

IBM Information

>>> On Demand

2006



Leveraging zIIP with DB2 for z/OS V8

Gopal Krishnan
DB2 for z/OS Performance
IBM Silicon Valley Lab,
San Jose, CA

Session 2218



TAKE BACK CONTROL

IBM INFORMATION ON DEMAND 2006
October 15 - 20, 2006
Anaheim Convention Center
Anaheim, California

Trademarks

AIX*	IBM eServer	z/VM*
CICS*	IBM logo*	zSeries*
DB2*	IMS	
DB2 Connect	On Demand Business logo	
DB2 Universal Database	Parallel Sysplex*	
DRDA*	System z	
FICON*	System z9	
GDPS*	WebSphere*	
HiperSockets	z/Architecture	
IBM*	z/OS*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of the Intel Corporation in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

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

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

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.

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

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. 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.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.



IBM INFORMATION ON DEMAND 2006

TAKE BACK CONTROL

Disclaimer

The information in this document has not been submitted to any formal IBM review and is distributed on an "as is" basis without any warranty expressed or implied. Use of this information or the implementation of any of these techniques is a user responsibility and depends on the user's ability to evaluate and integrate them into the user's operational environment. While each item may have been reviewed for accuracy in a specific situation there is no guarantee the same or similar results may be achieved elsewhere.



TOPICS

- zIIP Overview
- DB2 V8 workloads that leverage zIIP
- Monitoring zIIP & Estimation of zIIP redirect with DB2 V8
- Estimation of zIIP redirect prior to DB2 V8



zIIP

(z9 Integrated Information Processor)

Overview



IBM System z9, z/OS & DB2 for z/OS

- System z9 Integrated Information Processor (zIIP)
Enterprise Class & Business Class
- Enhanced Cryptography
- Channels (4 Gb & MIDAW)
- Faster Processors
- Up to 54 Processors EC
- More memory, better value;
64 bit virtual storage
- z/Architecture new instructions
- Parallel Sysplex
- IPv6
- SSL
- Java
- Decimal float
- Backup & restore
- Multilevel Security
- Unicode collation
- Compression
- System z Application Assist Processor (zAAP)
- WLM enhanced ...




Technology Evolution with Mainframe Specialty Engines

★ Building on a strong track record of technology innovation with specialty engines, IBM intends to introduce the System z9 Integrated Information Processor



IBM System z9 Integrated Information Processor (IBM zIIP) 2006




System z9 Application Assist Processor (zAAP) 2004



Integrated Facility for Linux® (IFL) 2001

Centralized data sharing across mainframes




Internal Coupling Facility (ICF) 1997

Support for new workloads and open standards

Incorporation of Java™ into existing mainframe solutions

Designed to help improve resource optimization for eligible data workloads within the enterprise



New IBM System z9 Integrated Information Processor (IBM zIIP)

- New specialty engine for the System z9 mainframe designed to help:
 - Customers integrate data across the enterprise
 - Improve resource optimization and lower the cost of ownership for eligible data serving workloads
- z/OS manages and directs work between the general purpose processor and the zIIP
 - Number of zIIPs per z9-109 not to exceed number of standard processors
 - No changes to DB2 for z/OS V8 applications
 - Lower price for zIIP engines
 - No IBM software charges on the zIIP – consistent with other specialty engines
 - Software cost benefit depends on peak workload characteristics and software license charge model
- DB2 for z/OS V8 will be first IBM exploiter of the zIIP with
 - System z9 EC or BC
 - z/OS 1.6 or later (with PTFs)
 - DB2 for z/OS V8 (with PTFs)



DB2 for z/OS V8 Workloads Leveraging zIIP



DB2 Workloads that leverage zIIP

- Portions of the following DB2 for z/OS V8 workloads may benefit from zIIP*
 1. ERP, CRM, Business Intelligence or other enterprise distributed applications
 - Via DRDA® over a TCP/IP connection



New Specialty Engine

2. Data warehousing applications*
 - Requests that utilize parallel queries including Star Schema
3. DB2 for z/OS V8 utilities*
 - Internal DB2 utility functions used to maintain index maintenance structures

How does the zIIP work

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 types of DB2 V8 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.

Example 2 = Complex parallel queries (BI)

Some complex 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

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



Example 1: Enterprise Distributed Applications

Enterprise Applications that access DB2 for z/OS V8 via DRDA 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



What is DRDA?

- DRDA = Distributed Relational Database Architecture
 - Developed by IBM
 - Enables relational data to be distributed among multiple platforms – ‘any app to any db and any db to any db’.
Applications and APIs accomplish the actual implementation



- DRDA is native to DB2 for z/OS. It reduces the need for additional gateway products that may affect performance and availability
- The Open Group adopted DRDA in 1998 as the open standard for database access interoperability
- An application uses DRDA application requestor or server to access a remote database (Examples of DRDA application requestors: DB2 Connect, DB2 T4 Universal Driver)
- Uses TCP/IP or SNA as a network protocol to flow DRDA commands

So.... regarding the zIIP: if DB2 for z/OS V8 workload comes over TCP/IP and is DRDA compliant, a portion of that DB2 workload is eligible to be redirected to the zIIP.



What are enclave SRBs?

- z/OS dispatches DB2 work in either TCB (Task Control Block) mode if request is local or SRB (Service Request Block) mode if request is distributed. Under these modes of operation the parallel tasks are assigned the same importance as the originating address space.
- Preemptable enclaves are used to do the work on behalf of the originating TCB or SRB address space. Enclaves are grouped by common characteristics and service requests and since they are preemptable, the z/OS dispatcher (and WLM) can interrupt these tasks for more important ones (ie manage a transaction end-to-end). There are two types of preemptable SRBs: client SRBs and enclave SRBs.
- If the DB2 for z/OS V8 request is coming in over distributed (ie DRDA over TCP/IP) then that work is executed in enclave SRBs.
- If the request is coming over local/ native connection, then that work is dispatched between TCBs, client SRBs, and enclave SRBs (parallel queries and index maintenance now use enclave SRBs)

- So..... regarding the zIIP, only the enclave SRB work is eligible to be redirected to the zIIP.
- DB2 V8 knows how its work is dispatched and directs z/OS to dispatch (redirect) a portion of the eligible work to 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.
- **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.**



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



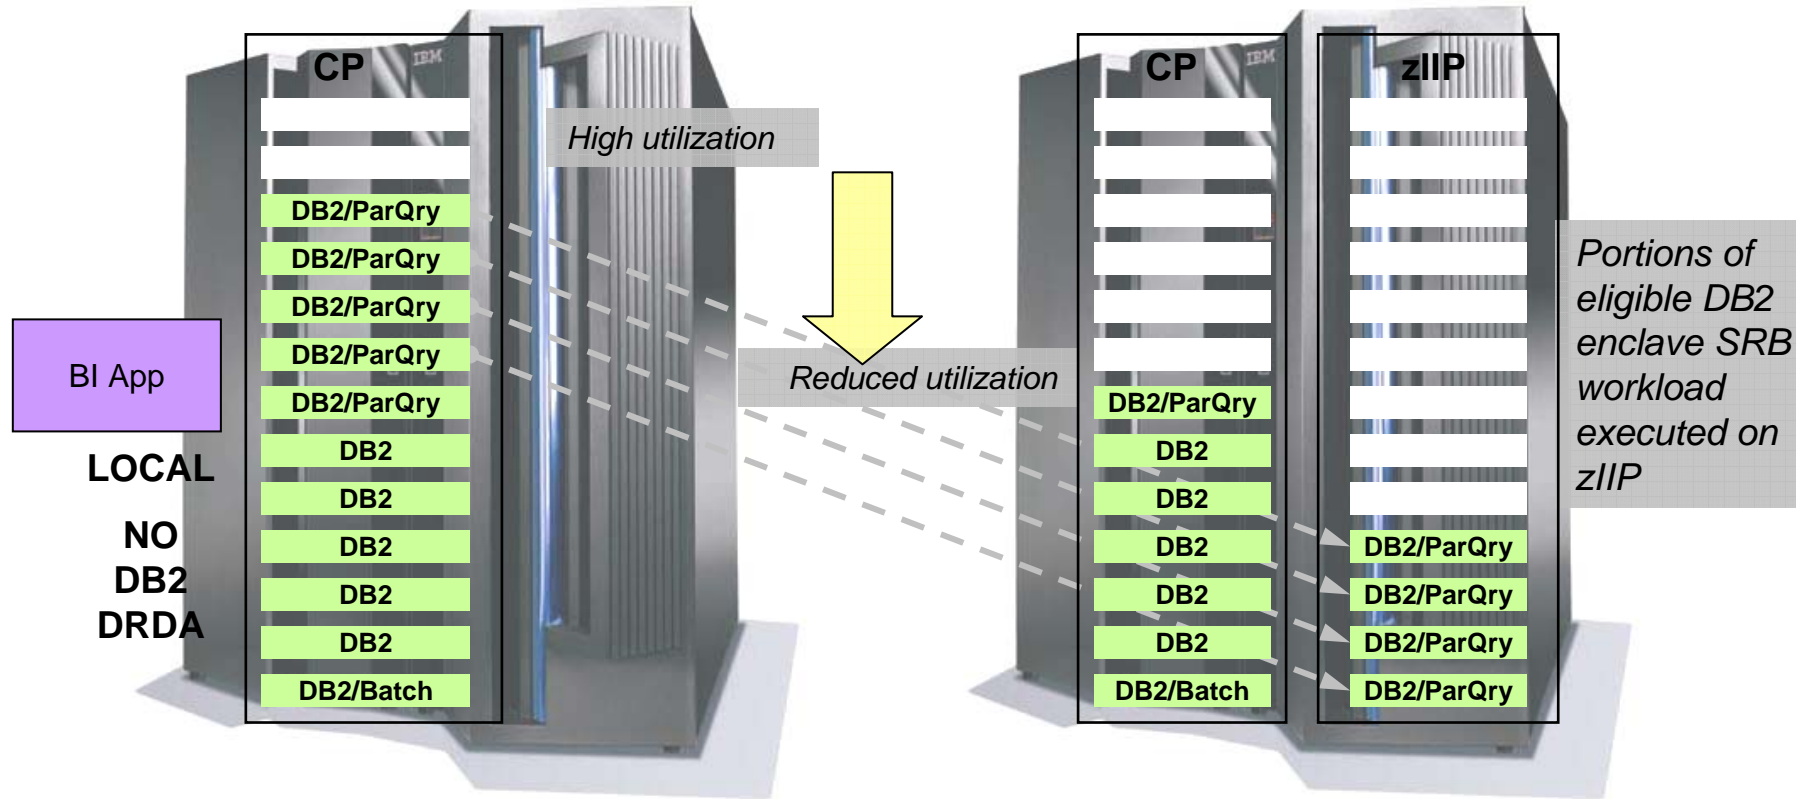
For illustrative purposes only

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



Example 2.5: Business Intelligence Applications (local - no DRDA)

Parallel queries via LOCAL 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 and how much parallelism is used



What are CPU Parallel Queries?

- Complex queries that Optimizer determines that they can benefit from CPU parallelism
 - The Query is run with several Child tasks under a Parent task
- Star Schema queries with a database schema for representing multidimensional data.
 - Sometimes graphically represented as a 'star'
 - Data stored in a central fact table
 - Surrounded by additional dimension tables holding information on each perspective of the data



- **So.... regarding zIIP:** if the work load uses DB2 for z/OS V8 complex parallel queries using Star Schema and non Star Schema, then a portion of that DB2 workload may be eligible to be redirected to the zIIP




Parallel Query zIIP Redirect Processing

- **Applicable to complex 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
 - Main tasks coming in via DRDA via TCP/IP may benefit from the DRDA zIIP redirect
- **The combined child & main tasks coming in through DRDA via TCP/IP is expected to yield additional zIIP redirect.**
- **Benefits queries consuming high CPU.**



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.



What is index maintenance?

- An index allows quick access to the rows in a table. Indices are created using one or more columns of a table.
- Over time, as data in a large database is manipulated indices can become less efficient. They need to be periodically maintained.
 - LOAD -- Loads Table
 - REORG – Reorganizes Index or Tablespace
 - REBUILD INDEX -- Creates or Rebuilds Index

- **So.... regarding the zIIP:** The BUILD phase of LOAD, REORG, and REBUILD Utilities performs Index maintenance. Most of the BUILD phase is eligible to be redirected to the zIIP.



DB2 Utilities zIIP Redirect processing

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

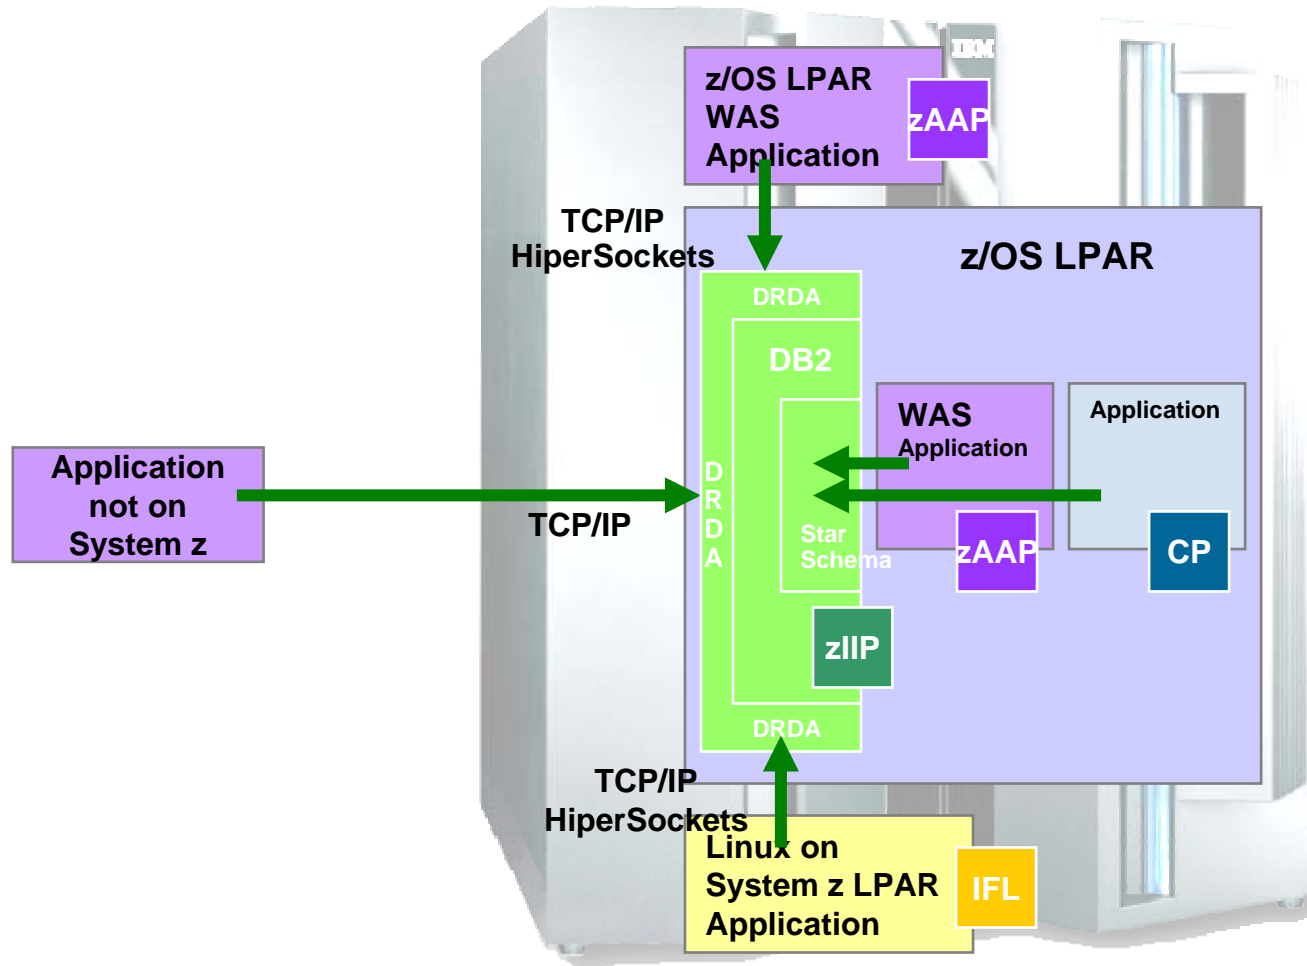


zIIP Redirect Software Maintenance

- DB2 V8 zIIP support :
 - PK18454 DRDA
 - PK19920, PK27712, PK30087 Utilities
 - PK19921, PK27578 Parallel Queries
 - II14219 Info Apar with additional information
- Support for zIIP related instrumentation changes in IFCIDs 3,147,148,231,239 for zIIP usage reporting, monitoring and projection:
 - PK25395 IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS
 - PK29966 DB2 Performance Expert V210
 - PK29967 DB2 Performance Monitor V810
- zIIP support maintenance info for z/OS, SDSF,RMF,WLM,BCP
 - <http://www.ibm.com/systems/z/ziip/gettingstarted/prereqs.html>
- RETAIN search keyword zIIP/K for zIIP related Apar/PTF information.



Co-existing with other Specialty Engines

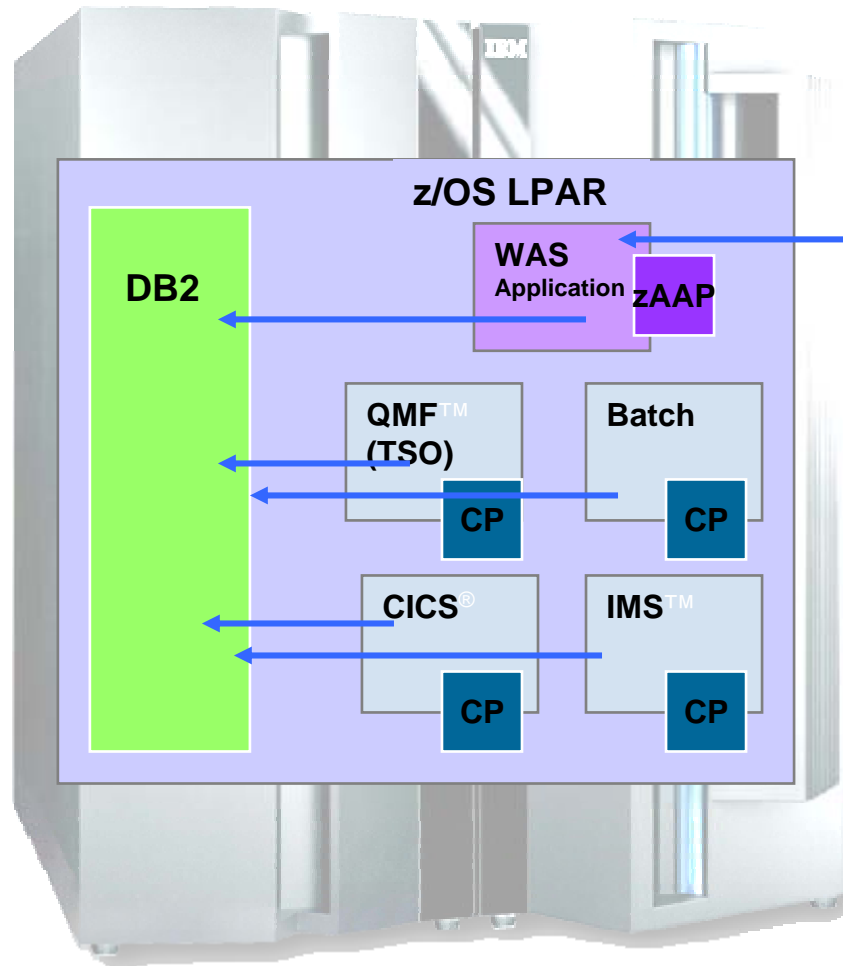


The IBM System z9 specialty engines can operate on the same machine together

(shown at left, instances where zIIP can be employed)



Some instances where zIIP would not be exploited



WAS applications may access with a local Type2 JDBC access. For performance reasons, this is the recommended scenario as the WAS code was optimized for this configuration. No zIIP would be needed here.



Monitoring & Estimation of zIIP Redirect with DB2 V8



SYS1.PARMLIB(IEAOPTxx) Parameters

PROJECTCPU = YES / NO

- **Set to YES for projecting zIIP redirect when zIIP is not installed**
 - **Shown under APPL% IIPCP in the RMF Workload Activity Report**
 - **Shown under IIPCP CPU in the Omegamon DB2PE Accounting Report**

ZIIPAWMT = 12000 (default value 12 ms)

- **Use the default value**
- **Alternate Weight Management – time spent waking up idle zIIP to compete for work or request for help from CP**



z9 Display CPU information with zIIP

```
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)
```



Monitoring System level zIIP redirect with zIIP installed

RMF CPU Report for CLI DRDA Workload :

C P U A C T I V I T Y				z/OS V1R7		SYSTEM ID H2	
				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		
AAP	AVERAGE		0.10		0.10	←	zAAP CPU %
3	IIP	100.00	32.47		32.47		
IIP	AVERAGE		32.47		32.47	←	zIIP CPU %

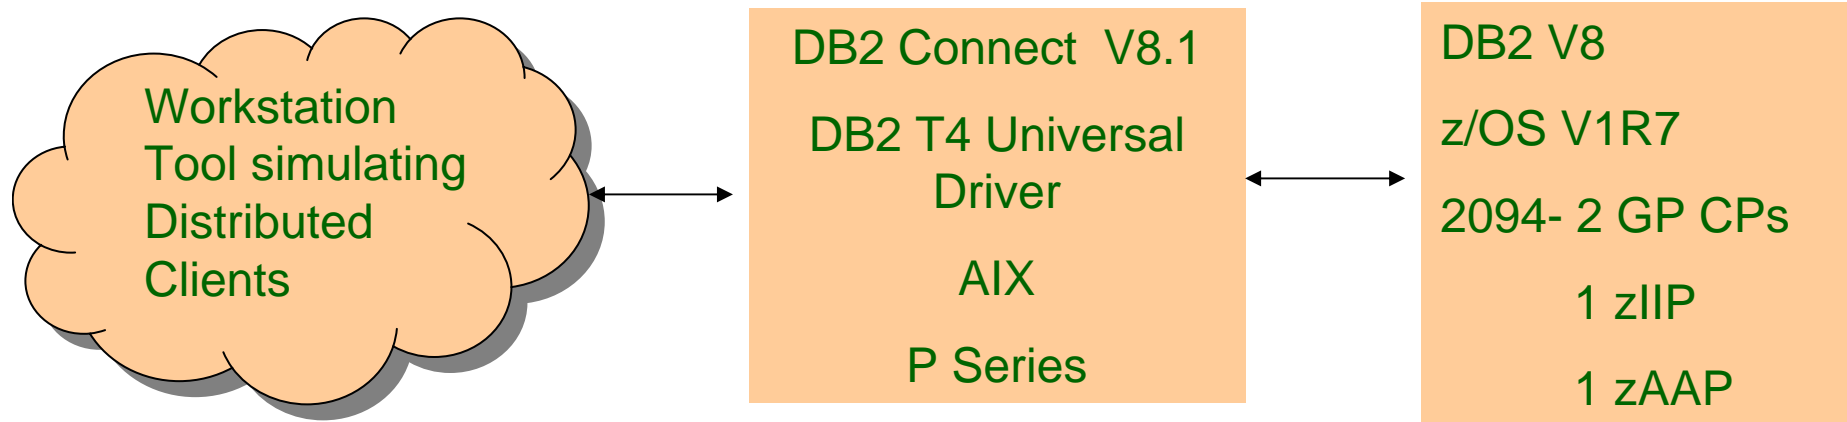
zIIP Redirect % at the LPAR level = 42%



Monitoring & Estimation of zIIP Redirect for **DRDA Workload** with DB2 V8



DRDA Workload Measurement Configuration



Workloads :

- ODBC/CLI SQL
- ODBC/CLI calling Stored Procedures
- T4 Driver JDBC Parallel Queries



RMF Workload Activity Report

Showing CLI SQL DRDA zIIP Redirect

REPORT BY: POLICY=DRDAIC1 WORKLOAD=DB2 **SERVICE CLASS=DDFWORK** RESOURCE GROUP=*NONE
 CRITICAL =NONE

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE----	SERVICE TIMES	---APPL %---
AVG	2.90	ACTUAL	14	SSCHRT 507.2	IOC	0	CP 24.02
MPL	2.90	EXECUTION	13	RESP 0.3	CPU 831425	SRB 0.0	AAPCP 0.00
ENDED	11384	QUEUED	0	CONN 0.2	MSO 0	RCT 0.0	IIPCP 0.00
END/S	207.84	R/S AFFIN	0	DISC 0.0	SRB 0	IIT 0.0	AAP 0.00
#SWAPS	0	INELIGIBLE	0	Q+PEND 0.1	TOT 831425	HST 0.0	IIP 29.49
EXCTD	0	CONVERSION	0	IOSQ 0.0	/SEC 15179	AAP 0.0	
AVG ENC	2.90	STD DEV	15			IIP 16.2	
REM ENC	0.00				ABSRPTN 5243		
MS ENC	0.00				TRX SERV 5243		

Service Times : CPU time includes IIP time

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

Using WLM 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%



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

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



RMF Workload Activity Report

Showing CLI SQL DRDA zIIP Estimate

REPORT BY: POLICY=DRDAIC1 ORKLOAD=DB2		SERVICE CLASS=DDFWORK		RESOURCE GROUP=*NONE		PERIOD=1 IMPORTANCE=2									
		CRITICAL =NONE													
TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD I/O--	---	SERVICE----	SERVICE TIMES	---	APPL %---	PAGE-IN RATES	---	STORAGE---				
AVG	2.89	ACTUAL	14	SSCHRT	494.1	IOC	0	CPU	30.2	CP	54.33	SINGLE	0.0	AVG	0.00
MPL	2.89	EXECUTION	14	RESP	0.3	CPU	857374	SRB	0.0	AAPCP	0.00	BLOCK	0.0	TOT	0.00
ENDED	11137	QUEUED	0	CONN	0.2	MSO	0	RCT	0.0	IIPCP	29.79	SHARED	0.0	CEN	0.00
END/S	200.22	R/S AFFIN	0	DISC	0.0	SRB	0	IIT	0.0	AAP	0.00	HSP	0.0	EXP	0.00
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.1	TOT	857374	HST	0.0	IIP	0.00	HSP MISS	0.0	SHR	0.00
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	15413	IFA	0.0			EXP SNGL	0.0		
AVG ENC	2.89	STD DEV	16					IIP	0.0			EXP BLK	0.0		
REM ENC	0.00					ABSRPTN	5326					EXP SHR	0.0		
MS ENC	0.00					TRX SERV	5326								

SERVICE TIMES	
CPU	30.2
SRB	0.0
RCT	0.0
IIT	0.0
HST	0.0
AAP	0.0
IIP	0.0

---APPL %---	
CP	54.33
AAPCP	0.00
IIPCP	29.79
AAP	0.00
IIP	0.00

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

Estimated Redirect % = 55% (APPL% IIPCP / APPL% CP)



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

CONNTYPE: DRDA		
AVERAGE	APPL(CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	0.002754	0.001726
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
IIP CPU TIME	0.000000	0.000000

← 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 % = 55% (Class 1 IIPCP / CP)



DRDA zIIP Redirect Measurement Summary

- Measured with CLI SQL and Stored Procedure distributed workloads.
 - CLI workload achieved expected redirect %
 - Stored Procedure achieved 13% redirect
 - Stored Procedure Call, Results set and Commit processing eligible for zIIP redirect.
 - Parallel Query workload achieved expected redirect %
- No noticeable CPU overhead or elapsed time increase for zIIP redirect processing.
- Positive feedback from Beta customers.
 - Using in production environment.



Monitoring & Estimation of zIIP Redirect for Parallel Query Workload with DB2 V8



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 Service Class BATCH_M and Reporting Class SSPQ1

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



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))



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 Service Class BATCH_M and Reporting Class SSPQ1

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)



Tivoli Omegamon DB2PE Accounting Report with Local Parallel Query zIIP Estimate

PLANNAME: DSNTEP81		
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 PROJECTCPU=YES or when zIIP processor is configured but offline

Estimated Redirect % = 70% (IIPCP / CP)



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.



Monitoring & Estimation of zIIP Redirect for **Utility Workload** with DB2 V8



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 Service Class BATCH_M and Reporting Class RBLDINDX

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



Tivoli Omegamon DB2PE Accounting Trace for Rebuild Index Utility zIIP Redirect

PLANNAME: DSNUTIL CORRNAME: RBLDV8

AVERAGE	APPL (CL.1)	DB2 (CL.2)
-----	-----	-----
CP CPU TIME	1:04.42090	24.169916
AGENT	0.016858	0.008149
NONNESTED	0.016858	0.008149
STORED PRC	0.000000	0.000000
UDF	0.000000	0.000000
TRIGGER	0.000000	0.000000
PAR.TASKS	1:04.40404	24.161766
IIPCP CPU	6.158141	N/A
IIP CPU TIME	16.235916	16.235916

← **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 % = 27% ((IIP +IIPCP) / (CP+IIP))
zIIP Redirect % = 20% ((IIP / (CP+IIP))
zIIP eligible but ran on CP = 7% ((IIPCP / (CP+IIP))

Total CPU : 80.6 sec, zIIP CPU = 16.2 sec



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))



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 Service Class BATCH_M and Reporting Class RBLDINDX

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

$$\text{Estimated Redirect \%} = \text{APPL\% IIPCP} / \text{APPL\% CP}$$



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)



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 offload



Estimation of zIIP Redirect Prior to DB2 V8



Estimating DRDA zIIP Redirect

- Ensure that WLM policy is setup with Service Class(es) for SUBSYSTEM TYPE=DDF
- Run the RMF Workload activity report (SYSRPTS) for the peak periods of interest showing DB2 DDF work related Service and Reporting Classes
- The APPL% CP under the DDF work Service or Reporting Class(es) for the DB2 Subsystem can be used to determine the DB2 enclave SRB CPU eligible to be redirected
 - Portion of the eligible work will be redirected to zIIP
 - If the eligible work is significant, contact IBM Tech Line / FTSS to help with the estimation for DRDA zIIP redirect
 - Indicate whether Stored Procedure and / or UDF (User Defined Function) are used extensively



Rule of Thumb for Parallel Query zIIP Redirect

- Gather DB2 Accounting trace Class(1,2,3) and Statistics Class(1) for the period of interest
- Run Accounting report using a tool such as DB2PM / DB2PE excluding Connect Type Utility.
- If the average Class 2 'PAR.TASKS' field in the Grand Total section has significant value (> 50 ms) then there may be a potential for zIIP redirect.
 - Run Accounting Trace
 - Using DB2PE performance database or another tool, filter the queries that use more 100 ms of PAR.TASKS CPU
 - Sum all the Class 2 PAR.TASKS CPU of such queries and subtract $0.1 * \text{number of parallel groups executed}$ (from Statistics report for the same period)
 - Significant portion of the result will be eligible for zIIP redirect
- More V8 parallelism potential with parallel sort and parallel multi column merge join.



Rule of Thumb for Utilities zIIP Redirect

Following are possible reduction in Utility Class 1 CPU with zIIP redirect 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



Summary

- zIIP specialty engine may enable cost effective solutions for some DB2 workloads.
 - Actual cost savings depends on Software pricing model used and the workloads that drive the peak CPU usage.
- Easy implementation
 - No DB2 application change
 - No external tuning options
- zIIP can be leveraged to grow or develop or port new distributed and business intelligence applications on DB2 for z/OS in a cost effective way.
- zIIP Reference Information: <http://www.ibm.com/systems/z/ziip/>

