Introduction to IBM Java Workload Engine zAAP
(zSeries Application Assist Processor)

Theresa Tai
08/25/05 Session: 8366
IBM System z9 New Technology Center
Poughkeepsie, New York
ttai@us.ibm.com

Understanding zAAP

- What is zAAP?
- Exploitation Requirements
- zAAP Architecture and Characteristics
- zAAP Configuration & Execution Options
- Projecting zAAP Eligibility
  - SMF/RMF Reporting metrics
  - zAAP Eligibility Projection Tool and Excel Workbook
- Reference Summary and Wrap-Up

Click on view and follow link to header & footer to enter Copyright and Author information

zSeries Application Assist Processor (aka IFA)

- A special-purpose processor on z990 and z890 hardware
  supporting z/OS Java workloads
- A specialized z/OS and z/OS.e Java execution environment for Java-based applications
  - With no anticipated modifications to Java application
- zAAP are attractively priced zSeries processors limited to execute z/OS Java workloads
- The processor capacity of the zAAP engines are 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 zAAPs

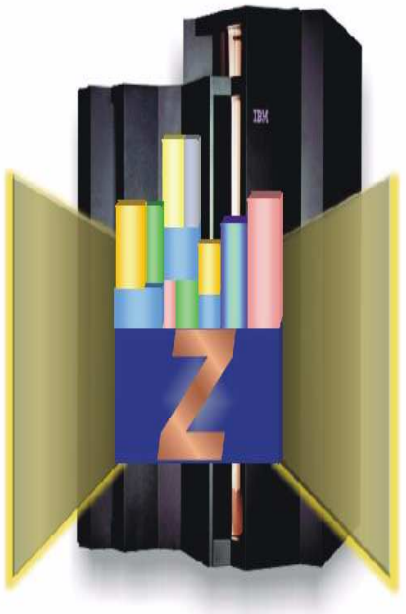

Requirements for zAAP Exploitation

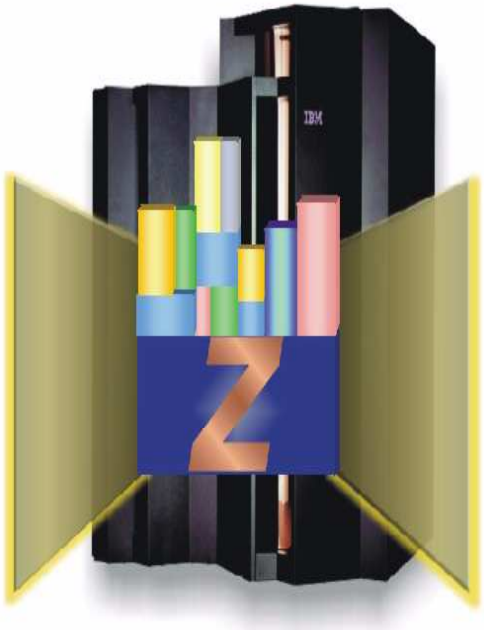

- Prerequisites:
  - z990 GA3 or z890
  - z/OS V1R6 or z/OS.e V1R6
  - IBM SDK for z/OS, Java 2 Technology Edition, V1.4 with APAR PQ86689
  - Middleware and Applications that are using SDK 1.4
    - WAS V5.1 +
    - CICS® /TS 2.3
    - DB2 V8
    - IMS™ V8
    - WebSphere WBI for z/OS
- Processor Resource/Systems Manager™ (PR/SM) must be enabled
  - zAAPs *must be jointly configured* with the General CPs
  - Using normal PR/SM™ Logical Partition Image Profile
- zAAP GA on 9/24/04 with z/OS V1R6

zAAP Objectives

- 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**

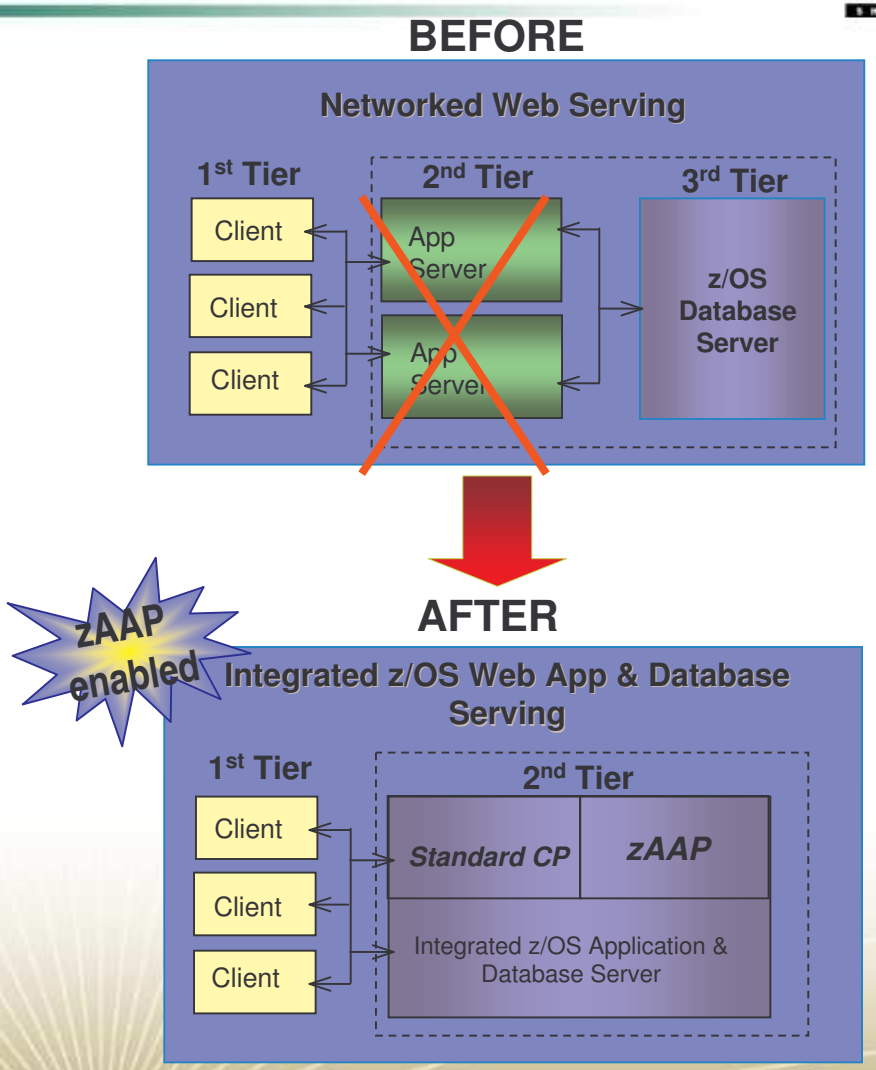

zAAP Objectives: A Simplified Example

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

80% Utilization

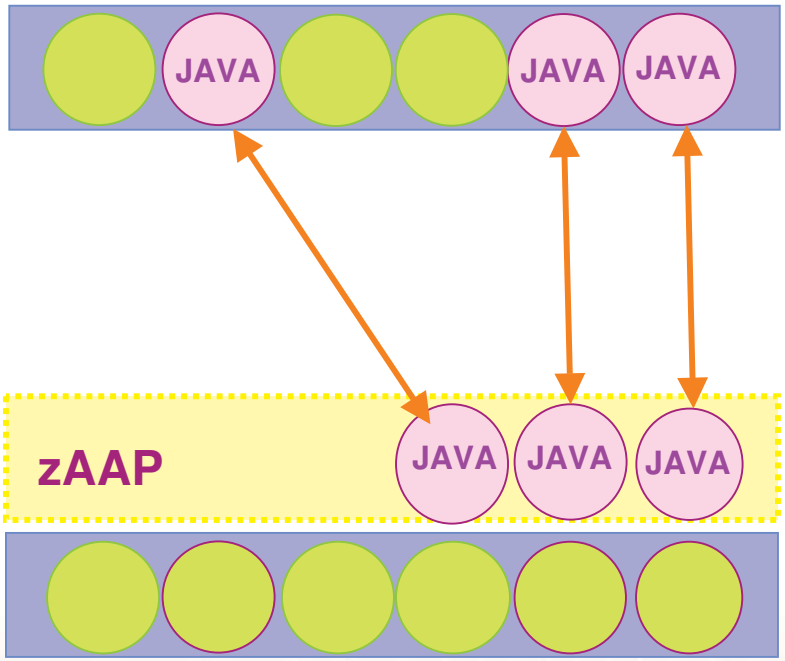

40% Utilization

1000 MIPS for WebSphere Application

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

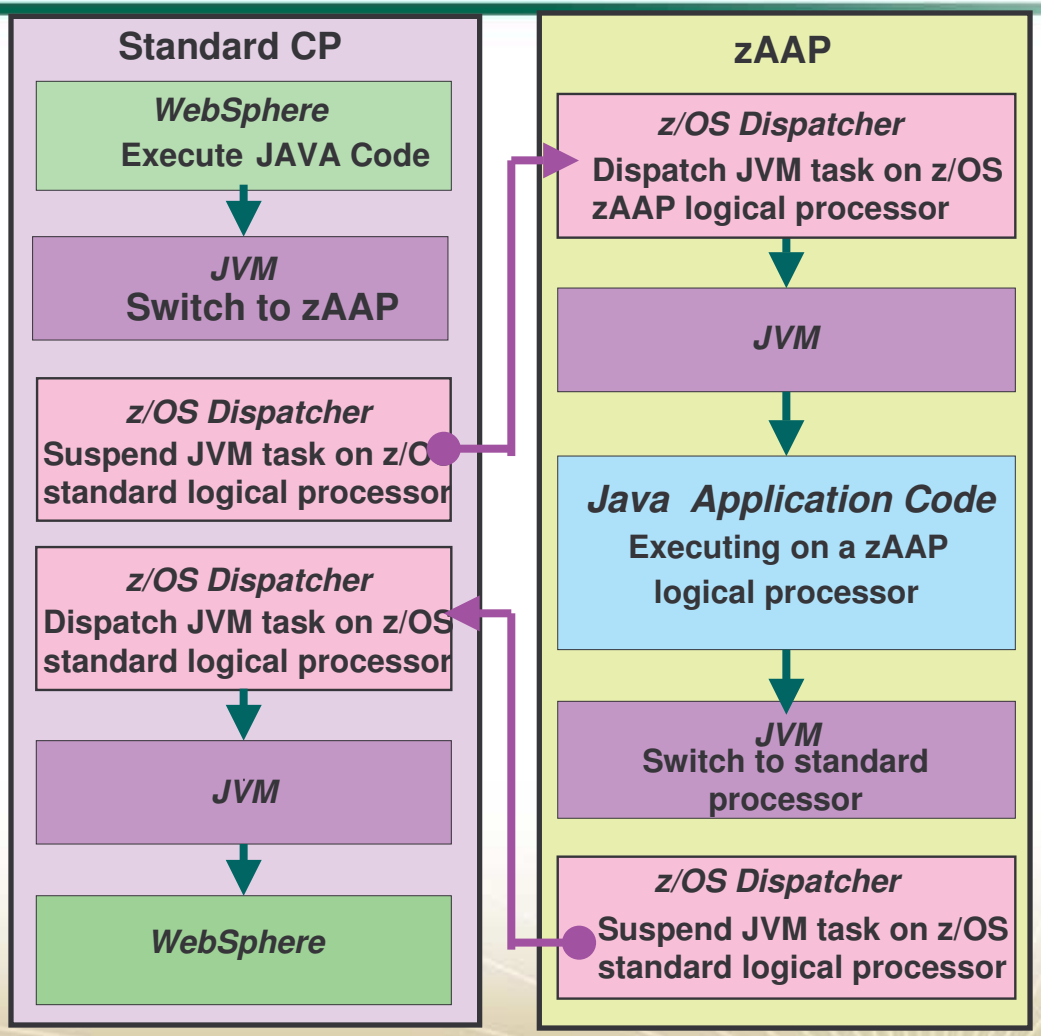

- IBM JVM, LE runtime, z/OS Supervisor, WLM, SMF/RMF components are being updated in support of the *zAAP feature*
- When *Java* is to be executed, the work unit is *"eligible"* to be dispatched on a *zAAP*
- A *Switch Service* is in place to work with the *z/OS Dispatcher,* managing the dispatching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

zAAP Characteristics

- Can not be IPLed
- Only executes z/Architecture™ mode instructions
- Do not support all manual operator controls
  - PSW Restart, LOAD or LOAD derivates (from file, CDROM, Server)
- Does not respond to SIGP requests unless enabled by z/OS
  that supports zAAPs
- The z/OS design accommodates processor diff
  - No I/O interrupts
  - No Clock Comparator interrupts
  - No affinity scheduling
- z890s CPs have 28 different capacity settings
  - zAAP on z890 always run at full speed of the appropriate n-way processor

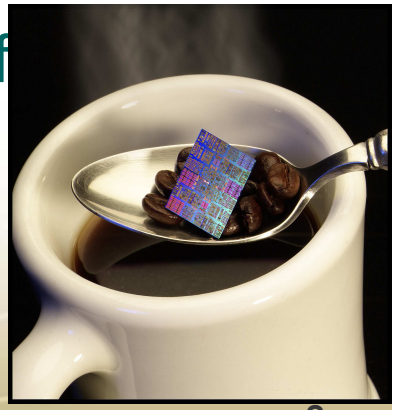

About z890

- z890s CPs have 28 different capacity settings
  - zAAP on z890 always run at full speed of the appropriate n-way processor
- Normalization factor
- Potentially faster then general purpose CPs on z890s
  - Where zAAP operates at the same speed as general purpose

zAAP Eligibility

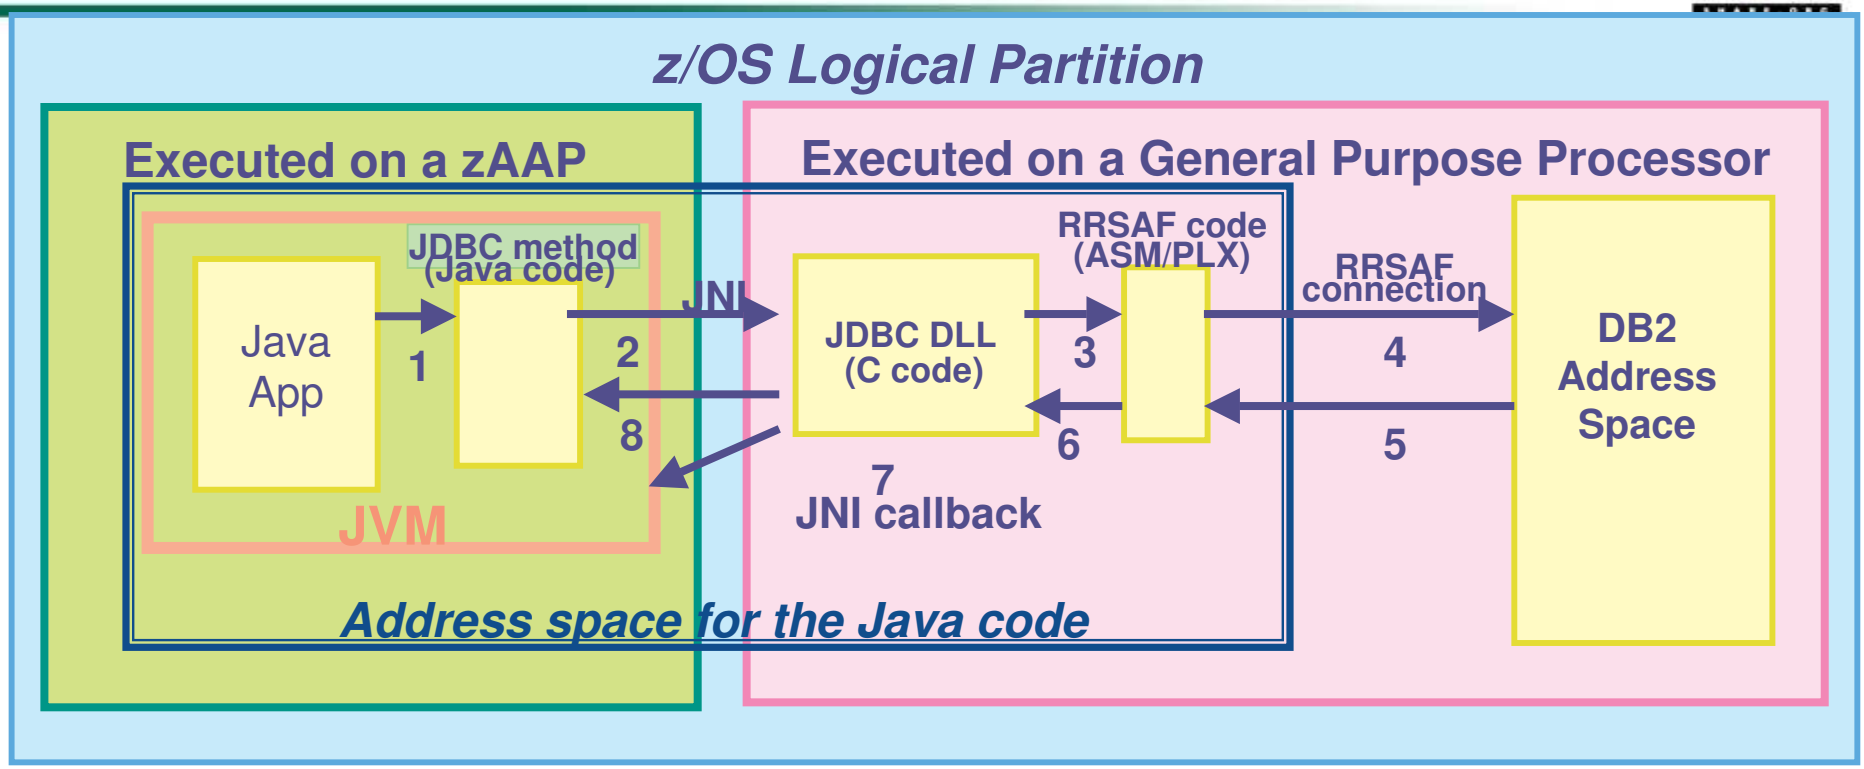

- Java application uses a JNI to request a z/OS DB2 database access are outside of JVM, therefore, execute only on the General Purpose Processor

PR/SM LPAR Configuration Panel

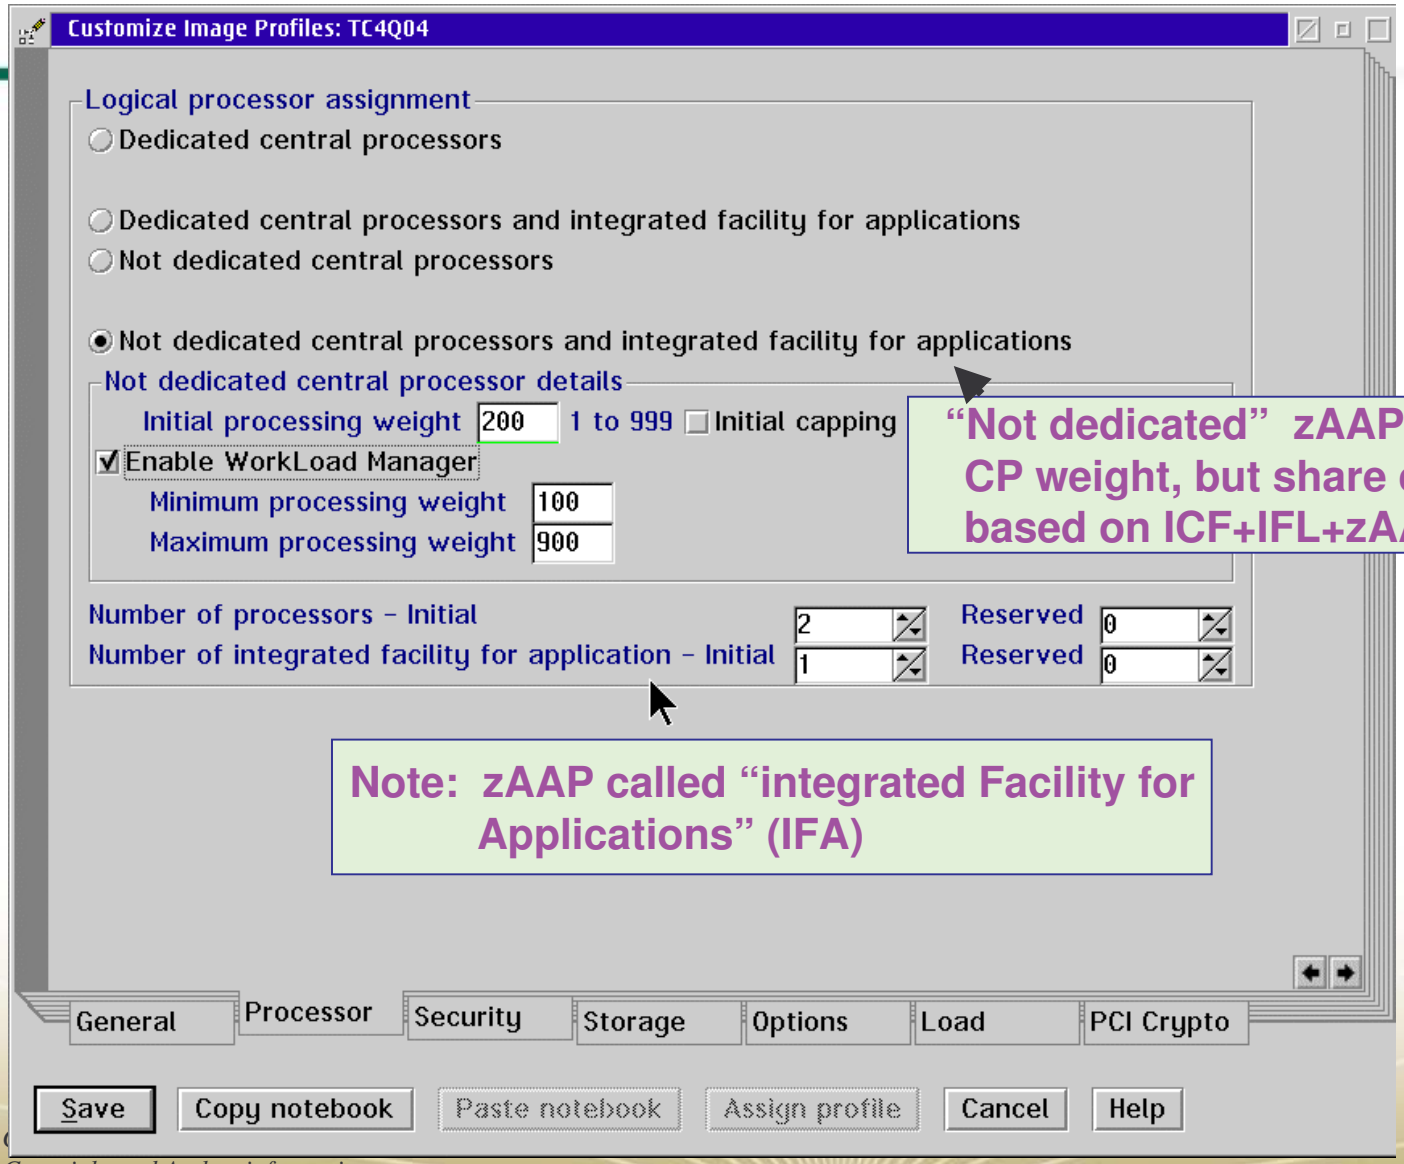

"Not dedicated" zAAP weight equals CP weight, but share calculation is based on ICF+IFL+zAAP weights.

Note: zAAP called "integrated Facility for Applications" (IFA)

zAAP Technical View: Two *zAAP Partitions*

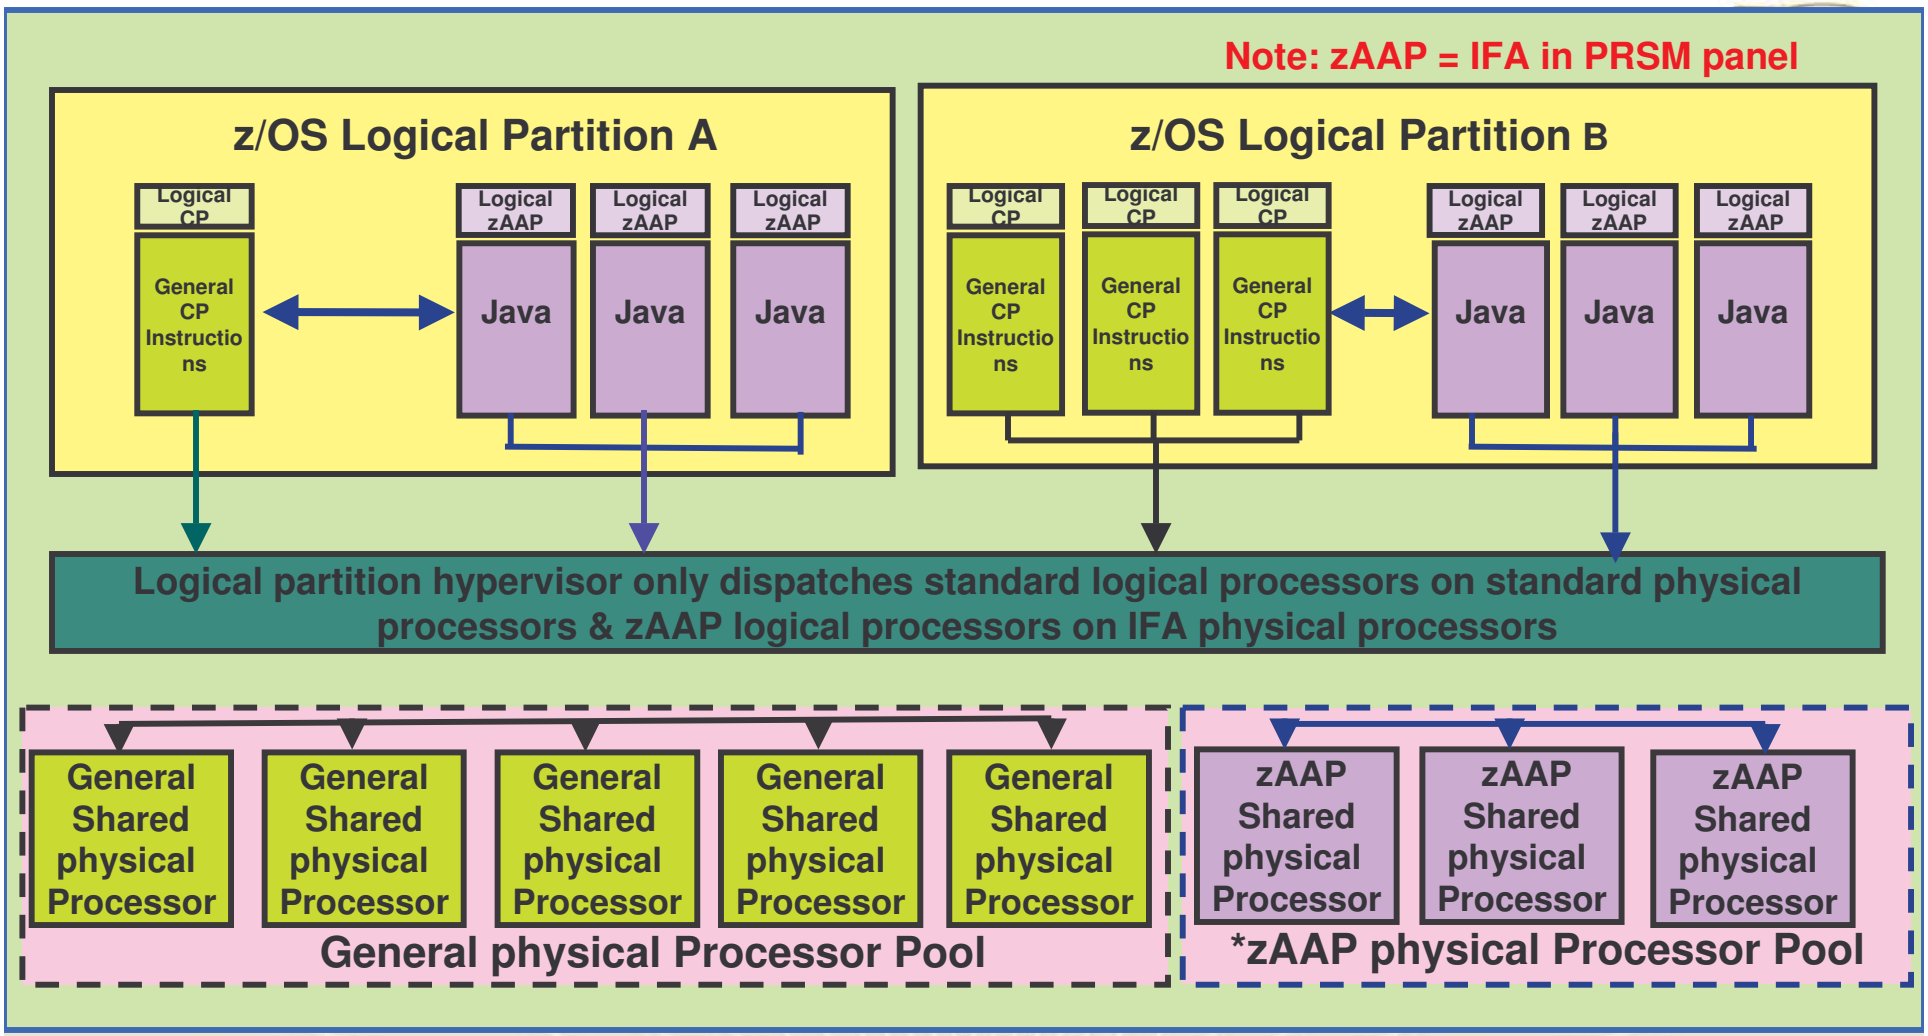

Note: You cannot install more physical zAAPs than physical CPs but you can assign more logical zAAPs than logical CPs to an LPAR

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 zAPPs 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

Controling zAAP Eligible Work

New parameters in IEAOPTxx of SYS1.PARMLIB

❖ IFACrossover=Yes

- zAAP-eligible work can be executed on both standard CPs and zAAPs
  - **IFAHonorPriority=Yes**
    - *Standard CPs will execute both zAAP and non-zAAP eligible work in priority order*
  - **IFAHonorPrority=No**
    - *zAAP-eligible work can run on standard CPs but at a priority lower than the non-zAAP work*

❖ IFACrossover=No

- Standard CPs execute only non-zAAP eligible work
- IFAHonorPriority settings are irrelevant

❖ Can be dynamically changed by the SET OPT command

Important: These parms can influence the dispatching of work on the standard CPs and capacity requirements for both standard CPs and zAAPs

zAAP Operator Interface

- D M=CPU

```
IEE174I 17.43.46 DISPLAY M
   PROCESSOR STATUS
   ID  CPU                SERIAL
   00  +                 136A3A2084
   01  +A                 136A3A2084
   02  +A                 136A3A2084
   . . .
   . . .
   A       ASSIST PROCESSOR
```

- CF CPU(*nn*),OFFLINE | ONLINE
- zAAPs are not WLM (IRD) managed, so there is no *+AW* or *–AW* status
- SDSF DA reflects zAAP usage (APAR PQ93310)
  - DA panel shows the address space service time on the CP, IFA and IFA service time on the CP

JVM Startup Options for zAAPs

-Xifa:on	Enables Java workloads to be run on the zAAP processors, if its available. (default setting)
-Xifa:off	Disable the use of zAAP processors
-Xifa:force	Forces Java attempting to use zAAP processors, even if there are none available (Valid on z/OS V1.6)
-Xifa:projectn	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 IFS 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

RMF™ Reporting

- RMF supports zAAP processors by extending the
  - Postprocessor *CPU activity report*
  - Postprocessor *Workload report*
  - Monitor III *Enclave report (pop-up panel for IFA Using and Delay samples)*
- 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 support is shipped as SPE APAR OA05371

zAAP Workload Reporting Samples

The Resource Consumption Section of the WLMGL report

```
TRANSACTIONS     TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ----SERVICE----   --SERVICE TIMES--   PAGE-IN RATES     ----
STORAGE----
AVG      4.42    ACTUAL              8.142    SSCHRT   6.5   IOC      56040    TCB         271.3   SINGLE     0.0   AVG
298.08
MPL      4.39    EXECUTION           8.142    RESP    53.6   CPU       1680K   SRB           9.2   BLOCK      0.0   TOTAL
1309.04
ENDED    1879    QUEUED                 0     CONN    20.0   MSO       2938K   RCT           4.4   SHARED     0.0   CENTRAL
1309.04
END/SEC  1.04    R/S AFFINITY           0     DISC     3.2   SRB      56695    IIT           2.0   HSP        0.0   EXPAND
0.00
#SWAPS   3154    INELIGIBLE             0     Q+PEND  25.7   TOT       4731K   HST           3.4   HSP MISS   0.0
EXCTD       0    CONVERSION             0     IOSQ     4.7   /SEC      2626    IFA          20.1   EXP SNGL   0.0   SHARED
0.00
AVG ENC  0.32    STD DEV             8.431                                     APPL% CP     15.0   EXP BLK    0.0
REM ENC  0.12                                                 ABSRPTN    598   APPL% IFACP   0.2   EXP SHR    0.0
MS ENC   0.01                                                 TRX SERV   594   APPL% IFA     1.1
```

IFA	IFA Service Time (in seconds)
APPL% CP	% of CPU time used by transactions running on regular CPs
APPL% IFACP	% of CPU time used by IFA transactions executed on regular CPs
APPL% IFA	% of CPU time on IFA processors used by IFA transactions

Note: If no IFAs/zAAPs configured, N/A is shown for the new fields.

zAAP CPU Activity Report Samples

CPU 2084 MODEL 316

---CPU--- NUM	TYPE	ONLINE TIME PERCENTAGE	LPAR BUSY TIME PERC	MVS BUSY TIME PERC	CPU SERIAL NUMBER	I/O TOTAL INTERRUPT RATE	% I/O INTERRUPTS HANDLED VIA TPI
0	CP	100.00	69.41	69.41	011511	58.67	0.00
1	CP	100.00	70.75	70.75	111511	233.6	0.00
2	CP	100.00	68.40	68.40	211511	254.2	0.00
3	CP	100.00	63.64	63.64	311511	63.49	0.00
4	CP	100.00	67.74	67.74	411511	1380	0.01
CP TOTAL/AVERAGE			67.99	67.99		1990	0.01
8	IFA	100.00	39.41	39.41	811511		
9	IFA	100.00	40.75	40.75	911511		
IFA AVERAGE			40.08	40.08			

- A new TYPE column indicates whether the processor belongs to the pool of regular CPs or IFAs
- The last two columns are only available for regular CPs
- A TOTAL/AVERAGE line is printed per pool

RMF Partition Data Report Sample


```
MVS PARTITION NAME              LP1          NUMBER OF PHYSICAL PROCESSORS      16
IMAGE CAPACITY                  167                    CP          8
NUMBER OF CONFIGURED PARTITIONS 6                      ICF         8
WAIT COMPLETION                 NO
DISPATCH INTERVAL               DYNAMIC

--------- PARTITION DATA ----------------  -- LOGICAL PARTITION PROCESSOR DATA --  -- AVERAGE PROCESSOR UTILIZA
            ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----  LOGICAL PROCESSORS  --- PHYS
NAME     S  WGT  DEF  ACT  DEF  WLM%  NUM  TYPE  EFFECTIVE      TOTAL      EFFECTIVE   TOTAL  LPAR MGM
LP1      A   50    0  167  NO    0.0    5  CP    00.29.26.356  00.29.27.505   99.94   100.0    0.02
LP2      A   50    0   33  NO    0.0    1  CP    00.05.53.443  00.05.53.501   99.98   100.0    0.00
LP4      A   50    0  268  NO    0.0    8  CP    00.47.08.000  00.47.08.008   100.0   100.0    0.00
*PHYSICAL*                                      00.00.00.237                 0.00
                                                ------------  ------------   ------
 TOTAL                                          01.22.27.801  01.22.29.251             0.03

ICF2     A   75                 8  ICF   00.04.53.443  00.04.53.501   99.98   100.0    0.00
IFL4     A   25                 3  ICF   00.24.08.000  00.24.08.008   100.0   100.0    0.00
LP1      A   50                 6  ICF   00.09.26.356  00.09.27.505   99.94   100.0    0.02
*PHYSICAL*                                      00.00.00.237                 0.00
                                                ------------  ------------   ------
 TOTAL                                          01.22.27.801  01.22.29.251             0.03
```

Partition Name & Weight match indicates IFA(s)

The ICF block contains logical processors of type ICF, IFL, IFA

zAAP Workload Reporting Samples

The WLMGL Goal and Actuals Section

```
RESPONSE TIME GOAL:    00.00.01.000 AVG

ACTUALS:  RESPONSE    EX    PERF   AVG    ----USING%----  ---------- EXECUTION DELAYS % --------- ---DLY%-- -CRYPTO%- ---
CNT%--    %

            TIME      VEL%  INDX  ADRSP    CPU  IFA  I/O   TOT  CPU  IFA  I/O  AUX  AUX SWIN      UNKN IDLE  USG  DLY  USG
DLY QUIE

        HH.MM.SS.TTT                                                    VIO PRIV

*ALL    00.00.01.854  31.8   1.1    5.7    3.6  2.2  2.6  13.4  8.5  4.3  0.3  0.2  0.1  0.1      58.1 22.7  1.1  3.1  0.0
0.0  0.0

LP1     00.00.01.999  30.5   1.3    2.1    3.5  2.1  2.3  13.5  9.0  4.2  0.2  0.1  0.0  0.1      60.3 20.5  0.2  1.1  0.0
0.0  0.0

LP2     00.00.01.001  49.4   1.0    1.9    3.3  2.0  3.1   6.8  1.8  4.3  0.3  0.1  0.0  0.0      63.1 24.0  1.4  4.1  0.0
0.0  0.0

LP4     00.00.01.003  24.1   1.0    1.8    3.9  2.3  2.3  10.1 14.8  4.3  0.3  0.1  0.4  0.1      50.3 23.7  0.3  0.3  0.0
0.0  0.0
```

New Using and Delay States

IFA Using	Work is found executing on an IFA
IFA Delay	Work is delayed for an IFA
IFA on CP Using	Contained in either CPU using samples or IFA using samples based on the dispatching priorities on selecting work from IFA work queue (for work is eligible to run on IFAs but executed on regular CPs)

Do I Need zAAP?

- Do you have considerable Java workload
- Anticipate growth in Java workloads
- Need more capacity
- Consolidating Java workloads
- Need to lower the cost of running Java workloads

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*
- Prior to z/OS V1R6 with SDK 1.3.1
  - Instrumented SDK 1.3.1 SR24 and the Excel workbook
    - Must be at PTF UQ94379
  - Instrumented SDK 1.3.1 SR22 with zAAP Projection Tool and the Excel workbook
    - Must be at PTF UQ84703 level
    - Must enable the Projection Tool for each address space

WP100431 - Obtaining the zAAP Usage Estimation Information in WebSphere for z/OS Version 5
WP100417 - z/OS Performance: Capacity Planning Considerations for zAAP Processors

Not All Java Applications Are Created Equal

- Some are good candidates
  - Heavy Java
- Some aren't good candidates
  - Light weigh Java
- 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

zAAP Projection Tool

- ~~SDK 1.3.1 SR22~~
  - ~~Must be at PTF UQ84703 level~~
- SDK 1.3.1 SR24: Projection Tool is integrated
  - SR24 is recommended for zAAP eligibility projection exercise
- Integrated as part of SDK 1.4 product
- Writes Output Message every 5 minutes
  - The information on processor time is provided as messages in standard out for the SDK which is available in the z/OS JOBLOG file
- A spreadsheet summarization tool is available to assist in the analysis
  of the zAAP Projection
  - zAAP projection tool workbook.xls (reads from JOBLOG)
    https://www6.software.ibm.com/dl/zosjava2/zosjava2-p

WP100431 Installation of the zAAP Projection Tool Instrumented SDK in WebSphere for z/OS V5
http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100431

Projection Tool Sample Output

IFA Projection data for system id=<SYSD.50594238> Starting at: 18:31:30 - Current address space CPU: 0.008068 sec.
<SYSD.50594238> Interval at: 18:36:30 Switches To/From IFA: 3717251 Java IFA: 99.3 sec. Java Standard CPU 101.86 sec. Interval address space CPU: 208.66 sec.
<SYSD.50594238> Interval at: 18:41:30 Switches To/From IFA: 3903114 Java IFA: 104.27 sec. Java Standard CPU 106.95 sec. Interval address space CPU: 219.09 sec.
<SYSD.50594238> Interval at: 18:46:30 Switches To/From IFA: 4176332 Java IFA: 111.57 sec. Java Standard CPU 114.44 sec. Interval address space CPU: 234.43 sec.
<SYSD.50594238> Interval at: 18:51:30 Switches To/From IFA: 3842225 Java IFA: 102.64 sec. Java Standard CPU 105.28 sec. Interval address space CPU: 215.68 sec.
<SYSD.50594238> TOTAL at: 18:51:30 Switches To/From IFA: 15638922 Java IFA: 418 sec. Java Standard CPU 429 sec. Total address space CPU: 878 sec.

Switches To/From IFA	State changes in processing of zAAP eligible work vs none eligible work for all Java threads in address space
Java IFA	Time accumulated for Java threads processing zAAP eligible work
Java Standard CPU	Time accumulated for Java threads processing none zAAP eligible work
Interval Address Space CPU	Time for all work in the Java address space across all dispatchable units including the Java threads

Excel Worksheet Example

SMF name	Instance or Group	RMF interval start	zAAP	CP	Space	%Time zAAP eligible	zAAP% engine eligible	Other Java% engine	Appl% engine	zAAP% w/capt ratio	ZAAPs w/wait
			Service Class		newwork		all LPARS			85%	75%
SYSD	test1	18:31:00	99	102	209	48%	33%	34%	70%	39%	52%
SYSD	test1	18:36:00	104	107	219	48%	35%	36%	73%	41%	55%
SYSD	test1	18:41:00	112	114	234	48%	37%	38%	78%	44%	58%
SYSD	test1	18:46:00	103	105	216	48%	34%	35%	72%	40%	54%

- Seconds of zAAP eligible processing, non zAAP-eligible (standard CP) processing, and total address space time for the JAVA space(s).
- Combines data from multiple address spaces (JVMs), service classes and LPARs
- Combines the data and aligns to intervals such as the RMF interval used.
- Ability to adjust zAAP utilization factoring in z/OS capture ratios
- zAAP and standard CP time expressed as a percent of the engine (single CP) that the data was collected on. This can be used as input to the projected number of zAAPs needed factoring in a target maximum utilization to ensure workload responsiveness

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 available since **01/05**
- IBM does not impose or impact software charges on zAAP capacity for the IBM WebSphere Application Server
- 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
- You should contact your ISVs directly to determine if their charges will be affected by zAAP
- Don't plan 100% busy time for zAAP engines as you would with the standard CPs

zAAP Enablement Resources Summary

- Performance White Paper
  - http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100417
- Techdoc WP100431 Installing zAAP Projection too
  - http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100431
- Techdoc WP100489 Mission: zAAP your costs Running WebSphere and Java on the zSeries Application Assist Processor
  - http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100489
- IBM Redbook on zAAP: SG24-6386
- IBM zAAP site
  - ibm.com/zseries/zaap
  - Frequently Asked Questions
  - Customer Brochure
- z990 Latest Enhancements Announcement Letter (104-118).
- z890 Announcement Letter (104-117).

w3.ibm.com to learn more on zAAP, Please …Search on "zAAP"

Enterprise IT management SYMPOSIUM

Application Development & integration SYMPOSIUM

SHARE User-Driven Training Event & Expo

Enterprise Systems management SYMPOSIUM

USERBLUE