IBM Inside Sales

International Technical Support Organization Global Content Services

zTour – DB2 for z/OS zSynergy and Query Accelerator Support

www.ibm.com/redbooks

Paolo Bruni
pbruni@us.ibm.com

Redbooks Workshop

IBM Inside Sales

International Technical Support Organization Global Content Services

Notices

Redbooks Workshop

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo)®
eServer™
iSeries®
i5/OS®
pureXML™
z/Architecture®
z/OS®
zSeries®
z9™
AIX®
Collation®
CICS®
Database 2™
Distributed Relational
DB2 Connect™
DB2 Universal Database
DB2®
DFSMS™
DFSMSdss™
DFSMShsm™
DRDA®
DS8000™
Enterprise Storage Server®
ESCON®
FlashCopy®
FICON®
Geographically Dispersed Parallel Sysplex™

GDPS®
Informix®
IBM®
IMS™
Language Environment®
MQSeries®
MVS™
Net.Data®
OMEGAMON®
OS/390®
Parallel Sysplex®
POWER™
PR/SM™
QMF™
Redbooks®
RACF®
RETAIN®
REXX™
RMF™
System z™
System z9™
System Storage™
System/390®
Tivoli®
VTAM®
WebSphere®

The following terms are trademarks of other companies:
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Performance Disclaimer

This document contains performance information based on measurements done in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the network configuration, I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the numbers stated here.

Introduction

Contents

This workshop discusses DB2 10 for z/OS functions related to performance and synergy with System z platform, system and disk level, and it highlights the integration of DB2 with the Netezza appliance to provide the IBM DB2 Analytics Accelerator for z/OS solution.

Agenda

DB2 10 for z/OS Snapshot

- **Fastest uptake**
  - **+2x customers vs. V9**
  - **+2.5x licenses vs. V9**
  - **25% coming from DB2 V8**
- **Adoption driven by:**
  - **Performance improvements without application changes**
  - **Virtual Storage Constraint relief for more threads**
  - **Security, RAS improvements**
  - **Bitemporal data**

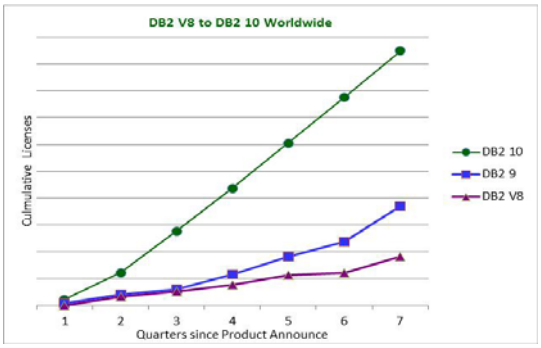

DB2 Deep Synergy with System z and DS8000

- Hardware-based Coupling Facility for efficient DB2 data sharing (availability and scale out)
- zIIP engines to offload certain DB2 workloads
- Hardware instructions for Unicode conversion
- Hardware data compression and encryption
- Cross-memory, memory protection keys
- Sorting
- z/OS Workload Manager (WLM)
- z/OS Security Server (RACF)
- 1 MB page size
- Volume-level and object-level FlashCopy
- Solid state disks (SSDs)
- z196 High Performance FICON (zHPF) with multi-track data transfer and FICON Express 8, prefetch, format write
- DS8000 6.2 enhancements

Top 10 in DB2 10 for z/OS

1. CPU reductions for transactions, queries, and batch
2. Ten times more users by avoiding memory constraints
3. More concurrency for catalog, utilities, and SQL
4. More online changes for data definition, utilities and subsystems
5. Improved security with more granularity
6. Bitemporal or versioned data – Time travel queries
7. SQL enhancements improve portability
8. pureXML performance and usability enhancements
9. Hash, index include columns, access path stability, skip migration, …
10. Productivity improved for database and systems administrators, and application programmers

Why migrate to DB2 10 for z/OS?

- Business needs to save money
  - **Reduce CPU time**
  - Service Oriented Architecture
- Application developers need improved productivity and integration
  - pureXML for a faster, more capable interface to XML data
  - Powerful new SQL temporal enhancements and portability
- Database Administrators need
  - Improved performance
  - Availability, scalability and memory management
  - Simpler security and regulatory compliance
  - More productive database administration

DB2 10 performance objective

- Historical goal under 5% performance regression
- Goal 5% -10% initial performance improvement
- Many customers reduce CPU time 10% - 20%

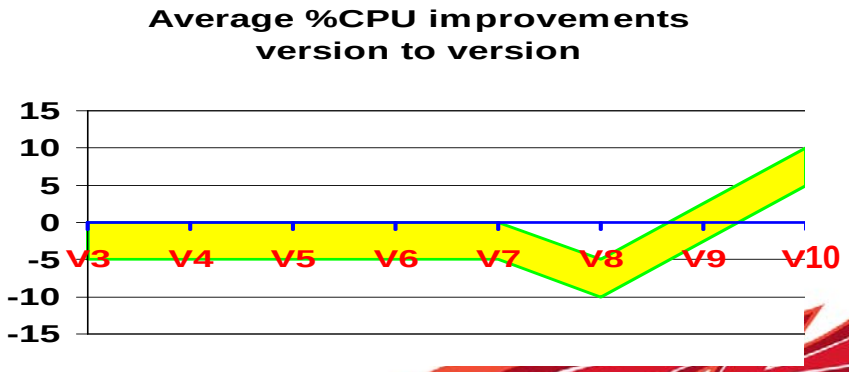

System level improvements

Key items that would improve performance without changing schema/applications

- Virtual storage constraint relief
- High performance DBATs
- Locking and latching contention reduction
- System z synergy
- Buffer pool enhancements
- Utilities, sort and zIIP support
- Optimizer enhancements (REBIND)

Performance enhancements requiring few changes (CM)

- SQL runtime improved efficiency
- Address space, memory changes to 64 bit, some REBINDs
- Faster single row retrievals via open / fetch / close chaining
- Distributed thread reuse High Performance DBATs
- DB2 9 utility enhancements in CM8
- Parallel index update at insert
- Workfile in-memory enhancements
- Index list prefetch
- Solid state disk use
- Buffer pool enhancements
  - Utilize z10 1 MB page frame size
  - "Fully in memory" option (ALTER BUFFERPOOL)

Performance enhancements requiring REBIND (CM)

- Most access path enhancements
- SQL paging performance enhancements
  - Single index access for complex OR predicates:
- IN list performance
  - Optimized Stage1 processing (single or multiple IN lists)
  - Matching index scan on multiple IN lists
- Query parallelism improvements
- More stage 2 predicates can be pushed down to stage 1
- More aggressive merge of views and table expressions
  - Avoid materialization of views
- REBIND enables further SQL runtime improvements
- If migrate from V8, get new RUNSTATS before mass rebind

Performance enhancements requiring NFM

- Efficient caching of dynamic SQL statements with literals
- Most utility enhancements
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- SQL Procedure Language performance improvements
- Workfile spanned records, PBG
- Insert improvement for UTS

Performance enhancements requiring NFM + DBA work

- Hash access path Alter + Reorg + rebind to activate
- Index include columns Alter + Rebuild + rebind to activate
- Inline LOBs Alter (need UTS and RRF)
- MEMBER CLUSTER for UTS
- DEFINE NO for LOB and XML columns
- Most security functions

DB2 10 application enablement and portability

- Data versioning by date (bitemporal support)
- pureXML enhancements
- Large object improvements
  - Allow non-NULL default values for inline LOBs
  - Loading and unloading tables with LOBs
    - LOBs in input/output files with other non-LOB data
- Improved portability and SQL consistency
  - Currently committed locking semantics
  - Implicit casting or loose typing
  - Timestamp with time zone
  - Variable timestamp precision – seconds to picoseconds
  - OLAP Moving Sum, Moving Average

DBM1 Virtual Storage Constraint Relief

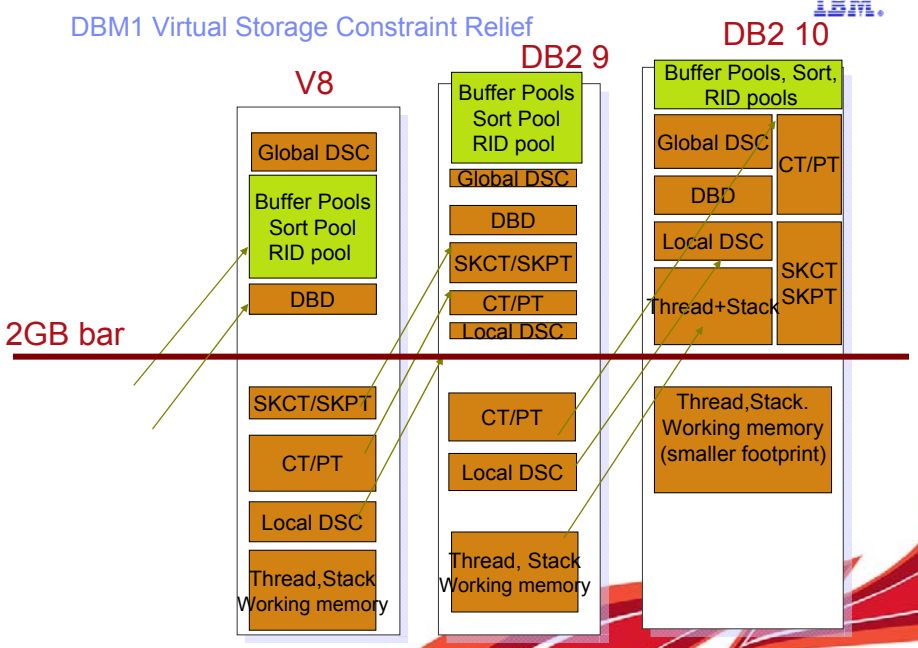

Relief in DBM1 Below The Bar

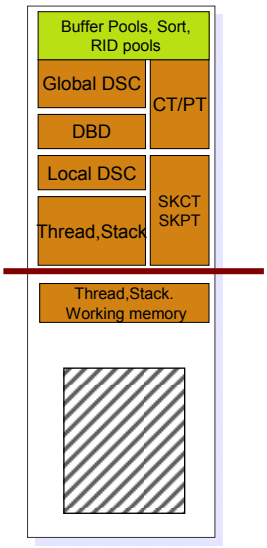

DB2 10

- DBM1 below 2GB
  - EDM storage - All above
  - Thread + Stack - 70-90% less usage in DB2 10 compared to DB2 9
    - xPROC (SPROC, IPROC, UPROC, etc) loaded in below the 2GB bar
    - Built in BIND time, shared at runtime
- More number of threads
- Reduce CPU time at the expense of storage
  - More thread reuse to avoid allocate/deallocate
  - Wider usage for bind option RELEASE(DEALLOCATE)
  - High Performance DBATs
  - Larger MAXKEEPD values for KEEPDYNAMIC=YES allow users to keep more short prepared statements

Exploitation of System z hardware

- **z10 and zEnterprise 196: Memory prefetch instruction**
- **Large page frame size (1 MB page frame) for buffer pools**
  - IEASYSXX LFAREA=(xx%| xxM | xxG| xxT) and RE-IPL
    - Backed by 256 contiguous 4K real frames and not page-able.
  - 1MB page is fixed, 64bit private pool
  - Buffer pools with long term page fix (PGFIX=YES)
    - If 1 MB page frames are available, DB2 will request 1MB first
- **Long Term Page Fix from DB2 V8 to reduce CPU cost for I/O operations even without 1 MB**
  - 1MB page frames to reduce CPU cost during get pages and release pages
    - Buffer pools with large variations of getpage activities
  - Observed 1 to 4% CPU reduction in workload level by using 1MB page frames with PGFIX(YES)
- **Larger processor cache (1.5 MB L2 per core, 24 MB L3 per chip, 129 MB L4)**
- **DB2 9 OLTP, Insert, Utility and query workloads observing 20% to 40% DB2 CPU reduction compared to z10 processors (faster processors and improved microcode).**
- **Typically 25 to 35 % and higher DB2 CPU reduction can be achieved as #of processors per LPAR increases → Best fit with DB2 10 scalability.**

zIIP enhancements

- Portion of RUNSTATS utility (Class 1 CPU)
  - Redirection rate varies depending on options used
- Parsing process of XML Schema validation (Class 2 CPU)
  - 100% of new validation parser is eligible
    - Can be zIIP, zAAP, or zAAP on zIIP
    - Retrofit into DB2 9 via PK90032 (preconditioning), PK90040 (enabling)
- Portion of DBM1 processes (asynchronous I/O)
  - Prefetch I/Os (DBM1 SRB)
  - Deferred write I/Os (DBM1 SRB)

Buffer pool related enhancements

- Large real and DB2 managed in memory buffer pool
  - z196 supports up to 3TB memory
  - PGSTEAL = NONE
    - Pre-load the data at the first open or at ALTER BPOOL
    - Avoid unnecessary prefetch request (similar to VPSEQT=0)
    - Avoid LRU maintenance -> no LRU latch
- Buffer pools allocation as needed
  - No more penalty for BP over-sizing
    - In DB2 9, an entire buffer pool is allocated when first used
    - If a defined size is bigger than actual used size and using PGFIX=YES
- Table space buffer pools are no longer allocated when index-only access
- Avoid exhaustive local BP scan
  - During p-lock negotiations in registering page or validity checking
  - Possible DB2 9 bypass
    - Increase PCLOSET/PCLOSEN
    - More and smaller buffer pools

Focus on UTS

- Wider adaption of Universal Table Spaces
  - Catalog Tables
  - New Functions: Hash Access, Inline LOB, Access Currently Committed
  - Alter Table Space types
- Segmented and UTS space search and CPU reduction
- Segmented and UTS space map page latch reduction
  - Reduce the time held for space map latch
- UTS with MEMBER CLUSTER option (NFM)
  - MEMBER CLUSTER does not maintain the clustering -> quicker space search in insert and remove the hot space map/data page in concurrent insert in data sharing
  - Same consequence during query due to loss of clustering

DDF enhancements

- **High Performance DBATs**
- DDF availability and monitoring functions
- DB2 supplied stored procedures like SYSPROC.ADMIN_INFO_SQL
- DDF internal optimization
  - Optimizes communication between DDF and DBM1
  - Bundle OPEN/FETCH../CLOSE as one network traffic
    - Just like SELECT statement
- Optimized special register and inactive thread processing
  - Improves performance of distributed application that set special register frequently
- REOPT(ONCE) CPU reduction
  - Reduced CPU cost of processing REOPT(ONCE) in distributed access
- UNICODEMGR support
  - Use Unicode for DRDA metadata to avoid EBCDIC->Unicode conversion

DB2 and z/OS version summary for migration

Version	PID	Generally available	OS prereq	Marketing withdrawal	End of service
V3	5685-DB2	December 1993	MVS V4R3	February 2000	January 2001
V4	5695-DB2	November 1995	MVS V4R3	December 2000	December 2001
V5	5655-DB2	June 1997	MVS V4R3	December 2001	December 2002
V6	5645-DB2	June 1999	OS/390 V1R3	June 2002	June 2005
V7	5675-DB2	March 2001	OS/390 V2R7	March 2007	June 2008
V8	5625-DB2	March 2004	z/OS V1R3	September 2009	April 2012
V9	5635-DB2	March 2007	z/OS V1R7	December 2012	June 2014
V10	5605-DB2	October 2010	z/OS V1R10	TBD	TBD
V11	5615-DB2	ESP March 2013	z/OS V1R13		

DB2 9 to 10 migration CPU reduction
IBM benchmark workloads

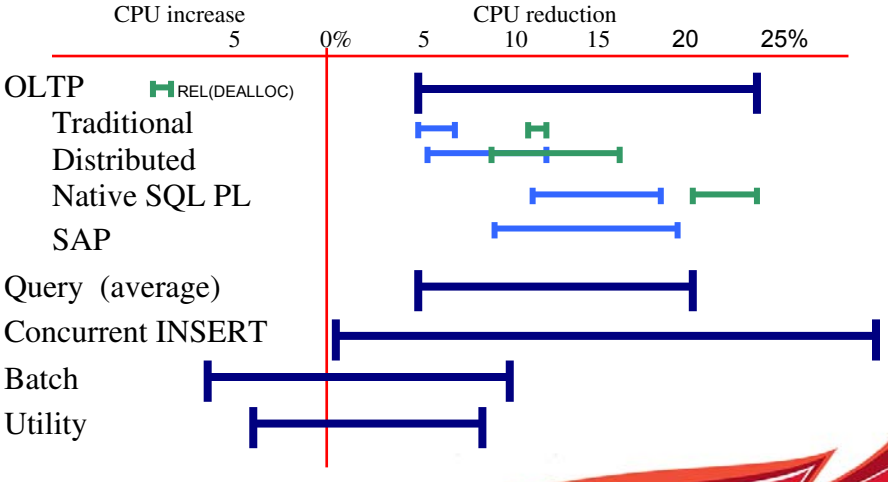

DB2 10 business value summary

- Business analytics
- Time travel queries with temporal (versioned) data support
- Integrated XML / relational support
- Compliance security enhancements
- Performance
- Manageability
- Extended distributed computing performance
- Virtual storage management.

IBM Inside Sales International Technical Support Organization Global Content Services

Information Management for System z

Continual investment and growth

2012
DB2 Analytics Accelerator
IMS 12
Active Active for z/OS
Industry Models for System z
Cognos 8 BI for z/OS
Smart Analytics Optimizer for z/OS
OMEGAMON DB2 end–to-end monitoring
2011
DB2 10 for z/OS
DB2 Sort for z/OS
Guardium for System z
SPSS for System z
ISAS 9600 for System z
IMS Tools Solution Packs for z/OS
Analytics
IMS 11 for z/OS
2010
InfoSphere Discovery for z/OS
InfoSphere Warehouse for System z
Information-Led Transformation
Cognos 8 BI for System z
InfoSphere MDM Server for System z
Business Intelligence
Optim Data Governance for System z
2009
Information Server for System z
Information Agenda
2008
Information On Demand
Data Warehousing

Redbooks Workshop

QMF V10 – *Enterprise Business Analytics for System z*

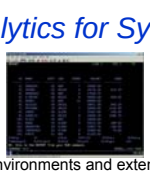

- Dramatically different than prior releases
  - ✓ We keep innovations coming
  - ✓ Existing customers can modernize and consolidate their BA environments and extend their current QMF investment at low cost
  - ✓ Workstation and WebSphere environments let you create and deploy visually rich reports, interactive dashboards, and much more
  - ✓ TCO is based on an enterprise-wide model. The greater the number of users and databases accessed, the lower the cost
- IBM DB2 Analytics Accelerator creates significant performance improvements and savings for QMF
  - ✓ Faster responses to queries
  - ✓ Act faster on results, dig deeper, generate personalized reports and dashboards
  - ✓ Make more timely, informed decisions
  - ✓ Provides a lower cost option for massive batch operations, one of QMF's known strengths
- What's new?
  - – ***New* QMF redbook**: *Complete Analytics with IBM DB2 Query Management Facility*
    http://www.redbooks.ibm.com/redbooks/pdfs/sg248012.pdf
  - – Join our next beta program!

    Ask your IBM representative to visit the following web page for information:

    https://w3-01.sso.ibm.com/software/support/trial/earlyprograms/form/nomination.wss?id=4233

Some not well known references

- Which JDBC driver (JCC) version corresponds with which DB2 release and Fix Pack level:

  https://www-304.ibm.com/support/docview.wss?uid=swg21363866
- The Cross-Platform SQL Reference Version 4.1

  http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
  - (DB2 for: z/OS V10, i V7.1, LUW V9.7)

Recent Redbooks publications

- **Data Studio and DB2 for z/OS Stored Procedures, REDP-4717**
- **Data Administration Tools with DB2 10 for z/OS, SG24-7916**
- **DB2 10 for z/OS: Configuring SSL for Secure Client-Server Communications, REDP-4799**
- **Extremely pureXML in DB2 10 for z/OS, SG24-7915**
- **DB2 10 for z/OS Technical Overview**
- **DB2 10 for z/OS Performance Topics, SG24-7942,**
- **Security Functions with DB2 10 for z/OS, SG24-7959**
- **Delivered in 2012**
  - Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS, SG24-8005
  - Complete Analytics with IBM DB2 Query Management Facility: Accelerating Well-Informed Decisions Across the Enterprise
  - DB2 for z/OS and List Prefetch Optimizer, REDP-4862
  - Streamline Business with Consolidation and Conversion to DB2 for z/OS, SG24-8044
  - IMS 12: The IMS Catalog, REDP-4812
- **Coming in 2012**
  - Tools for DB2 Utilities management
  - IMS Performance
  - WebSphere and DB2

DB2 10 Resources and Contacts

- International DB2 User Group
- DB2 for z/OS Exchange Forum
- World of DB2 for z/OS
- DB2 for z/OS group

- **DB2 Product Library** http://www.ibm.com/vrm/newsletter/11065
- **DB2 Best Practices** https://www.ibm.com/developerworks/data/bestpractices/db2zos/
- **Website** http://www.ibm.com/software/data/db2/zos/db2-10/
  - Case Studies, Customer statements, Demos, Brochures
- **Whitepapers**
  - Business Value Whitepaper – Julian Stuhler, Triton Consulting: "DB2 10 for z/OS: A Smarter Database for a Smarter Planet" *http://public.dhe.ibm.com/software/data/sw-library/db2/analystreports/tritonconsulting-db210forzos-smarterdatabase.pdf*
  - A Matter of Time: Temporal Data Management *http://public.dhe.ibm.com/software/data/sw-library/db2/papers/A_Matter_of_Time_-_DB2_zOS_Temporal_Tables_-_White_Paper_v1.4.1.pdf*
  - The Value of IBM's DB2 Utilities and Tools in 2011 – Phil Grainger http://public.dhe.ibm.com/software/data/sw-library/db2tools/value-of-IBM-DB2-Tools-Utilities.pdf
  - OMEGAMON Extended Insight Analysis: Where is your application spending its time? ftp://public.dhe.ibm.com/software/data/sw-library/db2imstools/extended_insight.pdf

(cont) DB2 10 Resources and Contacts

(cont) Whitepapers

- **Why DB2 for z/OS is Better than Oracle RAC** https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US&source=sw-infomgt&S_PKG=db2z-better-thank-oracle-rac-wp
- **zJournal article by Willy Favero** http://www.mainframezone.com/z-journal
- **SAP article on DB2 10 *(published by SAP)*** http://www.sdn.sap.com/irj/sdn/db2
- **SAP Best Practice Guide for Migrating to DB2 10 for z/OS *(published by SAP)*** https://websmp207.sap-ag.de/~sapidb/011000358700001414122010E
- **(Updated) Business Continuity Guide for Running SAP on System z – based on DB2 10 for z/OS, DB2 Connect 9.7 FP3a, SAP NetWeaver 7.10 and Tivoli Automation for z/OS V3.3** http://publibfp.dhe.ibm.com/epubs/pdf/iapacs03.pdf
- **DB2 10 for z/OS with SAP on IBM System z Performance Report – new techdocs white paper** http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101845
- **DB2 10 for z/OS – Optimized for SAP –** http://cattail.boulder.ibm.com/cattail/?source=s#view=andreas.r.mueller@de.ibm.com/files/3198290001883DDBA202FBE4093F23B6
- **SAP on DB2 10 for z/OS - Being More Productive, Reducing Costs and Improving Performance –** http://www.sdn.sap.com/irj/sdn/db2?rid=/library/uuid/005c6b33-aaf0-2d10-fcbb-b42e89ac5791
- **OMEGAMON Extended Insight Analysis: Where is your application spending its time?** ftp://public.dhe.ibm.com/software/data/sw-library/db2imstools/extended_insight.pdf
- **The Value of IBM's DB2 Utilities and Tools in 2011** http://public.dhe.ibm.com/software/data/sw-library/db2tools/value-of-IBM-DB2-Tools-Utilities.pdf

DB2 for z/OS Newsletter http://www.ibm.com/vrm/newsletter/11065
DB2 for z/OS e-Kit http://bit.ly/DB210e-Kit

DB2 for z/OS Migration Planning Workshops

Certifications so far...

- **1Q2011: SAP NetWeaver 7.30 and SAP R/3 4.6 is certified for DB2 10**
- **2Q2011: All SAP products based on NetWeaver 7.00/7.01 is certified for DB2 10**
- **2Q2011: PeopleSoft PeopleTools 8.50 and 8.51 is certified for DB2 10**
- **4Q2011: DB2 10 for z/OS certified for SAP NetWeaver 7.02 based products**
- **4Q2011: BusinessObjects XI 3.1 has been certified for DB2 10. The certification for BusinessObjects BI 4.x will follow**
- **1Q2012: Java upgrades starting from SAP Basis Release 6.40**

NEW

DB2 for z/OS Best Practices

www.ibm.com/developerworks/data/bestpractices/db2zos/

- Watch recorded presentations and read articles from your favorite DB2 experts.
- Get practical recommendations for areas such as:
  - Virtual storage tuning
  - Security
  - Software maintenance strategies
  - Migration
  - Setting up application servers to access DB2 for z/OS
  - Partitioning table spaces
  - Debugging stored procedures
  - And more!

Suggestions to db2zinfo@us.ibm.com

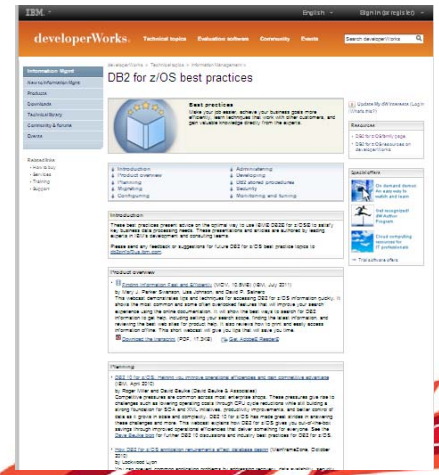

Subsystem and zSynergy

- Catalog restructure
- Latching contention relief
- Dynamic prefetch enhancements
- Buffer pool enhancements
- Workfile enhancements
- Logging enhancements
- I/O parallelism for index updates
- Space search improvement
- Log record sequence number spin avoidance
- Compression on insert
- z196 synergy
- 1 MB page frame support
- zIIP usage with DB2 10
- Open and close data set
- DASD enhancements
- SMF compression
- What's new

IBM Inside Sales International Technical Support Organization Global Content Services

Catalog restructure

© 2012 IBM Corporation

Redbooks Workshop

DB2 catalog restructure

- SMS managed
  - DSSIZE (>4GB) requires DFSMS Extended Addressability
- More LOBs
  - Several inline LOBs
- Links disabled in CM but really removed only in NFM
  - DSN1CHKR is no longer required
  - Dynamic SQL full PREPARE increase in class 2 CPU and elapsed time from 20% to 30%
  - Single DDL: CM9 no change, NFM slight less elapsed, up 10 to 40% CPU class 2
  - DROP DATABASE; CM9 slight less elapsed, up 10 to 30% CPU class 2. NFM less elapsed and more CPU
  - 5 concurrent DDL streams: almost linear for simple DDL streams with frequent COMMITs
  - Specify frequent COMMITs for complex DDL streams.
  - Serialize large CREATE and DROP database
  - Serialize aggressive CREATE of LOBs and XML

Catalog evolution

DB2 Version	Table spaces	Tables	Indexes	LOBs	Columns	Table check constraints
V1	11	25	27	0	269	N/A
V3	11	43	44	0	584	N/A
V5	12	54	62	0	731	46
V6	15	65	93	0	967	59
V7	20	84	118	2	1212	105
V8	22	85	132	2	1265	105
DB2 9	28	104	165	3	1643	119
DB2 10	95 (104-9)	134	233	18	1922	119

Single table table spaces
UTS PBG MAXPART 1
More LOBs, some inline
78 more indexes

Impact for migration

- DDL performance and concurrency
  - Indexes replace catalog links (active in NFM)
  - Increase in concurrency
  - Higher PREPARE time
  - Higher DDL time
  - Increased concurrency (NFM)
- BIND and REBIND stability
  - New table SYSPACKCOPY allows comparisons without SWITCH
  - Bind option APRETAINDUP to keep (default) or discard duplicates
- BIND performance
  - PLANMGMT=EXTENDED is the new default, impacts on CPU and elapsed
  - For data sharing, PM24721 uses space map cache and appends records in the DBD.
- Compression and inline LOBs
  - LOBs are more frequently used, some are inline
  - SPT01 uses LOBs, no compression until inlined (by PM27073/PM27811)

Latching contention relief

Performance Scalability - DB2 Latches (CM)

- Faster process on latch suspend/resume
- Most of known DB2 latches are addressed in DB2 10
  - LC12 : Global Transaction ID serialization
  - LC14 : Buffer Manager serialization
  - LC19 : Log write in both data sharing and non data sharing
  - LC24 : EDM thread storage serialization (Latch 24)
  - LC24 : Buffer Manager serialization (Latch 56)
  - LC25 : EDM hash serialization
  - LC27 : WLM serialization latch for stored proc/UDF
  - LC32 : Storage Manager serialization
- Internal contention relief
  - IRLM : IRLM hash contention
  - CML : z/OS Cross Memory Local suspend lock
  - UTSERIAL : Utility serialization lock for SYSLGRNG (*NFM) removed
  - Concurrent RE/BIND and most of DDL (*NFM)

Sample application with high frequency package allocation has up to 76% increased throughput in DB2 10

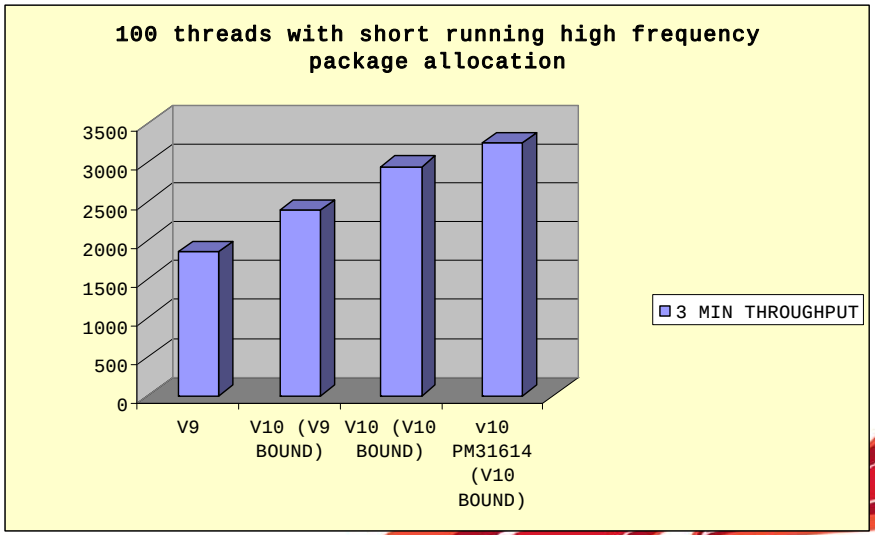

Elimination of UTSERIAL locks

- Each DB2 utility job stores its run time information in the DSNDB01.SYSUTILX directory table space. The UTSERIAL lock is an IRLM lock obtained by each utility when it wants to read or write from SYSUTILX. The granularity of this lock is such that it can cause contention when several concurrent utilities are run.
- DB2 10 NFM eliminates the UTSERIAL lock and merges tables SYSUTILX and SYSUTIL into one table called SYSUTILX.DB2 10 uses a new type of IRLM lock to serialize utility access to DSNDB01.SYSUTILX
  - This lock is a commit-duration lock that is used exclusively by DB2 utilities during utility compatibility checking, with a resource name (dbid.psid) for the database object that is targeted by the utility.
- These object level locks show up in lock traces as uobj.
- Measurement
  - 20 LOAD jobs concurrently on 2 members data sharing, both on DB2 9 and DB2 10
  - 20 table spaces each has 300 partitions and 6 indexes
  - IRLMRWT and UTIMOUT to smallest value
- On DB2 9 only 7 LOAD jobs completed
  - 13 LOAD abended RC00E40085 due to UTSERIAL lock contention
- On DB2 10 all 20 LOAD completed
  - No UTSERIAL locks
  - No lock contention

IBM Inside Sales International Technical Support Organization Global Content Services

Dynamic prefetch enhancements

© 2012 IBM Corporation

Redbooks Workshop

Dynamic Prefetch Enhancements

- The following 3 types of DB2 for z/OS prefetch are available:
  - Sequential prefetch: used for table scans. As soon as the target table space is accessed, two prefetch quantities of sequential pages are read into the buffer pool. .
  - List prefetch: used during query execution, RIDs for qualifying rows are obtained from one or more indexes. In most cases, DB2 sorts the RIDs, then issues asynchronous multi-page read requests against the base table space.
  - Dynamic prefetch: used at query execution. DB2 determines that the pattern of page access for a target table or an index is sequential enough to justify the activation of prefetch processing. It can issue two parallel prefetch I/Os.
- DB2 10 (CM no BIND) enhances dynamic prefetch processing as follows:
  - Index scans using list prefetch
  - Row level sequential detection (RLSD)
  - Progressive prefetch quantity

Disorganized Index Scans using List Prefetch

DB2's sequential detection algorithm detects whether index leaf pages are physically organized well enough to kick off dynamic prefetch

- If the index leaf pages are not well organized then DB2 performs synchronous I/Os
- Therefore a big incentive with DB2 9 to keep indexes organized

DB2 10 greatly mitigates any performance issues of a disorganized index by using list prefetch I/O for the leaf pages

- DB2 10 uses the N-1 level of the index to locate the leaf pages, whereas DB2 9 did not issue getpages for the non-leaf pages when doing an index scan

Disorganized Index scans using List Prefetch…

- Index scan can begin with an index probe to position someplace in a leaf page, then scans forward (or backwards) using a combination of dynamic and list prefetch
- Index manager switches from dynamic prefetch to list prefetch for a disorganized index
  - Index scan can be 2 to 6 times faster, CPU time is less too
  - Applies to index-to-data access and index-only access
- List prefetch on index leaf pages can greatly reduce the synchronous I/O waits for long running queries accessing disorganized indexes
- Utilities including REORG INDEX, CHECK INDEX, and RUNSTATS, can also benefit
  - Partition-level LOAD and REORG benefit also.
    - Prefetch and list prefetch are also used for updating the NPI which happens to perform an index scan
- There may be less of a need to run REORG, because your application performance might be able to tolerate accessing data that is more disorganized than previous versions of DB2.

Disorganized index scan (query time with index only access)

- PGFIX(YES)
- No data access
- SRB time zIIP eligible

DB2 10, 3 times faster with cold cache
6 times faster with hot cache

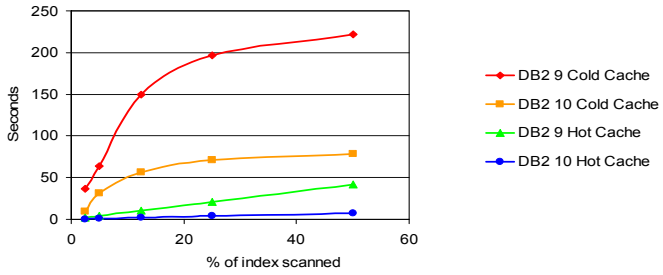

Row Level Sequential Detection (RLSD)

Page level sequential detection is used by dynamic prefetch in DB2 9

- When the next page is accessed, DB2 checks that it is within half the prefetch quantity forward or backwards of the previous page.
  - 16 pages if the prefetch quantity is 32 pages (not changed in DB2 10).
- If it is, then that page is noted as being page-sequential relative to the previous one. Otherwise it is not page-sequential.
- DB2 keeps a running counter that is the number of page-sequential getpages among the last eight getpages.
- If five out of the last eight getpages are page-sequential, a getpage can trigger prefetch, but only if the following conditions are true:
  - The getpage is in the *prefetch window*.
  - The counter is greater than 4.
  - The getpage is page-sequential.

Row Level Sequential Detection…

- DB2 remembers the previous two page numbers, not just the previous page number
- Rows are counted, not just getpages
- With RLSD, it is easy to trigger prefetch
- RLSD only applies to table spaces

Progressive prefetch quantity

- Since DB2 10 introduces row level sequential detection and thus triggers prefetch more quickly (after a minimum of 5 rows, as opposed to 5 pages in DB2 9), dynamic prefetch in DB2 10 also uses a progressive prefetch quantity.
- For example, with 4 KB pages the first prefetch I/O reads 8 pages, then 16 pages, then all subsequent I/Os will prefetch 32 pages (as in DB2 9).
- While row level sequential detection applies only to data, a progressive dynamic prefetch quantity is used for both index and data.
- Both list prefetch for index scan and row level sequential detection may relieve the need for REORG utilities as it provides an improved tolerance for disorganized indexes and data and therefore reduce the frequency of reorganization.

Measurements

- Row Level Sequential Detection
  - Read 10% of the 20 M rows (clustered and no free space) in key sequential order after successive inserts as shown in the next two foils.
  - For Row size = 49 bytes, page size = 4K (81 rows per page)
    - Once the cluster ratio reached 94%, DB2 9 stopped using dynamic prefetch altogether, but DB2 10 never stopped using dynamic prefetch. When the cluster ratio was 92%, DB2 10 reduced the elapsed time by about 45%
  - For Row size= 80 bytes, page size = 4K (40 rows per page)
    - With fewer (and larger) rows per page, dynamic prefetch performed better in DB2 9, because the number of dynamic prefetch did not begin to decrease until the cluster ratio fell to 96%.

Row size 49 bytes

Test case	Cluster ratio	Cardinality	NPAGES
1	100%	20,000,000	253167
2	98%	20,200,000	256024
3	96%	20,400,000	258882
4	94%	20,600,000	261740
5	92%	20,800,000	264598

Read 10% of the rows in key sequential order

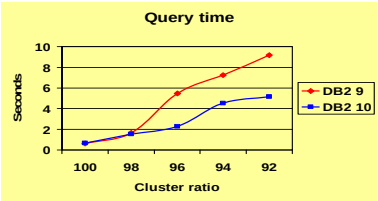

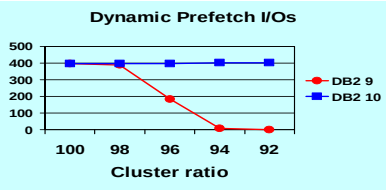

Row level sequential detection (RLSD) preserves good sequential performance for the clustered pages

Row size = 80 bytes, page size = 4K (40 rows per page)

Test case	Cluster ratio	Cardinality	NPAGES
1	100%	20,000,000	500,000
2	96%	20,400,000	520,000
3	92%	20,800,000	540,000
4	88%	21,200,000	560,000
5	84%	22,400,000	580,000

Read 10% of the rows in key sequential order

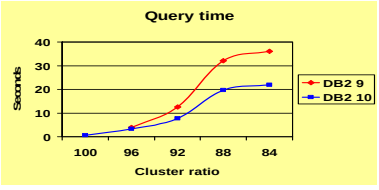

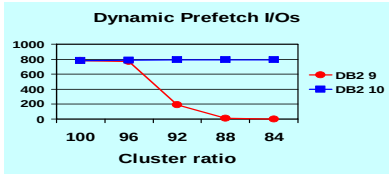

Row level sequential detection (RLSD) still preserves good sequential performance for the clustered pages but not as good because of fewer rows

Buffer pool enhancements

Buffer Manager enhancements

- Buffer pools are allocated on demand
  - In DB2 10, the buffer pool is allocated in chunks on demand
    - Other service task suspension, and DBM1 SRB time when buffer pool is growing
  - The buffer pool is deleted when the use count drops to zero
- Avoid OPEN pageset lock/unlock of table space in index-only access
- Buffer pool scan avoidance (next foil)
- 1 MB page frames
  - Applies to PGFIX(YES) if LFAREA specified in IEASYSxx parmlib member
  - Tables and indexes
- In-memory tables and indexes
  - Same performance as VPSEQT(0)
    - Avoid cost of LRU management and avoid cost of scheduling prefetch engine
  - PGSTEAL(NONE)
  - Objects loaded into buffer pool at Open time
  - Unpredictable results if the buffer pool is too small

Buffer pool scan avoidance

- Prior to DB2 10, DB2 members needed to execute complete scans of the buffer pool during certain transitions in inter-DB2 Read/Write (RW) interest for a page set/partition (p/p). For example;
  - When a p/p transitions from P-lock S state (inter-DB2 Read Only (RO)) to P-lock IS state (another member has declared RW interest for this p/p)
  - Then each member downgrading the P-lock on the p/p from S to IS must scan the buffer pool which this p/p is using to ensure that all the locally cached pages for this p/p are properly registered to the GBP for cross invalidation (XI).
  - The same is true for SIX to IX lock transition (this member used to be the sole updater for this p/p, but now another member had declared RW interest).
- DB2 10 eliminates all of these sequential buffer pool scans
- And in turn reduces LC14 contention

PGSTEAL(NONE) usage recommendations

- Use for small objects that incur a lot of synch I/Os
  - Consider the synch I/Os that occur after DB2 startup (if you care about startup performance) and during utility operations
- Especially good for indexes which are often randomly accessed (and they are smaller than table spaces)
  - Particularly good for aux indexes since aux indexes are almost always randomly accessed

Workfile enhancements

In-memory workfile enhancement (CM)

- When DB2 uses in memory workfile
  - Last sort in a top query block with ORDER BY or GROUP BY clause and the size of the sort records is less than 1 million bytes
    - SELECT c1, c2 FROM T1 ORDER BY c1;
    - SELECT MAX(c1), c2 FROM T1 GROUP BY c2;
  - Sort for join operations and the sort record length < 1000 bytes, the number of rows < 255 and the size of the sort records < 32K
- Performance Result
  - Up to 50% CPU time improvement for last sort in a top query block
  - Up to 2% CPU time improvement for sort for join operations
  - Improvement will be more noticeable in a multi threaded application due to the reduction of spacemap page contentions
- How to detect the in memory workfile usage
  - # of getpage reductions for workfile bufferpool
  - New IFCID 2 counters

Other Sort enhancements

- Increased the default sort pool storage to 10 MB
- Implemented a hash technique for GROUP BY queries.
  - Up to 18% CPU time decrease for individual query in internal query workload
  - 12 out of 87 queries had > 3% CPU time decrease
  - 10 out of 87 queries had > 3% CPU time increase, but they are short running queries.
  - IFCIDs 95 and 96 record the sort technique (RCYC and ESAG)
- Implemented a hash technique for sparse index to help improve probing.
- Removed padding of variable length data fields during the input phase.
  - Up to 1.5% CPU time decrease
- Improved FETCH FIRST n ROWS processing if sort record size < 128 KB for both ORDER BY and GROUP BY and when the sort record size > 128K, only the n rows to be fetched are written to the workfiles.

Workfile enhancements

- II14587 clarifies functions and maintenance for DB2 9 and 10
- DB2 9 introduced and favored the 32 KB pages and mixed workfiles with DGTT.
- IFCID 002 fields QISTWFP1/QISTWFP2 can be used to monitor and ZPARM MAXTEMPS controls how much space in the workfile database an agent can use
- PK70060 favors DB2-managed (STOGROUP) workfile table spaces with SECQTY 0 or user-managed table spaces (regardless of their secondary allocation) for non-DGTT work and DB2-managed table spaces with SECQTY > 0 or -1 (or omitted) for DGTT
- PM02528 introduces DSNZPARM WFDBSEP
- DB2 10 allows workfile records to be spanned
  - Workfile record length can now be up to 65529 bytes by allowing the record to span multiple pages.
  - The *maximum* limit for sort key length for sort is also increased from 16000 to 32000 bytes
- PM17336 and PM65767 add and fix the support of partition-by-growth table spaces in WORKFILE database. They require WFDBSEP YES.
  - Available in DB2 10 new-function mode

PBG Workfile Table Space

- Why PBG for workfile database (NFM)
  - Customer is able to control the physical storage growth by specifying MAXPARTITIONS and DSSIZE in the DDL.
  - To enable > 64G DGTT. (DGTT cannot span multiple workfile table spaces.)
- Workfile Type Preference
  - DB2 will prefer PBG for DGTTs
  - DB2 will prefer segmented for sort
  - If the preferred table space is not available, then the alternative table space will be picked
- DSNZPARM WFDBSEP
  - YES: DB2 always directs DGTT work only to DB2-managed (STOGROUP) workfile table spaces defined with a non-zero SECQTY and workfile work only to other workfile table spaces (DB2-managed table spaces defined with a zero SECQTY or user-managed table spaces)
  - If no table space with the preferred allocation type is available, DB2 issues an error (message DSNT501I or SQLCODE -904).
  - NO is the default: DB2 still attempts the above preferences. If no table space with the preferred allocation type is available, DB2 selects a table space with a non-preferred allocation type.: applies only to DB2 managed space
- Both performance and physical storage are comparable for concurrent DGTT (create/insert/select/delete) between PBG and segmented table space

Considerations

- Usage by DB2

  Code has been changed in the workfile table space selection algorithm to select PBG workfile table spaces in the following manner:

  1. DGTT, when the zparm WFDBSEP value is YES, will select PBG as the top choice (regardless of its SECQTY, so even if the PBG has a zero SECQTY). This is already followed today for the case of DGTT with WFDBSEP=NO.
  2. Non-DGTT (e.g. sort workfiles), when the zparm WFDBSEP value is NO, will select PBG as the last resort (regardless of its SECQTY, so even if the PBG has a zero SECQTY).
  3. Non-DGTT (e.g. sort workfiles), when the zparm WFDBSEP value is YES, will never select PBG (regardless of its SECQTY, so even if the PBG has a zero SECQTY).

- Create PBG table spaces for DGTTs. Only if they are not available, DB2 will look for segmented table spaces with non-zero SECQTY.
- Create segmented table space with 0 SECQTY for sort.
- Set WFDBSEP = YES for DB2 managed or user managed table spaces to segregate DGTTs and sort workfile usage but you must define segmented for use by sort if you have PBG for DGTT.

Logging enhancements

DB2 10 Logging Changes

- LC19 spin avoidance
  - DB2 10 holds the log latch for the minimum amount of time possible
  - Further LC19 spin avoidance enhancements
- Long term page fix log buffers
  - Check that OUTBUFF is reasonably sized
- Log I/O Enhancements
  - Overlapping commits and chained writes
  - New DASD technology allows DB2 to rewrite pages to the log1 and log2 in parallel rather than sequentially

DB2 Log Performance, sync write, **few** pages per commit

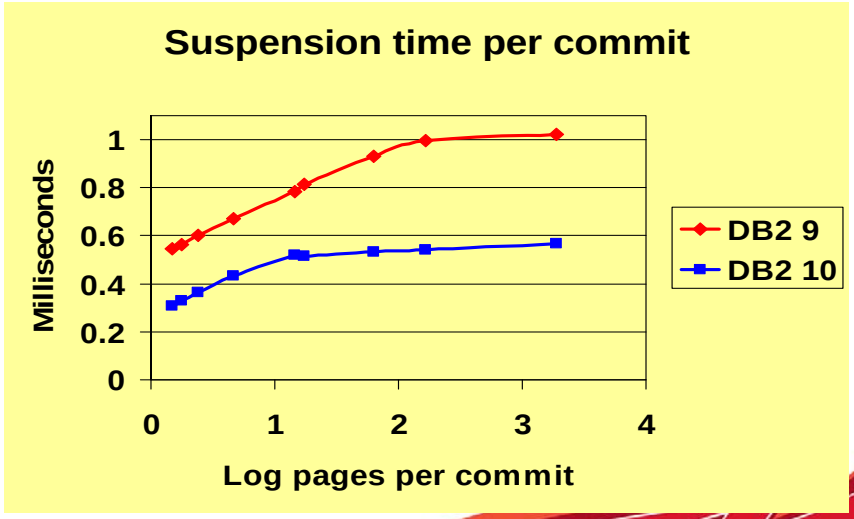

DB2 Log Performance, sync write, **many** pages per commit

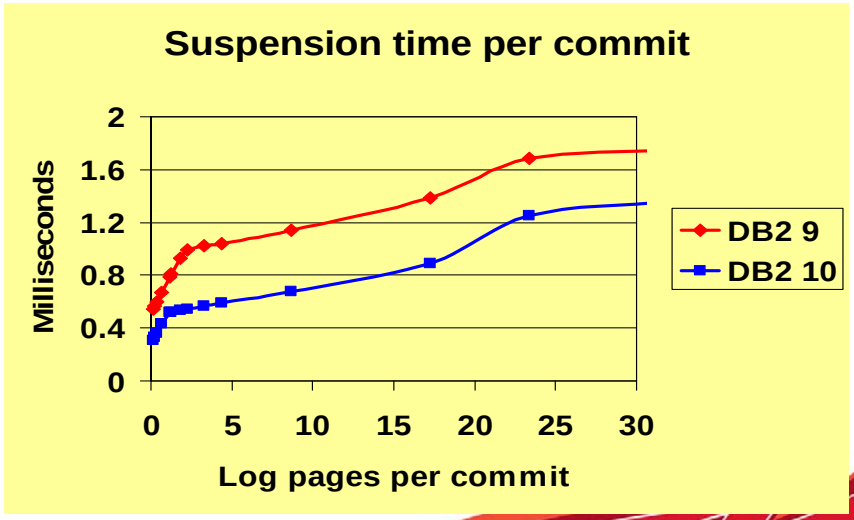

Logging performance, asynchronous write

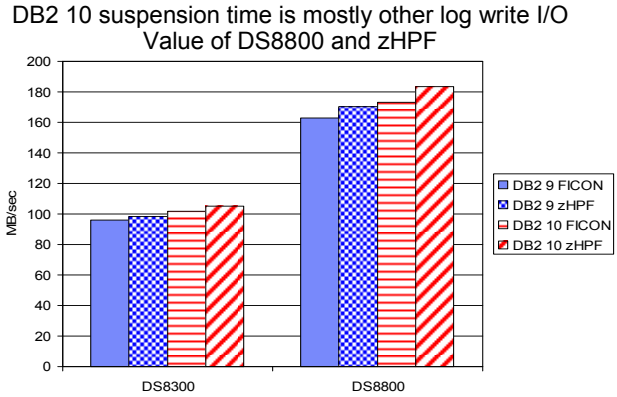

IBM Inside Sales
International Technical Support Organization Global Content Services

Compression on INSERT

© 2012 IBM Corporation
Redbooks Workshop

Compression on INSERT

- DB2 10 is able to build a compression dictionary automatically after a certain amount of data was inserted
  - INSERT / MERGE / LOAD SHRLEVEL CHANGE
  - In previous versions a LOAD or REORG are required to build the dictionary
- No measurable performance impact
- Since the first 1.2MB is not compressed, the percentage of rows that are not compressed is smaller than the Load utility, which begins using compression as soon as the compression dictionary is “filled up”

IBM Inside Sales International Technical Support Organization Global Content Services

Index I/O parallelism on update

© 2012 IBM Corporation

Redbooks Workshop

I/O parallelism for index update

- Index insert I/O parallelism manages concurrent I/O requests on different indexes into the buffer pool in parallel, with the intent of overlapping the synchronous I/O wait time for different indexes on the same table.
- DB2 must wait for I/O when reading the clustering index to find the candidate data page.
- Benefits tables with three or more indexes defined.
- I/O parallelism for index updates can also be disabled by setting the new online changeable DSNZPARM parameter INDEX_IO_PARALLELISM to NO. The default is YES.
- With index I/O parallelism, database I/O suspension is replaced by other read I/O suspension and class 2 CPU time is replaced by SRB time which is zIIP eligible.
- New IFCID 357 and IFCID 358 are available in DB2 10 to trace the start and end of index I/O parallel insert processing.

I/O parallelism for index updates (CM)

- DSNZPARM INDEX_IO_PARALLELISM (default ON)
  - Parallel read I/Os for additional indexes by using prefetch
  - Enabled only when there are index I/Os (buffer pool miss)
  - Applicable with all TS type except segmented TS
  - Enabled at 3rd (or 2nd if MC/APPEND) index update
- Elapsed time reduction
- Class 2 CPU time reduction with additional prefetch cost (DBM1 SRB)

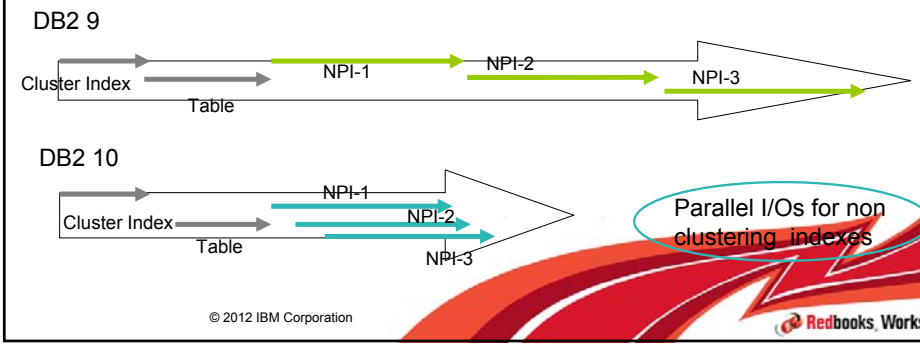

Insert index I/O parallelism

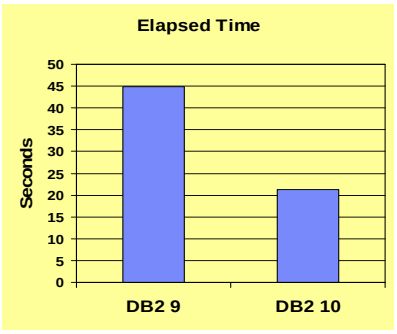

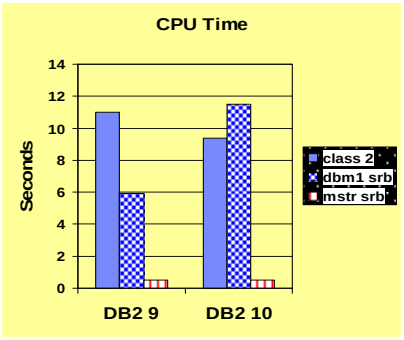

2000 random inserts on 6 indexes, 100% cache miss ratio
The result is: elapsed time decreases, class 2 CPU decreases, but the overall CPU increases. However DBM1 SRB time is zIIP eligible

IBM Inside Sales
International Technical Support Organization Global Content Services

zSynergy

© 2012 IBM Corporation
Redbooks Workshop

Exploitation of system z hardware/software

- z10 and zEnterprise 196 prefetch instruction
  - DB2 10 extensively uses to prefetch critical data structures ahead of time from real storage to the L2 CPU cache
- Large page frame size (1MB page frame) for buffer pools
  - Significant reduction of hit miss in TLB (translation lookaside buffer)
  - z10 and z196
- z10 specific instruction to improve performance
  - Cache line alignment
- System z10 processor supports HiperDispatch
  - A combination of hardware features, z/OS dispatching, and the z/OS Workload Manager that increases system capacity by increasing the probability of cache hits
  - Increases the system capacity by up to 10%
    - Varies according to the system configuration and workload
- More zIIP offload
- XML VSCR and XMF Fragment Validation
- z/OS 1.13 z/OS Batch Runtime Environment provides a managed environment for integration of Java and COBOL with DB2 for z/OS
- Exploits z/OS 1.12 new allocation functions to improve the performance of allocation, deallocation, open, and close of DB2 page sets. Also EAV.
- z/OS 1.11 for Security functions and DFSort zIIP eligibility
- WLM blocked workload support

z196

From the DB2 subsystem point of view, we expect more interaction and synergy with DB2 10. Faster CPUs, more CPUs, and more memory mean better DB2 performance, scalability.

Combined with DB2 10 improvements in buffer pool management, virtual storage constraint relief and latch contention reduction, DB2 applications can observe significant cost reductions and scalability improvements on zEnterprise.

- Compression hardware improvements provide better DB2 data compression performance.
- 192 MB L4 Cache benefit DB2 workloads
- Translation look-aside buffer (TLB) changes improve DB2 10 performance for 1 MB page sizes
- Hybrid architecture opens opportunities for query performance acceleration (IDAA requires z196 and DB2 9 or 10 for z/OS)
- zHPF improves I/O performance in general and new cache prefetch and WLM managed I/O with ES8800 6.2.

DB2 and zEnterprise 196

- Available since Sept 2010
- Larger processor cache (1.5 MB L2 per core, 24 MB L3 per chip, 129 MB L4)
- Higher DB2 CPU reduction can be achieved as #of processors per LPAR increases → Best fit with DB2 10 scalability
- More than 20% improvement with DB2 10 compared to DB2 9 on z196 64 way
- (SAP workloads show up to 1.8 times improvement in performance with DB2 10 on a z196 compared to DB2 9 on a z10 http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101845)
- DB2 9 OLTP, Insert, Utility and query workloads observing 20% to 40% DB2 CPU reduction compared to z10 processors. Typical range is 25 to 35 %.

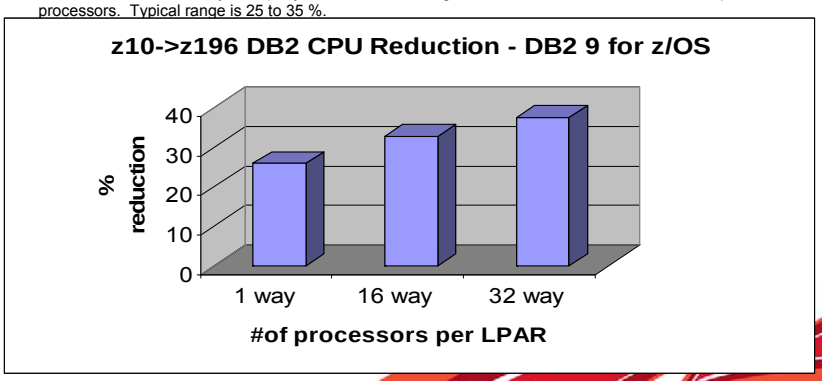

IBM Inside Sales
International Technical Support Organization Global Content Services

Delivers an end-to-end solution on a single platform

IBM DB2 Analytics Accelerator for z/OS

Capitalizing on the best of both worlds – System z and Netezza

The IBM DB2 Analytics Accelerator is a blending of System z and Netezza technologies that delivers unparalleled mixed workload performance for addressing complex analytic business needs.

- ***Performance, availability and scalability***
- ***Extreme Performance for Complex Analytics***
- ***Breakthrough technologies - Hardware acceleration***
- ***Transparent to DB2 applications***

Accelerating decisions to the speed of business

Memory and cache latencies

Cache and memory latency on a z10

- L1 Cache – 1 machine cycle (used separately for instructions and data)
- L1.5 Cache – 4 machine cycles
- L2 Cache – variable, 10's of machine cycles
- Real memory – ~ 600 machine cycles

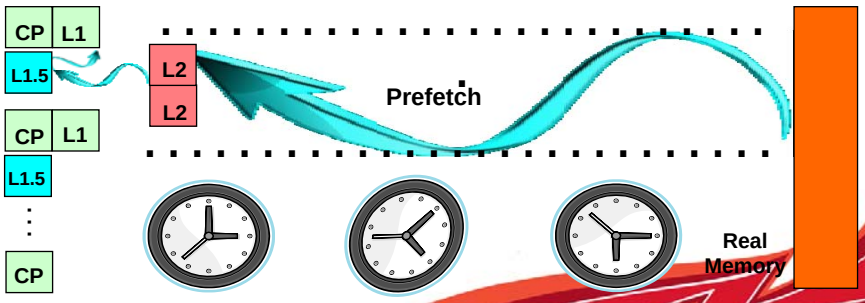

1 MB page frame support

1 MB page frames

- z/OS V1R10 introduced 1 MB real memory frames
- DB2 10 uses 1 MB real memory frames for buffer pools that are defined with the PGFIX(YES) attribute
  - 1 MB frames are not paged out by z/OS.
- Using 1 MB memory frames is intended to reduce the burden of virtual to real mapping memory management in z/OS by reducing the number of translation look-aside buffer (TLB) entries that are required for DB2 buffer pool pages
  - As shown on the following slides
- Results in better performance by decreasing the number of TLB misses that DB2 might incur
- The z/OS maintenance for large frames requires z/OS APAR OA31116, OA33702, and OA33529.

Why large pages?

- Significant reduction of hit miss in TLB
- z10 and z196
  - IEASYSXX LFAREA=(xx%| xxM | xxG| xxT) and RE-IPL
  - Backed by 256 contiguous 4K real frames and not page-able.
- 1 MB page is fixed, 64bit private pool
- Buffer pools with long term page fix (PGFIX=YES)
  - If 1MB page frames are available, DB2 will request 1MB first
- Long Term Page Fix from DB2 V8 to reduce CPU cost for I/O operations even without 1 MB
- 1 MB page frames to reduce CPU cost during get pages and release pages
  - Buffer pools with large variations of getpage activities
- Observed 1-4% CPU reduction in workload level by using 1 MB page frames

zIIP usage with DB2 10

IBM Inside Sales
International Technical Support Organization Global Content Services

Additional zIIP Redirect

- "Some" RUNSTATS, depending on the type of statistics
  - RUNSTATS with no additional parameters, the degree of zIIP eligibility can be up to 99.9%.
  - More complex statistics (for example, frequency statistics), the degree of zIIP
  - eligibility is less
- Most asynchronous I/O for buffer pools (dynamic prefetch, list prefetch, sequential prefetch, and deferred write processing)
  - Significant with index compression and Insert index I/O parallelism
- In z/OS V1R11, DFSORT is modified to allow additional zIIP redirect for DB2 utilities
  - in case of in memory object sort operations of fixed length records.
  - Included in the z/OS V1R11 base and is delivered to z/OS V1R10 through APAR PK85856
  - Included in the DB2 10 base and requires DB2 APAR PK85889 to be installed to function in DB2 for z/OS version 8 and 9.
- (And more queries can use parallelism)

Redbooks Workshop

More zIIP support for TCO improvement

- In DB2 10, the zIIP eligibility for DRDA workloads is increased from 55% to 60%
  - Also delivered in DB2 for z/OS Version 8 and 9 through APAR PM12256
- Parsing process of XML Schema validation (Class 2 CPU)
  - 100% of new validation parser is eligible
    - Can be zIIP, zAAP, or zAAP on zIIP
    - Retrofit into DB2 9 via PK90032 (preconditioning), PK90040 (enabling)

RUNSTATS eligibility by simple option

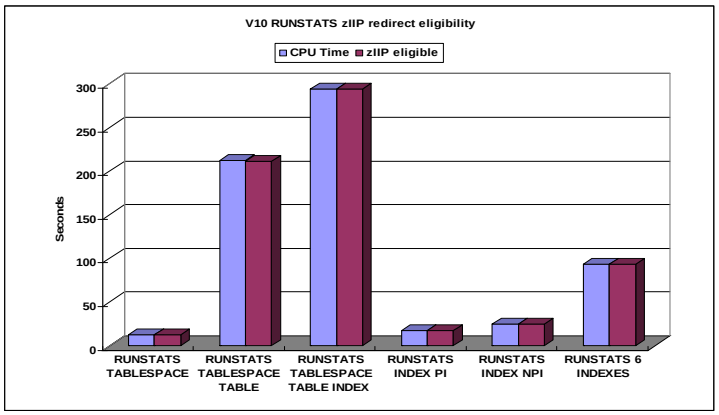

DB2 10 does not provide you in the Utility accounting with the CPU time which did not run on zIIP but is still eligible for. You need an RMF workload activity report. See details in the Redbook.

Asynchronous I/O reporting

- Buffer pool prefetch activities are asynchronously initiated by the database manager address space (DBM1) and are executed in a dependent enclave. Because asynchronous services buffer pool prefetch activities are not accounted to the DB2 client, they show up in the DB2 statistics report, where you can see CPU times accounted to zIIP in PREEMPT IIP SRB

CPU TIMES	TCB TIME	PREEMPT SRB	NONPREEMPT SRB	TOTAL TIME	PREEMPT IIP SRB
SYSTEM SERVICES ADDRESS SPACE	2.565794	0.825791	0.310654	3.702240	N/A
DATABASE SERVICES ADDRESS SPACE	0.187984	0.094044	0.003409	0.285437	0.486775
IRLM	0.000002	0.000000	0.065226	0.065228	N/A
DDF ADDRESS SPACE	0.005344	0.000000	0.000105	0.005448	0.000000
TOTAL	2.759124	0.919835	0.379394	4.058353	0.486775

IBM Inside Sales
International Technical Support Organization Global Content Services

Open and close data sets

© 2012 IBM Corporation
Redbooks Workshop

Shutdown and restart times, and DSMAX

- z/OS 1.12 provides a substantial reduction in allocation and deallocation of data sets and GRS improvements
  - New interface to open/close data sets
    - SYS1.PARMLIB ALLOCxx SYSTEM MEMDSENQMGMT
  - DB2 10 systems that use a large number of data sets (DSMAX of 100,000) will see improvements in shutdown and restart times.
  - Also, in normal operation, if a table space with a large number of partitions needs to have all partitions to be opened by DB2 at about the same time
- z/OS APARs OA32612 (reduce catalog searches) and OA33633 (ENQ improvement)
- z/OS APAR OA37697 helps DB2 DROP INDEX performance
- Z/OS APAR OA36354 helps DB2 CREATE INDEX performance.

z/OS 1.12 and DB2 Version 8, 9 and 10

- DB2 changes (Sept/Oct 2010)
  - DB2 10 with base code, DB2 V8 and 9 require APARs
    - PM00068 for support of 100K data set open
      - Without APAR, actual max was 60K+
    - PM17542
      - Enable new z/OS 1.12 allocation interface
      - Disable DD level SMF accounting
    - PM18557
      - Enable new z/OS 1.12 GRS interface
- And take one of the following actions
  - Update the ALLOCxx parmlib member to set the SYSTEM MEMDSENQMGMT value to ENABLE (strongly recommended)
  - Issue system command SETALLOC SYSTEM,MEMDSENQMGMT=ENABLE then restart DB2.

Open 100,000 data sets using 20 concurrent jobs

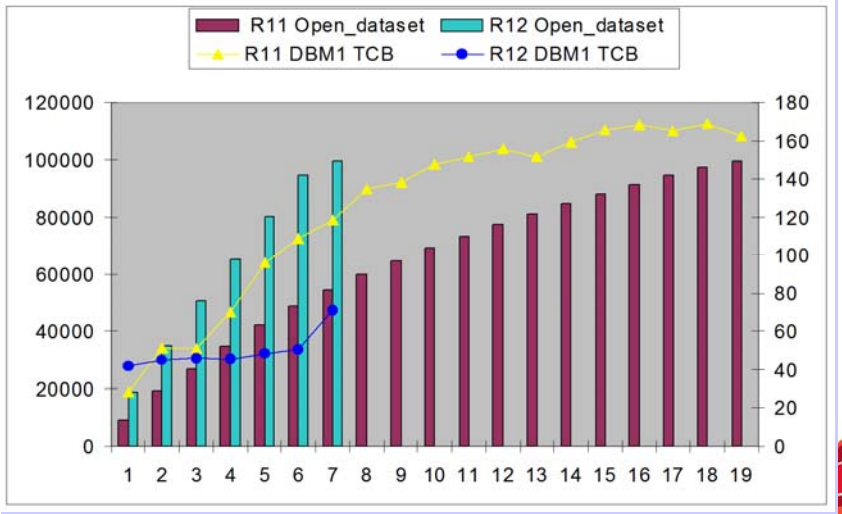

Improvement with z/OS 1.12 and DB2 APARs

- Initial open cost after DB2 restart
  - 100,000 data set open with DSMAX 100,000
    - z/OS 1.11 and 1.12, z10, DB2 9 20 parallel open, ECS enabled
    - 3 times less elapsed time with z/OS 1.12 and DB2 APARs
    - 6 times DB2 DBM1 CPU reduction with z/OS 1.12
- Open and Close with large DSMAX
  - 20,000 data set open with DSMAX 80,000 and 80K data sets are already open
    - z/OS 1.11 and 1.12, z10, DB2 9 20 parallel open, ECS enabled
    - 4 times less elapsed time and CPU time with z/OS 1.12

Disk storage enhancements

IBM Inside Sales
International Technical Support Organization Global Content Services

Synergy with disk storage

- Extended Address Volumes
  - Remove current limitations of the largest traditional 3390 model volume (3390-54) with a capacity of 65,520 cylinders (about 54 GB), and a maximum of 65280 I/O devices in a Parallel Sysplex.
  - Size limited to 223 GB (262,668 Max cylinders)
  - Supported in z/OS V1R10 and higher
- IBM System Storage DS8800
- High Performance FICON (zHPF)
- Other enhancements
  - Dynamic Volume Expansion

© 2012 IBM Corporation

Redbooks Workshop

DB2 support of EAV data sets

DB2 Objects	z/OS 1.10	z/OS 1.11	z/OS 1.12
Tables and Indexes	Yes	Yes	Yes
BSDS	Yes	Yes	Yes
Active Logs	Yes	Yes	Yes
Archive Logs	No	Yes, if EF Sequential	Yes
Utilities sequential input and output data sets	No	Yes, if EF Sequential	Yes
Utilities partitioned data sets and PDSEs	No	No	Yes
Sort work datasets	No	No	Yes, if DFSORT used by utilities
DB2 installation data sets (CLISTs, panels, samples, macros, etc.)	No	No	Yes
SDSNLINK SDSNLOAD	No	No	Yes

Introducing the DS8800

- **Faster hardware**
- **More efficient footprint**

	Frame	A	B	C	D	E
DS8700	#disks	128	256	256	256	128
	#DA ports	4	8	4	0	0
DS8800	#disks	240	336	480	N/A	N/A
	#DA ports	8	8	0	N/A	N/A

DS8800

- **Exceptional reliability for critical workloads**
- ***Generally Available on November 19, 2010***

DS8000 recent enhancements

- Higher CPU capacity requires greater I/O bandwidth and efficiency. High Performance FICON (zHPF) enhances the z/Architecture® and the FICON interface architecture to provide greater I/O efficiency. zHPF is a data transfer protocol that is optionally employed for accessing data from an IBM DS8000 storage subsystem.
- Both the DS8800 and the zHPF provide great improvements when used with DB2 for z/OS.
- Measurements were done comparing DS8800 with DS8300 with zHPF and FICON showing how DB2 functions can benefit from the I/O improvements. In this section we discuss:
  - Prefetch improvement via disk enhancement
  - DB2 logging and insert with disk enhancements
  - Utilities and storage enhancement
- z/OS APAR OA39087 helps DB2 LOAD RESUME on zHPF enabled device

Link protocol comparison (4 KB READ)

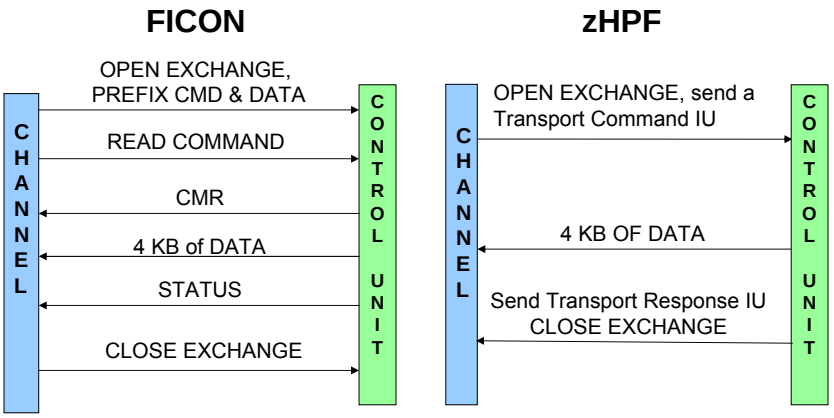

zHPF provides a much simpler link protocol than FICON

zHPF I/O eligibility

Type of DB2 I/O	z10	z196
Random single 4K read/write	YES	YES
Sequential prefetch and dynamic prefetch	NO	YES
DB2 workfiles, reads and update writes	YES	YES
List prefetch (i.e. scattered reads)	NO	NO
Log writes <=64K	YES	YES
Log writes > 64K	NO	YES
Log reads (with DB2 9 or 10)	NO	YES
Sequential update writes	NO	YES
Scattered writes	NO	NO
Contiguous update writes (>64K)	NO	YES
Format and preformat	NO	NO
Utility table space scans (sequential prefetch)	NO	YES
Sequential reads from DSORG=PS, EF data sets	NO	YES
Sequential reads from DSORG=PS, non-EF data sets	NO	NO
Sequential writes to DSORG=PS	NO	NO

Prefetch improvement via disk enhancement

Sequential prefetch reading 4 KB pages from cache

- Prefetch operations often read more than 64 KB bytes and therefore were not eligible for zHPF prior to the z196. The z196 makes sequential prefetch and dynamic prefetch both zHPF-eligible.
- The buffer pool is defined large enough to enable DB2 to read 256 KB per I/O while reading from disk cache.
- The results show that using FICON, the DS8800 channel throughput is 44% higher than the DS8300. Using zHPF, the DS8800 channel throughput is 55% higher than the DS8300. If we compare the DS8800 with zHPF to the DS8300 with FICON, the channel throughput increases 66%.

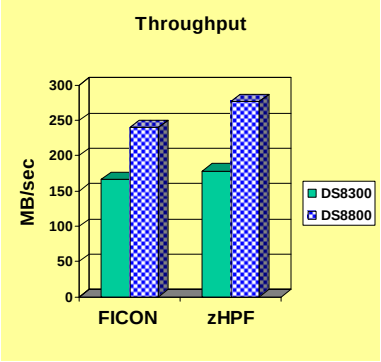

Prefetch improvement via disk enhancement…

- **Dynamic prefetch reading 4 KB pages from cache**
- The second set of measurements was done for dynamic prefetch using 4 KB pages on a DS8300 and a DS8800 attached to a z196 processor. The z196 processor enables DB2 prefetch I/Os to become zHPF eligible. The measurement conditions are similar to the sequential prefetch case where pages are read from the disk cache.
- The results show that using FICON, the DS8800 channel throughput is 33% higher than the DS8300. Using zHPF, the DS8800 channel throughput is 38% higher than the DS8300. And together comparing the DS8800 with zHPF to DS8300 with FICON, channel throughput increases 58%.

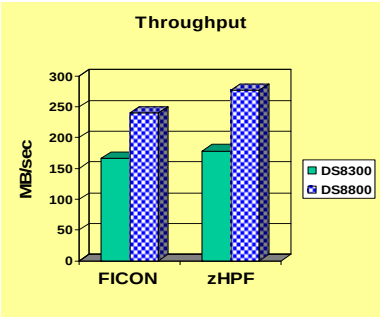

Prefetch improvement via disk enhancement…

- **Sequential prefetch reading 4 KB pages from disk**
- The third set of measurements was done with prefetch reading data from disk, not retrieving the data from the disk cache
- The results show that with sequential prefetch, DS8800 channel throughput is 33% higher than the DS8300. Using dynamic prefetch, DS8800 channel throughput is 29% higher than the DS8300.

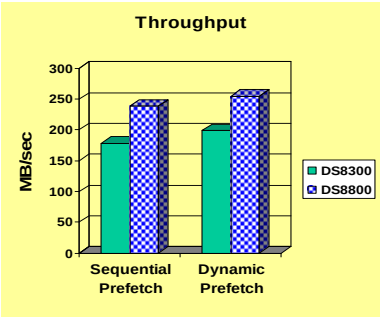

Prefetch improvement via disk enhancement…

- **Prefetch reading larger pages from disk and cache**
- The fourth set of measurements was done with sequential and dynamic prefetch reading data from disk and cache, for page sizes of 4 and 16 KB.
- A larger page size of 16 KB increases the throughput by 19% to 28%.
- When fetching from cache, dynamic is faster than sequential prefetch, but not so when fetching from disk (on the DS8800).

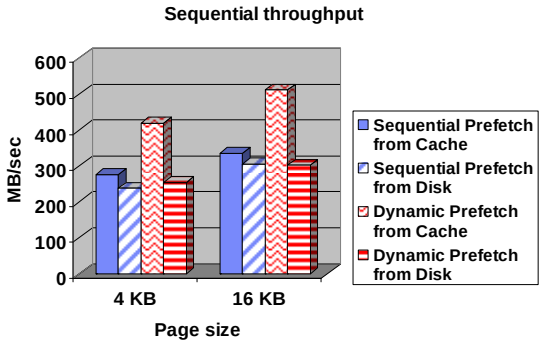

DB2 logging and insert with disk enhancements

Maximum log throughput

These measurements were done on a z196 which enables the log I/Os exceeding 64 KB to be zHPF eligible, providing an additional boost in performance. It shows that the DS8800 improves the maximum throughput by about 70% when the log buffer queues become large.

If the number of 4 KB log buffers is less than or equal to 16, these writes are eligible for zHPF on a z10. If the number of buffers is greater than 16, the z10 cannot use zHPF, but the z196 can. Measurements show that DB2 10 can achieve a slightly higher log throughput with zHPF compared to FICON. z196 enables log I/Os greater than 64 KB to use zHPF and DS8800 requires FICON Express 8 to achieve 8 Gbps speeds channel.

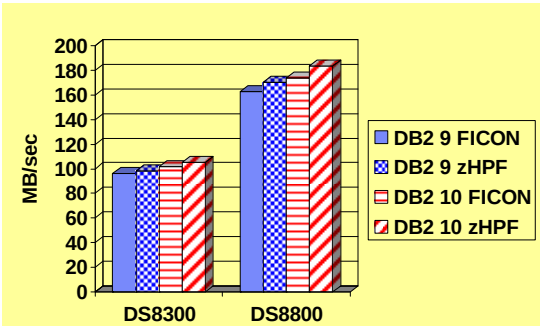

DB2 logging and insert with disk enhancements…

Load preformat

Load preformat is a convenient way to quantify preformat that results from insert processing. We measured write throughput both on the DS8300 and the DS8800.

The results show that the DS8800 preformat throughput is 23% higher than the DS8300.

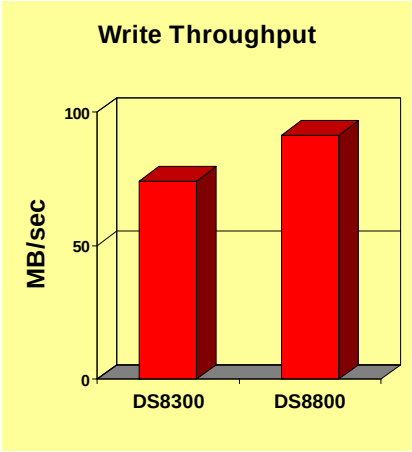

Utilities and storage enhancement

- Load utility
- The first set of measurements shows the Load utility performance using various page sizes on the DS8800 and the DS8300
- The results show significant improvements with the 16 KB and 32 KB pages, but not with 4 KB and 8 KB pages because of the zHPF not supporting format writes.

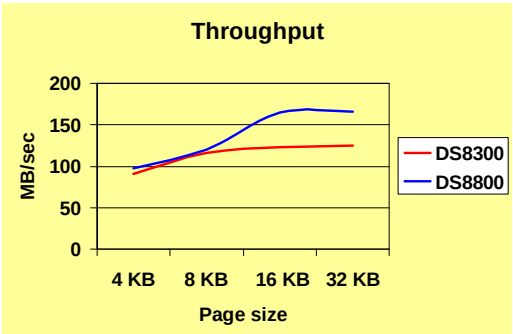

Utilities and storage enhancement…

- Load utility
- The second set of measurements shows the LOAD utility performance when loading small 20 KB LOBs using 4 KB pages table.
- The results show that with DB2 9, the DS8800 takes advantage of the format write. With DB2 10 both devices improve by about 50%, with the DS8800 not showing better performance than the DS8300 because DB2 is unable to drive the device utilization higher.

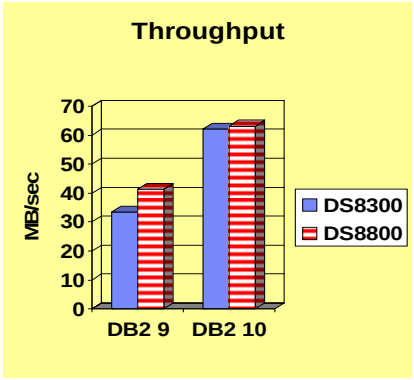

Utilities and storage enhancement…

- Unload utility
- A set of Unload utility measurements were done with extended format data sets both on the DS8800 and the DS8800 in DB2 9 and DB2 10
- The DS8800 shows about 19% more throughput than the DS8300 in DB2 9 and 33% in DB2 10.

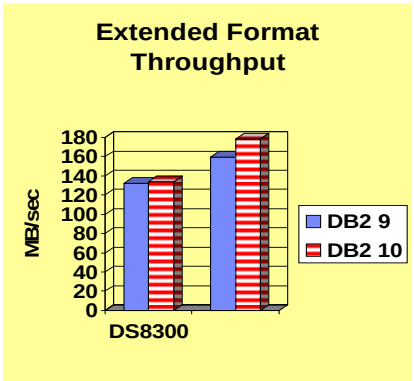

Utilities and storage enhancement…

- **EF BSAM enhancements with utilities**
- z/OS 1.9 introduced support for long-term page fixing of basic sequential access method (BSAM) buffers, and z/OS 1.10 introduced support for 64-bit BSAM buffers if the data set is in extended format. DB2 10 utilities exploit these recent z/OS enhancements, by offering the following enhancement
  - Allocating 64-bit buffers for BSAM data sets
  - Allocating more BSAM buffers for faster I/O
  - Long term page fixing BSAM buffers
  - DB2 10 utilities, increase MULTSDN from 6 to 10 and MULTACC from 3 to 5
- The measurement shows how the value of DB2 10 and new storage enhancements benefit DB2 when making use of enhancement such as increasing MULTACC from 3 to 5, the faster HA and DA in the DS8800, zHPF, multiple streams per raid rank for EF and non-EF data sets.
- See the DS8800 performance using a block size of 27966 on extended format BSAM data sets. For the extended format, the zHPF can be enabled with z196 processor.

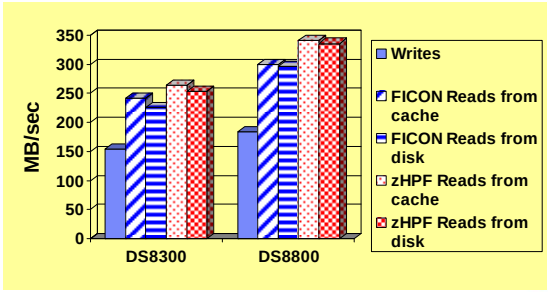

Disk synergy summary

- The DS8800 is the IBM's flagship enterprise disk platform. Built on 50+ years of enterprise class innovation, the DS8800 enables much higher performance and scalability than its predecessor models, while preserving client investments in prior DS8000 models. The DS8800 offers faster processor speeds, faster adapters and buses, and all with a smaller footprint in terms of floor space and energy consumption.
- For DB2 this means that all workloads will run faster: faster prefetch I/O, faster log I/O and faster utilities. The DS8800 also provides for much better OLTP performance.
- For more information, see IBM System Storage DS8800 Performance Whitepaper at:
  - Internal: http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101799
  - Partners: http://partners.boulder.ibm.com/src/atsmastr.nsf/WebIndex/WP101799

References

- **IBM zEnterprise 196 I/O Performance**
  ftp://public.dhe.ibm.com/common/ssi/ecm/en/zsw03169usen/ZSW03169USEN.PDF
- **DS8800 Technical Overview Redbook**
  http://www.redbooks.ibm.com/redpieces/pdfs/sg248886.pdf
- **IBM® System Storage™ DS8800™ Performance Whitepaper**
  http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101799
- **IBM® System Storage™ DS8700™ Performance Whitepaper**
  http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101614
- **IBM® System Storage™ DS8700™ Performance with Easy Tier®**
  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101675
- **DB2 for z/OS and List Prefetch Optimizer**
  - **http://www.redbooks.ibm.com/abstracts/redp4862.html?Open**
- **Ready to Access DB2 for z/OS Data on Solid-State Drives**
  http://www.redbooks.ibm.com/abstracts/redp4537.html

IBM Inside Sales
International Technical Support Organization Global Content Services

SMF compression

© 2012 IBM Corporation
Redbooks Workshop

SMF compression

- New DSNZPARM parameter SMFCOMP directs DB2 to request compression for SMF trace records 100, 101 AND 102 via standard z/OS Data Compression and Expansion Services CSRCESRV
- Trace data to GTF and OPx is not compressed.
- SMFCOMP is specified on the installation panel DSNTIPN in the field COMPRESS SMF RECS. The default value is OFF.
- APAR PM27872 provides you with decompression routine DSNTSMFD and sample JCL, DSNTEJDS..

```
DSNTIPN         INSTALL DB2 - TRACING PARAMETERS

===> _
Enter data below:

 1  AUDIT TRACE        ===> NO      > Audit classes to start. NO,YES,list
 2  TRACE AUTO START   ===> NO      > Global classes to start. YES,NO,list
 3  TRACE SIZE         ===> 64K       Trace table size in bytes. 4K-396K
 4  SMF ACCOUNTING     ===> 1       > Accounting classes to start. NO,YES,list
 5  SMF STATISTICS     ===> YES     > Statistics classes to start. NO,YES,list
 6  STATISTICS TIME    ===> 1         Time interval in minutes. 1-60
 7  STATISTICS SYNC    ===> NO        Synchronization within the hour. NO,0-59
 8  DATASET STATS TIME ===> 5         Time interval in minutes. 1-60
 9  MONITOR TRACE      ===> NO      > Monitor classes to start. NO,YES,list
10  MONITOR SIZE       ===> 1M        Default monitor buffer size. 1M-64M
11  UNICODE IFCIDS     ===> NO        Include UNICODE data when writing IFCIDS
12  DDF/RRSAF ACCUM    ===> 10        Rollup acctng for DDF/RRSAF. NO, 2-64K
13  AGGREGATION FIELDS ===> 0         Rollup acctng aggregation fields
14  COMPRESS SMF RECS  ===> OFF       Compress trace records destined for SMF

PRESS:   ENTER to continue   RETURN to exit   HELP for more information
```


DSNTSMFD output

```
*** DSNTSMFD *** STARTING     2011/03/14          13:02:13
-------------------------------------------------------------------------------
-
Total records read:.................................        2408
  Total DB2 records read:...........................        1307
    Total DB2 compressed records read:..............         197
    Total DB2 compressed records decompressed:......         197
  Total non-DB2 records read:.......................        1101

Aggregate size of all input records:................     9237888            8M
  Aggregate size of all input DB2 records:..........     2152260            2M
    Aggregate size of all DB2 compressed records:...      170830          166K
  Aggregate size of all output DB2 records:.........     2392486            2M
    Aggregate size of all DB2 expanded records:.....      411056          401K
  Aggregate size of all non-DB2 input records:......     7085628            6M

    Percentage saved using compression..............          58%

Details by DB2 subsystem
...
Subsystem ID:  DBOA
        Number of records:..........................         259
        Number of compressed records:...............         197
        Aggregate size of DB2 records:..............      302890          295K
        Aggregate size of DB2 compressed records:...      170830          166K
        Aggregate size of DB2 expanded records:.....      411056          401K
        Percentage saved using compression..........          58%


-------------------------------------------------------------------------------
*** DSNTSMFD *** ENDING       2011/03/14          13:02:13
```

What's new

IBM Inside Sales
International Technical Support Organization Global Content Services

DB2 10 News

Items delivered post GA, based on customer input

- Embedded XQuery support – PM47618
  - XQuery in IBM DB2 10 for z/OS, TIPS0896
- New built-in functions PACK and UNPACK – PM56631
  - Needed for SPSS in-database analytics
- Stored procedure monitoring improvements – PM53243
  - More easily identify a problematic stored procedure or a statement within that stored procedure
- LOAD SHRLEVEL CHANGE partition parallelism with a single input dataset – PM57632
  - Maintaining input datasets by part boundaries can be difficult after ALTER of the ranges or REORG REBALANCE
- New DSNZPARM for OPTIMIZE FOR 1 ROW to allow sort access plans – PM56845/PM57630
  - DB2 10 migration issue for some customers
- GBP Delete Name enhancement – PM67544
  - Performance enhancement for sysplexes that span longer distances
- SYSPACKAGE LASTUSED – PM66287
  - Remove RTS interference
- New SORTNPSI parm in REORG and ZPARM – PM55051
  - Performance for REORG partitions with NPSIs

Other Recent DB2 10 APARs

- New DEL_CFSTRUCTS_ON_RESTART DSNZPARM for auto-delete of CF structures on restart if no active connections exist – PM28925, PM31807 (IRLM support)
  - Useful for DR environments
- ALTER MAXROWS to set AREO* rather than AREOR – PM43597
  - Allows MAXROWS to take effect at partition level
- Ability to SELECT from SYSLGRNX – PM35190, PM42331
  - ISO(UR) enforced
- Increase FLA storage from 100 MB to 510 MB – PM31641
  - 5X increase in recommended limit on concurrent RECOVER jobs per DB2
- LOB pageset support for RECOVER BACKOUT YES – PM45650
- REORG to ignore free space – PM53254
- REORG parallel partitions – PM52012
- DB2 Sort 1.3 performance with DB2 utilities – PM69522
- PACK/UNPACK built-in function for scoring support – PM56631

Other updates

- Several new DSNZPARMs for FlashCopy use by utilities
- Several performance APARs
  - RETAIN (see also Appendix A in DB2 10 Performance Topics, SG24-7492)
  - Various Forums
- Diagnostic Guide
  - See APAR PM63485
- DB2 Accessories Suite for z/OS Version 2.2 (see II14686)
  - IBM Data Studio V3.1.1
  - IBM SPSS Modeler Server Scoring Adapter for DB2 on z/OS
  - IBM Spatial Support for DB2 for z/OS
  - The International Components for Unicode for DB/2 for z/OS
  - IBM Text Search for DB2 for z/OS, which replaces IBM Omnifind Text Search Server for DB2 for z/OS
  - IBM Installable Information Management Software for z/OS Solutions Information Center

DB2 and zEnterprise EC12

- **Faster CPU – 1.25x compared to z196**
  - Up to 1.3x measured with DB2 OLTP workload
  - Up to 1.35x measured with DB2 Query workload
  - Less compression overhead with DB2 data (1-15%)
- **50% More System Capacity to help consolidation**
  - Excellent synergy with DB2 10 scalability
- **New Features DB2 plans to exploit**
  - FLASH memory and pageable 1MB frames
    - Enables larger DB2 Buffer Pools
    - DB2 thread storage backed by pageable 1 MB for CPU reductions
  - 2 GB frame support
    - Enable even larger DB2 Buffer Pools
    - Larger frames expected to provide additional CPU savings, especially for very large memory

DB2 11 – "Sequoia" (Beta announced Oct. 3)

- Unmatched Reliability, Availability, and Scalability
  - Improved Data Sharing performance and efficiency
  - Even less downtime by removing growth limitations
  - Simplify management, improved autonomics, reduce planned outages with more online schema changes and utilities improvements
- Save money, Save time
  - Aggressive CPU reduction goals
  - 0-5% for many OLTP and heavy INSERT workloads
  - Up to 20-30% for key query workloads
  - Save time and resources with new autonomic and application development capabilities
- Simpler, faster migration
  - SQL compatibility, separate system migration from application migration
  - Access path stability improvements
  - Better application performance with SQL and XML enhancements
- Enhanced business analytics
  - Faster, more efficient performance for query workloads
  - Transparent archiving supports more efficient analysis
  - IDAA enhancements
  - More efficient inline database scoring enables predictive analytics

Extended RBA problem statement

- DB2's Relative Byte Address (RBA) for logging is 6 bytes
  - Gives 256 TB of log record addressing capacity per DB2 subsystem/member
- With heavy sustained logging rates, DB2 can exhaust the 6-byte RBA
  - DSNJ032I and DSNJ033E warning messages
    - Alert-level = 'WARNING' when RBA reaches x'F00000000000'
    - Alert-level = 'CRITICAL' when RBA reaches x'FFFF00000000'
  - Manual recovery actions are needed
    - Data Sharing: shut down the affected member and start a new member in its place
    - Non Data Sharing: reset all PGLOGRBA values back to zero (extended outage)
    - Documented in the DB2 Administration Guide
  - If alert-level reaches 'CRITICAL' then DB2 terminates to protect data integrity and force recovery actions
    - Reason code 00D10251
    - ACCESS(MAINT) restart allowed to prepare for recovery actions

Extended LRSN problem statement…

- The data sharing Log Record Sequence Number (LRSN) is derived from the 8-byte time-of-day clock which hits end of range in 2042
- However, some data sharing groups have a non-zero LRSN "delta" which gets added to the TOD clock
  - If a non-zero "delta" exists, then the LRSN will hit end of range prior to 2042
  - Use DSNJU004 to determine if you have a non-zero LRSN delta value
  - A "delta" value could be set when data sharing is enabled or re-enabled
    - Whenever the end-of-log RBA of the enabling member is past the TOD clock
- Some non data sharing customers have enabled data sharing to circumvent RBA nearing end-of-range
  - This would cause a non-zero LRSN delta, so LRSN hits end of range before 2042
- 6-byte LRSN value has precision to only 16 microseconds
  - Can cause LRSN 'spinning' which burns extra CPU and aggravated log latch contention
  - DB2 9 NFM addresses most LRSN spin situations, and DB2 10 NFM enhanced further. But some spins still exist due to the 16 microsec granularity (log latch not held, page latches are)

DB2 11 planned RBA/LRSN solution

- Expand the RBA and LRSN to 10 bytes
  - RBA addressing capacity of 1 yottabyte (2**80)
  - LRSN extended on left by 1 byte, on the right by 3 bytes
    - >30,000 years and 16Mx more precision
  - 8 bytes is not sufficient to solve LRSN issues and may not give sufficient capacity for the longer term
- NFM only (6 byte RBA/LRSN continues to be used in CM)
- Once in NFM, DB2 continues to use 6-byte values until you take action to convert
- Two conversion tasks:
  - Convert BSDSs to new format to enable logging with larger RBAs/LRSNs
  - Convert pagesets to new page format
- These tasks are optional
  - If you don't care about larger RBAs/LRSNs then you don't have to convert
  - But performance will be better if you convert BSDSs (avoid internal conversion overhead on log write)
- BSDSs can be converted without converting pagesets
- Pagesets can be converted in a piecemeal fashion
  - Expectation is that most customers will roll the conversion over a period of days/weeks/months

Some planned DB2 11 RAS improvements

- BIND / DDL concurrency with persistent threads
  - Use of persistent threads likely will increase in DB2 10 with vstor relief
- More online schema changes
  - Alter partitioning limit keys
  - DROP column
  - Point in time recovery support for deferred schema changes
- Cancel DDF Threads – new FORCE option
- Open data set limit raised to 200K
- Workfile space shortage warning new instrumentation and messages
- REORG avoidance
  - Automatic cleanup of index pseudo deleted entries
- Online REORG improvements – if a REORG is needed, then the goal is non-disruptive
  - SWITCH phase performance improvements
  - Drain improvements to improve concurrency of log apply and switch phases
  - REORG REBALANCE SHRLEVEL(CHANGE)
- Easier query performance management
  - Runstats input to optimizer
- Resource Limit Facility support for static SQL
- DB2/RACF authorization control enhancements
- Buffer pool management improvements
  - New FRAMESIZE BP attribute for direct control of z/OS large page frame usage
  - Max/min size for WLM system-managed BPs

DB2 11 some planned new application features

- Global variables
- SQL PL improvements (performance, manageability, function)
  - Autonomous transactions
  - Lift 2M SQL statement length limit
  - Array data type support
- Alias/synonym support for Sequence Objects
- Grouping sets
- DPSI performance improvements
- DGTT performance improvements
- Bitemporal data enhancements
  - Support for views
  - Special register support
- Transparent archive query
  - New DDL to relate the current table to the archive
  - Applications can query current + archive with no SQL changes
  - Archiving process is user-controlled
  - Leverages DB2 10 temporal constructs for archiving use cases

Easier DB2 Version upgrade

- SQL Compatibility – new option for enforcement
  - Limit SQL incompatibilities when possible
  - Provide mechanism to identify applications affected by SQL changes
  - Provide seamless mechanism to make changes at an application (package) level
    - This mechanism will enable support for up to two back level releases (N-2)
    - DB2 11 will be the initial deployment of this capability
    - DB2 10 will be the lowest level of compatibility supported
- Access path stability improvements

DB2 11 planning

- Dual mode migration (CM, NFM)
- Migration from DB2 10 only (no skip)
- z/OS 1.13 or above, z10 or above.
- No pre-V9 bound packages
- Sysplex query parallelism support is removed

Query Accelerator

- **Data warehouse challenges on z/OS**
- **Accelerating BI workloads with the IBM DB2 Analytics Accelerator**
- **Feasibility study/value assessment and Virtual Accelerator**
- **Installation quick start offering**
- **DB2 configuration and z/OS WLM settings**
- **IDAA Studio to define and load data**
- **Latency management**
- **Query Acceleration management and security considerations**
- **Monitoring IDAA and performance consideration**
- **Data sharing and failover scenarios**
- **What's new**

Why System z for Data Warehousing?

The operational data is native to the System z

- *Using the data locally lowers data latency, movement, and transformation enabling decisions to be driven with the most-current information.*

New BI trends map well to the strengths of DB2 for z/OS and System z

- *Core System z strengths map well to new Dynamic Warehouse and Operational BI initiatives that are driving the need for: mixed workload support, lower latency, increased reliability, availability, security, and compliance in a Warehouse DBMS.*

Consolidation and virtualization trends benefit from core System z strengths

- *Virtualization strengths of System z enable the consolidation of diverse application servers, to a single server, allowing for centralization and standardization of DW and BI environments.*

IBM specialty processors and the System z offer new ways to manage costs

- *zIIPs and IFLs are driving down hardware and software costs; DW/BI can make excellent use of these processors, ultimately driving cost advantages. System z processors deliver excellent speeds and feeds, at lower cost.*

Single System image view

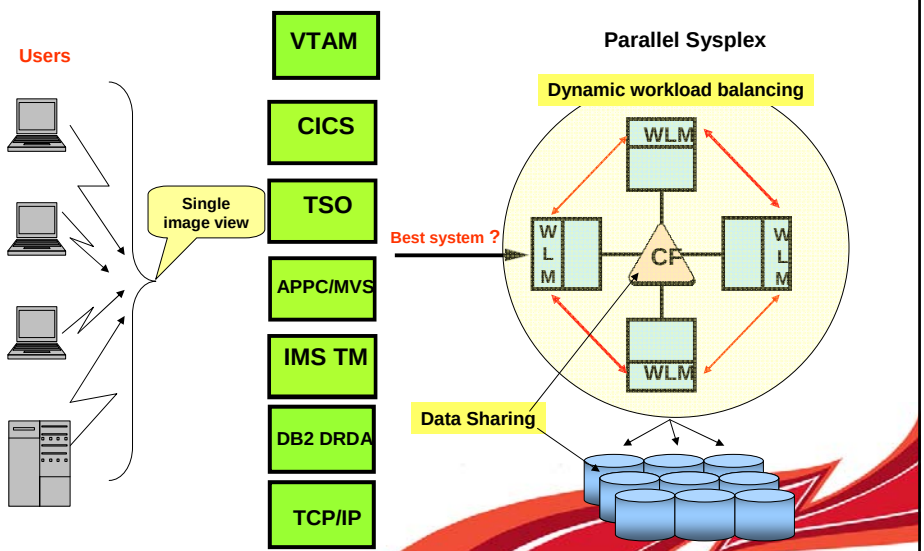

Prevent Query Monopolization with WLM

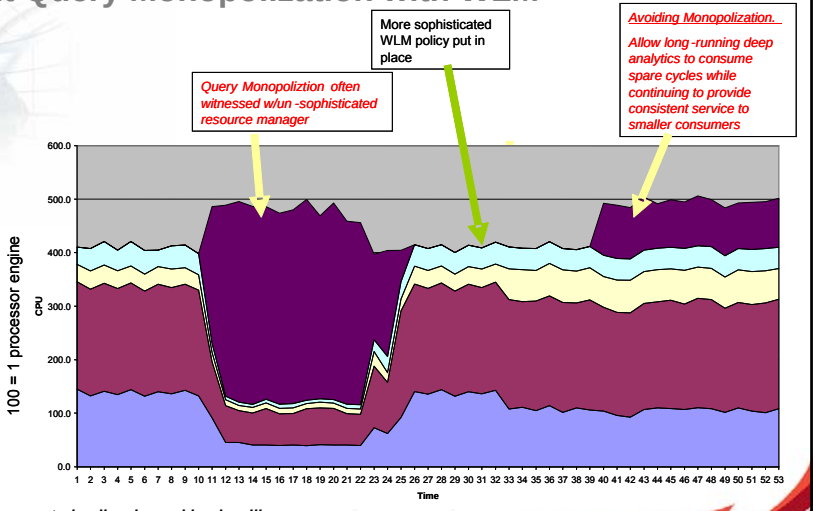

When you have unexpected spikes in workload, will you be able to handle it?
… Or do you have to micro-manage the system ?

IBM Inside Sales
International Technical Support Organization Global Content Services

Information-Led Transformation for System z

Business needs demand a strategic, centralized approach

- **Solution Edition for Data Warehouse**
- **IBM Smart Analytics System 9700 / 9710**
- **InfoSphere Warehouse on System z**
- **IBM Cognos for z/OS**
- **IBM DB2 Analytics Accelerator**
- **Master Data Management**
- **IBM Smart Analytics Cloud**
- **IBM SPSS**
- **QMF V10 enhancements**

DB2 Analytics Accelerator V2.1

Capitalizing on the best of both worlds – System z and Netezza

What is it?

The IBM DB2 Analytics Accelerator is a workload optimized, appliance add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.

How is it different

- ***Performance:*** Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- ***Integration:*** Deep integration with DB2 provides transparency to all applications.
- ***Self-managed workloads:*** queries are executed in the most efficient location
- ***Transparency:*** applications connected to DB2 are entirely unaware of the Accelerator
- ***Simplified administration:*** appliance hands-free operations, eliminating most database tuning tasks

Breakthrough Technology Enabling New Opportunities

Optimizing to the Workload

Marrying the best of each environment

Total solution remains centrally managed by System z…

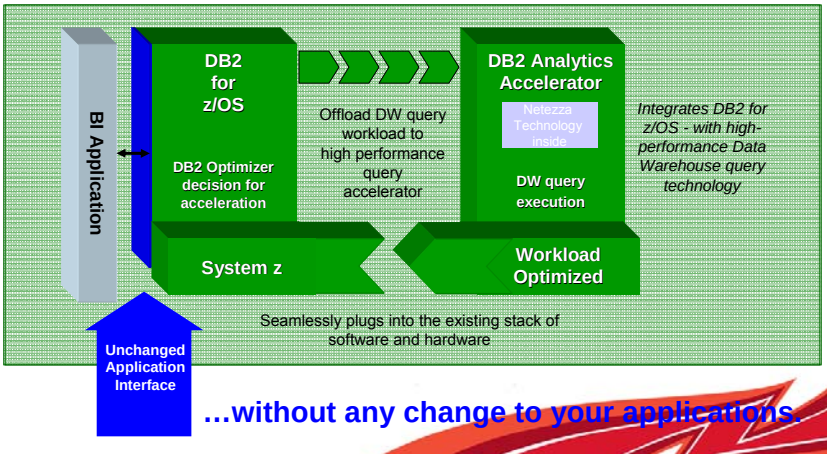

…without any change to your applications.

IBM DB2 Analytics Accelerator V2 Product Components

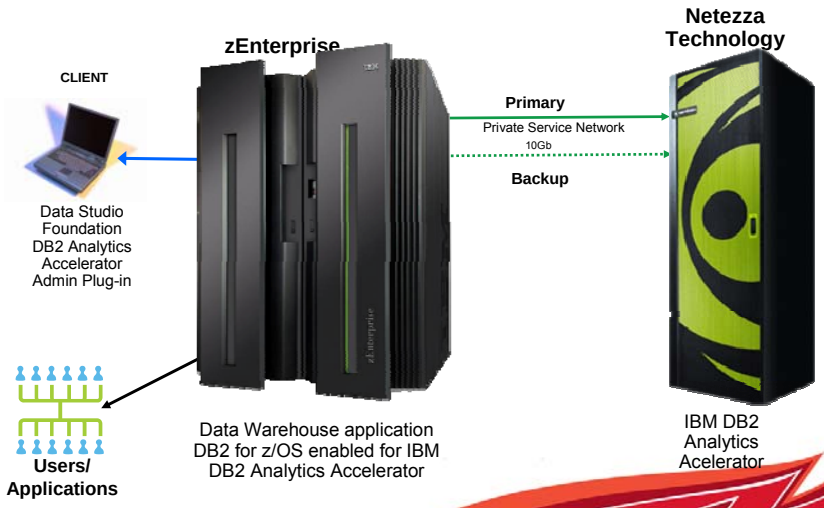

The DB2 side

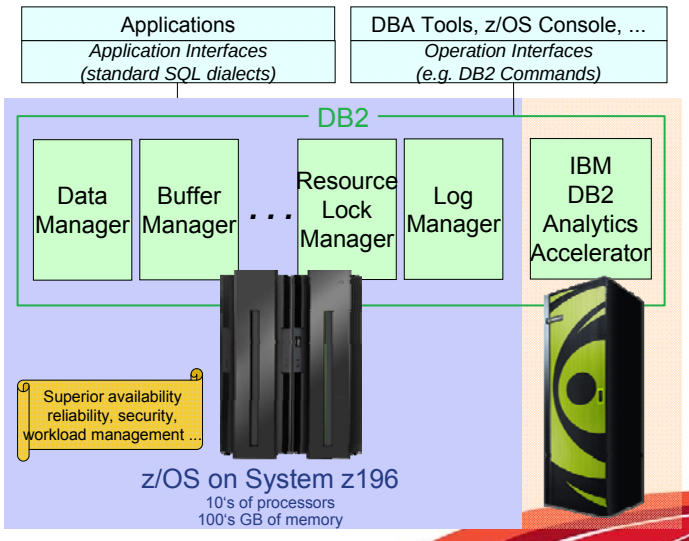

DB2 Analytics Accelerator V2

Powered by Netezza 1000 Appliance

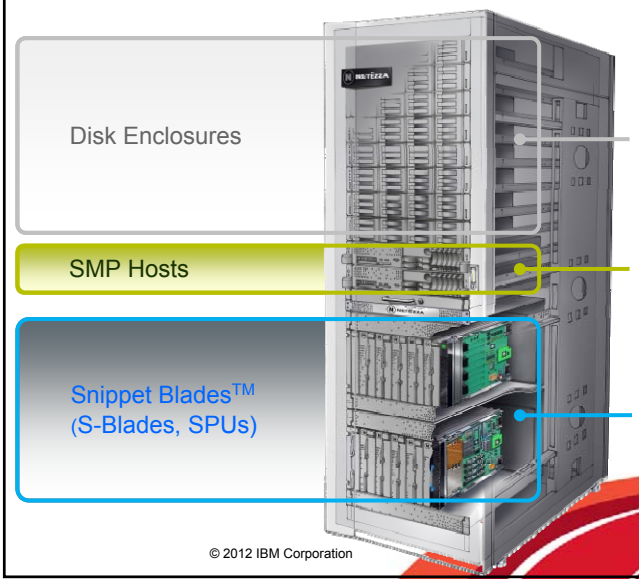

Slice of User Data
Swap and Mirror partitions
High speed data streaming
High compression rate
EXP3000 JBOD Enclosures
12 x 3.5" 1TB, 7200RPM, SAS (3Gb/s)
max 116MB/s (200-500MB/s compressed data)
e.g. TF12:
8 enclosures → 96 HDDs
32TB uncompressed user data (→128TB)

IDAA Server
SQL Compiler, Query Plan, Optimize
Administration
2 front/end hosts, IBM 3650M3
clustered active-passive
2 Nehalem-EP Quad-core 2.4GHz per host

Processor & streaming DB logic
High-performance database
engine streaming joins, aggregations,
sorts, etc. e.g. TF12: 12 back/end SPUs

IBM Inside Sales
International Technical Support Organization Global Content Services

The Appliance Connected to a System z

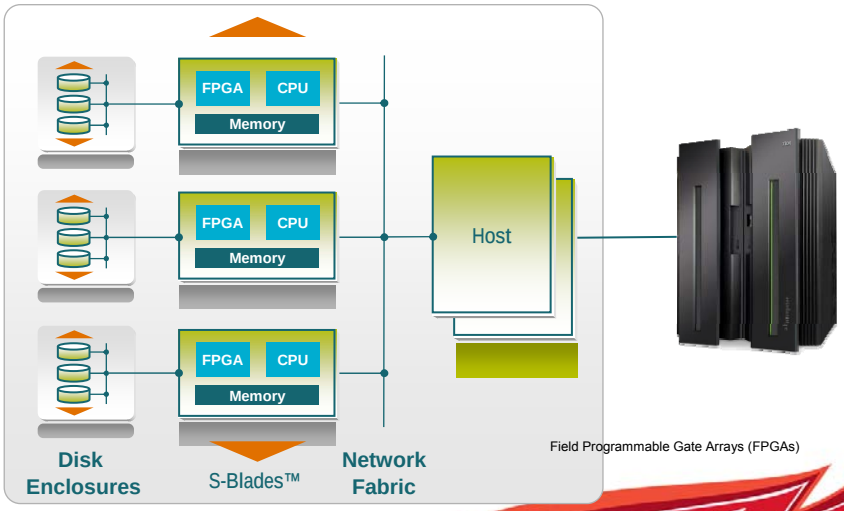

Netezza Appliance

The Key to the Speed

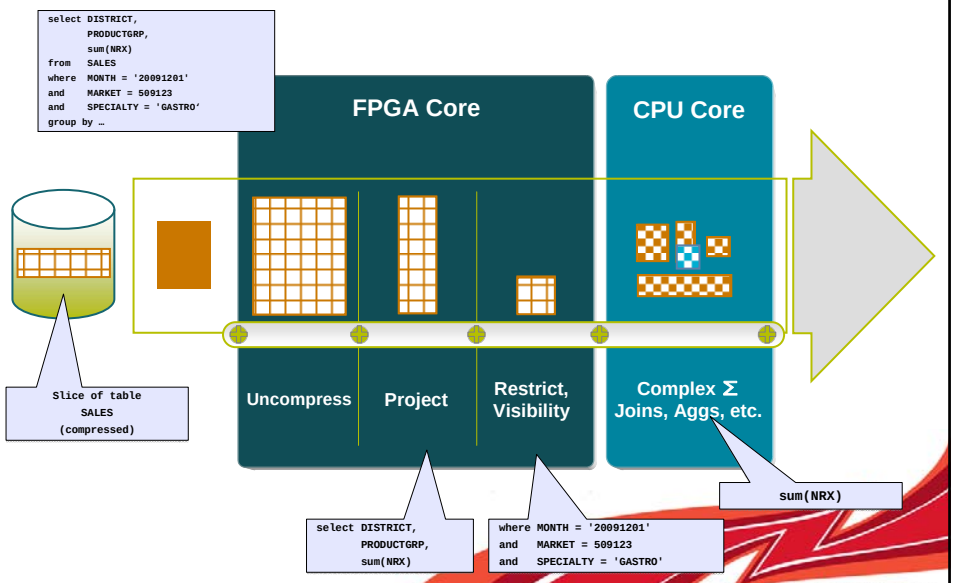

Why Both?

Marrying the best of each

IBM Netezza

Focused Appliance

IBM System z

Mixed Workload System

Capitalizing on the strengths of both platforms while driving to the most cost effective, centralized solution - destroying the myth that transaction and decision systems had to be on separate platforms

Very focused workload ⟷ **Very diverse workload**

Tailored to your needs

A Hybrid Solution

	IBM Netezza	IBM System z with IBM DB2 Analytics Accelerator
	Focused Appliance	Mixed Workload System
	• Appliance with a streamlined database and HW acceleration for performance critical functionality • Price/performance leader • Speed and ease of deployment and administration • Optimized performance for deep analytics, multifaceted, reporting and complex queries	• Mixed workload system z with operational transaction systems, data warehouse, operational data store, and consolidated data marts. • Unmatched availability, security and recoverability • Natural extension to System z to enable pervasive analytics across the organization. • Speed and ease of deployment and administration

Simplicity *The right mix of simplicity and flexibility* **Flexibility**

Business Analytics Life Cycle Architecture on System z

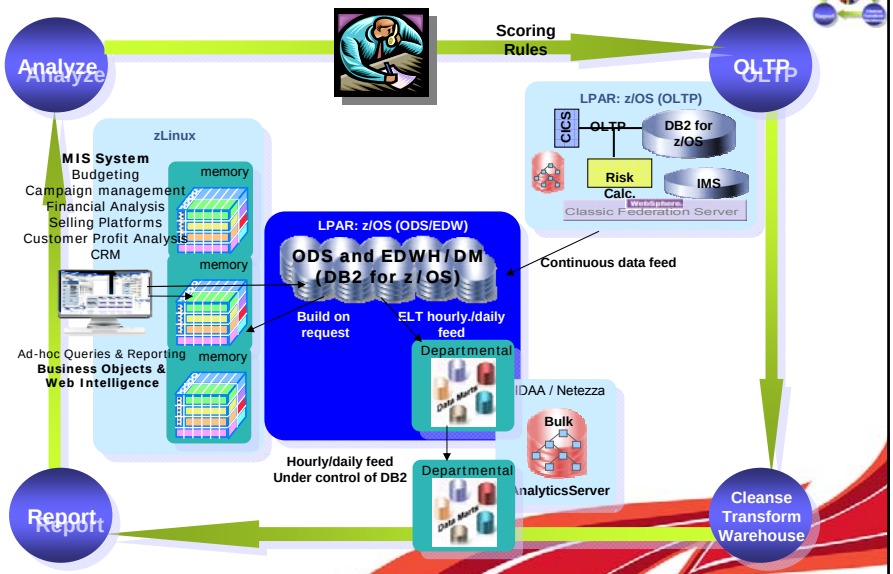

Workload-Optimized Systems

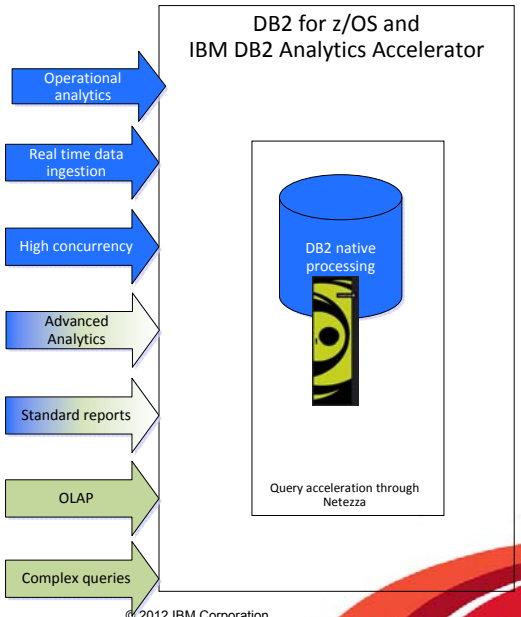

- Single and unique system for mixed query workloads
- Dynamic decision for most efficient execution platform
- Combines the strengths of both System z and Netezza
- Merging operational and data warehouse into a single optimized environment

Feasibility study and Value Assessment

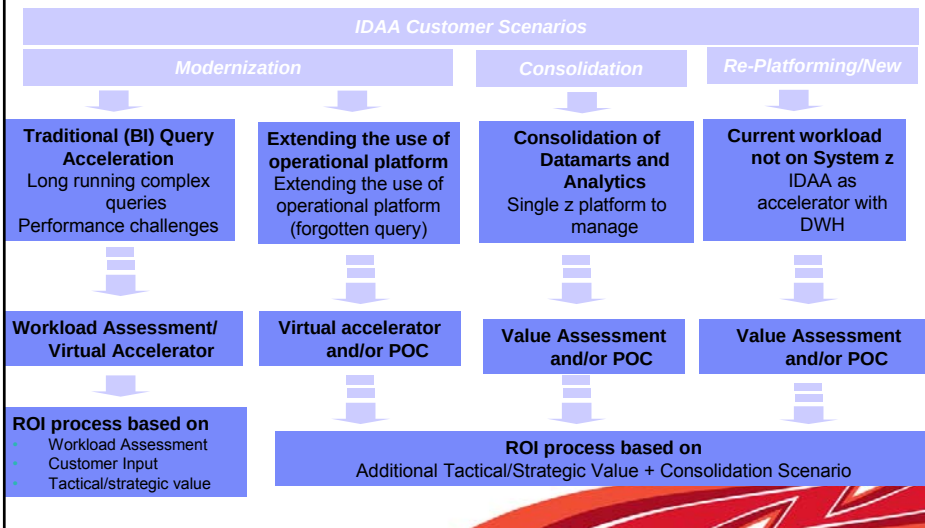

ROI Assessment Process for Great Outdoors Traditional BI Acceleration

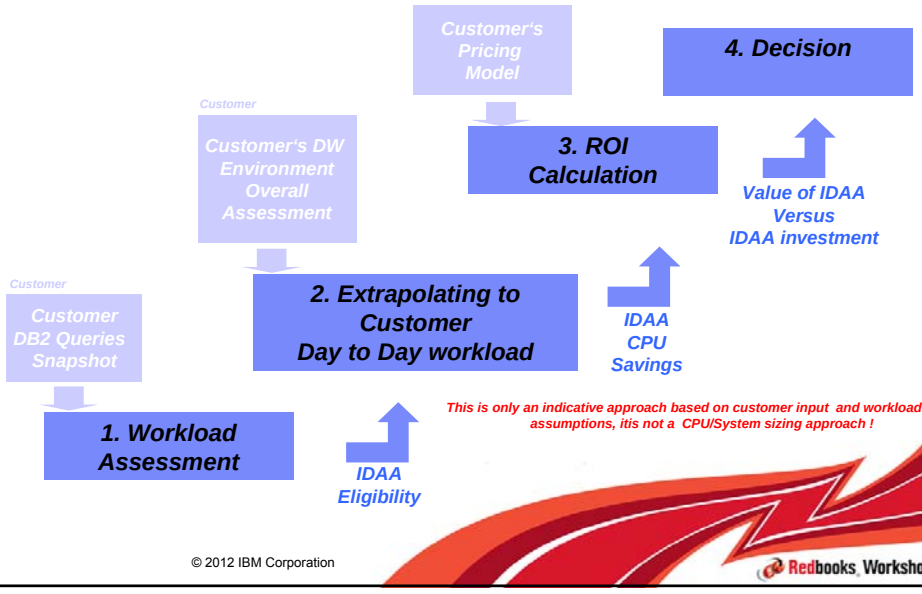

Virtual Accelerator (Explain only)

- To test workloads' access plan on a virtual accelerator, a real accelerator is not required.
- The results produced by virtual accelerators are written to EXPLAIN tables
  - make sure to create all the EXPLAIN tables including DSN_QUERYINFO_TABLE
- While the same information can also be obtained from real accelerators, virtual accelerators have the advantage that they do not require accelerator hardware.
  - You can thus check whether queries can be accelerated or not.
  - You can also calculate response time estimates without making extra demands on the IDAA hardware resources.
- You could install virtual accelerator with or without IDAA installed, and with or without DataStudio (IDAA Studio) installed.

Initial Installation

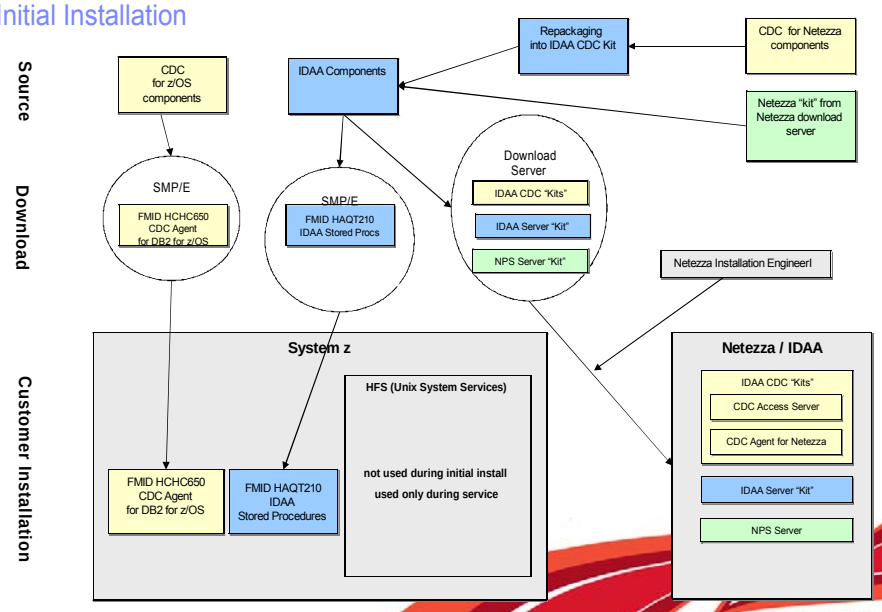

IBM Inside Sales
International Technical Support Organization Global Content Services

Network Connection Overview

- In general the Netezza hosts should be connected to the zEnterprise host with the usual HA concept, this means you need two 10GbE capable switches and at least two fiber cables from each host, for the zEnterprise you need two separate OSA cards (either OSD or OSX).

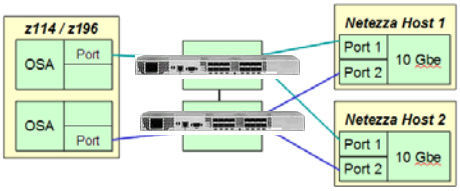

Netezza has SR ports only

The OSA cards on the zEnterprise

OSA Express3 LR or SR

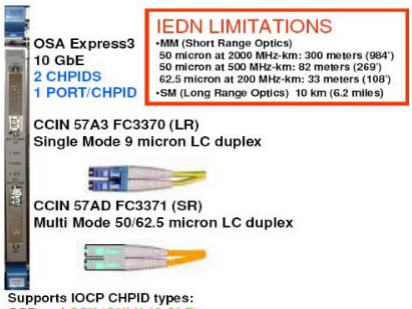

IBM zEnterprise Connection with zBX

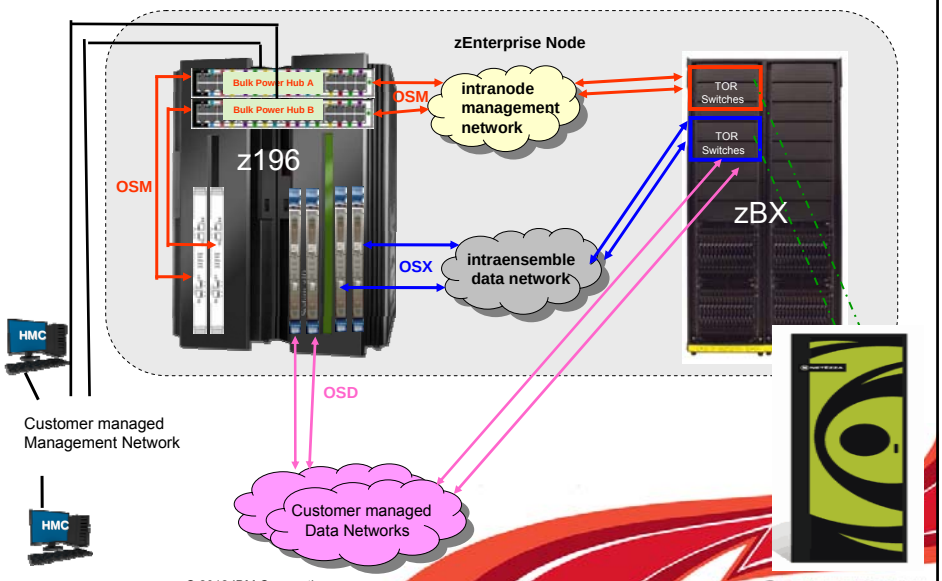

DB2 for z/OS and IDAA Configuration Overview

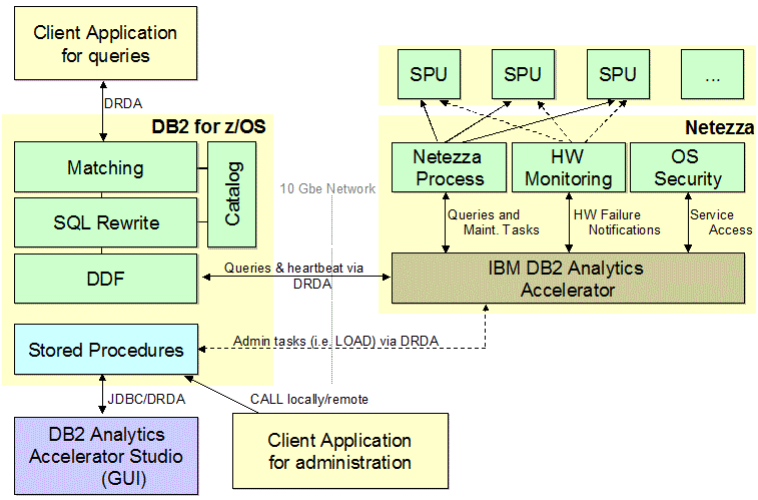

Settings

DB2

- Use SMP/E to install the required/enabling APARs/PTFs for your version of DB2 from the link below. http://www-01.ibm.com/support/docview.wss?uid=swg27022331
- Configure DB2 subsystem parameter for acceleration (i.e., set ZPARM ACCEL to COMMAND (or AUTO) and leave QUERY_ACCELERATION=NONE)
- If you are running DB2 9 for z/OS then in addition to the ACCEL DSNZPARM, you should also set the DSNZPARM ACCEL_LEVEL=V2.
- Customize and run DSNTIJAS installation job to create SYSACCEL tables.
- Create DSN_QUERYINFO_TABLE (new EXPLAIN table)
- Setup suitable WLM environment for the IDAA Stored Procedures

Network

- The network configuration of the SMP hosts and the wall IP is performed during installation by an IBM engineer.
- Configuring the OSA cards with VIPA is the responsibility of the client. See http://www.ibm.com/support/docview.wss?uid=swg27023654

Stored Procedure Components

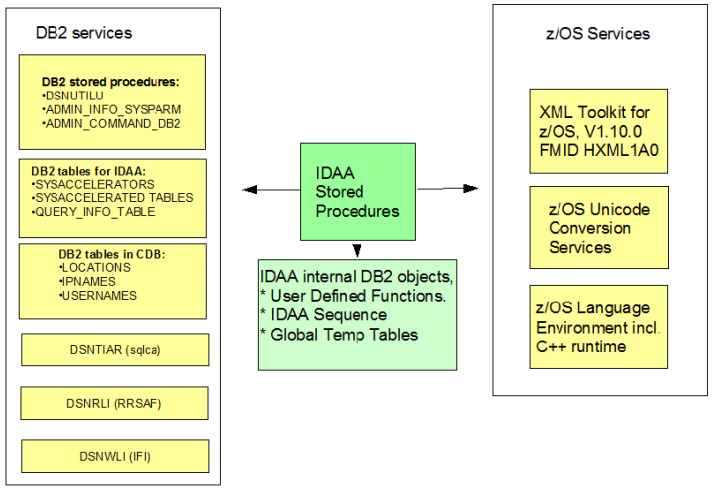

New/Enhanced DB2 Commands

- **The following DB2 commands have been introduced specifically to support accelerator.**
  - **-START ACCEL**
  - **-STOP ACCEL**
  - **-DISPLAY ACCEL**
- **The following DB2 commands have been enhanced to support accelerator.**
  - **-DISPLAY THREAD**
  - **-CANCEL THREAD**
  - **-DISPLAY LOCATION**
  - **-START PROFILE**
  - **-STOP PROFILE**

Caution: When you stop DB2 or DDF, the accelerator will be stopped too

IBM Inside Sales International Technical Support Organization Global Content Services

Display Accel-Command "-DIS ACCEL(DEMOACC) DETAIL“

```
DSNX810I  -D914 DSNX8CMD DISPLAY ACCEL FOLLOWS -
DSNX830I  -D914 DSNX8CDA
ACCELERATOR                     MEMB  STATUS  REQUESTS ACTV QUED MAXQ
------------------------------- ---- -------- -------- ---- ---- ----
EMU02                           D914 STARTED        53    0    0    0
LOCATION=EMU02 HEALTHY
DETAIL STATISTICS
  LEVEL  = AQT02010
  STATUS = ONLINE
  FAILED QUERY REQUESTS                              =        7
  AVERAGE QUEUE WAIT                                 =      116
  MAXIMUM QUEUE WAIT                                 =      315
  TOTAL NUMBER OF PROCESSORS                         =       20
  AVERAGE CPU UTILIZATION ON COORDINATOR NODES       =     1.00%
  AVERAGE CPU UTILIZATION ON WORKER NODES            =     2.00%
  NUMBER OF ACTIVE WORKER NODES                      =        5
  TOTAL DISK STORAGE AVAILABLE                       =  6564825 MB
  TOTAL DISK STORAGE IN USE                          =      668 MB
  DISK STORAGE IN USE FOR DATABASE                   =    38892 MB
TF12                            D914 STOPPED         0    0    0    0
TF3                             D914 STARTED         0   12    0    0
LOCATION=TF3 HEALTHY
DETAIL STATISTICS
***
```

© 2012 IBM Corporation

Redbooks Workshop

WLM Configuration

- WLM and Application environments

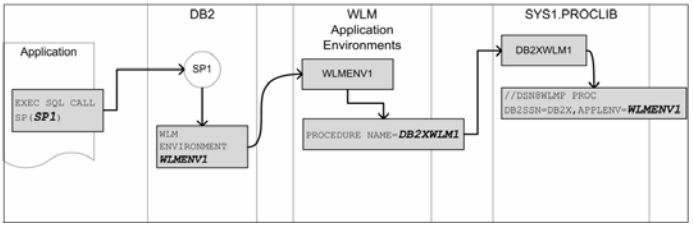

Suitable WLM setup is required for the IDAA DB2 stored procedures.

WLM configuration

- To avoid conflicts with environment variables that are set for stored procedures of other applications, use a dedicated WLM Application Environment for the IBM DB2 Analytics Accelerator for z/OS stored procedures
- The DB2-supplied stored procedures SYSPROC.ADMIN_INFO_SYSPARM, SYSPROC.DSNUTILU, and SYSPROC.ADMIN_COMMAND_DB2 must use a WLM environment that is separate from the one used by the IBM DB2 Analytics Accelerator for z/OS stored procedures. Verify that this and a few other requirements are met by following the steps here.
  - Verify that each, SYSPROC.ADMIN_INFO_SYSPARM, SYSPROC.DSNUTILU and SYSPROC.ADMIN_COMMAND_DB2, use a separate WLM environment.
  - Make sure that NUMTCB is set to 1 (NUMTCB=1) for the SYSPROC.ADMIN_INFO_SYSPARM and SYSPROC.DSNUTILU WLM environments.
  - WLM routes work to stored procedure address spaces based on the application environment name and service class associated with the stored procedure. The service class is assigned using the WLM classification rules. Stored procedures inherit the service class of the caller.
  - There is no separate set of classification rules for stored procedures.

Stored Procedures for Maintenance

ACCEL_ADD_ACCELERATOR	Pairing an accelerator to a DB2 subsystem
ACCEL_TEST_CONNECTION	Check of the connectivity from DB2 procedures to the accelerator
ACCEL_REMOVE_ACCELERATOR	Removing an accelerator from a DB2 subsystem and cleanup resources on accelerator
ACCEL_UPDATE_CREDENTIALS	Renewing the credentials (authentication token) in the accelerator
ACCEL_ADD_TABLES	Add a set of tables to the accelerator
ACCEL_ALTER_TABLES	Alter table definitions for a set of tables on the accelerator (only distribution and organizing keys)
ACCEL_REMOVE_TABLES	Remove a set of tables from the accelerator
ACCEL_GET_TABLES_INFO	List set of tables on the accelerator together with detail information
ACCEL_LOAD_TABLES	Load data from DB2 into a set of tables on the accelerator
ACCEL_SET_TABLES_ACCELERATION	Enable or disable a set of tables for query off-loading
ACCEL_CONTROL_ACCELERATOR	Controlling the accelerator tracing, collecting trace and detail of the accelerator (software level etc.)
ACCEL_UPDATE_SOFTWARE	Update software on the accelerator (transfer versioned software packages or apply an already transferred package, new: also list software both on z/OS and accelerator side)
ACCEL_GET_QUERY_DETAILS	Retrieve statement text and query plan for a running or completed Netezza query
ACCEL_GET_QUERY_EXPLAIN	Generate and retrieve Netezza explain output for a query explained by DB2
ACCEL_GET_QUERIES	Retrieve active and/or history query information from accelerator

1

IBM Inside Sales International Technical Support Organization Global Content Services

IBM DB2 Analytics Accelerator Studio positioning

- IDAA Studio:
  - Provides basic functionality included with the product for required usability functions
  - Available for Windows and Linux (check announcement letter for details of supported distributions)
  - Based on DataStudio 2.2.1 Standalone, DVD includes DataStudio 2.2.1
  - No additional charge – can be plugged into several Eclipse-based IBM products (DataStudio 2.2.1+, Optim Query Tuner 2.2.1+Optim Database Administrator)
  - Later Upgraded to DataStudio 3.1

Installation options

- 1) IDAA DVD or downloaded DVD image:
  - IBM DataStudio 2.2.1 (standalone) + IDAA Studio
  - one common InstallAnywhere installer to Install IDAA Studio
- 2) Eclipse Update Manager
  - Use Eclipse update site to install on top after installing any of the following supported base products
  - http://public.dhe.ibm.com/ibmdl/export/pub/software/data/db2/analytics-accelerator-studio/
    - IBM DataStudio 2.2.1 or 2.2.1.1 standalone
    - IBM DataStudio 2.2.1.1 IDE
    - IBM Optim Query Tuner 2.2.1 shell sharing with IBM Optim Database Administrator 2.2.3 (or either one shell sharing with DS IDE)

IBM DB2 Analytics Accelerator Studio function...

- Connection management
  - establish a connection from DataStudio to DB2 for z/OS (all administration is done thru DB2 for z/OS)
- Accelerator administration
  - establish a connection between an accelerator and a DB2 for z/OS subsystem (pairing)
  - start/stop acceleration for an accelerator in DB2
  - display accelerator status (like number of requests, online status)
  - transfer SW updates for accelerator and Netezza
  - activate SW updates for accelerator
- Define the data to load into the accelerator
  - define tables to be loaded from DB2 into the accelerator
  - define distribution and organizing keys for these tables
  - load the tables
  - update the tables
  - toggle the acceleration status of tables

IBM DB2 Analytics Accelerator Studio function...

- Monitoring and SQL execution
  - Display a list of active queries (including query properties and statement text)
  - Display the query history
  - Display the plan for a query
  - Execute a query from SQL script editor - SQL results view
  - Explain a query (visual explain for the DB2 part (dummy block) and the Netezza execution plan) to understand the impact of the choice of distribution and organizing keys
- Diagnostics
  - Configure trace
  - Collect trace (including Eclipse error log, IDAA SP, accelerator, software versions of all components)
  - Upload trace to PMR

IBM DB2 Analytics Accelerator Studio Architecture

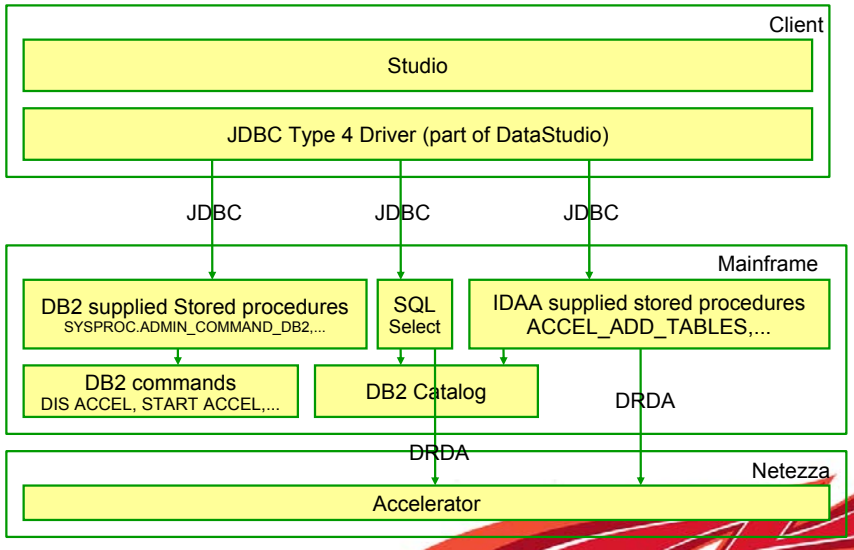

Description of User Interface Elements

- When you start IBM DB2 Analytics Accelerator Studio for the first time, you see the Welcome screen. To see the main window of IBM DB2 Analytics Accelerator Studio with the principal functions and controls, just close the Welcome screen by clicking the cross-shaped icon on the tab of the Welcome screen (usually positioned on the upper left).
- The Administration Explorer shows your database connection profiles and various objects like schemas, tables, and stored procedures that are defined in the database catalogs of your data servers.
  - you can create a connection profile for a DB2 Subsystem
  - you can open the SQL script editor to run or explain SQL statements for a defined connection
- The Object List Editor contains a table with an entry for each accelerator that has been configured. The Object List Editor opens when you select the Accelerators folder in the Administration Explorer.
- To display the tables and other accelerator details in the Accelerator view, double-click an accelerator name in the Object List Editor.

Latency Management

- Data in IDAA is available while replacing
- CDC process is available to reduce data latency
  - New partitions can be added without impacting old parts
- Stored procedures can also be used for automating various IDAA processes like LOAD in batch mode as well as using Data Studio
- Existing ETL or ELT process can be extended to integrate IDAA load process

Primary reasons why a query may not be routed

- Because it uses static SQL or a plan with DBRMs
- Because it uses CURRENT QUERY ACCELERATION = NONE (default)
- Because it is not read-only or the cursor is scrollable or rowset cursor
- Because it contains syntax that is not supported
- Because the accelerator or tables are disabled
- Because it references a table or column that is not loaded / enabled
  - (may be due to unsupported datatypes)
- Because the optimizer decides DB2 for zOS can do better
  - DB2 has short query heuristics to keep OLTP in DB2
    - e.g., Query with equality predicate on unique-indexed column is executed in DB2

DB2 decision to offload query

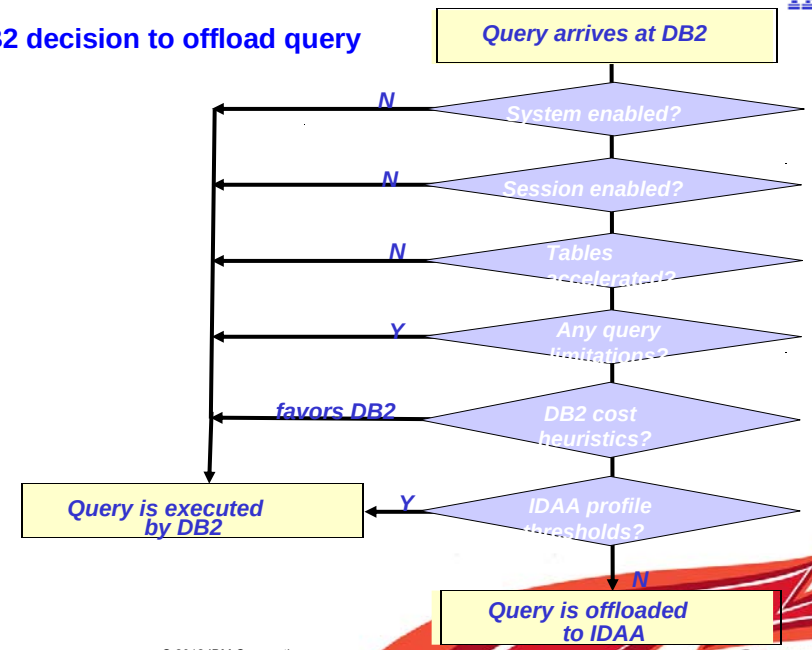

Access Plans

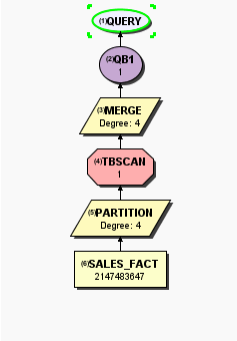

SET CURRENT QUERY ACCELERATION=NONE

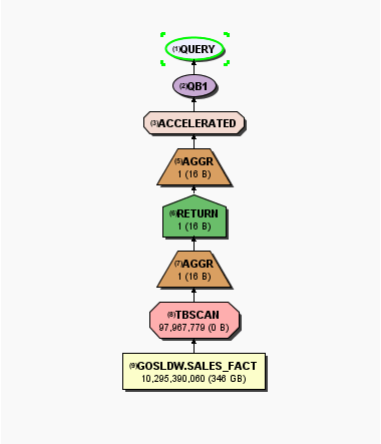

SET CURRENT QUERY ACCELERATION=ENABLE

Query Execution Flow

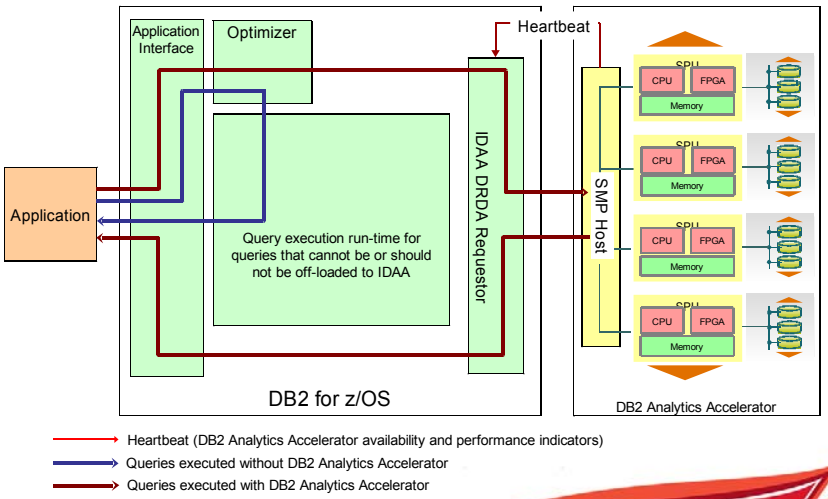

1

Distribution and organizing keys

- In general can use defaults (random/balanced distribution and no organizing keys)
- For large tables we have seen improvement by orders of magnitude by specifying organizing and or distribution keys

Distribution key:

- Determines how data is partitioned across worker nodes (data skew)

Organizing key:

- Determines the data sequence on disk within each node

Best practices for selecting distribution keys

- Use the following rules when selecting a unique or non-unique column as the distribution key for the table:
  - Use columns for the distribution key that **distribute table rows evenly** across the data slices. The more singular the values for a column, the more optimal their distribution.
    - e.g., do not use boolean keys, for example, True/False, I/0, or Y/N, because the system distributes rows with the same hash value to the same data slices; thus, the table would be divided across only two data slices.
  - For joins, base the column selection on an equality search, because if both tables distribute on the equality columns, **the system can perform the join operation locally.**
  - Select as few columns as possible for the distribution key to optimize the generality of the selection.

Best practices for organizing keys

- Review predicates on large tables and use columns referenced by (selective) predicates as organizing keys
- As a best practice, review the design and columns of your large fact tables and the types of queries that run against them. If you typically run queries on one dimension, such as date, you can load the data by date to take advantage of the zone maps. If you typically query a table by two dimensions, such as by storeId and customerID for example, then two organizing keys can help to improve the query performance against that table.
- The zone map has low key value and high key value for each 3 MB block similar to page range scanning functionality in DB2 during predicate evaluation.

Altering distribution key from Data Studio

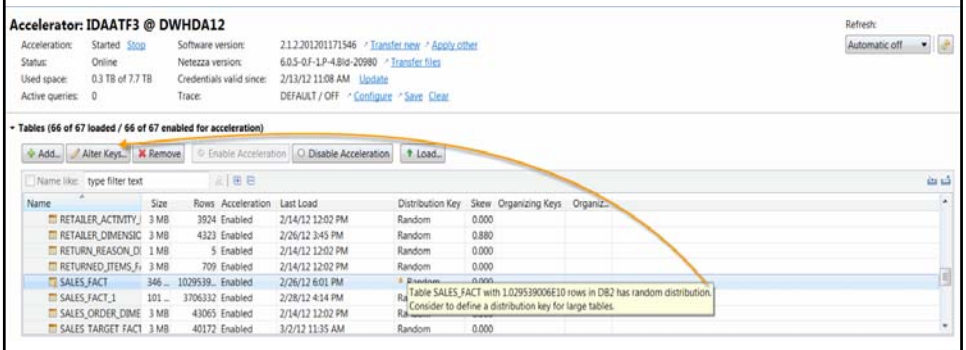

Altering distribution and organizing keys

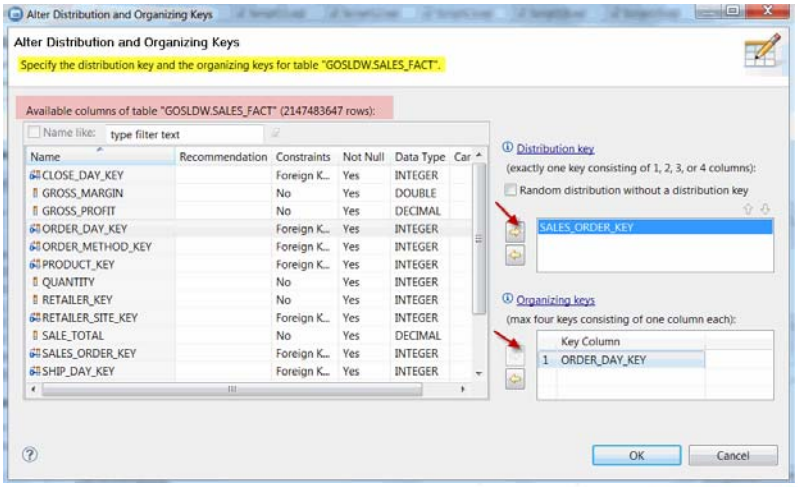

Access path before/after altering distribution key

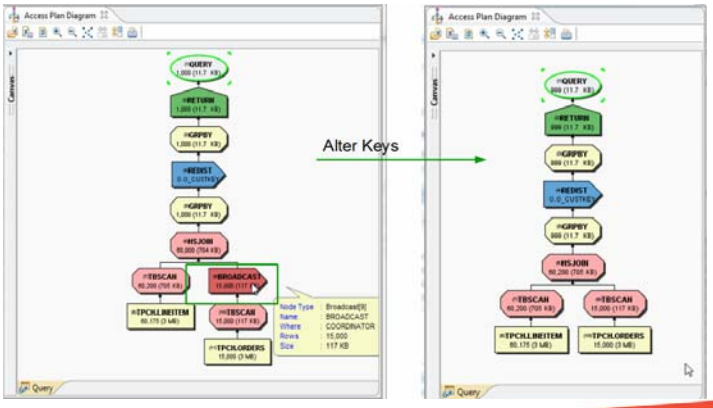

Scan Reduction with Clustered Base Tables (Organize on)

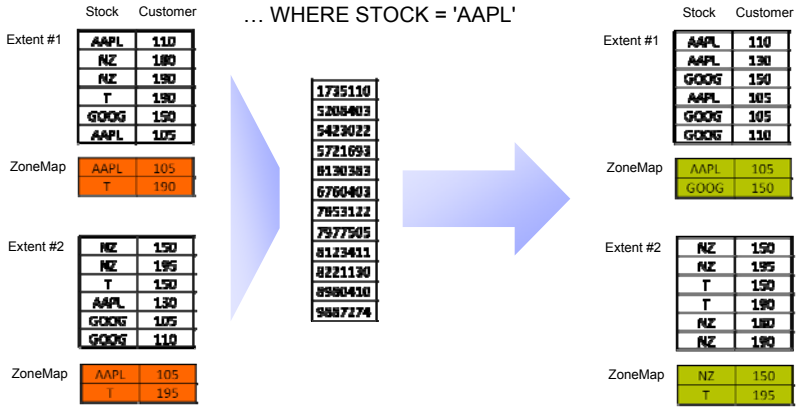

- **Multi-dimensional values mapped to single dimension for clustering**
- **Statistical clustering minimizes range of ZoneMap values for each extent**
- **Reduces scans without expanding storage requirements**

Security

- Data maintenance and security administration is still on DB2 for z/OS
- Support for tables with “AUDIT ALL”
- Private network
- IDAA functions can be restricted to selected users/group by restricting access to the appropriate stored procedures.
- Restrictions:
  - Editproc
  - MLS
  - Encrypt function

Query Monitoring option from Data Studio

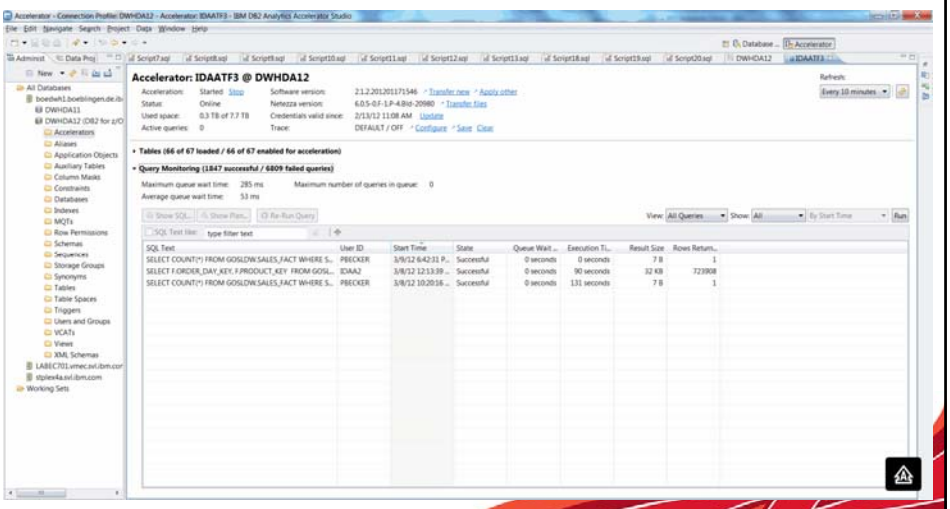

OMPE Accounting fields

```
Q8AC DS 0D
Q8ACNAME_OFF DS XL2 ACCELERATOR SERVER ID OFFSET
Q8ACPRID DS CL8 ACCELERATOR PRODUCT ID
Q8ACCONN DS XL4 # OF ACCELERATOR CONNECTS.
Q8ACREQ DS XL4 # OF ACCELERATOR REQUESTS.
Q8ACTOUT DS XL4 # OF TIMED OUT REQUESTS.
Q8ACFAIL DS XL4 # OF FAILED REQUESTS.
Q8ACBYTS DS XL8 # OF BYTES SENT.
Q8ACBYTR DS XL8 # OF BYTES RETURNED.
Q8ACMSGS DS XL4 # OF MESSAGES SENT.
Q8ACMSGR DS XL4 # OF MESSAGES RETURNED.
Q8ACBLKS DS XL4 # OF BLOCKS SENT
Q8ACBLKR DS XL4 # OF BLOCKS RETURNED.
Q8ACROWS DS XL8 # OF ROWS SENT
Q8ACROWR DS XL8 # OF ROWS RETURNED.
Q8ACSCPU DS XL8 ACCELERATOR SERVICES CPU TIME.(V1only)
Q8ACSELA DS XL8 ACCELERATOR SERVICES ELAPSED TIME.(V1)
Q8ACTCPU DS XL8 ACCELERATOR SVCS TCP/IP CPU TIME.
Q8ACTELA DS XL8 ACCELERATOR SVCS TCP/IP ELAPSED TIME.
Q8ACACPU DS XL8 ACCUMULATED ACCELERATOR CPU TIME.
Q8ACAELA DS XL8 ACCUMULATED ACCELERATOR ELAPSED TIME.
Q8ACAWAT DS XL8 ACCUMULATED ACCELERATOR WAIT TIME.
Q8ACEND DS 0F
```

Accounting Elapsed Times

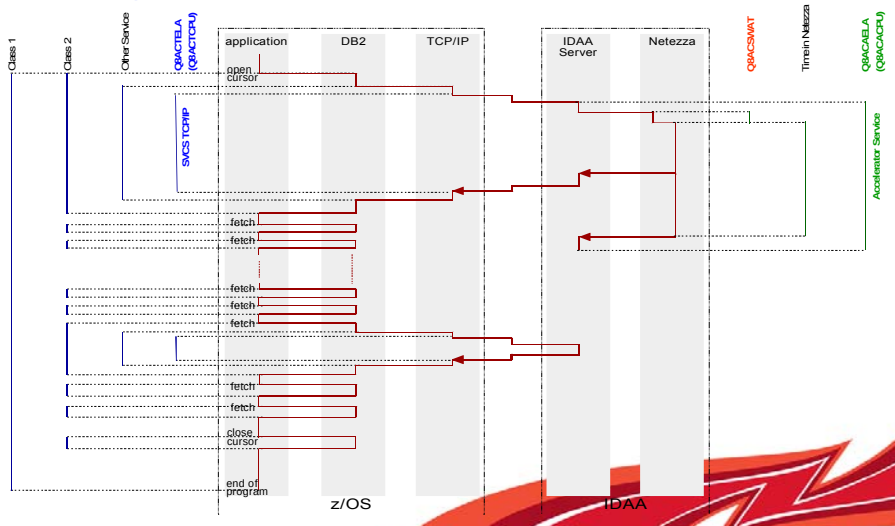

Workload performance comparison (GO W/L)

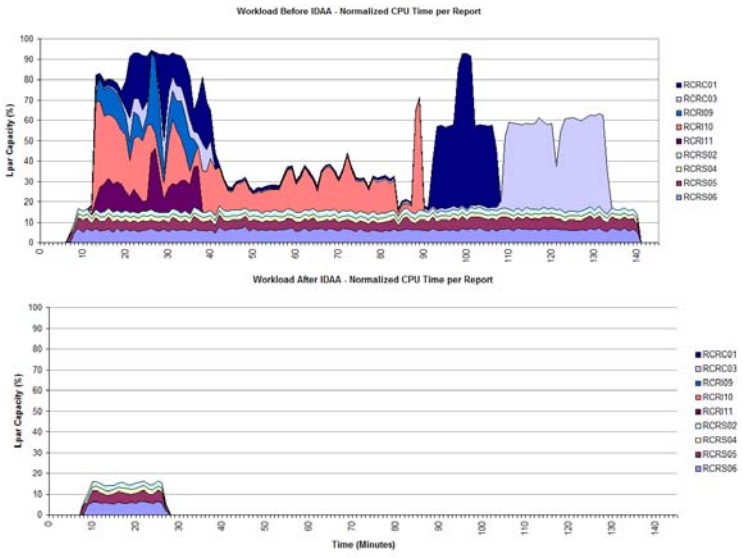

From SG24-8005

The results of the workload

Table 13-5 shows the response times of nine reports in our tests. The first five reports ran in DB2 (before IDAA) as well as in IDAA. The next four reports were short reports and were not directed to run in IDAA. They ran in DB2 in both test runs.

Table 13-5 Workload scenario: Elapsed Time before and after IDAA per report

Report	ET Before IDAA (secs)	ET After IDAA (secs)
RC01	1,382.10	105.48
RC03	2,294.14	134.46
RI09	283.50	64.22
RI10	764.54	112.90
RI11	294.48	128.83
RS02	1.98	1.88
RS04	0.04	0.06
RS05	11.84	12.20
RS06	4.10	4.00

These results were obtained from concurrency measurements. It is expected that the improvement ratios will be even more impressive in single query measurements. Since all of

Customer: Large insurance company

Adding value by Accelerating the Delivery of Business Reporting

				DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Qualifying Rows	Total Rows Returned	Hours	Sec(s)	Hours	Sec(s)	
Query 1	591,941,065	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	591,941,065	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	813,343,052	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	283,105,125	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	591,941,089	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	813,343,052	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	591,941,065	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	813,343,052	3,425,292	724	0:44	2,640	0.0	2	1,320
Query 9	813,343,052	4,130,107	137	0:42	2,520	0.1	193	13

With Accelerated Time to Value

- IBM DB2 Analytics Accelerator (Netezza 1000-12)
  Production ready - 1 person, 2 days
- Table Acceleration Setup in 2 Hours
  - DB2 "Add Accelerator"
  - Choose a Table for "Acceleration"
  - Load the Table (DB2 Loads Data to the Accelerator)
  - Knowledge Transfer
  - Query Comparisons
- Initial Load Performance
  400 GB Loaded in 29 Minutes
  570 Million Rows (Actual: Loaded 800 GB to 1.3 TB per hour)
- Extreme Query Acceleration - 1908x faster
  2 Hours 39 minutes to 5 Seconds
- CPU Utilization Reduction
  35% to ~0%

Customer Quote: "we had this up and running in days with queries that ran [illegible] times faster"

Real World Results – Major Financial Services Company

- Netezza 1000-6 set up within two days
- Day 1: Netezza Hardware installation and configuration
- Day 2: connection to z/OS and pairing DB2 with accelerator
- Tables loaded and first queries executed on day 2

First Impression

- Time to value
- Add tables, load first tables – first Queries accelerated all within one hour
- Easy to use

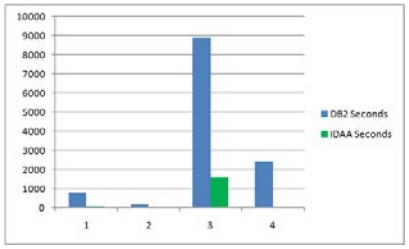

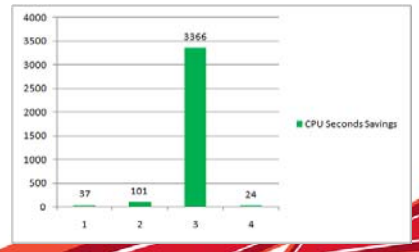

Real World Results – Major Healthcare Insurer

Query	Times Faster on IDAA
1	1789 x
2	1643 x
3	1193 x
4	1448 x
5	931 x
6	816 x
7	45 x
8	712 x
9	677 x
10	549 x

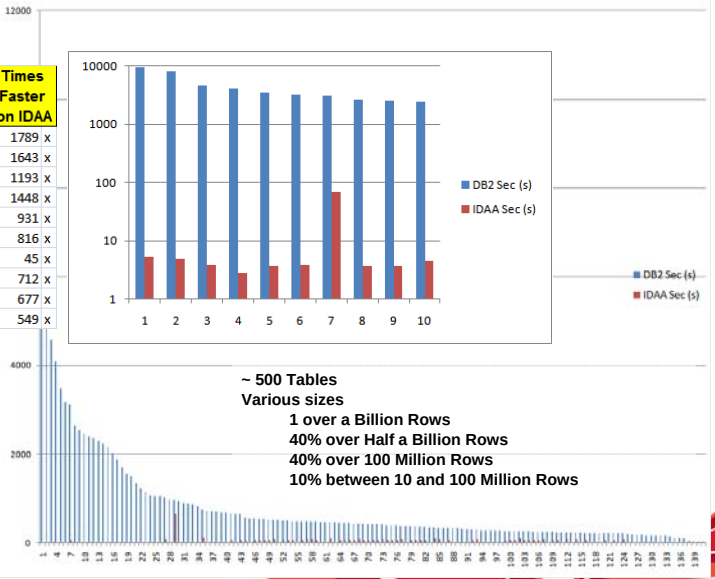

~ 500 Tables
Various sizes
1 over a Billion Rows
40% over Half a Billion Rows
40% over 100 Million Rows
10% between 10 and 100 Million Rows

2-way Data Sharing configuration

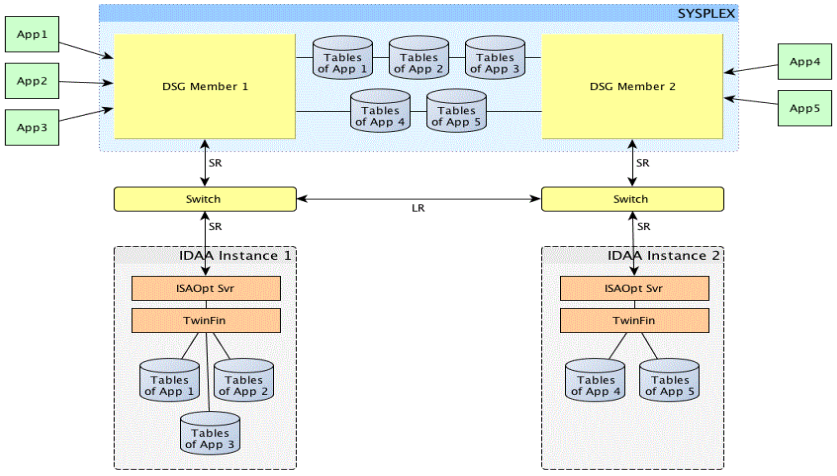

Failover Scenario #1 (IDAA tables being loaded after failure)

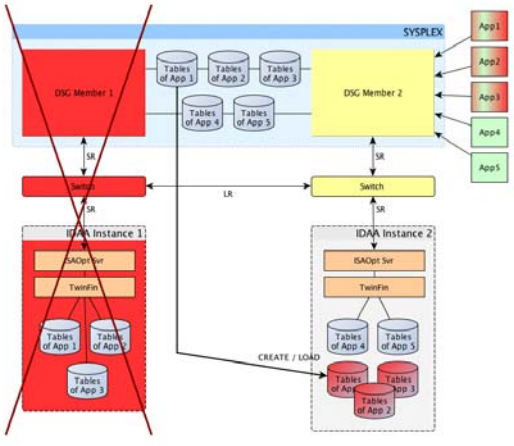

Failover Scenario

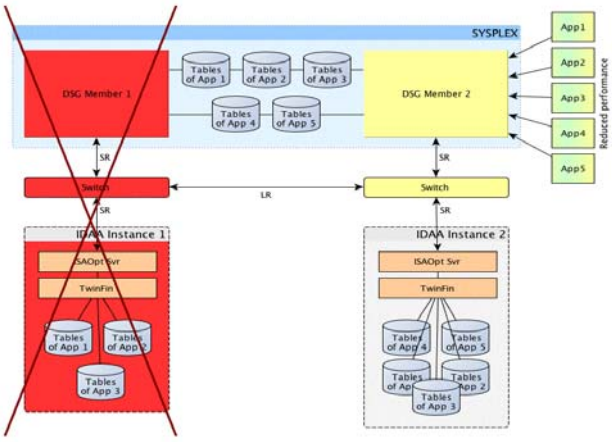

Failover scenario#2 (Tables loaded and ready but enabled after failure)

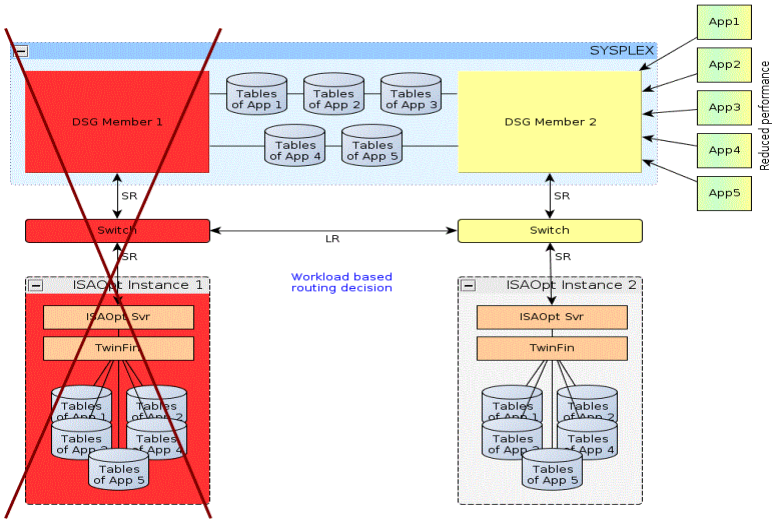

What's new?

- Enhancements via PTFs
- Announcement of IDAA V3.1
- zEC12 support
- 9700

Major enhancements via PTFs (idaav2r1/k)

PM60921
- Option to allow the fullselect portion of an INSERT statements to execute on IDAA. Enabled by a new option 2 added to DB2 subsystem parameter DSN6SPRM.QUERY_ACCEL_OPTIONS
- Use of UNLOAD Utility with INTERNAL format for decreasing CPU consumption in ACCEL_LOAD_TABLES
- Support of LONG VARCHAR, LONG VARGRAPHIC types in tables created prior to DB2 8
- Support of tables using EDIT PROCEDURES
- Limited support of mbcs string columns and graphic columns that are not encoded in UNICODE
- Support of CDC if installed on Netezza host (controlled availability)

PM57960
- The accelerator server was improved: Support for large query result set
- INTERNAL format of the DB2 Unload Utility
- Date and time format LOCAL of a DB2 subsystem is supported

PM60626
- Improvements and restructuring in physical catalog storage structure on the accelerator server. This allows future migration of server versions in both directions.
- Spill to disk on IBM DB2 Analytics Accelerator server
- Supporting multi-byte EBCDIC data encoding
- Enabling capability to cancel statements being executed in IBM DB2 Analytics Accelerator

PM60170
- Multi-byte character set EBCDIC encoding scheme
- DDL is added to create DSN_QUERYINFO_TABLE explain table in job DSNTIJSG and DSNTESC
- Multiple accelerator query offloading is added
- CPU reduction for query PREPARE when a table is not defined on the accelerator

IBM DB2 Analytics Accelerator V3.1 - Highlights

Further strengthen DB2 for z/OS competitive position to host mission critical operational BI workloads

- **High Performance Storage Saver**
  - Store a DB2 table or partition of data solely on the Accelerator. Removes the requirement for the data to be replicated on both DB2 and the Accelerator
- **Incremental Update**
  - Enables tables within the Accelerator to be continually updated throughout the day
- **UNLOAD Lite**
  - *Reduces z/OS MIPS consumption, by moving the preparation off System z*
- **zEnterprise EC12 Support**
  - Version 3 will support the zEnterprise EC12, z196 and z114 System z platforms z
- **Query Prioritization**
  - Brings System z workload management down to the individual query being routed to the Accelerator Platform Support
- **High Capacity**
  - Support has been extended to include the entire Netezza 1000 line *(1.28 PB)*

- General Availability: November 30th, 2012
- Beta program started July 2012
- Prerequisites: zEnterprise (z196, z114, zEC12), DB2 10

High Performance Storage Saver

Reducing the cost of high speed storage

Store historic data on the Accelerator only

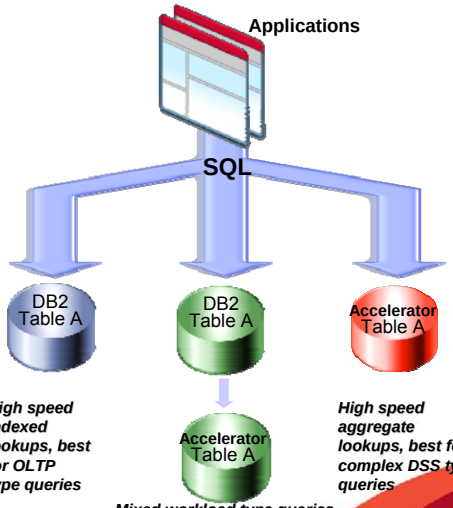

Tables can be resident on:

1. **DB2 Only**
2. **DB2 and Accelerator**
3. **Accelerator Only**

When data no longer requires updating, reclaim the DB2 storage

Special Registers control behavior

- **CURRENT QUERY ACCELERATION**
- **CURRENT GET_ACCEL_ARCHIVE**

Managed by zParms

Storage options to match data needs

Optimized in both price and performance for differing workloads

High Performance Storage Saver	Database Resident Partitions
Single Disk Store	**Dual Disk Store**
• Only stored on Accelerator storage (Less Cost) • Optimized performance for deep analytics, multifaceted, reporting and complex queries • Only full table update or full partition update from backup • Same high speed query access transparently through DB2	• Stored on both DB2 and Accelerator storage • Mixed query workload with transactions, single record queries and record updates with deep analytics, multifaceted, reporting and complex queries. • Full table, full partition update, Incremental update from DB2 data • Same high speed query access transparently through DB2

Cost *The right mix of cost and functionality* **Functionality**

Incremental Update

- **An alternative to a full table load or table partition load. Refreshes only the records of the table that have been recently modified in the data warehouse.**
  - This capability keeps the data on the DB2 Analytics Accelerator in sync with the data in the DB2 Analytic System.
  - Used for low volume highly critical information. Like CDC between the operation system and data warehouse some delay is present.
  - This is a fully integrated function that provides incremental update to tables and partitions.

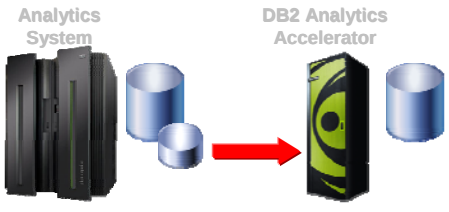

Incremental Update

Now expandable to 960 cores and 1.28 petabytes

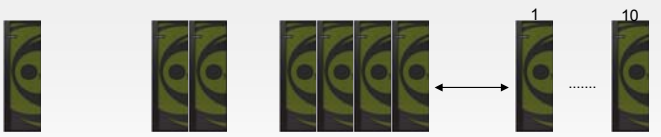

	TF3	TF6	TF12	TF24	TF36	TF48	TF72	TF96	TF120
Cabinets	1/4	1/2	1	2	3	4	6	8	10
Processing Units	24	48	96	192	288	384	576	768	960
Capacity (TB)	8	16	32	64	96	128	192	256	320
Effective Capacity (TB)*	32	64	128	256	384	512	768	1024	1280

Accelerator Platforms →

Predictable, Linear Scalability throughout entire family

Capacity = User Data space
Effective Capacity = User Data Space with compression

*4X compression assumed

Query Prioritization

Extending the value of System z workload management

Various priority end-user applications

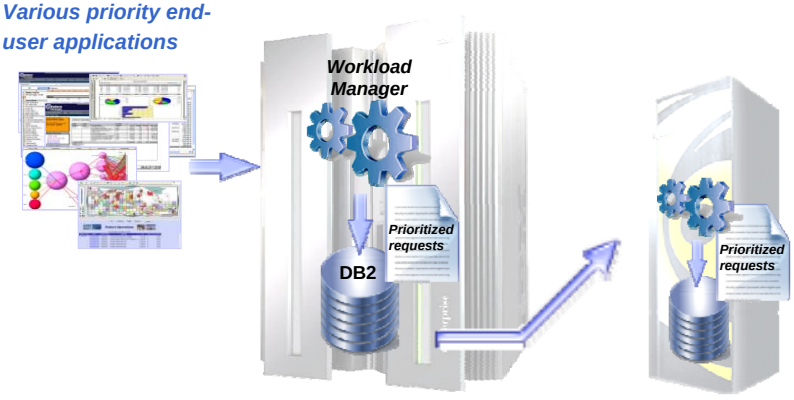

IBM Inside Sales

International Technical Support Organization Global Content Services

IBM

IBM zEnterprise Analytics System 9700

Flexibility in Critical Data Decision Systems

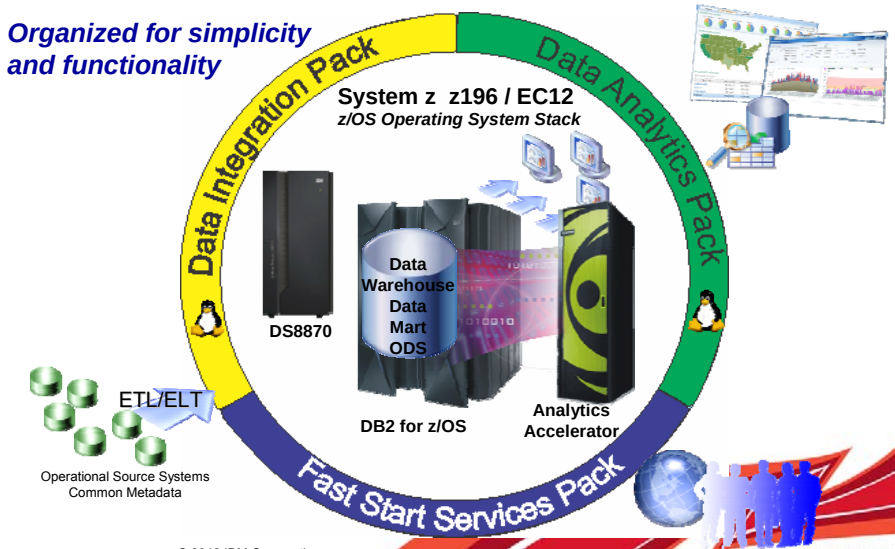

Redbooks Workshop

IBM Inside Sales

International Technical Support Organization Global Content Services

IBM

IBM zEnterprise Analytics System 9710

Cost Effective Critical Data Decision Systems

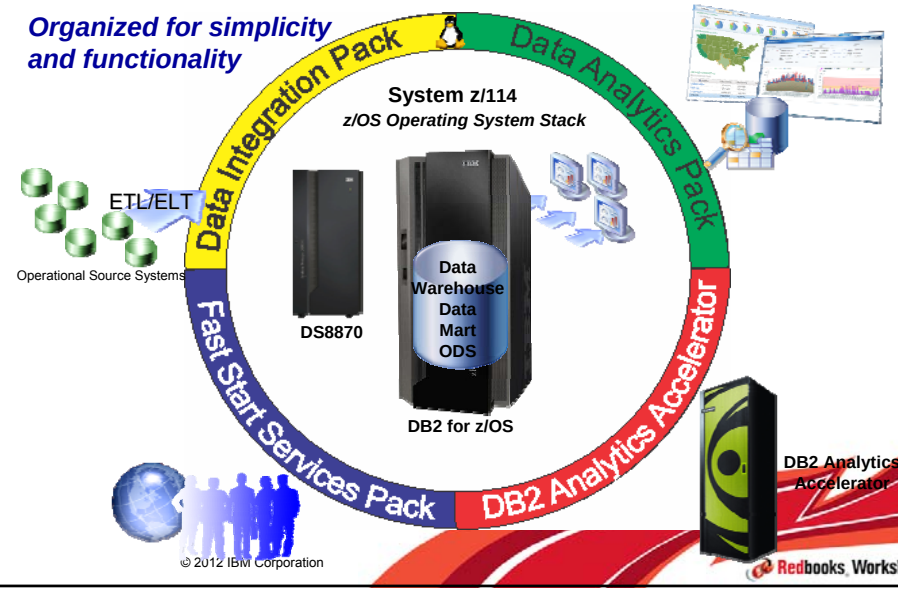

Redbooks Workshop

https://www.ibm.com/connections/blogs/basupportlink/entry/cut_the_cost_and_complexity_of_business_analytics4?lang=en_us

Cognos 10.2 for z/OS

Products
- Framework Manager (MS Windows only)
- Transformer User Interface (MS Windows only)
- Cognos Insight (MS Windows only)
- Cognos SDK
- Cognos Samples

Features
- Authentication via RACF
- z/OS Console Support
- Cognos Portal/Viewer
- Cognos Workspace
- Report/Query/Analysis Studios
- Event Studio / Human Task Service
- Cognos Mashup
- Active Reporting
- External Data within a Report
- Integrated Search Capability
- Data Access via DQM
- IBM Cognos Data Lineage Viewer
- Transformer Engine
- Installation via SMP/E
- Usage Based Pricing Support

Themes/Content

User Experience
- Interactive Discovery
- Accessibility

Performance & Scale
- Dynamic Cubes (DB2 z/OS only)
- Content Manager Scalability

Flexible Deployment
- Enhanced Data Sources
  - DB2 for z/OS V9 & 10
  - Teradata 14
  - Oracle 11g R2
  - Microsoft SQL Server 2011
  - Cognos PowerCube V8.X & V10.2
  - IMS V12
- BiDi Support

Ann 10/16
GA 10/19

Some key Redbooks

- Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS
  - http://www.redbooks.ibm.com/abstracts/sg248005.html
- The Netezza Data Appliance Architecture: A Platform for High Performance Data Warehousing and Analytics, REDP-4725
  - http://www.redbooks.ibm.com/abstracts/redp4725.html
- Complete Analytics with IBM DB2 Query Management Facility: Accelerating Well-Informed Decisions Across the Enterprise
  - http://www.redbooks.ibm.com/abstracts/sg248012.html
- Enterprise Data Warehousing with DB2 9 for z/OS
  - http://www.redbooks.ibm.com/abstracts/sg247637.html
- Co-Locating Transactional and Data Warehouse Workloads on System z
  - http:// www.redbooks.ibm.com/abstracts/sg247726.html
- DB2 for z/OS: Data Sharing in a Nutshell
  - http:// www.redbooks.ibm.com/abstracts/sg247322.html
- System Programmer's Guide To: Workload Manager
  - http:// www.redbooks.ibm.com/abstracts/sg246472.html
- Application Design for High Performance and Availability
  - http:// www.redbooks.ibm.com/abstracts/sg247134.html
- Workload Management for DB2 Data Warehouse, REDP-3927
  - http:// www.redbooks.ibm.com/abstracts/redp3927.html
- DB2 10 for z/OS Performance Topics
  - http:// www.redbooks.ibm.com/abstracts/sg247942.html

The End