



IBM Advanced Technical Support

# Everything zAAP!

Kathy Walsh  
Washington Systems Center  
Advanced Technical Support

# Trademarks and Disclaimers

AIX*	HiperSockets	S/390
AIX 5L*	IBM*	ServerProven*
BladeCenter	IBM eServer	Tivoli*
Chipkill	IBM logo*	TotalStorage*
DB2	iSeries	TotalStorage Proven
DB2 Universal Database	Lotus*	Virtualization Engine
Domino	OnForever*	X-Architecture
Enterprise Storage Server*	Parallel Sysplex*	xSeries*
e-business logo*	POWER	z/OS*
GDPS*	POWER5	z/VM*
Geographically Dispersed Parallel Sysplex	Predictive Failure Analysis*	zSeries*
	pSeries*	System z9*

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries or both.

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

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

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

\* 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


- What is a zAAP and how does it work
- New Performance Data and RMF Changes
- Capacity Planning
- Chargebacks



# The New zSeries Application Assist Processor (zAAP)

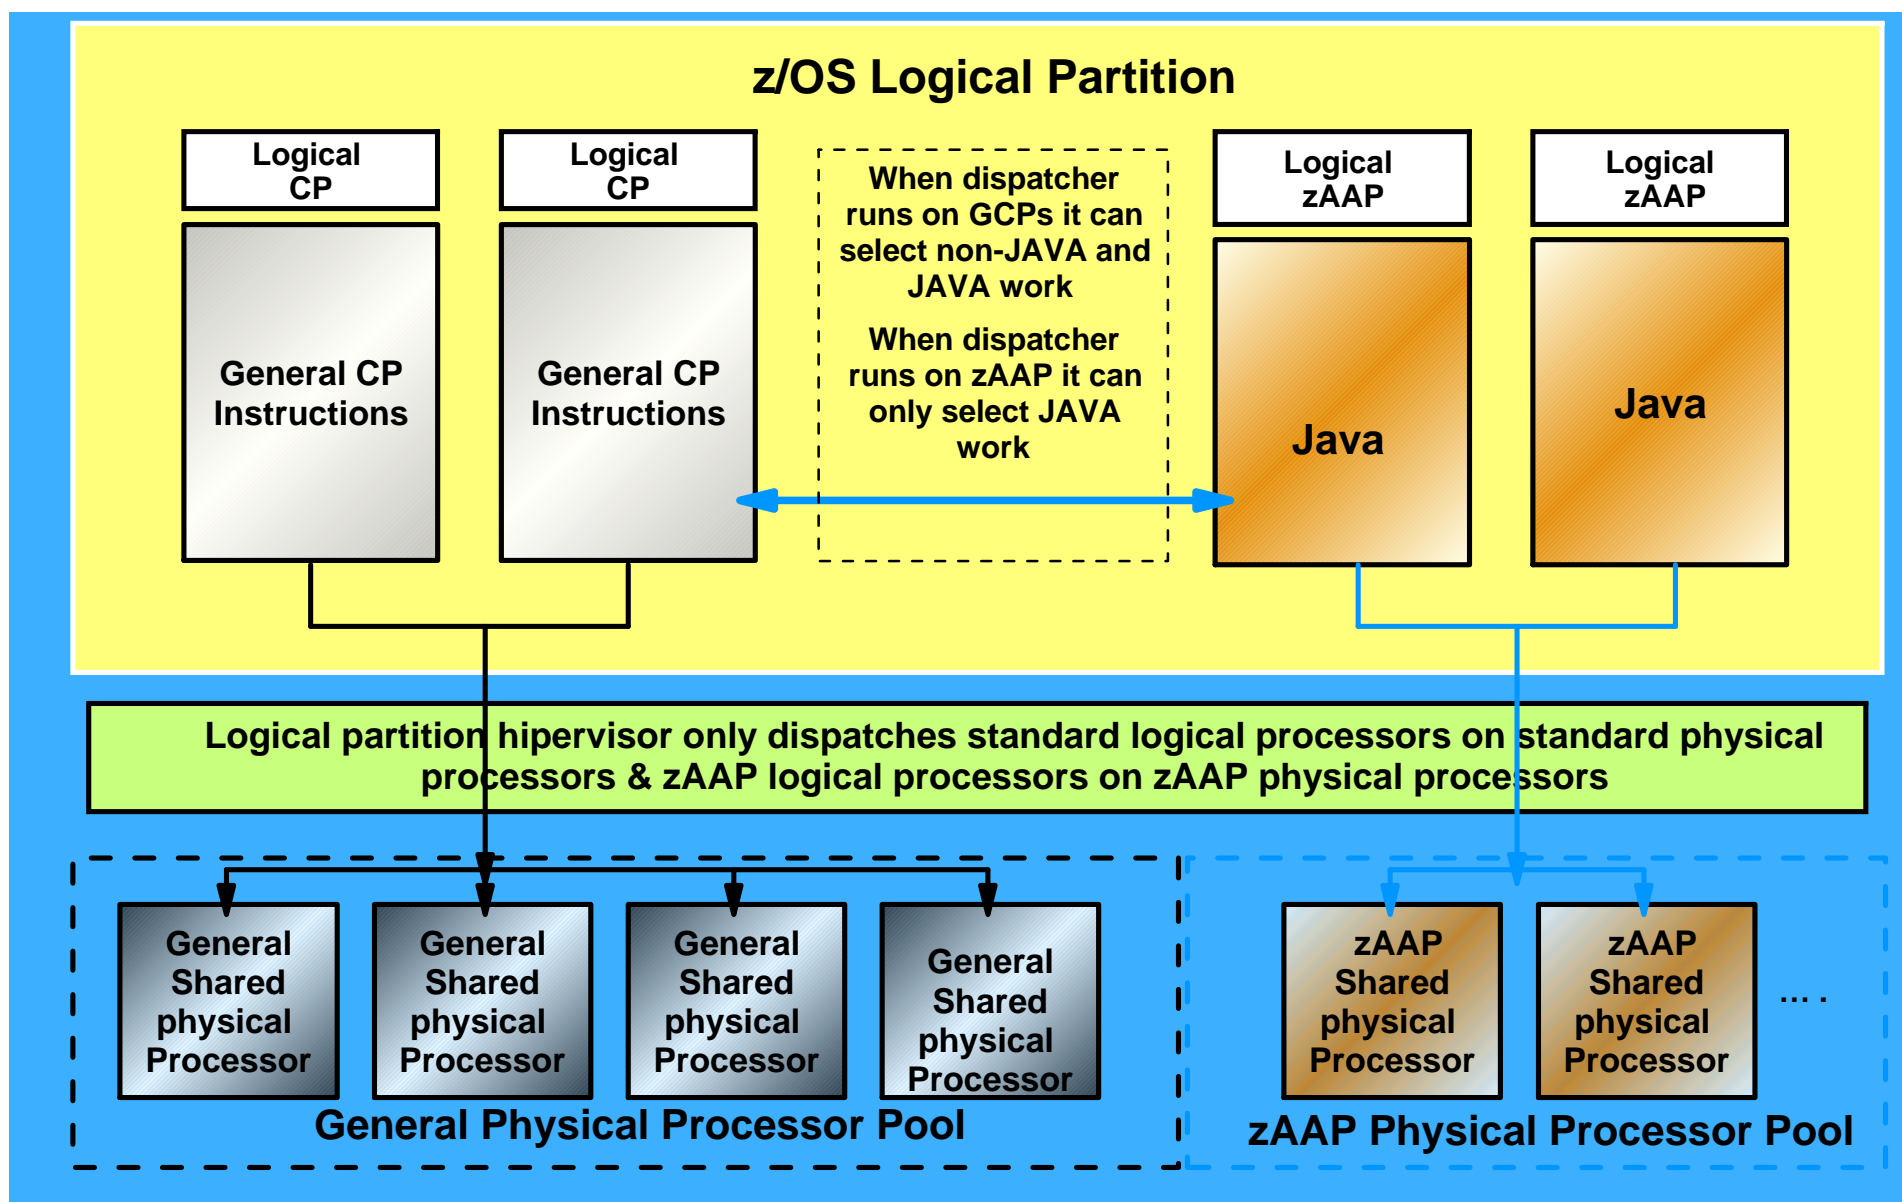
***New specialty assist processor dedicated exclusively to execution of Java workloads under z/OS® – e.g. WebSphere®, CICS, IMS, DB2***

- Available on IBM Server™ zSeries® 990 (z990), zSeries 890 (z890) and IBM System z9 servers
- Used by workloads with Java cycles, e.g. WebSphere, DB2®
  - Executes Java code with no changes to applications
- Attractively priced, much lower than standard CPs
- Significantly lower maintenance costs than standard CPs
- Traditional IBM System z software charges unaffected
- Sub-capacity eligible IBM software charges can be reduced
- Up to 1 zAAP per general purpose processor in a CEC



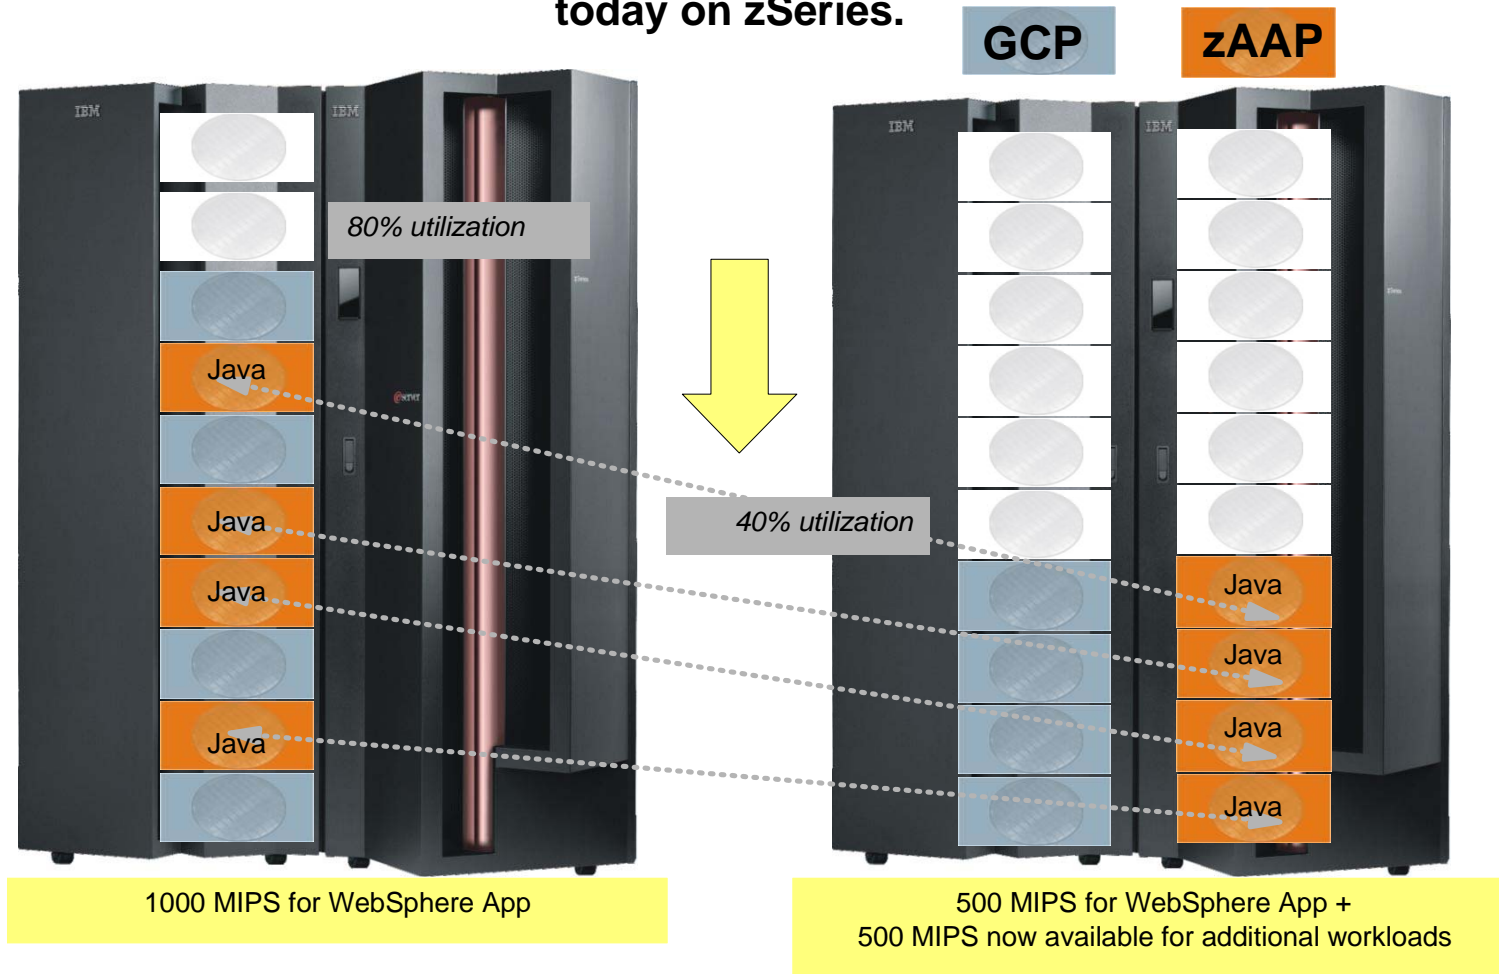
***Objective: Enable integration of new Java based Web applications with core z/OS backend database environment for high performance, reliability, availability, security, and lower total cost of ownership***

# zAAP Technical Overview: z/OS zAAP Partition



# zAAP Overview: A Simplified Example...

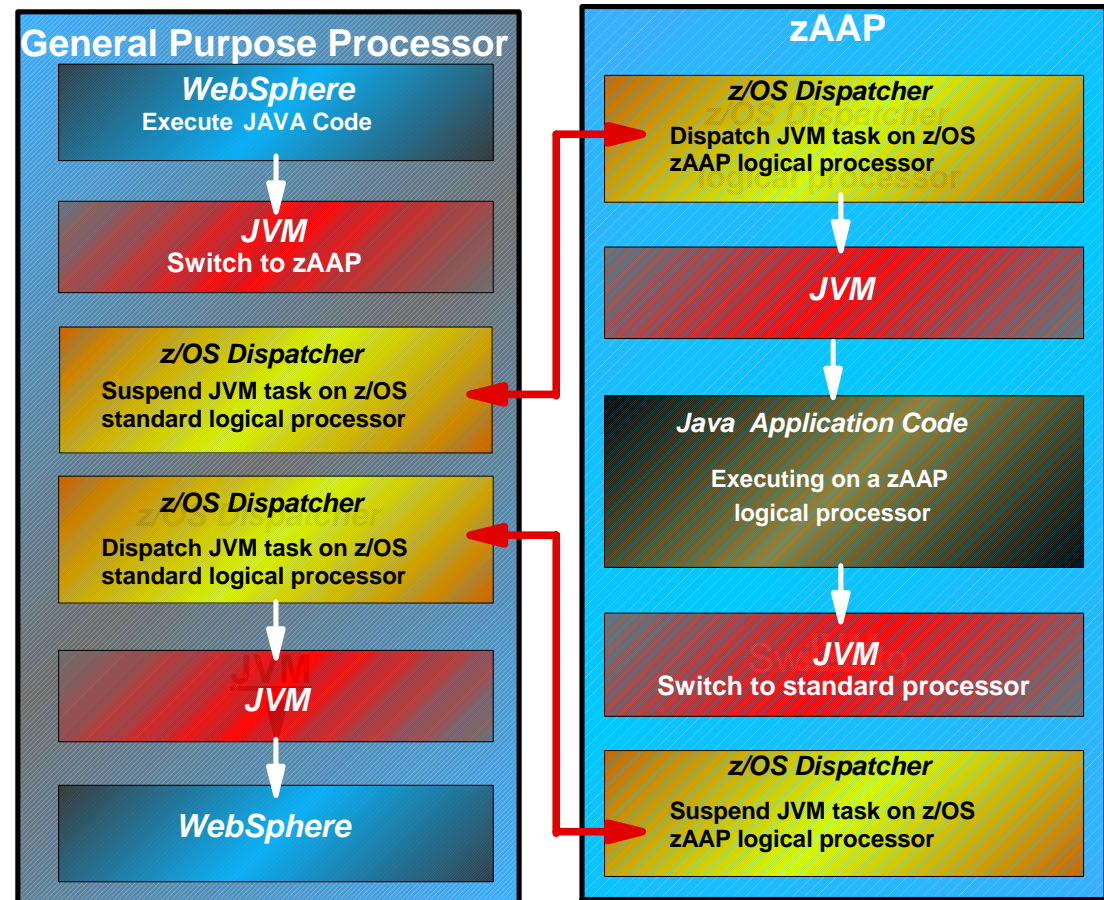
Consider a WebSphere Application that is transactional in nature and requires 1000 MIPS today on zSeries.



***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***

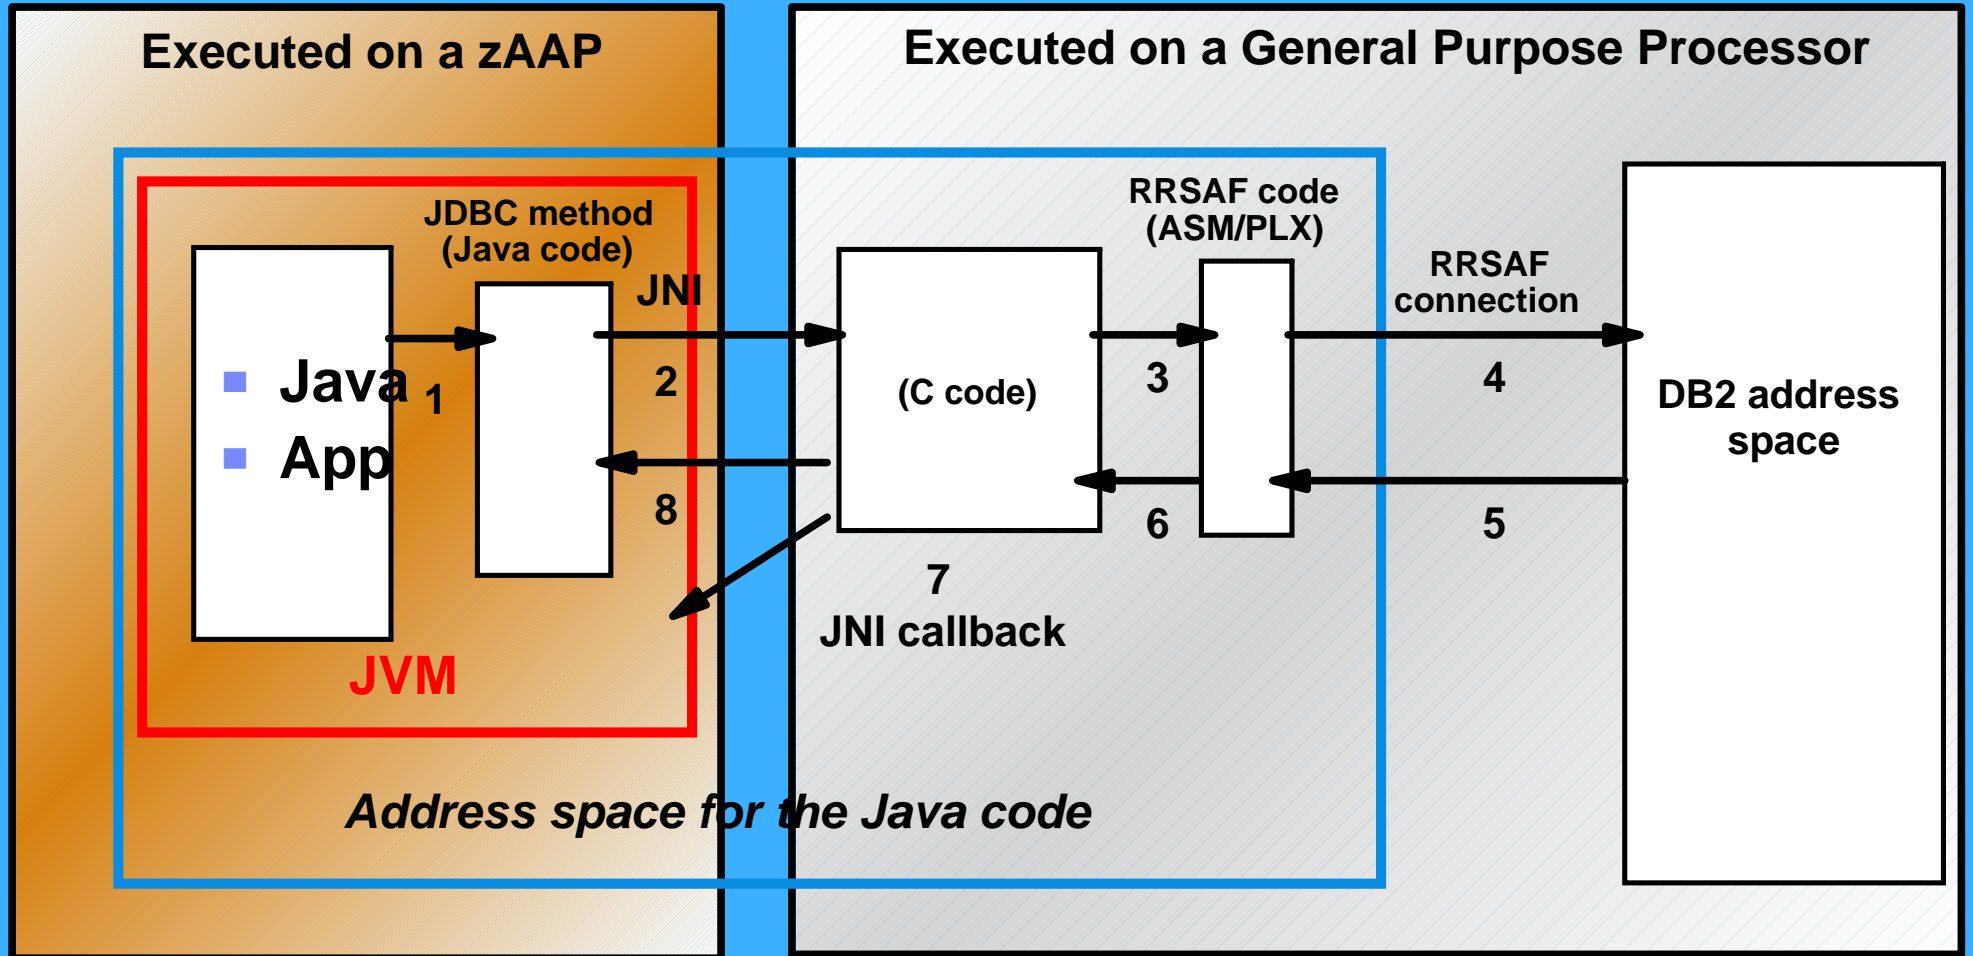
# zAAP Workflow: Executing Java under IBM JVM control

- 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 general purpose processors
- There are installation controls to limit the use of general purpose processors to execute zAAP eligible work



# zAAP Integration at Work: Java App calling DB2

## z/OS Logical Partition





## For IBM System z9 Processors

- Multiple Specialty (ICF) Pools Managed Independently
- Example: z9 processor with 10 CPs, 1 ICF, 2 IFLs, 3 zAAPs
  - ▶ CP Pool contains 10 engines
  - ▶ ICF pool contains 1 ICF
  - ▶ IFL pool contains 2 IFLs
  - ▶ zAAP pool contains 3 zAAPs
    - zAAP weight can be independently set from z/OS partition

# For @server zSeries z990 and z890 processors

## 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
  
- The ICF pool partition weights need to be updated to reflect the introduction of the zAAP

## IBM System z9 and eServer zSeries and zAAPs

- Partitions are created under the cover for LPARs which have zAAPs defined but these partitions do not contribute to the overall upper limit of the number of partitions on the CEC
- A zAAP PU failure will use an available spare to provide high availability
- Can define partition profiles with Initial and Reserved CPs
  - ▶ Use the z/OS command CF CPU=(cpid) to bring them onto the partition
- Capacity Backup (CBU)
  - ▶ System z9 only
- Capacity Upgrade on Demand (CUoD)
  - ▶ z990, z890
  - ▶ z9

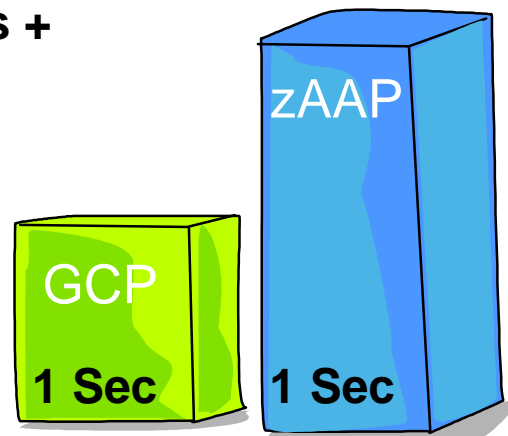
## Sub-Capacity GCPs and zAAPs

- zAAPs always run at full speed of processor model
- Sub-capacity GCPs are available on the z890, the z9 BC, and z9 EC processors
- Same z/OS image has CPs running at different speeds
  - ▶ Requires CPU seconds to be normalized

Example: zAAP is 8 times the speed of the GCP

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

**Execution Time = 1 second + (1 second \* (8)) = 8 seconds = 9 seconds**



- Normalization factor used is in RMF 72 subtype 3 record, R723NFFI
- Normalization factor used is in the SMF 30 record, SMF30ZNF
- When zAAP/zIIP and GCP are the same speed the normalization factor resolves to 1

# CPU Activity Report for with zAAPs - IBM System z9

z/OS V1R7

SYSTEM ID SYSD  
RPT VERSION V1R7 RMF

DATE 08/18/2006  
TIME 16.23.00

INTERVAL 00.59.999  
CYCLE 1.000

CPU	MODEL	750	H/W MODEL	S54					
---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	2.50		2.41		04B10E	27.15	0.49
1	CP	100.00	1.99		1.92		04B10E	24.37	0.14
2	CP	100.00	1.85		1.77		04B10E	38.32	0.04
3	CP	100.00	1.68		1.58		04B10E	28.18	0.12
CP	TOTAL/AVERAGE		2.01		1.92			118.0	0.18
6	AAP	100.00	40.19		40.96		04B10E		
7	AAP	100.00	0.31		0.29		04B10E		
AAP	AVERAGE		20.25		20.63				
8	IIP	100.00	0.29		0.17		04B10E		
9	IIP	100.00	0.00		0.00		04B10E		
IIP	AVERAGE		0.15		0.08				

No I/O Interrupts  
on zAAPs

# IBM System z9 Partition Data Report with zAAPs

PARTITION DATA							LOGICAL PARTITION PROCESSOR DATA		-- AVE PROC UTILIZATION PERCENT				
NAME	S	WGT	DEF	ACT	PROCESSOR- NUM	TYPE	DISPATCH TIME DATA		LOGICAL PROCESSORS		PHYSICAL PROC		
							EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR	MGMT	EFFECTIVE
SOSP4	A	500	0	4	4	CP	00.00.04.349	00.00.04.814	1.81	2.01	0.02	0.14	
....													
SOSP15	A	10	0	2	2	CP	00.00.01.593	00.00.02.541	1.33	2.12	0.03	0.05	
SOSP16	A	10	0	4	2	CP	00.00.04.305	00.00.05.208	3.59	4.34	0.03	0.14	
SOSP17	A	10	0	3	2	CP	00.00.02.671	00.00.03.614	2.23	3.01	0.03	0.09	
*PHYSICAL*								00.00.25.703			0.86		
TOTAL							00.20.33.086	00.21.04.443			1.05	41.10	
SOSP4	A	500			2	AAP	00.00.24.118	00.00.24.295	20.10	20.25	0.15	20.10	
SOSP15	A	10			1	AAP	00.00.00.353	00.00.00.876	0.59	1.46	0.44	0.29	
SOSP16	A	10			1	AAP	00.00.01.772	00.00.02.282	2.95	3.80	0.43	1.48	
SOSP17	A	10			1	AAP	00.00.01.357	00.00.01.875	2.26	3.13	0.43	1.13	
*PHYSICAL*								00.00.06.608			5.51		
TOTAL							00.00.27.601	00.00.35.939			6.95	23.00	
SOSP4	A	500			2	IIP	00.00.00.058	00.00.00.174	0.05	0.15	0.10	0.05	
*PHYSICAL*								00.00.02.002			1.67		
TOTAL							00.00.00.058	00.00.02.176			1.76	0.05	

\* RMF report has been modified to fit into the space

# zAAP z990 CPU Time Reporting

z/OS V1R6

SYSTEM ID SYSD  
RPT VERSION V1R5 RMF

DATE 09/09/2004  
TIME 18.55.00

CPU	2084	MODEL	314				
---	CPU---	ONLINE	TIME	LPAR	BUSY	MVS	BUSY
NUM	TYPE	PERCENTAGE	TIME PERC	TIME PERC	TIME PERC	CPU SERIAL	I/O TOTAL
0	CP	100.00	61.07	61.30	043A6A	4.04	
1	CP	100.00	77.84	78.20	043A6A	5.51	
CP	TOTAL/AVERAGE		69.45	69.75		9.55	
2	IFA	100.00	24.75	24.86	043A6A		
3	IFA	33.23	13.73	14.64	043A6A		
IFA	AVERAGE		22.00	22.31			

No I/O Interrupts on zAAPs

## PARTITION DATA REPORT

MVS PARTITION NAME	OOSP4	NUMBER OF PHYSICAL PROCESSORS	CP	ICF									
IMAGE CAPACITY	497		16	14									
NUMBER OF CONFIGURED PARTITIONS	30			2									
----- PARTITION DATA -----													
NAME	S	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
OOSP4	A	30	0	69	NO	0.0	2.0	CP	69.07	69.45	0.05	9.87	9.92
OOSPA	A	10	0	3	NO	0.0	2	CP	2.40	2.67	0.04	0.34	0.38
*PHYSICAL*											4.34		4.34
TOTAL											5.27	31.07	36.34
OOSP4	A	30					2	ICF	21.79	22.00	0.14	14.52	14.66
*PHYSICAL*											2.56		2.56
TOTAL											2.71	14.52	17.22

# zAAP CPU Times - RMF 72 Workload Activity

REPORT BY: POLICY=WLMPOL		WORKLOAD=BAT_WKL		SERVICE CLASS=BATSPEC		RESOURCE GROUP=BATMAXRG	
TRANSACTIONS	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---SERVICE---	SERVICE TIMES	---APPL %---
AVG	0.98	ACTUAL	6.520	SSCHRT	11.5	IOC 8326	<u>CPU</u> 24.7 <u>CP</u> 0.97
MPL	0.98	EXECUTION	6.128	RESP	7.0	<u>CPU</u> 662386	SRB 0.0 <u>AAPCP</u> 0.01
ENDED	10	QUEUED	391	CONN	6.9	MSO 0	RCT 0.0 IIPCP 0.00
END/S	0.17	R/S AFFIN	0	DISC	0.0	SRB 965	IIT 0.0
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.1	TOT 671677	HST 0.0 <u>AAP</u> 40.27
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC 11195	<u>AAP</u> 24.2 IIP 0.00
AVG ENC	0.00	STD DEV	0				IIP 0.0

GOAL: EXECUTION VELOCITY 35.0% VELOCITY MIGRATION: I/O MGMT 99.2% INIT MGMT 92.2%

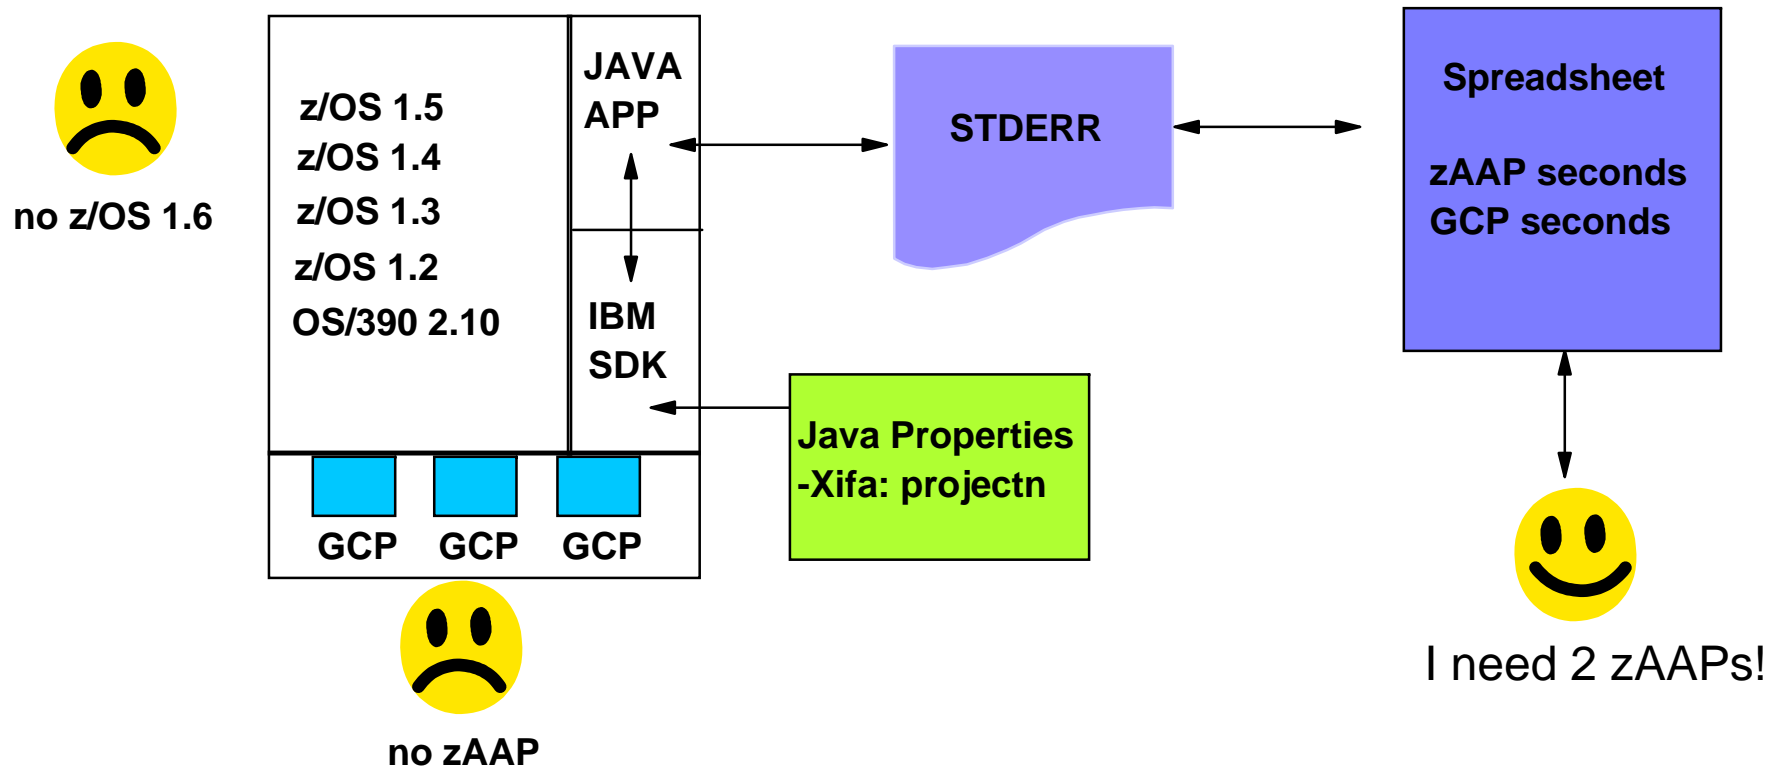
SYSTEM	RESPONSE TIME	EX VEL%	PERF INDX	AVG ADRSP	----- USING% -----	----- EXECUTION DELAYS % -----
					CPU <u>AAP</u> IIP I/O TOT CPU	
SYSD	--N/A--	99.2	0.4	1.0	0.8 45.9 0.0 3.9 0.4 0.4	

- CPU** - Includes GCP seconds and normalized zAAP seconds
- SRB** - Includes GCP seconds and normalized zAAP seconds
- AAP** - Includes only zAAP seconds
- APPL% CP** - Includes only non-zAAP (GCP) time
- APPL% AAPCP** - Time spent on a GCP which was zAAP eligible (Subset of APPL% CP)
- APPL% AAP** - Time spent on the zAAP
  
- CPU SERVICE** - Includes CPU Service Units accumulated on a zAAP



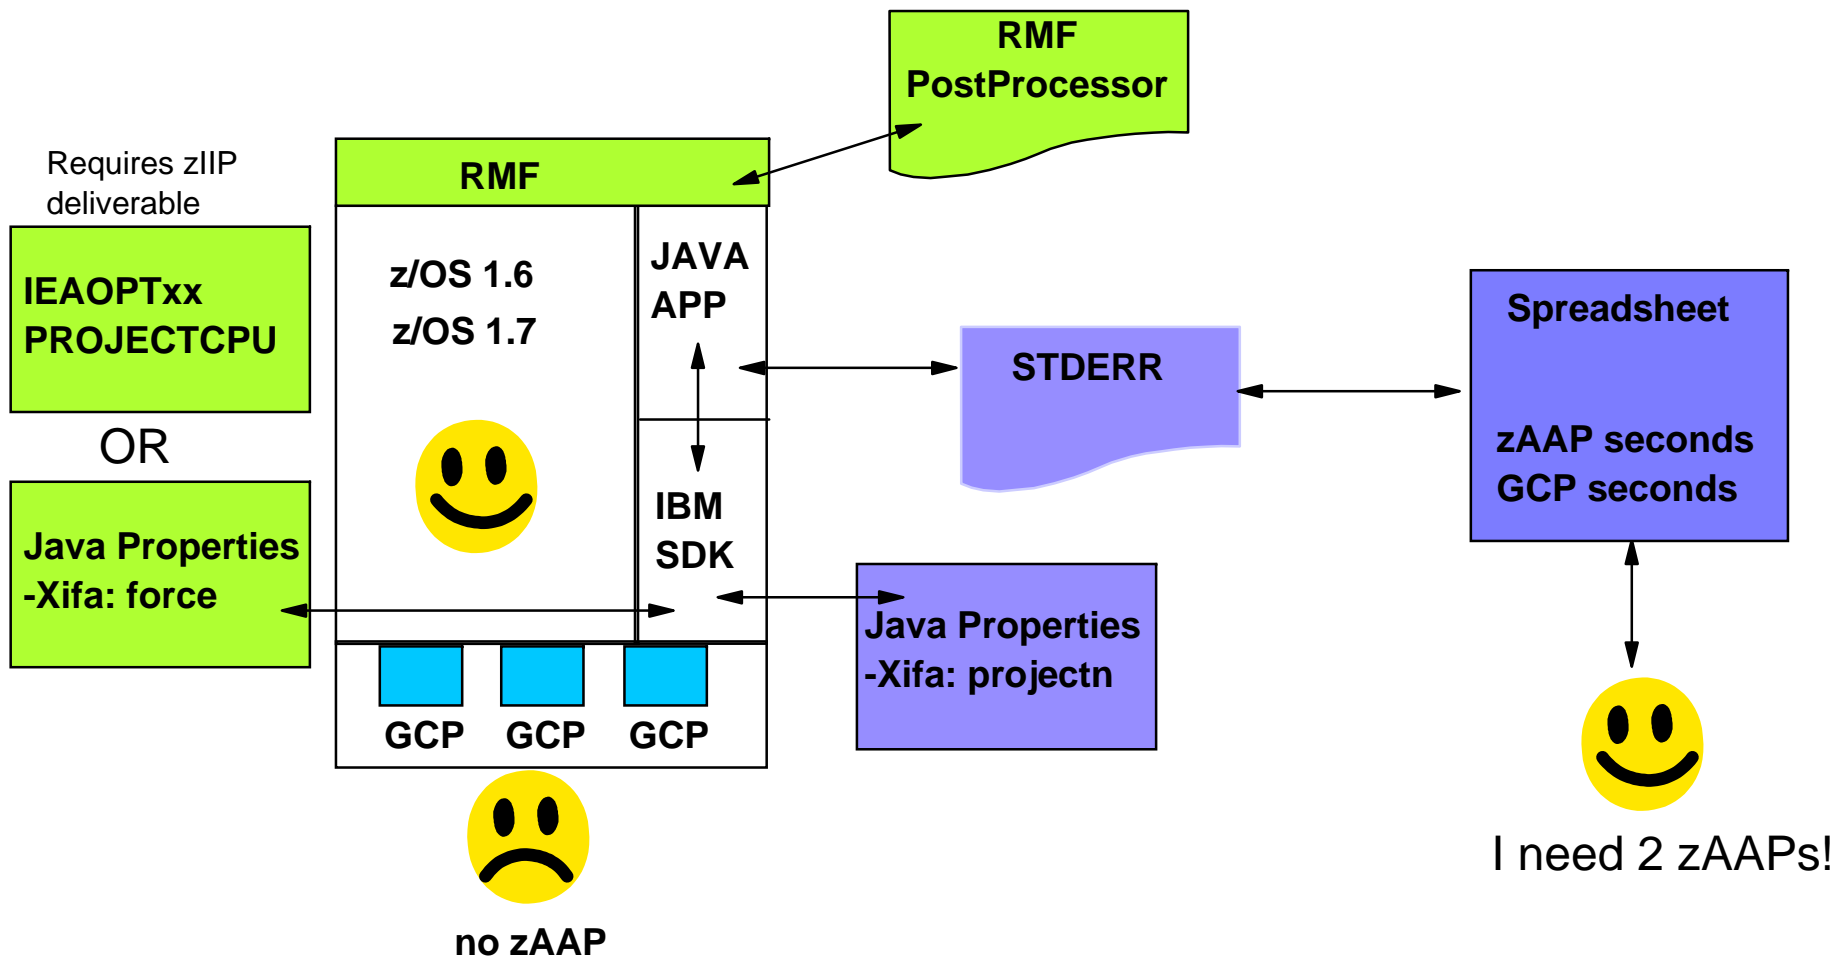
# Update Capacity Planning Prior to z/OS 1.6

- Evaluate applicability of a zAAP in the future
  - ▶ No RMF information available yet
  - ▶ Need to identify how much CPU time a workload is eligible to run on a zAAP



# Update Capacity Planning After z/OS 1.6

- Evaluate applicability of a zAAP in the future
  - ▶ Use RMF information to determine zAAP capacity



## zAAP Estimation Process

- Update the JVM to the right PTF levels
- Download the zAAP estimation tool
- Update the JVM option files to use instrumented JVM
- Run Test
  - ▶ Results will be sensitive to the contents of the test, make it as real as possible
- Download the STDERR file from the address space to a PC
  - ▶ Can edit the file to have just active intervals
- Process the flat file using the IBM supplied Excel workbook

## Subsystem and Minimum Java Levels for zAAP Estimation

Subsystem Version	IBM SDK for OS/390, Java 2 Technology Edition, V1.3.1 at PTF level UQ94379 (SR24 or later)	IBM SDK for z/OS, Java 2 Technology Edition, V1.4, at PTF level UQ88783 or later
WAS V4*, WAS V5.0*	X	
<b>WAS 5.1**</b>		X
<b>IMS™ V7**</b>	X	X
<b>IMS V8**</b>	X	X
<b>IMS V9**</b>		X
CICS® 2.2*	X	
<b>CICS 2.3**</b>		X
<b>DB2® V7**</b>	X	X
<b>DB2 V8**</b>	X	X

\* Configurations which support zAAPs

\*\* zAAP projection tool can be used to assist in capacity planning

## What are the Tuning Options?

- Specified in IEAOPTxx member of parmlib
  - ▶ IFAHonorPriority=YES|No
  - ▶ IFACrossOver=YES|No
- Can have an impact on the:
  - ▶ Performance of Java work
  - ▶ The capacity plan for the zAAP



Multiple Changes ....

All for the better

## zAAP IEAOPTxx Options have Changed

- Java eligible work tends to be high priority (e.g. Websphere, CICS)
- GCPs tend to outnumber zAAPs (3:1 or 4:1)
  - ▶ With IFACrossover=YES, GCP tend to take most of the zAAP work, hurting the effectiveness of the solution
  - ▶ With IFACrossover=NO, spikes in Java workloads can saturate the zAAP, introducing response time issues when available GCP capacity exists
- Need to provide a change which allows zAAPs to process all available work, and indicate to GCPs when **"they need help"**
  - ▶ Changes in z/OS 1.6 and later releases via APAR OA14131
  - ▶ WSC Flash 10432 for more information
    - [www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)



Changed one more time....

## New zIIP / zAAP SYS1.Parmlib IEAOPTxx Support

- Available with z/OS 1.8 and for z/OS 1.6 and 1.7 as part of the zIIP FMID
  - ▶ z/OS 1.6 FMID is JBB77S9 and z/OS 1.7 FMID is JBB772S
- Three enhancements provided:
  - ▶ PROJECTCPU – Projects zIIP / zAAP consumption without hardware
    - For zAAPs no need to tailor JVM with Xifa:Projectn
  - ▶ ZIIPAWMT – Specify zIIP alternate wait management threshold
  - ▶ ZAAPAWMT – Specify zAAP alternate wait management threshold
- The IFACROSSOVER=YES option is no longer honored
  - ▶ This previous zAAP configuration option has become obsolete, due to lack of customer demand / function complexity
  - ▶ IFAHONORPRIORITY= keyword applies to zAAP work but not to zIIP work
    - All zIIPs are statically configured with the equivalent of IFAHONORPRIORITY=YES (and no IFACROSSOVER)

# Parameter Options Summary

## ★ zAAP

### IEAOPTxx

PROJECTCPU=YES|NO

ZAAPAWMT=

IFAHONORPRIORITY=YES|NO

### JVM Options (Application Level)

-Xifa:projectn

-Xifa:ON|OFF

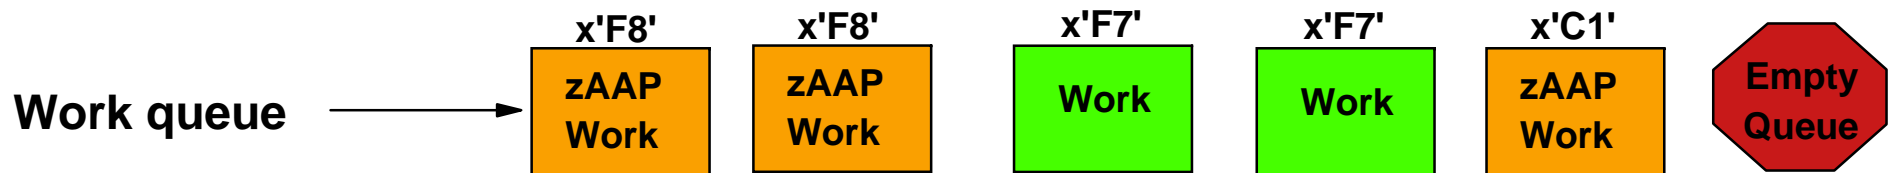
-Xifa:force

Application selection is probably not needed but can be obtained with use of  
-Xifa:Off

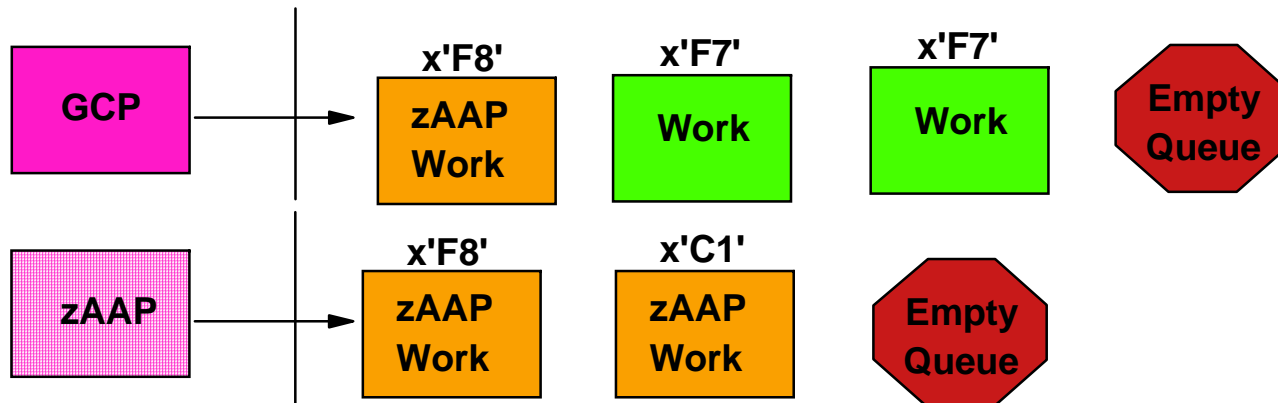


# IFAHONORPRIORITY (HP)

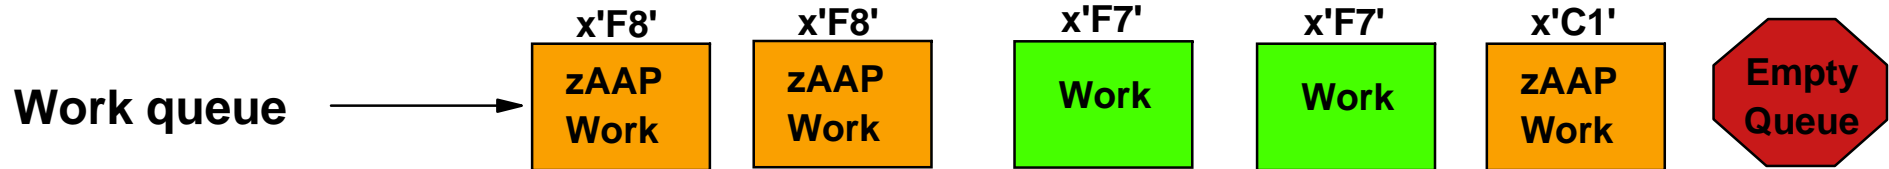
- Refers to WLM set dispatch priority
- Describes the behavior of the general purpose CPs
  - ▶ zAAPs always process work in dispatch priority order
- Alternate Wait Mgmt should not be disabled
  - ▶ IEAOPTxx parms CCCAWMT and zAAPAWMT



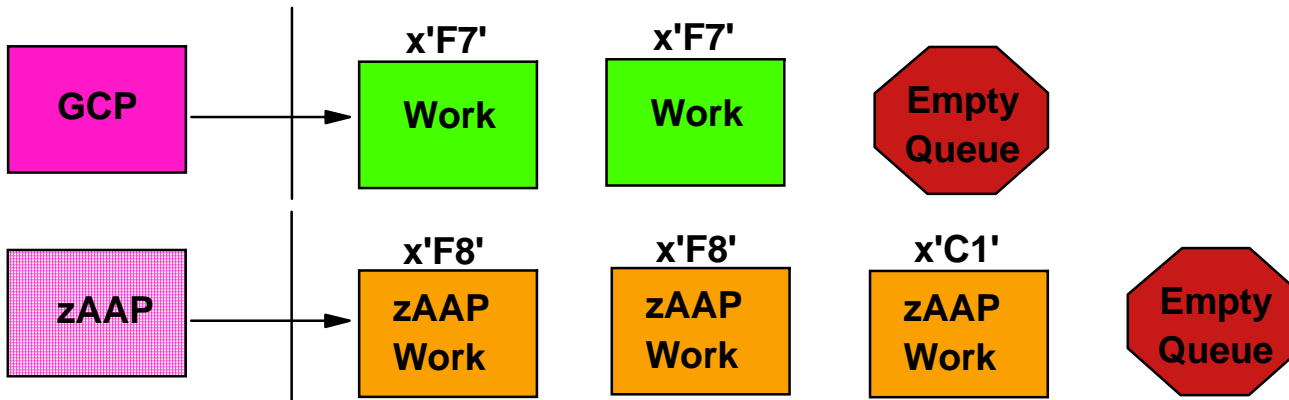
## IFAHONORPRIORITY=YES and zAAP [needs help](#)



# IFAHONORPRIORITY (HP)



IFAHONORPRIORITY=YES and zAAP [doesn't need help](#)



## Demonstration of APAR OA14131's Effectiveness

- Trade 2 LSPR workload - Websphere V5.1 - single LPAR
- HP=Yes, XO=Yes
  - ◆ Discretionary BR15 jobs are run so GCPs never run out of non zAAP work
  - ★ CP% and APPL% do not include BR15 GCP time - represents wait state

CP	IFA	ETR	ETR +/-	Resp ms	CP % ★	IFA %	Avg %	Level	APPL % CP	APPL % IFACP	APPL % IFA	% IFACP
4	4	1984		15	99	39	69	base z/OS 1.6	372	58	146	11%
4	4	2247	13%	11	96	62	79	APAR OA14131	369	10	229	2%
6	2	2360	19%	10	79	97	83	APAR OA14131	436	69	173	11%
7	1	2292	16%	11	78	98	80	APAR OA14131	519	151	90	25%

## zAAPs and WLM Resource Groups

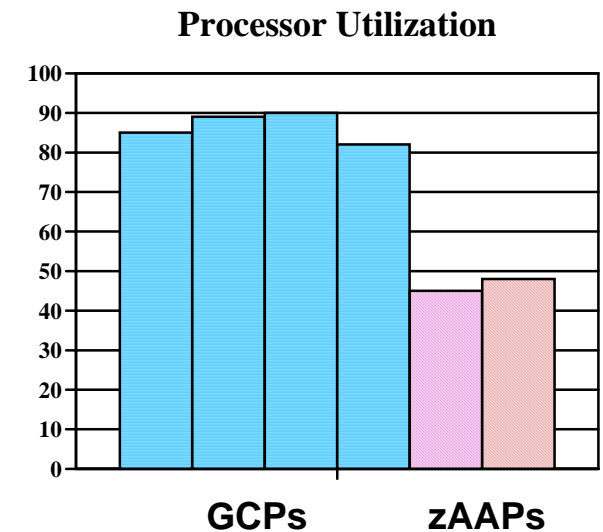
- zAAP CPU time does not contribute to the resource group maximums or minimums
  - ▶ May want to readjust resource groups which have work which is zAAP eligible to reflect changes due to zAAP offload if you want the sum of GCP and zAAP service to be the same
  - ▶ APPL% AAPonCP time will contribute to resource group targets
  
- Service class with a Resource Group Maximum which is enforced will cause work on a zAAP to also stop running
  - ▶ WLM enforces Resource Group Maximums by setting the work in the service classes non-dispatchable
  - ▶ When the address space is non-dispatchable the work on the zAAP cannot run

## zAAPs and PR/SM Hard Caps

- z990, z890
  - ▶ If the z/OS partition is hard capped the zAAP is also capped
  
- IBM System z9
  - ▶ z/OS partition GCPs and zAAP CPs are in different PR/SM pools
  - ▶ Hard cap state is set independently by pool
    - So z/OS partition can be capped to control software costs while the zAAP CP is not capped
  
- The work must still be made zAAP eligible
  - ▶ Implies work must run to some extent on a z/OS GCP and the hard cap may impact making the work eligible for a zAAP
  - ▶ Higher the switch rate between the GCP and the zAAP the greater the potential for impact

# Update Capacity Planning for zAAP Usage

- Prior to the zAAP, CPU capacity planning for z/OS dealt with only one processor type
  - ▶ RMF reported processor utilization as well as CPU consumption by service class
  
- We now have two types of processors which are used by z/OS
- New questions: How much do I need and which type should I get?
  - ▶ RMF in Version 1.6 of z/OS will report:
    - General purpose processor utilization as well as CPU consumption by service class
    - zAAP utilization as well as consumption by service class
    - CPU time run on general purpose CP but eligible to run on a zAAP (AAPCP)



# Capacity Planning Factors in a zAAP Environment

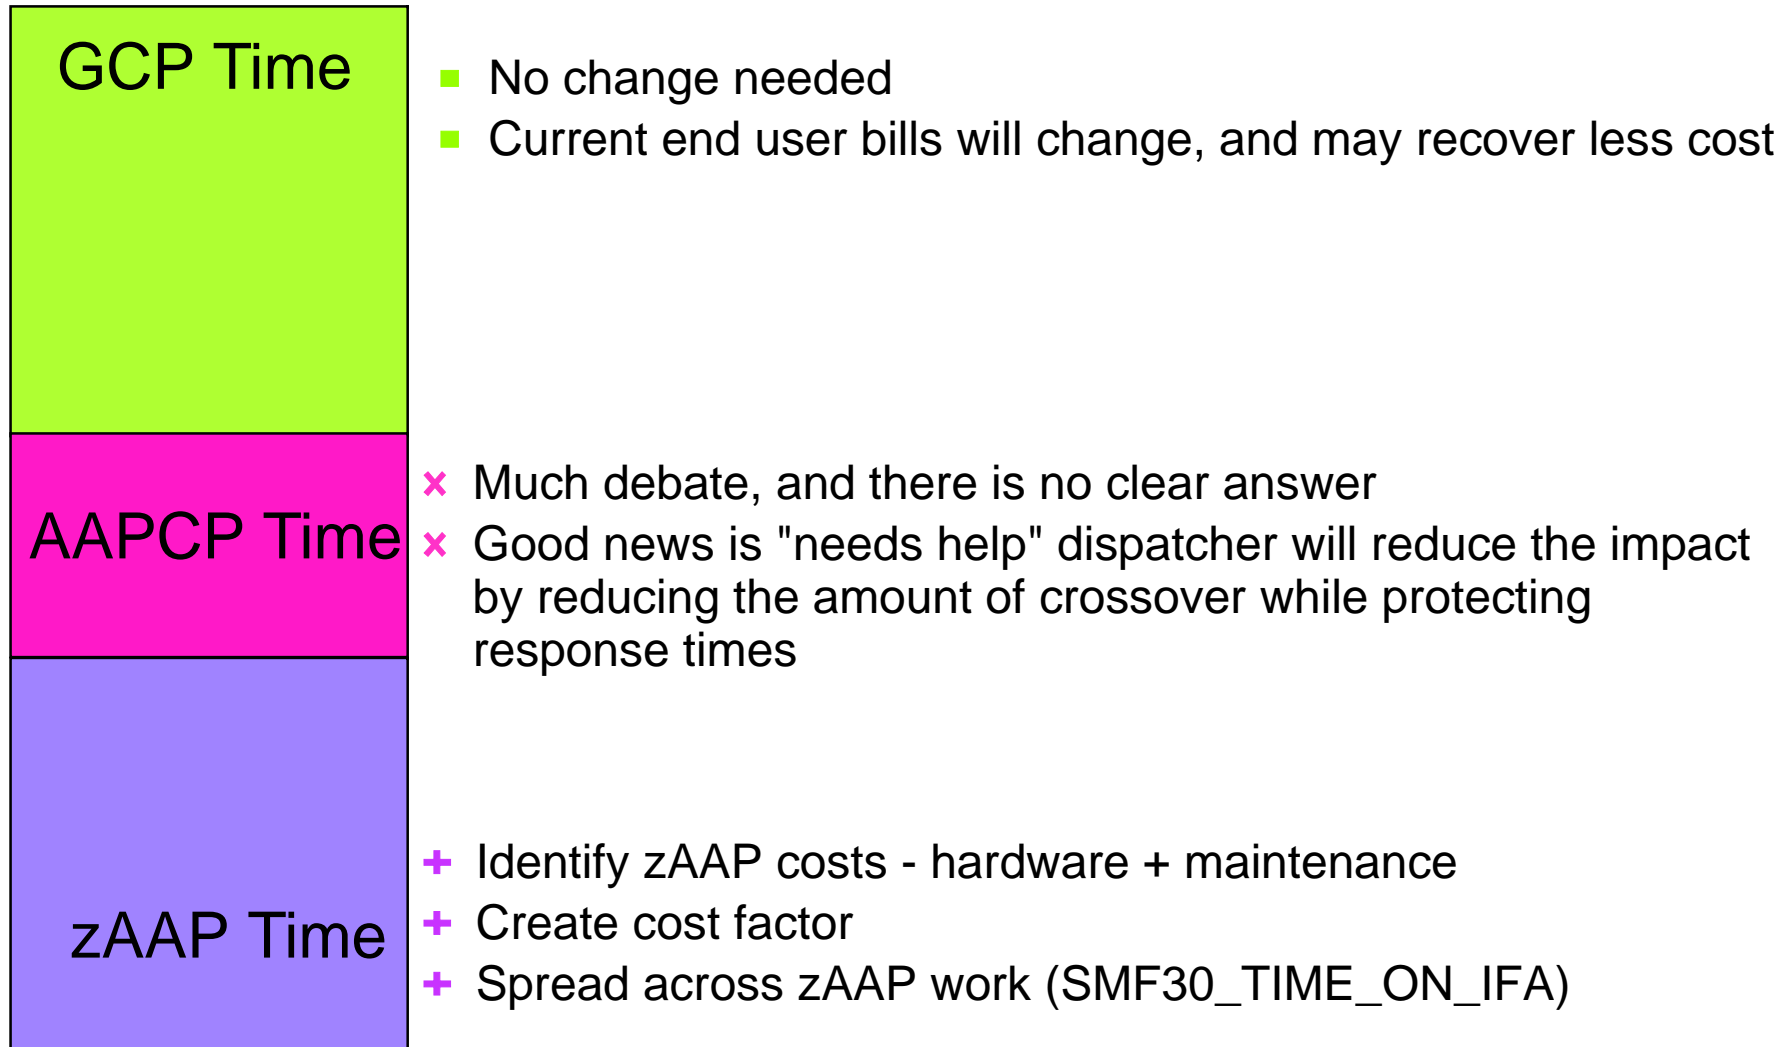
- Software
  - ▶ MP effects on z/OS
  - ▶ Switch request rate
  - ▶ Utilization of the zAAP
- Hardware
  - ▶ Number of books
  - ▶ Cost of additional partition
- Use zPCR tool to estimate the costs
  - ▶ Lightly used environments will not incur the full cost of the above planning factors
  - ▶ zPCR gives a conservative estimate
    - Being refined all the time
- Download zPCR and training on the tool at:  
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS1381>

## Accounting Methods with zAAPs

- Some facts:
  - ▶ SMF30CPT will not include zAAP time, new CPU fields for zAAPs
  - ▶ Service units will contain normalized zAAP time
    - Normalization factor is now in the SMF 30
    - Service Unit coefficients are now in the SMF 30
    - Can calculate the zAAP service units
  
- The issues:
  - ▶ How do you recover the zAAP costs
  - ▶ How do you handle impacts of crossover on billing
    - Not something the application does, it's something the system does to the application
    - Highly variable, hence not repeatable



# Chargeback Issues



# Handling IFACP - Depends on the OPT settings

JOB	Seconds Used	GCP Seconds \$6 per second	IFACP seconds	zAAP Seconds \$1 per second	Bill	Change
JOB 1	16	13	3	3	\$81	
JOB 2	20	10	0	10	\$70	
JOB 3	25	15	5	10	\$100	
JOB 1	16	16	6	0	\$96	18%
JOB 2	20	12	2	8	\$80	14%
JOB 3	25	11	1	14	\$80	25%

IFAHonorPriority=NO	IFAHonorPriority=YES
<ul style="list-style-type: none"> <li>◆ zAAP work does not run on the GCPs</li> </ul>	<ul style="list-style-type: none"> <li>◆ zAAP work runs at Priority on GCP when zAAP needs help - represents using premium GCP MIPS                             <ul style="list-style-type: none"> <li>■ GCP is a <u>safety net</u> to protect response times</li> </ul> </li> <li>◆ How to Charge - At GCP rates or zAAP rates?                             <p>cost of safety net = (cost of GCP per second - cost of zAAP per second) * AAPCP seconds</p> </li> <li>◆ If cost of the safety net &gt;= cost of amortized zAAP then buy a zAAP</li> <li>◆ If not then need to identify how to spread the cost of the safety net to recapture the cost</li> </ul>

## Summary

- zAAP has the ability to offload Java cycles from the general purpose CPs
- Estimation techniques are available to help understand the potential before hardware and software pre-reqs are met
- Estimation tool can help identify good Java candidates
- New "needs help" dispatcher improves zAAP utilization while protecting response times of high priority work
- Capacity planning methods need to be updated