



Introduction to IBM Specialty Engine zAAP (zSeries Application Assist Processor) for Java Workloads

Theresa Tai
IBM System z New Technology Center
Poughkeepsie, New York
ttai@us.ibm.com

February 26, 2008
Session: 8366



Understanding Specialty Engine zAAP

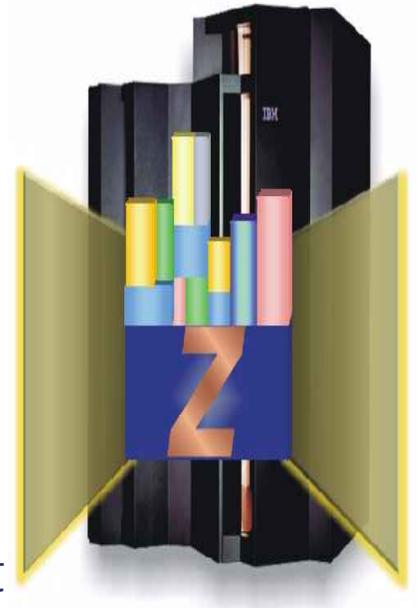


- ❖ What is IBM Specialty Engine zAAP?
- ❖ Exploitation Requirements
- ❖ zAAP Architecture and Characteristics
- ❖ Latest zAAP Configuration & Execution Options
- ❖ Projecting zAAP Eligibility Update
- ❖ What is New and Exciting: zAAP in z/OS V1R9
- ❖ Summary and Wrap-Up

zSeries Application Assist Processor



- ❖ A special-purpose processor on z890, z990, z9 hardware supporting *z/OS Java workloads*
- ❖ A specialized z/OS Java execution environment for Java-based applications
 - With no anticipated modifications to Java applications
- ❖ zAAP are attractively priced zSeries processors limited to execute z/OS Java workloads
- ❖ The processor capacity of the zAAP engines are not included when determine capacity-based software license charges from IBM software
 - The amount of savings will vary based on the amount of Java code actually executed by zAAP processors



zAAP Exploitation Requirements



❖ Prerequisites:

- z990 GA3, z890 or z9-109
- z/OS V1R6 or z/OS.e V1R6
- IBM SDK for z/OS, Java 2 Technology Edition, V1.4 with APAR PQ86689
- Middleware and Applications that are using IBM SDK 1.4
 - WebSphere Application Server V5.1 +
 - CICS® /TS 2.3
 - DB2 V8
 - IMS™ V8

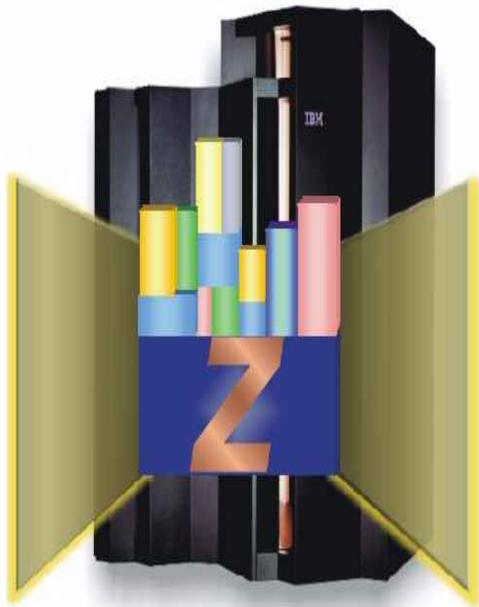
❖ Processor Resource/Systems Manager™

PR/SM must be enabled

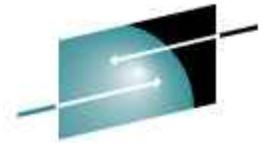
- zAAPs *must be jointly configured* with the General CPs
- Using normal PR/SM™ Logical Partition Image Profile

❖ zAAP operation enhancements

- APAR OA14131 and OA13953



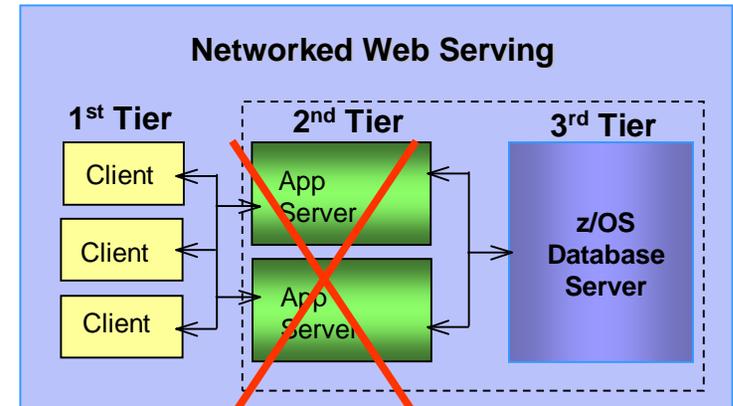
zAAP Objectives



SHARE
Technology • Connections • Results

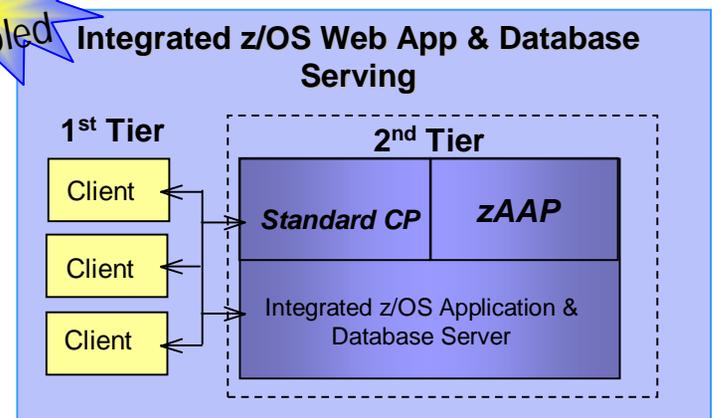
- ❖ Help simplify and reduce server infrastructures and improve operational efficiencies
- ❖ Help improve standard CP and system productivity
- ❖ Leverage on a single zSeries tier vs multi-tier front and backend data server solution
- ❖ zAAPs can deliver significant TCA savings

BEFORE



AFTER

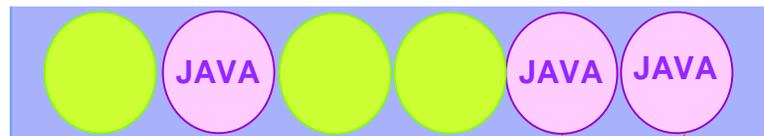
zAAP enabled



zAAP Objectives: A Simplified Example

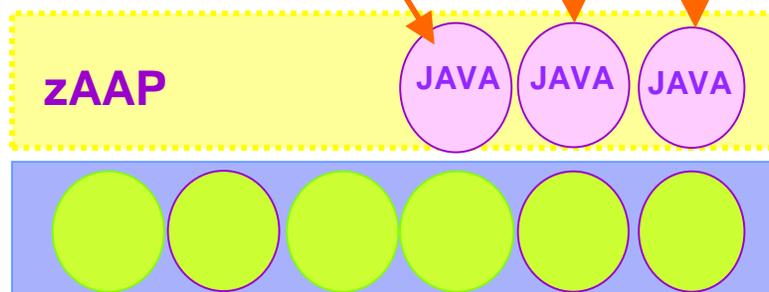
Consider a WebSphere Application that is transactional in nature and requires 1000 MIPS

80% Utilization



1000 MIPS for
WebSphere Application

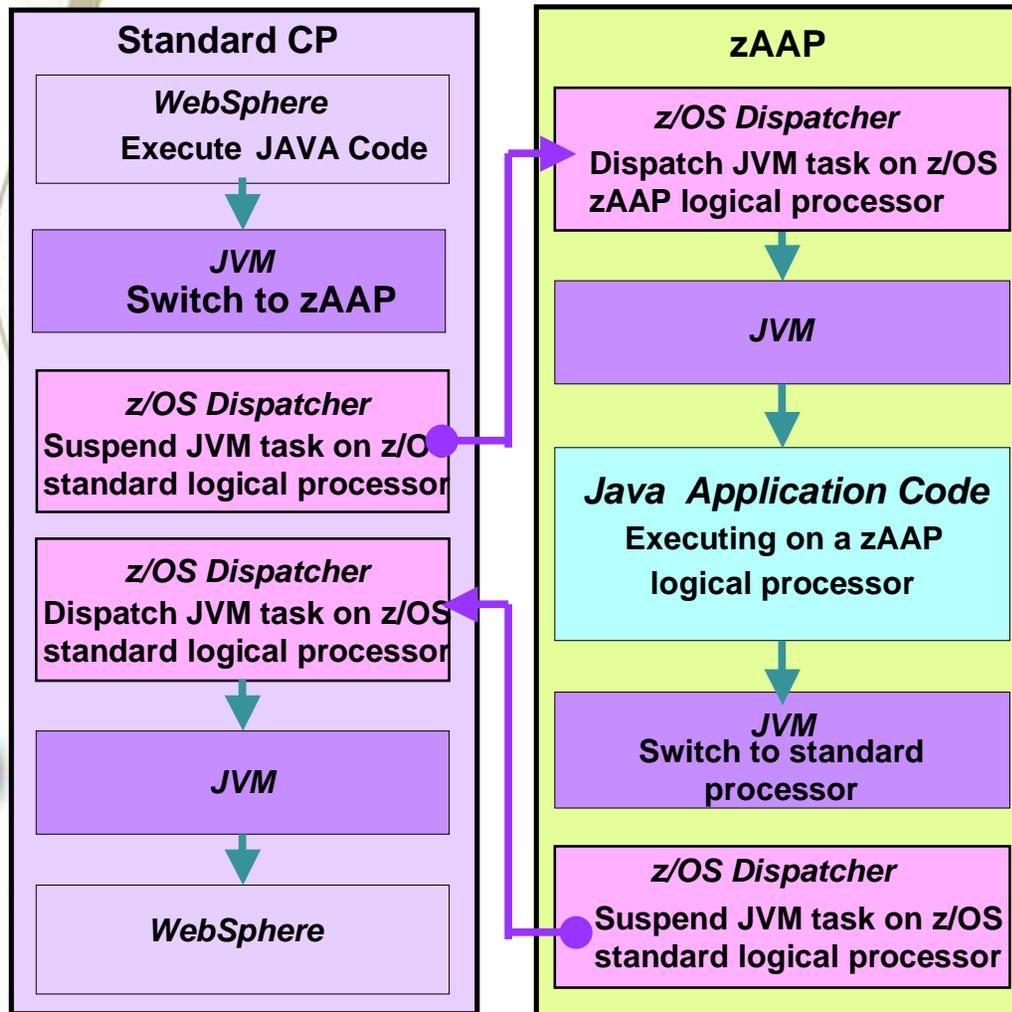
40% Utilization



500 MIPS for WebSphere App
Plus
500 MIPS Now available for
Additional Workloads

With the zAAP engines, we can reduce the standard CP capacity requirement for the Application to 500 MIPS or at a 50% reduction.

zAAP Architecture and Process Flow

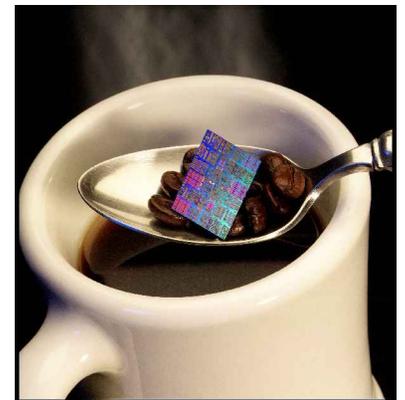


- ❖ IBM JVM, LE runtime, z/OS Supervisor, WLM, SMF/RMF components are being updated in support of the *zAAP feature*
- ❖ When *Java* is to be executed, the work unit is "*eligible*" to be dispatched on a *zAAP*
- ❖ A *Switch Service* is in place to work with the *z/OS Dispatcher*, managing the dispatching of *zAAP eligible* work between the standard CPs and the *zAAP Engines*

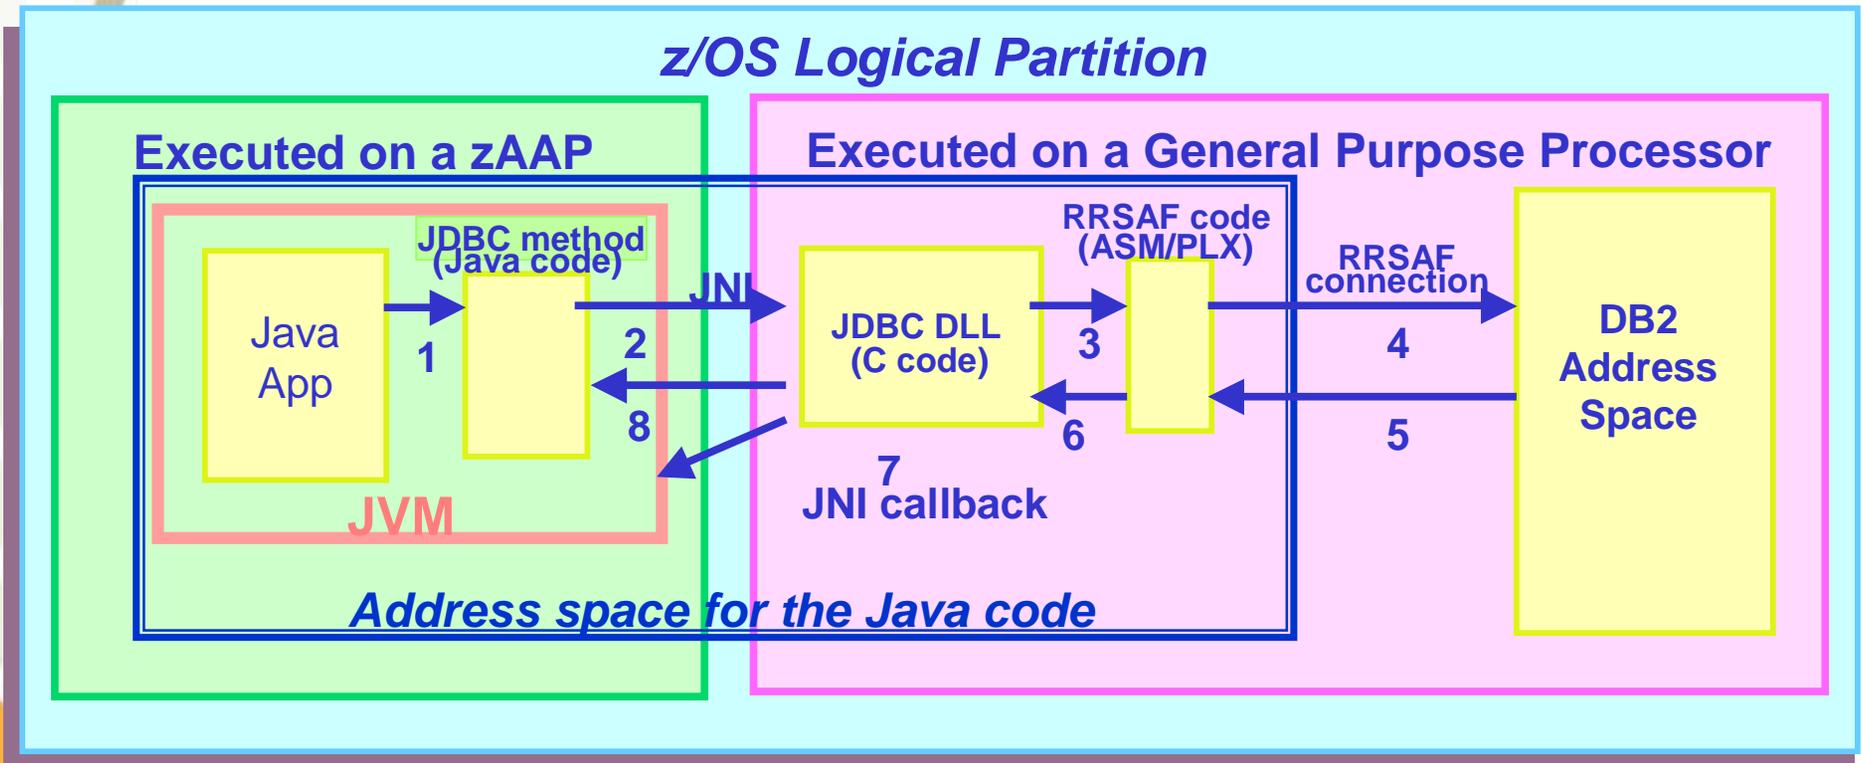
zAAP Characteristics



- ❖ zAAP can not be IPLed
- ❖ Only executes z/Architecture™ mode instructions
- ❖ Do not support all manual operator controls
 - PSW Restart, LOAD or LOAD derivatives (from file, CDROM, Server)
- ❖ Does not respond to SIGP requests unless enabled by z/OS that supports zAAPs
- ❖ The z/OS design accommodates processor differences
 - No I/O interrupts
 - No Clock Comparator interrupts
 - No affinity scheduling

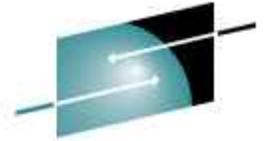


zAAP Eligibility



- ❖ Java application uses a JNI to request a z/OS DB2 database access are outside of JVM, therefore, execute only on the General Purpose Processor

PR/SM LPAR Configuration Panel



SHARE
Technology • Connections • Results

Customize Image Profiles: TC4Q04

Logical processor assignment

- Dedicated central processors
- Dedicated central processors and integrated facility for applications
- Not dedicated central processors
- Not dedicated central processors and integrated facility for applications

Not dedicated central processor details

Initial processing weight 1 to 999 Initial capping

Enable WorkLoad Manager

Minimum processing weight

Maximum processing weight

Number of processors - Initial Reserved

Number of integrated facility for application - Initial Reserved

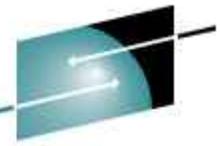
General Processor Security Storage Options Load PCI Crypto

Save Copy notebook Paste notebook Assign profile Cancel Help

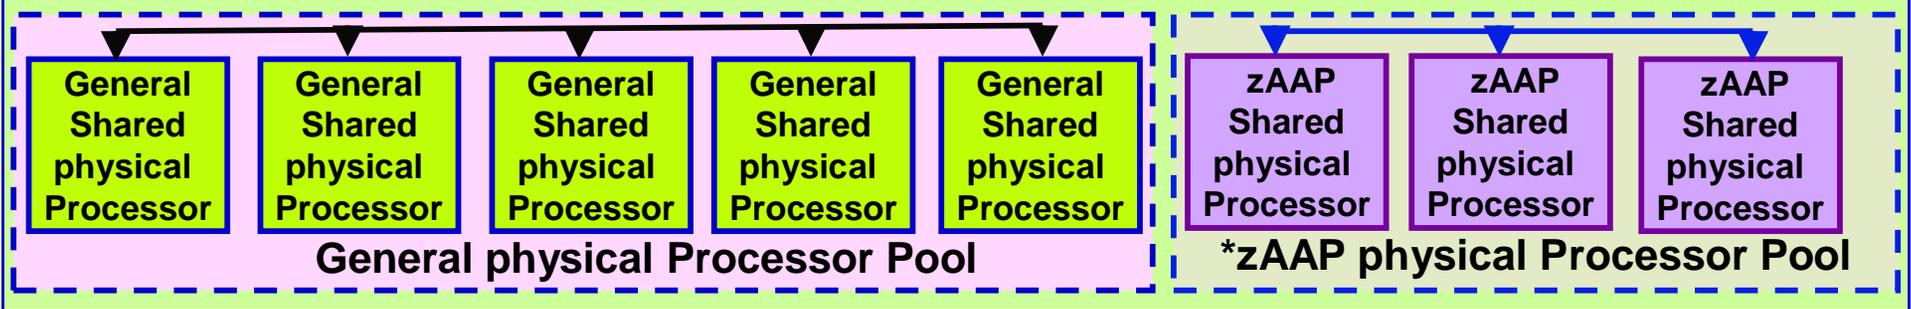
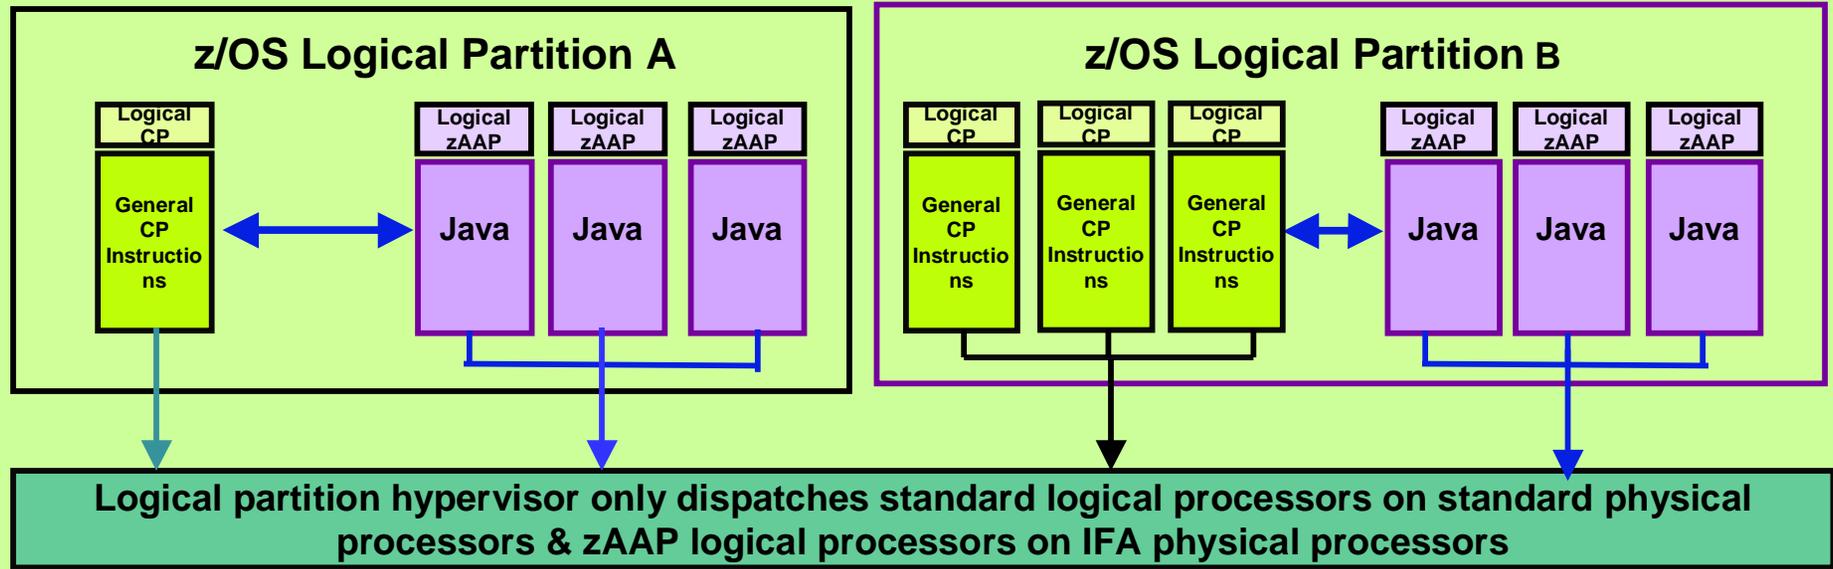
“Not dedicated” zAAP weight equals CP weight, but share calculation is based on ICF+IFL+zAAP weights.

Note: zAAP called “integrated Facility for Applications” (IFA)

zAAP Technical View: Two zAAP Partitions



Note: zAAP = IFA in PRSM panel



Note: You cannot install more physical zAAPs than physical CPs but you can assign more logical zAAPs than logical CPs to an LPAR

Single Shared ICF Pool Considerations

- ❖ zAAPs, CFs, and Linux partitions all use ICF CPs which are managed out of a single pool of capacity
 - Managed independently from the General CP pool
- ❖ zAAPs will acquire their characteristics from the z/OS partitions using the zAAPs
 - If z/OS uses dedicated CPs, the zAAPs defined to the partition will be dedicated
 - If z/OS uses shared CPs, the zAAPs defined to the partition will use shared CPs and the weight given to the zAAPs will be equal to the z/OS partitions weight

Important: The ICF pool's partition weights need to be updated to reflect the introduction of the zAAP processors

Managing zAAP Eligible Work

(FLASH10432)



SRM parameters in IEAOPTxx of SYS1.PARMLIB

- ❖ IFACrossover=Yes
 - zAAP-eligible work can be executed on the standard CPs
 - When the standard processor entered a wait state (and zAAP need help)
 - Processed at a lower priority than standard discretionary work
- ❖ IFAHonorPriority=Yes (the default, recommended)
 - Standard CPs may execute both zAAP-eligible and non-eligible work in priority order
 - If zAAP processors are unable to execute all zAAP-eligible work
- ❖ IFAHonorPriority=No
 - zAAP-eligible work can run on standard processors but at a lower priority than the non-zAAP work
- ❖ IFACrossover=No
 - Standard CPs will not execute zAAP-eligible work
- ❖ Can be dynamically changed by the SET OPT command

zAAP Option Supporting Matrix



Crossover	Honor Priority	General Purpose Processors Behavior
No	No	No zAAP work on General Purpose Processors
No	Yes	zAAP work on General Purpose Processors only when help is needed by zAAP engine/processors
Yes	No	zAAP work on General Purpose Processors only when no non-zAAP work is ready
Yes	Yes	zAAP work on General Purpose Processors when - <i>help is needed</i> and - when <i>no non-zAAP work ready</i>

Note: This change would allow more zAAP eligible work to run on zAAP processors while still remaining responsive to the zAAP demand.

In the event the last zAAP processor becomes unavailable to process zAAP work, the settings of the IFACROSSOVER and IFAHONORPRIORITY parameters are ignored as if no zAAP processors had been defined to the LPAR. The zAAP work is run in priority order with all other work by the GPPs.

Reference: [FLASH10432](#)

JVM Startup Options for zAAPs (SDK 1.4 only)

<i>-Xifa:on</i>	Enables Java workloads to be run on the zAAP processors, if its available. (default setting)
<i>-Xifa:off</i>	Disable the use of zAAP processors
<i>-Xifa:force</i>	Forces Java attempting to use zAAP processors, even if there are none available (Valid on z/OS V1.6 or later)
<i>-Xifa:projectn</i>	Tracks projected zAAP CPU usage and made available to SMF/RMF reporting (Valid on z/OS V1.2, V1.3, V1.4 and V1.5)

- ❖ *-Xifa:force* option allow the customers to use SMF 72 records for capacity planning to figure out how many IFAs they would need for their Java workloads
- ❖ *-Xifa:projectn* option will help customers to track the "Would- have-been" IFA CPU time (where n is interval length, default value is 15)

Important: JVM startup options which are only processed at JVM startup time

IEAOPTxx Options

- ❖ The PROJECTCPU parameter
 - PROJECTCPU=YES|NO
 - PROJECTCPU=YES (the default)
 - Will cause the switching to happen
 - This new support is available as part of the zIIP web deliverable
 - z/OS 1.6 FMID is JBB77S9, z/OS 1.7 FMID is JBB772S
- ❖ SDK V5 does not support Xifa:projectn option
- ❖ The zAAP estimation information can be obtained from the RMF Workload Activity Report
 - Need to set the following option in the IEAOPTxx member of SYS1.PARMLIB:
 - PROJECTCPU=YES
 - We recommend the use PROJECTCPU=YES for zAAP projection
 - It is simple and easy to use, if you are on z/OS 1.7 or later

SDK V5 zAAP Projection Support

- ❖ Only the Xifa:force option is supported
 - Identify potential zAAP usage when zAAP is not available
 - RMF/SMF data will be collected when this option is set
 - The setting is valid for z/OS V1.6 or later

- ❖ Instead of Xifa:projectn
 - The zAAP estimation information can be obtained from the RMF Workload Activity Report
 - Need to set the following option in the IEAOPTxx member of SYS1.PARMLIB
 - PROJECTCPU=YES

Note: The Java5 JVM continues to support the -Xifa:force option. The downside to this approach is it requires updating the JVM arguments for all copies of the JVM and subsequently removing the option when finished.

RMF Workload Activity Report

PROJECTCPU=YES



SHARE

Technology • Connections • Results

REPORT BY: POLICY=WLMPOL

REPORT CLASS=RPXENCL

HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS T3ENC

TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	...	---SERVICE---	SERVICE	TIMES	---	APPL	%---	
AVG	0.30	ACTUAL	3	...	IOC	0	CPU	24.4	CP	20.32
MPL	0.30	EXECUTION	2	...	CPU	581337	SRB	0.0	AAPCP	17.11
ENDED	14237	QUEUED	0	...	MSO	0	RCT	0.0	IIPCP	0.00
END/S	118.64	R/S AFFIN	SRB	0	IIT	0.0		
#SWAPS	0	INELIGIBLE	0	...	TOT	581337	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0	...	/SEC	4844	AAP	0.0	IIP	0.00
AVG ENC	0.30	STD DEV	11	...				IIP	0.0	
REM ENC	0.00			...	ABSRPTN	16K				
MS ENC	0.00			...	TRX SERV	16K				

- ❖ Report above estimate the zAAP usage for each WLM service class
- ❖ The enclaves running in this report class consumed 20.32% (CP) of an engine in the reported period, 17.11% (AAPCP) of an engine is zAAP eligible
- ❖ 84% of the enclave activity in this period is zAAP eligible
17.11% divided by 20.31% equal 84%
- ❖ For **all** service classes in your z/OS image (in addition to setting PROJECTCPU=YES)
SYSRPTS (WLMGL (SYSNAME(<system_name>))
SYSRPTS (WLMGL (POLICY))

Note: AAPCP is for zAAP and IIPCP is for zIIP

RMF™ Reporting for zAAP

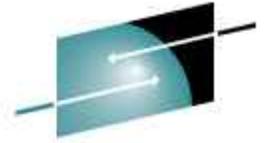
- ❖ RMF supports zAAP processors by extending the
 - Postprocessor *CPU activity report*
 - Postprocessor *Workload report*
- ❖ Monitor I
 - CPU Activity, Including Partition Report
 - Workload Activity Report
- ❖ Monitor III
 - *Enclave report (pop-up panel for IFA Using and Delay samples)*
 - *CPC Capacity Report*
- ❖ The Internals
 - Distinguishes between standard CP and zAAP processors where necessary
 - Collects and reports about *zAAP service times*
 - Collects and reports about *zAAP using and delay states* for service and report class periods

zAAP CPU Times – SMF 30 Records

- ❖ zAAP CPU time is not included in SMF30CPT
- ❖ New fields for zAAP timings

FIELD	Description
SMF30_TIME_ON_IFA	CPU time spent on IFA
SMF30_ENCLAVE_TIME_ON_IFA	Enclave time spent on IFA
SMF30_DEP_ENCLAVE_TIME_ON_IFA	Dependent enclave time spent on IFA
SMF30_TIME_IFA_ON_CP	CPU time spent running IFA eligible work on a GCP (already in SMF30CPT)
SMF30_ENCLAVE_TIME_IFA_ON_CP	IFA Enclave time spent on a GCP (already included in SMF30CPT)
SMF30_DEP_ENCLAVE_TIME_ON_CP	IFA dependent enclave time spent on a GCP (already in SMF30CPT)

- ❖ This would ensure proper billing for the new Java workloads
 - zAAPs are assist processors and have different cost structures (lower cost, does not carry IBM software charges, and have lower maintenance costs)
 - If zAAP time was in SMF30CPT you would mix CPU seconds with different cost structures
 - Current billing programs do not have to be changed



SHARE

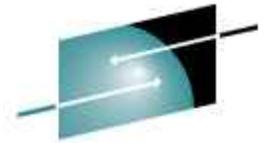
Technology • Connections • Results

What is New and Exciting



New zAAP Eligibility in z/OS 1.9

- ❖ The z/OS XML component has been enhanced to leverage IBM System z™ Specialty Engine zAAP
- ❖ Allow certain DB2 task mode XML processing cycles to be redirected to the zAAP engines
 - NOT ALL task mode DB2 stuff are zAAP offloadable
 - XML validation is not offloadable
- ❖ IBM middleware and applications requesting z/OS XML System Services will be zAAP eligible
- ❖ DB2 9 processing via local call/connection
 - Leverage pureXML capabilities of DB2 9 to be able to store XML data in DB2
 - Any individual XML document inserts, XML updates from store procedures and bulk table load
- ❖ The enhancements for z/OS V1.9 are rolled back to V1.7 & V1.8



z/OS XML zAAP Eligibility

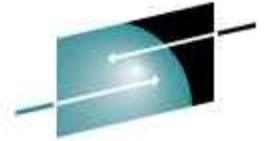
DB2 V9 XML Workload From	zAAP Eligible
CICS®	Yes
IMS	Yes
TSO	Yes
WebSphere® for z/OS (JCC T2)	Yes
DB2 Stored Procedure/UDF/Trigger	Yes
Native SQL Store Procedure	Yes
Call Attach (CAF)	Yes
RRS Attach	Yes
Local LOAD Utility (Data)	Yes

What Can You Expect on Eligibility

- ❖ You can expect 100% of XML parsing operation is zAAP eligible
- ❖ Looking ahead
 - IBM intend to extend the use of z/OS XML for additional enhancements:
 - XML Toolkit processing to exploit zAAP
 - Include validating parsing workloads
 - DB2 9 plans to implement new accounting field for zAAP CPU detail in SMF 101 records
 - via DB2 APAR PK50575

Can I Benefit From zAAP and zAAP Eligibility Projection





SHARE
Technology • Connections • Results

Can I Benefit from zAAP?

- ❖ Do you have considerable Java workload
 - As of z/OS v1.9 it will also include XML-based workloads
- ❖ Anticipate growth in Java workloads
- ❖ Need more capacity
- ❖ Consolidating Java workloads
- ❖ Need to lower the cost of running Java workloads

Not All Java Applications Are Created Equal

- ❖ Some are good candidates
 - Heavy Java or XML-based
- ❖ Some aren't good candidates
 - Light weight Java
- ❖ The synergy of all Java workloads
 - Heavy and the light weight Java content
- ❖ The cost of dispatching between zAAP and Standard CPs
 - It costs more to get there than being there
 - Look at the “Switch Rate” and “zAAP eligible microseconds per switch” under Excel workbook

How Many zAAPs Do I Need?

Projecting zAAP eligibility for my Java workload

- ❖ z/OS V1R6 with SDK 1.4.2 (SR2A)
 - SMF Type 30 and 72 records
 - JVM property option *-Xifa:force option (Plus the Excel workbook)*
- ❖ zAAPCON Tool / Questionnaire
 - Identify zAAP eligibility for Java workloads from competitive processors (none z)
- ❖ *z/OS Performance: Capacity Planning Considerations for zAAP* White Paper
 - Describes the zAAP Projection Tool, the prototype measurements
 - Describes Capacity Planning Methodology
 - May help with sizing consolidation of distributed Java workloads onto zSeries and zAAPs
- ❖ zAAP Sizing Analysis (Document # PRS2840)
 - <http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS2840>

The zAAPCON Sizing Tool



w3 Technical Sales Support, Americas

BluePages Search HelpNow Feedback

Sizewise
The Techline Sizing & Capacity Planning Portal

General Information IBM Software ISV Software IBM Hardware Other Sizing Resources News AP & Europe Techline

Welcome to the Techline Sizewise website. Here you will find access to all of Techline's Sizing and Capacity Planning Offerings. Please click on the links below or select the desired tab.

- **Sizing support for IBM Software Solutions.** Using the latest sizing information and tools available, Techline analyzes the customer's requirements and projects the resource needed for the new application. The results are then returned to you in a high quality customer deliverable. Techline also provides consultative support to assist customers in making platform deployment decisions. Sizing support for [IBM applications](#) is currently available for select IBM Data Management, Lotus, Tivoli and WebSphere products. Click on the IBM Software tab to see the list of IBM applications supported. Then download a sizing questionnaire to get started.
- **Sizing Support for ISV Applications** provides hardware sizing for select Customer Resource Management (CRM), Enterprise Application Systems (EAS), and Supply Chain Management (SCM) applications supported on IBM server platforms. Click on the [ISV Software](#) tab to see the list of ISV applications supported. Then download a sizing questionnaire to get started.
- **Solution Sizing Sheets** for self-service help projecting hardware requirements for ISV & IBM Integrated Solutions. [Solution Sizing Sheets](#) are available to provide assistance in making preliminary and budgetary planning decisions during the early phases of the opportunity engagement. Once the customer is interested more detailed analysis can be made using Techline's formal sizing process.
- **IBM Hardware** tab provides access to all [hardware capacity planning, sizing, and channel analysis offerings](#). Capacity analysis takes the customer's current workload and makes projections onto the new processor(s) being considered. We use the latest tools to model the performance characteristics to provide the estimates. The results are returned to you in a high quality customer deliverable. Techline also provides consultative support to assist customers in making platform decisions.

- ❖ A sizing/capacity planning tool (questionnaire)
- ❖ Identify zAAP eligibility
 - For Java application workloads that are currently running on Distributed platforms
- ❖ A cost saving option to consolidate Java workloads to z/OS
- ❖ Contact IBM Techline or your local IBM team

[Techdocs Library](#)[Flashes](#)[Presentations
& tools](#)[Hints, tips
& Technotes](#)[FAQs](#)[White papers](#)[SSPDs - Solution
Scenario Profiles](#)[CSPs - Customer
Support Plans](#)[QPPs - Quick
Proposals](#)[RFIs and RFPs](#)[TSGs - Technical
Support Guides](#)[Skills Transfer](#)[Sizings](#)[List by product](#)[List by date](#)[List by doc ID](#)[Active last
month](#)[Active last 12](#)

z/OS Server Consolidation with zAAP Sizing Questionnaire

Document Author: Viola Berg

Document ID: **PRS1873**

Doc. Organization: Techline

Document Revised: 05/05/2006

Product(s) covered: # 2084; # 2086; # 2094; Java; z9-109; zSeries 890; zSeries 990; z/OS; IBM System z9

Abstract: The z/OS Server Consolidation with zAAP Sizing Questionnaire is for Server Consolidation of Java workloads to z/OS zSeries (or z9s) with zAAP engines. The sizing is intended to provide a rough estimate of the zSeries GP and zAAP capacity needed to match the capacity currently being used on distributed platform servers when the work running on those servers is ported to z/OS. This implies that there will not be major changes to the application software or middleware when the work is moved to zSeries. It is a pre-sales effort based on information available at a point in time providing an entry into understanding the customer's hardware requirements. Customer's actual experiences will vary from the sizing estimate for many reasons including network activity, application design, and availability factors. The degree of variability can range from small to very significant.

It is important to understand that the sizing estimate is a pre-sales effort mainly based on benchmark performance data. The sizing estimate should not replace capacity planning for installed systems. The estimate is intended to help you roughly size the zSeries processor capacity needed for JAVA applications currently running on other servers, assuming that the work will move to z/OS zSeries

Customer results may vary, and IBM assumes no liability for actual results that differ from the sizing.



zOS_zAAP_Questionnaire 1.0.doc

Things You Should Know . . .



- ❖ zAAP capability can be exploited by any Java application using the IBM JVM
- ❖ Number of zAAPs may not exceed the number of permanently purchased CPs (including z990 unassigned CPs or z890 Downgrade – Record Only CPs) on a given machine model
- ❖ All Java applications runs under z/OS are eligible to execute on zAAP engines
 - Java workloads for zLinux are not eligible to run on zAAP engines
- ❖ RMF will use the term IFA (Integrated Facility for Applications) in all reports and panels. The term IFA will also be seen in PR/SM™ Logical Partition Image Profile
- ❖ zPCR support for zAAP capacity planning tool
- ❖ Capacity planning - consider SDP of 75% instead of the typical 90%
- ❖ You should contact your ISVs directly to determine if their charges will be affected by zAAP
- ❖ Although the zAAP engines do not contribute to the rated MSU capacity of a system, provision is made for the customer to do capacity planning, performance management and chargeback related to zAAP processor utilization
- ❖ zAAP Projection Tool workbook can be downloaded from the same website as the zAAP Projection Tool. www6.software.ibm.com/dl/zosjava2/zosjava2-p
 - This workbook is provided on an “as is” basis.

zAAP Summary



- ❖ zAAP ...an industry first
 - Only specialized processing units for Java code today
 - Supported by IBM Middleware such as WebSphere, CICS, DB2...
 - Helps reduce demands on general purpose processors
 - Make them available for other work

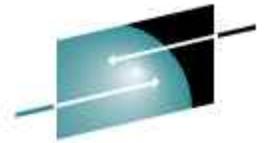
- ❖ zAAPs for e-business Integration and Infrastructure Simplification
 - Integrate Java technology-based applications with mission-critical data
 - Helps reduce infrastructure complexity for multi-tier applications

- ❖ zAAPs Provide Investment Flexibility
 - Extend the value of existing zSeries investments
 - Cost-effective, specialized Java execution environment
 - Low Total Cost of Acquisition
 - Helps reduce Total Cost of Ownership

zAAP Enablement Resources Summary



- ❖ Performance White Paper
 - <http://www-1.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100417>
- ❖ Techdoc WP100489 Mission: zAAP your costs Running WebSphere and Java on the zSeries Application Assist Processor
 - <http://www-1.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100489>
- ❖ New Dispatcher Enhancements for zAAP (Flash10432 by Kathy Walsh)
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10432>
- ❖ IBM **Redbooks** on zAAP
 - SG24-6386 zSeries Application Assist Processor (zAAP) Implementation
 - SG24-7177 Java Stand-alone Applications on z/OS, Volume I
- ❖ Getting Started with zAAP
 - <http://www-03.ibm.com/systems/z/zaap/gettingstarted/>
- ❖ IBM zAAP site
 - <http://www-03.ibm.com/systems/z/zaap/>



SHARE

Technology • Connections • Results

THANK
YOU