



IBM Software Group

Leveraging zIIP, zAAP Specialty Engines with DB2 for z/OS

SHARE in Boston

Greg Dyck
DB2 for z/OS Development
gad@us.ibm.com

August 2010

© IBM Corporation

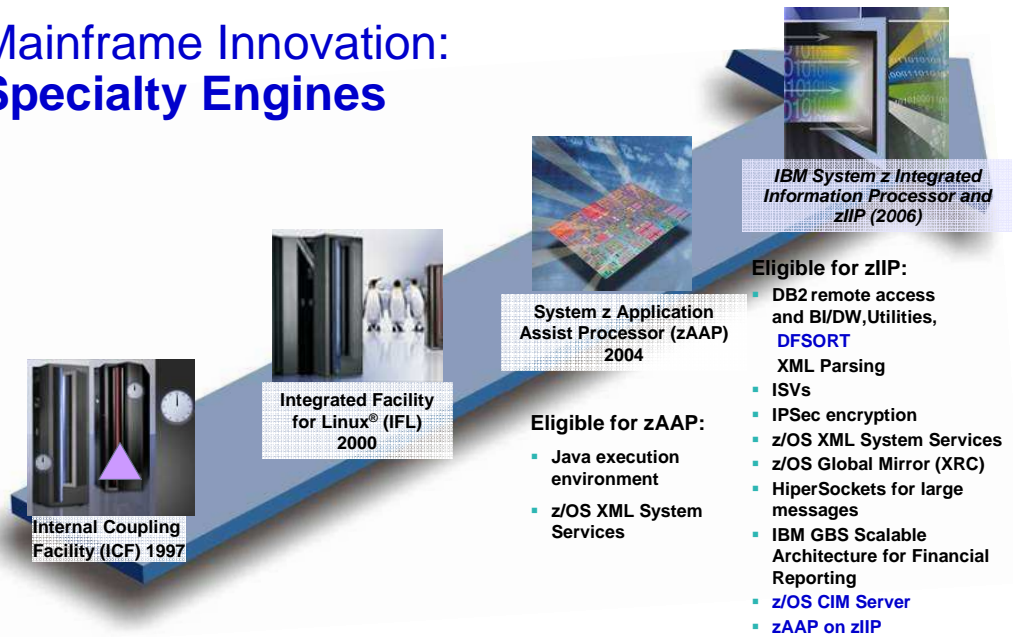
TOPICS

- zIIP & zAAP Overview
- DB2 workloads that leverage zIIP & zAAP
- Estimation & Monitoring of zIIP & zAAP redirect
- Recent zIIP & zAAP related enhancements

zIIP & zAAP

Overview

Mainframe Innovation: Specialty Engines



How zIIP and zAAP are similar

- **Both run asynchronously to general purpose processors**
- **Neither can run z/OS or be IPLed with z/OS (or any other operating system)**
- **Both receive eligible work from z/OS**
- **Both supported with z/OS 1.6 or later**
- **IBM does not impose any IBM software charges on either**
- **As of the date of this publication there are no known instances where an ISV has imposed software charges on either**
- **Both have the same price**
- **Both have the same technology dividend**
- **PROJECTCPU tool can measure the eligible workload for both**
- **RMF™ monitors both zAAP and zIIP activity**
- **WLM manages both zAAP and zIIP workloads**

How zIIP and zAAP are different

zAAP	zIIP
Introduced in 2004 .	Introduced in 2006 .
System z Application Assist Processor (originally the zSeries Application Assist Processor). Available on System z9 and System z10 servers, and IBM eServer™ zSeries® 990 and 890 (z990, z890).	System z Integrated Information Processor Available on System z9 and System z10 servers.
Intended to help implement new application technologies on System z, such as Java and XML.	Intended to help integrate data and transaction processing across the enterprise and on to System z9 and System z10.
<p>Exploiters include:</p> <ul style="list-style-type: none"> ▪ Java via the IBM SDK (IBM Java Virtual Machine (JVM)), such as portions of: <ul style="list-style-type: none"> – WebSphere Application Server – IMS™ – DB2 – CICS® – Java batch – CIM Client applications ▪ z/OS XML System Services, such as portions of: <ul style="list-style-type: none"> – DB2 9 (New Function Mode) – Enterprise COBOL V4.1 – Enterprise PL/I V3.8 – IBM XML Toolkit for z/OS, V1.9 and later – CICS TS V4.1 	<p>Exploiters include, portions of:</p> <ul style="list-style-type: none"> ▪ DB2 V8 for z/OS, DB2 9 for z/OS <ul style="list-style-type: none"> – Data serving – Data Warehousing – Select utilities ▪ z/OS Communications Server <ul style="list-style-type: none"> – Network encryption – HiperSockets for large messages ▪ z/OS XML System Services <ul style="list-style-type: none"> – DB2 9 New Function Mode ▪ z/OS Global Mirror (XRC), System Data Mover (SDM) ▪ IBM GBS Scalable Architecture for Financial Reporting ▪ z/OS CIM server



What is the zAAP on zIIP capability?

- **A new capability that can enable System z Application Assist Processor (zAAP) eligible workloads to run on System z Integrated Information Processors (zIIPs).**
 - For customers with no zAAPs and zIIPs
 - **The combined eligible workloads may make the acquisition of a single zIIP cost effective.**
 - For customers with only zIIP processors
 - **Makes Java and z/OS XML System Services -based workloads eligible to run on existing zIIPs – maximizes zIIP investment.**
 - Available September 25, 2009 with z/OS V1.11 and z/OS V1.9 and V1.10 (with PTF)
 - **This new capability is not available for z/OS LPARS if zAAPs are installed on the server.**

z/OS V1.11 is enhanced with a new function that can enable System z Application Assist Processor (zAAP) eligible workloads to run on System z Integrated Information Processors (zIIPs). This function can enable you to run zIIP- and zAAP-eligible workloads on the zIIP.

This new capability is ideal for customers without enough zAAP- or zIIP-eligible workload to justify a specialty engine today; the combined eligible workloads may make the acquisition of a zIIP cost effective. This new capability is also intended to provide more value for customers having only zIIP processors by making Java and XML-based workloads eligible to run on existing zIIPs.

Customers who have already invested in zAAP, or have invested in both zAAP and zIIP processors, should continue to use these as this maximizes the new workload potential for the platform.

This capability is available with z/OS V1.11 (and z/OS V1.9 and V1.10 with PTF for APAR OA27495) and all System z9 and System z10 servers; some additional restrictions apply. This capability does not provide an overflow so additional zAAP eligible workload can run on the zIIP, it enables the zAAP eligible work to run on zIIP when no zAAP is defined. Therefore, this new capability is not available for z/OS LPARS if zAAPs are installed on the server.

This new capability does not remove the requirement to purchase and maintain one or more general purpose processors for every zIIP processor on the server. This part of the IBM terms and conditions surrounding the IBM System z specialty engines is unchanged.

How to enable the zAAP on zIIP capability

- **The capability ships default enabled with z/OS V1.11.**
 - Parameter in SYS1.PARMLIB(IEASYSxx) : ZAAPZIIP = YES (default in z/OS V1.11)
 - If you wish to disable the function for any reason, you must IPL with ZAAPZIIP=NO in the IEASYSxx Parmlib member.

- **Also available with z/OS V1.9 and V1.10**
 - With PTF for APAR OA27495, and
 - Enabled with ZAAPZIIP=YES in the IEASYSxx Parmlib (the default is NO)

- **This new capability does not remove the requirement to purchase and maintain one or more general purpose processors for every zIIP processor on the server.**

The capability ships default enabled with z/OS V1.11. If you wish to disable the function for any reason, you must IPL with ZAAPZIIP=NO in the IEASYSxx Parmlib member.

The capability is available in z/OS V1.9 and V1.10 by PTF for APAR APAR OA27495, but installing the function by the appropriate PTF will not automatically enable the function. If you wish to enable the capability, you must explicitly enable the function by adding ZAAPZIIP=YES in the IEASYSxx Parmlib member after installing the appropriate service.

To summarize:

Parameter in IEASYSxx: ZAAPZIIP=NO | YES. A value of NO disables this function. A value of YES enables the function.

The default in z/OS V1.11 is “YES.”

The default in z/OS V1.9 and V1.10 is “NO.” If you wish to enable this support on z/OS V1R9 or z/OS V1R10, the parameter is called ZZ in IEASYSxx. This value is also accepted on z/OS V1R11 on which ZZ is a synonym for the ZAAPZIIP parameter.

This new capability is not available if zAAPs are installed on the server and z/OS is operating in an LPAR. Otherwise, this capability is available with z/OS V1.11 (and z/OS V1.10 and V1.9 with PTF for APAR OA27495) and all System z9 and System z10 servers.

In addition, this new capability does not change IBM terms and conditions surrounding IBM System z specialty engines. Current System z specialty engine terms and conditions are unchanged. You must purchase and maintain one or more general purpose processors for every zIIP processor on the server. You should plan for this capability accordingly.

DB2 Workloads that leverage zIIP

- Portions of the following DB2 for z/OS workloads will benefit from zIIP
 1. ERP, CRM, Business Intelligence or other enterprise applications
 - Via DRDA over a TCP/IP connection (enclave SRBs)
 - Workloads using DB2 Connect, T4 JCC Universal driver, Data Server Client and CLI/ODBC, JDBC, SQLJ, .NET, pureQuery APIs
 - Remote native SQL procedures (DB2 9 for z/OS NFM)
 - XML parsing processing (DB2 9 for z/OS NFM)



2. Data warehousing / Business Intelligence applications
 - CPU intensive parallel queries, including star schema queries
 - DB2 9 higher percentage of parallel queries are eligible
 - DB2 10 still more queries eligible for parallelism
 3. DB2 for z/OS utility functions used to maintain index structures and DFSORT sort for some utilities, DB2 10 Runstats
 4. DB2 10 Buffer Manager pre-fetch processing
- Usage of zIIP is transparent to applications
 - No changes to applications
 - Supported in DB2 V8 & DB2 9 Compatibility & New Function Modes

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Not all of this work will be run on zIIP. z/OS will direct the work between the general processor and the zIIP. The zIIP is designed so a software program can work with z/OS to dispatch workloads to the zIIP with no anticipated changes to the application – only changes in z/OS and DB2.

IBM DB2 for z/OS version 8 will be the first IBM software able to take advantage of the zIIP. Initially, the following workloads can benefit:

- SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP), CRM (Siebel), or business intelligence and are expected to provide the primary benefit to customers. Stored procedures and UDFs run under TCBs, so they are not generally eligible, except for the call, commit and result set processing. V9 remote native SQL Procedure Language is eligible for zIIP processing.
 - BI application query processing utilizing DB2 parallel query capabilities – including the star schema queries and adding other large parallel queries; and
 - functions of DB2 utilities LOAD, REORG and REBUILD that perform index maintenance.
 - With DB2 10, Buffer Manager pre-fetch processing will also be executed on a zIIP.
- For more, see <http://www.ibm.com/systems/z/ziip/>

How does zIIP work ?

The zIIP is designed so that a program can work with z/OS to have a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The types of DB2 work listed below are those executing in enclave SRBs, portions of which can be redirected to the zIIP.

Example 1 = Distributed SQL requests (DRDA)

Workload that access DB2 for z/OS V8 via DRDA over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP.

- Includes DRDA DB2 9 Native (non-WLM) SQL Stored Procedures and XML parsing.

Example 2 = CPU intensive parallel queries

CPU intensive parallel queries will now use enclave SRBs. z/OS directs a portion of this work to the zIIP.

Example 3 = DB2 utilities for index maintenance and DFSORT

DB2 Utilities LOAD, REORG, and REBUILD will use enclave SRBs for the portion of the processing that is related to index maintenance and DFSORT. z/OS directs a portion of this work to the zIIP.

z/OS dispatches work in either TCB (Task Control Block) mode or SRB (Service Request Block) mode. DB2 parallel tasks use SRB mode and are assigned the same importance as the originating address space. Preemptible enclaves are used to do the work on behalf of the originating address space TCB or SRB. Enclaves are grouped by common characteristics and service requests and since they are preemptible, the z/OS dispatcher (and WLM) can interrupt these tasks for higher priority ones.

There are two types of preemptible SRBs: Client SRBs and Enclave SRBs

If the DB2 for z/OS V8 request is coming through DRDA over TCP/IP, then most of that work (other than stored procedures and user-defined functions SQL work) is executed in enclave SRBs. If the request is coming over local or a native connection, then that work is dispatched between TCBs, client SRBs and SRBs.

Portions of parallel queries and some utility index maintenance also use enclave SRBs. Only the enclave SRB work is eligible to be redirected to zIIP. The client SRB work, non-preemptible SRB work, or TCB work is not eligible to be redirected to the zIIP. DB2 directs z/OS to dispatch or redirect a portion of the eligible work to the zIIP. The above types of DB2 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. zIIP is designed so a subsystem program, like DB2, can work with z/OS to dispatch a portion of its enclave SRB work directed to the zIIP with no anticipated changes to the application – only changes to DB2 and z/OS.

Customers must be current on hardware (System z9), current on software (z/OS 1.6 or later, DB2 V8 or later) and have a workload peak using the types of work supported by the zIIP:

- Remote SQL processing of DRDA network-connected applications over TCP/IP

These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business intelligence running on other platforms. Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. Stored procedures and UDFs run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing. V9 remote native SQL Procedure Language Stored Procedure SQL processing (in DBM1 address space) is eligible for zIIP processing.

- Parallel queries

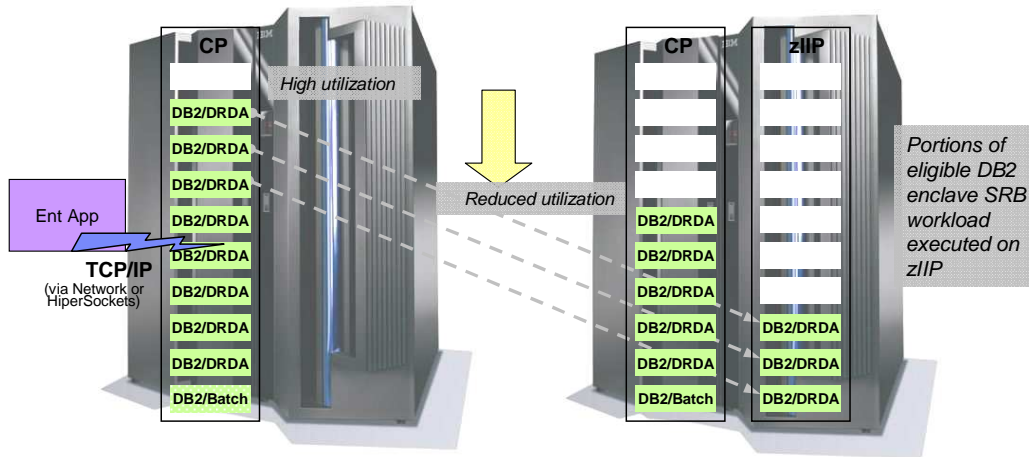
CPU intensive parallel queries after their parallel group CPU consumption exceeds certain threshold, they will be eligible for zIIP redirect. If the work comes in remotely over DRDA using TCP/IP it will benefit by the DRDA zIIP redirect as well.

- DB2 utility index processing

Functions of the LOAD, REORG and REBUILD DB2 utilities that perform index maintenance are eligible for zIIP. This is not a common peak capacity constraint, but could be useful in reducing CPU charges.

Example 1: Enterprise Distributed Applications

Enterprise Applications that access DB2 for z/OS V8 via DRDA, including DB2 9 Native SQL Stored Procedures and XML parsing, over a TCP/IP connection will have portions of these SQL requests directed to the zIIP.



For illustrative purposes only
Actual workload redirects may vary

Distributed Relational Database Architecture (DRDA) is the infrastructure for remote access that DB2 supports between the requesting RDBMS and the serving RDBMS. The other remote access type is DB2 private protocol access but no enhancements are planned for private protocol..

An application uses DRDA application requestor or server to access a remote database. DB2 Connect and JCC Type 4 Universal Driver are examples of a DRDA application requestor. Regarding the zIIPs, if DB2 for z/OS workload comes over TCP/IP and is DRDA compliant, a portion of that DB2 workload is eligible to be redirected to the zIIP. Both TCP/IP and DRDA are needed.

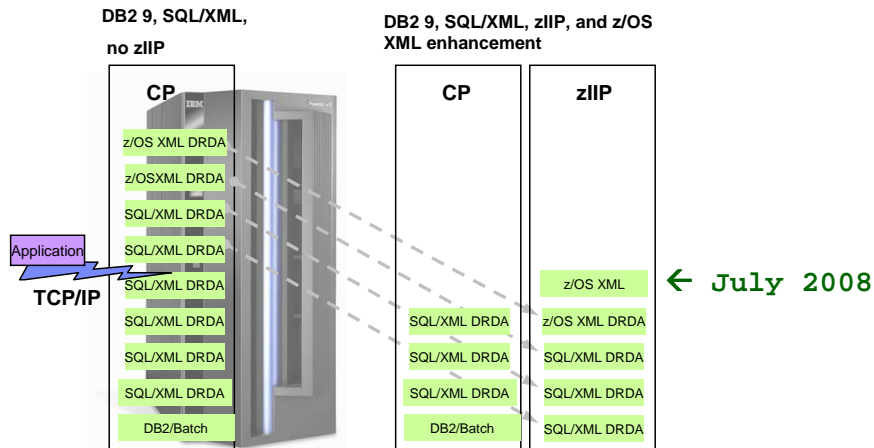
Transactions / Queries that access DB2 via DRDA over TCP/IP connections are dispatched within z/OS as enclave SRBs. z/OS directs a portion of this work to the zIIP. Only DRDA workloads that are coming via TCP/IP are zIIP eligible. DRDA workloads that come via SNA are not zIIP eligible.

Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. WLM managed Stored Procedures and UDFs SQL processing run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing.

DB2 9 remote native SQL Procedure Language Stored Procedure SQL processing that run in DBM1 address space are eligible for zIIP processing.

Example 1A: z/OS XML and DB2 9 DRDA and zIIP

Enterprise Applications that access DB2 9 for z/OS via DRDA over a TCP/IP connection can have all of enclave SRB SQL/XML System Services requests directed to the zIIP



For illustrative purposes only. Single application only. Actual workload redirects may vary
 DB2 9 for z/OS XML enclave SRB processing eligible to run on zIIP
 DB2 9 all z/OS XML System Services processing eligible for zIIP (July 2008 – OA23828 z/OS 1.8, 1.9)

eXtensible Markup Language (XML) is a W3C (World Wide Web Consortium) standard for the interchange of structured data. While XML has a look and feel similar to HTML (Hypertext Markup Language), the purpose of XML is dramatically different. HTML describes how information is to be presented (for example, HTML is used to create the display of the webpages you look at), but XML describes its actual meaning.

XML is not a fixed format database (although it can contain data and database elements) rather it is a language that uses tags to describe data. Since XML is not fixed format, it can combine many diverse types of data into many diverse structures. This lack of fixed format makes it easy to exchange data between different companies and different hardware and software technologies.

XML Parsing helps determine what is data and what is the meta data (description of the data).

Parsing is very often the most CP-intensive part of XML. DB2 will use z/OS XML System Services for parsing.

With July 2008 Apar OA23828 (z/OS 1.8, 1.9) all z/OS XML System Services parsing that is executed in enclave SRB mode from zIIP-eligible enclaves will be redirected to eligible to run on the zIIP.

DB2 DRDA zIIP Redirect processing

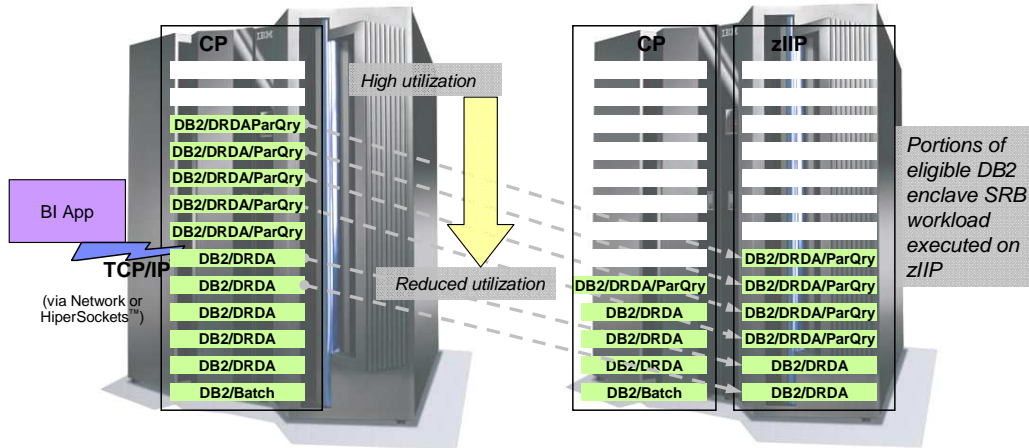
- Applicable to DRDA workload over TCP/IP connection
 - Only work done under enclave SRB is eligible.
 - Portion of the eligible work will be redirected.
- WLM managed Stored Procedure & UDF (User Defined Function) SQL processing under TCB, are not eligible for zIIP redirect
 - Stored Procedure Call, Results set and Commit processing that run under enclave SRB, are eligible for zIIP redirect.
- DB2 9 DRDA Native SQL Stored Procedure SQL processing and DRDA XML parsing is eligible for zIIP redirect.
- DB2 to DB2 TCP/IP DRDA Server processing is eligible for zIIP redirect
 - Requester DB2 local processing is not eligible.
 - Processing under TCB

The amount of DRDA via TCP/IP to be redirected might be affected by the DB2 WLM managed stored procedures and user-defined functions in the application. DB2 WLM managed stored procedures and user-defined functions run under TCBs, not SRBs, so they are not eligible for zIIPs. The call to the stored procedures, the commit processing and the result of processing may be redirected but the primary process which is in TCB mode cannot be redirected. A measurement with WLM managed Stored Procedure workload showed about 10% zIIP redirect.

In DB2 9 remote native SQL Stored Procedure SQL calls that comes in via TCP/IP DRDA will run in the DBM1 address space using enclave SRBs and hence are eligible for zIIP redirect.

Example 2.0: Business Intelligence Applications

Parallel queries via DRDA over a TCP/IP connection will have portions of this work directed to the zIIP



For illustrative purposes only

Actual workload redirects may vary depending on how long the queries run, how much parallelism is used, and the number of zIIPs and CPs employed

Business Intelligence applications that send requests that utilize DB2 for z/OS complex parallel queries may have portions of these SQL requests directed to the zIIP when DB2 gives z/OS the necessary information. If you have a warehouse that uses parallel processing and significant CPU time, then the zIIP may provide a benefit. If the work comes in remotely over DRDA using TCP/IP, then the initial work is eligible as remote work. If the work comes via local connection and includes parallel queries, after the initial transaction processing time, the parallel processing threads are eligible and can process on the zIIP.


Eligible query work has been broadened beyond just the initial star schema parallel queries to all CPU intensive parallel queries. Any long running CPU intensive parallel query can now benefit from zIIP. Star join is a special join technique that DB2 uses to efficiently join tables that form a star schema. A star schema is a logical database design that is included in decision support applications. Complex parallel queries now use enclave SRBs. z/OS directs a portion of this work to the zIIP, after the CPU usage exceeds certain threshold for each parallel group.

Parallel query workloads may benefit from two redirected tasks, the main task which is a DRDA request and a child task which is the parallel query. The child and main tasks are additive which means a combination of them is expected to yield a larger amount of redirect than that of just DRDA via TCP/IP alone. Longer running CPU intensive parallel queries see higher benefit. Benefits to a data warehousing application may vary significantly depending on the details of that application.

This is an example of a business intelligence application type in which complex parallel queries via DRDA are being used.

Example 2.5 : Business Intelligence Applications – No DRDA

Parallel queries via LOCAL connection will have portions of this work directed to the zIIP



This chart illustrates zIIP redirect for local (non- DRDA) parallel queries.

Notice that zIIP benefit does not have DRDA zIIP benefit compared to the previous chart.

CPU intensive parallel queries from any environment (BATCH, IMS, CICS, etc...) can benefit from zIIP redirect.

Parallel Query zIIP Redirect Processing

- Applicable to CPU intensive parallel queries using Star Schema and non Star Schema
 - Portion of the child task processing will be redirected after certain CPU usage threshold has exceeded for each parallel group.
 - Main tasks coming in via DRDA via TCP/IP will benefit from the DRDA zIIP redirect as well.
- The combined child & main tasks coming in through DRDA via TCP/IP is expected to yield additional zIIP redirect.
- Increased zIIP redirect potential with Star Join dynamic Index ANDing enhancement in DB2 9.
 - Lab workloads achieved additional 10% zIIP redirect
- Benefits parallel queries consuming high CPU.

Enabling Parallelism:

Two zPARMs :

Current Degree CDSSRDEF = 1 (Disabled) or ANY (Enabled) at the subsystem level - recommend setting to 1 and enable parallelism selectively.

Max Degree PARAMDEG = 0 to 254 per parallel group (OK to set it to 0 and let DB2 decide – use a number to limit the degree if needed to limit CPU and storage usage by parallel tasks)

Static Queries : DEGREE = 1 (Disabled) or ANY (Enabled) in Plan / Package BIND

Dynamic Queries : SET CURRENT DEGREE = 1 (Disabled) or ANY (Enabled)

Make sure BP seq threshold VPSEQT is non zero (default 80%).

The actual work of complex parallel queries redirects may vary depending on the following:

How long the queries run and how CPU intensive they are.

How much parallelism is used

The table design being used

The number of zIIPs and CPs employed

The following factors are taken into consideration by DB2 for the exploitation of query parallelism:

DB2 attempts to exploit sysplex parallelism if the underlying DB2 instance is a DS member (not supported for star join queries)

DB2 attempts to exploit CP parallelism if there are more than one CP available.

DB2 attempts to exploit I/O parallelism if neither sysplex parallelism nor CP parallelism is eligible (no zIIP benefit).


The estimated cost of each query block

Number of CPs if CPU-bound (in general, parallelism is disabled if only 1 CP and no zIIP)

Number of partitions of the leading table if I/O bound (in general, I/O parallelism is disabled if table is not partitioned)

Example 3: DB2 for z/OS Utilities

zIIP redirect is applicable to DB2 for z/OS utilities used to maintain index structures



For illustrative purposes only, actual workload redirects may vary.

Only the portion of the DB2 utilities used to maintain index structures (within LOAD, REORG, and REBUILD) is redirected.

An index allows quick access to the rows in a table. Indexes are created using one or more columns of a table. Over time, as data in a large database is manipulated indexes can become less efficient. They need to be updated and maintained. This can be a very big task.

The DB2 utilities LOAD, REORG, and REBUILD INDEX now use enclave SRBs for the portion of the process that is eligible to execute in a zIIP. The LOAD utility loads tables and the REORG utility improves index performance while the REBUILD INDEX utility creates or rebuilds your indexes. Only the build portion of the LOAD, REORG, and REBUILD DB2 utility processing related to index maintenance and can be redirected to zIIPs.

The amount of workload eligible for zIIPs depends on the degree of index build processing for the Utility workload.

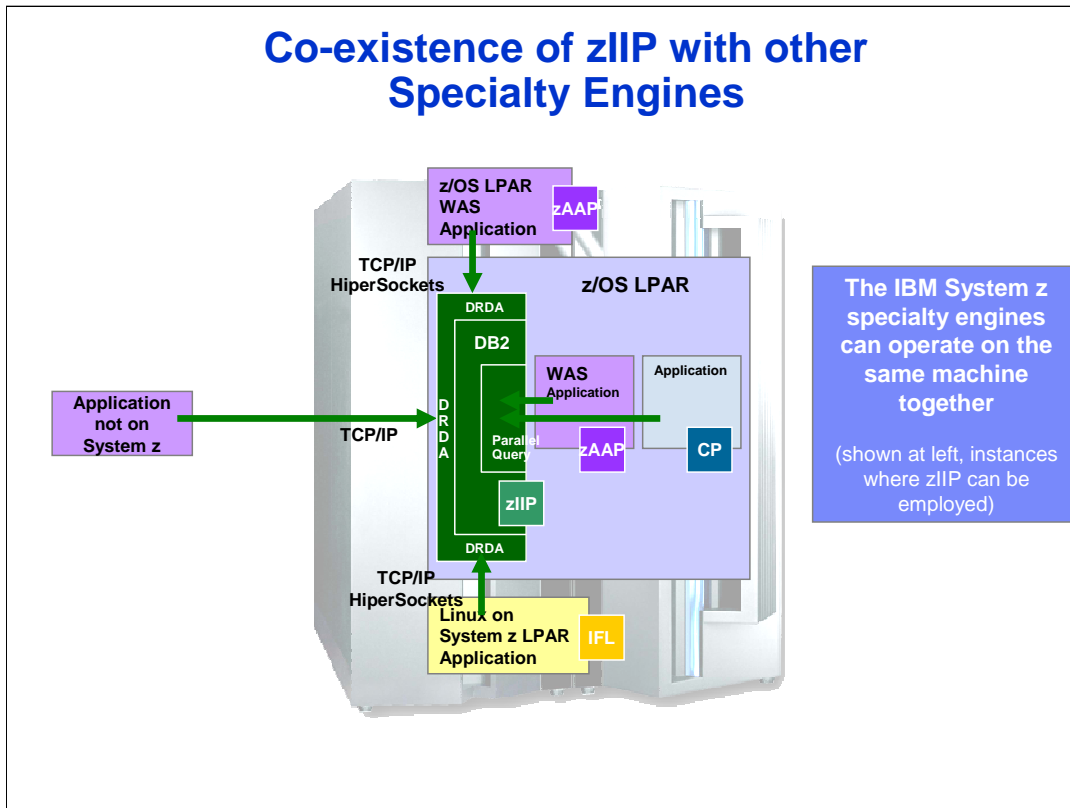
DB2 Utilities zIIP Redirect processing

- Portions of DB2 Utility (LOAD, REORG, & REBUILD) processing related to Index maintenance are eligible to be redirected
- Redirect benefit depends on:
 - How many Indices are defined on the Table
 - How many Partitions are in the Table for Partition Utility
 - Number of Columns, Column data type etc.
- Higher end of range is expected with:
 - Tables with many Indices or many partitions for Partition Utility
- Lower end of range is expected with:
 - Tables with fewer Indices
 - Fewer partitions for Partition Utility
 - Compression used

DB2 Utilities DFSORT zIIP redirect

- **Introduced in Aug 2009 for zIIP redirect for DFSORT processing for some DB2 Utilities**
 - Applicable to in-memory fixed length record sort processing in DFSORT
- **Pre-requisites :**
 - z/OS 1.10
 - DB2 PK85889 (V8, DB2 9)
 - DFSORT PK85856
- **Utilities that benefit :**
 - LOAD, REORG and BUILD Index for Index Sort
 - RUNSTATS for COLGROUP processing

Co-existence of zIIP with other Specialty Engines




This chart shows how zIIP can co-exist with other specialty engines in the same LPAR.

Linux and WebSphere for z/OS applications that use IBM JCC Type 4 driver for SQLJ and JDBC that uses TCP/IP DRDA can benefit from DB2 zIIP redirect in addition to zAAP.

WebSphere for z/OS applications that use IBM JCC Type 2 driver uses RRS attach (not DRDA) to communicate to DB2 and hence is not eligible for zIIP redirect. For some applications this configuration will provide better performance than using IBM JCC Type 4 driver.

Some instances where zIIP would not be exploited



Exceptions : DB2 parallel queries (zIIP) and XML parsing (zAAP) eligible.

One objective with the zIIP is to help bring the costs of network access to DB2 more closely in line with the costs of running similar workloads under CICS®, IMS™, or Batch on the mainframe. Database workloads such as CICS, IMS, WLM managed DB2 stored procedures and batch have become increasingly efficient and cost effective on the mainframe. Today, customers are looking to further leverage their data on the mainframe, and are turning to the mainframe more and more frequently for new application workloads. These application workloads, such as

Enterprise Resource Planning, Customer Relationship Management and Business Intelligence often use DB2 as a database server.

Added system costs for network access to DB2 results in hardware and software charges that are substantially higher for remote applications, when compared to the same workload running as a local batch application. The zIIP is aimed at helping customers more cost effectively run these database serving applications on the mainframe.

CPU intensive parallel queries from CICS, IMS, QMF or Batch are eligible for zIIP redirect.

DB2 Workloads that leverage zAAP

- SQLJ & JDBC applications using IBM JCC Type 2 driver on a z/OS LPAR
 - Connects to DB2 via RRS Attach
- Java based (SQLJ, JDBC) DB2 External Stored Procedures running in a z/OS WLM Application Environment
- z/OS XML System Services called via TCB for parsing DB2 9 for z/OS XML documents

zIIP & zAAP Enablement Process

zIIP & zAAP Software Enablement Process

- **Install z/OS zIIP & zAAP support maintenance**
 - DB2 info Apar II4219, Retain keyword zIIP/K
- **Install DB2 for z/OS support maintenance**
- **Set up SYS1.PARMLIB(IEAOPTxx) member**
 - When zIIP or zAAP hardware is not installed, set PROJECTCPU=YES for projecting zIIP or zAAP redirect
 - No need to be on z9 or z10 Processor for projection
zIIP redirect projection / estimation is shown under APPL% IIPCP / AAPCP in the RMF Workload Activity Report

Recommend taking default values for the following parameters related scheduling algorithms :

 - IIPHONORPRIORITY, ZIIPAWMT for zIIP
 - IFAHONORPRIORITY, ZAAPAWMT for zAAP

Need zIIP z/OS PTFs for z/OS 1.6 and 1.7. zIIP support is in the base code in z/OS 1.8. Need zIIP DB2 PTFs for DB2 z/OS V8. zIIP support is in the base code in DB2 9.

The PROJECTCPU=YES option (also available on z/OS V1R6 and z/OS V1R7 as part of the zIIP FMIDs) now also allows zAAP projection to occur, without requiring any JVM configuration changes. Previously, each impacted JVM had to be individually configured to

cause zAAP statistics to be collected in RMF and SMF.

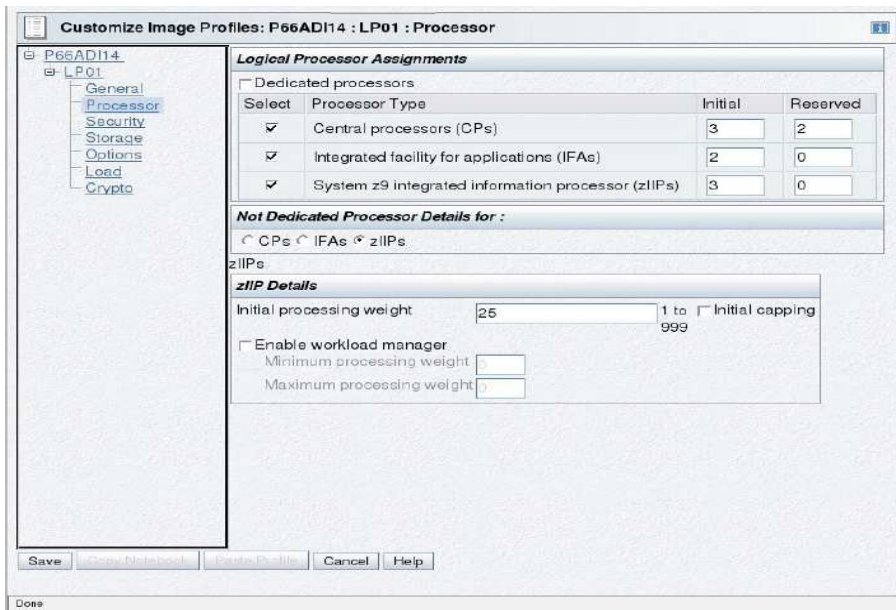
To aid in determining the number of zIIP engines required to satisfy a specific customer's usage, this new parmlib option is available once all the software updates have been applied. The PROJECTCPU=YES parameter enables z/OS to collect zIIP usage as if there was one configured, when the target workload is being run.

This projection capability can be run at any time after the zIIP support software pre-requisites are installed, on a production environment if desired. RMF and IBM Tivoli Omegamon DB2 Performance Expert Accounting reports show this calculated zIIP time so that an accurate zIIP capacity plan can be made. The projection can be done on a non-z9 hardware.

IIPHONORPRIORITY= YES|NO (default is YES). YES Specifies that standard processors run both zIIP processor eligible and non-zIIP processor eligible work in priority order when the zIIP processors indicate the need for help from standard processors. NO Specifies that standard processors will not examine zIIP processor eligible work regardless of the demand for zIIP processors unless contention on a suspend lock requires a zIIP eligible unit of work to be dispatched. Recommend default value is YES.

ZIIPAWMT: default value is 12000 (12 millisecond) , 3200 with Hiperdispatch – time a zIIP eligible work will wait for a available zIIP processor before it will be redirected back to a normal CP. Recommend default value

zIIP & zAAP Hardware Enablement Process



This chart shows the z9 HMC (Hardware Maintenance Console) screen that is used to configure the number of CPs, zAAPs (IFAs) and zIIPs for the LPAR. The processor configuration takes effect after the LPAR is activated and then the LPAR IPL will be performed.

Display CPU information with zIIP & zAAP

```
D M=CPU
IEE174I 10.37.03 DISPLAY
PROCESSOR STATUS
ID CPU SERIAL
00 + 02B29E2094
01 + 02B29E2094
02 +A 02B29E2094
03 +I 02B29E2094

CPC ND = 002094.S28.IBM.02.00000004B29E
CPC SI = 2094.724.IBM.02.000000000004B29E
CPC ID = 00
CPC NAME = SYSS01
LP NAME = STLABH2 LP ID = 2
CSS ID = 0
MIF ID = 2


+ ONLINE - OFFLINE . DOES NOT EXIST W WLM-MANAGED
N NOT AVAILABLE

A APPLICATION ASSIST PROCESSOR (zAAP)
I INTEGRATED INFORMATION PROCESSOR (zIIP)
```

This chart shows the CPU configuration display on the z/OS console with the D M=CPU command.

This display shows a configuration with 2 CPs, 1 zAAP and 1 zIIP processor. + sign indicates that the processor is online.

DRDA Workload Measurement Configuration



Workloads :

- ODBC/CLI SQL
- ODBC/CLI calling Stored Procedures
- IBM JCC T4 Driver for JDBC Parallel Queries
- DB2 Utilities
- XML LOAD and INSERT workload

Monitoring System level zIIP & zAAP Redirect with zIIP & zAAP installed

RMF CPU Report for CLI DRDA Workload :

C P U A C T I V I T Y				SYSTEM ID H2		
z/OS V1R7				RPT VERSION V1R7 RMF		
CPU	2094	MODEL 724	H/W MODEL S28			
---	---	---	---	---	---	---
	CPU---	ONLINE TIME	LPAR BUSY		MVS BUSY	
NUM	TYPE	PERCENTAGE	TIME PERC		TIME PERC	
0	CP	100.00	22.49		22.49	
1	CP	100.00	21.72		21.72	
CP	TOTAL/AVERAGE		22.11		22.11	← CP CPU %
2	AAP	100.00	0.10		0.10	← zAAP CPU %
AAP	AVERAGE		0.10		0.10	
3	IIP	100.00	32.47		32.47	← zIIP CPU %
IIP	AVERAGE		32.47		32.47	

zIIP Redirect % at the LPAR level = 42%

RMF report SYSIN :

```
REPORTS(CPU) for CPU Activity Report
SYSRPTS(WLMGL(SCLASS,RCLASS,POLICY,SYSNAM(XXXX))) for Workload Activity Report
```

Chart shows the CPU utilization of the different processors in the LPAR generated by the RMF batch CPU report. Bottom of the chart shows the RMF batch report control cards to generate the RMF CPU report and the Workload activity reports.

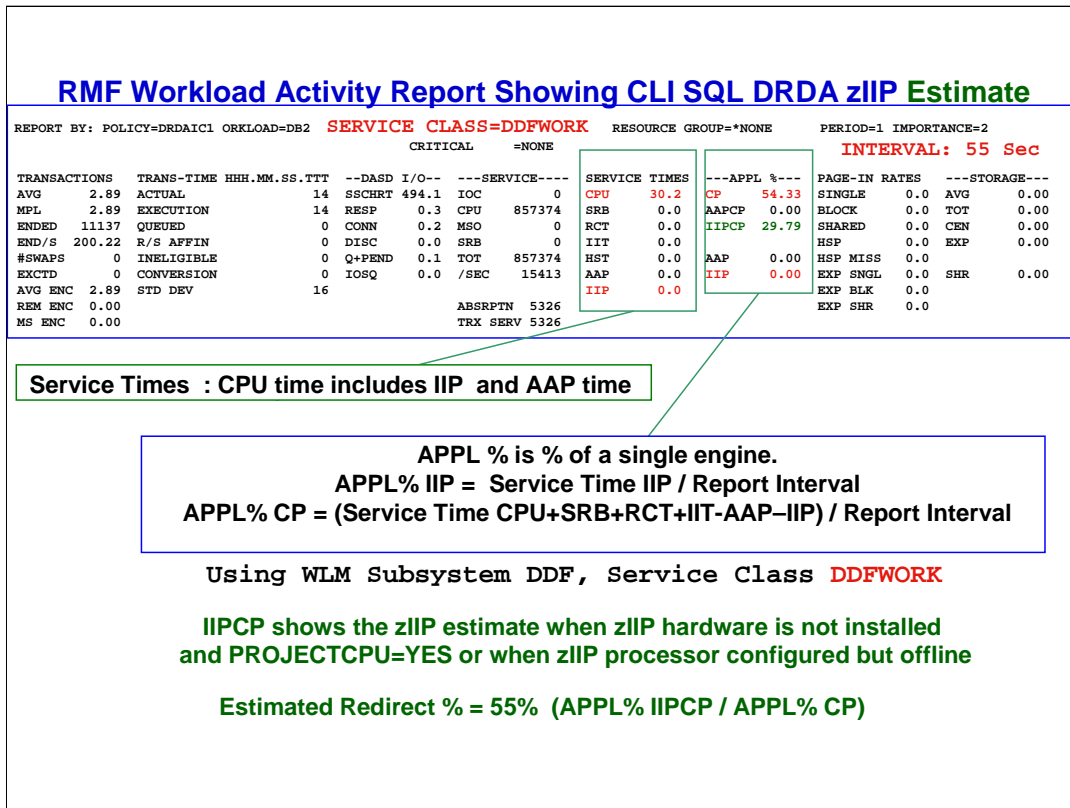
In the RMF workload activity report do a find on 'Service Policy' to look at the zIIP redirect at the WLM Policy level. You can use the SYSNAM(XXXX) in the SYSRPTS control card to get the zIIP redirect for a specific LPAR..

The RMF CPU activity report shows 2 CPs, 1 zAAP and 1 zIIP engines.

The report is for a distributed ODBC/CLI workload showing 42% zIIP redirect at the LPAR level.

This is derived by using the formula : $\text{zIIP CP\%} / \text{Total CP\%} = 32.47 / (22.49 + 21.72 + 0.10 + 32.47)$

**Estimation & Monitoring
of zIIP Redirect for
DRDA Workload**



The chart shows the RMF workload activity report showing the zIIP redirect estimate for the DRDA CLI SQL workload before installing the zIIP hardware and using the SYS1.PARMLIB(IEAOPTxx) PROJECTCPU=YES parameter. SMF Record Type 72 Subtype 3 is used to generate this report.

The WLM policy has been setup for Subsystem DDF and Service Class DDFWORK. WLM Policy can be setup to monitor individual DB2 subsystems by Service Class or Reporting Class and Classification qualifier type SI.

Service Times block shows the information in seconds for the interval. CPU time include IIP and AAP time. IIPCP (zIIP eligible on CP) is not shown under Service Times.

APPL % values are % of a single engine and is calculated from the Service times and dividing by the report interval. APPL% CP includes AAPCP and IIPCP.

Estimated redirect % = APPL% IIPCP / APPL% CP.

Estimate shows 55% redirect for the DB2 DRDA SQL CLI workload.

Tivoli Omegamon DB2PE Accounting Report with CLI SQL DRDA zIIP Redirect Estimate

AVERAGE	CONNTYPE: DRDA APPL (CL.1)	DB2 (CL.2)	
CP CPU TIME	0.002754	0.001726	← Includes IIPCP CPU time. Does not include IIP CPU time.
AGENT	0.002754	0.001726	
NONNESTED	0.002754	0.001726	
STORED PRC	0.000000	0.000000	
UDF	0.000000	0.000000	
TRIGGER	0.000000	0.000000	
PAR.TASKS	0.000000	0.000000	
IIPCP CPU	0.001534	N/A	← zIIP eligible work run on CP
IIP CPU TIME	0.000000	0.000000	← CPU time on zIIP

**IIPCP shows the zIIP estimate when zIIP hardware is not installed
and PRJECTCPU=YES or when zIIP processor is configured but offline**

Estimated Redirect % = 55% (Class 1 IIPCP / CP)

Note: 'IIP' changed to 'SE'(Specialty Engine) with Omegamon DB2PE Apar PK51045

This chart shows how one can estimate the zIIP redirect using the IBM Tivoli Omegamon DB2 Performance Expert accounting report using PRJECTCPU=YES parameter when zIIP hardware is not installed. The estimation uses Class 1 CPU times. Class 1 is used because it includes DDF CPU time and DB2 Class 2 CPU time.

IIP CPU time is the CPU time on zIIP. The IIP CPU time is zero since zIIP was not installed for this measurement.

IIPCP CPU time will show any zIIP eligible processing that ran on CP because zIIP was not installed or when the installed zIIP processors are too busy to handle the scheduled work. IIPCP CPU time is non-zero because zIIP is not installed and PROJECTCPU=YES.

The example shows zIIP redirect estimation of 55% at the DB2 DRDA workload level.

RMF Workload Activity Report Showing CLI SQL DRDA zIIP Redirect

REPORT BY: POLICY=DRDAIC1 WORKLOAD=DB2 SERVICE CLASS=DDFWORK RESOURCE GROUP=*NONE
 INTERVAL: 54 Sec

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE---	SERVICE TIMES	---APPL %---
AVG	2.90	ACTUAL	14	SSCHRT 507.2	IOC	0	CPU 29.3
MPL	2.90	EXECUTION	13	RESP 0.3	CPU	831425	CP 24.02
ENDED	11384	QUEUED	0	CONN 0.2	MSO	0	SRB 0.0
END/S	207.84	R/S AFFIN	0	DISC 0.0	SRB	0	RCT 0.0
#SWAPS	0	INELIGIBLE	0	Q+PEND 0.1	TOT	831425	IIT 0.0
EXCTD	0	CONVERSION	0	IOSQ 0.0	/SEC	15179	HST 0.0
AVG ENC	2.90	STD DEV	15				AAP 0.0
REM ENC	0.00						IIP 29.49
MS ENC	0.00						
					ABSRPTN	5243	
					TRX SERV	5243	

Service Times : CPU time includes IIP and AAP time

APPL % is % of a single engine.
 APPL% IIP = Service Time IIP / Report Interval
 APPL% CP = (Service Time CPU+SRB+RCT+IIT-AAP-IIP) / Report Interval

Using WLM Subsystem DDF, Service Class **DDFWORK**
 Redirect % = Service Time IIP / Service Time CPU
 = APPL% IIP / (APPL% CP+APPL% IIP)
 = 55% for this DRDA workload
 zIIP Redirect % at the LPAR level = 42%

This chart shows the zIIP direct % when zIIP is being used.

It shows the redirect % of 55% at the DRDA workload level using the APPL% formula.

The effective redirect % for this workload at the LPAR level is 42% as shown in chart 27. It is lower at the LPAR level because of the CPU consumed by other non DB2 DRDA components (other DB2 address spaces, TCP/IP etc).

The DRDA redirect % can be calculated using the Service times also.

The formula is : Service Times IIP / Service Times CPU.

Tivoli Omegamon DB2PE Accounting Report with CLI SQL DRDA zIIP Redirect

CONNTYPE: DRDA		
AVERAGE	APPL(CL.1)	DB2 (CL.2)
CP CPU TIME	0.001197	0.000751
AGENT	0.001197	0.000751
NONNESTED	0.001197	0.000751
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	0.000000	0.000000
IIPCP CPU	0.000000	N/A
IIP CPU TIME	0.001480	0.000911

← Chargeable CPU time.
Includes IIPCP CPU time.
Does not include IIP CPU time.

← zIIP eligible work run on CP

← CPU time on zIIP

IIPCP value of zero indicates that 100% of the zIIP eligible work ran on zIIP

$$\begin{aligned}
 \text{Redirect \%} &= \text{Class 1 IIP CPU} / (\text{CP CPU} + \text{IIP CPU}) \\
 &= 55 \% \text{ for this workload}
 \end{aligned}$$

Note: 'IIP' changed to 'SE'(Specialty Engine) with Omegamon DB2PE Apar PK51045

The chart shows the DRDA workload zIIP redirect % using the DB2 Performance Expert accounting report for Connect Type DRDA.

IIP CPU time is the CPU time on zIIP.

IIPCP CPU time will show any zIIP eligible processing that ran on CP because zIIP processor was busy.

High non-zero value will indicate a need to configure more zIIP processors. In this example the zero value for IIPCP CPU indicates that there is no need to configure additional zIIP processors.

DRDA zIIP Redirect Measurement Summary

- **Measured with CLI SQL and Stored Procedure distributed workloads.**
 - CLI workload achieved expected redirect %
 - WLM Managed Stored Procedure achieved 10% redirect
 - Stored Procedure Call, Results set and Commit processing eligible for zIIP redirect.
- **Parallel Query workload achieved expected redirect %**
- **DB2 9 DRDA Native SQL Procedure Language Stored Procedure SQL processing and XML parsing is eligible for zIIP redirect**
- **No noticeable CPU overhead or elapsed time increase for zIIP redirect processing.**
- **Positive feedback from Customers using in production**

This chart summarizes the DB2 DRDA zIIP redirect measurements.

Parallel query via TCP/IP DRDA will get the redirect benefit from DRDA zIIP redirect for the main task and parallel query zIIP redirect for the parallel query child tasks.

Native SQL Stored Procedure calls from TCP/IP DRDA applications will run under enclaves in DBM1 (instead of under WLM) and hence eligible for zIIP redirect.

Several customers are running successfully in production.

**Estimation & Monitoring
of zIIP Redirect for
Parallel Query Workload**

RMF Workload Activity Report Showing Local Parallel Query zIIP Redirect Estimate

REPORT BY: POLICY=DRDAIC1		REPORT CLASS=SSPQ1										
		HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS BATCH_M										
TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---	SERVICE---	---	SERVICE	TIMES--	APPL	%---	
AVG	0.20	ACTUAL	3.57	.786	SSCHRT	0.4	IOC	94	CPU	129.2	CP	10.75
MPL	0.20	EXECUTION	3.56	.910	RESP	8.1	CPU	3559K	SRB	0.0	AAPCP	0.00
ENDED	1	QUEUED	875		CONN	2.9	MSO	0	RCT	0.0	IIPCP	8.46
END/S	0.00	R/S AFFIN	0		DISC	1.0	SRB	23	IIT	0.0		
#SWAPS	1	INELIGIBLE	0		Q+PEND	0.2	TOT	3559K	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0		IOSQ	4.0	/SEC	2961	AAP	0.0	IIP	0.00
AVG ENC	0.00	STD DEV	0						IIP	0.0		
REM ENC	0.00						ABSRPTN	15K				
MS ENC	0.00						TRX SERV	15K				

**Using WLM Subsystem JES, Service Class BATCH_M , Reporting Class SSPQ1
With Classification Qualifier TN for Job Name**

**IIPCP shows the zIIP estimate when zIIP hardware is not installed
and PROJECTCPU=YES or when zIIP processor configured but offline**

Estimated Redirect % = 79% (APPL% IIPCP / APPL% CP)

This chart is showing the zIIP redirect estimate for batch parallel query workload. The WLM policy was setup for Subsystem JES, Service class BATCH_M and reporting class SSPQ1 for the parallel query Job name.

The WLM classification qualifier TN was used for the Job name. Wildcard (NAME%%%%) can be used to group several jobs under one WLM reporting Class.

The example shows zIIP redirect estimate of 79% for the parallel query.

Tivoli Omegamon DB2PE Accounting Report with Local Parallel Query zIIP Estimate

PLANNAME: DSNTPE81		
AVERAGE	APPL (CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	54.689704	54.681809
AGENT	6.774643	6.766781
NONNESTED	6.774643	6.766781
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	47.915061	47.915027
IIPCP CPU	38.242719	N/A
IIP CPU TIME	0.000000	0.000000

← Chargeable CPU time. Includes IIPCP CPU time. Does not include IIP CPU time.

← zIIP eligible work run on CP

← CPU time on zIIP

IIPCP shows the zIIP estimate when zIIP hardware is not installed and PRJECTCPU=YES or when zIIP processor is configured but offline

Estimated Redirect % = 70% (IIPCP / CP)

Note: 'IIP' changed to 'SE'(Specialty Engine) with Omegamon DB2PE Apar PK51045

This chart shows the zIIP redirect estimate for all the parallel queries run under the Plan name DSNTPE81.

The estimated redirect % using the accounting report is 70%/

RMF Workload Activity Report Showing Local Parallel Query zIIP Redirect

REPORT BY: POLICY=DRDAIC1		REPORT CLASS=SSPQ1 HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS BATCH_M									
TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---	SERVICE---	SERVICE	TIMES	---	APPL	%---
AVG	0.19	ACTUAL	3.52.930	SSCHRT	0.4	IOC	94	CPU	129.1	CP	2.23
MPL	0.19	EXECUTION	3.52.074	RESP	8.9	CPU	3556K	SRB	0.0	AAPCP	0.00
ENDED	1	QUEUED	856	CONN	3.1	MSO	0	RCT	0.0	IIPCP	0.01
END/S	0.00	R/S AFFIN	0	DISC	1.5	SRB	28	IIT	0.0		
#SWAPS	1	INELIGIBLE	0	Q+PEND	0.1	TOT	3556K	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0	IOSQ	4.2	/SEC	2845	AAP	0.0	IIP	8.11
AVG ENC	0.00	STD DEV	0					IIP	101.3		
REM ENC	0.00					ABSRPTN	15K				
MS ENC	0.00					TRX SERV	15K				

**Using WLM Subsystem JES, Service Class BATCH_M , Reporting Class SSPQ1
With Classification Qualifier TN for Job Name**

**Redirect % = Service Time IIP / Service Time CPU (more accurate)
= APPL% IIP / (APPL% CP+APPL%IIP)
= 78 % for this Query**

This chart shows the actual parallel query zIIP offload for the WLM reporting class SSPQ1 when the zIIP is installed .

Tivoli Omegamon DB2PE Accounting Report with Local Parallel Query zIIP Redirect

PLANNAME: DSNTEP81		
AVERAGE	APPL(CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	19.373768	19.365788
AGENT	6.779348	6.771411
NONNESTED	6.779348	6.771411
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	12.594420	12.594377
IIPCP CPU	2.813831	N/A
IIP CPU TIME	35.886951	35.886951

← Chargeable CPU time.
Includes IIPCP CPU time.
Does not include IIP CPU time.

← zIIP eligible but ran on CP

← CPU time on zIIP

Total zIIP eligible work % = 70% ((IIP +IIPCP) / (CP+IIP))
zIIP Redirect % = 65% (IIP / (CP+IIP))
zIIP eligible but ran on CP = 5% (IIPCP / (CP+IIP))

Note: 'IIP' changed to 'SE'(Specialty Engine) with Omegamon DB2PE Apar PK51045

This chart shows the actual zIIP direct for all the parallel queries run under the Plan name DSNTEP81 when the zIIP is installed.

Parallel Query zIIP Redirect Measurement Summary

- **Measurement done with local and distributed Star Schema and non Star Schema parallel queries.**
 - Distributed parallel queries benefit from the DRDA zIIP redirect for the Main task as well.
- **No significant increase in Total CPU (CP +zIIP) and elapsed time.**
- **IFCID 231 has been enhanced to provide zIIP related CPU information.**
- **More V8 parallelism potential with parallel sort and parallel multi column merge join.**
- **Increased zIIP redirect potential with Star Join dynamic Index ANDing enhancement in DB2 9.**

This chart summarizes the zIIP measurements for the parallel query.

CPU intensive parallel queries after their parallel group CPU consumptions exceeds certain threshold then subsequent child task processing will be scheduled to run under enclave SRB and a portion of it will be redirected to zIIP.

Parallel queries coming in via TCP/IP DRDA will get the DRDA zIIP redirect for the main task and also the parallel child task zIIP redirect after a certain parallel group CPU threshold is reached.

IFCID 231 has been enhanced to show the CP and IIP CPU time for the parallel child tasks.

**Estimation & Monitoring
of zIIP Redirect for
Utility Workload**

RMF Workload Activity Report Showing Rebuild Index Utility zIIP Redirect Estimate

REPORT BY: POLICY=DRDAIC1

REPORT CLASS=RBLDINDX

DESCRIPTION =DB2 REBUILD INDEX

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE---	SERVICE	TIMES	---	APPL	%---	
AVG	0.17	ACTUAL	3.29.961	SSCHRT	312.3	IOC	176	CPU	82.3	CP	17.44
MPL	0.17	EXECUTION	1.18.230	RESP	0.3	CPU	2267K	SRB	0.0	AAPCP	0.00
ENDED	1	QUEUED	2.11.731	CONN	0.2	MSO	0	RCT	0.0	IIPCP	4.56
END/S	0.00	R/S AFFIN	0	DISC	0.0	SRB	50	IIT	0.0		
#SWAPS	1	INELIGIBLE	0	Q+PEND	0.1	TOT	2267K	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	4804	AAP	0.0	IIP	0.00
AVG ENC	0.00	STD DEV	0					IIP	0.0		
REM ENC	0.00					ABSRPTN	29K				
MS ENC	0.00					TRX SERV	29K				

Using WLM Subsystem JES, Service Class BATCH_M , Reporting Class RBLDINDX
With Classification Qualifier TN for Job Name

IIPCP shows the zIIP estimate when zIIP hardware is not installed
and PRJECTCPU=YES or when zIIP processor configured but offline

Estimated Redirect % = APPL% IIPCP / APPL% CP
= 26%

This chart shows the zIIP redirect estimate formula for the DB2 Rebuild Utility when the zIIP processor is not installed.

The WLM policy has been setup for Subsystem JES, Service Class BATCH_M and reporting class RBLDINDX for the DB2 Rebuild Utilities.

The WLM Classification qualifier TN was used for the job name. Wild card (RBLD%%) could be used to group several jobs under one WLM Reporting Class.

Tivoli Omegamon DB2PE Accounting Report with Utility Workload zIIP Estimate

PLANNAME:DSNUTIL or CONNTYPE:UTILITY		
AVERAGE	APPL(CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	1:03.92512	31.245707
AGENT	14.005918	11.460791
NONNESTED	14.005918	11.460791
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	49.919203	19.784917
IIPCP CPU	16.045606	N/A
IIP CPU TIME	0.000000	0.000000

← Chargeable CPU time.
Includes IIPCP CPU time.
Does not include IIP CPU time.

← zIIP eligible work run on CP

← CPU time on zIIP

IIPCP shows the zIIP estimate when zIIP hardware is not installed and PROJECTCPU=YES or when zIIP processor is configured but offline

Estimated Redirect % = 25% (IIPCP / CP)

Note: 'IIP' changed to 'SE'(Specialty Engine) with Omegamon DB2PE Apar PK51045

This chart shows the zIIP redirect estimate for all the DB2 Utilities that were run under Plan name DSNUTIL or Connect type Utility, when the zIIP hardware is not installed.

DB2 Class 1 IIPCP and CP CPU time is used to estimate the zIIP redirect.

RMF Workload Activity Report Showing Rebuild Index Utility zIIP Redirect

REPORT BY: POLICY=DRDAIC1		REPORT CLASS=RBLDINDX									
										HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS BATCH_M	
TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---	SERVICE----	SERVICE	TIMES	---	APPL %---	
AVG	0.17	ACTUAL	3.01.033	SSCHRT	357.0	IOC	178	CPU	81.5	CP	15.84
MPL	0.17	EXECUTION	1.08.519	RESP	0.3	CPU	2313K	SRB	0.0	AAPCP	0.00
ENDED	1	QUEUED	1.52.514	CONN	0.2	MSO	0	RCT	0.0	IIPCP	1.47
END/S	0.00	R/S AFFIN	0	DISC	0.0	SRB	51	IIT	0.0		
#SWAPS	1	INELIGIBLE	0	Q+PEND	0.1	TOT	2313K	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	5603	AAP	0.0	IIP	3.91
AVG ENC	0.00	STD DEV	0					IIP	16.1		
REM ENC	0.00					ABSRPTN	34K				
MS ENC	0.00					TRX SERV	34K				

**Using WLM Subsystem JES, Service Class BATCH_M , Reporting Class RBLDINDX
With Classification Qualifier TN for Job Name**

**Redirect % = Service Time IIP / Service Time CPU (Accurate)
= APPL% IIP / (APPL% CP+APPL%IIP)
= 20 % for this Rebuild Index Utility**

This chart shows the actual zIIP redirect % for the DB2 Rebuild Index Utilities under the WLM reporting class RBLDINDX.

Tivoli Omegamon DB2PE Accounting Report for Utility Workload zIIP Redirect

PLANNAME: DSNUTIL or CONNTYPE: UTILITY		
AVERAGE	APPL(CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	52.070150	19.363503
AGENT	13.315781	10.777834
NONNESTED	13.315781	10.777834
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	38.754370	8.585669
IIPCP CPU	3.808629	N/A
IIP CPU TIME	12.759936	12.759936

← **Chargeable CPU time.**
Includes IIPCP CPU time.
Does not include IIP CPU time.

← **zIIP eligible but ran on CP**

← **CPU time on zIIP**

Total zIIP eligible work % = 26% ((IIP + IIPCP) / (CP + IIP))
zIIP Redirect % = 20% (IIP / (CP + IIP))
zIIP eligible but ran on CP = 6% (IIPCP / (CP + IIP))

Note: 'IIP' changed to 'SE' (Specialty Engine) with Omegamon DB2PE Apar PK51045

This chart shows the actual zIIP redirect for all the DB2 Utilities that were run (Plan name DSNUTIL or Connect Type UTILITY) when the zIIP processor is installed.

Notice non zero value for IIPCP CPU indicating zIIP was processor was unable to process all the zIIP eligible work.

Utility zIIP Redirect Measurement Summary

- **Measured LOAD, REBUILD INDEX and REORG Utilities.**
- **zIIP redirect % depends on % CPU consumed by the Build Index phase of the Utility.**
- **Observed Class 1 CPU reduction for configuration with 4 CPs and 2 zIIPs with fixed length Index key :**
 - 5 to 20% for Rebuild Index
 - 10 to 20% for Load or Reorg of a Partition with one Index only, or Load of entire Table, or Reorg of entire Tablespace
 - 40% for Rebuild Index of logical Partition of Non Partitioning Index
 - 40 to 50% for Reorg Index
 - 30 to 60% for Load or Reorg of a Partition with more than one Index
- **CPU overhead incurred during execution unit switch from TCB to enclave SRB during Index Rebuild phase**
 - Typically less than 10%
 - Eligible for zIIP redirect

This chart summarizes the zIIP redirect measurements for the DB2 Utilities. The zIIP redirect % is proportional to amount of build index processing. CPU cost associated with SORT and Compression are not eligible for zIIP redirect.

To support zIIP redirect the build index processing was changed from TCB processing to enclave SRB processing which introduces execution unit switch CPU overhead. The zIIP redirect % shown in this chart have been adjusted for the CPU overhead. The measurement were done with a Tablespace with 10 partitions and varying number of indices up to 6.

Utility zIIP Redirect with DFSORT

- **Introduced in Aug 2009 for zIIP redirect for DFSORT processing for some DB2 Utilities**
 - Applicable to in-memory fixed length record sort processing in DFSORT
- **Utilities that benefit :**
 - LOAD, REORG and BUILD Index for Index Sort
 - RUNSTATS for COLGROUP processing
- **Measured zIIP redirect benefit**
 - 30% to 60% of DFSORT CPU
 - 10% to 40% of total Utility CPU
 - Varies with number of Indices
 - More benefit with more Indices
 - Measurement with up to 6 Indices

DB2 XML related zIIP,zAAP enhancements

- **z/OS XML System Services using zAAP (Sep 2007- OA20308)**
 - Supported in z/OS 1.7
 - Will benefit DB2 9 NFM XML Parsing
 - XML Applications (Insert, Update)
 - XML LOAD Utility
 - zAAP redirect for TCB processing
- **z/OS XML System Services zIIP redirect for enclave SRB processing (July 2008 – OA23828)**
- Whitepaper on DB2 9 use of XML System Services :
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101088>

IBM has announced the following additional zIIP and zAAP enablement in April 2007

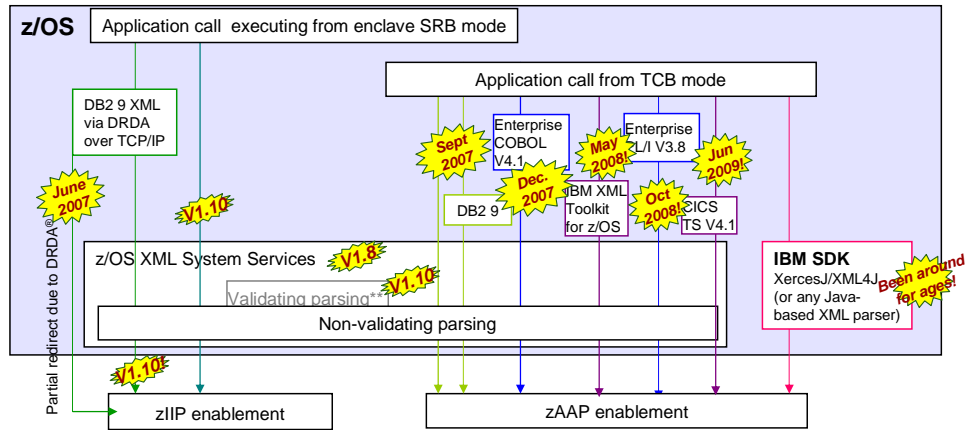
- z/OS 1.8 Communication Server takes advantage the zIIP processor for the IPSec network end to end encryption security protocol processing. IPSec AH (Authentication Header) and ESP (Encapsulation Security Payload) protocol routine CPU processing will be eligible for zIIP redirect.(August 2007).

-z/OS XML Services for Parsing : Available on z/OS 1.7, 1.8 and 1.9. When DB2 for z/OS 9 applications do inserts or updates XML data, DB2 invokes the z/OS XML System Services to parse the data. DB2 for z/OS 9 XML data is loaded by Utilities, the data needs to be parsed and it invokes z/OS XML System Services to parse the data.. zAAP redirect will be used for TCB processing.

-Use of zIIP for 100% of the enclave SRB processing for XML toolkit and system services

z/OS XML System Services zIIP or zAAP eligibility summary

ALL validating and non validating parsing performed by z/OS XML System Services is eligible for zAAP or zIIP!



- How much work is eligible for the zAAP will depend on amount of XML data being processed.
- ** XML System Services validation parsing – first IBM exploiter is the IM XML Toolkit for z/OS, V1.10

We've been rolling out z/OS XML system service eligibility for zIIP and zAAP over time and with all the exploiters and bits and pieces it can get confusing... so, in general, what XML workloads can be eligible for zIIP or zAAP ... and the answer is ALL validating and non validating parsing performed by z/OS XML System Services is eligible for zAAP or zIIP

z/OS XML System Services (z/OS XML) High speed low latency XML parser in z/OS base (1.8) we've had XML parsing on z/OS for quite some time .. there is XML parsing in COBOL and PL1, and C++ but this new parser is intended to provide improved performance and parsing services for the platform.

Lets step thru the exploiters....

The first exploiter was DB2 9 remotely attached ... enclave SRBs and in reality this was due to the fact that it was DRDA and the amount of eligibility was the same as DRDA and you really did not have to do any XML magic to get this.

Next (Sept 2007 – with the availability of z/OS V1.9... we had all z/OS XML system services parsing executing locally or from TCB mode was eligible for zAAP and the first exploiter of that was DB 9 (those are those next 2 green arrows).

Later in 2007 in December Enterprise COBOL also announced z/OS XML system services parsing eligibility (thats the blue arrow)

XML System Services Parsing Measurements Summary

- **DB2 9 measurements with LOAD of XML Tables**
 - 17 to 36% of XML System Services Parsing processing was redirected to zAAP
 - Higher zAAP redirect with
 - Larger doc size, # of nodes
 - Fewer indices
- **zIIP redirect for DRDA TCP/IP XML Insert, Update applications**
 - Depends on the degree of XML parsing
 - Redirect is in addition to the base DRDA redirect
 - 63% zIIP redirect for Lab XML INSERT workload

New Tivoli Omegamon DB2 PE Accounting Report Layout Showing XML LOAD zAAP Redirect

AVERAGE	APPL(CL.1)	DB2 (CL.2)
ELAPSED TIME	11:36.6837	11:27.1818
NONNESTED	11:36.6837	11:27.1818
STORED PROC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
CP CPU TIME	3:29.67361	3:23.08585
AGENT	3:29.53441	3:23.08558
NONNESTED	3:29.53441	3:23.08558
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	0.139205	0.000270
SECP CPU	0.000079	N/A
SE CPU TIME	57.497246	57.497246
NONNESTED	57.497246	57.497246
STORED PROC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	0.000000	0.000000
SUSPEND TIME	0.000000	5:43.35838
AGENT	N/A	5:43.35838
PAR.TASKS	N/A	0.000000
STORED PROC	0.000000	N/A
UDF	0.000000	N/A
NOT ACCOUNT.	N/A	1:23.24058

Changes introduced by Apars :

DB2 9 : PK50575
Omegamon DB2PE : PK51045

←-- New - Replaces IIPCP CPU
←-- New - Replaces IIP CPU Time

New report fields :

SE CPU : Includes both zIIP & zAAP
CPU usage

SECP : zIIP eligible work that ran on
CP due zIIP overflow or with
PROJECTCPU=YES

Does not show similar zAAP
information - use to RMF report
AAPCP info.

RMF Workload Activity Report for XML LOAD Utility

REPORT BY: POLICY=POL_XML WORKLOAD=BATCH SERVICE CLASS=BATCHMED RESOURCE GROUP=*NONE PERIOD=1 IMPORTANCE=3
 CRITICAL =NONE

-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---SERVICE---	--SERVICE TIMES--	---APPL %---	-----STORAGE-----
AVG	6.50	ACTUAL	9.50.231	SSCHRT 665.2	IOC 884057 CPU 1448.496	CP 188.65	AVG 3751.48
MPL	6.50	EXECUTION	8.46.172	RESP 0.3	CPU 398897K SRB 0.862	AAPCP 0.04	TOTAL 24368.79
ENDED	4	QUEUED	1.04.059	CONN 0.2	MSO 0 RCT 0.000	IIPCP 0.00	SHARED 0.00
END/S	0.01	R/S AFFIN	0	DISC 0.0	SRB 237449 IIT 0.155		
#SWAPS	0	INELIGIBLE	0	Q+PEND 0.1	TOT 400019K HST 0.000	AAP 53.16	--PAGE-IN RATES--
EXCTD	0	CONVERSION	0	IOSQ 0.0	/SEC 667455 AAP 318.578	IIP 0.05	SINGLE 0.0
AVG ENC	0.00	STD DEV	0		ABSRPTN 103K IIP 0.329		BLOCK 0.0
REM ENC	0.00				TRX SERV 103K PROMOTED 0.000		SHARED 0.0
MS ENC	0.00						HSP 0.0

GOAL: EXECUTION VELOCITY 40.0% VELOCITY MIGRATION: I/O MGMT 73.9% INIT MGMT 73.9%

XML LOAD job was run under Service Class BATCHMED.

zAAP redirection is 22% (= 53.16/(188.65+53.16)*100), (APPL% AAP / (CP + AAP))

Similar to the Omegamon DB2PE redirection % (SE CPU / (CP CPU + SE CPU))

AAPCP of 0.04 is showing that there was a small overflow from zAAP

Non-DB2 related zIIP zAAP Usage

- **z/OS 1.8 CS TCP/IP IPsec VPN protocol using zIIP (Aug 2007)**
 - End to end network encryption security Protocol

- **z/OS Communication Server zIIP redirect for HiperSockets multiple write outbound z/OS large TCP/IP messages (Aug 2008)**
 - Will benefit workloads with large messages based on XML, HTTP, SOAP, Java and large file transfers
 - Available on z/OS 1.10 and z10

- **zIIP enablement for data extract portion of Scalable Architecture for Financial Reporting (SAFR) performance engine (Aug 2008)**
 - IBM Global Business Services Solution
 - SAFR provides high volume Business Intelligence reports
 - Available on z/OS 1.8 , z9 and z10

Non-DB2 related zIIP zAAP Usage

- **zIIP assisted z/OS Global Mirror (XRC Extended Remote Copy) – Feb 2008**
 - z/OS 1.8, 1.9 (OA23174) DFSMS allows a portion of the System Data Mover (SDM) processing to be eligible for zIIP

- **z/OS 1.11 CIM(Common Information Model)Server (z/OS 1.11)**
 - CIM is an industry standard for defining and sharing information between devices and programs (z/OS 1.7)
 - Client application requests CIM Servers to return information on z/OS resources such as RMF, WLM, DFSMS and BCP
 - CIM Server processing enabled to use zIIP
 - Java based CIM Client applications on System z use zAAP.

Summary

- **Easy implementation**
 - No DB2 application change or tuning options
- **Additional DB2 9 zIIP redirect capabilities with :**
 - DRDA Native SQL Procedure workload
 - DRDA SQL/XML Parsing
 - Enhancements to Star Join parallel queries using dynamic Index ANDing
- **Additional DB2 9 zAAP redirect with XML Parsing**
- zIIP & zAAP can be leveraged to grow or develop or port new distributed and business intelligence and XML applications on DB2 for z/OS in a cost effective way.
 - Frees up general purpose CP cycles for other workload processing
- Reference Information:
<http://www.ibm.com/systems/z/specialtyengines/>

Disclaimer

© Copyright IBM Corporation [current year]. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, and DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Thanks.