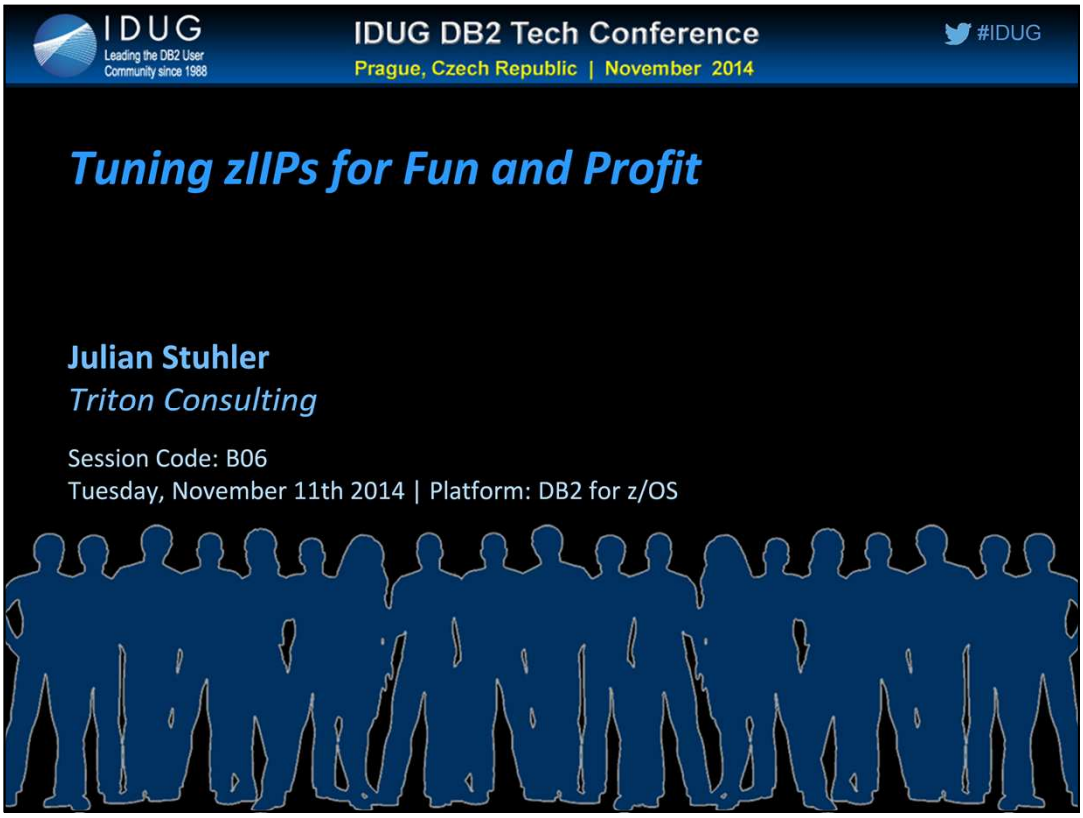

In the past few releases, DB2 for z/OS has steadily increased its exploitation of zIIP speciality processors. In addition to providing significant potential for mainframe cost reduction, zIIP offload can have a positive (or negative) impact on DB2 query performance. This presentation will review DB2 zIIP usage by release, and provide the attendee with some practical guidance on zIIP capacity planning and usage monitoring.

Acknowledgements

- John Lemmon, Triton Consulting

- Adrian Burke, IBM
- Florence Dubois, IBM
- Willie Favero, IBM
- James Gill, Triton Consulting
- Martin Packer, IBM

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- How Does DB2 Exploit zIIPs?
- Performance Considerations
- zIIP Usage Measurement and Capacity Planning
- Case Study
- Summary

Introduction

- DB2 consultant with Triton Consulting, based in the UK
- 26 years DB2 experience
  - Database Administration
  - Systems Programming
  - Application Development
- IBM Gold Consultant
- IBM Champion
- IBM White Papers, Redbooks, Flashbooks, etc
- IDUG Best Speaker and Past President

Topics

- Introduction
- **What Are zIIPs and How Can They Help?**
- How Does DB2 Exploit zIIPs?
- Performance Considerations
- zIIP Usage Measurement and Capacity Planning
- Case Study
- Summary

What Are zIIPs?

- zIIP – z/OS Integrated Information Processor
- IBM System z speciality processor (SP)
  - Able to execute work that has been redirected from a Central Processor (CP) under certain conditions
  - Work that runs on a speciality processor does not count towards z/OS Monthly Licence Charge (MLC) software fees
  - Other speciality processors include zAAP, IFL, ICF
- zIIP designed to support enclave SRB workloads

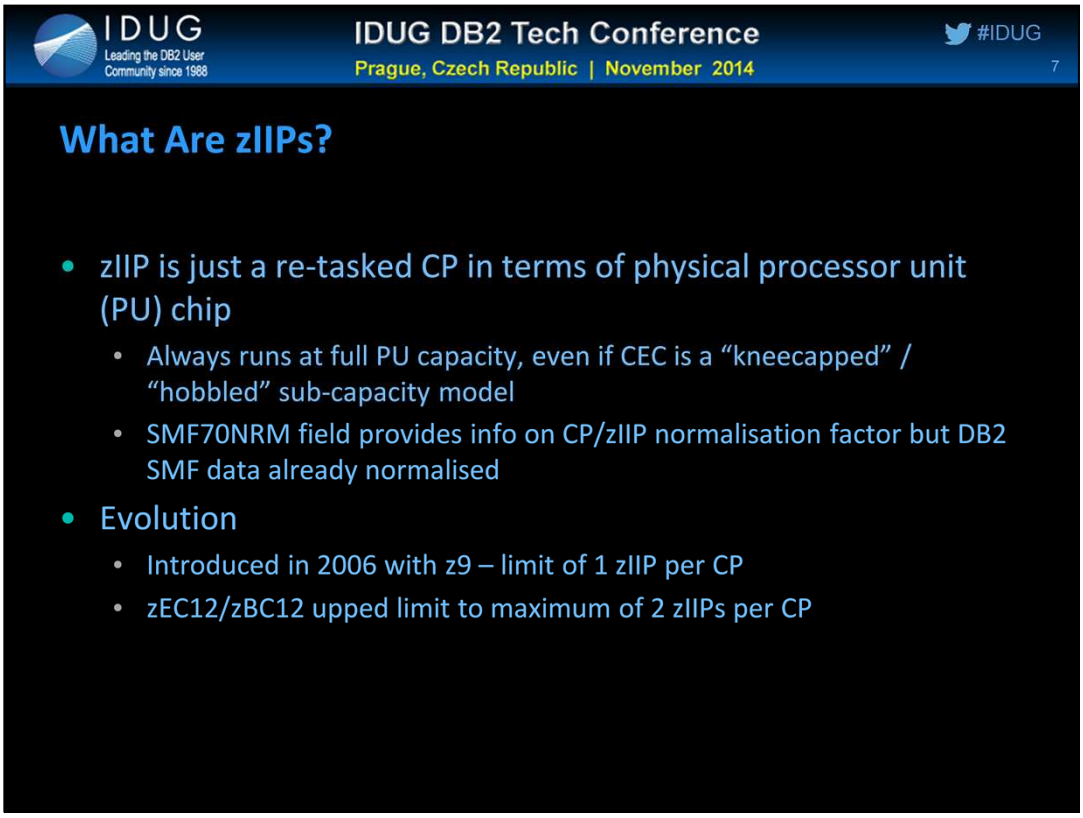

Field SMF70NRM gives us a number which we can divide by 256 and see how much faster the specialty engines are than the CPs

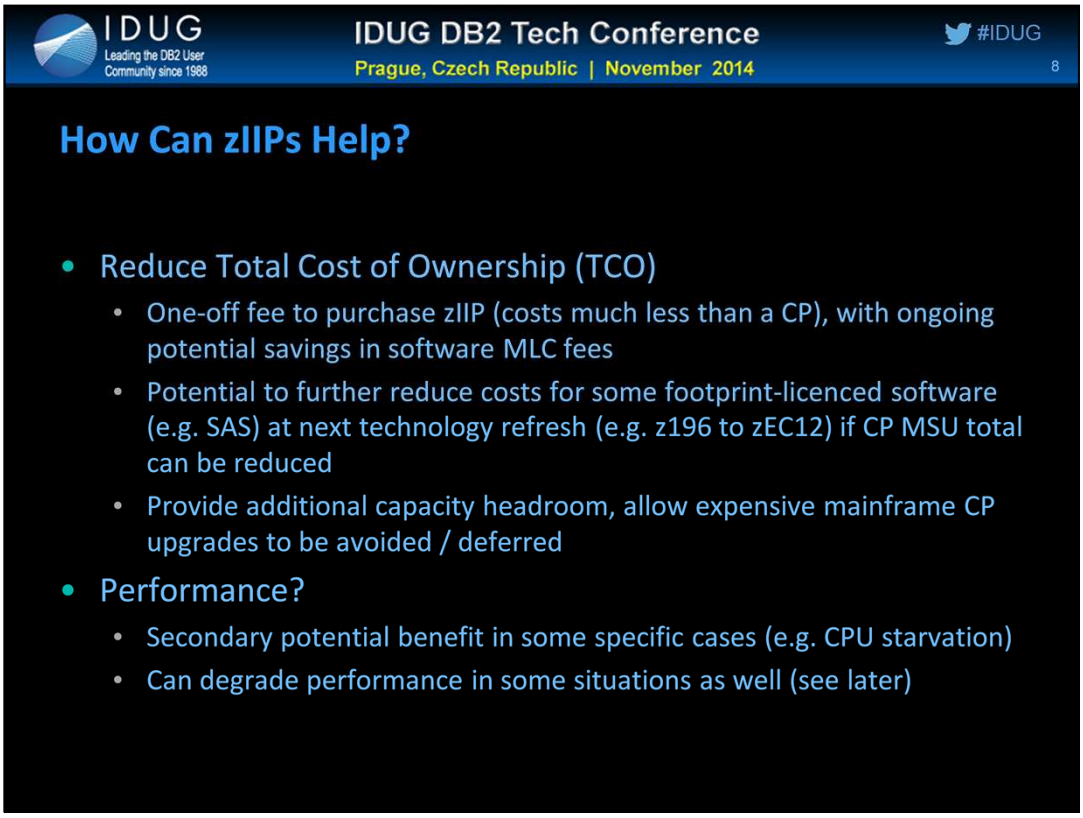

Also allows next technology refresh (e.g. z196 => zEC12) to provide less total MSUs and therefore a reduction in some ISV software – e.g. SAS is footprint licence.

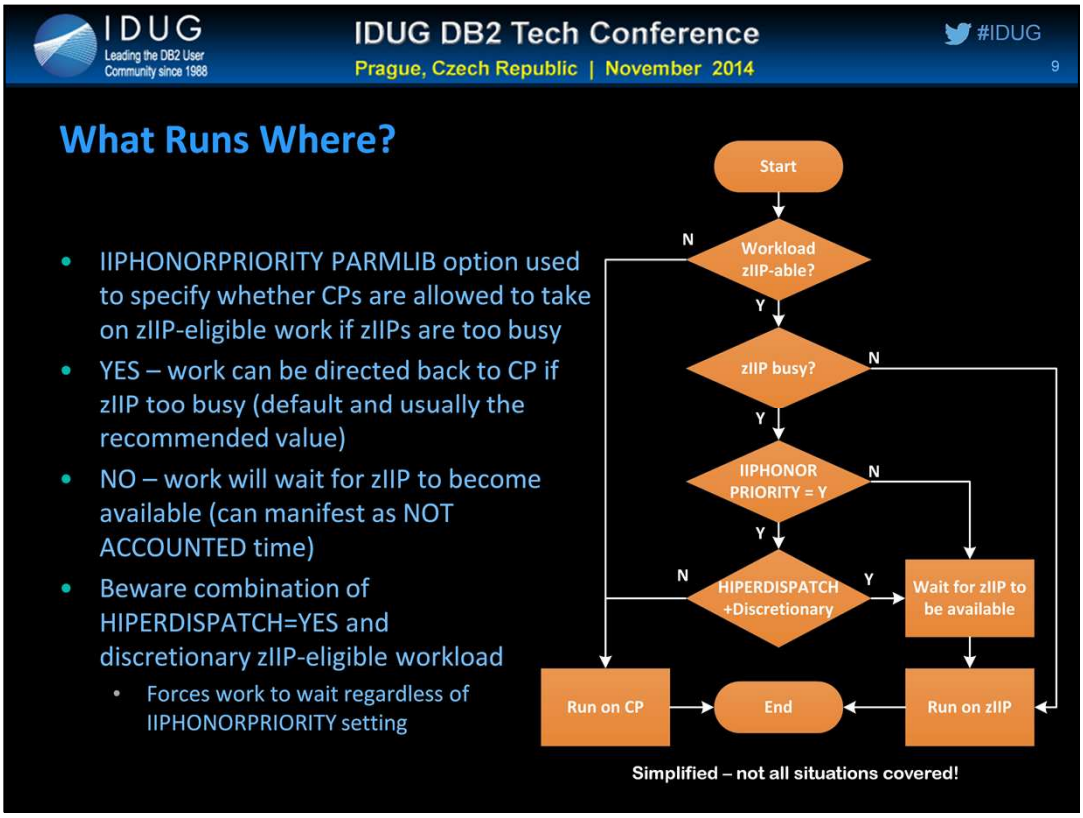

HiperDispatch is a workload dispatching feature found in the newest IBM mainframe models (the System z10 and IBM zEnterprise System processors) running recent releases of z/OS. HiperDispatch was introduced in February, 2008.

One of the engineering challenges with large SMP server designs is to maintain near-linear scalability as the number of CPUs increases. Performance and throughput do not double when doubling the number of processors. There are many overhead factors, including contention for cache and main memory access. These overhead factors become increasingly difficult to mitigate as the number of CPUs increases. The design goal for delivering maximum performance is to minimize those overhead factors. Each new mainframe model supports a higher maximum number of CPUs (up to 64 main processors in a single System z10 mainframe for example), so this engineering challenge becomes ever more important.

HiperDispatch helps address the problem through a combination of hardware features, z/OS dispatching, and the z/OS Workload Manager. In z/OS there may be tasks waiting for processing attention, such as transaction programs. Each task often requires access to memory. In a large SMP design such as System z, some CPUs are physically "closer" with faster access to cache memory that might hold supporting data for particular tasks. HiperDispatch exploits this fact and steers tasks to the CPUs most likely to have the fastest access to relevant data already in cache. If that particular CPU is busy, HiperDispatch will, at first, wait for it to

finish its other task, even if another less favorable CPU is idle. However, there are limitations to how patient HiperDispatch will be, as governed by Workload Manager goals. If z/OS Workload Manager senses that there's a risk the pending task will miss its service level (responding within a certain number of milliseconds to a user request for example), Workload Manager and HiperDispatch will send the task over to an idle CPU for processing, even if that CPU must fetch data from slower main memory.

What Runs Where?

- ZIIPAWMT PARMLIB option used to specify how often WLM checks to see if a zIIP needs help
  - Fixed at 3200 for HIPERDISPATCH=YES, so delay of up to 3.2 m/s if zIIP too busy and workload has to be directed back to a CP
- ZIIPMAXQL PARMLIB option used to specify maximum number of dispatchable units that will queue waiting for a zIIP processor
  - Default is 7
- IBM advice is not to change either of these unless advised by L2 support

Terminology

- zIIP-able
  - Workload that is eligible for zIIP offload
- zIIP-ed
  - zIIP-able workload that is actually executed on a zIIP
- Un-zIIP-ed
  - zIIP-able workload that ends up being executed on a CP

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- **How Does DB2 Exploit zIIPs?**
- Performance Considerations
- zIIP Usage Measurement and Capacity Planning
- Case Study
- Summary

DB2 zIIP Exploitation

- DB2 has steadily made more and more work zIIP-able from V8 onwards, and trend is expected to continue
  - Significant jump in exploitation in DB2 10 onwards, capacity planning becomes much more important!
- Some consistent themes for offloadable work
  - "New" workloads – distributed access, XML
  - Utilities
  - Async system tasks (prefetch, pseudo-deleted index clean-up)
- Good instrumentation is available
  - Both general and DB2-specific SMF records – see later

DB2 zIIP Exploitation

- DB2 is currently the major (but not the only) zIIP exploiter
  - DFSORT
  - z/OS Communications Server (IPSec encryption, HiperSockets Multiple Write)
  - z/OS Global Mirror (aka XRC)
  - zAAP on zIIP is driving usage (WAS on z/OS, Java stored procedures)
  - Many ISV products / functions
- DB2 is the first exploiter to begin making performance critical tasks zIIP-able

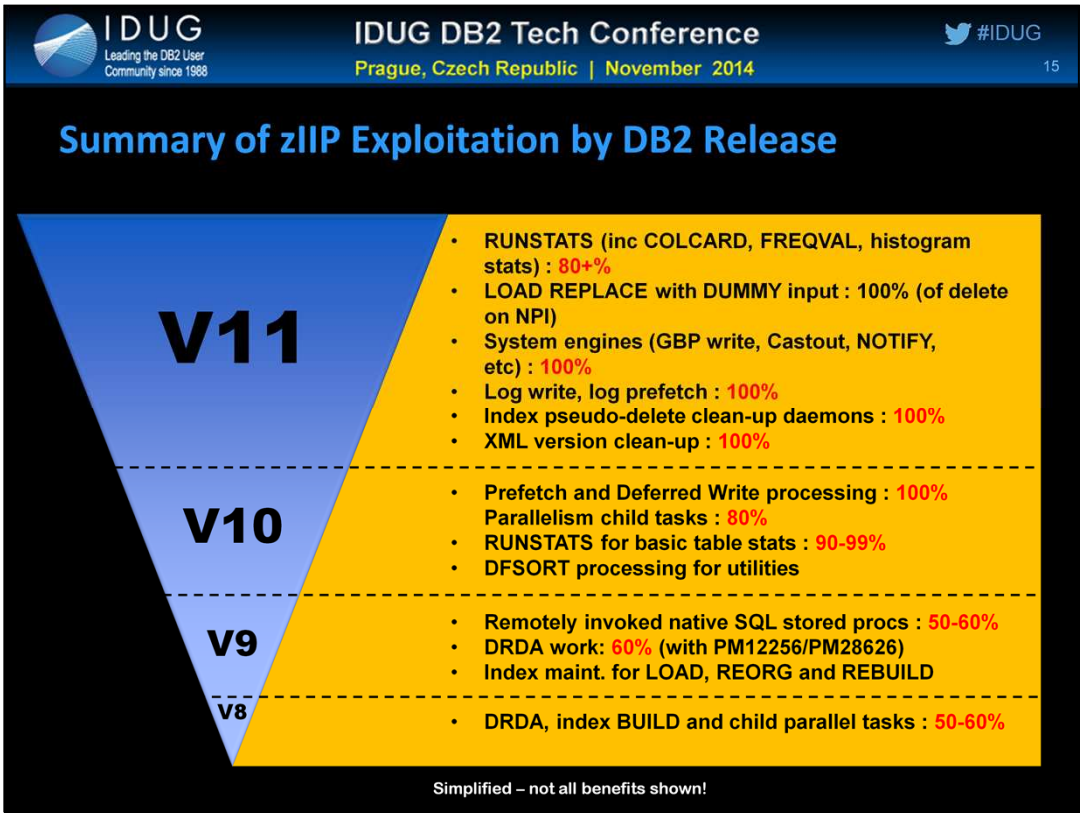

Prefetch and Deferred Write processing : 100% (c. 70% of DBM1 SRB time)

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- How Does DB2 Exploit zIIPs?
- **Performance Considerations**
- zIIP Usage Measurement and Capacity Planning
- Case Study
- Summary

Performance Considerations

- zIIPs are primarily intended as a means of reducing TCO by managing MLC costs, not a performance boost
- BUT… zIIPs can certainly have in impact on performance!
- Performance can sometimes be improved if you are driving your CPs very hard and workload is suffering from CPU starvation
  - Enabling offload to zIIP effectively increases overall compute capacity, thereby reducing CPU starvation and elapsed time
- Performance can often be degraded if you're driving your zIIPs too hard
  - Wait for zIIP, or wait to be redirected back to CP

zIIP Utilisation – How Much is Too Much?

- Many customers drive their CPs very hard (in the 90%-100% utilisation range), and that's OK
  - Entire infrastructure is designed to cope with very high utilisation levels
- zIIPs typically cannot be driven nearly as hard without significant performance issues
  - Aiming for average of 30-50% utilisation per RMF interval is good rule of thumb (peaks can of course be higher)
- DB2 10 made several performance-critical system tasks zIIP-able (and DB2 11 adds some more)
  - Performance impact of insufficient zIIP capacity is increasing with each release

zIIP Utilisation – How Much is Too Much?

- We take a triple hit if redirecting back to CP (un-zIIP-ed)
  - Delay caused by waiting for zIIP then going back to CP
  - Workload will run on CP, so contributes to 4HRA and could impact MLC fees
  - Work may take longer to execute if being redirected back to sub-capacity CP (zIIPs always run at full PU capacity)
- Disabling zIIP (i.e. taking offline from LPAR) can be an effective short-term measure
  - May have MLC implications!
  - Proper long-term solution is probably to purchase more zIIP capacity

Specific Situations to Avoid

- DB2 10 and above, HIPERDISPATCH=YES, IIPHONORPRIORITY=YES, zIIPs too busy (averaging >> 30-50%)
  - Potential for significant slowdown in prefetch activity, leading to widespread DB2 performance issues
  - See Flo's blog: https://www.ibm.com/developerworks/community/blogs/22586cb0-8817-4d2c-ae74-0ddcc2a409bc/entry/december_17_2012_6_07_am3?lang=en
- zIIP-able work, HIPERDISPATCH=YES, IIPHONORPRIORITY=YES, performance critical workload incorrectly classified as discretionary, over-stressed zIIP
  - zIIP-able work will still be forced to run on zIIP, potentially leading to significant elapsed time issues
  - See Willie's blog: http://db2onlinehandbook.com/?p=566

Specific Situations to Avoid

- Low-weight LPARs with low zIIP usage, sharing a single zIIP
  - Possible latency issues waiting for logical zIIP to be dispatched if used by other LPARs

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- How Does DB2 Exploit zIIPs?
- Performance Considerations
- **zIIP Usage Measurement and Capacity Planning**
- Case Study
- Summary

zIIP Capacity Planning is important

- DB2 10 and DB2 11 allow significantly more offloading
  - Time to get serious about zIIP Capacity Planning if you want to optimise Cost Reduction opportunities
- Understanding how IBM software pricing works isn't essential, but it will certainly help
- Tools
  - General SMF/RMF (30, 70, 72)
  - DB2 SMF (100, 101)
  - Other (SCRT)
- Make sure you set PARMLIB PROJECTCPU=YES, this directs z/OS to record the amount of work eligible for zIIP even if you don't have one installed

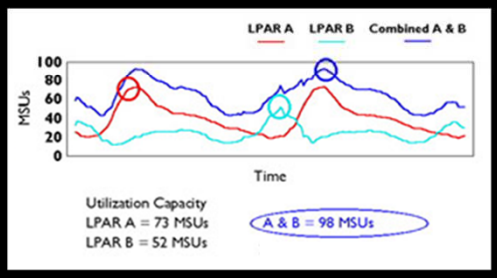

SMF 30 Data

- SMF 30 (Common Address Space Work / Activity)
  - 30.2 & 30.3 – Written at end of interval
  - 30.4 – Written at end of step
  - 30.5 – Written at end of job
- CPU times are in hundredths of seconds (.01 seconds)
- Useful zIIP fields
  - SMF30_TIME_ON_ZIIP – actual zIIP offload (zIIP-ed)
  - SMF30_TIME_ZIIP_ON_CP – zIIP-eligible workload redirected to CP (un-zIIP-ed)
- Broken down by LPAR, Address Space, WLM name, Service Class, Userid, Account-code…

SMF 30 data can tell us…

- Current zIIP offload (zIIP-ed) in relation to total standard CP CPU
- zIIP engines always run at full capacity, so to calculate CPU in standard CP, you'll need to normalise
  - You'll use more CPU on sub-capacity standard CP

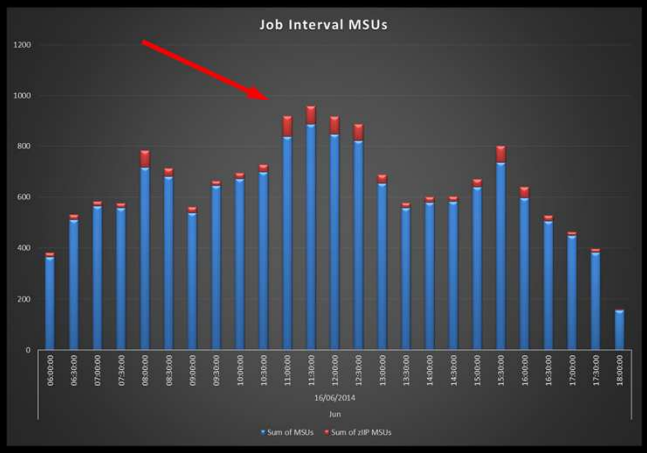

SMF 30 data can tell us…

- Current zIIP offload (zIIP-ed) broken down by Service Class
- SSYHI looks like it may contain DB2 in this example…

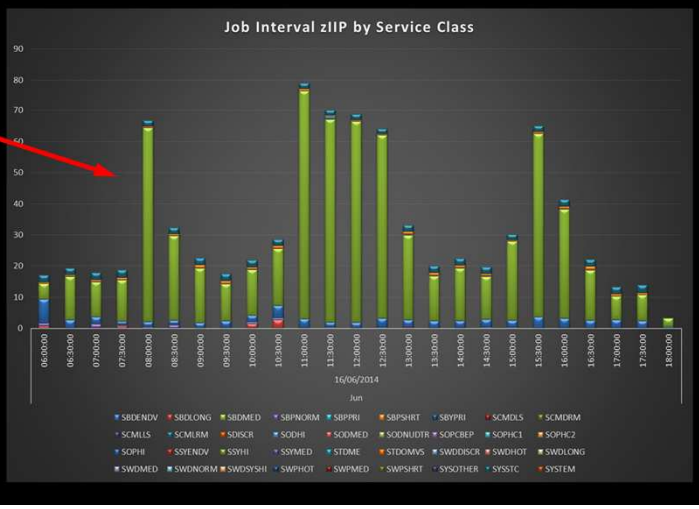

SMF 30 data can tell us…

- To confirm, we can then chart the same data by PGM prefix
- Same data as shown on previous chart
  - DSN* and DXR* for DB2

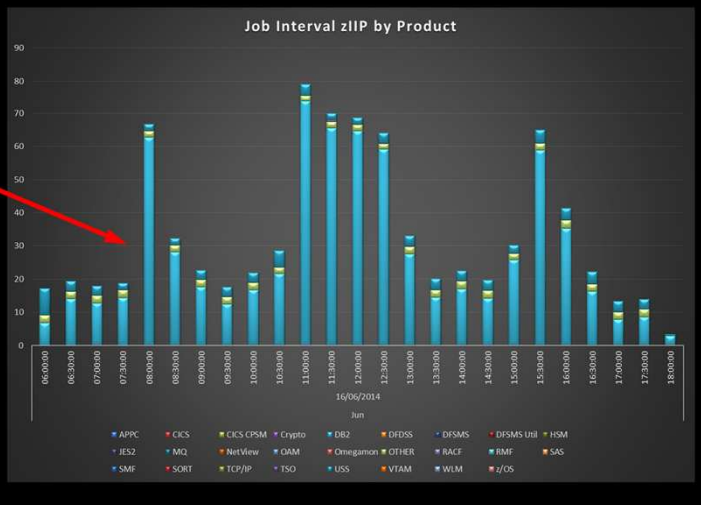

SMF 30 data can tell us…

- Overall zIIP utilisation
- Shown as
  - MSU (remember to normalise)
  - zIIP %busy (remember, ROT is to not exceed 35-50% busy)

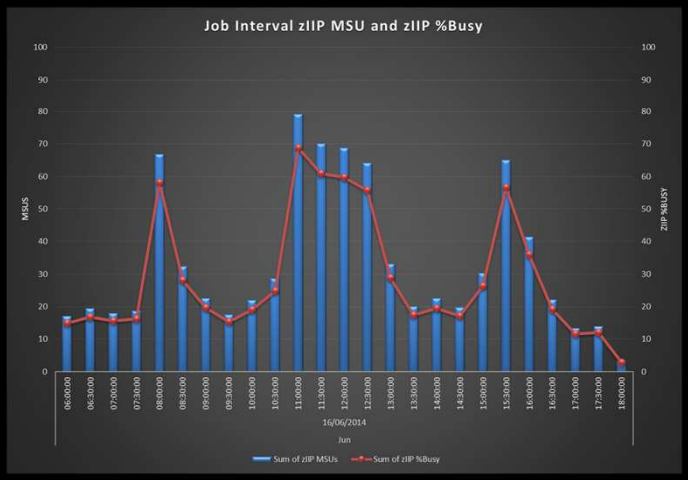

SMF 30 data for DBM1 can tell us…

- If you're still on V9, measuring current zIIP utilisation won't give you much of a handle on what will happen in V10
  - Prefetch / deferred write not reported as IIPCP (un-zIIPed)
- Different approach needed
  - Up to 100% of DBM1 address space becomes eligible at v10+, so use SMF30 to measure total DBM1 MSU usage to get zIIP estimate

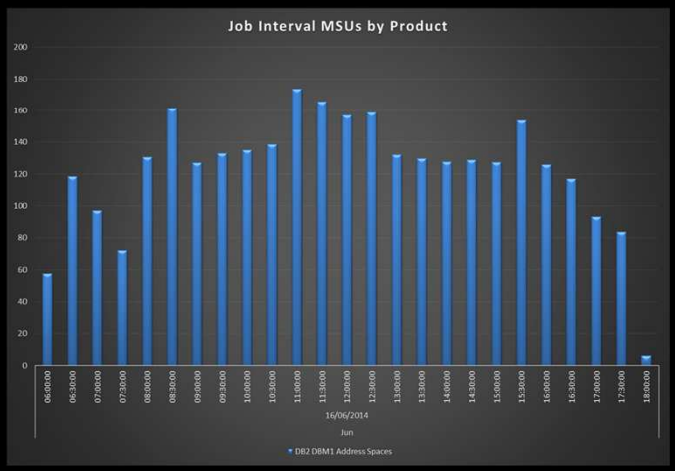

SMF 72 data can tell us…

- RMF Report - SYSRPTS(WLMGL(SCPER))

```
REPORT BY: POLICY=PRIMARY    WORKLOAD=WDB2      SERVICE CLASS=SDDFHI     RESOURCE GROUP=*NONE      PERIOD=2 IMPORTANCE=3
                                                CRITICAL     =NONE

-TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE---   SERVICE TIME  ---APPL %---  --PR
AVG       2.35  ACTUAL           35.548  SSCHRT  1144  IOC         0   CPU 2554.262  CP     227.09  BLK
MPL       2.35  EXECUTION        35.541  RESP     1.7  CPU    858575K  SRB    0.000  AAPCP    0.00  ENQ     0.000  TOTAL      0.00
ENDED       16  QUEUED               6  CONN     0.7  MSO         0  RCT    0.000  IIPCP  166.15  CRM     0.000  SHARED     0.00
END/S     0.02  R/S AFFIN            0  DISC     0.5  SRB         0  IIT    0.000                  LCK     0.001
#SWAPS       0  INELIGIBLE           0  Q+PEND   0.5  TOT    858575K  HST    0.000  AAP       N/A  SUP     0.000  -PAGE-IN RATES-
EXCTD        0  CONVERSION           0  IOSQ     0.0  /SEC    954292  AAP      N/A  IIP     56.81                 SINGLE     0.0
AVG ENC   2.35  STD DEV        2.15.178                                 IIP  511.138                               BLOCK      0.0
REM ENC   0.00                                        ABSRPTN   407K                                             SHARED     0.0
MS ENC    0.00                                        TRX SERV  407K                                             HSP        0.0
```

Un-zIIP-ed work – 1.66 engines.

zIIP-ed offload – 0.56 engines.

SMF 70 data can tell us…

- RMF Report – REPORTS(CPU)

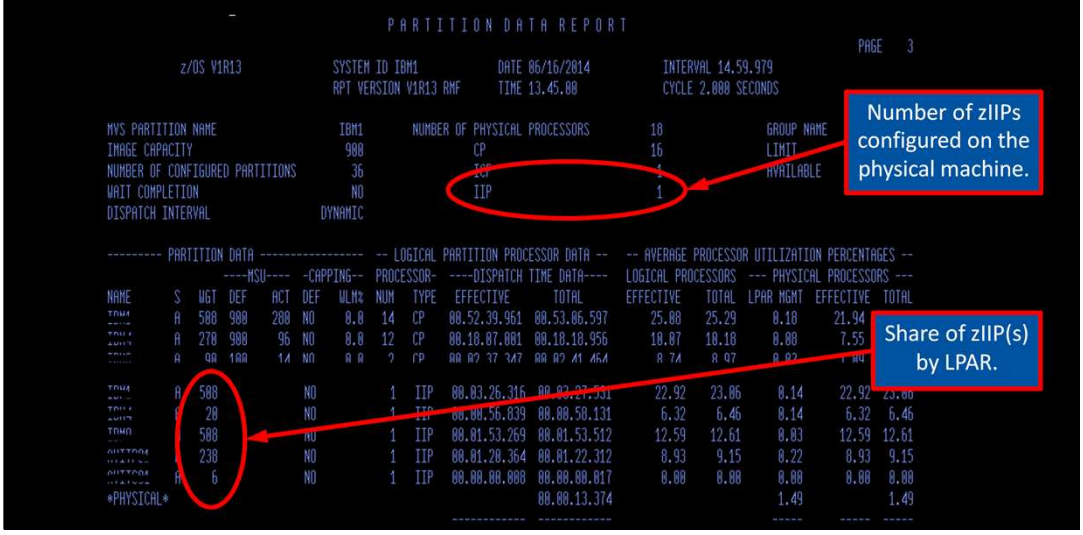

```
                          P A R T I T I O N   D A T A   R E P O R T                    PAGE    3

          z/OS V1R13              SYSTEM ID IBM1           DATE 06/16/2014         INTERVAL 14.59.979
                                  RPT VERSION V1R13 RMF    TIME 13.45.00           CYCLE 2.000 SECONDS

MVS PARTITION NAME                      IBM1       NUMBER OF PHYSICAL PROCESSORS       18            GROUP NAME
IMAGE CAPACITY                           900                  CP                       16            LIMIT
NUMBER OF CONFIGURED PARTITIONS           36                  ICF                       1            AVAILABLE
WAIT COMPLETION                           NO                  IIP                       1
DISPATCH INTERVAL                    DYNAMIC

--------- PARTITION DATA ----------------  -- LOGICAL PARTITION PROCESSOR DATA --   -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
                  ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----   LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
NAME      S  WGT  DEF   ACT  DEF  WLM%  NUM  TYPE   EFFECTIVE      TOTAL          EFFECTIVE   TOTAL  LPAR MGMT  EFFECTIVE  TOTAL
TPM4      A  588  900   288  NO    0.0   14  CP   00.52.39.961  00.53.06.597       25.08   25.29     0.18      21.94
TPM7      A  278  900    96  NO    0.0   12  CP   00.18.07.881  00.18.18.956       10.07   10.18     0.00       7.55
TPM8      A   98  100    14  NO    0.0    2  CP   00.02.37.347  00.02.41.464        0.14    0.97     0.03       1.00

TPM4      A  588         NO             1  IIP   00.03.26.316  00.03.27.531       22.92   23.06     0.14      22.92   23.06
TPM7      A   28         NO             1  IIP   00.00.56.839  00.00.58.131        6.32    6.46     0.14       6.32    6.46
TPM0      A  588         NO             1  IIP   00.01.53.269  00.01.53.512       12.59   12.61     0.03      12.59   12.61
OUTTPM4   A  238         NO             1  IIP   00.01.20.364  00.01.22.312        8.93    9.15     0.22       8.93    9.15
OUTTPM1   A    6         NO             1  IIP   00.00.00.000  00.00.00.017        0.00    0.00     0.00       0.00    0.00
*PHYSICAL*                                      00.00.13.374                                1.49                          1.49
                                                ------------  ------------                  -----                 -----   -----
```

Number of zIIPs configured on the physical machine.

Share of zIIP(s) by LPAR.

DB2 Stats can tell us…

- DB2 statistics reports show zIIP offload from system address spaces
- DB2 10
  - Prefetch and deferred write processing for DBM1
  - PM30468 ensures zIIP-ed workload correctly reported under DBM1
- DB2 11
  - Log write, log prefetch MSTR
  - Index pseudo-delete clean-up daemons for DBM1

zIIP-ed offload

CPU TIMES	TCB TIME	PREEMPT SRB	NONPREEMPT SRB	CP CPU TIME	PREEMPT IIP SRB
SYSTEM SERVICES ADDRESS SPACE	0.120789	0.037248	0.027924	0.185961	0.000000
DATABASE SERVICES ADDRESS SPACE	0.403656	0.218449	0.009574	0.631680	0.000000
IRLM	0.000043	0.000000	0.224940	0.224983	0.000000
DDF ADDRESS SPACE	0.005795	1.122965	0.007127	1.135886	0.000000
TOTAL	0.530282	1.378663	0.269566	2.178511	0.000000

DB2 Stats can tell us…

- Worth monitoring LC24 counter as possible indication of zIIP starvation issues from V10 CM onwards
  - LC24 counter covers serialisation latches for EDM thread storage (latch 24) but also Buffer Manager (latch 56)
  - Reported in IFCID 001, QVLSLC24
  - Significant LC24 activity with no corresponding EDM pool activity suggests potential issue with zIIP prefetch

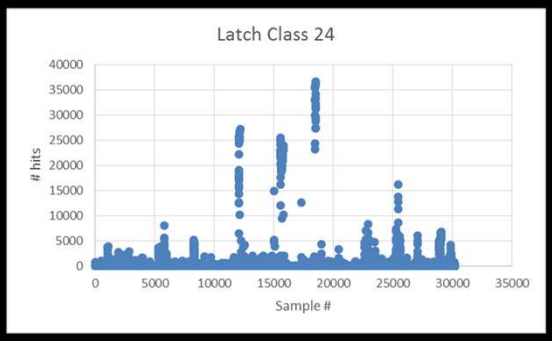

DB2 Accounting can tell us…

- DB2 accounting data shows actual and eligible zIIP offload for each plan / package
- Values are already normalised to CP speed, so no need for further adjustment
- SECP / Un-zIIP-ed figure originally removed in DB2 10 accounting data, but PM57206 re-introduces it
- Delays due to prefetch waiting for over-committed zIIP will appear as Class 3 “wait for other read” (QWACAWTR)

AVERAGE	APPL(CL.1)	TIMES/EVENTS	APPL(CL.1)
ELAPSED TIME	21.452714	ELAPSED TIME	6:37.60226
NONNESTED	21.452714	NONNESTED	6:37.60226
STORED PROC	0.000000	STORED PROC	0.000000
UDF	0.000000	UDF	0.000000
TRIGGER	0.000000	TRIGGER	0.000000
CP CPU TIME	7.217101	CP CPU TIME	2:41.61476
AGENT	7.172907	AGENT	2:41.37611
NONNESTED	7.172907	NONNESTED	2:41.37611
STORED PRC	0.000000	STOR	
UDF	0.000000	UDF	
TRIGGER	0.000000	TRIG	
PAR.TASKS	0.044194	PAR.TASKS	0.238051
SECP CPU	0.000000	SECP CPU	0.000000
SE CPU TIME	4.034602	SE CPU TIME	19.274049
NONNESTED	4.034602	NONNESTED	19.274049
STORED PROC	0.000000	STORED PROC	0.000000
UDF	0.000000	UDF	0.000000
TRIGGER	0.000000	TRI	
PAR.TASKS	0.000000	PAR.TASKS	0.000000
SUSPEND TIME	0.000000	SUSPEND TIME	0.000000
AGENT	N/A	AGENT	N/A
PAR.TASKS	N/A	PAR.TASKS	N/A
STORED PROC	0.000000	STORED PROC	0.000000
UDF	0.000000	UDF	0.000000

Un-zIIP-ed workload

zIIP-ed offload

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- How Does DB2 Exploit zIIPs?
- Performance Considerations
- zIIP Usage Measurement and Capacity Planning
- **Case Study**
- Summary

Case Study – Estimating the Benefit

- Large customer with significant DB2 workload, running on zEC10 hardware
- Single zIIP engine was configured on CEC
- SMF 30 chart shows approx. 120-170 MSU of CP time used by DBM1 during most of the day
  - Contributing to 4HRA
  - zIIP Busy line shows **expected** zIIP usage at V10: 100-160% busy = 2 fully loaded zIIPs!

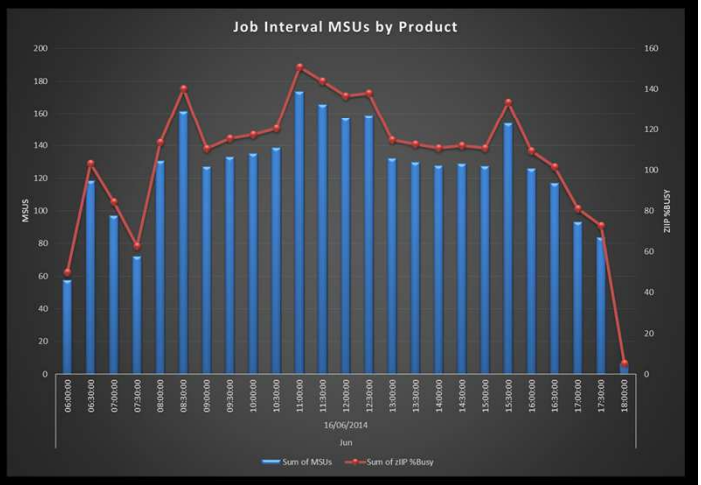

Case Study – Quantifying the Benefit

- Sub-Capacity Reporting Tool (SCRT) report shows which LPARs contribute to the DB2 MLC

==P5==========

PRODUCT MAX CONTRIBUTORS

Product N	Product ID	Highest	Date/Time	LPAR ABC1	LPAR ABC2	LPAR ABC3	LPAR XYZ1	LPAR XYZ4	LPAR XYZ8	LPAR XYZ9	LPAR DEV1
z/OS V1	5694-A01	977	07 Apr 2014 - 11:00 U	100	4	3	344	153	17	353	5
DB2 V9 fo	5635-DB2	949	07 Apr 2014 - 11:00 U	100	0	0	344	152	0	353	0
CICS TS fo	5655-M15	949	07 Apr 2014 - 11:00 U	100	0	0	344	152	0	353	0
MQSeries	5655-L82	849	07 Apr 2014 - 11:00 U	0	0	0	344	152	0	353	0
Tivoli Net	5697-ENV	970	07 Apr 2014 - 11:00 U	100	4	0	344	152	17	353	0
IBM Enter	5655-S71	316	16 Apr 2014 - 22:00 U	84	0	0	0	232	0	0	0

Case Study – Quantifying the Benefit

- Using SMF70, group the LPARs that contribute to the DB2 MLC to verify that DBM1 peaks correspond to DB2 MLC monthly peaks so we know we'll see savings
- The addition of 2 zIIPs on this machine would allow offload of DBM1 address space and reduce peak 4HRA by 120+ MSUs
- 120 MSUs less for software on those LPARs

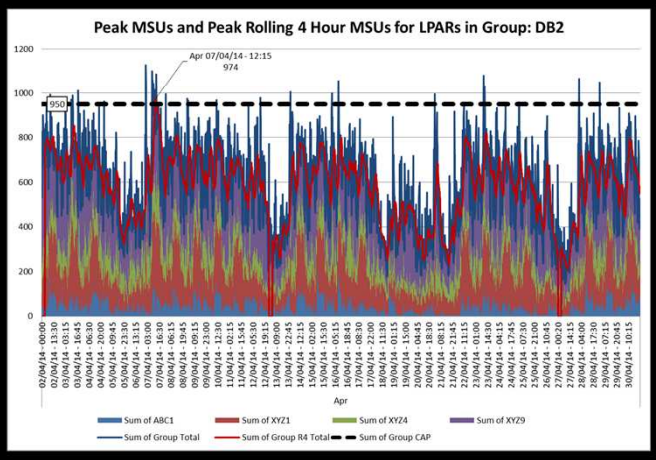

Case Study – The Payback

- Other CEC showed very similar workload characteristics, so two additional zIIPs were purchased for each CEC (total of 4 zIIPs) prior to DB2 10 upgrade
- DB2 10 showed the expected zIIP offload following upgrade a few weeks later
- Annual MLC savings estimated at approx. £800k per annum
- ROI of less than two months

Topics

- Introduction
- What Are zIIPs and How Can They Help?
- How Does DB2 Exploit zIIPs?
- Performance Considerations
- zIIP Usage Measurement and Capacity Planning
- Case Study
- **Summary**

Summary

- zIIPs are a predominantly a TCO management tool **but**
  - They can improve performance in a CPU-starved environment
  - They can degrade performance if they are over-utilised
- DB2 10 starts directing performance-critical workload to zIIPs for the first time
  - Impact of over-utilised zIIP becomes much bigger
  - Importance of monitoring is correspondingly higher
- You have plenty of instrumentation available to tell you exactly what's happening on your systems
- Work with your z/OS Sysprog / Capacity planner to see if you can save your company big money

IDUG
Leading the DB2 User Community since 1988

IDUG DB2 Tech Conference
Prague, Czech Republic | November 2014

#IDUG

Julian Stuhler
Triton Consulting
julian.stuhler@triton.co.uk

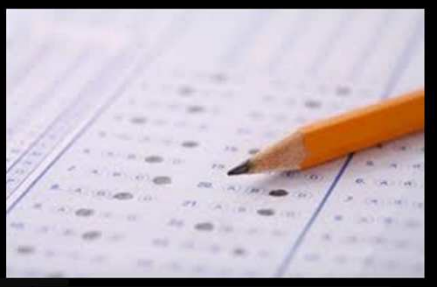

Session B06
Tuning zIIPs for Fun and Profit

Please fill out your session evaluation before leaving!