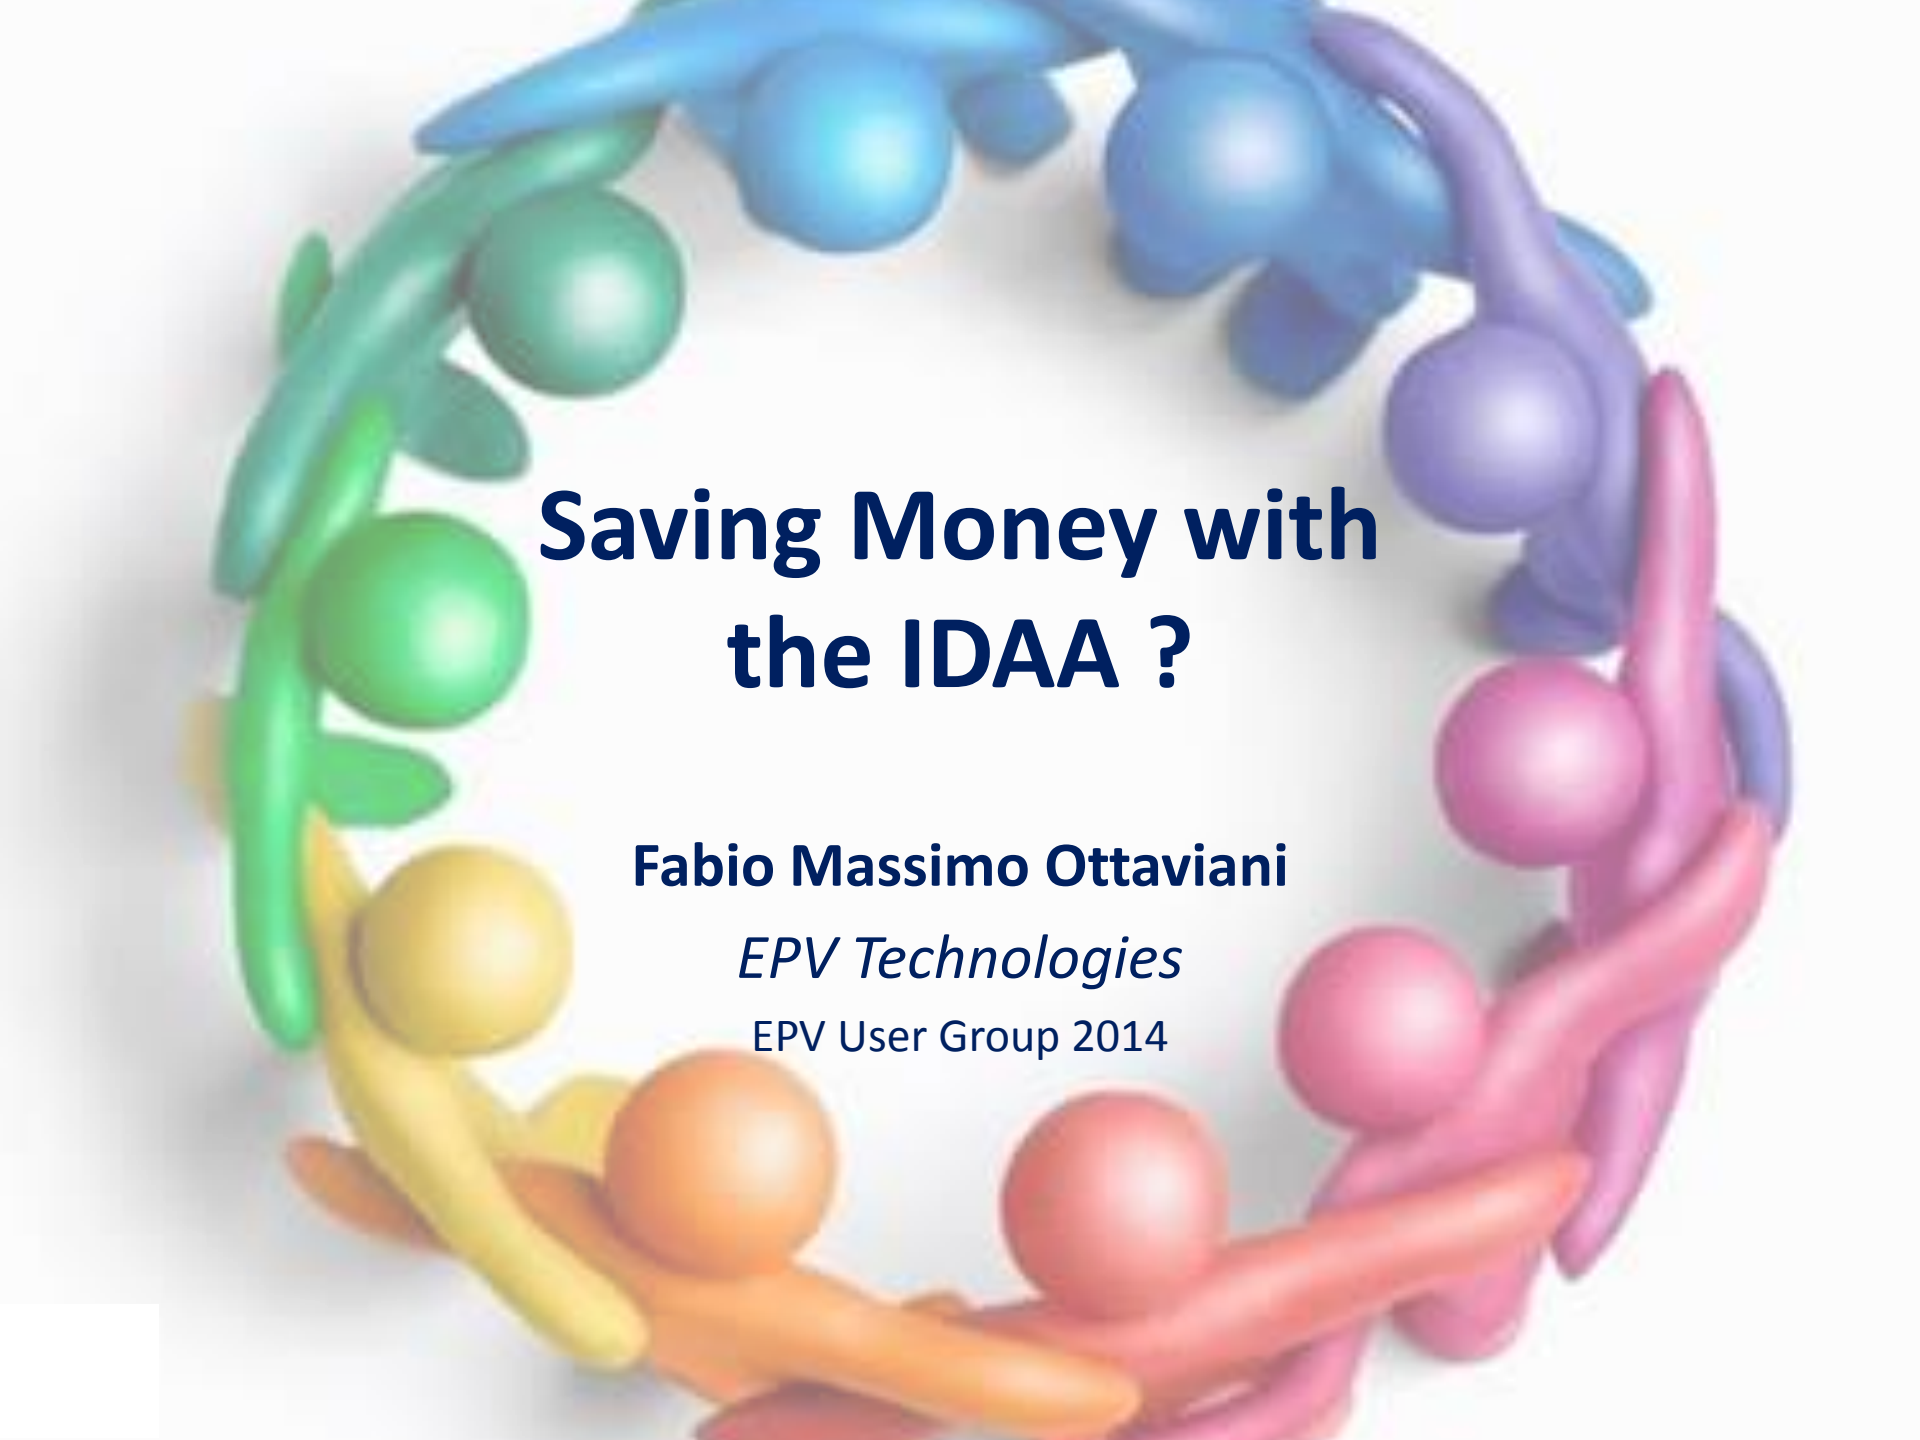

# Saving Money with the IDAA ?

**Fabio Massimo Ottaviani**

*EPV Technologies*

EPV User Group 2014

# Disclaimer, copyright and trademarks

- Disclaimer:

THE INFORMATION CONTAINED IN THIS PRESENTATION HAS NOT BEEN SUBMITTED TO ANY FORMAL REVIEW AND IS DISTRIBUTED ON AN "AS IS" BASIS WITHOUT ANY WARRANTY EITHER EXPRESS OR IMPLIED. THE USE OF THIS INFORMATION OR THE IMPLEMENTATION OF ANY OF THESE TECHNIQUES IS A USER RESPONSIBILITY AND DEPENDS ON THE USER'S ABILITY TO EVALUATE AND INTEGRATE THEM INTO THE USER'S OPERATIONAL ENVIRONMENT. WHILE EACH ITEM MAY HAVE BEEN REVIEWED FOR ACCURACY IN A SPECIFIC SITUATION, THERE IS NO GUARANTEE THAT THE SAME OR SIMILAR RESULTS WILL BE OBTAINED ELSEWHERE. USERS ATTEMPTING TO ADAPT THESE TECHNIQUES TO THEIR OWN ENVIRONMENTS DO SO AT THEIR OWN RISK.

- Copyright Notice:



- Trademarks:

All the trademarks mentioned here belong to their respective companies.

# EPV Technologies

- Since 2003 EPV Technologies has developed and distributes the EPV products suite for Capacity Management
- Products' goal is to control and optimize hardware and software to reduce costs and maintain and often improve service levels
- EPV zParser, EPV SMF2XL, EPV for z/OS, EPV Graph for z/OS, EPV for DB2, EPV for WMQ, EPV for zLINUX, EPV for UNIX, EPV Graph for UNIX, EPV for VMware, EPV for WINDOWS, MyEPV
- Customers in Italy, Belgium, Brazil, France, Germany, Israel, Japan, Luxembourg, Switzerland, UK and USA
- WWW.EPVTECH.COM

# Agenda

- Introduction
- IDAA Accounting
- IDAA Statistics
- Accelerator Modeling
- What about Money ?
- Conclusions

# Introduction

- The IBM DB2 Analytics Accelerator (IDAA) is an appliance that comes as additional hardware and software to be connected to a z196, z114 or zEC12 system
- Initially based on zBX, since V2 it's powered by Netezza technology
- Offloading queries to the IDAA avoids CPU usage on z/OS and guarantees great performance improvements
- Current version is V4; it was released in November 2013 (ENUS213-427)

# Introduction

- Most important new features in V4
  - Static SQL and multi-row FETCH support
  - Automated workload balancing across multiple accelerators
  - Improved incremental update performance
  - Many High Performance Storage Saver enhancements
  - DB2 V11 support

# Introduction

- Software requirements for the IDAA:
  - z/OS V1.11, or later
  - DB2 V10.1 for z/OS , or later
  - DB2 Utilities Suite for z/OS , V10.1, or later
  - XML Toolkit for z/OS , V1.10, or later
  - IBM Ported Tools for z/OS , V1.2.0, or later

# Introduction

- IDAA is fully integrated in DB2 and it is used as an access path under the control of the Optimizer
- DB2 is still the owner of the data
- Users and applications have to access the IDAA through DB2; they don’t need to be aware of it
- No application change is required
- IDAA utilization has to be allowed by DSNZPARM parameters and the CURRENT QUERY ACCELERATION special register

# Introduction

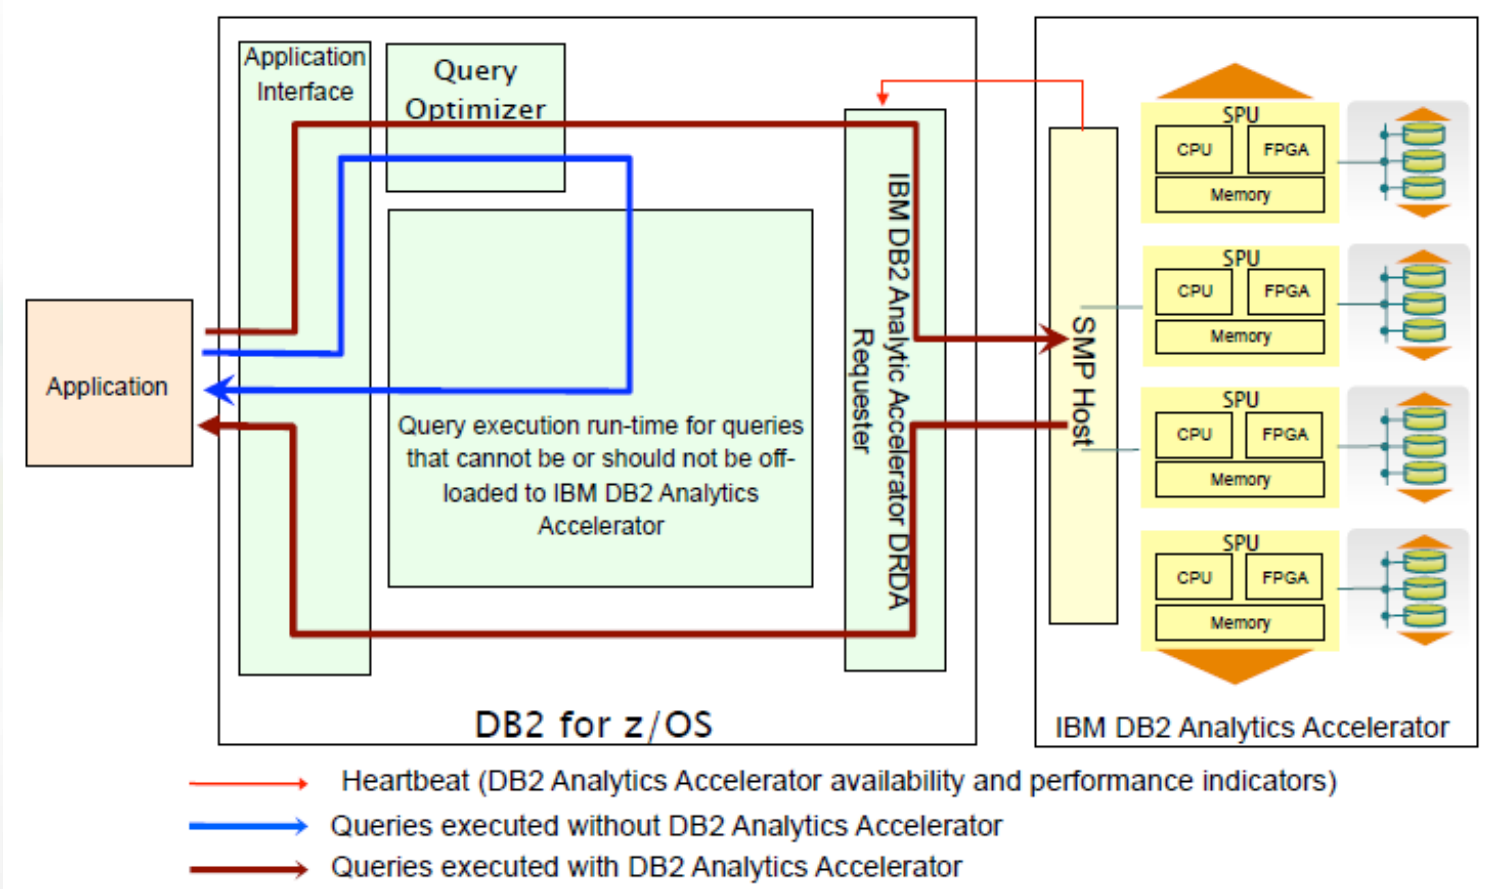


from 'Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS'
IBM Redbooks SG24-8005-00

# Introduction

- IDAA hardware resources (N2001 full rack):
  - 2 SMP hosts (also called coordinator nodes)
  - 7 S-Blades (also called worker nodes)
  - 7 * 16 CPU cores (processing units)
  - 7 * 16 FPGA (Field Programmable Gate Array)
  - 7 * 128 GB Memory
  - 240 disks for user data (RAID1 protected) + 34 spare disks + 14 disks for swap/log space required by the Linux operating system running on the S-Blades (2 disks / S-Blade)
  - 48 TB uncompressed data (4:1 compression expected)

# Introduction

- DB2 only talks with the SMP hosts
- SMP hosts send the query to the S-Blades
- Each S-Blade processes those slices of data that are stored on the disks assigned (40 or 32)
- Each FPGA, inside the S-Blade, decompresses a slice of data, removes all columns that are not needed for further processing and filter the data to reduce the amount of data passed back to the CPU to a minimum
- Each CPU, inside the S-Blade, completes its part of the query processing

# Introduction

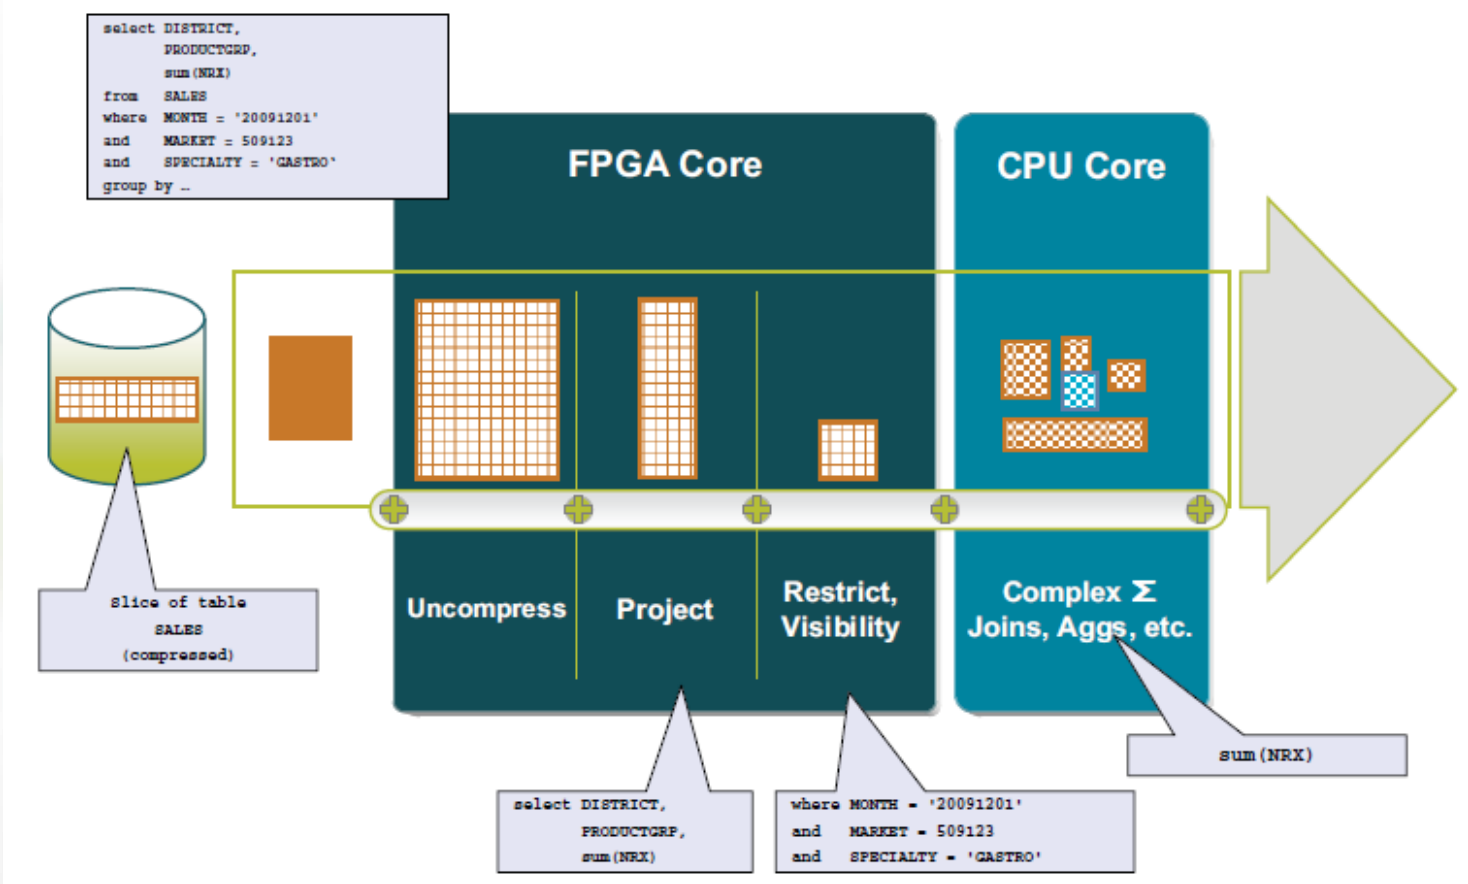


from 'Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS'
IBM Redbooks SG24-8005-00

# Introduction

- New set of metrics in Q8AC section of IFCID 3 (DB2 accounting)
- New set of metrics in Q8ST section of IFCID 2 (DB2 statistics)
- Let's start with accounting …

# IDAA accounting

# Accounting Elapsed Times

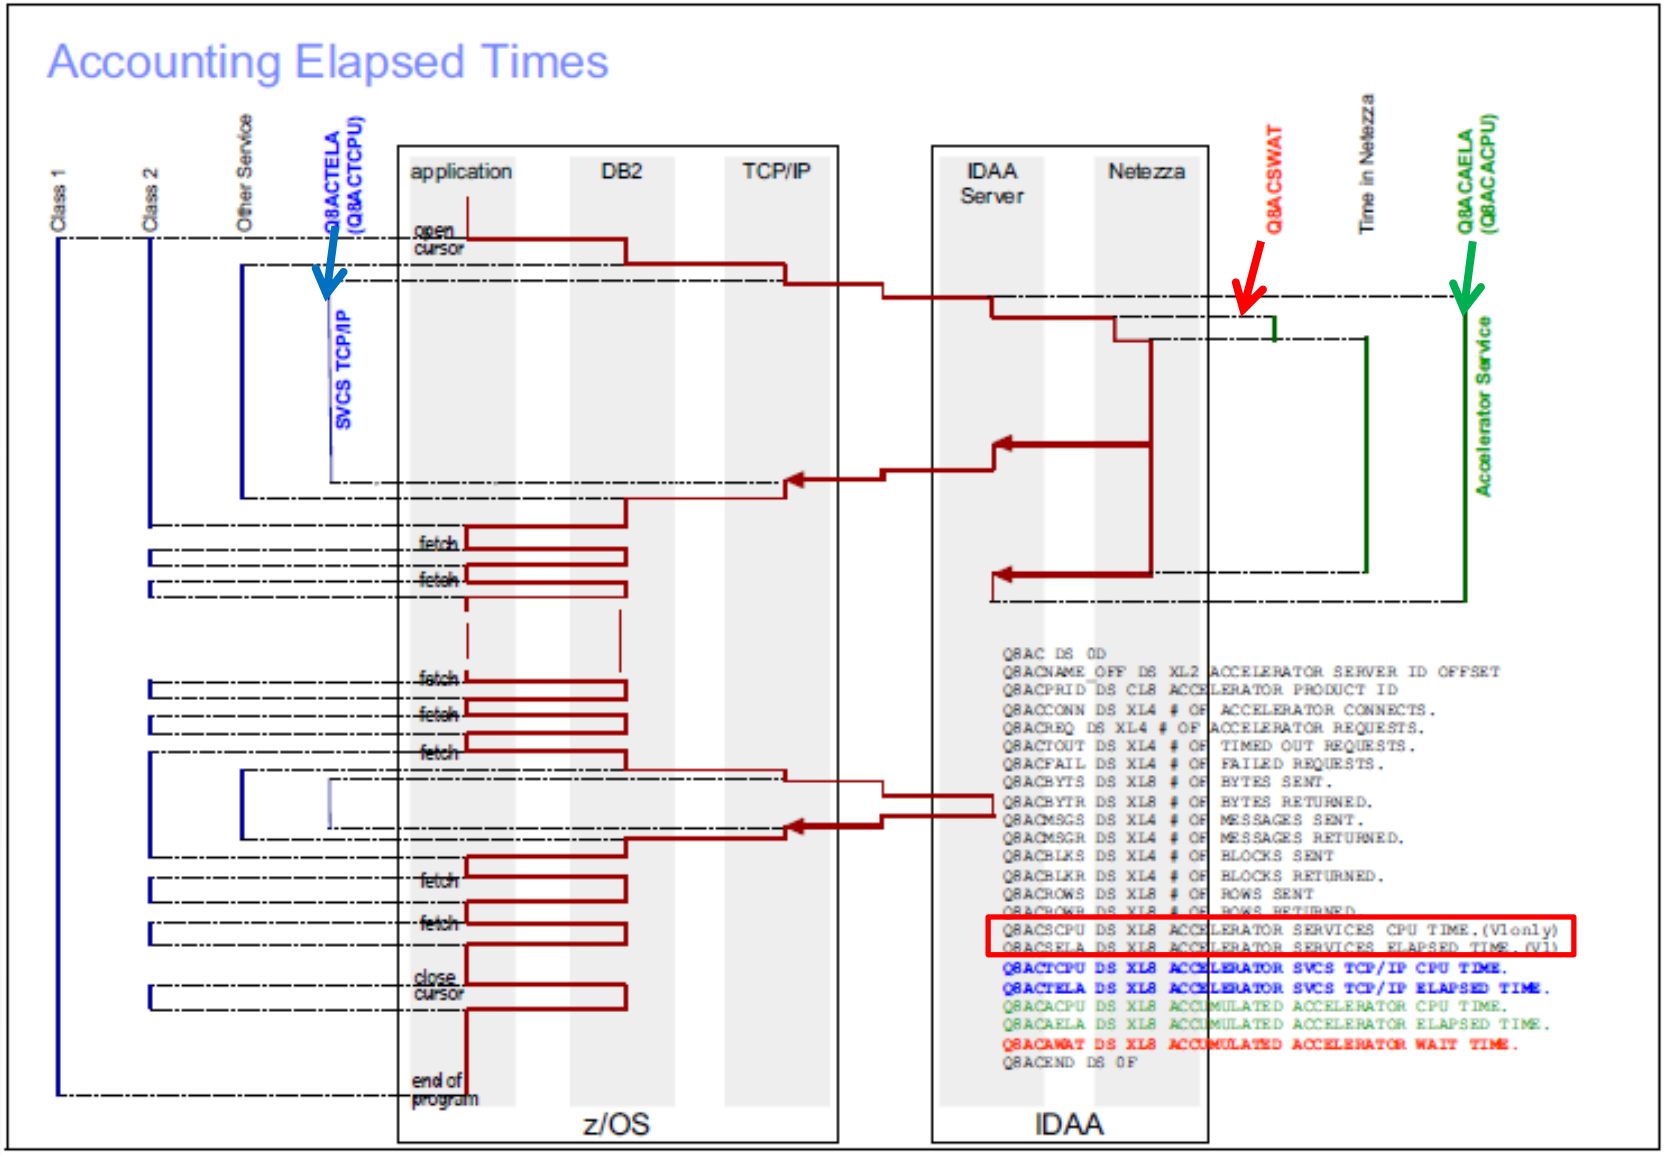


from 'Hybrid Analytics Solution using IBM DB2 Analytics Accelerator for z/OS V3.1'
IBM Redbooks SG24-8151-00

# IDAA Accounting Metrics

- Q8ACNAME, accelerator name
- Q8ACPRID, accelerator product id
- Q8ACCONN, connects to accelerator
- Q8ACREQ, requests sent to accelerator (Q8ACFAIL for failed requests and Q8ACTOUT for timed out requests are also provided)

# IDAA Accounting Metrics

- Q8ACTELA, accelerator services TCP/IP elapsed time measured in DB2; it starts when sending the requests to an accelerator and ends when receiving the results from an accelerator
- Q8ACTCPU, The accelerator services TCP/IP CPU time measured in DB2 for the amount of CPU consumed by the DDF service task to perform the SEND and RECEIVE to an accelerator

# IDAA Accounting Metrics

- Q8ACAELA, elapsed time spent in the accelerator when executing requests from the DB2 subsystem
- Q8ACACPU, The CPU time spent in the accelerator when executing requests from the DB2 subsystem
- Q8ACAWAT, The wait time spent in the accelerator when executing requests from the DB2 subsystem

# IDAA Accounting Metrics

- Q8ACBLKR, number of blocks returned from accelerator
- Q8ACBLKS, number of blocks sent to accelerator
- Q8ACBYTR, number of bytes returned from accelerator
- Q8ACBYTS, number of bytes sent to accelerator
- Q8ACMSGR, number of messages returned from accelerator
- Q8ACMSGS, number of messages sent to accelerator
- Q8ACROWR, number of rows returned from accelerator
- Q8ACROWS, number of rows sent to accelerator

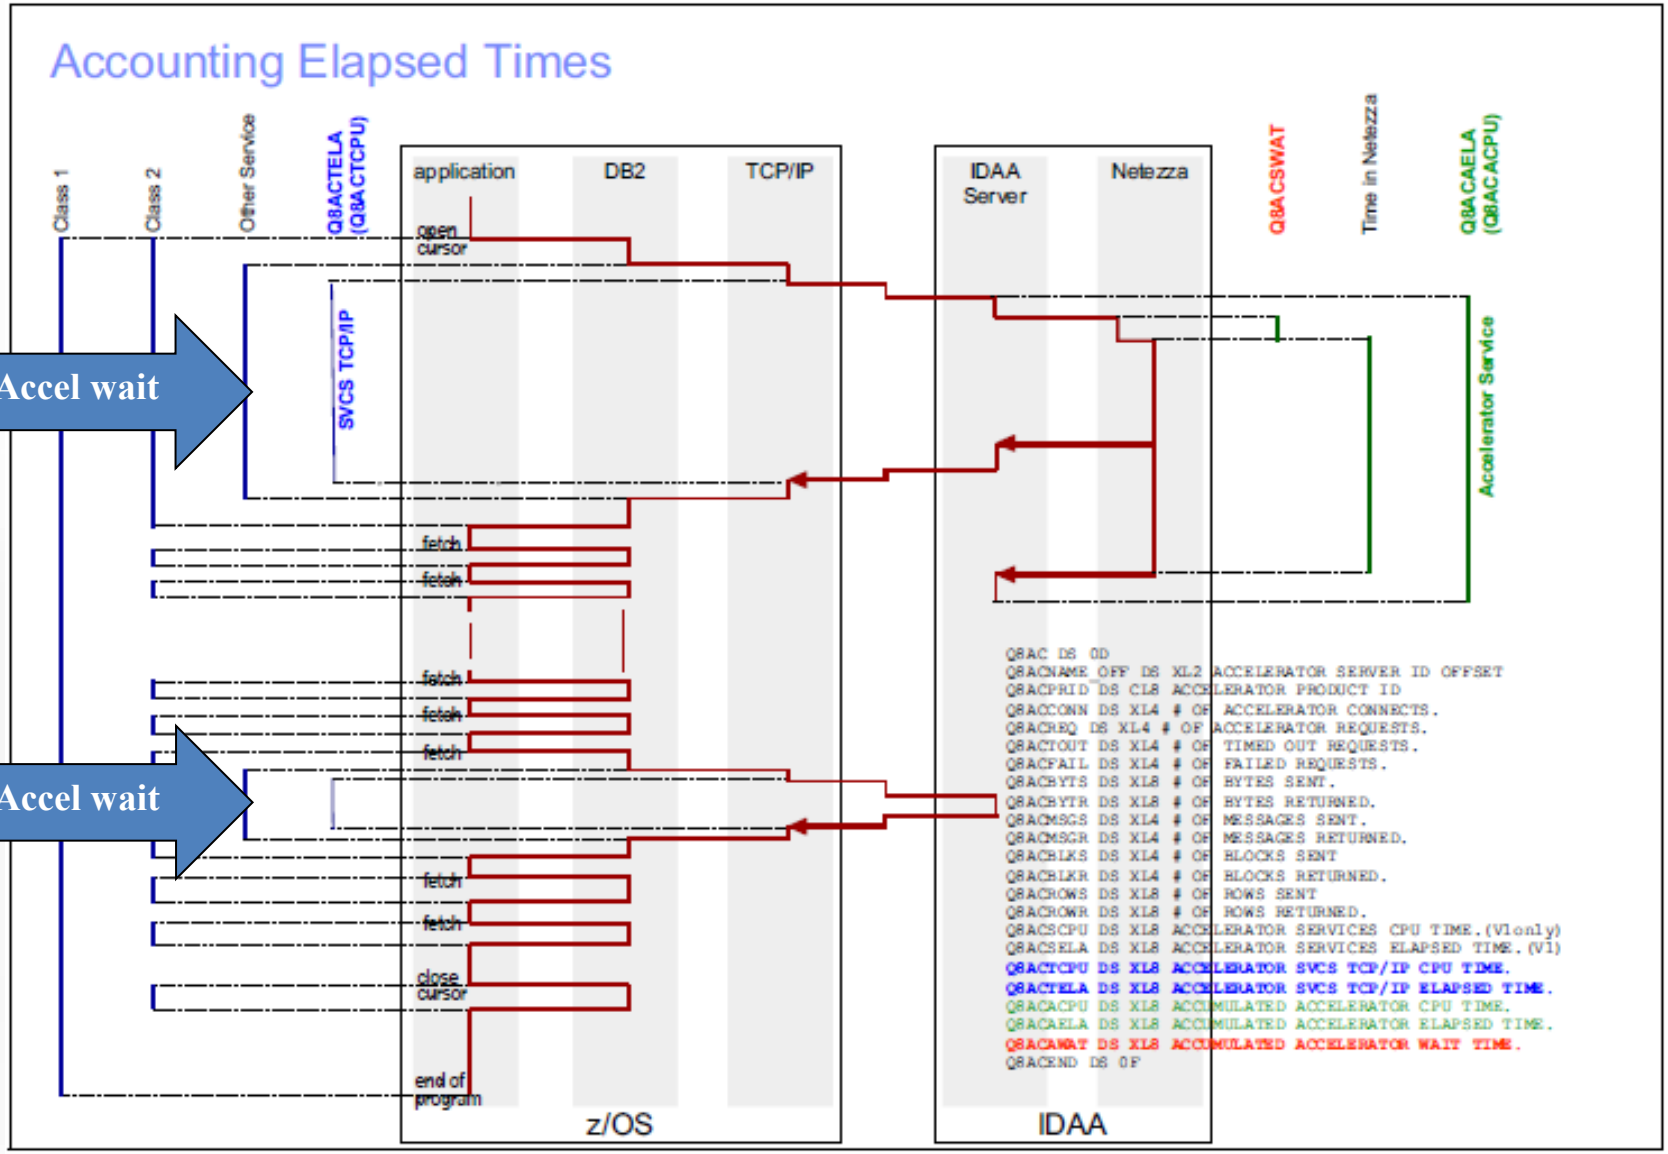


from 'Hybrid Analytics Solution using IBM DB2 Analytics Accelerator for z/OS V3.1'
IBM Redbooks SG24-8151-00

# IDAA Accounting Metrics

- Time spent in accelerator has to be considered as suspension time (Class 3 metrics)
- If you don’t consider the new metrics about accelerator suspension time you may get a, misleading, high NOT ACCOUNTED time for queries exploiting the accelerator

# IDAA Accounting Metrics

- High % NOT ACC for business critical applications should be investigated

*% NOT ACC = (DB2 elapsed – DB2 CPU – DB2 zIIP – DB2 wait) / DB2 elapsed*

| PLAN NAME | CONNTYPE | CONN | REQS | DB2 CPU | DB2 IIP | DB2 ELAPSED | DB2 WAIT | % NOT ACC |
|---|---|---|---|---|---|---|---|---|
| MACAA452 | BATCH | BATCH | 3 | 2.064,58 | 0 | 3.060,93 | 56,64 | 30,7 |
| XXPFOCIC | CICS ATTACH | CICSPA13 | 452 | 1.088,64 | 0 | 4.374,80 | 2.794,59 | 11,2 |
| AEVISCIC | CICS ATTACH | CICSPA11 | 18.558 | 789,54 | 0 | 3.598,34 | 2.605,02 | 5,7 |
| KO2PLAN | RRSAF ATTACH | RRSAF | 304 | 546,39 | 0 | 1.902,92 | 1.281,32 | 4 |
| MJRCSCIC | CICS ATTACH | CICSPA13 | 398 | 414,43 | 0 | 5.530,86 | 5.017,50 | 1,8 |
| TRPECCIC | CICS ATTACH | CICSPA11 | 732 | 152,84 | 0 | 314,01 | 102,54 | 18,7 |
| PGRATCIC | CICS ATTACH | CICSPA11 | 7.407 | 142,57 | 0 | 726,03 | 546,69 | 5,1 |

# IDAA Accounting Metrics

- CLASS 3: ACCELERATOR SUSPENSION TIME

The accumulated wait time for requests to an accelerator

Field Name: QWACAACW

- CLASS 3: ACCELERATOR EVENTS

The number of wait trace events processed for requests to an accelerator

Field Name: QWACAACC

# IDAA Accounting Metrics

- CLASS 8: PACKAGE ACCELERATOR SUSPENSION TIME

The accumulated wait time for requests to an accelerator

Field Name: QPACAACW

- CLASS 8: PACKAGE ACCELERATOR EVENTS

The number of wait trace events processed for requests to an accelerator

Field Name: QPACAACC

# IDAA statistics

# IDAA Statistics

- Quick evolution of Q8ST metrics
- Issues to be aware of:
  - some Q8ST metrics are only valid with IDAA V4.1
  - some Q8ST metrics refer to Accelerator, some to DB2 SS
  - some Q8ST metrics are accumulated
  - some Q8ST metrics are obsolete

# IDAA Statistics – Accelerator Scope

| Variable Name | EPV Label | Scope | Accum |
|---|---|---|---|
| Q8STPRID | ACCEL_PRODUCT_ID | AC | N |
| Q8STNAME | ACCEL_SERVER_ID | AC | N |
| Q8STTATE | ACCEL_CURRENT_STATE | AC | N |
| Q8STTART | ACCEL_SERVER_START_AC | AC | N |
| Q8STTATC | ACCEL_STATUS_CHANGE_AC | AC | N |
| Q8STNQSA | ACCEL_SUCCESSFUL_QUERIES_AC | AC | Y |
| Q8STNQFA | ACCEL_FAILED_QUERIES_AC | AC | Y |
| Q8STACTV_64 | ACCEL_CURR_EXECUTING_QUERIES_AC | AC | N |
| Q8STMAXA_64 | ACCEL_MAX_EXECUTING_QUERIES_AC | AC | N |
| Q8STTCQA | ACCEL_CPU_EXEC_QUERIES_AC | AC | Y |
| Q8STTCMA | ACCEL_CPU_MAINT_OPER_AC | AC | Y |
| Q8STCQL | ACCEL_CURR_QUEUE_LENGTH_AC | AC | N |
| Q8STMAXQ_64 | ACCEL_HWM_QUEUE_LENGTH_AC | AC | N |
| Q8STQUEW | ACCEL_AVG_QUEUE_WAIT_AC | AC | N |
| Q8STQUEM | ACCEL_MAX_QUEUE_WAIT_AC | AC | N |
| Q8STCORS | ACCEL_PROCESSORS_WORKER_AC | AC | N |
| Q8STCCPU_64 | ACCEL_AVG_CPUBUSY_COORD_AC | AC | N |
| Q8STWCPU_64 | ACCEL_AVG_CPUBUSY_WORKER_AC | AC | N |
| Q8STWNOD_64 | ACCEL_ACTIVE_WORKER_NODES_AC | AC | N |
| Q8STDSKA | ACCEL_DISK_STORAGE_AVAILABLE_AC | AC | N |
| Q8STDSKU | ACCEL_DISK_INUSE_WORKER_AC | AC | N |
| Q8STDSA | ACCEL_DISK_STORAGE_INUSE_AC | AC | N |
| Q8STNMDS | ACCEL_DATA_SLICES_AC | AC | N |

Throughput

Queuing

Capacity

# IDAA Statistics – Subsystem Scope

| Variable Name | EPV Label | Scope | Accum |
|---|---|---|---|
| Q8STSREQ | ACCEL_SUCCESSFUL_QUERIES | SS | Y |
| Q8STFREQ | ACCEL_FAILED_QUERIES | SS | Y |
| Q8STNQCS | ACCEL_CURR_EXECUTING_QUERIES | SS | N |
| Q8STMNQS | ACCEL_MAX_EXECUTING_QUERIES | SS | N |
| Q8STTCQS | ACCEL_CPU_EXEC_QUERIES | SS | Y |
| Q8STTCMS | ACCEL_CPU_MAINT_OPER | SS | Y |
| Q8STCONN | ACCEL_CONNECTS | SS | Y |
| Q8STREQ | ACCEL_REQUESTS | SS | Y |
| Q8STFAIL | ACCEL_FAILED_REQUESTS | SS | Y |
| Q8STTOUT | ACCEL_TIMED_OUT_REQUESTS | SS | Y |
| Q8STBLKR | ACCEL_BLOCKS_RETURNED | SS | Y |
| Q8STBLKS | ACCEL_BLOCKS_SENT | SS | Y |
| Q8STBYTR | ACCEL_BYTES_RETURNED | SS | Y |
| Q8STBYTS | ACCEL_BYTES_SENT | SS | Y |
| Q8STMSGR | ACCEL_MESSAGES_RETURNED | SS | Y |
| Q8STMSGS | ACCEL_MESSAGES_SENT | SS | Y |
| Q8STROWR | ACCEL_ROWS_RETURNED | SS | Y |
| Q8STROWS | ACCEL_ROWS_SENT | SS | Y |
| Q8STTCPU | ACCEL_TCP/IP_CPU_TIME | SS | Y |
| Q8STTELA | ACCEL_TCP/IP_ELAPSED_TIME | SS | Y |
| Q8STAELA | ACCEL_ELAPSED_TIME | SS | Y |
| Q8STAWAT | ACCEL_WAIT_TIME | SS | Y |
| Q8STACPU | ACCEL_CPU_TIME | SS | Y |
| Q8STDSKB | ACCEL_DISK_STORAGE_INUSE | SS | N |

# IDAA Statistics – Obsolete

| Variable Name | EPV Label |
|---|---|
| Q8STAVGQ03 | ACCEL_AVG_QUERY_QUEUE_LAST3H |
| Q8STAVGQ24 | ACCEL_AVG_QUERY_QUEUE_LAST24H |
| Q8STFINV | ACCEL_FAILED_QUERIES_INVALID_STATE |
| Q8STMIPS | ACCEL_PROCESSING_CAPACITY |
| Q8STSKEW | ACCEL_DATA_SKEW |
| Q8STCPMU | ACCEL_AVG_MEM_USAGE_COORD |
| Q8STWPMU | ACCEL_AVG_MEM_USAGE_WORKER |
| Q8STWSMA | ACCEL_AVG_SHRMEM_AVAILABLE_WORKER |
| Q8STWSMM | ACCEL_AVG_SHRMEM_USAGE_WORKER |
| Q8STWSMU | ACCEL_MAX_SHRMEM_USAGE_WORKER |
| Q8STACTV | ACCEL_CURR_EXECUTING_QUERIES_AC |
| Q8STMAXA | ACCEL_MAX_EXECUTING_QUERIES_AC |
| Q8STMAXQ | ACCEL_HWM_QUEUE_LENGTH_AC |
| Q8STCCPU | ACCEL_AVG_CPUBUSY_COORD_AC |
| Q8STWCPU | ACCEL_AVG_CPUBUSY_WORKER_AC |
| Q8STWNOD | ACCEL_ACTIVE_WORKER_NODES_AC |
| Q8STSELA | ACCEL_SERVICES_ELAPSED_TIME |
| Q8STAELA | ACCEL_ELAPSED_TIME |

# IDAA Statistics – Replication

| Variable Name | EPV Label | Scope | Accum |
|---|---|---|---|
| Q8STNIA | ACCEL_REPL_INSERT_ROWS_AC | AC | |
| Q8STNUA | ACCEL_REPL_DELETE_ROWS_AC | AC | |
| Q8STNDA | ACCEL_REPL_UPDATE_ROWS_AC | AC | |

accelerator scope variables

| Variable Name | EPV Label | Scope | Accum |
|---|---|---|---|
| Q8STNLRS | ACCEL_LOG_RECORDS_READ | SS | |
| Q8STNLTS | ACCEL_LOG_RECORDS_APPLICABLE | SS | |
| Q8STNBS | ACCEL_LOG_BYTES_PROCESSED | SS | |
| Q8STTCCS | ACCEL_REPL_CPU_TIME | SS | |
| Q8STCSS | ACCEL_REPL_STATUS | SS | |
| Q8STLSC | ACCEL_REPL_STATUS_CHANGE | SS | |
| Q8STNIS | ACCEL_REPL_INSERT_ROWS | SS | |
| Q8STNUS | ACCEL_REPL_DELETE_ROWS | SS | |
| Q8STNDS | ACCEL_REPL_UPDATE_ROWS | SS | |
| Q8STCRL | ACCEL_REPL_LATENCY | SS | |

subsystem scope variables

# Accelerator Modeling

# Accelerator Modeling

- The new ACCELMODEL parameter (which can be set to YES or NO) provides indicators for possible CPU and elapsed time savings if IDAA was available
- It does not require presence of the accelerator
- If set to YES, DB2 accounting records include projected CPU, zIIP and elapsed time savings but …
- Some additional actions are needed

# Accelerator Modeling

- *For dynamic SQL statement assessment, use the RUNSTATS utility for target table spaces or restart DB2 to invalidate corresponding dynamic statement cache entries. ACCELMODEL needs to be set to YES during statement execution for Accelerator modeling.*

from ‘Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1’
IBM Redbooks SG24-8213-00

# Accelerator Modeling

- *For static SQL statement assessment, you must issue a BIND or REBIND PACKAGE for the DB2 packages you plan to include for this workload assessment. You may use option APREUSE(ERROR) with the REBIND PACKAGE statement to make sure that the current access plan remains unchanged. ACCELMODEL=YES is required during BIND/REBIND only, not at static SQL statement execution time.*

from 'Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1'
IBM Redbooks SG24-8213-00

# Accelerator Modeling

- Functionality delivered via two DB2 10 APARs
  - PM90886: for IDAA V3 (dynamic SQL support)
  - PM95035: for IDAA V4 (add static SQL support)
- Functionality delivered via one DB2 11 APAR
  - PM96478: for IDAA V4 (add static SQL support)

# Accelerator Modeling

- Three new accounting fields will be added to the QWAC section of the IFCID 3:
  - QWAC_ACCEL_ELIG_ELA, the accumulated elapsed time spent processing SQL in DB2 that may be eligible for execution on an accelerator
  - QWAC_ACCEL_ELIG_CP, the accumulated CPU time spent processing SQL in DB2 that may be eligible for execution on an accelerator
  - QWAC_ACCEL_ELIG_SE, the accumulated CPU time consumed on an IBM specialty engine while processing SQL in DB2 that may be eligible for execution on an accelerator

# Accelerator Modeling

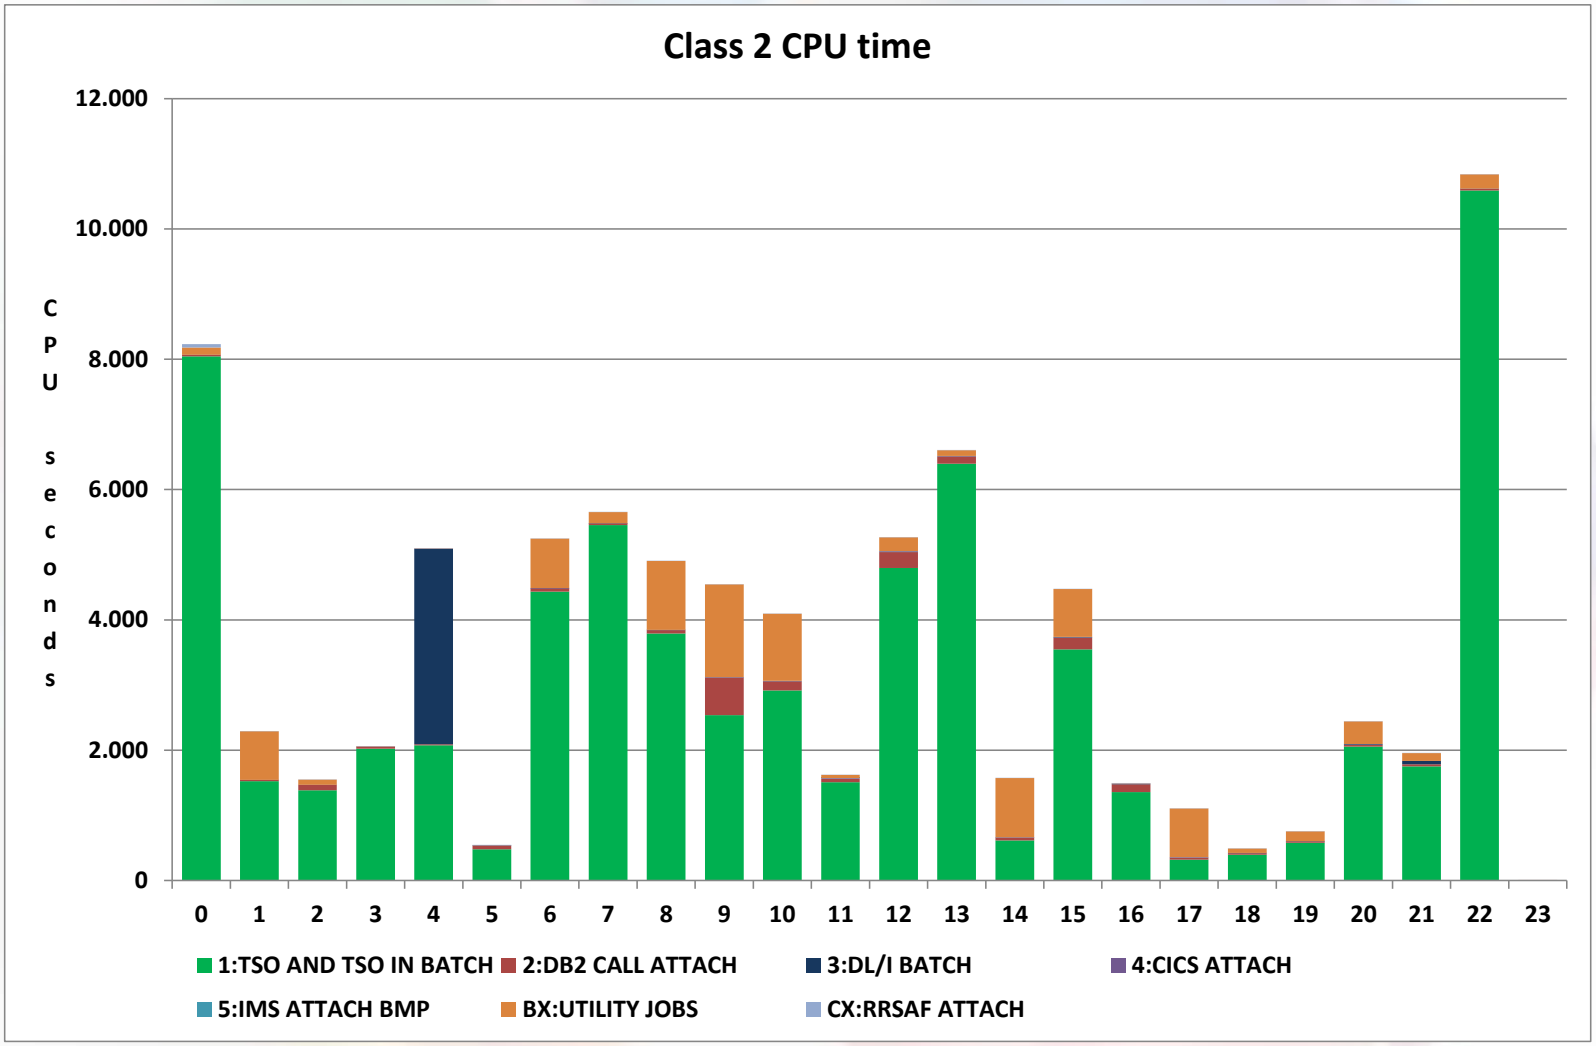


# Accelerator Modeling

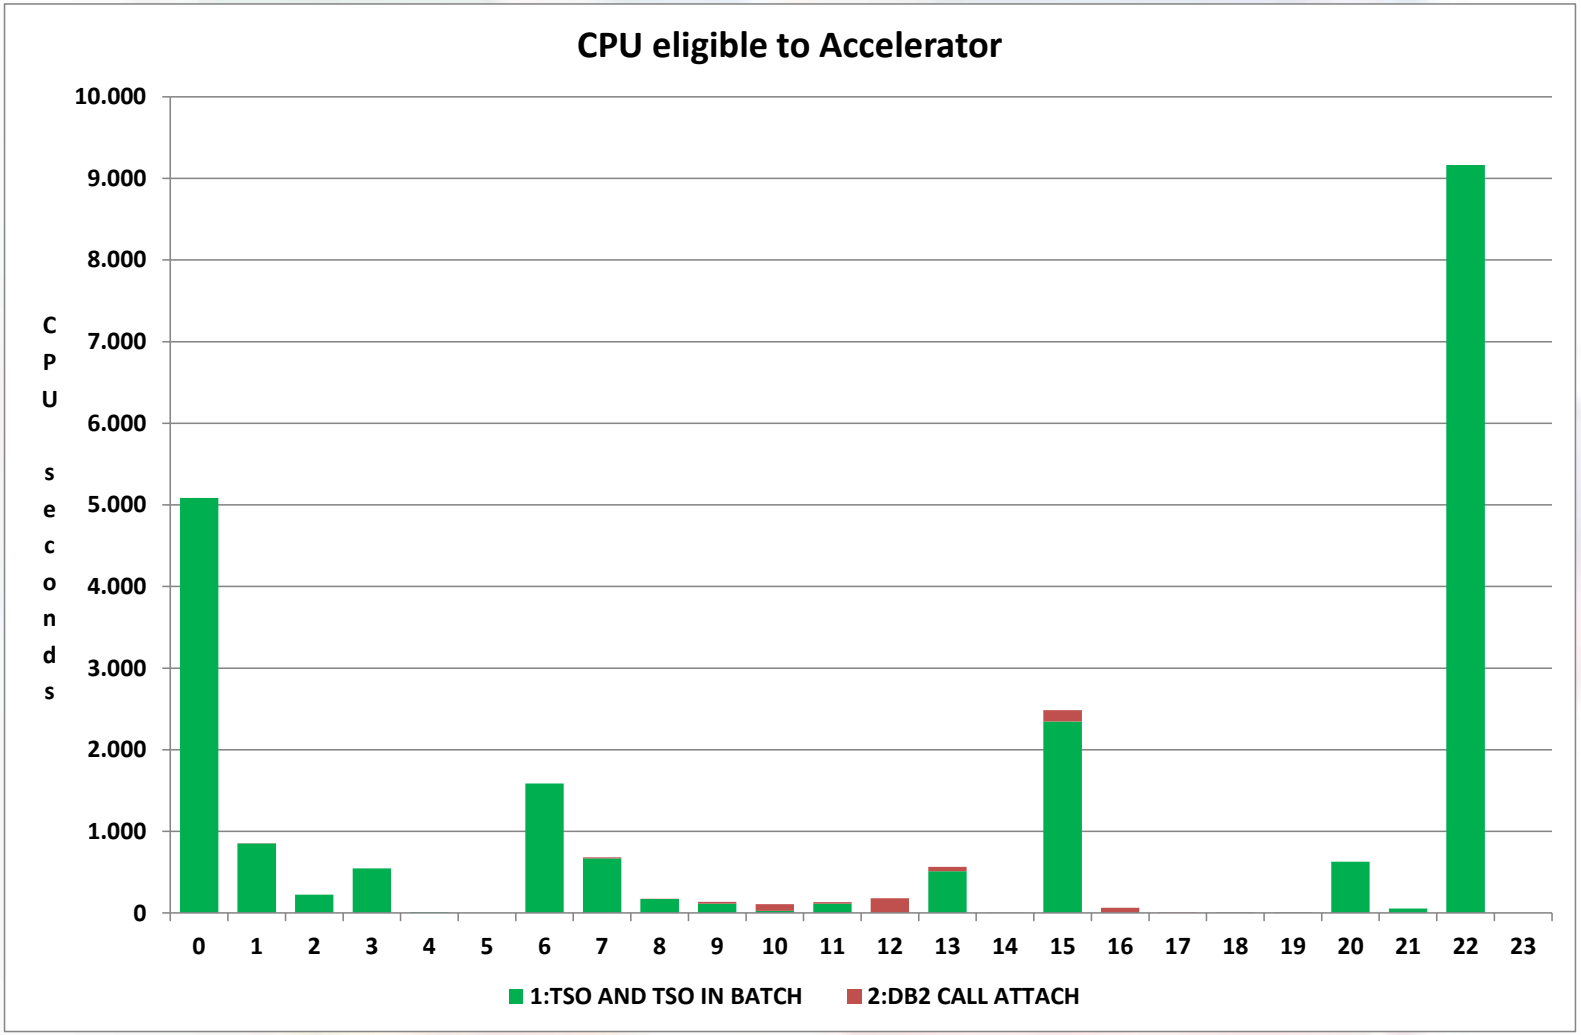


# What about Money ?

# What about Money ?

- IDAA is not free
- So the question is: IDAA CPU savings will justify the additional IDAA cost ?
- Every company should answer this question; the answer will mostly depend on workload characteristics and software licensing policies

# What about Money ? - WLC

- In this case we have to evaluate the reduction which could be obtained in the monthly peak of the MSU used in the 4 hour rolling average
- MSU used by replication are not considered here

# What about Money ? - WLC

- The following steps need to be performed to estimate MSU savings:
  - Activate accelerator modeling on the appropriate workloads
  - Collect SMF 101 data for an adequate period of time
  - Distribute the CPU consumptions of long running threads across the elapsed hours
  - Calculate the hourly CPU seconds used by summing all address space or service class consumption across all the CEC LPARs
  - Calculate the hourly CPU seconds eligible to the IDAA by summing QWAC_ACCEL_ELIG_CP across all the LPARs

# What about Money ? - WLC

- Increment both values in order to account for the uncaptured CPU
- Transform CPU seconds in MSU
- Subtract the IDAA eligible MSU from the total MSU
- Calculate the monthly peak of the MSU used in the 4-hour rolling average
- Compare this value with the «real» value of the monthly peak of the 4-hour rolling average MSU
- Convert the MSU reduction, if any, to money savings

# What about Money ? - WLC

- No WLC savings for the first customer (slides 37, 38)
- We tested the process with a second customer
- It seems that interesting MSU savings could be expected

WLC WITH IDAA DAILY MSU

CEC : CEC1

| DATE | DAY | TYPE | MODEL | MSU | MSU WLC | MSU W/IDAA |
|---|---|---|---|---|---|---|
| 2014-06-24 | Tue | 2827 | 712 | 1.709 | 721 | 668 |
| 2014-06-23 | Mon | 2827 | 712 | 1.709 | 704 | 633 |
| 2014-06-26 | Thu | 2827 | 712 | 1.709 | 688 | 617 |
| 2014-06-25 | Wed | 2827 | 712 | 1.709 | 638 | 581 |
| 2014-06-27 | Fri | 2827 | 712 | 1.709 | 603 | 546 |

WLC WITH IDAA DAILY MSU

CEC : CEC2

| DATE | DAY | TYPE | MODEL | MSU | MSU WLC | MSU W/IDAA |
|---|---|---|---|---|---|---|
| 2014-06-26 | Thu | 2827 | 712 | 1.709 | 795 | 703 |
| 2014-06-23 | Mon | 2827 | 712 | 1.709 | 720 | 654 |
| 2014-06-24 | Tue | 2827 | 712 | 1.709 | 700 | 618 |
| 2014-06-27 | Fri | 2827 | 712 | 1.709 | 685 | 628 |
| 2014-06-25 | Wed | 2827 | 712 | 1.709 | 609 | 554 |

# What about Money ? – Full Capacity

- In this case we have to evaluate the reduction which could be obtained in the CPU capacity needed to cover the peak period used for capacity planning
- MSU used by replication are not considered here

# What about Money ? – Full Capacity

- The following steps need to be performed to estimate savings:
  - Activate accelerator modeling on the appropriate workloads
  - Collect SMF 101 data for an adequate period of time
  - Distribute the CPU consumptions of long running threads across the elapsed hours
  - Calculate the CPU seconds used by summing all address space or service class consumption across all the CEC LPARs aggregated as of your capacity planning rules
  - Calculate the hourly CPU seconds eligible to the IDAA by summing QWAC_ACCEL_ELIG_CP across all the LPARs aggregated as of your capacity planning rules

# What about Money ? – Full Capacity

- Increment both values in order to account for the uncaptured CPU
- Subtract the IDAA eligible CPU seconds from the total CPU seconds
- From the obtained CPU seconds estimate a new hardware configuration
- Evaluate a possible hardware reduction
- Transform this reduction, if any, to hardware and software money savings

# Conclusions (1/2)

- We found a growing interest in many companies towards the IDAA especially after V4 has been released
- Many SMF metrics are available to control IDAA performance and for IDAA capacity planning
- However their meaning, especially for the statistics included in the IFCID 2, is not really straightforward and has to be carefully understood before using them
- Some additional information could be useful to better understand what's happening inside the IDAA

# Conclusions (2/2)

- The major interest of many companies is the possibility of cutting z/OS costs through CPU usage reduction
- This possibility has to be carefully verified taking into account:
  - workload characteristics
  - service level agreements
  - software licensing policy
  - IDAA replication cost

# Questions ?