Understanding zAAPs and WSC Experiences

Kathy Walsh
IBM
Washington Systems Center
Advanced Technical Support

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

CICS*
DB2*
e-business logo*
e-business on demand
IBM*
IBM eServer
IBM logo*
IMS
InfoPrint*
OS/390*
Parallel Sysplex*
PR/SM
Processor Resource/Systems Manager
RACF*
RMF
WebSphere*
z/Architecture
z/OS*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft is a registered trademarks of Microsoft Corporation in the United States, other countries or both.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* 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.

Agenda

- **Introducing the zAAP**
- **Technical Overview**
- **Capacity Planning Information**
- **WSC Experiences**

Advanced Technical Support - Washington Systems Center

The New zSeries Application Assist Processor (zAAP)

New specialty assist processor dedicated exclusively to execution of Java workloads under z/OS® – e.g. WebSphere®, CICS, IMS, DB2

- Available on IBM Server™ zSeries® 990 (z990) and zSeries 890 (z890) and future zSeries servers only
- Used by workloads with Java cycles, e.g. WebSphere, DB2®
  - Executes Java code with no changes to applications
- Attractively priced at $125K USD per zAAP engine
- Significantly lower maintenance costs than standard CPs
- Up to 1 zAAP per general purpose processor in a CEC
- Traditional IBM zSeries software charges unaffected
- Sub-capacity eligible IBM software charges can be reduced
- zAAP feature available now; software exploitation planned for September 24th, 2004 with z/OS 1.6

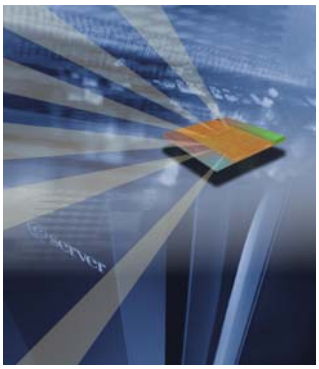

Objective: Enable integration of new Java based Web applications with core z/OS backend database environment for high performance, reliability, availability, security, and lower total cost of ownership

zAAP Overview: A Simplified Example…

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

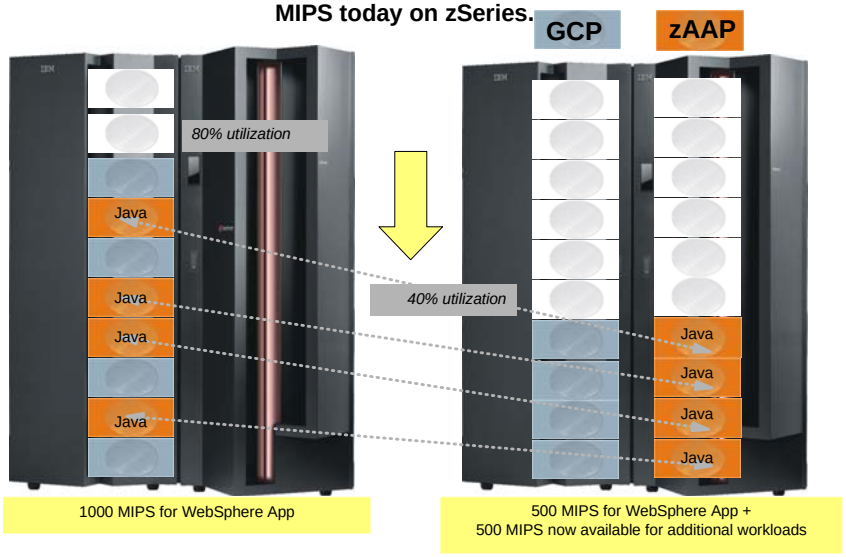

In this example, with zAAP, we can reduce the standard CP capacity requirement for the Application to 500 MIPS or a 50% reduction. ** For illustrative purposes only*

zAAP Technical Overview: *z/OS zAAP Partition*

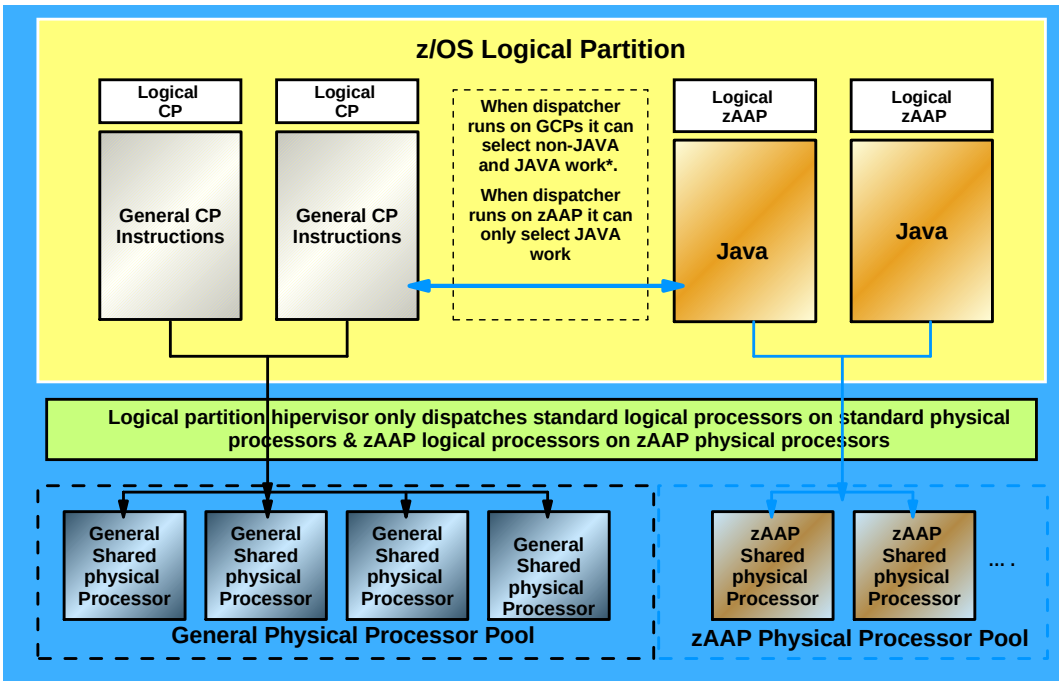

* Subject to installation controls

zAAP Workflow: Executing Java under IBM JVM control

- IBM JVM, parts of LE runtime, and z/OS Supervisor are needed to support JVM execution 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
- zAAP ineligible work is only dispatched on standard processors
- There are installation controls to limit the use of standard processors to execute zAAP eligible work

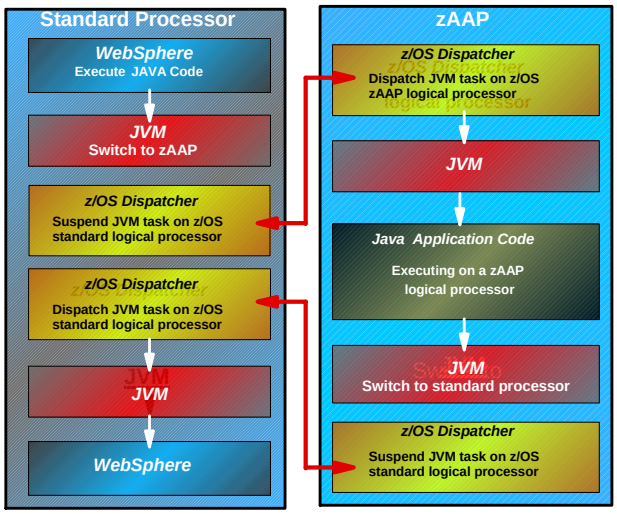

zAAP Integration at Work: Java App calling DB2

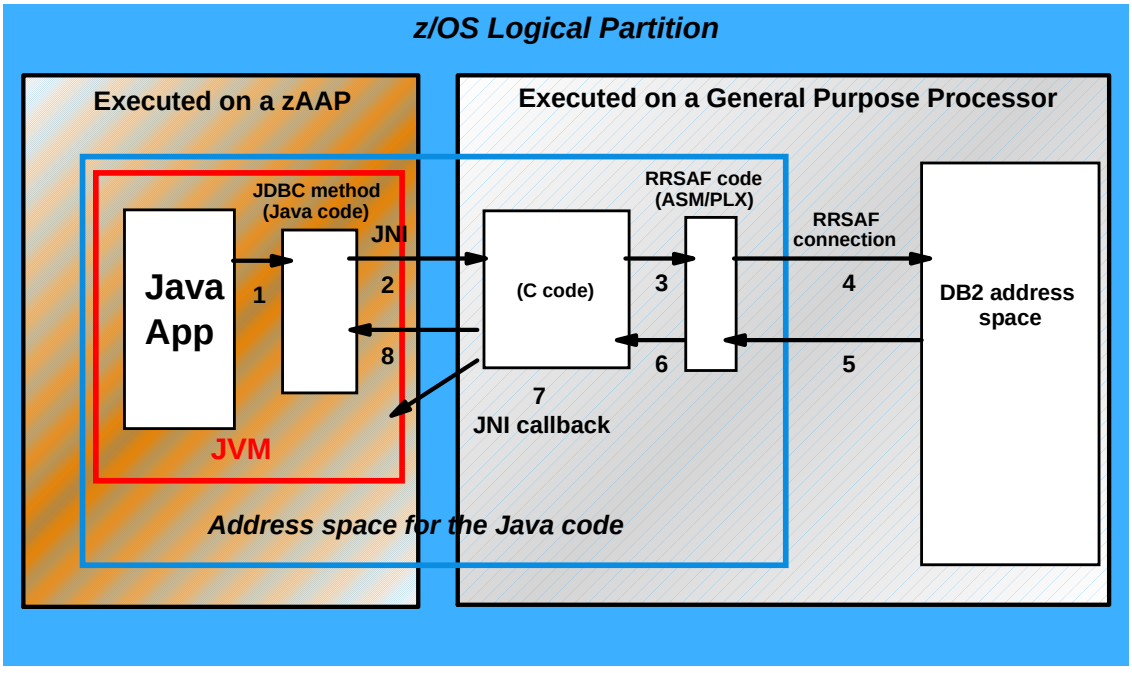

Increased e-business Integration and Infrastructure Simplification

- **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 stack
- **Leverage core zSeries strengths and manage Java Workloads automatically with z/OS**
  - Security
  - Workload Manager (WLM)
  - Availability
  - Scalability
  - Flexibility

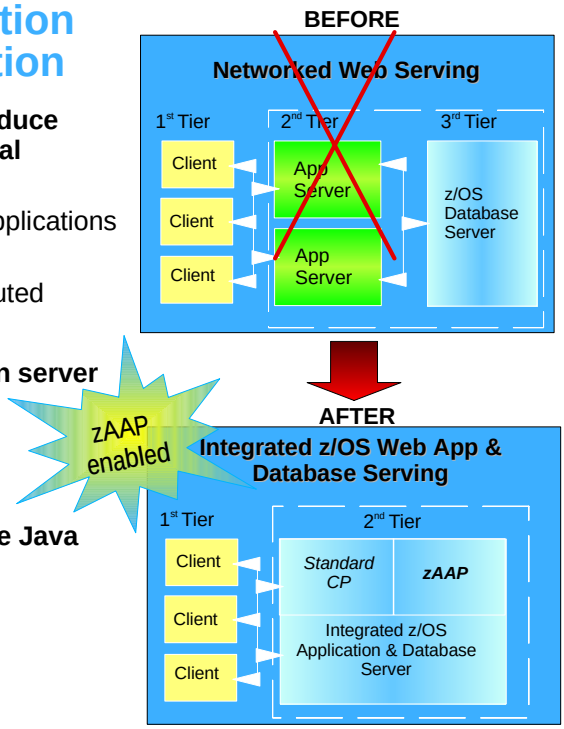

Advanced Technical Support - Washington Systems Center

How zAAPs Differ from General Purpose Processors

- **Some zAAPs Limitations**
  - zAAPs cannot be IPLed
  - zAAPs only execute z/Architecture™ mode instructions
  - zAAPs do not support all manual operator controls
    - No: PSW Restart, LOAD or LOAD derivatives (load from file, CDROM, Server)
  - zAAPs don't respond to SIGP requests unless enabled by a z/OS which supports zAAPs
  - Additional architecture differences are anticipated in future implementations
    - e.g., Java specific performance enhancements
- **The z/OS design accommodates processor differences for zAAPs:**
  - No I/O interrupts
  - No Clock Comparator interrupts
  - No affinity scheduling

Requirements for zAAP Exploitation

- **Available on z990, z890 and follow-on models only**
- **z/OS and z/OS.e exclusive**
- **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 exploit zAAPs automatically:**
  - WAS 5.1
  - CICS® /TS 2.3
  - DB2 V7, DB2 V8
  - IMS™ V7, IMS™ V8, IMS™ V9
  - 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 (including z990 unassigned CPs or z890 Downgrade - Record Only CPs) on a given machine model.

Advanced Technical Support - Washington Systems Center

Subsystem and Minimum Java Levels for zAAP Estimation

Subsystem Version	zAAP Projection Tool for Java 2 Technology Edition, SDK 1.3.1	IBM SDK for z/OS, Java 2 technology Edition, V1.4, with PTF for APAR PQ86689
WAS V4*, WAS V5.0*	X	
WAS 5.1**		X
IMS™ V7**	X	X
IMS V8**	X	X
IMS V9**		X
CICS® 2.2*	X	
CICS 2.3**		X
DB2® V7**	X	X
DB2 V8**	X	X

*** Configurations which support zAAPs**
**** zAAP projection tool can be used to assist in capacity planning**

Potential zAAP Impact on Full Capacity Pricing Scheme

BEFORE zAAP:
- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Full-Capacity EWLC @ 647 MSUs
- Average Prime Shift Machine Utilization = 80%

AFTER zAAP:
- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Java cycles executed on zAAPs
- Full Capacity EWLC remains @ 647 MSUs
- Average Prime Shift Machine Utilization = 50%

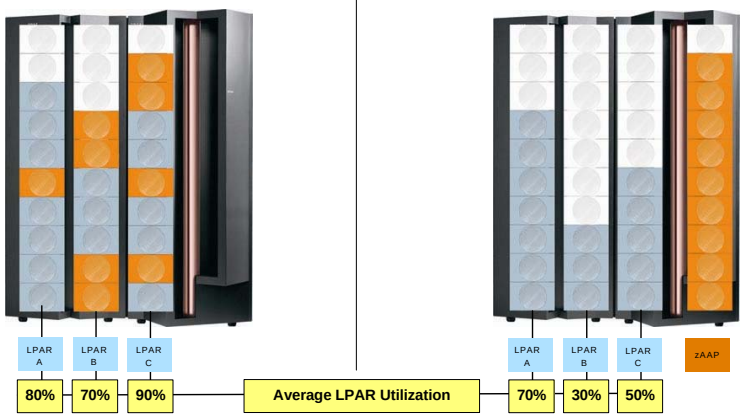

Potential zAAP Impact on Sub-Capacity Pricing Scheme

BEFORE zAAP:
- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Sub-Capacity Pricing based on;
  - LPAR A rolling 4hr avg @ 233 MSUs
  - LPAR B rolling 4hr avg @200 MSUs
  - LPAR C rolling 4hr avg @ 267 MSUs
- Rolling 4hr avg of Machine = 547 MSUs
- Average Prime Shift Machine Utilization = 80%

AFTER zAAP:
- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Java cycles executed on zAAPs
- New Sub-Capacity Pricing on *reduced* rolling 4hr avg
  - LPAR A rolling 4hr avg @ 213 MSUs
  - LPAR B rolling 4hr avg @ 100 MSUs
  - LPAR C rolling 4hr avg @ 167 MSUs
- New Rolling 4hr avg of Machine = 480 MSUs
- Average Prime Shift Machine Utilization = 50%

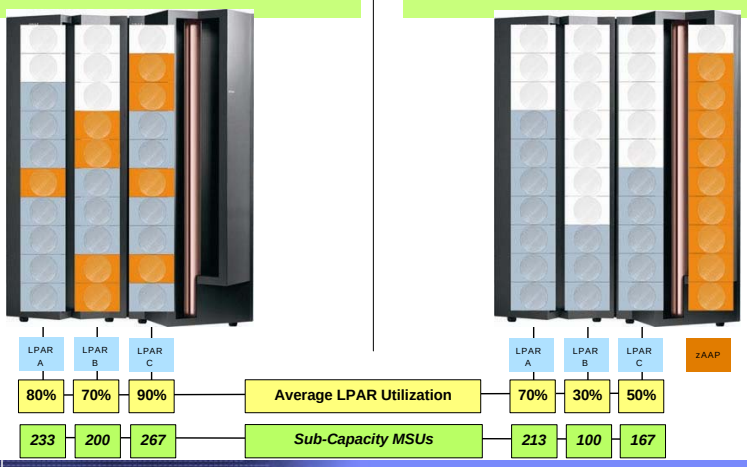

Technical Overview

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 zAPPs will be equal to the z/OS partitions weight
- **The ICF pool's partition weights need to be updated to reflect the introduction of the zAAP**

Updating CF and Linux Partition Weights with a zAAP

Base Case

Hardware	Partition	4 GCPs, 1 ICF
GCP Pool	MVSA	750
	MVSB	250 GCP pool = 1000
ICF Pool	ICF1	95
	ICF2	5
		ICF pool = 100
LPAR Weights Allowed Resource	MVSA	75%
	MVSB	25%
	ICF1	95%
	ICF2	5%

Add a zAAP

Partition	4 GCPs, 1 ICF, 1 zAAP
MVSA	750
MVSB	250 GCP pool = 1000
ICF1	95
ICF2	5
MVSA	750 ICF pool = 850
MVSA	75%
MVSB	25%
ICF1	22%
ICF2	1%
MVSA	176% (100%)

Partition	4 GCPs, 1 ICF, 1 zAAP
MVSA	750
MVSB	250 GCP pool = 1000
ICF1	712
ICF2	38
MVSA	750 ICF pool = 1500
MVSA	75%
MVSB	25%
ICF1	95%
ICF2	5%
MVSA	100%

Setting zAAP Weights

- **Always start with the zAAP weight (fixed value, with no direct installation control)**
- **Sum up the partitions which will have shared zAAPs**
- **Total share / # of shared zAAPs gives the weight of 1 CP in the ICF pool**
- **Use this value to reset the weight of the others to the same ratio you had before the zAAP**

2 SHARED zAAPs in the ICF Pool, and 1 IFL in the ICF Pool					
	Weight	Shared zAAP /IFLs	Fair Share	New Weight	Share of zAAP / IFLs
ZOS1	300	2			75% (1.5 zAAPs)
ZOS2	300	0			0%
ZOS3	100	1			25% (0.5 zAAPs)
IFL1	25	1	50%	100	50%
IFL2	25	1	50%	100	50%

zAAP weight = 400
shared zAAPs = 2
IFL CPs = 1

1 CP of weight = 400 / 2 = 200
IFL1 = 200 (1 CP share) * 50% (previous share) = 100
IFL2 = 200 (1 CP share) * 50% (previous share) = 100

z890 Considerations

- **z890s have 28 different capacity settings**
- **z890 have 4 PUs characterizable as GCP, IFL, ICF, and now zAAPs**
- **zAAPs on z890 run at full speed of the appropriate n-way**
- **Will need to understand CPU normalization**

Example: zAAP is 2 times the speed of the GCP

Execution Time = GCP seconds + (zAAP seconds * normalization factor / 256)

Execution Time = 1 second + (1 second * (2)) = 3 seconds

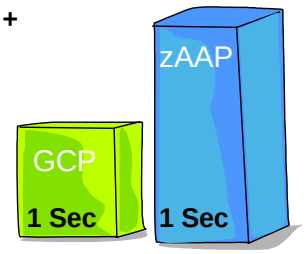

- **Normalization factor used is in RMF 72 subtype 3 record, R723NFFI**
- **If zAAP and GCP are the same speed the normalization factor is 1**

- **New SYS1.PARMLIB Options**
  - IFACROSSOVER = YES | NO
  - IFAHONORPRIORITY = YES | NO

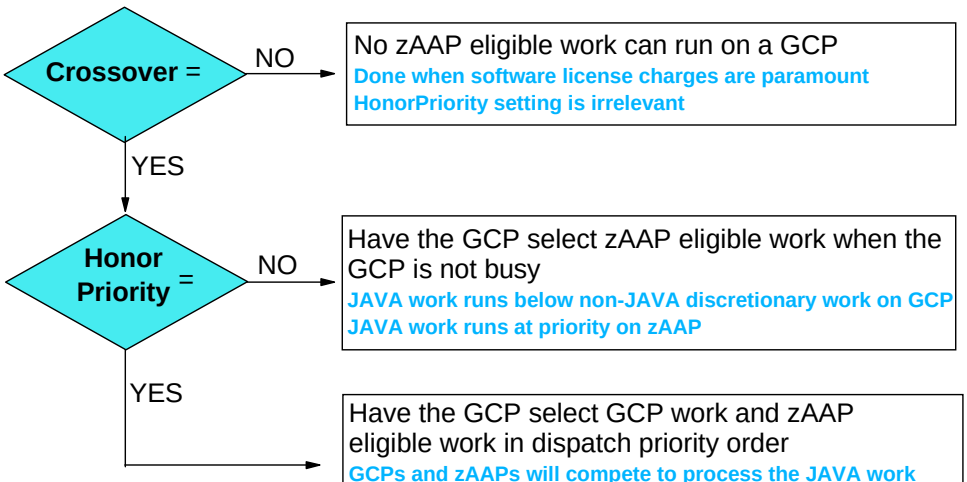

GCP - General CP
IFA - zAAP CPs

RMF Support for zAAPs

- **RMF Monitor I**
  - CPU Activity, Including Partition Report
  - Workload Activity Report
- **RMF Monitor III**
  - Meaning of CPU UTIL% in Sysinfo and Workflow/Exception report changes: Only regular CPs are included and no IFAs
  - CPC Capacity Report
  - Enclave Report
- **All RMF reports where CPU activity or delay is measured, or CPU activity and delay influences job workflow, the zAAP CPU activity is included**
- **RMF support described in APAR OA05731**

Advanced Technical Support - Washington Systems Center

zAAP CPU Time Reporting

```
            z/OS V1R6   SYSTEM ID SYSD         DATE 07/24/2004        INTERVAL 04.59.998
                        RPT VERSION V1R5 RMF   TIME 22.50.00          CYCLE 1.000 SECONDS
CPU  2084   MODEL  315
---CPU---  ONLINE TIME LPAR BUSY  MVS BUSY   CPU SERIAL  I/O TOTAL       % I/O INTERRUPTS
NUM  TYPE  PERCENTAGE  TIME PERC  TIME PERC  NUMBER      INTERRUPT RATE  HANDLED VIA TPI
 0   CP    100.00      98.31      99.19      043A6A      13.96             0.02
 1   CP    100.00      98.35      99.26      043A6A      13.55             0.02
CP   TOTAL/AVERAGE     98.33      99.23                  27.51             0.02
 2   IFA   100.00      19.79      20.39      043A6A
IFA  AVERAGE           19.79      20.39
```

- **In RMF and SMF documentation zAAPs are called IFAs**
- **New fields in the RMF product section is added to indicate the presence of zAAPs**
  - SMFxxPRF, bit 4 indicates IFA processors are available

zAAP CPU Times - RMF 72 Workload Activity

```
TRANSACTIONS    TRANS.-TIME  SS.TTT   --DASD I/O--      ---SERVICE----   --SERVICE TIMES--
AVG      1.00   ACTUAL             0   SSCHRT   0.1      IOC         3   TCB          532.9
MPL      1.00   EXECUTION          0   RESP     1.3      CPU     11059K  SRB            0.0
ENDED       0   QUEUED             0   CONN     0.8      MSO      1252M  RCT            0.0
END/S    0.00   R/S AFFINITY       0   DISC     0.3      SRB       219   IIT            0.0
#SWAPS      0   INELIGIBLE         0   Q+PEND   0.1      TOT      1263M  HST            0.0
EXCTD       0   CONVERSION         0   IOSQ     0.0      /SEC     4210K  IFA           57.6
AVG ENC  0.00   STD DEV            0                                     APPL% CP     158.4
REM ENC  0.00                                           ABSRPTN  4210K  APPL% IFACP    0.0
MS ENC   0.00                                           TRX SERV 4210K  APPL% IFA     19.2

         RESP TIME EX    PERF  AVG    --- USING% --- ----- ----- EXECUTION DELAYS % ---------
SYSTEM             VEL% INDX ADRSP    CPU  IFA  I/O  TOT   CPU  IFA

SYSD        N/A    89.0  0.6   1.0   79.5  9.5  0.0 11.0   7.6  3.4
```

TCB	- Includes GCP seconds and normalized IFA seconds
SRB	- Includes GCP seconds and normalized IFA seconds
IFA	- Includes only IFA seconds
APPL% CP	- Includes only non-IFA (GCP) time
APPL% IFACP	- Time spent on a GCP which was IFA eligible (Subset of APPL% CP)
APPL% IFA	- Time spent on the IFA

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)

- **Done to ensure proper billing for new Java workloads**
  - zAAPs are assist processors and have different cost structures, (lower cost, don't 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

How Do I Know if zAAP is Right for My Workloads?

- **"zAAP Projection Tool for Java 2 Technology Edition, SDK 1.3.1"**
  - Available with Excel Summary Workbook
  - Runs in test environment
  - Gathers usage information on % of Java in your workloads that could execute on zAAP
  - Useful in predicting number of zAAPs necessary for optimum configuration
- ***z/OS Performance: Capacity Planning Considerations for zAAP* White Paper**
  - Describes the zAAP Projection Tool
  - Describes the prototype measurements
  - Describes Capacity Planning Methodology
- **Size 390**
  - Provides special assistance for the sizing methodology described in White Paper
  - May help with sizing consolidation of distributed Java workloads onto zSeries and zAAPs
- **Information is on the web**
  - http://www-1.ibm.com/servers/eserver/zseries/zaap/gettingstarted/

WSC Experiences

WSC Testing

- **Installed z/OS 1.6 as part of the ESP**
- **Ordered and installed 1 zAAP**
  - 2084-315 and 1 zAAP (Eventually added a second zAAP)
- **Series of benchmarks**
  - Ran a Java program (JSP) which calculated Pi to user defined positions (it was used as a soaker program for the zAAP
  - Ran CICS 2.3, with Java transactions which called a DB2 V7 database
  - Ran USS Java workload which called DB2
- **Methodology**
  - Ran the workloads with SDK 1.3 and z/OS 1.5 to get a baseline and create the trace data needed for the estimation tool
  - Ran the workloads on z/OS 1.6 varying the IEAOPTxx parameters for IFACROSSOVER and IFAHONORPRIORITY
  - Measured:
    - Response Time
    - CPU per tran
    - Transaction Count
    - Various RMF Appl% values
    - zAAP busy

Advanced Technical Support - Washington Systems Center

WSC Early Results

- **Importance of switch rate on performance and throughput**
  - **CICS workload had an extremely high switching rate due to the method of calling DB2 and using the result set**
    - **JNI call backs to set variables in the JVM environment**
  - **Highlighted the need to review the switch rate when evaluating applications**
    - **Introduction of a new metric - Microseconds per switch**
    - **Amount of Java execution time per switch**
  - **Updated the jni call back functions to identify additional switch points which were not productive**
    - **May reduce the need for microseconds per switch**
  - **High switch rates can influence the accuracy of the estimation tool**
    - **On each switch 2 TIMEUSED macros are used to get the CPU time**
    - **One of the macro invocations is included in the zAAP offloadable time**

WSC Early Results

- **IEAOPTxx settings highly influence the success of utilizing the zAAP and each setting introduces additional sensitivities**
  - **IFACrossover = YES**
    - **Higher priority of the JAVA workload, the less likely the zAAP will get busy**
    - **Number of GCPs vs the number of zAAPs will reduce zAAP busy**
    - **Utilization of the GCPs will influence zAAP busy**
  - **IFACrossover = NO**
    - **Utilization of the zAAPs may need to be lower than traditional zSeries CPs**
    - **May need an increased number of zAAPs**
- **The estimation tool will most closely match your results when you run with IFACrossover = NO**
- **Changes in GCP capacity can influence the zAAP busy when IFACrossover = YES**
  - **May have the effect of causing the zAAPs to become less busy**

Advanced Technical Support - Washington Systems Center

WSC Early Results

- **Need to review and understand the impact of zAAP CPs on monitor's utilization numbers**

```
SDSF DA SYSD  SYSD     PAG    0 SIO   147 CPU  68/ 66  LINE 1-10 (10)
COMMAND INPUT ===>                                          SCROLL ===> CSR
NP   JOBNAME  SrvClass DP C StepName Workload JobID      CPU%    SIO Server
     KMWDL    BATLOW   F3 L STEP1    BAT_WKL  JOB27402   0.00   0.00 NO
     KMW7L    BATLOW   F3 L STEP1    BAT_WKL  JOB27405   0.00   0.00 NO
     KMW8L    BATLOW   F3 L STEP1    BAT_WKL  JOB27406   0.00   0.00 NO
     KMW9L    BATLOW   F3 L STEP1    BAT_WKL  JOB27407   0.00   0.00 NO
```

  - **System had 2 GCPs and 1 zAAP, system was 66% busy but work was not running.... why**
  - **SDSF incorporated the zAAP CP into CPU% field. So 2GCPs were 100% busy and the zAAP was 0% busy so the CPU% was 200% / 300% = 66%**
- **IFAHONORPRIORITY and IFACROSSOVER parmlib settings are recorded in the field R723MFLG, offset 1**
  - **If Bit 0 is ON, CROSSOVER=YES is set**
  - **If Bit 1 is ON, HONORPRIORITY=YES is set**

WSC Early Results

- **Need to evaluate and update any performance management reports which report utilization**
- **Update capacity planning methods for zAAPs**
- **Update capture ratio calculations to include zAAP times**

$$CR = \frac{(\Sigma \text{ Service Class APPL\%}) / \#\text{ Logical CPs} + (\Sigma \text{ Service Class APPL\% IFA}) / \#\text{ Logical IFAs}}{\text{LPAR Busy} + \text{IFA Busy}}$$

- **Update IEFACTRT exit to add information on zAAP time**
  - **IEF374I STEP/CICS /STOP 2004224.1105 CPU 238MIN 13.05SEC message contains non-IFA and normalized IFA time**

Advanced Technical Support - Washington Systems Center

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 and lowers total cost of ownership
  - Cost-effective, specialized Java execution environment
  - Low Total Cost of Acquisition ($125K USD per zAAP)
  - Helps reduce Total Cost of Ownership (software and maintenance Savings)