



Specialty engines on System z9 - zAAPs and zIIPs (and IFLs and ICFs)

Tapio Koskinen
IT specialist
IBM Systems and Technology Group



© 2006 IBM Corporation


IBM Systems



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

Centralized data sharing across mainframes



Internal Coupling Facility (ICF) 1997



Integrated Facility for Linux (IFL) 2001

Support for new workloads and open standards



System z Application Assist Processor (zAAP) 2004

Incorporating Java into existing mainframe solutions



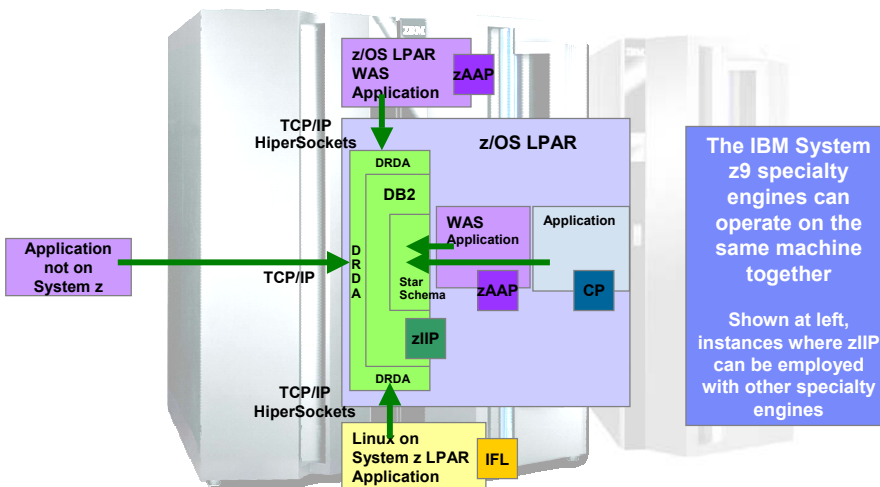
IBM System z9 Integrated Information Processor (IBM zIIP) 2006

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

Specialty engines

- **IBM System z9 Integrated Information Processor (IBM zIIP)**
 - ▶ Designed to improve resource optimization
 - ▶ Can help lower cost of computing for eligible workloads
 - ▶ Requires z/OS 1.6
 - First IBM exploiter will be DB2 UDB for z/OS V8
- **System z Application Assist Processor (zAAP)**
 - ▶ zAAPs support Java code execution
 - z/OS Java Virtual Machines (JVMs) assist with the execution of code from standard processors to zAAPs
 - JVM executes the Java code on the zAAP
 - ▶ Designed to provide a Single Tier integrated application and database serving environment
 - ▶ Requires z/OS 1.6
- **Integrated Facility for Linux (IFL)**
 - ▶ Provides additional processing capacity exclusively for Linux workloads
 - ▶ Runs Linux native or as a guest of z/VM® V4 and V5
- **Internal Coupling Facility (ICF)**
 - ▶ Provides additional processing capacity used for coupling to other processors


Specialty engines



Introducing the System z9 Application Assist Processor (or zAAP)

A specialty assist processor dedicated exclusively to execution of Java workloads under z/OS® – e.g. WebSphere®



- ▶ Available on IBM System z9 – EC & BC
 - Also Server™ zSeries® 990 (z990) and zSeries 890 (z890)
- ▶ Leveraged by workloads with Java cycles, e.g. WebSphere, DB2®
- ▶ Up to 1 zAAP per general purpose processor in an LPAR
- ▶ Executes Java Code with no changes to applications
- ▶ Traditional IBM zSeries software charges unaffected



Objective: Enable integration of new Java based Web applications with core z/OS backend database environment for high performance, reliability, availability, security, and lower total cost of ownership

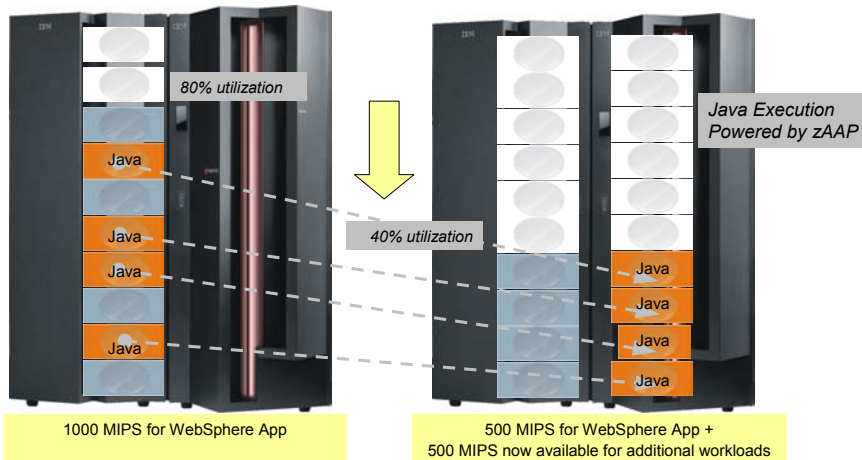
Leveraging zAAPs for e-business Integration and Infrastructure Simplification

- zAAPs can help consolidate, simplify and reduce server infrastructure and improve operational efficiencies.
 - ▶ Enables strategic integration of e-business applications with mission-critical database workloads
 - ▶ Potential operational advantages over distributed multi-tier solutions
- Eliminates separate tier to handle application server workload
 - ▶ Remove one hardware tier
 - ▶ Remove one TCP/IP link
- Leverage core zSeries strengths and manage Java Workloads automatically with z/OS
 - ▶ zSeries Security
 - ▶ zSeries Workload Manager (WLM)
 - ▶ zSeries Availability
 - ▶ zSeries Scalability
 - ▶ zSeries Flexibility



zAAP Concept Overview: A Simplified Example...

Consider a WebSphere Application that is transactional in nature and requires 1000 MIPS today on zSeries.



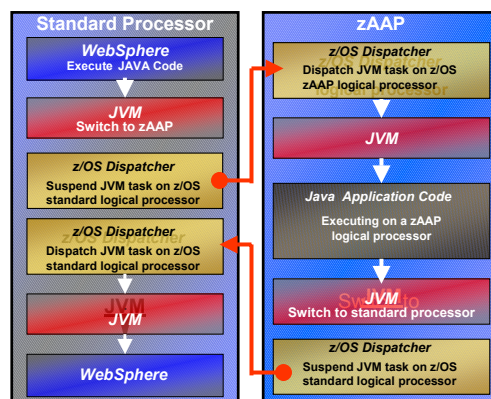
In this example, with zAAP, we can reduce the standard CP capacity requirement for the Application to 500 MIPS or a 50% reduction. * For illustrative purposes only

7

IBM Systems

zAAP Architecture and Workflow: Executing Java under IBM JVM control

- IBM JVM, parts of LE runtime, and z/OS Supervisor needed to support JVM execution can operate on zAAPs
- IBM JVM communicates to z/OS dispatcher when Java code is to be executed
 - ▶ When Java is to be executed, the work unit is "eligible" to be dispatched on a zAAP
- z/OS dispatcher attempts to dispatch zAAP eligible work on a zAAP (when present)
 - ▶ zAAP ineligible work only dispatched on standard processors
- If there is insufficient zAAP capacity available, or standard processors are idle, the dispatcher may dispatch zAAP eligible work on a standard processor
 - ▶ There is an installation control to limit the use of standard processors to execute zAAP eligible work (see Java code execution options)



8

IBM Systems

Requirements for zAAP Exploitation

- Available on System z9 – EC & BC
- Prerequisites:
 - ▶ z/OS 1.6 (or z/OS.e 1.6)
 - ▶ IBM SDK for z/OS, Java 2 Technology Edition, V1.4 with PTF for APAR PQ86689
 - ▶ Processor Resource/Systems Manager™ (PR/SM) must be enabled.

Subsystems and Apps using SDK 1.4 will exploit zAAPs automatically:

- ▶ WAS 5.1
- ▶ CICS® /TS 2.3
- ▶ DB2 V8
- ▶ IMS™ V8
- ▶ WebSphere WBI for z/OS

zAAPs must be jointly configured with general purpose processors within z/OS LPARs

- ▶ Number of zAAPs may not exceed the number of permanently purchased CPs



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 IBM System z9 not to exceed number of standard processors
 - ▶ No changes anticipated to DB2 UDB for z/OS V8 applications
 - ▶ Price for each zIIP is \$125k (US) for the z9 EC and \$95k (US) for the z9 BC.*
 - ▶ No IBM software charges on the zIIP – consistent with other specialty engines
- **DB2 UDB for z/OS V8 will be first IBM exploiter of the zIIP with**
 - ▶ System z9 EC or z9 BC
 - ▶ z/OS 1.6 or later (with PTFs)
 - ▶ DB2 UDB for z/OS V8 (with PTFs)

* Prices may vary outside of the US



DB2 V8 exploitation of IBM zIIP can add value to database workloads

- Portions of the following DB2 UDB for z/OS V8 workloads may benefit from zIIP*

- ERP, CRM, Business Intelligence or other enterprise applications
 - Via DRDA® over a TCP/IP connection
 - Examples could be SAP using DB2Connect, and WebSphere applications that use type 4 connectors to access DB2 UDB for z/OS



New Specialty Engine

- Data warehousing applications*
 - Requests that utilize complex star schema parallel queries
- DB2 UDB for z/OS V8 utilities*
 - Internal DB2 utility functions used to maintain index maintenance structures

- Utilization of the zIIP is expected to be transparent to the application.


- No anticipated changes to applications that access DB2 UDB for z/OS V8 data

* 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 above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Stored procedures and user-defined functions do not use SRBs and so are not eligible.

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’




- DRDA is implemented in DB2 UDB for z/OS V8 and 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. DB2 Connect is an example of a DRDA application server. The universal driver is an example of a DRDA application requestor
- DRDA is network independent. It can use TCP/IP or SNA as a network protocol to flow DRDA commands. Connections using SNA are not eligible for zIIP.

So.... regarding the zIIP: If DB2 for z/OS V8 work load comes over TCP/IP and is DRDA compliant (and not stored procedures or user-defined functions), a portion of that DB2 workload is eligible to be redirected to the zIIP – you need BOTH TCP/IP and DRDA but not stored procedures or user-defined functions.

What is Star Schema?

DB2 V8 zIIP redirect enabling APAR :
PK19921 for Star Schema

- **Star schema = a relational 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
 - ▶ Example: store "facts" of the sale (units sold, price, ..) with product, time, customer, and store keys in a central fact table. Store full descriptive detail for each keys in surrounding dimension tables. This allows you to avoid redundantly storing this information (such as product description) for each individual transaction
- **Complex star schema parallel queries include the acts of joining several dimensions of a star schema data set (like promotion vs. product).**



- So.... regarding zIIP: **if the workload uses DB2 UDB for z/OS V8 to join star schemas, then portions of that DB2 workload will be eligible to be redirected to the zIIP.**

What is index maintenance?


- **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.**
 - ▶ LOAD – loads your tables
 - ▶ REORG – improves your index performance
 - ▶ REBUILD INDEX – creates or rebuilds your indexes

- 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 V8 zIIP redirect enabling APAR :
PK19920 for Utilities


System z9 and DB2 UDB for z/OS V8 are an ideal data serving platform

- **Data consolidation helps reduce:**
 - ▶ Multiple copies, disparate data
 - ▶ Cost and complexity of back up and recovery
 - ▶ Network traffic
 - ▶ Amount of storage
 - ▶ DB administration and management
 - ▶ Risk associated with distributed privacy, security, and audit policies
- **Leverage System z technology**
 - ▶ Parallel Sysplex clustering for scalability AND availability AND performance
 - ▶ Data sharing = single view of the data
 - ▶ Data compression for TCO
 - ▶ Centralized backup, recovery, privacy, security, and audit policies
- **IBM zIIP specialty engine designed to help:**
 - ▶ Customers integrate data across enterprise
 - ▶ Improve resource optimization
 - ▶ Lower the TCO for data serving workloads



Example 1: Enterprise Applications

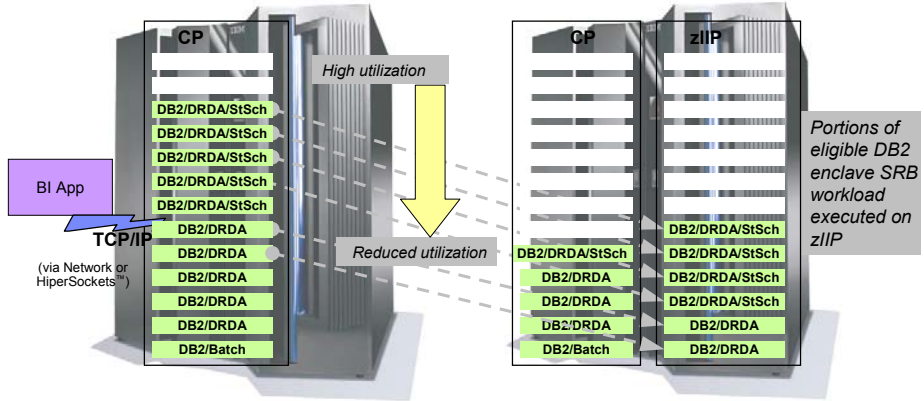
Enterprise Applications that access DB2 UDB for z/OS V8 via DRDA over a TCP/IP connection will have portions of these SQL requests directed to the zIIP



For illustrative purposes only. Single application only. DB2 V8 zIIP redirect enabling APAR: PK18454 for DRDA

Example 2.0: Business Intelligence Applications

Complex star schema parallel queries via DRDA over a TCP/IP connection will have portions of this work directed to the zIIP




For illustrative purposes only. Single application 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)

Complex star schema parallel queries via LOCAL connection will have portions of this work directed to the zIIP




For illustrative purposes only. Single application only.

Actual workload redirects may vary depending on how long the queries run and how much parallelism is used

Example 3: DB2 UDB for z/OS V8 utilities

DB2 UDB for z/OS V8 utilities used to maintain index structures



For illustrative purposes only, single application 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.

How to Estimate the Redirect Potential of DRDA (M #1)

- Two Methodologies (M) are currently available:
 - ▶ Method #1: Analyze RMF type 70 and 72 records
 - First cut to determine total amount of DRDA work running in a partition
 - Minimal amount of data
 - Requires good granularity of DRDA work in WLM service classes or report classes
 - Can provide complete analysis if customer is NOT using stored procedures or UDF (User Defined Function) or SNA
 - This method should be used for all scenarios

How do I Estimate the redirect Potential of DRDA (M #2)

- ▶ Method #2: Analyze SMF 101 (DB2 accounting) records
 - Required if customer makes extensive use of Stored Procedures or UDF
 - Can use a significant amount of data
 - IBM provides a tool which can be run on the customer machine to process the data and create a small file to be sent to IBM for further analysis
 - DB2 for z/OS DRDA zIIP Redirect Data Collector
 - Available from Techline
 - Results of this tool are integrated with the EDF Model in zCP3000.
 - Can process DB2 V7 or V8 data
 - Results can provide more accurate estimates with V8 data

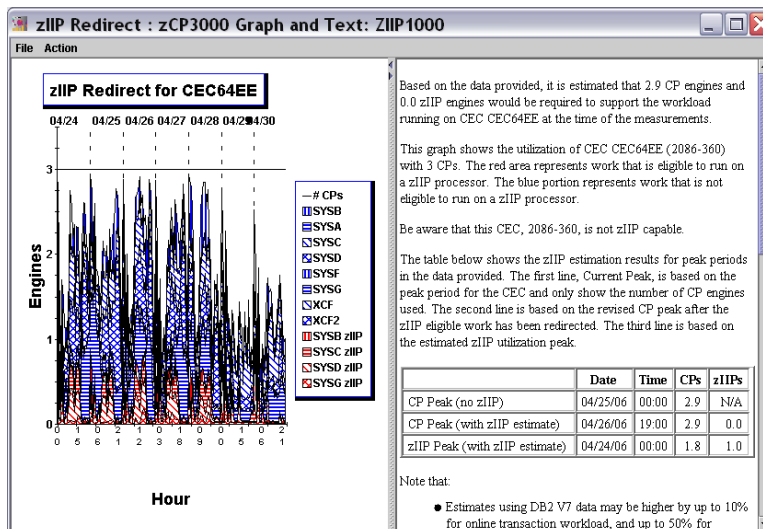
Measuring zIIP activity

- **Once a zIIP is installed (with z/OS R6 or R7 (w/ PTF) and DB2 V8 (w/ PTFs)), monitoring zIIP activity will be similar to monitoring zAAP activity**
 - ▶ Set up WLM policy with Service Class(es) for SUBSYSTEM TYPE=DDF
 - ▶ RMF Monitor 1 Type 70 Record will monitor overall zIIP activity:
 - Logical processor busy as seen by z/OS is reported
 - Physical processor busy as seen by LPAR is reported
 - ▶ RMF Monitor 1 Type 72 Record will show more detail:
 - The amount of time spent executing on zIIP processors is reported
 - Usage and Delay sample counts for zIIP eligible work is reported
 - ▶ In addition, DB2 accounting trace records can provide information on the zIIP. IBM Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS, DB2 Performance Expert or IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS can be used to monitor the zIIP information.

What is needed to estimate zIIP usage... (without the zIIP installed)


- For all LPARs whose workload has zIIP potential, RMF and DB2 traces need to be collected for certain time intervals. Afterwards, reports need to be generated from this data.
- Time intervals:
 - ▶ Data needs to be collected over time intervals of approx. 30 min. In general, there should be at least an interval representative for online and batch peak times, respectively.
- When SAP is running on System z...
 - ▶ ...and application servers on z/OS are used besides external ones, the address spaces used for the internal application servers and the DRDA- (or ICL-) address spaces should be defined in different report classes in WLM
- What data to collect?
 - ▶ RMF I data for CPU-activity report and workload activity report
 - SMF-records 70,72
 - ▶ DB2 Traces (SMF-records 101):
 - DB2 Accounting Class(1,2)

zIIP Analysis using zCP3000 (M #1 & #2)



zIIP and overall machine utilization

The overall impact the zIIP will make on your LPAR or machine or sysplex will depend on your workloads



For illustrative purposes only. Your machine utilization will vary

A vision for System z advanced data serving System z Enterprise Hub for Mission Critical Data

- ❖ With a strong foundation for transaction processing, built on 40+ years of technology innovation, System z servers with z/OS and DB2 can provide a premier platform for data serving, today and into the future*
- ❖ IBM plans to continue to invest in new solutions to address customers' strategic information on demand goals*



Today's Capabilities

- Industry-leading data integrity and security
- Data sharing solution for centralized view of data
- Scalability and availability for enterprise class workloads
- Comprehensive systems and data management environment



Extension of capabilities*

- New specialty engine (zIIP) with DB2 exploitation - for mission critical ERP, CRM, and Data Warehousing workloads*
- Database support improves regulatory compliance and autonomies
- Support of encryption capability (tape subsystem) with z/OS centralized key mgmt
- Data protection to achieve highest levels of security certifications



Future direction*

- Additional zIIP exploitation
- DB2 enhancements to help improve usability and reduce complexity and management costs.
- DB2 table scan acceleration via DS8000
- Support of encryption capability (disk subsystem) with z/OS centralized key mgmt
- Handle larger volumes of data, with improved scalability

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

For more information

- **zIIPs**

- ▶ <http://www-03.ibm.com/systems/z/about/>

- Click on data serving
 - FAQ, Whitepapers

- **zAAPs**

- ▶ <http://www-03.ibm.com/systems/z/zaap/>

- FAQ, Getting started, sizing information, whitepapers

- **IFLs**

- ▶ <http://www-03.ibm.com/servers/eserver/zseries/os/linux/>

- FAQ, Whitepapers, Tuning hints and tips