OLDS Response Time	1.727	1.443	-16.44%	0.399	-76.89%	0.280	-83.79%

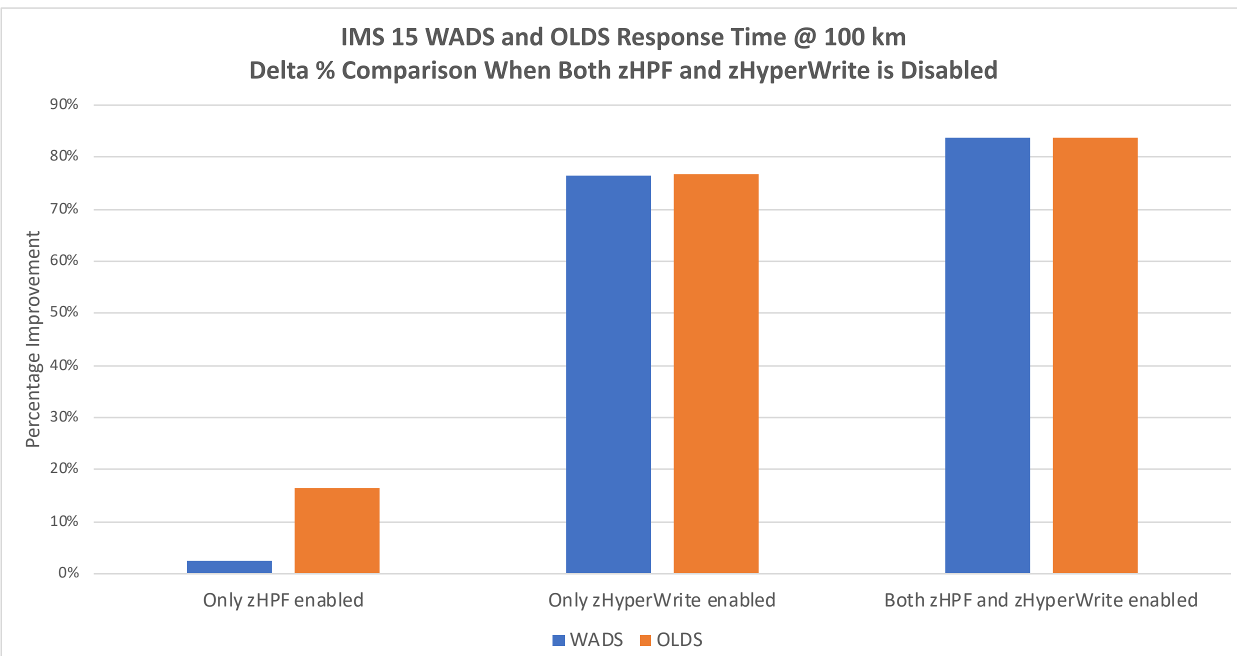

Figure 47: IMS 15 WADS and OLDS Response Time at 100km

7.4 IMS 15 Fast Path Encryption

Data sets that are accessed by DFSMS access methods (BSAM, QSAM, VSAM) are eligible for data set encryption. Encrypted data sets must be SMS-managed extended format data sets. All supported releases of IMS (IMS 13, IMS 14, and IMS 15 at the time of writing this paper) support encryption of IMS data sets accessed with standard access methods (for example, VSAM databases, image copies, the IMS OLDS, and SLDS). In addition, IMS supports encrypting the logger WADS and Fast Path DEDB area data sets with IMS 15.

This section describes the performance of the IMS 15 Fast Path workload with and without the following data sets encrypted.

- Write Ahead Data Set (WADS) via APAR PI84947
- Online Log Data Set (OLDS)
- Fast Path DEDB VSAM via APAR PI83756 (Note: DEDB encryption requires Fast Path 64-bit buffers. This is specified by FPBP64=Y in the FASTPATH section of the DFSDFxxx member in the IMS PROCLIB data set)

Running the Fast Path workload with encryption on z14 will take advantage of the enhanced performance of the on-chip cryptographic coprocessors as well as the new Crypto Express card enabling pervasive encryption without the need for any application changes.

7.4.1 System Configuration

The Fast Path workload, encrypted with an AES-256 bit key associated with a 64-byte key label, was used for testing the encryption capability on the z14 system configured in a two LPAR environment as shown in Figure 48.

- LPAR 1 hosts IMS, 375 IFP regions, and five IMS Connects with seven General Purpose Engines
- LPAR 2 hosts TPNS driving a total of 8,000 IMS Connect clients via TCP/IP with 10 General Purpose Engines

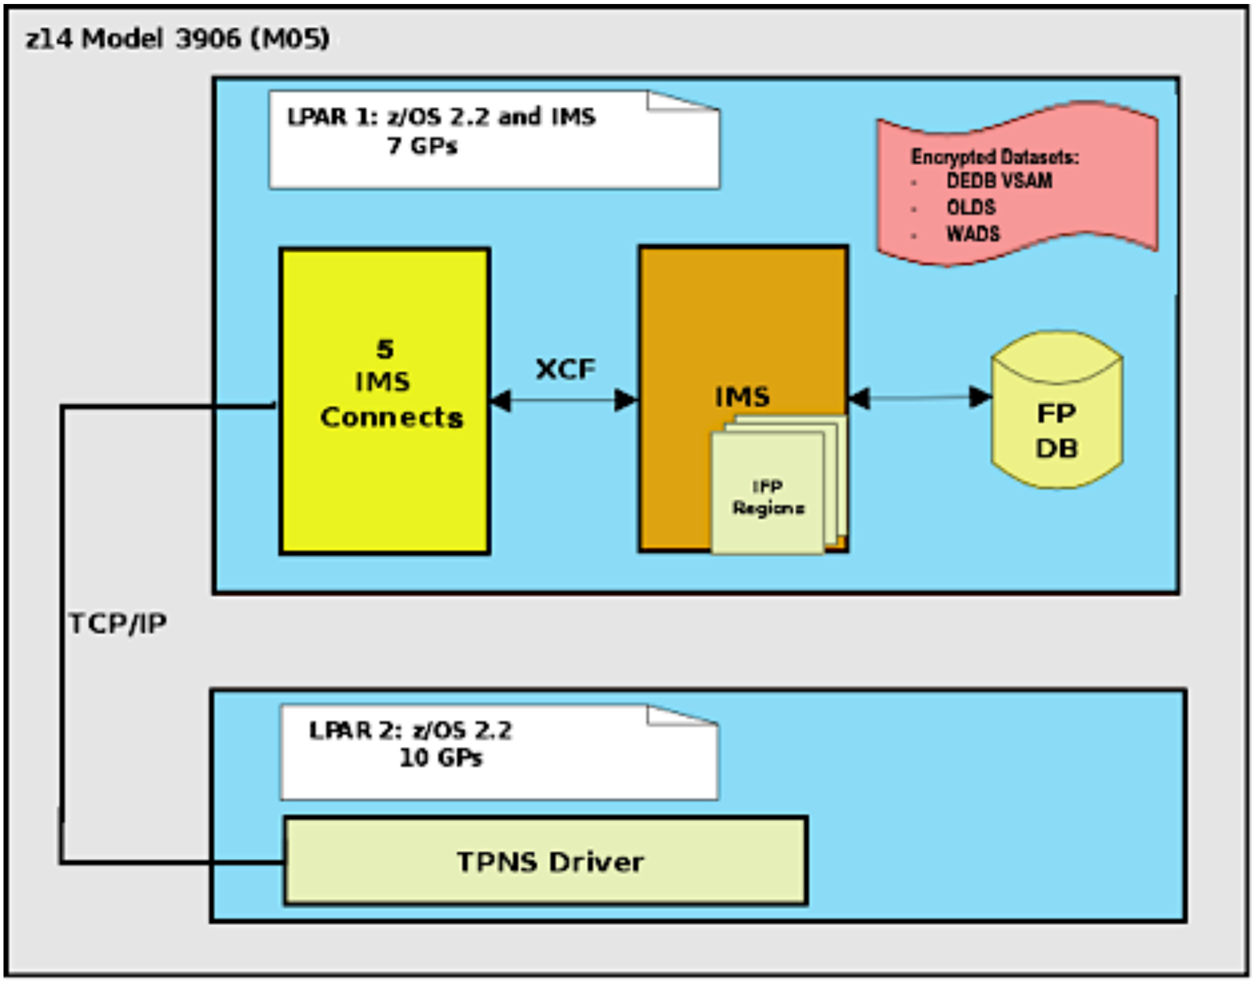

Figure 48: IMS 15 Fast Path Workload Encryption Configuration

7.4.2 Evaluation Results

The performance evaluation showed the ITR degraded by about 3.6% when encrypting all of Fast Path DEDBs, OLDs, and WADs running on the z14 system. The additional total general CPU per transaction to encrypt on z14 was approximately 3.7% as shown in Table 32 below.

Table 32: IMS 15 Fast Path Workload Encryption Results

Fast Path with WADS Encryption				
	IMS 15 (w/o encryption)	IMS 15 (with encryption)	Delta	Delta %

CPU % Busy	81.50%	84.70%	3.20%	3.93%
ETR (Tran/Sec)	98733.87	98894.04	160.17	0.16%
ITR	121145.85	116758.02	-4387.83	-3.62%
Total IMS Response Time (ms)	0.642	0.654	0.012	1.87%
Total General CPU µs/Tran	57.782	59.953	2.171	3.76%
IMS CPU Service Time/Tran (µs)	11.196	12.340	1.144	10.22%
ICON CPU Service Time/Tran (µs)	17.464	18.491	1.027	5.88%
IFP CPU Service Time/Tran (µs)	18.024	18.904	0.880	4.88%
Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	368	368	0	0.00%
Avg. CSA Above 16M Key 7 (M)	715	716	1	0.14%
Private Storage IMS Control Region				
Avg. LSQA Private (K)	1280	1336	56	4.38%
Avg. LSQA EPrivate (M)	15.1	15.2	0.1	0.66%
Avg. USER Private (K)	2272	2272	0	0.00%
Avg. USER EPrivate (M)	96.2	87.6	-8.6	-8.94%
Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	352	348	-4	-1.14%
Avg. LSQA EPrivate (M)	8.89	8.89	0.00	-0.04%
Avg. USER Private (K)	628	628	0	0.00%
Avg. USER EPrivate (M)	265	265	0	0.00%
Private Storage IMS Connect Region				

Avg. LSQA Private (K)	360.8	360.8	0.0	0.00%
Avg. LSQA EPrivate (M)	15.44	15.12	-0.32	-2.07%
Avg. USER Private (K)	56	56	0	0.00%
Avg. USER EPrivate (M)	206	205.4	-0.6	-0.29%

A series of scaling tests were run to compare IMS 15 with and without encryption at various CPU percent busy values. Figure 49 shows the ITR vs Transaction Rate, Figure 50 shows the IMS Response Time vs Transaction Rate, and Figure 51 shows the CPU % Busy vs Transaction Rate.

The Fast Path encryption workload on IMS 15 showed degradation in ITR and Total IMS response time between encryption and non-encryption.

- ITR degraded by 3.62%
- Total IMS response time degraded by 1.87%

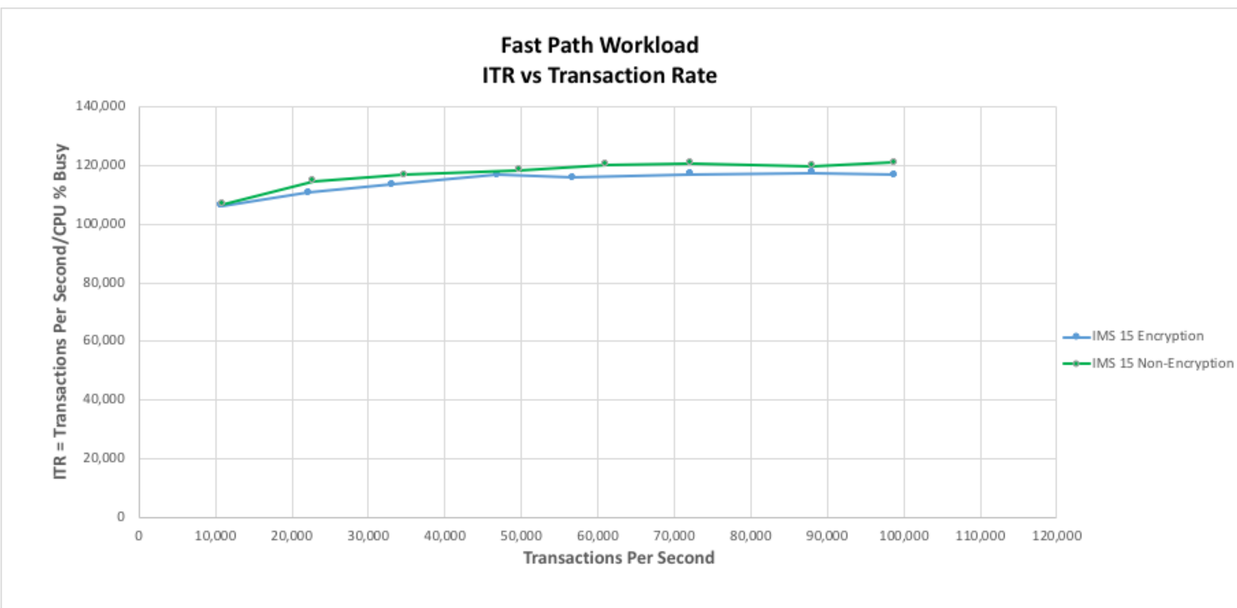

Figure 49: FP Workload Encryption ITR vs Transaction Response Rate Comparison

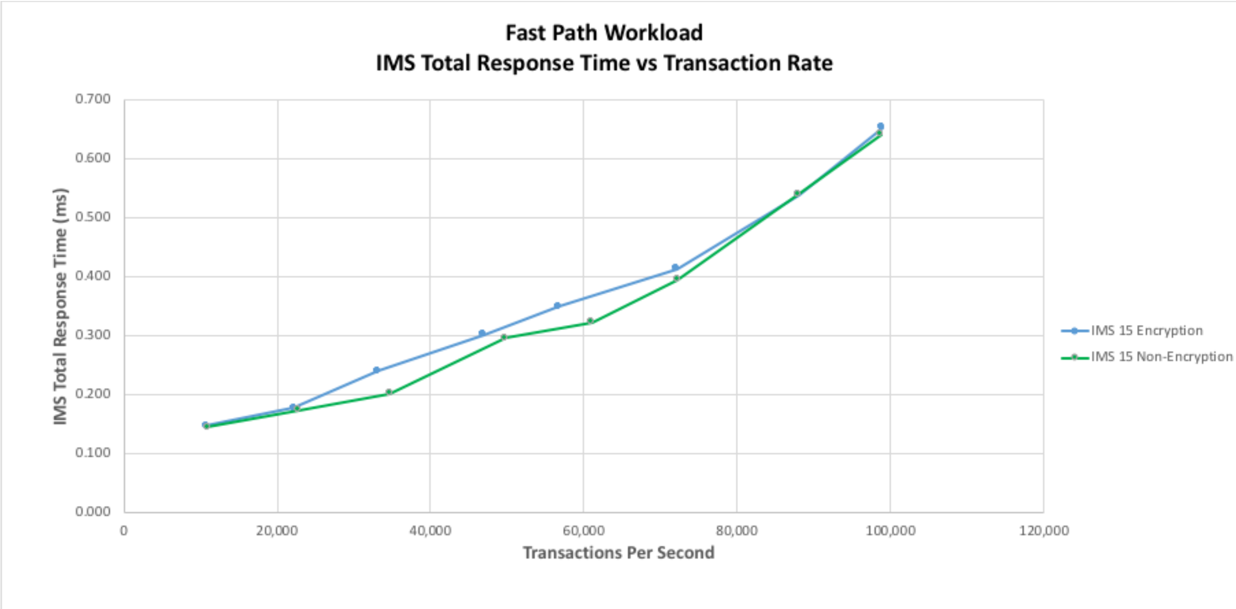

Figure 50: FP Workload Encryption IMS Total Response Time vs Transaction Response Rate Comparison

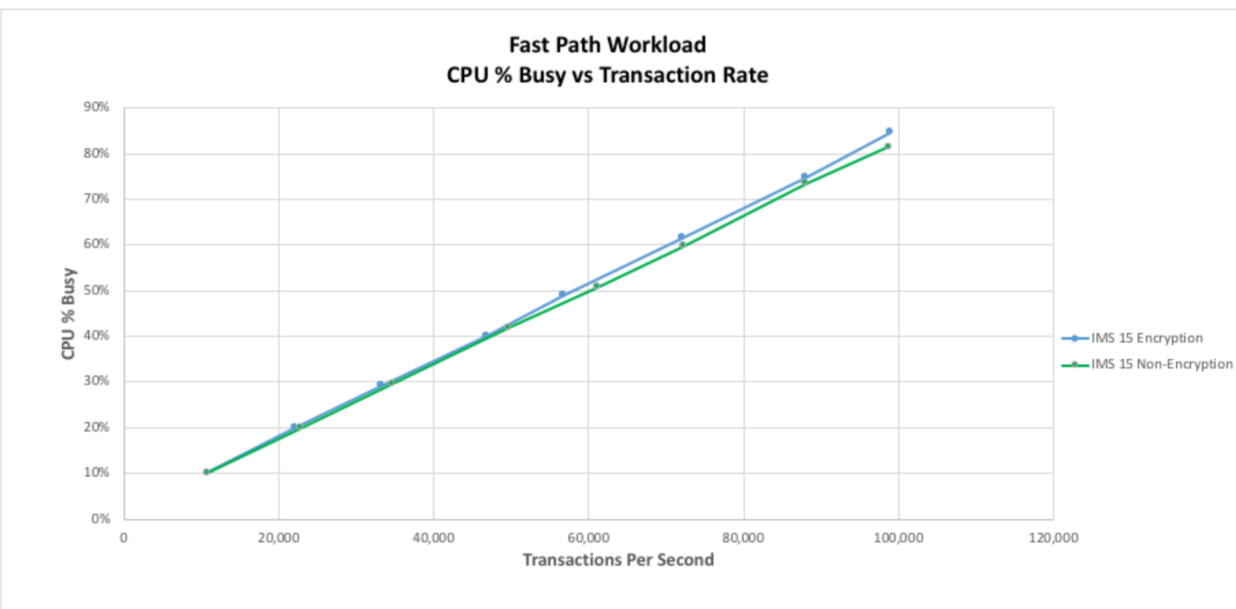

Figure 51: FP Workload Encryption CPU % vs Transaction Response Rate Comparison

7.5 IMS 15 Network Security Credential Propagation

IMS 15 provides an expanded auditable and accountable enterprise security environment by allowing the distributed network security credentials to be passed and logged into IMS for all inbound and outbound messages.

The distributed network security credentials can include a network user ID and a network session ID. These credentials can be included in the RACF SMF records by specifying LOGSTR=YES in the OTMA client descriptor.

The main objective for the network security propagation performance test is to determine the overhead, if any, in propagating up to 500 bytes of distributed network credentials (session ID and user ID) in IMS Connect and IMS log records when compared to no ID propagation.

7.5.1 System Configuration

The system configuration for the network security credential propagation performance testing was the same as the z/OS Connect EE IMS SP described in Section 6.9.

- LPAR 1 hosts IMS, 350 MPP regions, and IMS Connect with eight General Purpose Engines
- LPAR 2 hosts a Java-based workload driver driving 50 clients via TCPIP and 10 General Purpose Engines.

The following size of network security credentials were used within a 1K message:

- Base test with no security credential
- 25 bytes of security credential
- 50 bytes of security credential
- 100 bytes of security credential
- 200 bytes of security credential
- 250 bytes of security credential
- 500 bytes of security credential

7.5.2 Evaluation Results

The performance data for the network security propagation shows ITR decreasing as the size of the network security credential increases. Table 33 shows a less than 1% degradation for a credential size of 25 bytes and about a 10% degradation for 500 bytes when compared to zero bytes of network credential size for ITR. Figure 52 shows a graphical representation of the ITR as the sizes of network credential increases.

Table 33: Network Security Credential Propagation Performance ITR Result

Network Security Credential Propagation			
	ITR	**Delta vs 0 Bytes**	**Delta % vs 0 Bytes**
0 Bytes	25281.05	N/A	N/A
25 Bytes	25093.50	-187.55	-0.74%
50 Bytes	25023.66	-257.39	-1.02%
100 Bytes	24631.52	-649.53	-2.57%
200 Bytes	24080.92	-1200.13	-4.75%
250 Bytes	23956.50	-1324.55	-5.24%
500 Bytes	22635.28	-2645.77	-10.47%

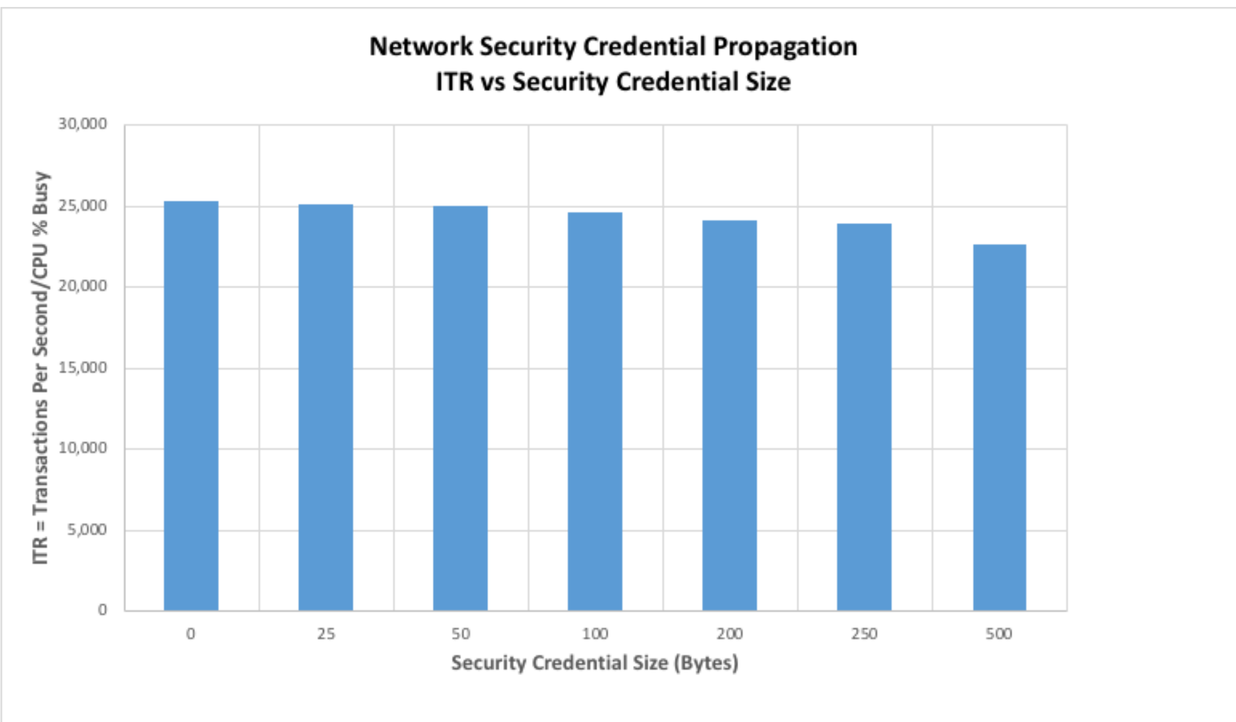

Figure 52: Network Security Credential Propagation ITR vs Security Credential Size Comparison

Table 34 and Figure 53 shows the total service time increasing as the size of the network credentials increase. The total service time increased less than 1% for 25 and 50 bytes and about 10% for 500 bytes.

Table 34: Network Security Credential Propagation Performance Total Service Time Result

Network Security Credential Propagation			
	Total Service Time/Tran (µsec)	**Delta vs 0 Bytes**	**Delta % vs 0 Bytes**
0 Bytes	297.832	N/A	N/A
25 Bytes	299.800	1.97	0.66%
50 Bytes	299.732	1.90	0.64%
100 Bytes	303.762	5.93	1.99%
200 Bytes	310.542	12.71	4.27%
250 Bytes	311.506	13.67	4.59%
500 Bytes	327.518	29.69	9.97%

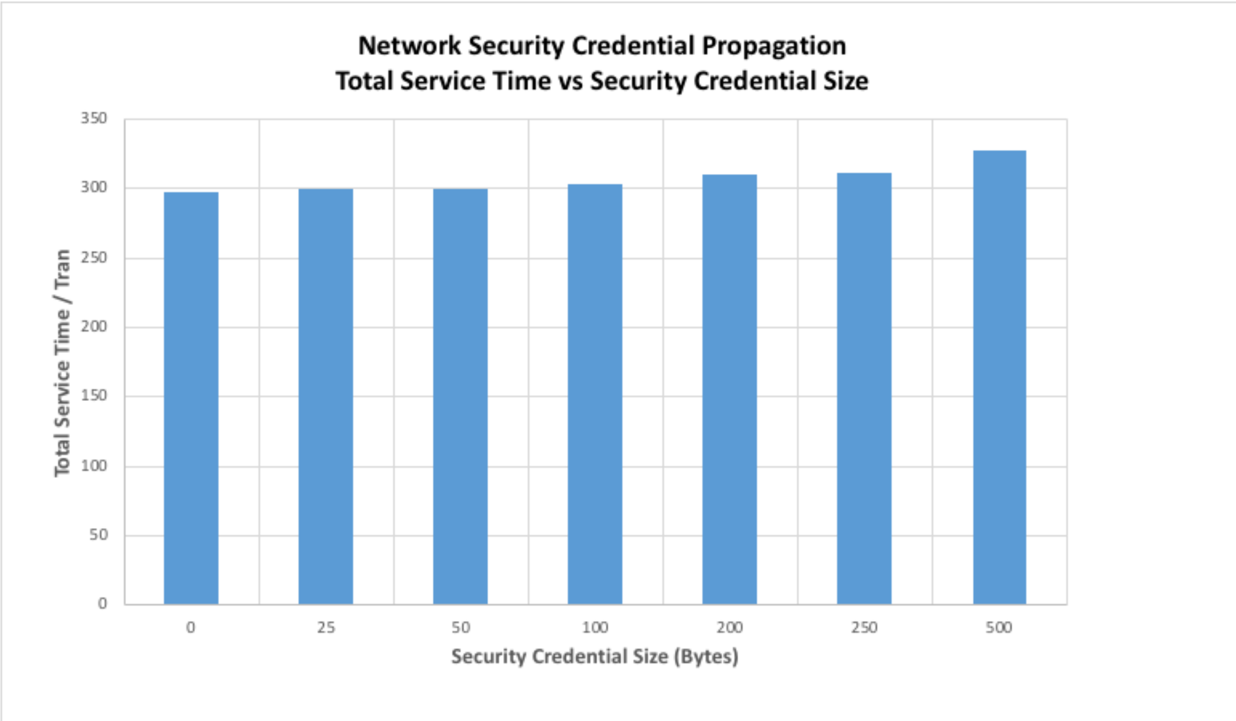

Figure 53: Network Security Credential Propagation Total Service Time vs Security Credential Size Comparison

8 IMS 14 Enhancements

8.1 Introduction

The release of IMS 14 contained a variety of new features and enhancements as described in the IMS Version 14 Release Planning guide. This section presents the performance results of three specific enhancements from IMS 14. These enhancements are:

- IMS 14 Open Transaction Manager Access (OTMA) Resume TPIPE Parallelism
- IMS 14 Java Message Processing (JMP) 64-bit Support
- IMS 14 External Subsystem Attach Facility (ESAF) Connection Pooling

The IMS 14 enhancements were evaluated on the z14 system running in an IMS 15 Managed ACB configured environment.

8.2 IMS 14 Open Transaction Manager Access (OTMA) Resume TPIPE Parallelism

In IMS 14, OTMA TPIPE was enhanced to allow multiple Resume TPIPEs to be concurrently active for parallel message processing to multiple clients. Previously, OTMA allowed only a single Resume TPIPE to be active for a given client. The parallelism of Resume TPIPE alleviates the restrictions and bottlenecks when processing messages at high rates.

The OTMA Resume TPIPE parallelism is activated by specifying the parameter MULTIRTP=Y in the IMS Connect data store definition or the OTMA client descriptor in the DFSYDTx member of the IMS proclib. Specifying a non-zero LIMITRTP value in the OTMA descriptor also enables the Resume TPIPE parallelism. The LIMITRTP= parameter controls the number of Resume TPIPEs that can be active for a minimum value of 10 and a maximum value of 4095. When the Resume TPIPEs reaches the LIMITRTP value, the incoming messages will be queued for the next available TPIPE in a FIFO order.

The objective of the OTMA Resume TPIPE Parallelism performance test is to validate the performance improvement of the resume TPIPE parallelism enhancement and study the impact of the LIMITRTP parameter in a synchronous callout environment.

8.2.1 System Configuration

The OTMA ICAL TPIPE evaluation was executed on z14 in a two LPAR configuration as shown in Figure 54:

- LPAR 1 hosts IMS, a single MPP region, and IMS Connect with five General Purpose Engines
- LPAR 2 hosts TPNS driving a total of 1,000 clients via SNA with 10 General Purpose Engines

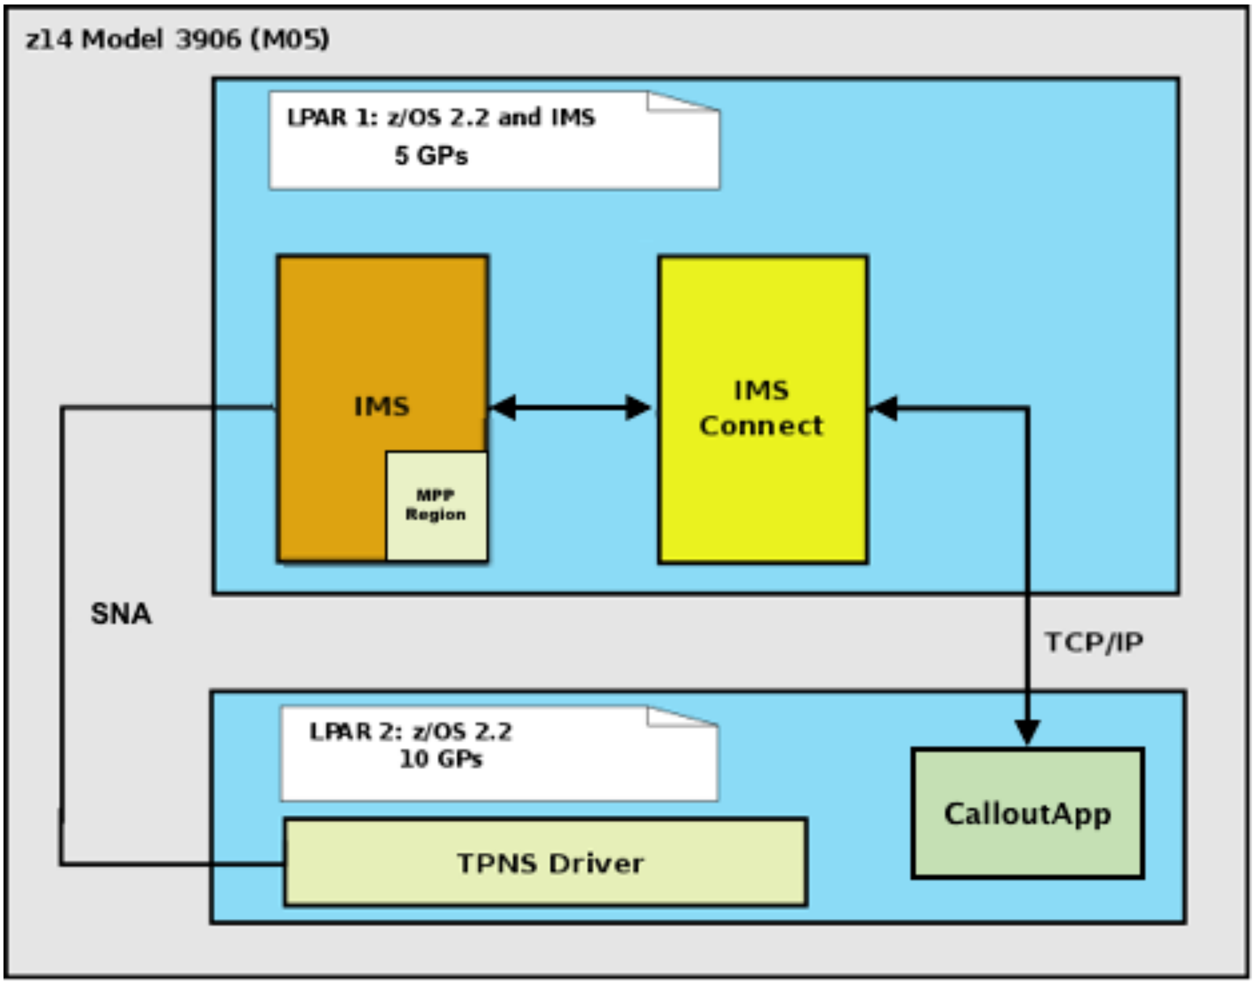

Figure 54: OTMA Resume TPIPE Parallelism Configuration

The figure above shows the TPNS driver invoking the OTMA callout program in IMS. The callout program then sends an 8,000-byte message request through the OTMA TPIPE and IMS Connect to the Java web service CalloutApp. The CalloutApp processes the input message and echoes the 8,000-byte message back to IMS OTMA.

8.2.2 Evaluation Results

The OTMA Resume TPIPE Parallelism performance evaluation focused on the IMS Call (ICAL) response time and throughput. The performance results in Table 35 shows that the response time decreased significantly when TPIPE parallelism was enabled versus non-parallelism for each of the TPIPE values. The transaction rate also increased when TPIPE parallelism was enabled allowing more throughput.

Figure 55 shows the OTMA Resume TPIPE Parallelism ICAL Response Time vs LIMITRTP value, Figure 56 shows the OTMA Resume TPIPE Parallelism ICAL Transactions per Second vs LIMITRTP value, and Figure 57 shows the OTMA Resume TPIPE Parallelism ITR vs LIMITRTP value.

Table 35: OTMA Resume TPIPE Parallelism Evaluation Results

OTMA Resume TPIPE Parallelism Evaluation			
	ICAL Response Time (ms)		
	1 TPIPE	5 TPIPE	10 TPIPE
Non-Parallel	109.75	40.539	24.625
LIMITRTP	1 TPIPE	5 TPIPE	10 TPIPE
10	23.167	11.936	13.838
25	13.198	12.394	14.418
50	12.414	13.648	15.358
75	12.606	13.569	16.054
100	12.885	14.390	15.987
	ICAL Transaction Rate per Second		
Non-Parallel	2730.49	7383.69	12150.10
LIMITRTP	1 TPIPE	5 TPIPE	10 TPIPE
10	12918.48	24641.25	21333.58
25	22547.84	23635.06	20515.42
50	23539.59	21448.80	19186.84
75	23020.65	21584.50	18364.71

100	22571.32	20433.16	18503.89
	ITR		
Non-Parallel	27655.97	32030.96	30562.99
LIMITRTP	**1 TPIPE**	**5 TPIPE**	**10 TPIPE**
10	30999.95	29438.41	28726.73
25	28669.34	29278.44	27397.84
50	29382.03	27926.99	26851.56
75	29101.86	28262.86	26107.72
100	29164.57	27179.71	26160.67

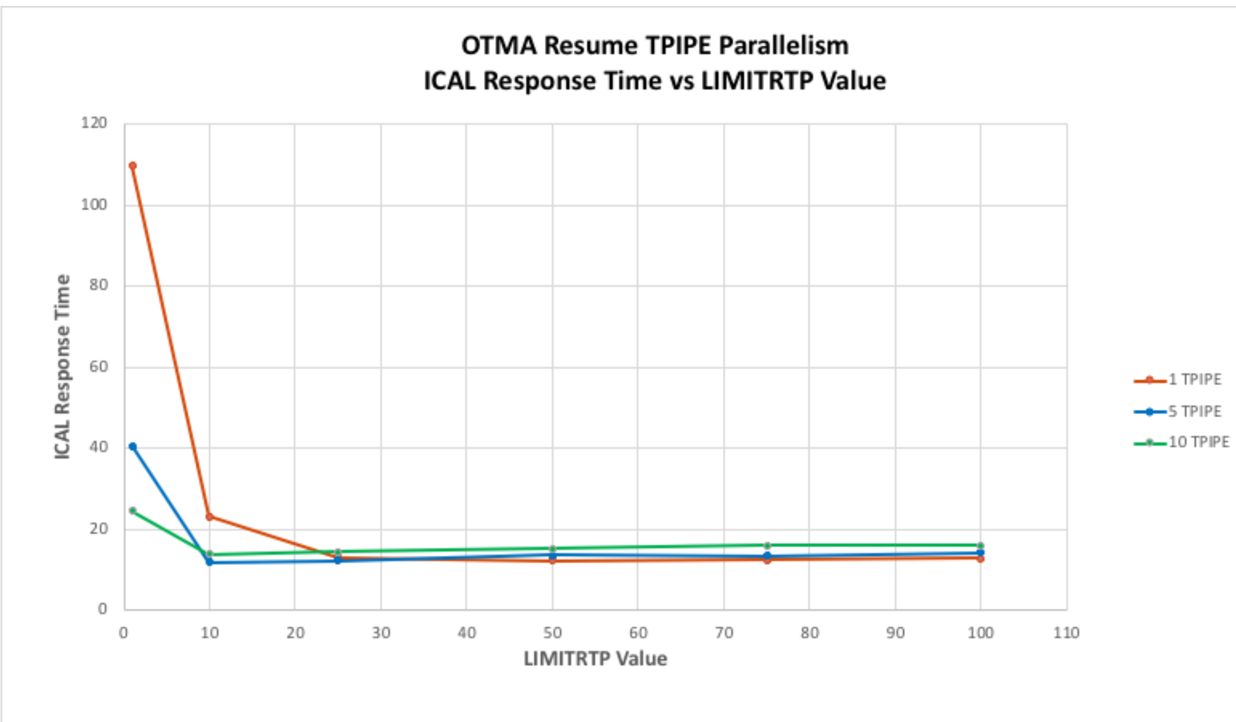

Figure 55: OTMA Resume TPIPE Parallelism ICAL Response Time vs LIMITRTP Value

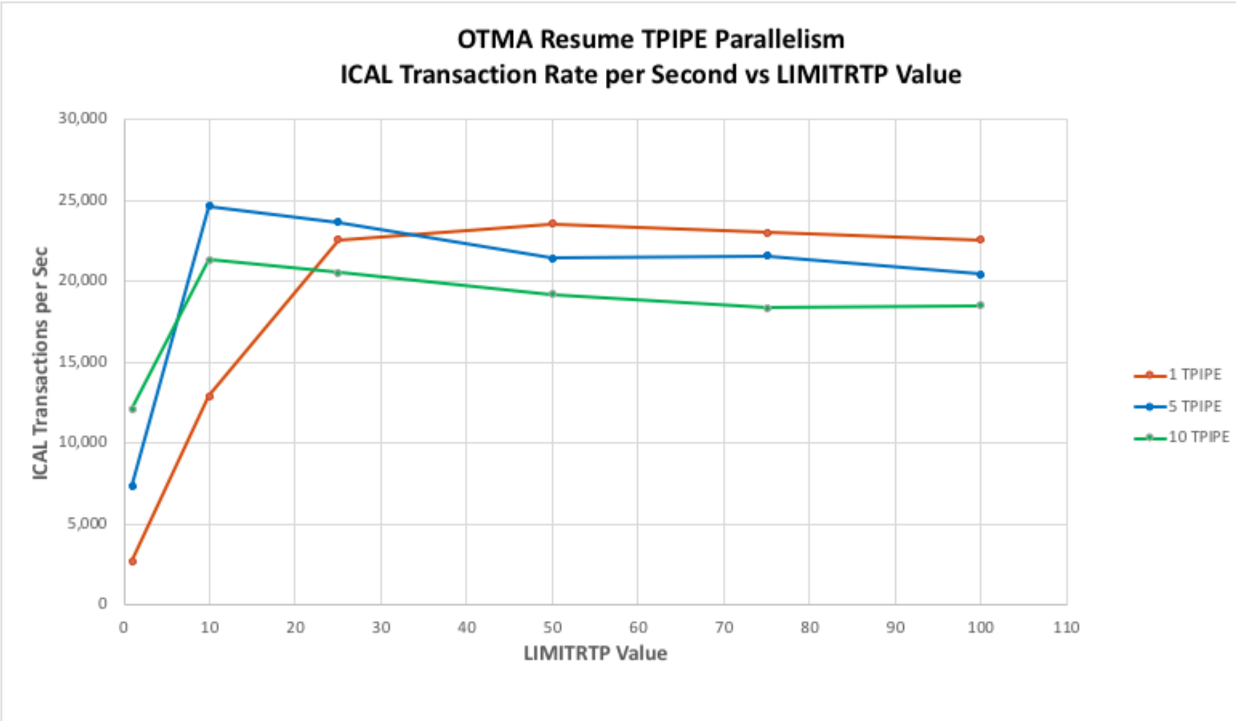

Figure 56: OTMA Resume TPIPE Parallelism ICAL Transactions per Second vs LIMITRTP Value

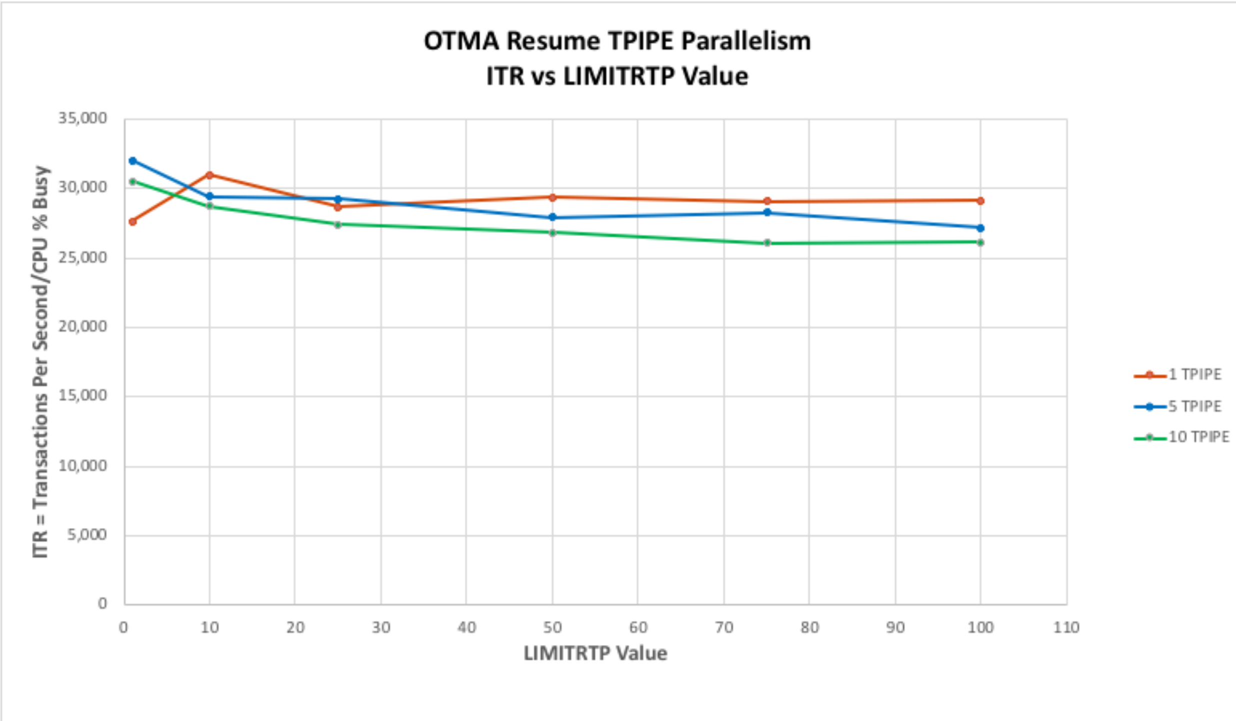

Figure 57: OTMA Resume TPIPE Parallelism ITR vs LIMITRTP Value

8.3 IMS 14 Java Message Processing (JMP) 64-bit Support

In IMS 14, the Java Message Processing (JMP) and Java Batch Processing (JBP) dependent regions were modernized to have the option of running with 64-bit Java Virtual Machines (JVMs). Prior to IMS 14, the JMP and JBP dependent regions only supported 31-bit JVMs. In order for the Java dependent regions to use 64-bit JVMs, a “JVM=64” parameter must be specified on the DFSJBP and/or DFSJMP procedure EXEC statement otherwise 31-bit mode is used by default.

Some attention is required when converting the JVM from a 31-bit to a 64-bit mode. In the 31-bit mode, IMS will perform an implicit check for Java exceptions before executing a SYNC call but not in the 64-bit mode. Also, the Language Environment supports C, C++, and assembly language interoperability in the 64-bit addressing mode but does not support COBOL or PL/I interoperability and may cause system or user abends if the Java application invokes COBOL or PL/I in the 64-bit mode. For more information, refer to the online IMS documentation.

8.3.1 System Configuration

The JMP 64-bit evaluation workload was executed on the z14 in a two LPAR configuration as shown in Figure 31 in Section 6.8.1:

- LPAR 1 hosts IMS, IMS Connect, and 200 JMP regions with six General Purpose Engines
- LPAR 2 hosts a Java-based workload driver driving 38 clients via TCPIP and 10 General Purpose Engines

8.3.2 Evaluation Results

The ITR performance of the JMP workload with 64-bit addressing mode in IMS 15 showed a degradation of 1.87% compared with the JMP workload with 31-bit addressing mode. Table 36 shows the overall comparison between JMP workloads with 31-bit and with 64-bit addressing modes.

Table 36: JMP 64-bit Performance Result

Java Message Processing 64-bit Support				
	IMS 15 (31-bit)	**IMS 15 (64-bit)**	**Delta**	**Delta %**
CPU % Busy	78.88%	79.16%	0.28%	0.35%
ETR (Tran/Sec)	29031.51	28590.40	-441.11	-1.52%
ITR	36804.65	36117.23	-687.42	-1.87%
Total IMS Response Time (ms)	0.757	0.780	0.023	3.04%
Total General CPU µs/Tran	163.02	166.13	3.11	1.90%
IMS CPU Service Time/Tran (µs)	12.12	11.81	-0.31	-2.48%
ICON CPU Service Time/Tran (µs)	16.088	15.777	-0.311	-1.93%
JMP CPU Service Time/Tran (µs)	121.827	125.166	3.339	2.74%
Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	284	284	0	0.00%

Avg. CSA Above 16M Key 7 (M)	20.7	20.7	0.0	0.00%
Private Storage IMS Control Region				
Avg. LSQA Private (K)	608	608	0	0.00%
Avg. LSQA EPrivate (M)	13.8	13.7	-0.1	-0.72%
Avg. USER Private (K)	2272	2268	-4	-0.18%
Avg. USER EPrivate (M)	57.0	57.2	0.1	0.18%
Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	384	384	0	0.00%
Avg. LSQA EPrivate (M)	9.1	9.1	0.0	0.00%
Avg. USER Private (K)	628	628	0	0.00%
Avg. USER EPrivate (M)	264	264	0	0.00%
Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.8	13.9	0.1	0.72%
Avg. USER Private (K)	56	56	0	0.00%
Avg. USER EPrivate (M)	176	176	0	0.00%

A series of scaling tests were run to compare IMS 15 with 31-bit and 64-bit addressing modes at various CPU percent busy values. Figure 58 shows the ITR vs Transaction Rate comparison, Figure 59 shows the IMS Response Time vs Transaction Rate comparison, and Figure 60 shows the CPU % Busy vs Transaction Rate.

The JMP workload with 64-bit addressing mode showed a degradation in ITR, Total IMS response time, and JMP CPU service time per transaction compared to the JMP workload with 31-bit addressing mode. This degradation showed that there is a cost in moving from a 31-bit to a 64-bit addressing mode.

- ITR degraded by 1.87%
- Total IMS response degraded by 3.04%
- JMP CPU service time per transaction degraded by 2.74%

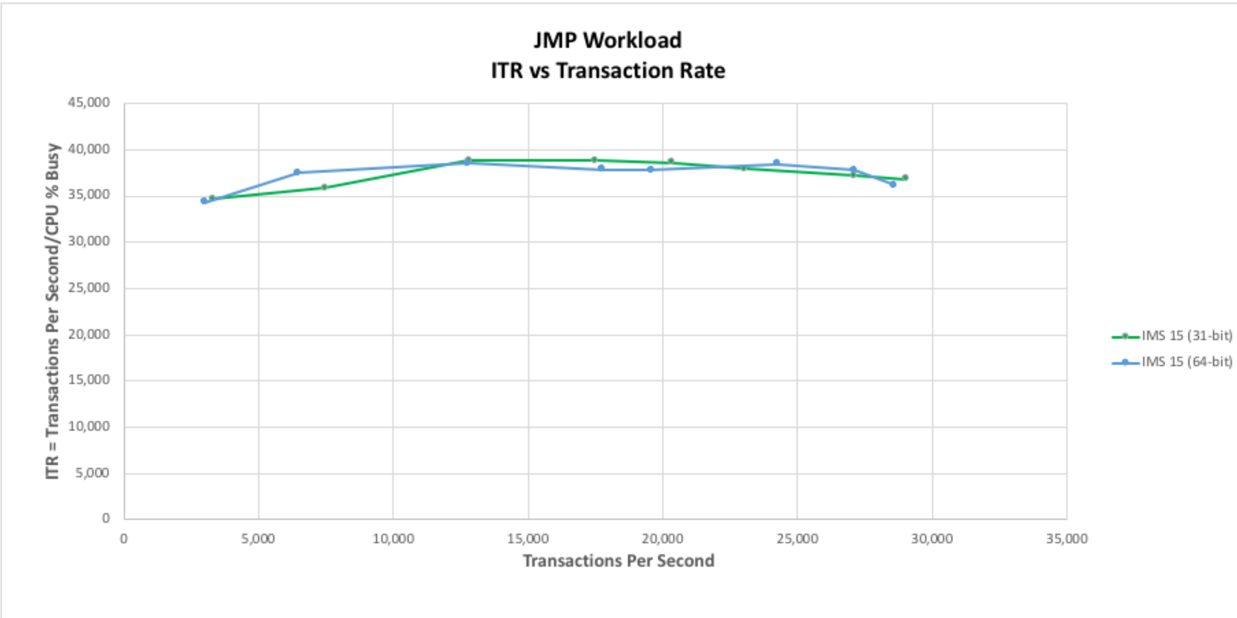

Figure 58: JMP ITR vs Transaction Rate Comparison

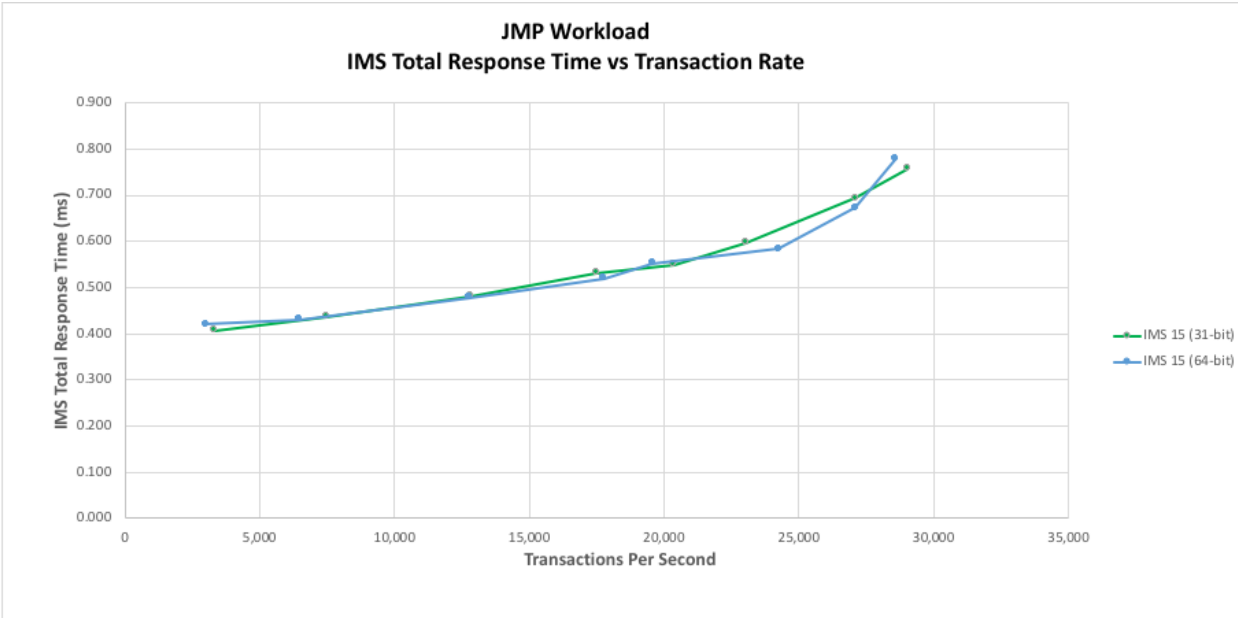

Figure 59: JMP IMS Total Response Time vs Transaction Rate Comparison

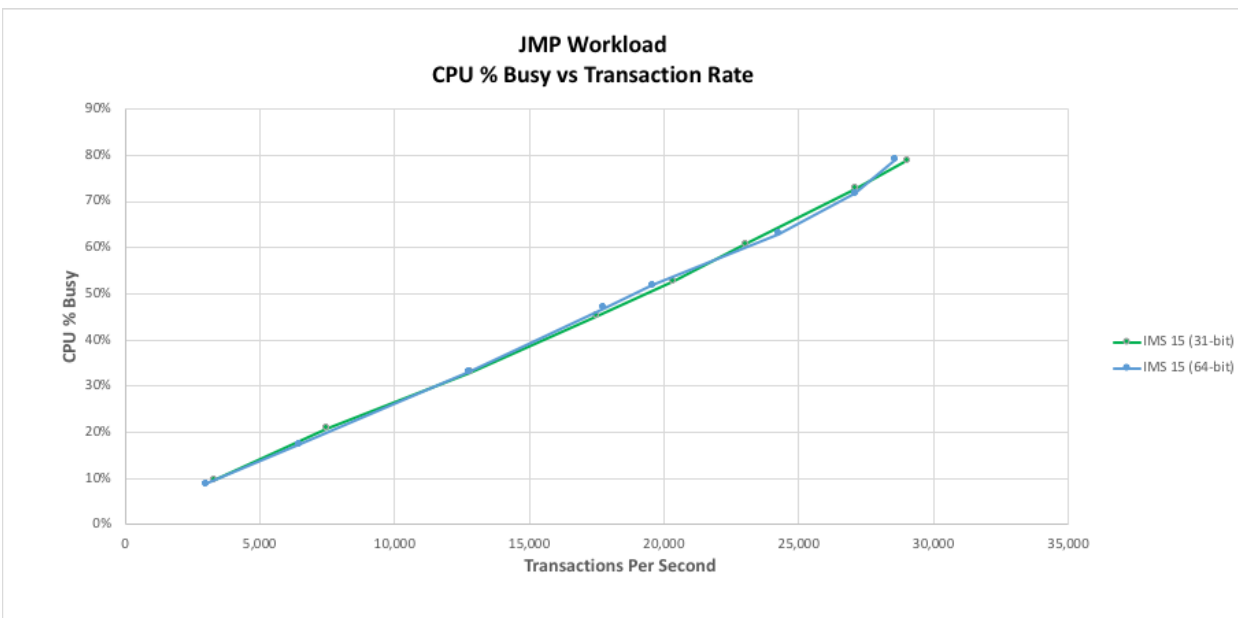

Figure 60: JMP CPU % Busy vs Transaction Rate Comparison

8.4 IMS 14 External Subsystem Attach Facility (ESAF) Thread Connection Pooling

Db2® for z/OS® added connection pooling support for threads that use the IMS-Db2 attachment facility. This feature is enabled by using the IMS SSM PROCLIB member to set the external subsystem module table (ESMT) parameter to DSNMIN20 instead of DSNMIN10. When this feature is enabled, Db2 maintains a separate connection pool (up to a maximum of 50) for each IMS dependent region such that when the same IMS transaction is executed, Db2 can reuse the resources that were associated from the previous pooled thread. This configuration avoids the process of deallocating and allocating the connection threads resulting in reduced scheduling overhead.

The IMS TM Db2 IRWW workload running PL/I applications with and without the IMS ESAF connection pooling enabled was compared in the study:

- Specifying DSNMIN20 instead of DSNMIN10 as the ESMT parameter to enable ESAF connection pooling.
- Set the Db2 system parameter CACHEDYN=YES to enable the global dynamic statement cache. Size the DSC pool by using the EDMSTMTC system parameter.
- Bind the Db2 packages with the KEEPDYNAMIC(YES) option.

8.4.1 System Configuration

The system configuration used to evaluate ESAF thread connection pooling is similar to the configuration used in the IMS TM-Db2 IRWW (Full PL/I) Workload Performance Evaluation in Section 6.7 except 68 MPP regions were used instead of 66 IFP and 2 MPP regions to process IMS transactions.

- LPAR 1 hosts IMS, 68 MPP regions, Db2, and ESAF with 10 General Purpose Engines
- LPAR 2 hosts a Java-based workload driver with varying client threads to achieve 85% CPU % busy and two General Purpose Engines

The IMS TM-Db2 workload was executed with the IMS processing limit count (PLCT) values of 0, 5, and 65535 to compare the performance of the ESAF thread connection pooling.

8.4.2 Evaluation Results

There was a significant improvement in ITR and total IMS response time when the ESAF thread connection pooling was enabled. Although the ESAF connection pooling demonstrated improvements in ITR, there was an increase of 54% in the usage of CSA storage above the 16M line due to the persistence of the connection pooling threads. In this workload there were 3300 Db2 active threads. The storage increase may vary depending on the number of active connection pooling threads. The results for PLCT = 0, 5, and 65535 are shown in Table 37, Table 38, and Table 39 respectively.

When PLCT is set to 0, Table 37 shows the ITR improved by 13% and total IMS response time decreased by 81%.

Table 37: ESAF Thread Connection Pooling with PLCT = 0

ESAF Thread Connection Pooling with PLCT = 0				
	Non-Connection Pooling	**Connection Pooling**	**Delta**	**Delta %**
CPU % Busy	85.61%	85.12%	-0.49%	-0.57%
ETR (Tran/Sec)	2785.21	3130.73	345.52	12.41%
ITR	3253.37	3678.02	424.65	13.05%
Total IMS Response Time (ms)	68.214	12.850	-55.364	-81.16%
Total General CPU µs/Tran	3073.736	2718.855	-354.881	-11.55%
IMS CPU Service Time/Tran (µs)	28.865	26.956	-1.909	-6.61%
ICON CPU Service Time/Tran (µs)	31.064	26.509	-4.555	-14.66%
MPP CPU Service Time/Tran (µs)	1986.700	1789.667	-197.033	-9.92%
Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	264	264	0	0.00%
Avg. CSA Above 16M Key 7 (M)	33.3	51.3	18.0	54.05%
Private Storage IMS Control Region				
Avg. LSQA Private (K)	652	656	4	0.61%
Avg. LSQA EPrivate (M)	14.5	14.5	0.0	0.00%
Avg. USER Private (K)	2376	2376	0	0.00%

Avg. USER EPrivate (M)	59.4	57.6	-1.8	-3.03%
Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	348	348	0	0.00%
Avg. LSQA EPrivate (M)	8.96	8.95	-0.01	-0.11%
Avg. USER Private (K)	624	624	0	0.00%
Avg. USER EPrivate (M)	263	263	0	0.00%
Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.3	13.5	0.2	1.50%
Avg. USER Private (K)	56	56	0	0.00%
Avg. USER EPrivate (M)	179	176	-3	-1.68%

When PLCT is set to 5, Table 38 shows the ITR improved by 35% and total IMS response time decreased by 86%.

Table 38: ESAF Thread Connection Pooling with PLCT = 5

ESAF Thread Connection Pooling with PLCT = 5				
	Non-Connection Pooling	**Connection Pooling**	**Delta**	**Delta %**
CPU % Busy	84.75%	85.16%	0.41%	0.48%
ETR (Tran/Sec)	3447.45	4685.86	1238.41	35.92%
ITR	4067.79	5502.42	1434.63	35.27%
Total IMS Response Time (ms)	69.532	9.814	-59.718	-85.89%
Total General CPU µs/Tran	2458.339	1817.383	-640.956	-26.07%

IMS CPU Service Time/Tran (µs)	27.780	22.977	-4.803	-17.29%
ICON CPU Service Time/Tran (µs)	28.889	23.570	-5.319	-18.41%
MPP CPU Service Time/Tran (µs)	1601.013	1192.163	-408.850	-25.54%
Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	264	264	0	0.00%
Avg. CSA Above 16M Key 7 (M)	33.8	52.1	18.3	54.14%
Private Storage IMS Control Region				
Avg. LSQA Private (K)	652	652	0	0.00%
Avg. LSQA EPrivate (M)	14.6	14.6	0.0	0.00%
Avg. USER Private (K)	2376	2376	0	0.00%
Avg. USER EPrivate (M)	60.3	57.9	-2.4	-3.98%
Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	348	348	0	0.00%
Avg. LSQA EPrivate (M)	8.96	8.95	0.00	-0.04%
Avg. USER Private (K)	624	624	0	0.00%
Avg. USER EPrivate (M)	263	263	0	0.00%
Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.4	13.4	0.0	0.00%
Avg. USER Private (K)	56	56	0	0.00%
Avg. USER EPrivate (M)	180	176	-4	-2.22%

When the PLCT is set to 65535, Table 39 shows the ITR improved by 27% and total IMS response time decreased by 86%.

Table 39: ESAF Thread Connection Pooling with PLCT = 65535

ESAF Thread Connection Pooling with PLCT = 65535				
	Non-Connection Pooling	Connection Pooling	Delta	Delta %
CPU % Busy	84.35%	85.65%	1.30%	1.54%
ETR (Tran/Sec)	3610.30	4681.59	1071.29	29.67%
ITR	4280.14	5465.95	1185.81	27.70%
Total IMS Response Time (ms)	77.464	10.167	-67.297	-86.88%
Total General CPU µs/Tran	2336.371	1829.507	-506.864	-21.69%
IMS CPU Service Time/Tran (µs)	27.135	23.132	-4.003	-14.75%
ICON CPU Service Time/Tran (µs)	28.926	22.596	-6.33	-21.88%
MPP CPU Service Time/Tran (µs)	1512.538	1203.591	-308.947	-20.43%
Common Storage Below and Above 16MB for Avg. Key 7				
Avg. CSA Below 16M Key 7 (K)	264	264	0	0.00%
Avg. CSA Above 16M Key 7 (M)	33.5	51.6	18.1	54.03%
Private Storage IMS Control Region				
Avg. LSQA Private (K)	652	652	0	0.00%
Avg. LSQA EPrivate (M)	14.5	14.5	0.0	0.00%
Avg. USER Private (K)	2376	2376	0	0.00%
Avg. USER EPrivate (M)	61.0	58.6	-2.4	-3.93%
Private Storage IMS DL/I Region				
Avg. LSQA Private (K)	348	348	0	0.00%

Avg. LSQA EPrivate (M)	9.01	9.01	0.00	0.00%
Avg. USER Private (K)	624	624	0	0.00%
Avg. USER EPrivate (M)	263	263	0	0.00%
Private Storage IMS Connect Region				
Avg. LSQA Private (K)	360	360	0	0.00%
Avg. LSQA EPrivate (M)	13.4	13.6	0.2	1.49%
Avg. USER Private (K)	56	56	0	0.00%
Avg. USER EPrivate (M)	181	176	-5	-2.76%

Figure 61 shows the ESAF Thread Pooling vs Non-Pooling comparisons for ITR and Figure 62 shows the ESAF Thread Pooling vs Non-Pooling comparisons for IMS Total Response Times.

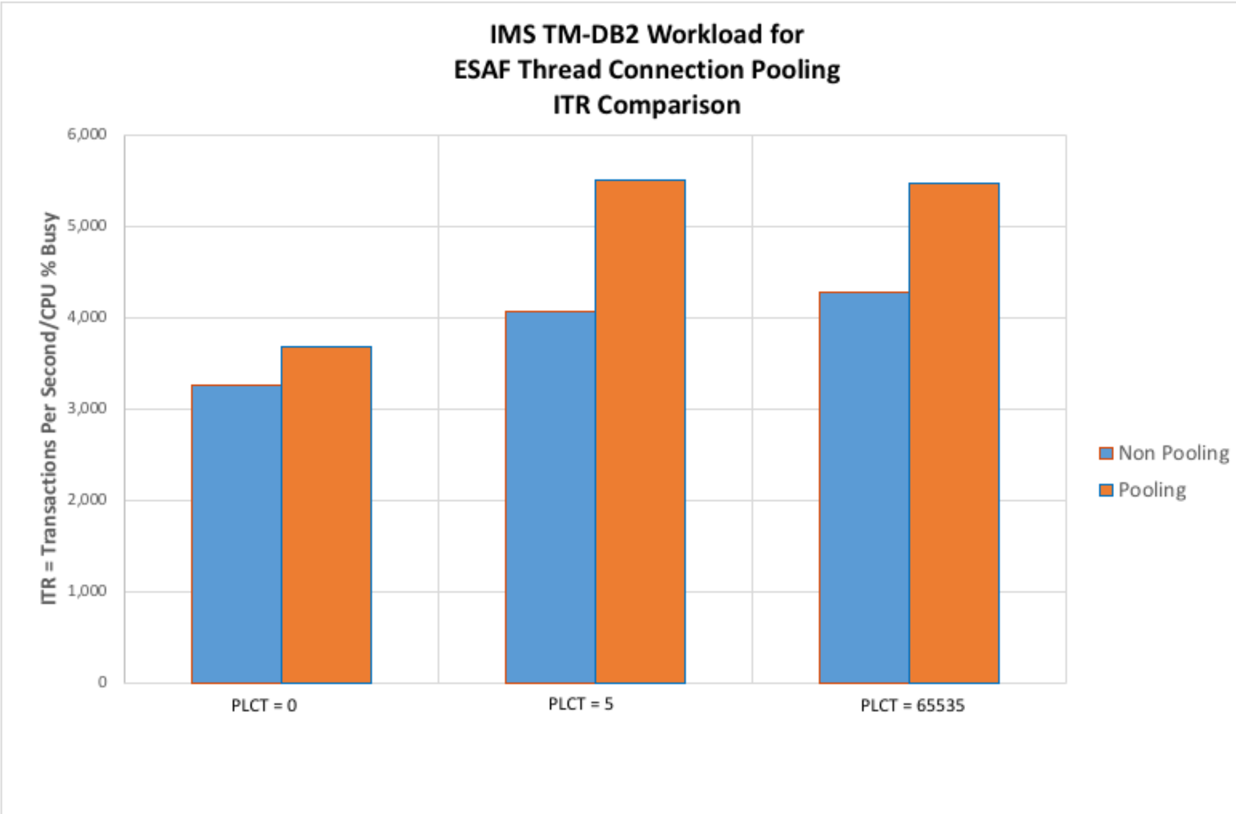

Figure 61: IMS TM-Db2 Workload ESAF Thread Connection Pooling ITR Comparison

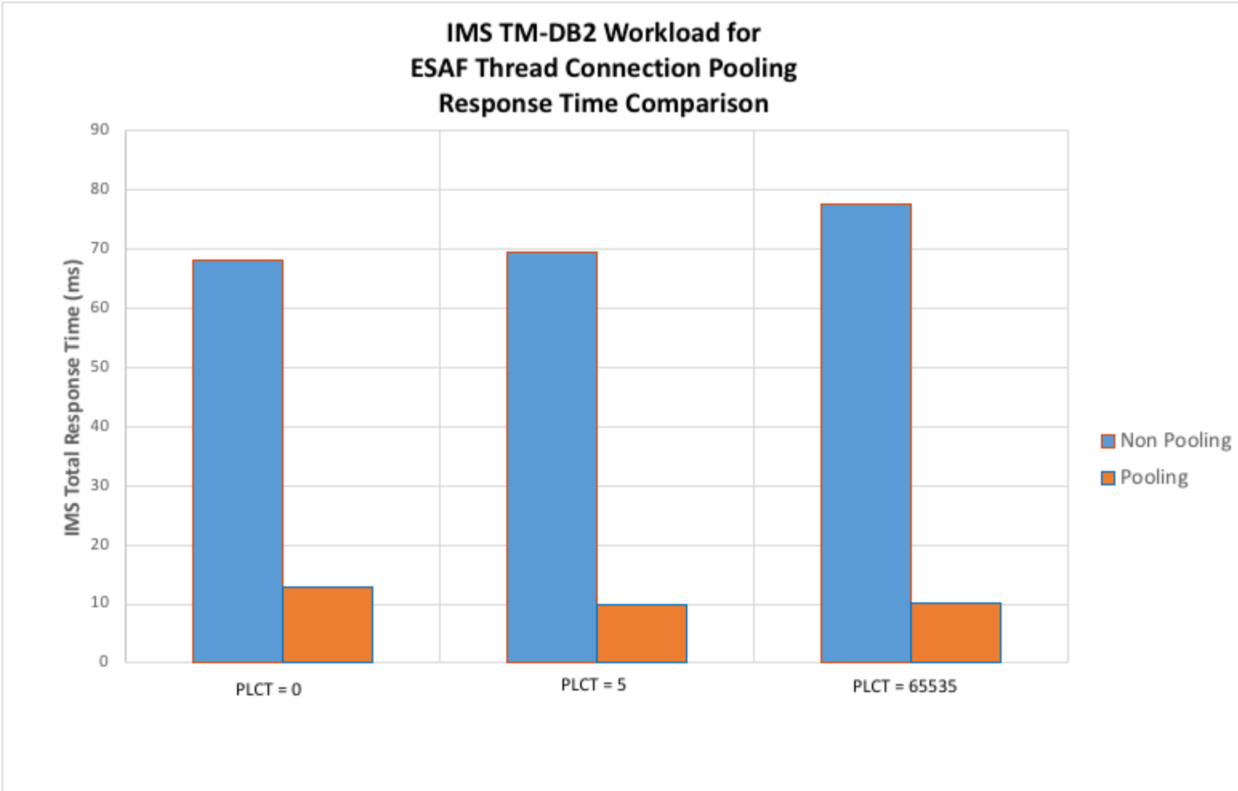

Figure 62: IMS TM-Db2 Workload ESAF Thread Connection Pooling IMS Response Time Comparison

9 Conclusion

Our performance evaluation for all the base workloads showed an increase in ITR of up to 4% between IMS 15 Managed ACB vs IMS 14. The ITR differences between the IMS 15 performance with and without Managed ACB were in the range of -0.95% degradation to 2.53% improvement.

The IMS 15 enhancements of IMS Logger Media Manager, Fast Path encryption, and Network Security Credential Propagation showed some noteworthy performance outcomes.

- The IMS Logger Media Manager had more than a 75% improvement in WADS and OLDS response times when utilizing both zHPF and zHyperWrite.
- The encryption of the Fast Path workload running with the new z14 encryption hardware cost 3.62% in ITR and 1.87% in Total IMS response time without the need for any application changes.

The IMS 14 enhancements also had some positive performance outcomes.

- The OTMA Resume TPIPE Parallelism showed that configuring OTMA with parallel TPIPEs reduced response time and throughput compared to non-parallel TPIPE.
- The added support for 64-bit JVM caused a degradation of about 1.87% in ITR when compared to 31-bit.
- The ESAF thread connection pooling increased the ITR by about 15% when compared to non-connection pooling.

Overall, IMS 15 continues to provide excellent performance, scalability, availability, and security as with past IMS releases. IMS customers can continue to benefit from these and other version enhancements for all areas of Transaction Manager, Database Manager, and Systems components.

10 IMS Resources

IMS Knowledge Center Links for IMS Data Set Encryption Support:

IMS V13:
https://www.ibm.com/support/knowledgecenter/SSEPH2_13.1.0/com.ibm.ims13.doc.sag/system_admin/ims_dataset_encryption.htm

IMS V14:
https://www.ibm.com/support/knowledgecenter/SSEPH2_14.1.0/com.ibm.ims14.doc.sag/system_admin/ims_dataset_encryption.htm

IMS V15:
https://www.ibm.com/support/knowledgecenter/SSEPH2_15.1.0/com.ibm.ims15.doc.sag/system_admin/ims_dataset_encryption.htm

IMS V15 Announcement Letter:
https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=gpateam&supplier=897&letternum=ENUS217-398

Data Set Encryption for IBM z/OS V2.2 Frequently Asked Questions:
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ131494

Video on setting up encryption
https://www.youtube.com/watch?v=zdSXRUSmkb4

Description of changes in z/OS and DFP to support data set encryption:

http://publibz.boulder.ibm.com/zoslib/pdf/OA50569.pdf

IBM z14 Announce and Information:

- **Announcement Letter:** https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS117-044&appname=USN
- **Introduction Page:** https://www.ibm.com/us-en/marketplace/z14
- **Press Release:** https://www-03.ibm.com/press/us/en/pressrelease/52805.wss
- **Technical Specs:** https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSD03046USEN&
- **Redbook:** http://www.redbooks.ibm.com/redbooks.nsf/pages/z14?Open

z/OS 2.3 Announcement Letter: https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS217-246&appname=USN

IMS 14 Performance Evaluation on IBM z14: https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=54013754USEN&

IMS 14 Information Center Link:
https://www.ibm.com/support/knowledgecenter/SSEPH2_14.1.0/com.ibm.ims14.doc/ims_product_landing_v14.html

IMS Home Page: https://www.ibm.com/it-infrastructure/z/ims

Recommended Publication: An Introduction to IMS - Your Complete Guide to IBM Information Management Systems

11 Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

Trademarks

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

The terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside of your enterprise, without the express consent of IBM.

Rights