



IBM Advanced Technical Support

Everything zAAP!

zSTSU, 2005

Kathy Walsh
IBM
Advanced Technical Support
Washington Systems Center



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)

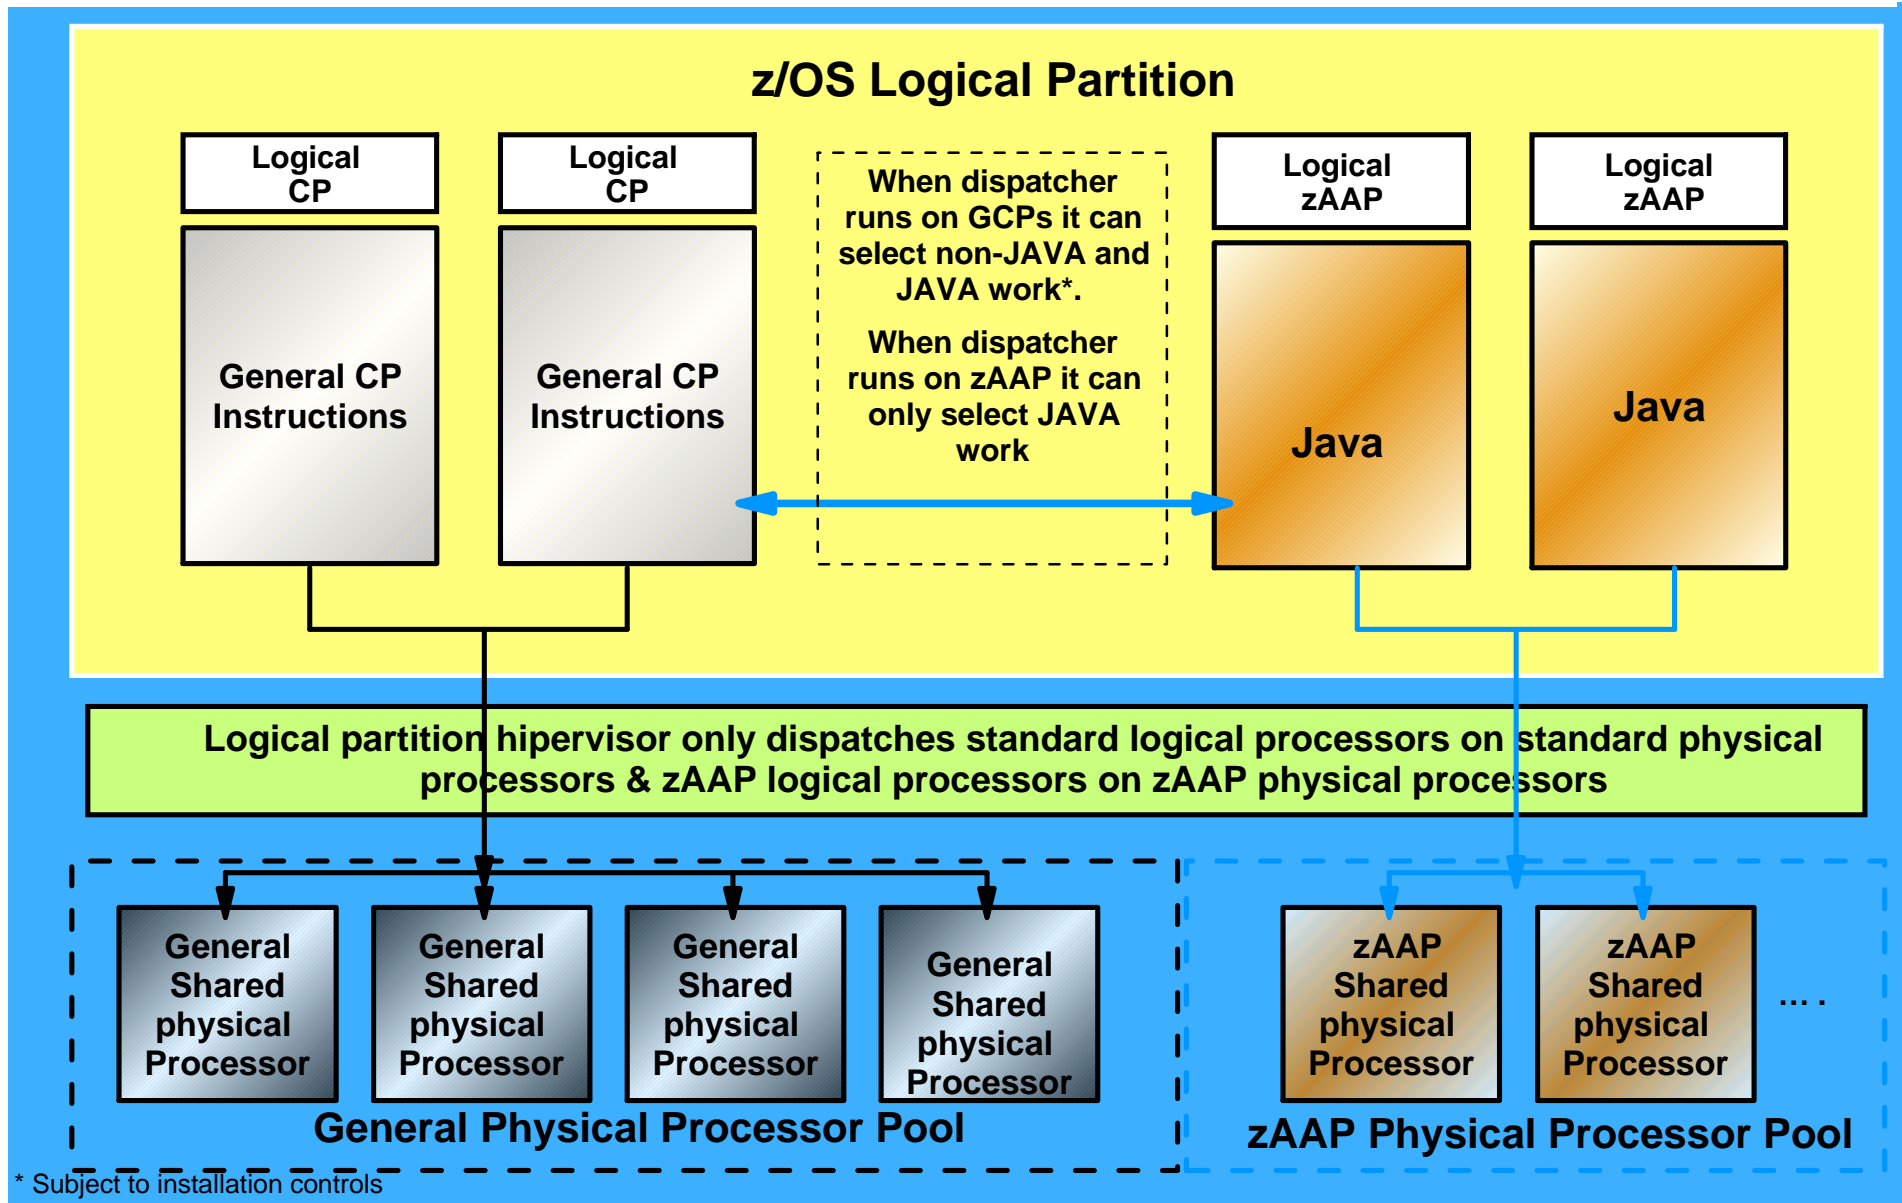
New specialty assist processor dedicated exclusively to execution of Java workloads under z/OS® – e.g. WebSphere®, CICS, IMS, DB2

- Available on IBM Server™ zSeries® 990 (z990), zSeries 890 (z890) and System z9-109 servers
- 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
- Traditional IBM zSeries software charges unaffected
- Sub-capacity eligible IBM software charges can be reduced
- Up to 1 zAAP per general purpose processor in a CEC
- zAAP feature available now



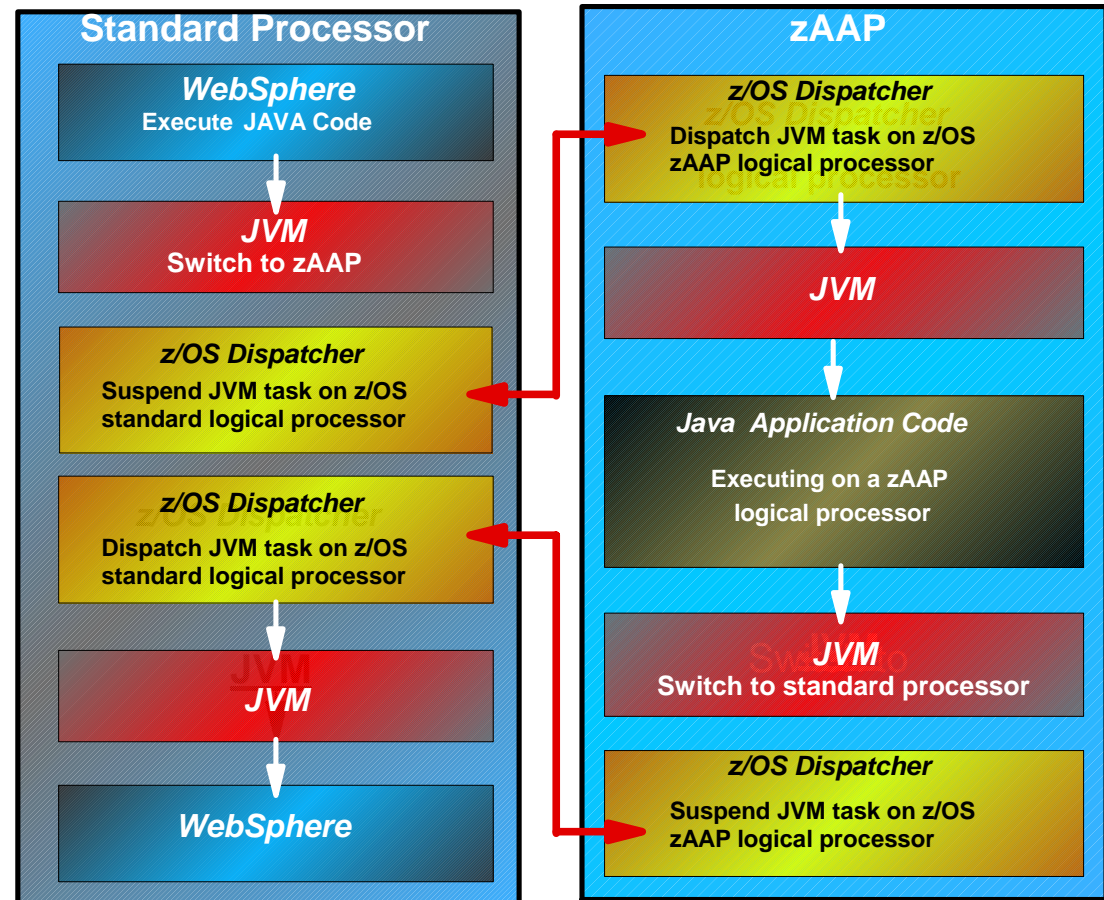
Objective: Enable integration of new Java based Web applications with core z/OS backend database environment for high performance, reliability, availability, security, and lower total cost of ownership

zAAP Technical Overview: z/OS zAAP Partition



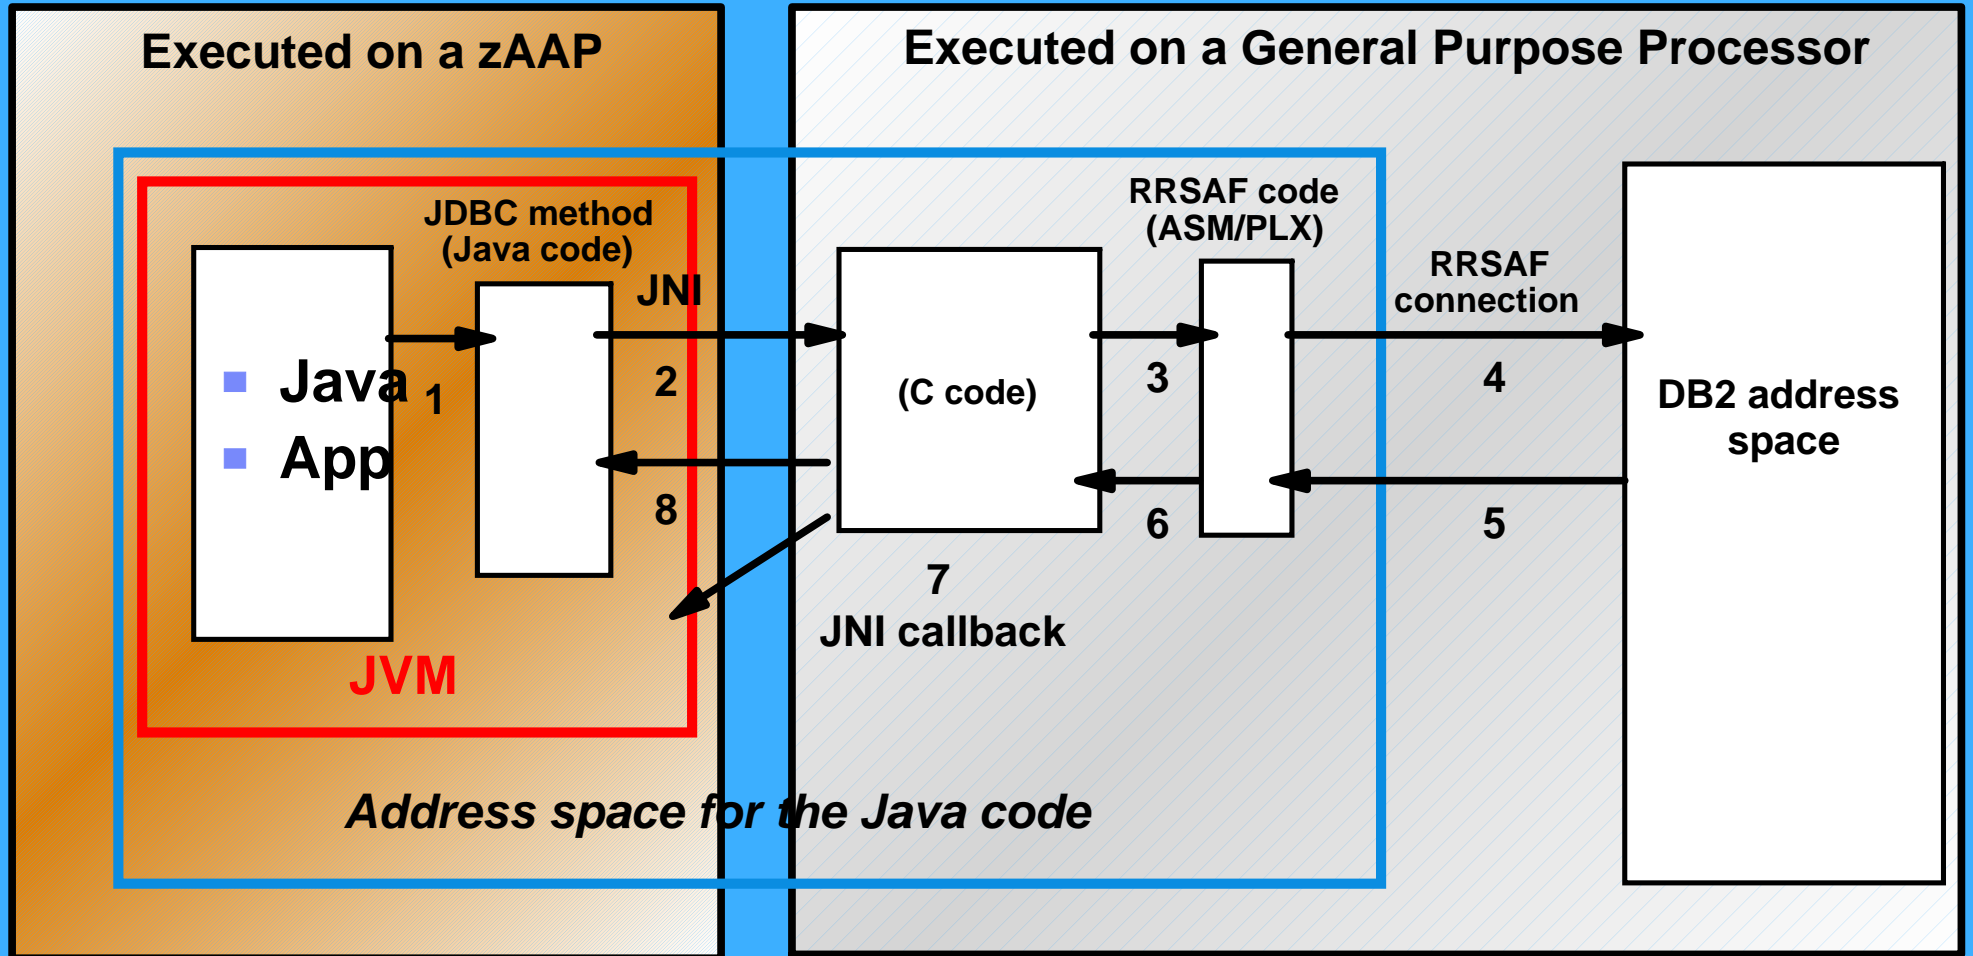
zAAP Workflow: Executing Java under IBM JVM control

- 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



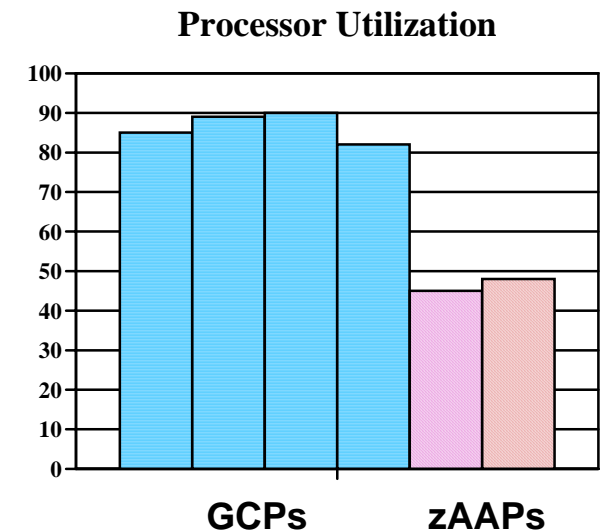
zAAP Integration at Work: Java App calling DB2

z/OS Logical Partition



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?
 - ▶ 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



zAAP 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	CPU	SERIAL	I/O	TOTAL
NUM	TYPE	PERCENTAGE	TIME PERC	TIME PERC	TIME PERC	TIME PERC	NUMBER	NUMBER	INTERRUPT	RATE	
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							

PARTITION DATA REPORT

MVS PARTITION NAME	OOSP4	NUMBER OF PHYSICAL PROCESSORS	16											
IMAGE CAPACITY	497	CP	14											
NUMBER OF CONFIGURED PARTITIONS	30	ICF	2											
----- PARTITION DATA -----														
-- LOGICAL -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --														
-----MSU----- -CAPPING-- PROCESSOR- LOGICAL PROCESSORS --- PHYSICAL PROCESSORS ---														
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

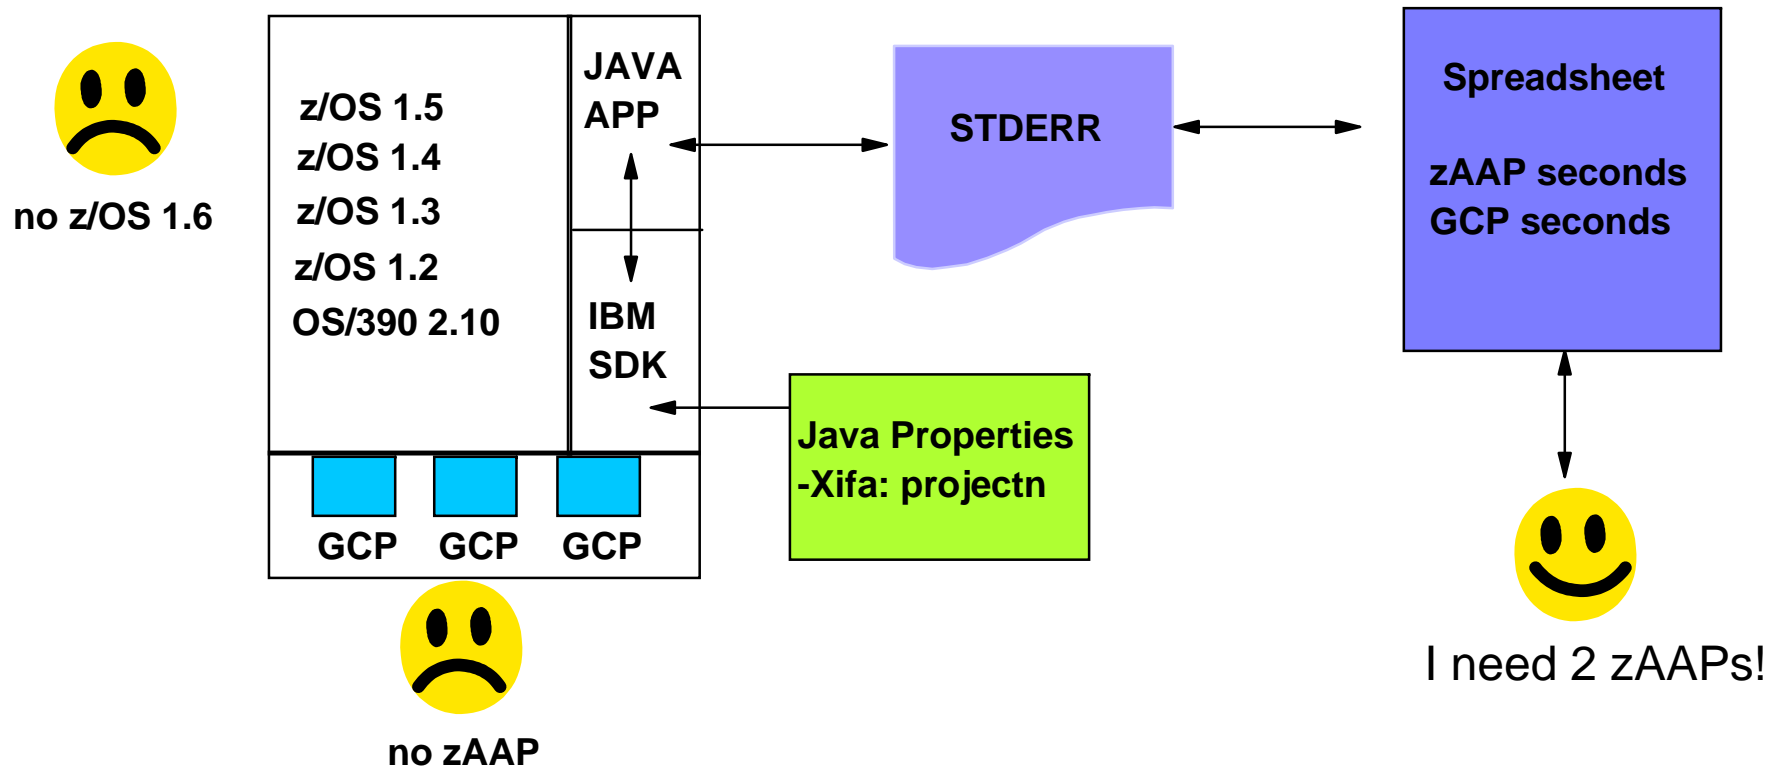
TRANSACTIONS	TRANS.-TIME	SS.TTT	--DASD I/O--	---SERVICE----	--SERVICE TIMES--
AVG 1.00	ACTUAL	0	SSCHRT 0.1	IOC 3	<u>TCB</u> 532.9
MPL 1.00	EXECUTION	0	RESP 1.3	CPU 11059K	SRB 0.0
ENDED 0	QUEUED	0	CONN 0.8	MSO 1252M	RCT 0.0
END/S 0.00	R/S AFFINITY	0	DISC 0.3	SRB 219	IIT 0.0
#SWAPS 0	INELIGIBLE	0	Q+PEND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

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

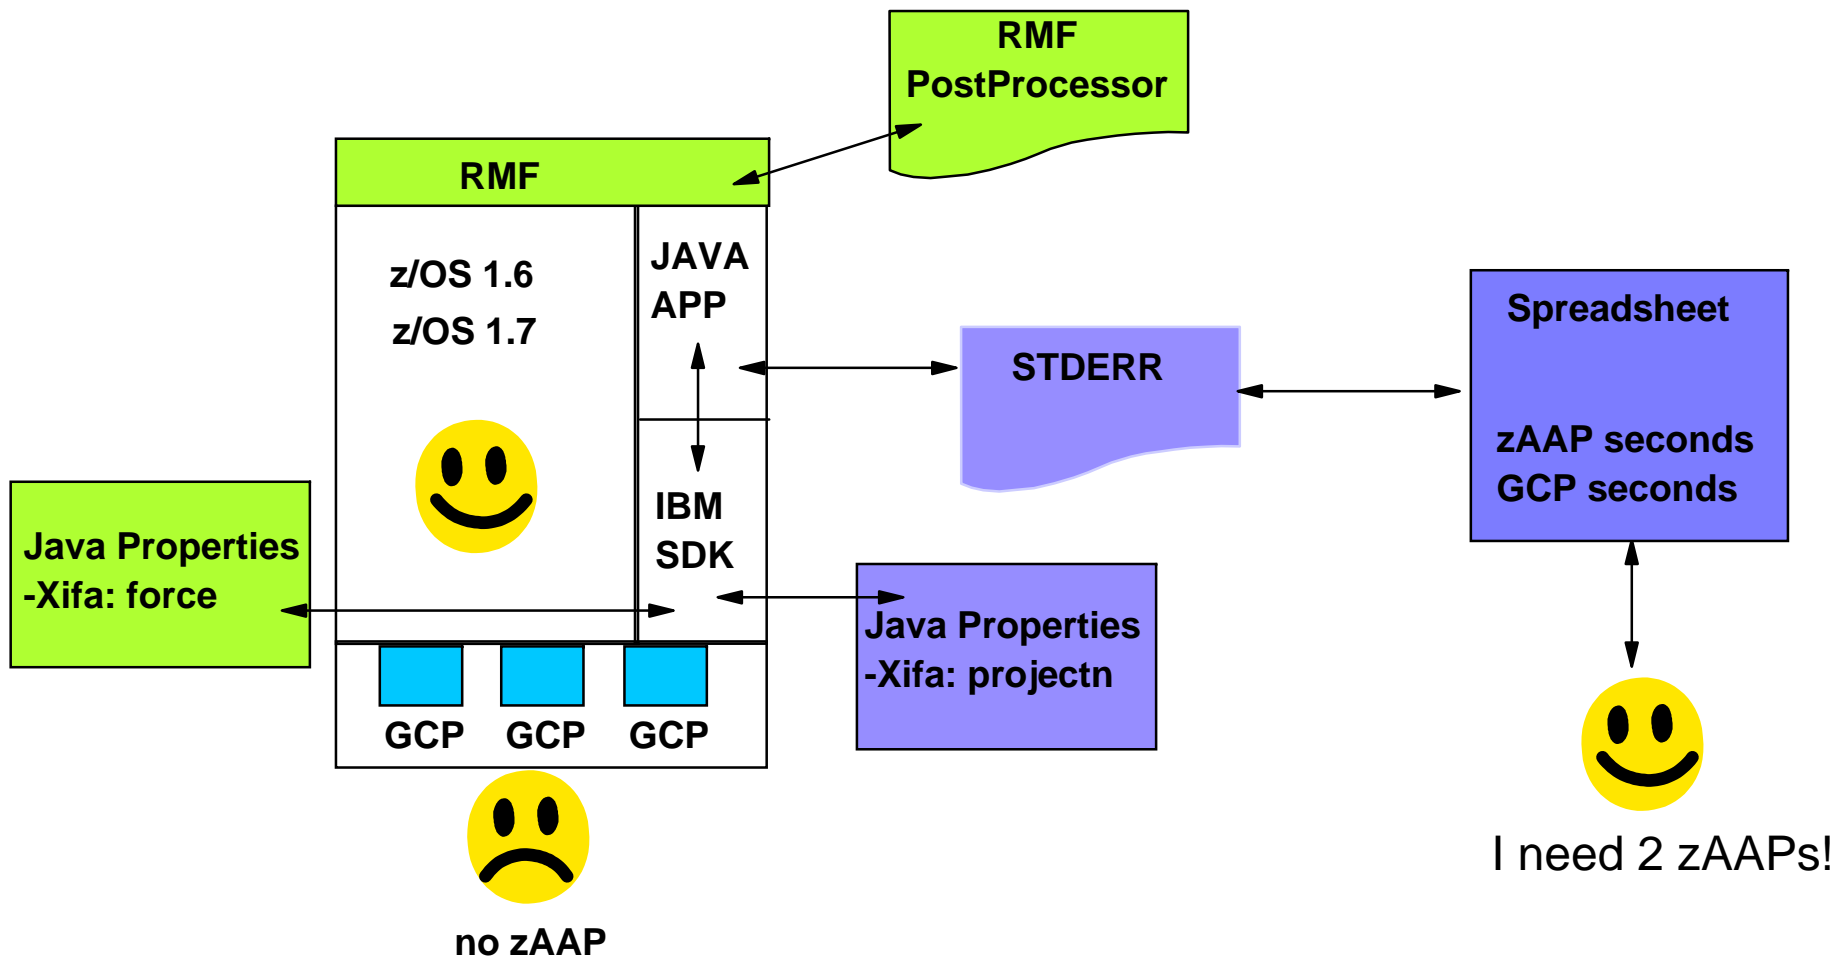
Update Capacity Planning Prior to z/OS 1.6

- 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	
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

Where's the Data?

- Ugly
 - ▶ No SMF record
 - ▶ No log file
 - ▶ Print image

```
<WRC1.50397458> Interval at: 10:07:07 Switches To/From IFA: 242857 Java IFA: 18.745860 sec. Java Standard CPU 0.861216 sec. Interval address space CPU: 23.120377 sec.  
<WRC1.50397458> Interval at: 10:09:08 Switches To/From IFA: 587782 Java IFA: 5.083801 sec. Java Standard CPU 2.023745 sec. Interval address space CPU: 7.706499 sec.  
<WRC1.50397458> Interval at: 10:11:09 Switches To/From IFA: 613917 Java IFA: 13.415511 sec. Java Standard CPU 2.094004 sec. Interval address space CPU: 16.643484 sec.
```

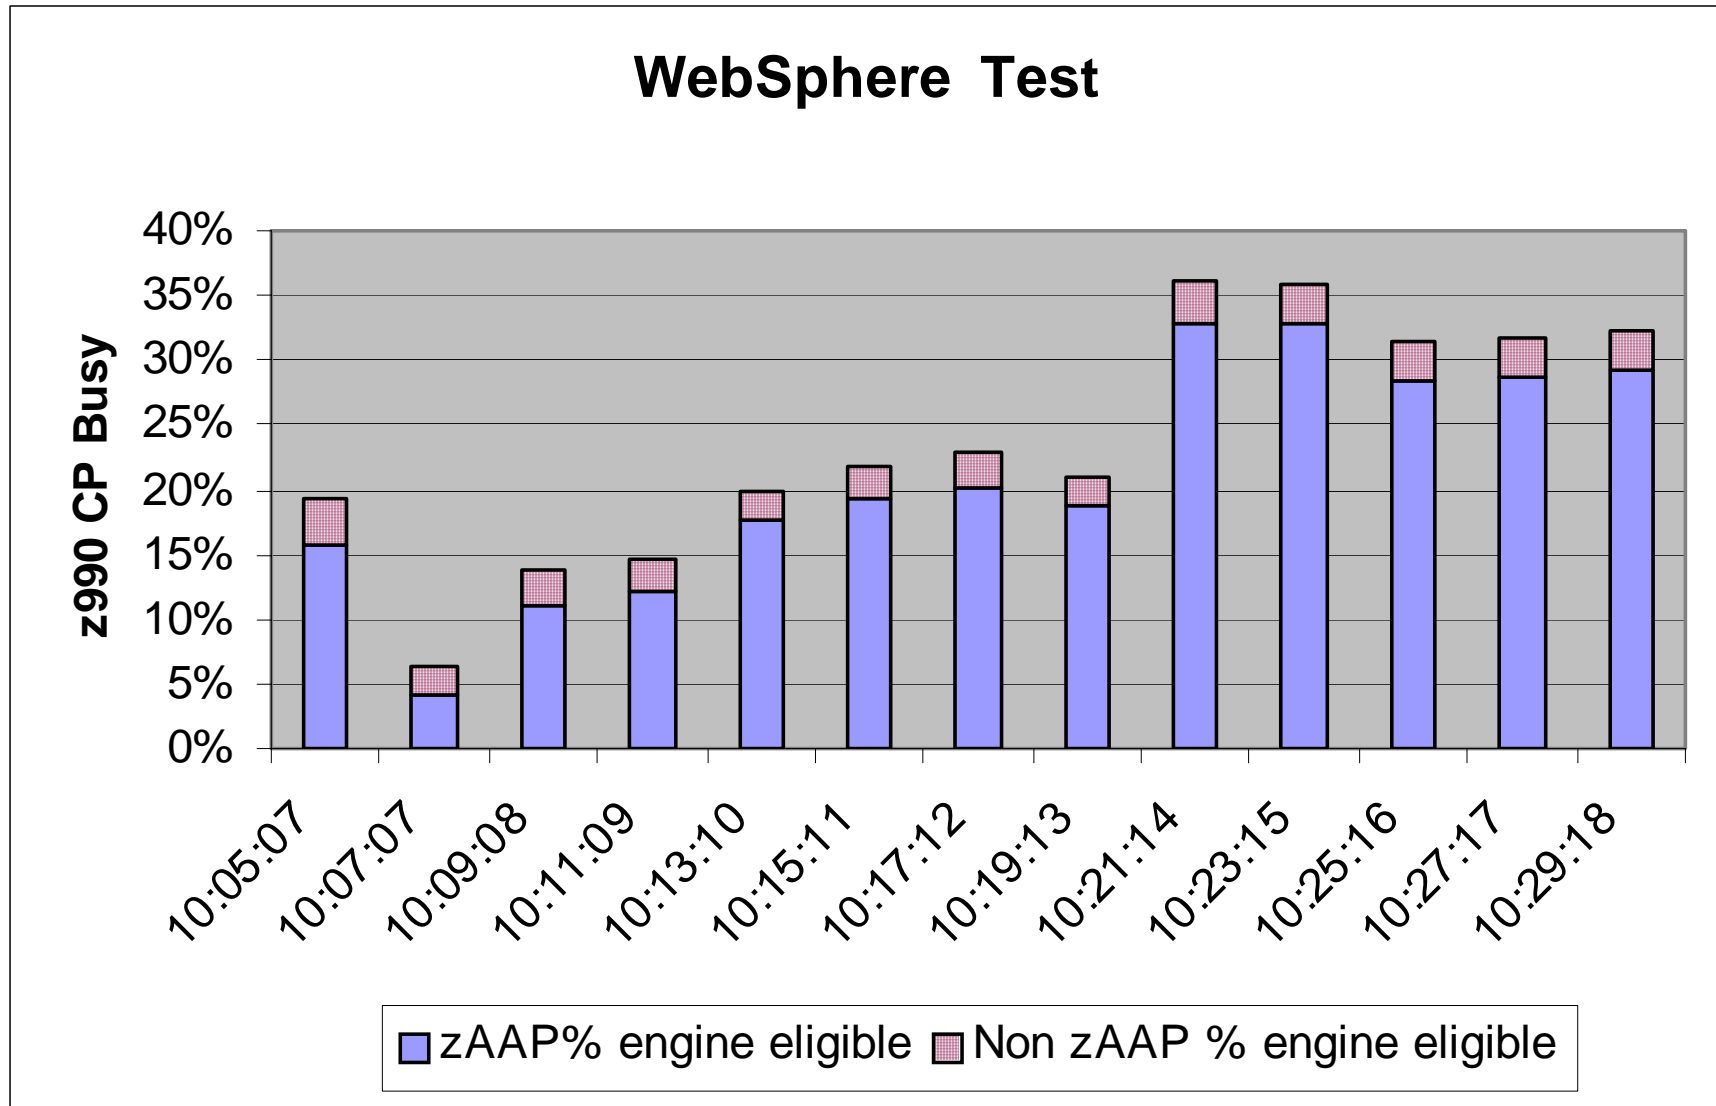
- Written at user defined interval
- Summary line written when JVM shuts down
- If the application never exits the JVM the recorded intervals can be much greater than the requested interval
 - ▶ Happens when the application is JAVA intensive

How Can I Process Ugly Data?

- IBM provides an Excel workbook to help process the output of the instrumented JVM
 - ▶ Available at
<https://www6.software.ibm.com/dl/zosjava2/zosjava2-p>

Time at start of interval	zAAP eligible seconds	Java not eligible seconds	Space CPU seconds	%Time zAAP eligible	zAAP% engine eligible	Appl% engine	zAAP% w/capt ratio	ZAAPs w/wait
			Service Class				90%	80%
10:07:07	5	2	8	66%	4%	6%	5%	6%
10:09:08	13	2	17	81%	11%	14%	12%	15%
10:11:09	15	2	18	84%	12%	15%	14%	17%
10:13:10	21	2	24	88%	18%	20%	20%	24%
10:15:11	23	2	27	89%	19%	22%	22%	27%
10:17:12	24	2	28	88%	20%	23%	22%	28%
10:19:13	23	2	25	89%	19%	21%	21%	26%
10:21:14	40	3	44	91%	33%	36%	36%	46%
10:23:15	40	2	44	91%	33%	36%	37%	46%
10:25:16	35	2	38	91%	29%	32%	32%	40%
10:27:17	35	2	38	91%	29%	32%	32%	40%
10:29:18	35	2	39	91%	29%	32%	32%	41%

Use the Capability of a Spreadsheet



What is the Objective for the zAAP

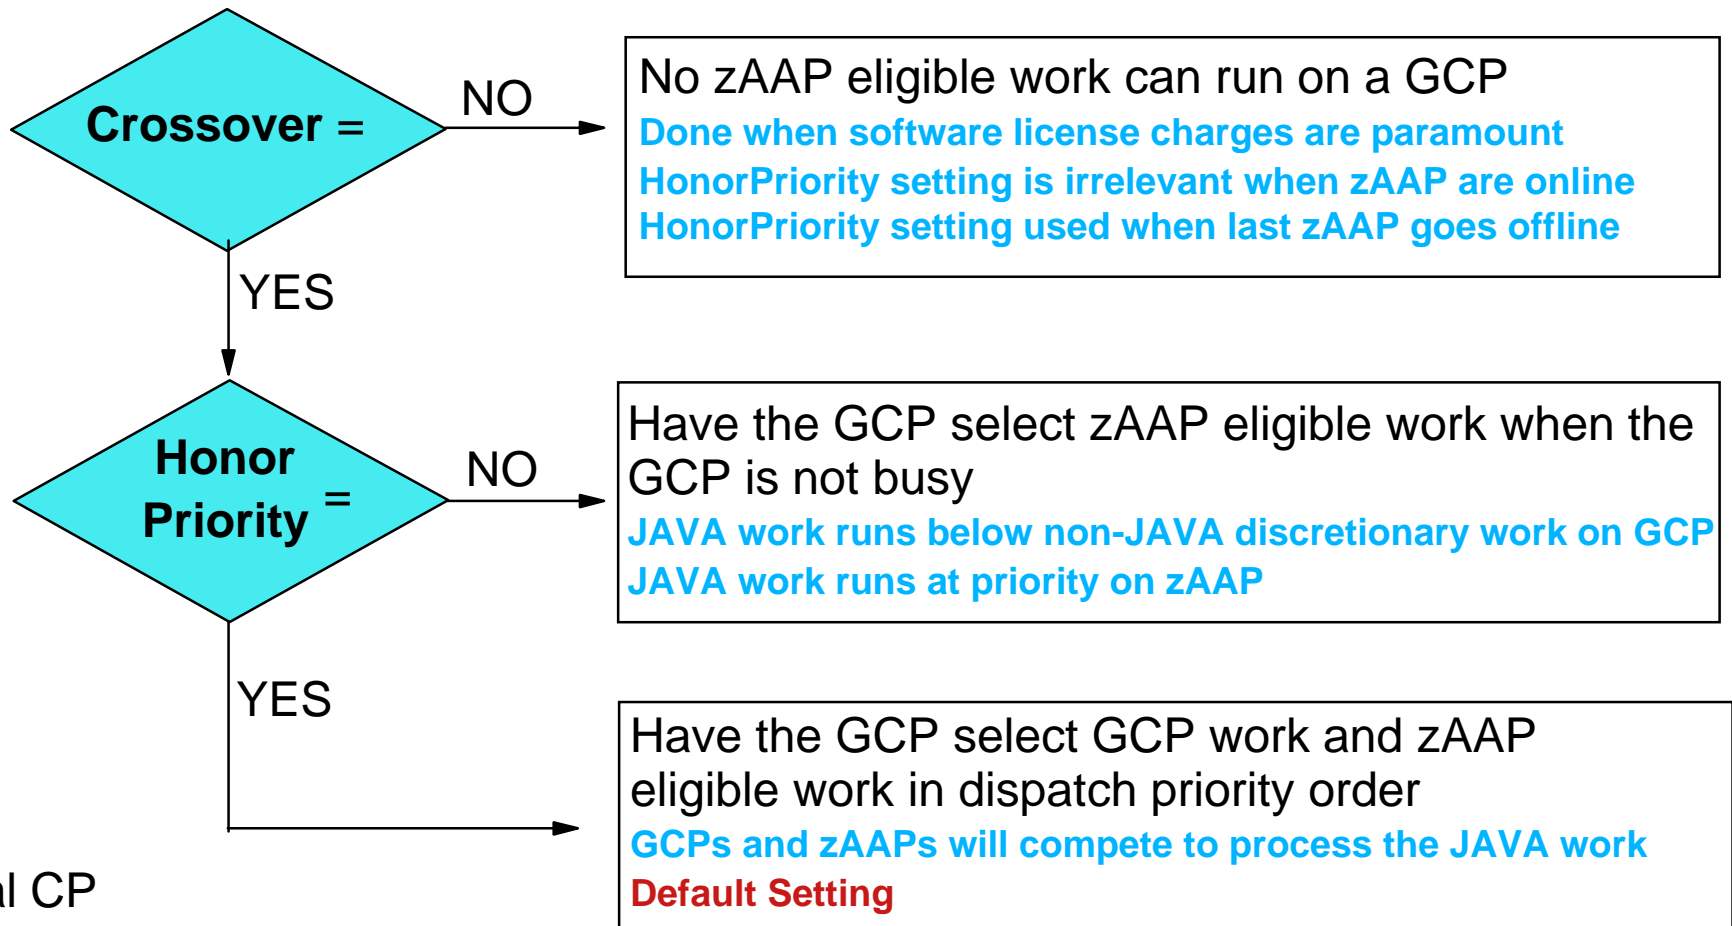
- Reduce software charge with possible improvement in performance
- Reduce the software charge and allow excess general purpose CP capacity to support Java workload
- Reduce software charge while maintaining current performance

What are the Tuning Options?

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

■ New SYS1.PARMLIB Options

- ▶ IFACROSSOVER = YES | NO
- ▶ IFAHONORPRIORITY = YES | NO



GCP - General CP
IFA - zAAP CPs

How to Set the Options

- IFACrossOver=Yes, IFAHonor Priority=Yes
 - ▶ Provides best performance for high priority work regardless of workload type
 - ▶ Probably reduces the estimated amount of work running on the zAAP

- IFACrossOver=Yes, IFAHonor Priority=No
 - ▶ Allows all processors to run ready work, but gives the zAAP better chance of running Java work
 - ▶ Helps achieve estimated offload of Java work

- IFACrossOver=No
 - ▶ Provides maximum offload of Java work onto the zAAP
 - ▶ May cause Java work to queue to a zAAP when a general purpose CP is available

Application Considerations

- Application selection is probably needed
 - ▶ Not all Java applications will be a good candidate for zAAP

- Use of zAAP controlled at the JVM level
 - ▶ Controls include:
 - On
 - Off
 - Force
 - ▶ Are all applications running under the same JVM good candidates?

Application Selection

- Maximize the amount of time on the zAAP
 - ▶ Java intensive application - good candidate
 - ▶ Java trivial application - bad candidate
 - Estimation tool provides guidance

Time at start of interval	zAAP eligible seconds	Java not eligible seconds	Space CPU seconds	%Time zAAP eligible	Duration
			Service Class		
10:05:22	137	102	240	57%	120
10:07:22	167	197	365	46%	121
10:09:23	164	199	364	45%	121
10:11:24	163	197	361	45%	121
10:13:25	162	194	357	45%	121

Application Selection

- Minimize the cost of getting on the zAAP
 - ▶ Software switching rate
 - Low switch rate minimizes the software cost
 - High switch rate increases the software cost
 - Estimation tool provides guidance

Time at start of interval	Switch rate	zAAP eligible seconds	Java not eligible seconds	Space CPU seconds
7:04:20	7,806	18	11	51
7:09:21	2,835	9	5	28
7:14:22	3,022	8	5	27
7:19:23	2,906	15	5	32
7:24:25	2,512	7	4	24
7:29:26	2,687	7	5	23

Application Selection

- Java eligible time per switch is probably a good indicator of application applicability for a zAAP
 - ▶ Value can be calculated using the estimation tool

Time at start of interval	Switch rate	zAAP eligible seconds	Duration	zAAP eligible microseconds per switch
10:05:22	19,084	137	120	60
10:07:22	31,681	167	121	44
10:09:23	32,913	164	121	41
10:11:24	29,986	163	121	45
10:13:25	28,798	162	121	46
10:15:26	29,763	161	121	45

Capacity Planning Factors in a zAAP Environment

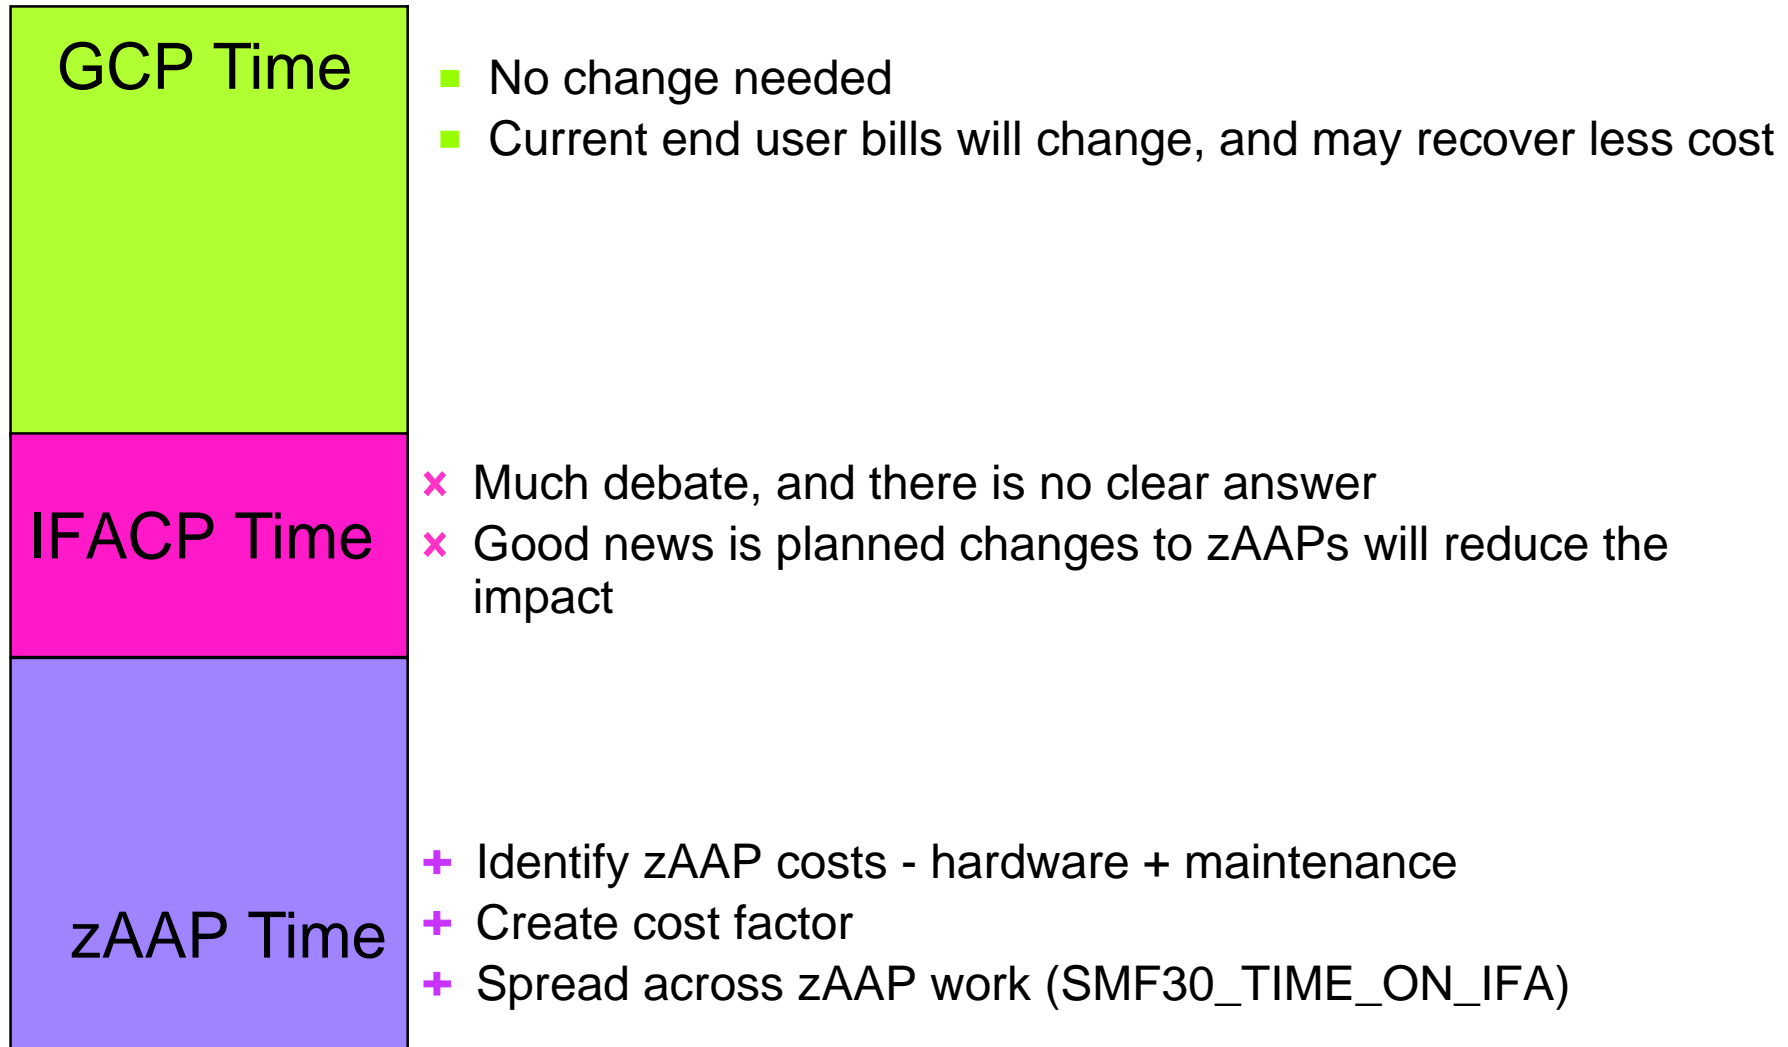
- Software
 - ▶ MP effects on z/OS
 - ▶ Switch request rate
 - ▶ Utilization of the zAAP
- Hardware
 - ▶ Number of books
 - ▶ Cost of additional partition
- Use zPCR to estimate the costs
 - ▶ Lightly used environments will not incur the full cost of the above planning factors
 - ▶ zPCR gives a conservative estimate
 - Being refined all the time

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%

Crossover=Yes, HonorPriority=NO	Crossover=Yes, HonorPriority=YES
<ul style="list-style-type: none"> ◆ zAAP work runs after Discretionary - represents GCP MIPS which would have been lost anyway ◆ Charge IFACP time at zAAP rate <ul style="list-style-type: none"> ■ This is money you wouldn't get in this interval otherwise because the zAAP was busy ■ Plus zAAP in next interval is less busy and can be charged there because we used "free" cycles. It's less busy because if we didn't use GCP to run the work we would have to run it on the zAAP eventually 	<ul style="list-style-type: none"> ◆ zAAP work runs at Priority on GCP when zAAP needs help - represents using premium GCP MIPS <ul style="list-style-type: none"> ■ GCP is a <u>safety net</u> to protect response times ◆ How to Charge - At GCP rates or zAAP rates? cost of safety net = (cost of GCP per second - cost of zAAP per second) * IFACP seconds ◆ If cost of the safety net >= cost of amortized zAAP then buy a zAAP ◆ If not then need to identify how to spread the cost of the safety net to recapture the cost

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
- Amount of actual offloaded cycles can be affected by zAAP options selected
 - ▶ And soon new functions via APAR
- Estimation tool can help identify good Java candidates
- Capacity planning methods need to be discussed
- More information to follow as we learn more about zAAPs

IFA OPT Options to change with PTF planned for 3Q2005

- Java eligible work tends to be high priority (e.g. Websphere, CICS)
- GCPs tend to outnumber zAAPs (3:1 or 4:1)
 - ▶ With Crossover=YES, GCP tend to take most of the zAAP work hurting the effectiveness of the solution
 - ▶ With Crossover=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 GCP when "they need help"

IFA IEAOPTxx Planned Changes

XO	HP	IFA NH**	Description
Y	Y	Y	Standard processors run both zAAP eligible and non zAAP eligible work in priority order
Y	Y	N	Standard processors run only non-zAAP eligible work in priority order and run zAAP eligible work as if it has a priority less than non zaap eligible discretionary work
Y	N	Y	Standard processors run non zAAP eligible work in priority order and run zAAP eligible work as if it has a priority less than non zaap eligible discretionary work
Y	N	N	Standard processors run non zAAP eligible work in priority order and run zAAP eligible work as if it has a priority less than non zaap eligible discretionary work
N	Y	Y	Standard processors run both zAAP eligible and non zAAP eligible work in priority order
N	N	Y	Standard processors do not run zAAP eligible workloads
N	Y	N	Standard processors do not run zAAP eligible workloads
N	N	N	Standard processors do not run zAAP eligible workloads

XO: Crossover

HP: HonorPriority

* zAAP Processors run all work in priority order

** IFA NH: IFA Needs Help (determined by Alternate Wait State Mgmt)

Demonstration of planned PTF'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	PTF XXXXXX	369	10	229	2%
6	2	2360	19%	10	79	97	83	PTF XXXXXX	436	69	173	11%
7	1	2292	16%	11	78	98	80	PTF XXXXXX	519	151	90	25%