

# DB2 9 and z/OS XML System Services synergy for distributed workloads

Akiko Hoshikaya – DB2 for z/OS Development  
Judy Ruby-Brown – ATS Dallas Systems Center

The previous versions of this white paper focused on initial measurements for pureXML™ technology using DB2® Version 9.1 for z/OS (DB2 9 hereafter) as an exploiter of z/OS XML System Services.

It is recognized that XML, which is widely adopted across industries for its flexibility and portability, is nonetheless verbose and its manipulations incur more overhead than its relational counterparts. Even though z/OS XML System Services is an efficient system level XML parser on z/OS, additional general CPU cost reduction is possible. With the proper software and hardware, local DB2 calls to z/OS XML System Services are redirected to a IBM System z™ Application Assist Processor (zAAP) specialty engine. Included are invocations from the attachments that run in TCB mode. DB2 already has functions that can be redirected to the System z Integrated Information Processor (zIIP) specialty engine.

This addendum describes the support provided in z/OS Version 1.10 that enables redirection of XML parsing to the zIIP specialty engines for enclave SRBs. This paper should be read in conjunction with “*DB2 9 and z/OS XML System Services Synergy Update*” at <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101227>, as it contains most of the technical details about the workload and methodology used, as well as prior results. The “future direction” statement that z/OS V1.10 fulfills is found at the end of this paper.

## ***zIIP support for distributed workload***

This measurement demonstrates the effect of z/OS 1.10 enhancements for zIIP redirection for z/OS XML System Services parsing. Prior to this enhancement, a portion of a distributed application had been zIIP-eligible, as it executes under an enclave SRB. Its invocation of z/OS XML System Services was not eligible for redirection. zAAP redirection is only available in TCB mode. With this enhancement, 100% of the z/OS XML System Services parsing execution can be redirected to the zIIP, if capacity exists.

## ***Prerequisites***

The primary prerequisites for 100% zIIP redirection with z/OS XML System Services parsing are:

## Hardware:

A zIIP processor, available on IBM System z10™ Enterprise Class (z10 EC™), IBM System z10™ Business Class (z10 BC), IBM System z9® Enterprise Class (z9™ EC, formerly z9-109) and IBM System z9 Business Class (z9 BC) systems.

## Software:

z/OS Version 1.10. Service to enable this capability for earlier z/OS releases is:

z/OS V1.9      OA23828/UA41774

z/OS V1.8      OA23828/UA41773

DB2 9 in New Function Mode (NFM).

DB2 9 already contains the proper fields for zIIP redirection on the accounting reporting, as documented in the previous white paper. There are no other changes required in DB2 to support this capability.

## Environment

The TPoX workload is a simplification of a real world security brokerage application. It has three different types of XML documents: Order, Security and CustAcc. CustAcc includes a customer with all his/her accounts. The information about holdings is included in the account data. The objective of this workload is to demonstrate the CPU usage of z/OS XML System Services in insert operations. In the previous paper, measurements were taken for the following two scenarios described in the TPoX benchmark:

- INSERT1: A customer places a new order to buy a stock
  - Result: Insert a new Order document in the collection of order documents
- INSERT2: A new customer signs up for an online brokerage account
  - Insert a new CustAcc document in the collection of CustAcc documents
  - Document size varies from 4K to 20K

The exercise demonstrates the effect of z/OS V1.9 (and above) on z/OS XML System Services redirection for a distributed application with heavy insert processing. This measurement reran only “INSERT2”. This is similar to the measurement described in “XML Insert comparison of z/OS Versions 1.8 to 1.9” in the previous paper.

## Configuration used:

Processor	IBM System z9 Enterprise Class (z9 EC)
LPAR configuration:	2 General Purpose CPs, dedicated; 2zIIPs
Memory:	24GB memory
Storage	IBM DS8300
Operating system	z/OS Version 1.9 plus OA23838 (and pre-reqs)
DB2	DB2 9 May 2008 PTF level

The measurements are performed in a controlled environment where there is no other workload running in the LPAR.

## Measurement Results

There were 40 concurrent threads inserting into the TPoX Customer Account table from simulated remote clients. A commit was issued after every insert. In “XML Insert comparison of z/OS Versions 1.8 to 1.9”, a commit was issued after every 10 inserts.

1. Case1-1 : Base run - without z/OS APARs to enable zIIP redirection for z/OS XML System Services.
2. Case1-2 : XML/SS zIIP run - with z/OS APARs to enable zIIP redirection for z/OS XML System Services.

In Case1-1, DB2 DRDA zIIP redirection can be utilized. As it shown in the table, about 54.3% of workload is eligible to run on zIIP processors.

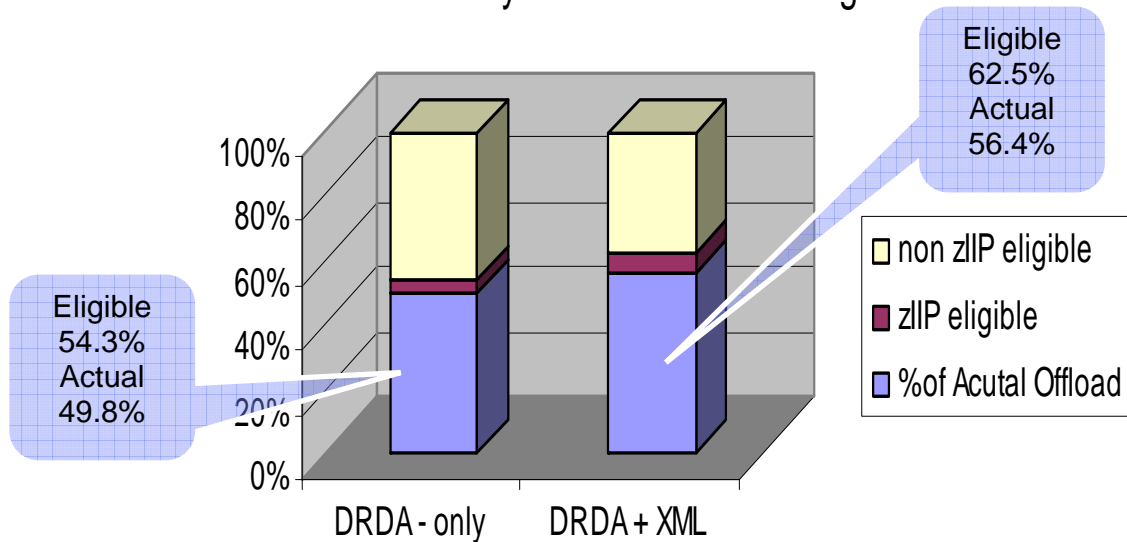
In Case1-2, in addition to DB2 DRDA zIIP redirection, 100% of XML System Service processing is eligible for zIIP. Thus, 62.5% of work (or additional 8.2%) becomes zIIP eligible.

Values	Case1-1	Case1-2
(time : in milliseconds)	DRDA only (no Z/OS XML SYSTEM SERVICES zIIP)	Both DRDA/XML zIIP
#of CPs	2	2
#of zIIP	2	2
Z/OS XML SYSTEM SERVICES OA23828	No	Yes
#of user threads	40	40
Transactions per second	1692	2016
CPs CPU utilization	89.06%	72.25%
zIIP CPU utilization	69.77%	68.83%
DB2 elapsed time (ms) per transaction	19.669	19.216
DB2 CP CPU time (ms) per transaction	0.75	0.575
DB2 zIIP CPU time (ms) per transaction	0.738	0.75
DB2 total CPU time (ms) per transaction (ms)	1.488	1.325
%of zIIP offloaded	49.60%	56.60%
DB2 total suspend time (ms) per transaction	16.737	16.399
WLM report - % of zIIP eligible work	54.29%	62.51%
- % of Actual redirected	49.78%	56.42%

The “% of zIIP offloaded” is calculated based on the DB2 accounting trace reports, while “WLM report - % of Actual redirected” comes from the RMF Workload Activity report. These numbers are very close. Customers can use the DB2 accounting traces for DB2 thread level monitoring, while RMF reports the system level of redirection.

The pertinent information is represented graphically:

## XML System Service zIIP usage



For Case 1-2., the RMF Workload Activity report is shown, as well as the matching accounting report formatted with IBM Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS, Version 4, (hereafter called OMPE).

### RMF Workload Activity from Case 1-2:

REPORT BY: POLICY=POL_XML		WORKLOAD=DB2		SERVICE CLASS=DB2DDF		RESOURCE GROUP=*NONE		PERIOD=1 IMPORTANCE=1					
				CRITICAL =NONE									
-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE---	---SERVICE	TIMES--	---APPL %---	-----STORAGE-----				
AVG	32.99	ACTUAL	18	SSCHRT	0.1	IOC	0	CPU	1153.804	CP	102.55	AVG	0.00
MPL	32.99	EXECUTION	18	RESP	0.2	CPU	327320K	SRB	0.000	AAPCP	0.00	TOTAL	0.00
ENDED	869793	QUEUED	0	CONN	0.1	MSO	0	RCT	0.000	IIPCP	14.32	SHARED	0.00
END/S	1773.96	R/S AFFIN	0	DISC	0.0	SRB	0	IIT	0.000				
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.1	TOT	327320K	HST	0.000	AAP	0.00	--PAGE-IN RATES--	
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	667575	AAP	0.000	IIP	132.77	SINGLE	0.0
AVG ENC	32.99	STD DEV	17					IIP	651.003			BLOCK	0.0
REM ENC	0.00					ABSRPTN	20K					SHARED	0.0
MS ENC	0.00					TRX SERV	20K	PROMOTED	0.000			HSP	0.0

The pertinent CPU percentage fields are:

1. CP – total executing on the general purpose CPs
2. IIP – total spent executing on the zIIP
3. AAP – total spent executing on the zAAP
4. IIPCP – total that could have been executed on a zIIP, but there was either not enough capacity or a zIIP was not configured
5. AAPCP – total that could have been executed on a zAAP, but there was either not enough capacity or a zAAP was not configured

Fields #4 and #5 are included in #1. Fields #2 and #3 are excluded from #1.

zIIP eligible :  $IIP+IIPCP / CP+IIP = 62.5\%$   
 $(132.77+14.32) / (102.55+132.77)$   
zIIP actual redirect :  $IIP / CP+IIP = 56.4\%$   
 $132.77/(102.55+132.77)$

Accounting from Case 1-2 :

AVERAGE	APPL (CL.1)	DB2 (CL.2)
ELAPSED TIME	0.019216	0.017801
NONNESTED	0.019216	0.017801
STORED PROC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
<b>CP CPU TIME</b>	<b>0.000575</b>	0.000536
AGENT	0.000575	0.000536
NONNESTED	0.000575	0.000536
STORED PROC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	0.000000	0.000000
<b>SECP CPU</b>	<b>0.000078</b>	N/A
<b>SE CPU TIME</b>	<b>0.000750</b>	0.000701
NONNESTED	0.000750	0.000701
STORED PROC	0.000000	0.000000

The fields

- SECP CPU can reflect one of two things:
  - It is zIIP (not zAAP) eligible work that ran on a general purpose CP because a zIIP was not configured
  - The capacity of zIIP processors had been reached and the “overflow” could not be accommodated.
- Statistics for zAAP-eligible work that runs on a general purpose CP are not available to DB2.
- SE CPU TIME represents the CPU time consumed on all specialty engines combined.

SE eligible :  $SE\ CPU\ time + SECP\ CPU / CP\ CPU\ time + SE\ CPU\ TIME = 62.5\%$   
 $(.000750+.000078) / (.000575+.000750)$   
SE actual redirection :  $SE\ CPU\ time / CP\ CPU\ TIME + SE\ CPU\ TIME = 56.6\%$   
 $.000750 / (.000575+.000750)$

## Conclusion

XML System Services work requested from distributed threads has been already eligible partially for zIIP engines. With z/OS support to enable zIIP for XML System Services parsing, distributed DB2 XML workload with insert/update could enjoy further redirection to zIIP engines.

As discussed in previous papers, the ratio of zIIP redirection is highly dependent of the XML document that is being parsed, particularly its size and complexity.

## **Satisfaction of future direction statement**

The previous white paper listed a capability called “Additional zIIP Exploitation” under a Future Directions statement for z/OS Version 1.10:

*“z/OS XML System Services plans include additional zIIP exploitation, specifically enabling all z/OS XML parsing in enclave SRB mode to be eligible for zIIP. For example, with respect to DB2, z/OS XML processing may be partially directed to zIIPs when utilized as part of a distributed request (like DB2 DRDA) today. This enhancement can help further benefit DB2 pureXML workloads by optionally directing all z/OS XML System Services parsing that is executed in enclave SRBs to the zIIP. This function is planned to be available on z/OS V1.8 and V1.9 with PTF for APAR OA23828. Delivery of this function satisfies the statements of direction in Hardware Announcement 107-190, dated April 18, 2007, and Software Announcement 207-175 dated August 7, 2007.”*

### **Special Notice**

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer’s ability to evaluate and integrate them into the customer’s operational environment. While IBM may have reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Anyone attempting to adapt these techniques to their own environments do so at their own risk

Any performance data contained in this document were determined in various controlled laboratory environments and are for reference purposes only. Customers should not adapt these performance numbers to their own environments as system performance standards. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.



Copyright IBM Corporation 2008

IBM Corporation  
New Orchard Rd.  
Armonk, NY 10504  
U.S.A

Produced in the United States of America  
03/08

All Rights Reserved

CICs, DB2, DB2 Connect, DRDA, IBM, IBM eServer, IBM logo, IMS, OMEGAMON, pureXML, RMF, System z, System z9, System z10, Tivoli, WebSphere, z9, z10 EC, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corporation of the United States, other countries or both.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

Information concerning non-IBM products was obtained from the suppliers of their products or their published announcements. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.