


What's New in DB2 V9 for z/OS

Presenter: Mike Biere
IBM WW Marketing Mgr.
IBM Silicon Valley Lab
mbiere@us.ibm.com
Session: 002



IBM Information
ON Demand 2010

INFORMATION-LED
TRANSFORMATION



January 21 - SINGAPORE • January 26 – MALAYSIA • January 28 - THAILAND



What's new in DB2 for z/OS?

- Synergy with DB2 for z/OS and System z
- What does DB2 9 have for me?
- What are DB2 9 performance characteristics?
- How is database administration improved?
- What are the improvements for application programming?
Is XML for me?
- What is the vision beyond DB2 9? What early planning can I do for DB2 9 and beyond?



DB2 Deep Synergy With System z

Key integration points include:

- Data sharing (availability and scale out)
- Processor instructions and memory structure
- Hardware data compression & encryption
- zIIP specialty engines
- Unicode conversion
- Encrypted TCP/IP communication (SSL), encrypted data
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 1 MB page size (z10)
- Decimal float arithmetic (z10)
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF)
- z/OS RRS integrated commit coordinator
- Solid state disks



IBM z10 Extends Scale and Value



- G4 - 1st full-custom CMOS S/390®
- G5 - IEEE-standard BFP; branch target prediction
- G6 – Copper Technology (Cu BEOL)
- z900 - Full 64-bit z/Architecture®
- z990 - Superscalar CISC pipeline
- z9 EC - System level scaling
- z10 EC – Architectural extensions

DB2 for z/OS *Lowering TCO*

→ Maximum value for dollar investment

- Hardware pricing
 - CPU saving specialty engines (zIIP, zAAP..)
 - Compression of disk space (data, index)
- Software pricing
 - Reduction for tiers
 - Parallel Sysplex aggregation
 - z990, z9, z10 technology dividend
 - 10% reduction in charge units for each step
 - zNALC, Value Unit Edition, Subcapacity pricing



→ CPU + Memory + I/O and disk + Software + Energy and floor space + People = Improved Total Cost of Ownership (TCO)

Helping to drive down the cost of IT

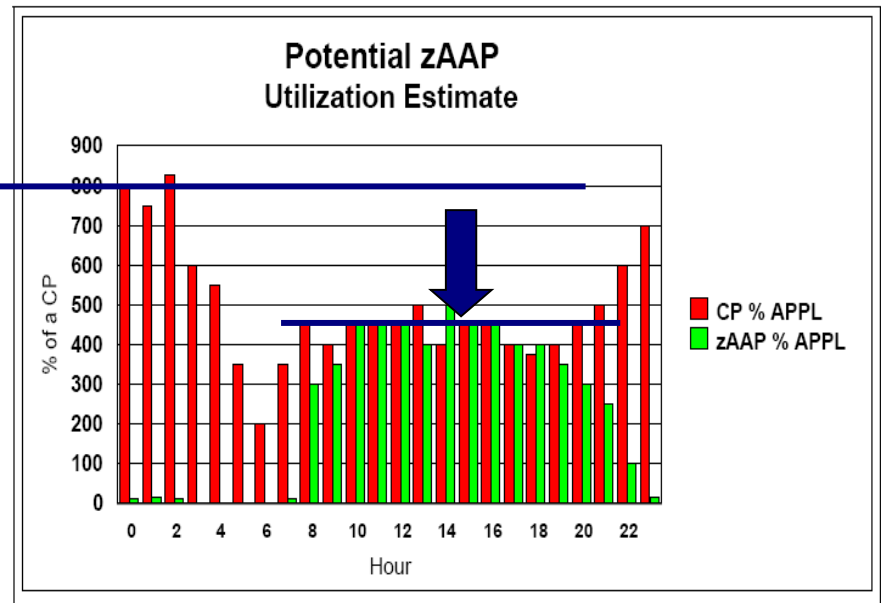
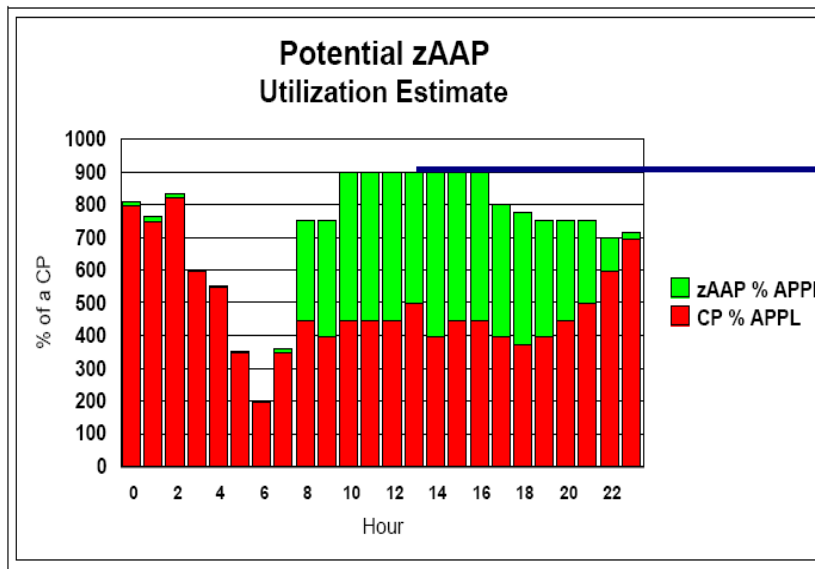
Now even more workloads can benefit from zIIP

- Integrate data across the enterprise, optimize resources and lower the cost of ownership
 - DB2 remote SQL, parallel, **utilities**
 - Network encryption
 - Serving XML data
 - z/OS Global Mirror
 - Use by ISVs
 - zIIPs offer economics to help you
 - **PLUS** zIIP price same for z10 EC as z9 EC
- IPsec encryption
 - HiperSockets
 - Financial Reporting
 - **z/OS CIM Server**
 - **DB2 sort utility**
 - **zAAP on zIIP**



**IBM System z10 Integrated Information Processor and
IBM System z9 Integrated Information Processor**

Why do zIIPs, zAAPs and IFLs Reduce Cost?



1. Hardware costs: By moving workload from general purpose processors to zIIP, zAAP and IFL processors (higher cost to lower cost processors).

2. Software Costs: license/maintenance costs based on number of and usage of general purpose central processors. Specialty engines can reduce number of CP's.

No z/OS software charges based on zIIP, zAAP and IFL processors or usage.

DB2 9 for z/OS at a glance

Application Enablement

- pureXML
- Optimistic locking for WebSphere
- LOB performance, usability
- Native SQL procedure language
- SQL improvements that simplify porting

RAS, Performance, Scalability, Security

- More online schema changes
- Online REBUILD INDEX, Online REORG improvements, Clone tables
- Trusted context and ROLES
- Parallel Sysplex clustering improvements
- 64-bit virtual storage improvements

Simplification, Reduced TCO

- Index compression
- Partition By Growth tables
- Package stability
- Volume based backup / recovery
- Automatic object creation

Dynamic Warehousing

- Many SQL improvements
- Dynamic index ANDing
- Histogram statistics
- New built-in OLAP expressions
- Optimization Service Center

DB2 SQL 2004

z z/OS V8

common

lw Linux, Unix & Windows V8.2



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning

c
o
m
m
o
n

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

l
w

Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE or DELETE, multi-site join, MDC

DB2 SQL 2007

z z/OS 9

common

luw Linux, Unix & Windows 9



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from MERGE, index compression**

c

o

m

m

o

n

l

u

w

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, **UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitioning, compression**

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, **XQuery**

DB2 SQL 2008 cross platform reference

z z/OS 9

common

luw Linux, Unix & Windows 9.5



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, TRUNCATE, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, index compression

c
o
m
m
o
n

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitions, data compression, **session variables, DECIMAL FLOAT, optimistic locking, ROLE**

l
u
w

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, **XML enhancements, array data type, global variables, vendor syntax**

DB2 SQL 2009

z z/OS 9

common

luw Linux, Unix & Windows 9.7



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE

c
o
m
m
o
n

Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, **TRUNCATE, index & XML compression, created temps**

l
u
w

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, XML enhancements, array data type, global variables, **even more vendor syntax, LOB & temp table compression, inline LOB, administrative privileges, implicit casting, date/time changes, currently committed**

DB2 SQL

z z/OS X

common

luw Linux, Unix & Windows 9.7



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, **data versioning, access controls**

c
o
m
m
o
n


Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, OmniFind, spatial, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, TRUNCATE, index & XML compression, created temps, **inline LOB, administrative privileges, implicit casting, date/time changes, currently committed, moving sum & average**

l
u
w

Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, more Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, XML enhancements, array data type, global variables, even more vendor syntax, LOB & temp table compression,

What You Can Do with pureXML

- Create tables with XML columns or alter table add XML columns
- Insert XML data, optionally validated against schemas
- Create indexes on XML data
- Efficiently search XML data
- Extract XML data
- Decompose XML data into relational data
- Construct XML documents from relational and XML data
- All the utilities and tools support for XML




Universal Table Space

- Combination of segmented with partitioning options
 - Better space management
 - Support of mass deletes / TRUNCATE
 - Better answer when there is no good partitioning key
 - Better answer when size is unknown
- If partitioned
 - Still must be one table per table space
 - Can choose **Range Based partitioning** (as before: PBR)
 - Can choose **Partitioned By Growth** (PBG)
- DROP / CREATE to migrate existing page sets
- Simple table spaces can not be created
 - Default table space is now Segmented (CM) or PGB (NFM)

LOB Improvements in DB2 9: Faster & Easier

- ➔ Progressive Streaming for LOB Locator Values
 - DB2 uses LOB size to determine whether to send LOB data to Java or DB2 CLI clients in one (<32KB), in chunks (<1MB) or as LOB locator (>=1MB) [Transparent to application using LOB locators]
- ➔ Elimination of LOB locks for improved availability and performance
- ➔ Utility Changes
 - REORG LOB reclaim space
 - Logging for > 1GB LOBs
 - Online CHECK LOB and DATA
 - LOB column lengths > 32KB in utilities
 - LOAD, UNLOAD, Cross load V7, V8 APARs
- ➔ File reference variables allow direct transfer of LOB data between DB2 and the file named in the variable
- ➔ Implicit object creation
- ➔ FETCH CONTINUE allows applications to retrieve LOB/XML data in pieces without using locators

IBM
**LOBs with DB2 for z/OS:
Stronger and Faster**




Indexing Enhancements

- ➔ Larger index pages allow for more efficient use of storage
 - Fewer page splits for long keys
 - More key values per page
- ➔ Index compression provides page-level compression
 - Data is compressed to 4K pages on disk
 - 32K/16K/8K pages results in up to 8x/4x/2x disk savings
 - No compression dictionaries
 - Compression on the fly
 - No LOAD or REORG required
- ➔ Find unused indexes with real time statistics
- ➔ Rebuild Index SHRLEVEL CHANGE
- ➔ Define RANDOM index keys to avoid hot spots with multiple processes inserting sequential keys

Information Management software

IBM

Jeff Berger
Paolo Bruni

Data Sharing DB2 9 Enhancements

- Log latch contention relief
- Restart performance enhancements
 - Reduced impact of retained locks – released as rollbacks are completed
 - Open data sets ahead of log apply
- Command to remove GBP-dependency at object level
 - ACCESS DB MODE(NGBPDEP)
 - Typical usage would be before batch run
 - Command to “prime” open data set
 - ACCESS DB MODE(OPEN) [PART]
- Auto-recover GRECP/LPL objects on group restart
 - Useful in Disaster Recovery or GDPS scenarios
- DB2 overall health taken into account for WLM routing
- Balance group attach connections across multiple members on same LPAR (V7 & V8 usermod)
- Group wide outage no longer needed for new LOB locking protocol (apar)

Utilities Highlights

- More online utilities
 - Rebuild Index SHRLEVEL CHANGE
 - Reorg LOB now supports SHRLEVEL REFERENCE (space reclamation)
 - Check data, LOB and repair locate ... SHRLEVEL CHANGE
 - Check index SHRLEVEL REFERENCE supports parallel for > 1 index
 - Clones for “online LOAD REPLACE”
- Online REORG BUILD2 phase elimination
- Substantial CPU reductions
- REORG parallelism for UNLOAD, RELOAD, LOG phases
- Utility TEMPLATE switching
- UNLOAD SKIP LOCKED DATA option

Utilities Highlights...

- RECOVER to any point-in-time with consistency
- MODIFY Recovery enhancements
 - “*Retain*” keyword added to improve management of copies
 - LAST(n), LOGLIMIT, GDGLIMIT
- Volume-based COPY/RECOVER (BACKUP SYSTEM/RESTORE SYSTEM)
 - RECOVER modified to enable object-level recovery from volume FlashCopy
 - Full integration of tape into BACKUP/RESTORE SYSTEM utilities
 - Incremental FlashCopy, APAR PK41001
- Truncate log based on timestamp
- RECOVER RESTOREBEFORE to use an earlier image copy
- Display progress of RECOVER during log apply
- COPY CHECKPAGE option always active
 - “Copy Pending” avoided if broken page encountered
- COPY SCOPE PENDING to copy only objects in “Copy Pending”

Newest major changes in APARs

- Reordered row format parameter & utility syntax
- Workfile separation
- WLM management for buffer pools
- -ACCESS command wildcards
- Ability to have compression on SPT01
- 64 bit ODBC
- Explain format to current release
- Help with removing private protocol
- Precompiler like V7 on DB2 9

Why Migrate to DB2 9 for z/OS?

→ Business needs

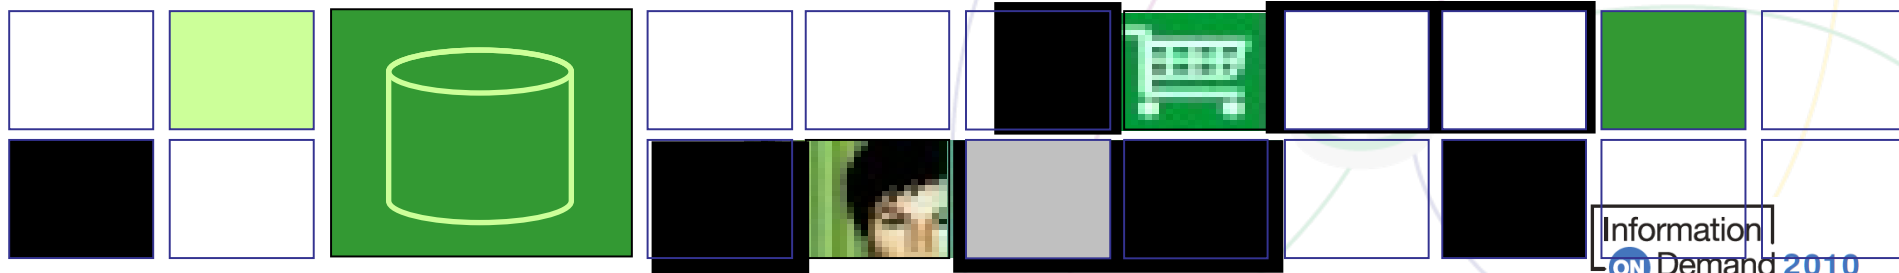
- Reduce CPU time & disk space
- Improve business agility
- Service Oriented Architecture

→ Application developers need

- PureXML for a powerful SQL and XML interface to XML data
- Powerful new SQL enhancements
- Portability with SQL and data definition compatibility


→ Database Administrators need

- Improve availability and performance
- More flexible security and easier regulatory compliance
- Better web application & data warehouse function and performance
- LOB function, performance, usability



Why is migration easier to DB2 9 for z/OS?


- Migration process enhancements: ENFM speed, CM*
- Much less performance regression:
 - Earlier improvements
 - Package stability & tools for avoiding access path issues
- CCSIDs and old product issues resolved in V8
- Simpler virtual storage considerations
- Less impact from incompatible changes
- Earlier deliveries from vendors





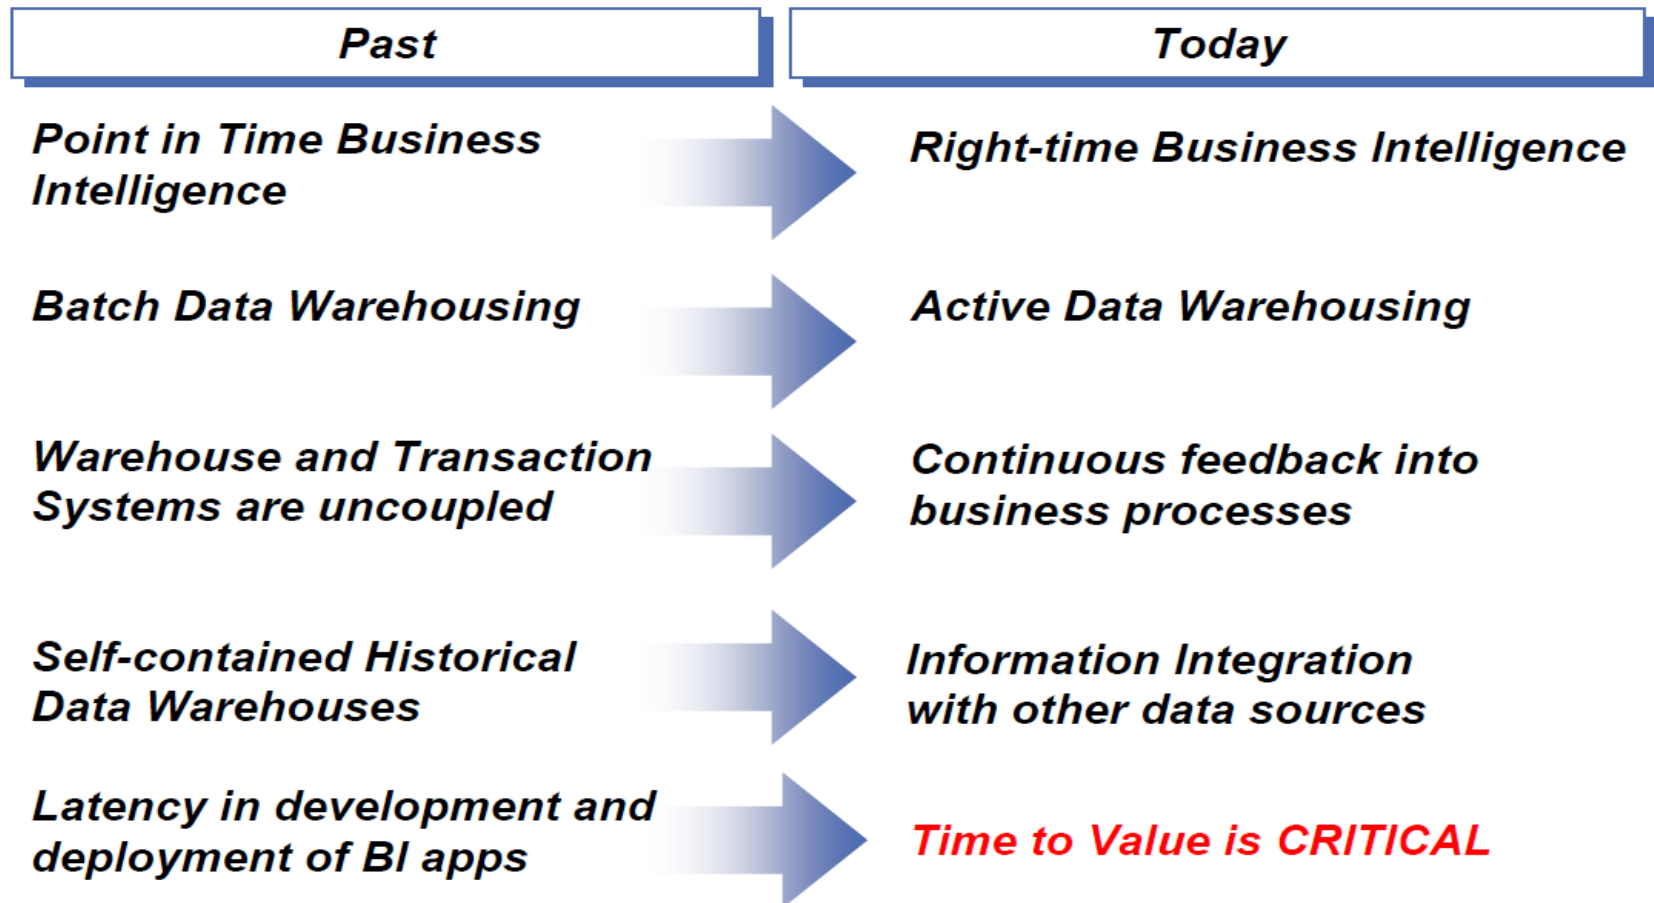
Most consumable DB2 9 improvements

- CM very little to no action:
 - Optimization Service Center, Data Studio, & Optim Query Tuner
 - Utility CPU reductions
 - Archive log striping, compression
 - Larger prefetch, write & preformat quantities
 - LOB performance
 - DDF VSCR
 - Index lookaside
 - Changed online REORG
- Package stability
- Improved RUNSTATS
- Optimization improvements, EDMPOOL VSCR, increased parallel & zIIP
- NFM
- LOB lock avoidance
- Data sharing logging
- Improved index leaf page split
- Reordered row format, native SQL
- Index: larger page sizes, compression, index on expression



Data Warehouse & Business Intelligence Trends

Business Intelligence Becoming Mission-Critical



Data Warehousing on z/OS – What is driving this?

- Customer commitment to the z platform
 - Customers want to protect their significant investment in System z
 - TCO can be reduced through the utilization of existing processors, people, practices
 - TCO may also be achieved through a consolidation approach
- New BI trends are changing the DBMS landscape
 - The distinction is blurring between warehouse and OLTP databases based on new trends such as Dynamic Warehouse and Operational BI, driving:
 - The need for increased reliability, availability, security, and compliance in a DWH DBMS
 - The need for very current warehouse data, where proximity to the source provides an advantage
- Many z customers already have a warehouse on DB2 for z/OS
 - This drives requirements into hardware and software, which in turn drives a trend
 - DB2 has responded with increased functionality and performance; hardware changes are driving down costs
- Specialty processors provide new ways to optimize TCO
 - zIIPs and IFLs are driving down hardware and software costs; DWH/BI can make excellent use of these processors, ultimately driving TCO advantages

New Information On Demand Software for System z

Better business decisions,
faster and with a lower overall TCO

- DB2 9 for z/OS
- InfoSphere Information Server for System z
- InfoSphere MDM Server for System z
- Cognos 8 BI for System z
- Cognos Now! For Linux on System z
- InfoSphere Warehouse on System z
- IBM Smart Analytics Optimizer
- SPSS ... watch this space

DB2 Developer Workbench → Data Studio → Optim

before

IBM DB2 Developer Workbench V9.1

- SQL Query Editor
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

**Data Studio is a full replacement of
DB2 Developer Workbench
plus much more**

- **DB2 for Linux, Unix, Windows v9.1, v9.5, v9.7**
- **DB2 for z/OS v7, v8, 9**
- **DB2 for i v5r2, v5r3, v5r4, v6**
- **Informix Dynamic Server (IDS) v9, v10, v11**

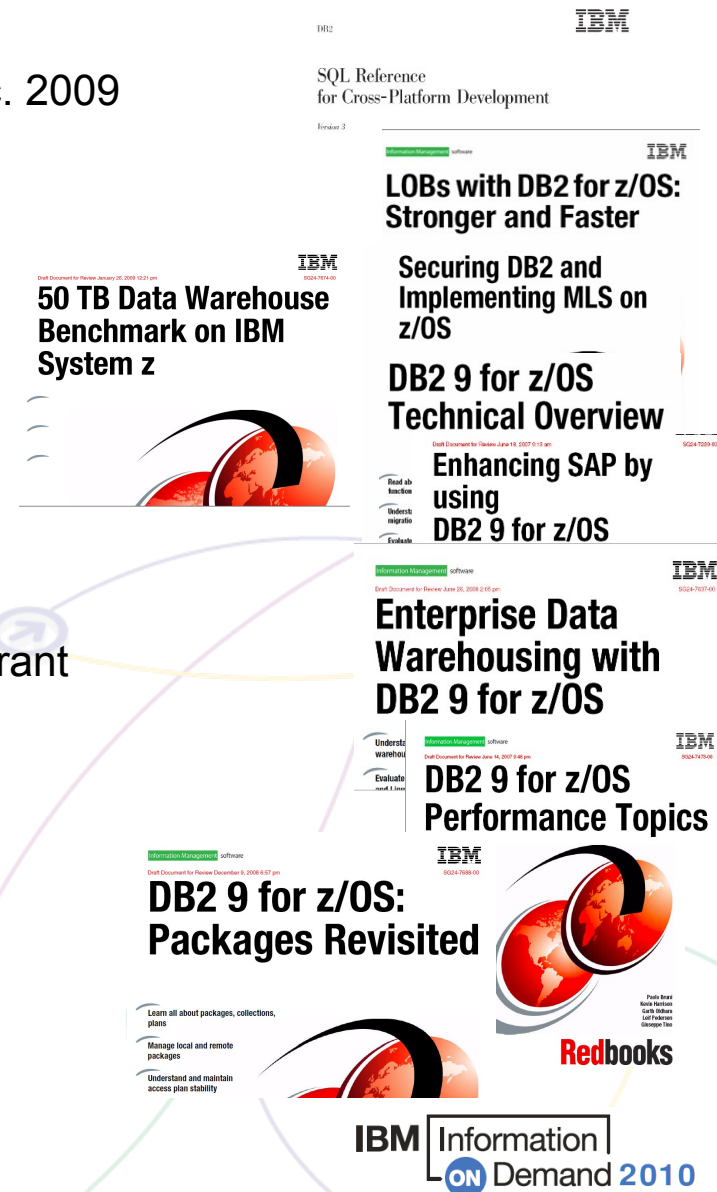
now

IBM Data Studio

- Integrated Query Editor – SQL + XQuery
 - SQLJ Editor
 - SQL Builder
 - XQuery Builder
 - SQL Routine Debugger
 - Java Routine Debugger
 - XML Editor
 - XML Schema Editor
 - Data Management
 - Visual Explain
 - Project Management
-
- ER Diagramming
 - Data Distribution Viewer
 - Object Management
 - Browse & Update Statistics
 - Security Access Control
 - Connection Management integration with Kerberos and LDAP
 - Data Web Services
 - IDS Server Support
 - Health Monitoring DB2 for LUW 9.5 and DB2 9 for z/OS

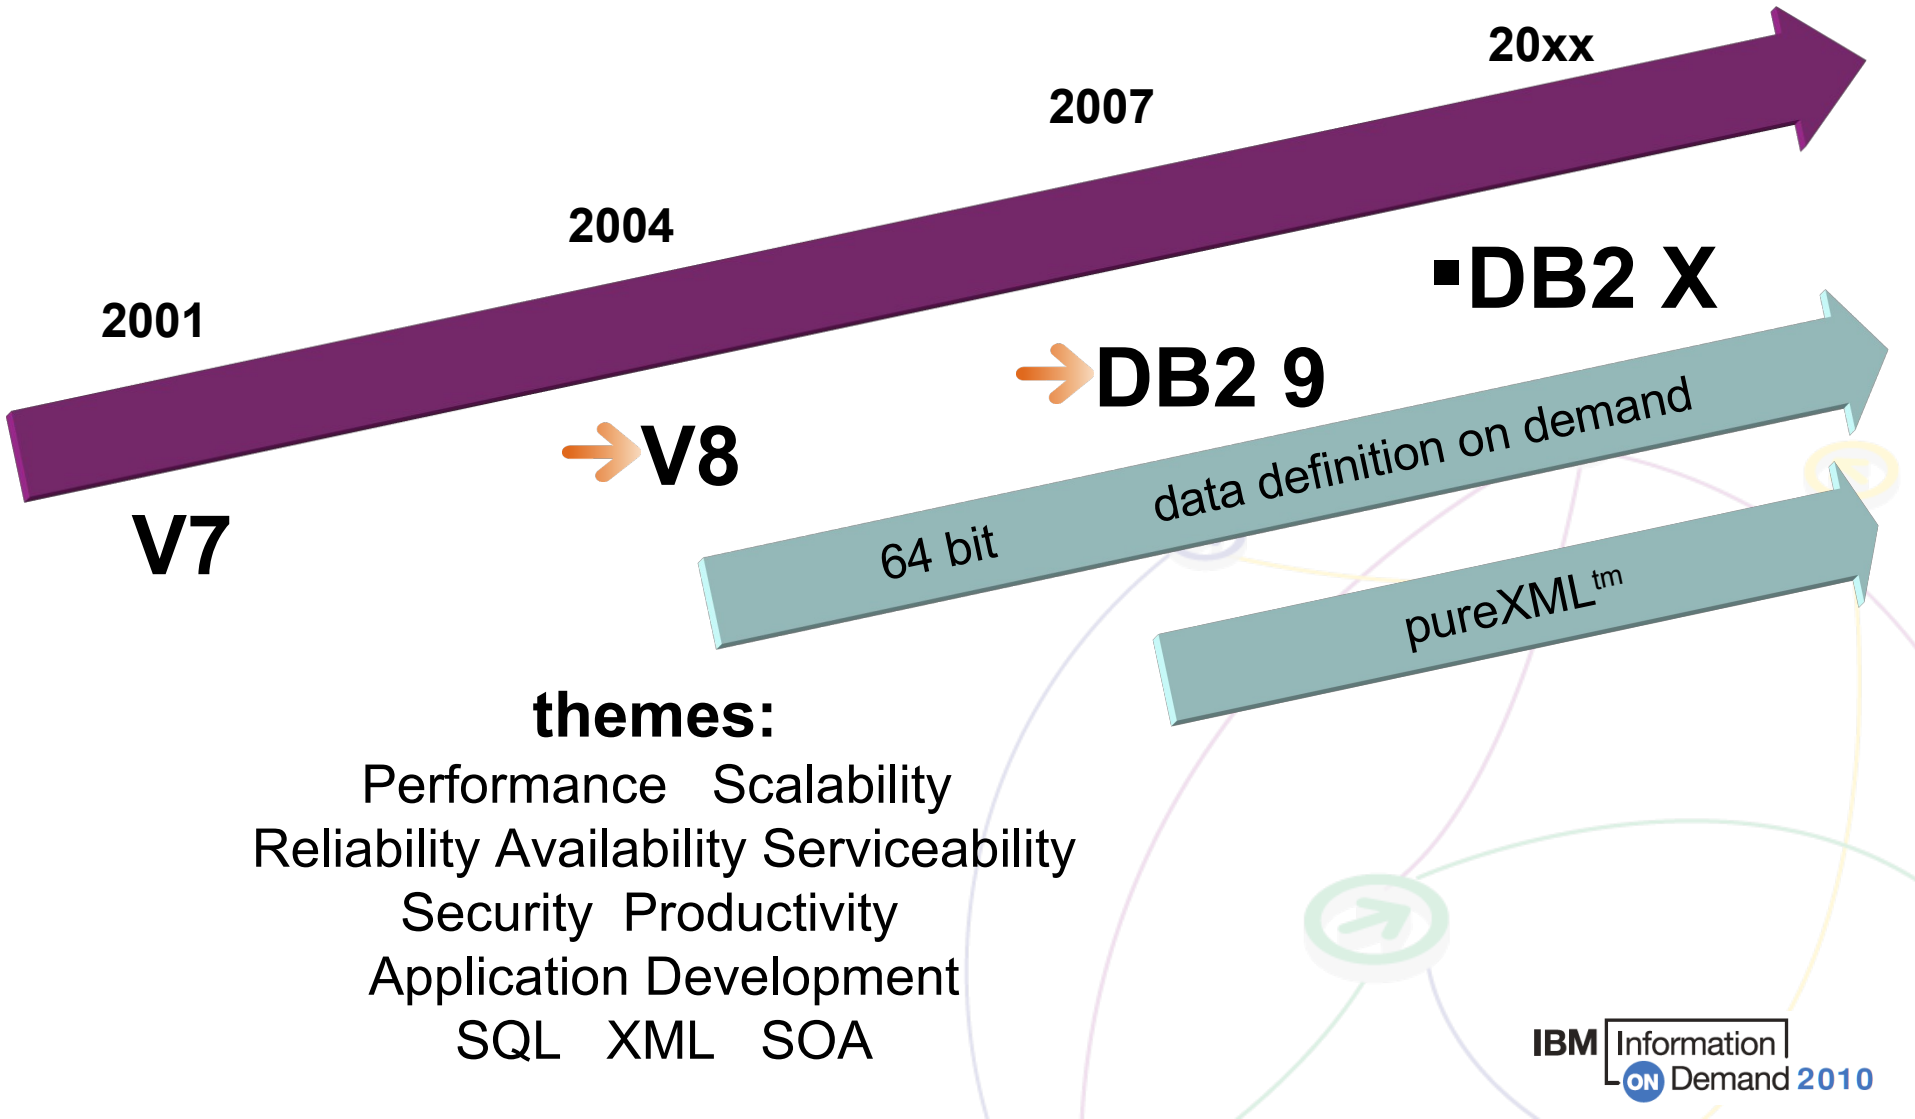
DB2 9 in IBM Redbooks Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473 updated Dec. 2009
3. DB2 9 Stored Procedures SG24-7604
4. Index Compression with DB2 9 for z/OS redp4345
5. SQL Reference for Cross-Platform Development
6. Enterprise Database Warehouse, SG24-7637
7. 50 TB Data Warehouse on System z, SG24-7674
8. New Tools for Query Optimization SG24-7421
9. LOBs with DB2 for z/OS SG24-7270
10. Deploying SOA Solutions SG24-7663
11. Enhancing SAP - DB2 9 SG24-7239
12. SAP Application on Linux z SG24-6847
13. Best practices SAP BI - DB2 9 SG24-6489-01
14. Data Sharing in a Nutshell, SG24-7322
15. Securing DB2 & MLS z/OS SG24-6480-01
16. Data Sharing: Distributed Load Balancing & Fault Tolerant Configuration redp4449
17. Considerations on Small & Large Packages redp4424
18. Backup and Recovery Considerations redp4452
19. Powering SOA with IBM Data Servers SG24-7259
20. Packages Revisited, SG24-7688
21. Data Studio V2.1 Web Services redp4510
22. Ready to Access Solid-State Drives redp4537
23. Distributed Functions SG24-6952-01 Sept 2009
24. Buffer Pool Monitoring & Tuning redp4604
25. Securing & Auditing Data SG24-7720
26. Serialization and Concurrency SG24-4725-01 new
27. Utilities SG24-6289-01 draft



DB2 for z/OS Into the Future

Delivering Customer Value



DB2 X for z/OS What's exciting?

Efficiency

- CPU reduced: transactions & queries
- Ten times more concurrent users

Resiliency

- More online schema changes
- Concurrency for catalog & utilities
- Improved security controls and audit

Applications


- Versioned data or temporal queries
- pureXML and SQL enhancements

→ Productivity improved for DBAs, application programmers, & systems


DB2 Version X Performance Objective

- Historical goal <5 % version-to-version performance regression
- Goal of 5% -10% initial performance improvement for X
- Many customers expected to reduce CPU time by 10% - 20%

Average %CPU improvements version to version



DB2 X: Preliminary Performance Plan → Significant CPU Reductions: Best on z10



Transactions	DB design changes
Batch	Hash access
REBIND	application changes
	SQL adjustments

Your situation will vary. Less CPU is better.


Processors z890, z990, z9 and later

z/OS 1.10 and later

Technology Preview: IBM Smart Analytics Optimizer

What is it?

- A special purpose, network-attached appliance that is an add-on to a DB2 for z/OS system
- Offloads typical DW/BI queries resulting in predictable and orders-of-magnitude faster query response times while reducing overall TCO



Business Value

- Dramatically lowers the cost for query and reporting on System z
- Advanced in-memory scale-out cluster technologies that keep the complete system centrally managed without having to change any requirements for BI applications
- Complements the many new Data Warehousing features in DB2 9 for z/OS
- Leverages the many new warehousing and business intelligence solutions now available on System z

Targeted Uses for DB2 for z/OS customers:

- Requirements to accelerate a subset of their warehouse or reporting queries
- Looking for more insight and business intelligence from operational data
- Needs to consolidate datamarts or data stores into one enterprise warehouse

Thank YOU



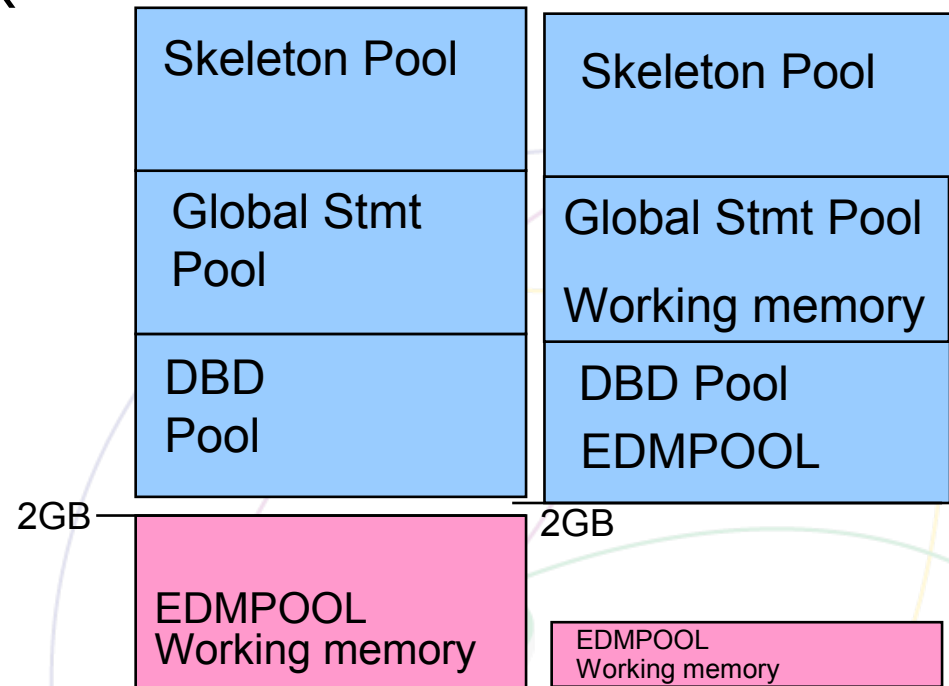
Additional Reference Slides for DB2x, utilities, etc.



DB2 X: Scalability 64 bit Evolution (Virtual Storage Relief)

Scalability: Virtual storage constraint is still an important issue for many DB2 customers.



- DB2 9 helped (~ 10% – 15%)
- DB2 X expect to have 5 X to 10 X threads, move 80% - 90%
 - More concurrent work
 - Reduce need to monitor
 - Able to consolidate LPARs
 - Reduced cost
 - Easier to manage
 - Easier to grow



Running a Large Number of Active Threads

Today



Coupling Technology



- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s / LPAR

DB2 X

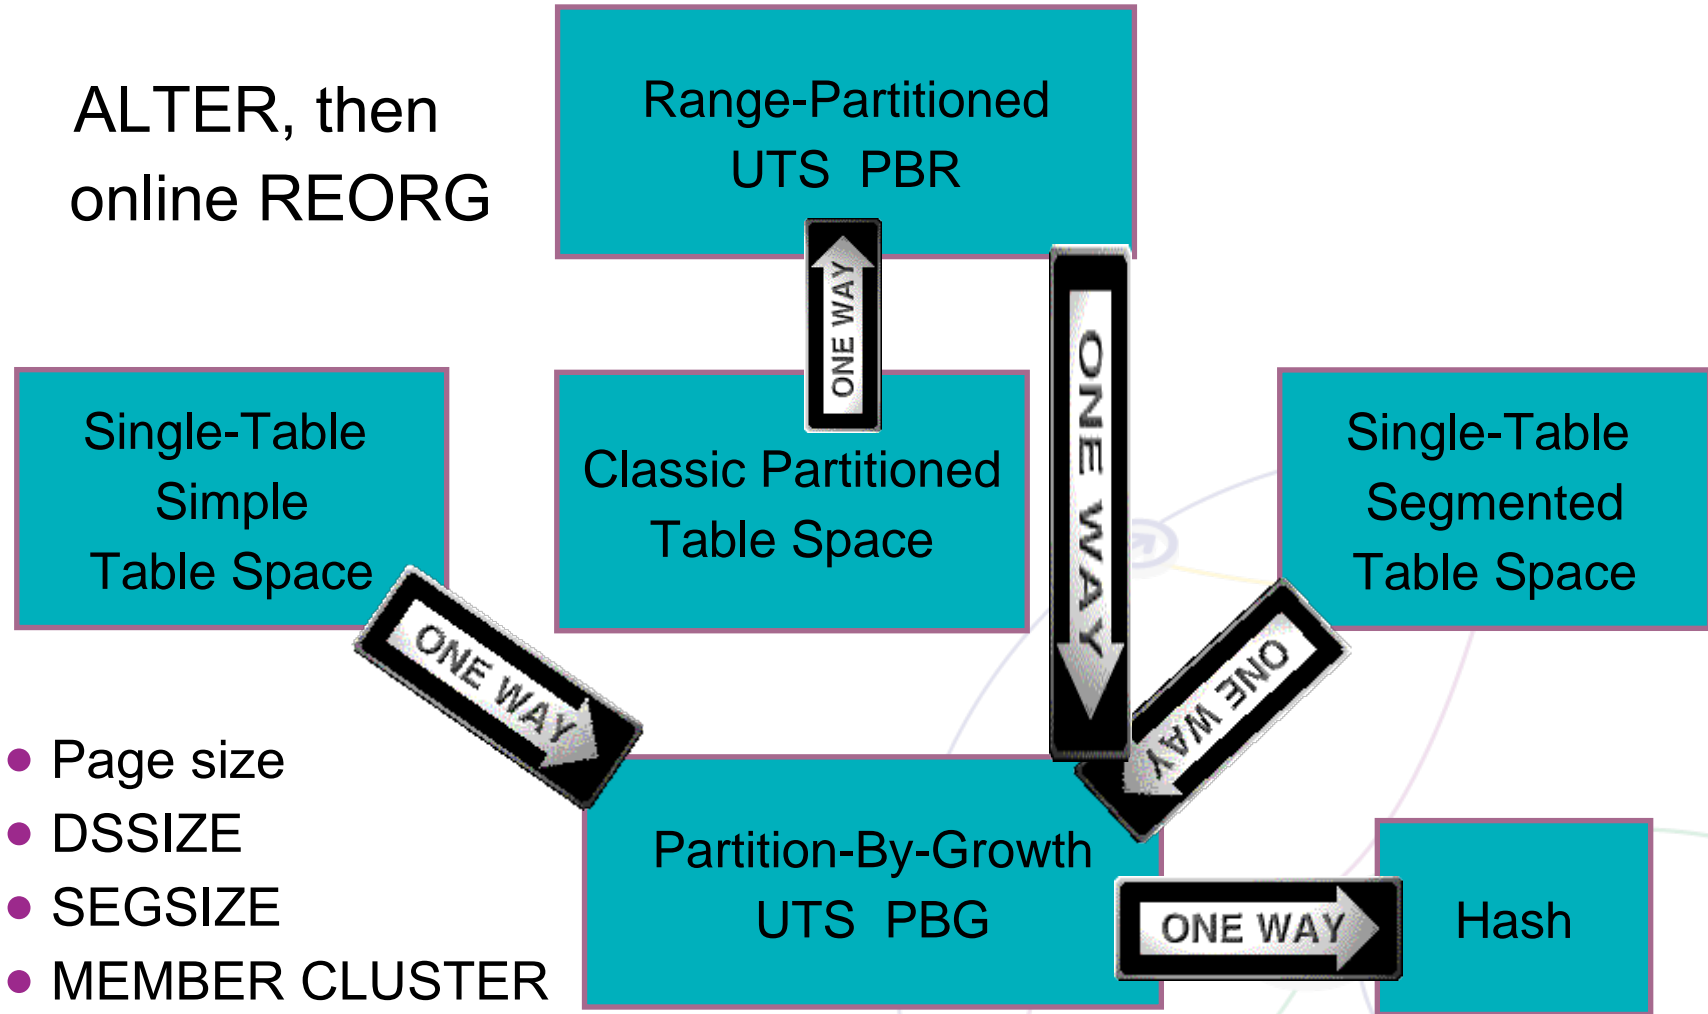
Coupling Technology



- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for HA and XXL scale
- Rule of thumb: save 1/2% CPU for each member reduced, more on memory

DB2 X Availability: ALTER table spaces

ALTER, then
online REORG



- Page size
- DSSIZE
- SEGSIZE
- MEMBER CLUSTER

DB2 X: Business Security & Compliance

- Protect sensitive data from privileged users
 - SYSADM without data access, DBADM for all DB
- Separate authority to perform security related tasks
- Allow EXPLAIN without execute privilege or ability to access data
- Audit privileged users
- More granular administrative authorities
- Row and column access control
 - Allow masking of value
 - Restrict user access to individual cells



Use disk encryption

DB2 X: Productivity – Doing More with Less!

- Auto statistics collection
- Compress 'on the fly'
 - Avoid need to run utility
- Reduce contention, more online processing
- Simpler memory management
- Automatic config of IBM supplied UDFs and SPs
- Enhancements for monitoring

Name	Monitoring Status	Data Server Status	Critical	Warning	Alert	CPU Usage	Disk Space	Memory Usage	Locking	SQL Performance	Connections	Transactions	Logging	Maintenance
Production	■	■	3	8	■	■	■	▲	■	■	■	■	■	■
Web	■	■	1	1	■	■	■	▲	■	■	■	■	■	■
eCommerce	■	■	0	0	■	■	■	■	■	■	■	■	■	■
Support	■	■	1	1	■	■	■	▲	■	■	■	■	■	■
Retail	■	■	0	0	■	■	■	■	■	■	■	■	■	■
New York	■	■	0	0	■	■	■	■	■	■	■	■	■	■
Los Angeles	■	■	0	0	■	■	■	■	■	■	■	■	■	■
Accounts	■	■	2	3	■	■	■	▲	■	■	■	■	■	■
Marketing	■	■	0	4	■	■	■	▲	■	■	■	■	■	■
Test	■	■	0	0	■	■	■	■	■	■	■	■	■	■
Development	■	■	0	0	■	■	■	■	■	■	■	■	■	■

Manual invocation of

- **RUNSTATS**
- **COPY/BACKUP SYSTEM**
- **QUIESCE**
- **MODIFY RECOVERY**
- **REORG**

Many Features improve for SAP

→ Autonomics

- Compress on the fly on INSERT
- Auto-statistics
- Access path stability and hints enhancements
- Access path lock-in and fallback for dynamic SQL
- Automatic checkpoint interval
- Automated installation, configuration & activation of DB2 supplied stored procedures & UDFs
- Data set FlashCopy in COPY & inline copy
- Inline image copies for COPY YES indexes
- UNLOAD from FlashCopy backup
- REORG enhancements
- Reduce need for reorganizations for indices
- Performance
- CPU reductions
- Hash access path

- Optimizer enhancements, paging through result sets
- Parallel index update at insert
- Faster single row retrievals
- Inline LOBs
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- DEFINE NO for LOBs and XML
- MEMBER CLUSTER for UTS
- Query parallelism enhancements: lifting restrictions
- Dynamic Index ANDing Enhancements
- Option to avoid index entry creation for NULL value
- Index include columns
- Buffer pool enhancements
- Scalability
- Many more threads
- Reducing latch contention
- Workfile spanned records, PBG support, and in-memory enhancements

→ Availability

- More online schema changes for table spaces, tables and indexes via online REORG
- Online REORG for LOBs
- Online add log
- Automatically delete CF structures before/during first DB2 restart

→ Portability

- Allow non-NULL default values for inline LOBs
- Loading and unloading tables with LOBs in stream
- 'Last committed' locking semantics
- Default SAP settings for DB2

→ Security

- More granular DBA privileges

pureXML improved performance and usability

- XML schema validation in the engine for improved usability and performance
- Binary XML exchange format improves performance
- XML multi-versioning for more robust XML queries
- Allow easy update of sub-parts of XML document
- Stored proc, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK XML utility



Data Warehousing

- Moving Sum, Moving Average
- Enhanced query parallelism technology for improved performance
 - Remove query parallelism restrictions
- In-memory techniques for faster query performance
- Advanced query acceleration techniques
 - IBM Smart Analytics Optimizer

Key details about DB2 X

- CM, ENFM, NFM modes
- Prerequisites
 - z/OS V1.10
 - DB2 9 for z/OS in NFM
 - System z10, z9, z890, z990, and above (no z800, z900)
- Items deprecated in earlier versions eliminated:
 - Private protocol → DRDA (new help in DSNTDP2DP, PK64045)
 - Old plans and packages V5 or before → REBIND
 - Plans containing DBRMs → packages PK62876
 - ACQUIRE(ALLOCATE) → ACQUIRE(USE)
 - Old Plan table formats → DB2 V8 or 9 format (59 columns)
 - XML Extender → XML type
 - DB2 MQ XML user-defined functions and stored procedures → XML functions
 - DB2 Management Clients feature (DB2 Administration Server, Control Center, & Development Center) → IBM Data Studio application & administration services
 - msys for Setup DB2 Customization Center → install panels
 - BookManager use for DB2 publications → Info Center, pdf

DB2 X for z/OS At a Glance

Application Enablement

- Versioned data or Temporal
- pureXML enhancements
- Last Committed reads
- SQL improvements that simplify porting

RAS, Performance, Scalability, Security

- Wide range of performance improvements
- Hash access to data
- More online schema changes
- Catalog restructure for improved concurrency
- Row and column access control
- Administrator privileges with finer granularity

Simplification, Reduced TCO

- 5 – 10 times more threads per DB2 image
- Auto statistics
- Data compression on the fly
- Query stability enhancements
- Reduced need for REORG
- Utilities enhancements



Dynamic Warehousing

- Moving sum, moving average
- Many query optimization improvements
- Query parallelism improvements
- Advanced query acceleration

Get updated books latest so far December 2009

<http://publib.boulder.ibm.com/infocenter/imzic/>

- ✓ Administration Guide
- ✓ Data Sharing: Planning and Administration
- ✓ Performance Guide
- ✓ Utility Guide and Reference
- ✓ Application Programming Guide & Reference
- ✓ Application Programming Guide & Reference for JAVA™
- ✓ ODBC, Spatial, ...
- ✓ Redbooks
- ✓ Installation Guide
- ✓ RACF Access Control Module Guide
- ✓ Messages
- ✓ Codes



Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR
- ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

IBM, the IBM logo, ibm.com, **DB2 and WebSphere** are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml
Other company, product, or service names may be trademarks or service marks of others.

Data Management Communities for DB2

- IDUG – the worldwide community of DB2 users
 - Membership is FREE – join today! www.idug.org

- Data Management Community – share and interact with peers around the world
 - www.ibm.com/software/data/management/community.html

- Information Champions – recognizes individuals who have made the most outstanding contributions to the Information Management community
 - www.ibm.com/software/data/champion

Landscape – Customer Challenges

- Tremendous regulatory compliance pressures to demonstrate adequate institutional controls including audit reporting.
- Current DB2 on z/OS environment typically has minimal auditing
- Manual effort requiring interaction by DBA's
- Reactive in nature with the implication that you only find information post event, or after the first breach
- Home grown process can provide some level of access reporting, however:
 - Application managed code you have to maintain
 - Exposure as a lack of robust application change controls can allow disabling of audit processing
- Overhead (perceived or actual) in many cases drive decision to not audit DB2 on z/OS data
- DB2 trace based processes are managed by DBA's
 - The DBA's are responsible for generating audit data with which they are in turn audited, this constitutes a significant security risk and exposure.**



DB2 Audit Trace



DB2 Audit Trace versus RACF

Why Audit when Production is Locked Down?

→ Common arguments:

- "We don't need to audit, we have controls surrounding who can access data"
- "We control who is connected to the DB2 SYSADM group and we know what those people are authorized to do"

→ Counter arguments:

– RACF does two things:

- Prevents people from accessing a resource that is not essential or appropriate for their jobs
- Allows people access to the necessary data to do their jobs

– But RACF does NOT:

- prevent a malicious update if the user has authority to the data.
- prevent an authorized user from accessing sensitive data that is **NOT** within the scope of their job.
 - E.g. a bank teller looks up the CEO's bank balance or personal customer information
- provide meaningful information about access to protected DB2 resources (authorized or not).

DB2 Audit Trace versus RACF

→ Key Points:

- RACF provides significant controls to protect access to resources, but does little in the way of meaningful access reporting
- DB2 Audit trace will do nothing to protect data, but provides data to help understand what type of access has occurred.
 - Auditing is about ensuring that the appropriate controls are in place to identify inappropriate access and use of production data
 - You need some form of audit facility to watch your privileged users who have RACF and/or DB2 authority and users that have access to sensitive data within the scope of their job
 - Understanding how trusted (privileged) users access sensitive information is essential to ensuring that data is indeed protected

What to Audit – A busy slide

- **Closed Application Environment (*Probably not a candidate*)**
 - Traditional Application controls well defined and comprehensive
 - CICS and IMS TM – Signon and Transaction Access secured via RACF
 - Production Batch – Controlled via program pathing / Job Scheduling
- Data warehouse – no risk of update but access audit might be needed
- Adhoc execution environment – QMF, SPUFI, etc. Constitutes exposure
 - SPUFI Plan can be restricted but ALL use should be audited
- Privileged ID's (DBA/Sysadmin) should be audited
- Distributed Application Environment
 - Use of SQLESETI can provide granularity with credential population to IFI extensions
 - End User Workstation Name
 - End User Workstation Process
 - End User Workstation Userid
 - Implement RACF Enterprise Identify Mapping Feature
 - <http://www-03.ibm.com/servers/eserver/security/eim>
- Data may not be as granular as you think
 - Depending on how you configured your connections into DB2 – CICS attach, SAP, or CICS users with unique id's, and distributed transactions. May get all audit data but may not be meaningful because of attach environments. Group versus AUTHID. SQLESETI implementation can help
- “Offline” Utilities and certain tools are used outside of DB2
 - RACF dataset access defined controls
 - “Trigger” based audit
 - Use of DSN1COPY should be restricted

Audit data sources

→ DB2 catalog

- SQL queries on catalog, other data
- audit, accounting and performance traces
- recovery log, current & historical data
- RACF audit facility, other SMF data, ...

→ Audit tools and techniques

- tracing: audit, performance, accounting, monitor
- formatting the traces: OMPE or PM, others
- TCIM, DB2 Audit Management Expert, others
- DSN1SMFP, others
- log formatting: tools, DSN1LOGP, Log Analyzer
- various recovery and cloning techniques
 - triggers
- REPORT RECOVERY
- RACF print, unload

What actions are needed to start the Audit trace?

→ -DSN START TRACE (AUDIT) CLASS (1,2,4,5,8) DEST (SMF)

- Requires one of the following privileges:
 - SYSOPER
 - SYSCTRL
 - SYSADM
 - TRACE
- In addition, Class 4 and 5 events will only be collected for objects (tables) with the audit attribute turned on via ALTER:
 - AUDIT CHANGES – enables collection of changes in conjunction with CLASS (4)
 - AUDIT ALL – enables collection of changes and / or reads with CLASS 4 and/or 5 active
- Note: When ALTER AUDIT is performed, plan and package invalidation occurs which requires a rebind to be performed

Audit class Events that are traced


1. Access attempts that DB2 denies because of inadequate authorization. This class is the default.
2. Explicit GRANT and REVOKE statements and their results. This class does not trace implicit grants and revokes.
3. CREATE, ALTER, and DROP statements that affect audited tables, and the results of these statements. This class traces the dropping of a table that is caused by DROP TABLESPACE or DROP DATABASE and the creation of a table with AUDIT CHANGES or AUDIT ALL. ALTER TABLE statements are audited only when they change the AUDIT option for the table.
4. Changes to audited tables. Only the first attempt to change a table, within a unit of recovery, is recorded. (If the agent or the transaction issues more than one COMMIT statement, the number of audit records increases accordingly.) The changed data is not recorded, only the attempt to make a change is recorded. If the change is not successful and is rolled back, the audit record remains; it is not deleted. This class includes access by the LOAD utility.

Accesses to a dependent table that are caused by attempted deletions from a parent table are also audited. The audit record is written even if the delete rule is RESTRICT, which prevents the deletion from the parent table. The audit record is also written when the rule is CASCADE or SET NULL, which can result in deletions that cascade to the dependent table.

5. All read accesses to tables that are identified with the AUDIT ALL clause. As in class 4, only the first access within a DB2 unit of recovery is recorded. References to a parent table are also audited.
6. The bind of static and dynamic SQL statements of the following types:
 - INSERT, UPDATE, DELETE, CREATE VIEW, and LOCK TABLE statements for audited tables. Except for the values of host variables, the audit record contains the entire SQL statement.
 - SELECT statements on tables that are identified with the AUDIT ALL clause. Except for the values of host variables, the audit record contains the entire SQL statement.
7. Assignment or change of an authorization ID because of the following reasons:
 - Changes through an exit routine (default or user-written)
 - Changes through a SET CURRENT SQLID statement
 - An outbound or inbound authorization ID translation
 - An ID that is being mapped to a RACF ID from a Kerberos security ticket
8. The start of a utility job, and the end of each phase of the utility.

Suggested Audit traces on DB2 for z/OS DB2 Common Criteria

- IFCIDs for Audit
- Accounting
 - 0003 successful access
- Audit
 - 0140: Audit all authorization failures
 - 0141: Audit all grants & revokes
 - 0142: Audit DDL Create / Alter / Drop
 - 0143: Audit First Write
 - 0144: Audit First Read
 - 0145: Audit DML Statement
 - 0314: Authorization Exit Parameters
- Performance
 - 0004: Trace Start
 - 0005: Trace Stop
 - 0023: Utility Start
 - 0024: Utility Change
 - 0025: Utility End
 - 0106: System Parameters
 - 0247: input host variables
 - 0350: SQL Statement



Suggested Audit traces – The “Bare Bones Minimum”

→ DB2 security audit suggestions:

- Catalog table queries
- Audit class 1, 2, 3
 - 0140: audit all authorization failures
 - 0141: audit all grants & revokes
- DB2 9 audit class 10: audit trusted context
 - 0269: establish trusted connection and switch user
 - 0270: CREATE & ALTER TRUSTED CONTEXT statements
- Performance
 - 0004: Trace Start
 - 0005: Trace Stop
 - 0106: System Parameters

Auditing utilities which act outside of DB2

The audit gap

- When a 3rd party unload is executed against the DB2 VSAM data sets instead of through DB2, the IBM audit record has no knowledge of data access. However, the 3rd party utility “history” table will contain the date and time of the utility with the relevant utility id. The utility activity at run time is kept in another “in-flight” table. But the records are deleted upon completion of the utility.

Closing the Gap

- A DB2 trigger is deployed on the “in-flight” table that checks against the list of sensitive tablespaces. If it is one of our audited objects, the after trigger fires to insert this information into the DBA version of the in-flight table.

→ CREATE TRIGGER

→ xxxx.trigger name

→ AFTER

→ INSERT

→ ON xxxxx.DBA_UTILITY_INFLIGHT

→ REFERENCING

→ NEW AS N

→ FOR EACH ROW

→ MODE DB2SQL

→ WHEN (N.NAME2 IN ('TS1', 'TS2', 'TS3', 'TS4','TS5')) BEGIN

→ ATOMIC INSERT INTO xxxxx. DBA_UTILITY_INFLIGHT (UTILID, NAME1, NAME2, KIND,

→ PARTITION, UTILNAME, SHRLEVEL, STATUS, XCOUNT, DDNAME,

→ BLOCKS, ORIG_STATUS, EXTRBA, STATE) VALUES (N.UTILID, N.NAME1,

→ N.NAME2, N.KIND, N.PARTITION, N.UTILNAME, N.SHRLEVEL,

→ N.STATUS, N.XCOUNT, N.DDNAME, N.BLOCKS, N.ORIG_STATUS, N.EXTRBA,

→ N.STATE) ; END

- In DBA_UTILITY_INFLIGHT, the record will not be deleted and so the audit trail is left in tact. A separate query of this table will yield all 3rd party unload activity.

Audit Trace Overhead

- The performance impact of auditing is directly dependent on the amount of audit data produced. When the audit trace is active, the more tables that are audited and the more transactions that access them, the greater the performance impact. The overhead of **audit trace is typically less than 5% but workload dependent**.
- When estimating the performance impact of the audit trace, consider the frequency of certain events. For example, security violations are not as frequent as table accesses. The frequency of utility runs is likely to be measured in executions per day. Alternatively, authorization changes can be numerous in a transaction environment.
 - Following is the summary of results of the DB2 V8 Audit trace measurements :

The measurements were done with Audit trace class(*) on and all the tables in the workload were enabled for 'Audit All'.

For OLTP measurement with distributed IRWW SQL CLI workload with 9 Tables, 3 PI, 8 NPI and 7 transactions running at 493 transactions per second, the **DB2 Class 2 CPU increase was +7.2%**.

For Utility measurements with LOAD, Rebuild Index, Reorg Table, Reorg Index utilities using 1 Table, 10 partitions, 1 PI and 5 NPI, there was no measurable CPU increase.

- Weigh auditing requirements against workload and anticipated impacts to application service levels and performance objectives carefully.
- Don't underestimate impact on SMF activity and associated overhead

V9 Trace Extensions – START TRACE

→ Qualifications by:

- LOC
 - Location-Name
 - LUName
 - IPAddress
- PLAN
- PACKAGE
 - PKGLOC
 - PKGCOL
 - PKGPROG
- Workstation Identifiers
 - USERID
 - APPLNAME
 - WRKSTN
- Miscellaneous
 - CORRID
 - CONNID
 - ROLE

→ Exclude by:

- LOC
 - XLOC
- PLAN
 - XPLAN
- PACKAGE
 - XPKGLOC
 - XPKGCOL
 - XPKGPROG
- Workstation Identifiers
 - XUSERID
 - XAPPLID
 - XWRKSTN
- Miscellaneous
 - XCORRID
 - XCONNID
 - XROLE

V9 Trace Extensions - Wildcards

→ Tracing threads using the * wildcard:

–You can use the wildcard suffix, “*” to filter threads. For example, if you specify “-START TRACE PLAN (A,B,C*)”, DB2 will trace, and then return A, B, CDE, CDEFG, CDEFGH, and so on. It will trace threads “A”, “B” and all threads starting with “C”.

→ Tracing threads using the positional, (_) wildcard:

–You can utilize the positional wildcard, which is represented by the, “_” character, to trace threads when you want the wildcard in the middle, or when you want to trace threads of a specific length. For example, if you specify “-START TRACE PLAN (A_C)”, all threads will be traced that are three characters that have “A” as the first character, and “C” as the third.

→ Tracing multiple threads at once using wildcards:

–You also have the option of tracing multiple threads based on multiple trace qualifications. For example, you can specify, “-START TRACE PLAN (A*, B*, C*)” to simultaneously trace ALL threads for plan that start with “A”, “B”, and “C”. The wildcard character, “*” will trace all threads.

–You have the ability to filter multiple threads at the same time, setting specific criteria for the trace: For example, you can specify “-START TRACE PLAN (A) USERID (B)”. This will trace the threads where the plan thread is A, and the user ID is B.

V9 Trace Extensions – Some Restrictions

- When tracing threads, you can only specify more than one thread criteria for one filter per “-START TRACE” command.
 - For example, you can specify “-START TRACE PLAN (A,B) USERID (B) WRKSTN (E),” but you cannot specify “-START TRACE PLAN (A, B) USERID (A, B) WRKSTN (E).
- If you use one or no values for PLAN, AUTHID, or LOCATION, the START TRACE command starts a single trace. If you use multiple values for PLAN, AUTHID, or LOCATION, the command starts a trace for each plan, authorization ID, or location. There can be a total of up to 32 traces going at one time (**all trace types**).
- You must use a privilege set of the process that includes one of the following privileges or authorities:
 - TRACE privilege
 - SYSOPR authority
 - SYSCTRL authority
 - SYSADM authority

DSN1SMFP offline utility

- The DSN1SMFP utility processes DB2 trace data into reports.
- DSN1SMFP accepts data that SMF collects in standard SMF format and produces from one to fifteen reports. DSN1SMFP accepts all SMF record types, but it processes only type 101 (DB2 Accounting) and 102 (DB2 Performance) records.
- DSN1SMFP checks each type 101 and 102 record for DB2 audit trace types of these DB2 IFCIDs:
 - 003: Accounting - DDF Data by Location (security-relevant fields only)
 - 004: Trace Start
 - 005: Trace Stop
 - 023: Utility Start
 - 024: Utility Change
 - 025: Utility End
 - 106: System Parameters (security-relevant fields only)
 - 140: Audit Authorization Failures
 - 141: Audit DDL Grant/Revoke
 - 142: Audit DDL Create/Alter/Drop
 - 143: Audit First Write
 - 144: Audit First Read
 - 145: Audit DML Statement
 - 350: SQL Statement

DSN1SMFP – Sample Report Outputs

IFCID – 141 Audit Grant/Revoke Report

```
GRANTOR : SYSADM          REASON : SYSADM          RETURN: 0000000000
OBJECT  : STORAGE GROUP   OPTIONS: X'0400000000000000'
SQL STMT: GRANT USE OF STOGROUP DSN86810 TO PUBLIC
```

IFCID – 106 System Parameters Report

```
COMMON CRITERIA ENVIRON : NO
SYSADM ID 2              : SYSADM
ENABLE DB2 AUTHORIZATION: YES
PACK AUTH CACHE         : 0000032768
ONL SYSPARM CORID       :

MISCELLANEOUS INSTALLATION PARAMETERS
DDL REGISTRATION FLAG: X'30'
SITE TYPE             : LOCAL
CACHE DYNAMIC SQL    : NO
DBADM CREATE VIEW    : NO
ONL SYSPARM USER ID  :

INSTALL SYSADM : SYSADM
SYSOPER ID     : SYSOPR
AUTH. CACHE SIZE: 01024
EDM STMT CACHE : 0005120000
ONL SYSPARM TIME: 08:26:40

DEFAULT USERID : IBMUSER
SYSOPER ID 2   : SYSOPR
HOP SITE AUTHORIZ.: YES
ONL SYSPARM TYPE : N/A
```

OMEGAMON XE for DB2 Performance Monitor/Expert for z/OS

- Real-time monitoring
 - Threads and Statistics monitoring
 - DB2 Connect monitoring
 - Object Analysis
 - Data Sharing/Sysplex data (DB2Plex data)
- Near-term history
- Trace collection (**also as part of the PWH process support**)
- Reporting
 - Accounting, Statistics, SQL Activities, Locking, I/O Activity, Audit, Utilities, Record Trace
 - Executable as separate jobs or via PWH process engine
- Performance Warehouse with expert analysis support
- Buffer Pool Analysis, expert advice, and simulation (**only with the OMEGAMON XE for DB2 Performance Expert**)

DB2 OMEGAMON Performance Expert Audit Report Set

- Not strictly a performance report.
- Reports information about usage of auditable objects and authorization management.
 - Authorization changes
 - Authorization control (GRANTS and REVOKEs of privileges)
 - Authorization failures
 - DML statements against auditable DB2 tables at bind time
 - DDL operations against auditable DB2 tables
 - Read/write access against auditable DB2 tables
 - Utility executions against auditable DB2 tables
- Traces show individual events.
- Reports show audit information for an aggregation of DB2PE identifiers, e.g. primauth-planname-objects.

The OMPE "File" Report
command is used to create DB2
Load compatible record formats

OMPE "File" report
commands

OMPE Audit
Detail Report

```

MSG. ID.      DESCRIPTION
-----
FPEC2001I    COMMAND INPUT FROM DDNAME SYSIN
              AUDIT
              REPORT
              LEVEL (DETAIL)
              TYPE (DDL DML)
              DDNAME (AUDITDD)
              FILE
              TYPE (DDL)
              DDNAME (AUFILDD1)
              FILE
              TYPE (DML)
              DDNAME (AUFILDD2)
              FILE
              TYPE (AUTHFAIL)
              DDNAME (AUFILDD3)
              EXEC
    
```

```

LOCATION: NDCDB203                OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V3)                PAGE: 1-1
GROUP: N/P                      AUDIT REPORT - DETAIL                REQUESTED FROM: NOT SPECIFIED
MEMBER: N/P                    ORDER: PRIMAUTH-PLANNAME                TO: NOT SPECIFIED
SUBSYSTEM: DSNC                SCOPE: MEMBER                ACTUAL FROM: 09/06/06 01:47:43.60
DB2 VERSION: V8                TO: 09/06/06 01:49:38.83
PRIMAUTH CORRNAME CONNTYPE
ORIGAUTH CORRNMBR INSTANCE
PLANNAME CONNECT                TIMESTAMP TYPE                DETAIL
-----
SYS248  SYS248  DB2CALL  01:47:43.60 DML  TYPE : 1ST READ
SYS248  'BLANK'  BF5CF720228D  DATABASE: SYS248SA  TABLE OBID: 5
ETIPLAN1 DB2CALL  PAGESET : SYS248TS  LOG RBA : X'000000000000'

SYS248  SYS248  DB2CALL  01:48:22.56 DML  TYPE : 1ST WRITE
SYS248  'BLANK'  BF5CF7454387  DATABASE: SYS248SA  TABLE OBID: 5
ETIPLAN1 DB2CALL  PAGESET : SYS248TS  LOG RBA : X'00036FBEA220'

SYS248  SYS248  DB2CALL  01:48:22.56 DML  TYPE : 1ST WRITE
SYS248  'BLANK'  BF5CF7454387  DATABASE: SYS248SA  TABLE OBID: 5
ETIPLAN1 DB2CALL  PAGESET : SYS248TS  LOG RBA : X'00036FBEA3DA'
    
```

Invoking the DB2 load utility to populate the DB2 Performance DB with Audit data.

Load Control sample statements located in RKO2SAMP

```

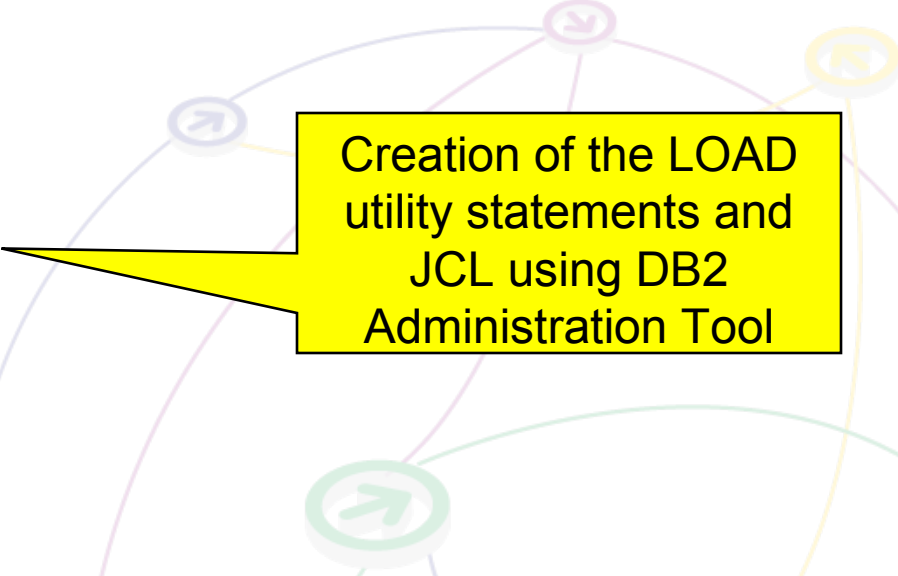
File Edit Edit Settings Menu Utilities Compilers Test Help
EDIT      SYS248.SPFTEMP2.CNTL      Columns 00001 00072
000052 LOAD INDDN SYSREC
000053 RESUME NO
000054 REPLACE
000055 INTO TABLE DB2PMFAUDT_DML
000056 WHEN (251;259) = 'DML N'
000057 (DB2PM_REL POSITION(3) SMALLINT,
000058 DB2_REL POSITION(9) CHAR(2),
000059 LOCAL_LOCATION POSITION(11) CHAR(16),
000060 GROUP_NAME POSITION(27) CHAR(8),
000061 SUBS_ID POSITION(35) CHAR(4),
000062 MEMBER_NAME POSITION(39) CHAR(8),
000063 NET_ID POSITION(47) CHAR(8),
000064 LUNAME POSITION(55) CHAR(8),
000065 INSTANCE_NBR POSITION(63) CHAR(12),
000066 LUW_SEQNO POSITION(75) SMALLINT,
000067 REQ_LOC_NAME POSITION(87) CHAR(16),
000068 ENDUSER POSITION(103) CHAR(16),
000069 WSNNAME POSITION(119) CHAR(18),
Command ==> Scroll ==> CSR_
F1=Help F2=Split F3=Exit F5=Rfind F6=Rchange F7=Up
F8=Down F9=Swap F10=Left F11=Right F12=Cancel
    
```

```

DB2 Admin ----- DSNB Specify Utility Options - LOAD ----- 08:20
Option ==>
Top of data
Execute utility on table SYS248.DB2PMFAUDT_DML
using the following options:

Utility ID ==> LOADAUD (Name identifying this utility to DB2)
Unloaded Data ==> SYS248.OMPE.AUFIL2 (Name of data set containing unloaded data)
Unloaded How? ==> U (U=Unload Utility, R=Reorg Utility)
Table/Col Info ==> CANDLET.XEGA.DEMOMVS.RKO2SAMP (DG0XLDML) (Name of data set containing table/column info)
RESUME ==> NO (Yes/No, load recs into non-empty table space)
SHRLEVEL ==> (None/Change, concurrent table space access)
REPLACE ==> YES (Yes/No, empty table space/index before load)
COPYDDN1 ==> (DDname identifying primary copy data set)
COPYDDN2 ==> (DDname identifying backup copy data set)
RECOVERYDDN1 ==> (DDname identifying primary ds @ recovery site)
RECOVERYDDN2 ==> (DDname identifying backup ds @ recovery site)

TABLE ALL ==> (Yes/No, info for all columns in table space)
F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIND F6=RCHANGE
F7=UP F8=DOWN F9=SWAP F10=LEFT F11=RIGHT F12=RETRIEVE
    
```



Creation of the LOAD utility statements and JCL using DB2 Administration Tool

A view of the audit data stored in the OMPE performance warehouse using DB2 Control Center

Log RBA can be used to locate details about other actions for the LUW

Open Table - DB2PMFAUDT_DML

DSNC - DSNC - AUDITDB - SYS248 - DB2PMFAUDT_DML

E	PRIMAUTH	ORIGAUTH	TIMESTAMP	IFCID	DATABASE_DBID	PAGESET_OBID	TABLE_OBID	DATABASE_NAME	PAGESET_NAME
	SYS248	SYS248	Sep 6, 2006 1:47:41 AM 602771	144	307	2	5	SYS248SA	SYS248TS		
	SYS248	SYS248	Sep 6, 2006 1:48:22 AM 560444	143	307	2	5	SYS248SA	SYS248TS	00036FBEA220	EA220
	SYS248	SYS248	Sep 6, 2006 1:48:22 AM 564498	143	307	2	5	SYS248SA	SYS248TS	00036FBEA3DA	EA3DA
	SYS248	SYS248	Sep 6, 2006 1:48:28 AM 130075	144	307	2	5	SYS248SA	SYS248TS		
	SYS248	SYS248	Sep 6, 2006 1:48:58 AM 571847	143	307	2	5	SYS248SA	SYS248TS	00036FBEEA62	EA62
	SYS248	SYS248	Sep 6, 2006 1:48:58 AM 579028	143	307	2	5	SYS248SA	SYS248TS	00036FBEEAC1C	EAC1C
	SYS248	SYS248	Sep 6, 2006 1:49:06 AM 253828	144	307	2	5	SYS248SA	SYS248TS		
	SYS248	SYS248	Sep 6, 2006 1:49:38 AM 826482	143	307	2	5	SYS248SA	SYS248TS	00036FBEEADD6	EADD6
	SYS248	SYS248	Sep 6, 2006 1:49:38 AM 831367	143	307	2	5	SYS248SA	SYS248TS	00036FBEEB000	EB000
	SYS248	SYS248	Sep 6, 2006 1:49:38 AM 838245	143	307	2	5	SYS248SA	SYS248TS	00036FBEEB1BA	EB1BA

Commit Roll Back Filter Fetch More Rows

Automatically commit updates 10 row(s) in memory Close Help

Table OBD will require join with DB2 Catalog SYSTABLES for meaningful reporting

Limitations of the audit trace

- The audit trace does not record everything, as the following list of limitations indicates:
 - The auditing that is described in this information takes place only when the audit trace is on.
 - The trace audits only the tables that you specifically choose to audit.
- The trace does NOT capture before/after change data because the DB2 log records this information.
 - If an agent or transaction accesses a table more than once in a single unit of recovery, the audit trace records only the first access.
- The audit trace does not audit some utilities. The trace audits the first access of a table with the LOAD utility, but it does not audit access by the COPY, RECOVER, and REPAIR utilities. The audit trace does not audit access by stand-alone utilities, such as DSN1CHKR and DSN1PRNT.
- You cannot audit the catalog tables because you cannot create or alter catalog tables.
- 3rd Party DB2 utilities (run outside of DB2) will not be caught with the AUDIT CLASS 8
- Dynamic SQL host variable data not collected
- This auditing coverage is consistent with the goal of providing a moderate volume of audit data with a low impact on performance. However, when you choose classes of events to audit, consider that you might ask for more data than you are willing to process.
- Depending on AUDIT classes active, and workload mix, significant increases in SMF activity might be experienced. One customer scenario, with CLASS (1-6) a 12% increase in SMF was observed.

Separation of Roles and Responsibilities

→ DB2 trace based processes are managed by DBA's

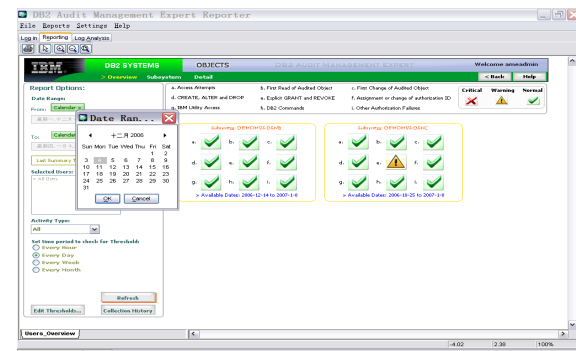
- The DBA's are responsible for generating audit data with which they are in turn audited, this constitutes a significant security risk and exposure
- Trace data collection can be interfered with or turned off completely
 - DBA can issue –DSN Stop Trace
 - Use IFASMFDMF to selectively filter SMF data based on timestamp
 - Use DB2PM (Or Equivalent) filter such as DATE/TIME/EXCLUDE to filter selected records
- **Having the DBA involved in the collection of audit data is viewed as weak from a compliance and control perspective**

→ Security and Auditors with system privileges

- Also viewed as problematic from a compliance perspective
- Requires additional technical skills not within their core competencies
- Misuse of privileges without coordination can result in performance and availability issues
 - Turning on traces without proper filtering to reduce overhead or quantity of trace data collected
 - Altering objects to AUDIT without ensuring that plan/package invalidation is not an issue

Audit Management Expert - Monitor and Audit

- ➔ **Helps auditors answer:**
 - Who, What, Where, Why, When, How
- ➔ **Centralizes the audit data**
 - Pulls together disparate data sources from all the systems into a central repository
- ➔ **Automates auditing process**
 - Eliminates all home grown processes
- ➔ **Creates segregation of duties**
 - Gives auditors the business activity collected without being reliant on the technical personnel they need to monitor
- ➔ **Flexible Reporting**
 - Drill down from overview to detail for forensic analysis



Audit Management Expert Overview

- Auditors will be able to Access:
 - SELECT, INSERT, UPDATE, and DELETE activity by user or by object
 - **SQL Text and Host Variable value for each statement**
 - **Row count that SQL statement affects**
 - CREATE, ALTER, and DROP operations against an audited object
 - Explicit GRANT and REVOKE operations
 - Utility access to an audited object
 - DB2 commands entered
 - Assignment or modification of an authorization ID
 - Authorization failures
- **Provides auditors with flexible options for examining the data in the audit repository**
 - Audit Trace Data, **Audit SQL Collector (ASC)**, Log Analysis data
 - V2.1 no longer needs to alter objects to 'AUDIT ALL' for read/update
 - DB2 Catalog Objects can now be audited for SQL read/update

Security and separation of roles

- Supporting internal and external auditors in collection and reporting of DB2 audit data
 - Does not require auditors to be DB2 defined users within the monitored DB2 system(s)
 - Does not require the auditors to log on to the operating system where the monitored system is running
 - Does not require extensive interaction between the auditor and the system support personnel (DBA/Sys admin)
- Auditor will not be able to directly manipulate any DB2 resources
- Provide complete visibility of all auditable objects to an administrator level user
- Provide controls for limiting visibility to auditors of auditable objects
- Removes DBA from audit data collection process. With V2.1 removes the “ALTER for AUDIT” requirement

DB2 Audit Management Expert Components

→ Audit server

- Started task or batch job
- central control point for an Audit Management Expert network
- single audit server can support data collection from multiple agents on multiple z/OS systems

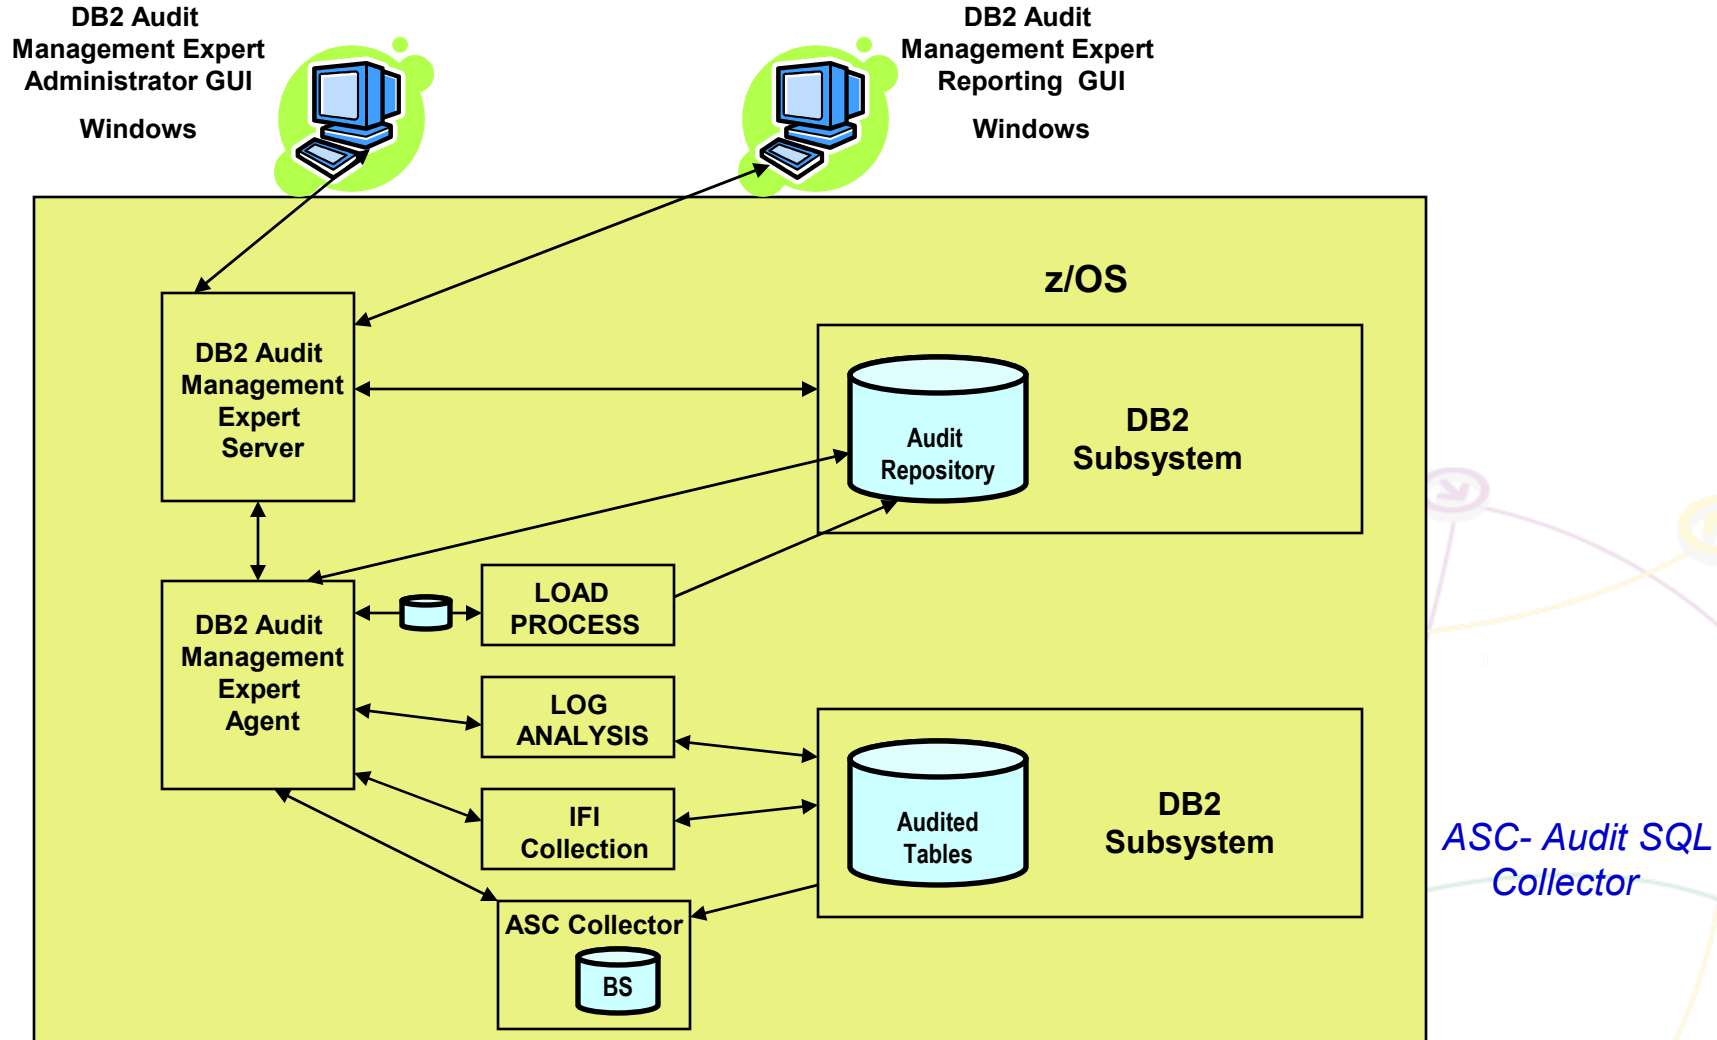
→ Agent

- Started Task or batch job
- responsible for communications in an Audit Management Expert environment
- acts as a "container" to run the various collectors
- One per DB2 to Audit

→ CLIENT User interfaces

- Audit Management Expert Reporter
- Audit Management Expert Administration
- Windows

DB2 Audit Management Expert Architecture



ASC- Audit SQL
Collector

DB2 Audit Management Expert Profiles

→ Profiles are created/maintained via **Administration UI**

– Collection Profile

- records the details for what audit data is stored to the Audit Management Expert repository

– Agent Profile

- Select ASC collection method
- Configure General settings
 - Retention count, interval length
- DB2 Load utility parameters
- Define Job cards for load and log analysis

– User Profile

- contain information specific to an individual Audit Management Expert user such as: the user type, configurable privileges, and associated user groups

AME and Enterprise Wide Auditing - Challenges

- Existing appliance technology based on data feeds from primarily 2 sources
 - Event log from DB2 trace events written to SMF (agent)
 - Network “Sniffer” implementation (appliance)
- Restrictions and challenges with DB2 Trace versus a superior low overhead data collection approach with AME’s ASC
- Network traffic based audit feeds challenged by
 - Encrypted Data Streams
 - Local Attachments (Batch, TSO, etc.)
 - Stored Procedures
 - Performance impact to network throughput due to indiscriminate examination of all network flows
- Strong requirement to view and manage Audit events across the enterprise from a single UI

AME Extract File Enhancement via maintenance stream (PTFs UK41519, UK41521, UK41523)

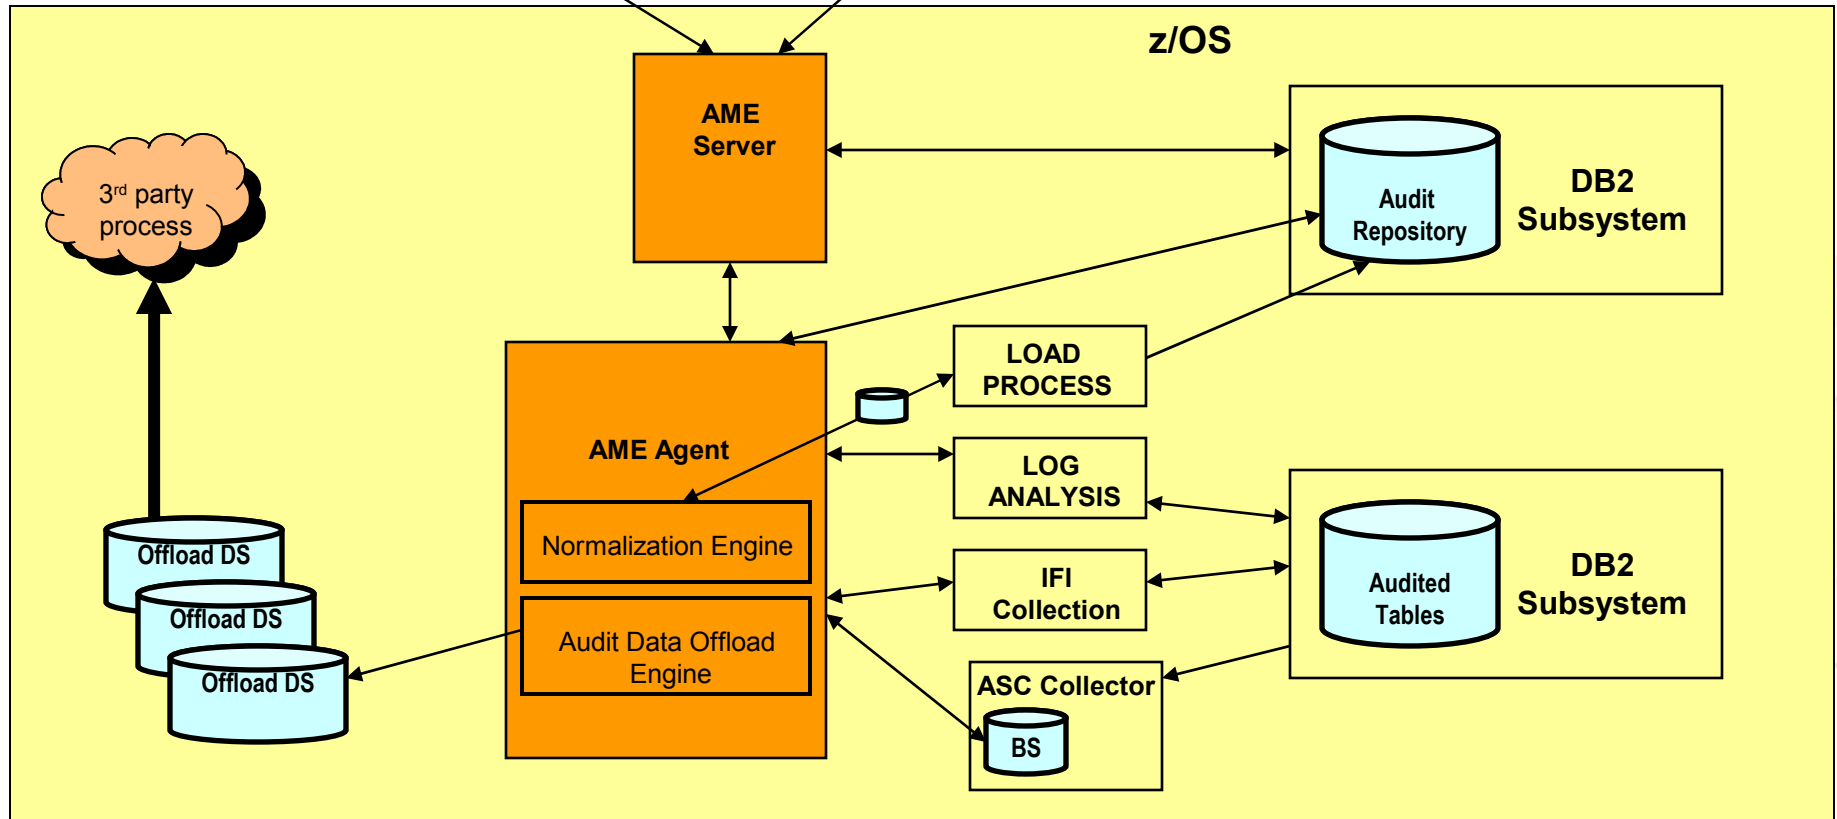
- AME will provide an option to generate audit log data sourced from either DB2 Trace or ASC (Audit SQL Collector) in an documented extract file format.
- Extract files will be standard physical sequential datasets.
- Exploiters will be responsible for transporting (via secure FTP for example) data to appliance server environment.
- Management of Extract files (archiving, deleting, etc.) will be the responsibility of exploiters
- Data will be not be aggregated (normalized), this is to reduce overhead of data collection.
- Static SQL statement collection will be optional, this is to avoid the overhead of accessing the catalog with static SQL statement number to extract SQL statement text.
- Exploiters to include
 - Tivoli Consul Insight Manager (coming Q1 2010)
 - Tizor - Mantra
 - Imperva - SecureSphere
 - Others anticipated at a future date

DB2 Audit Management Expert Architecture Dual Mode

DB2 Audit
Management Expert
Administrator GUI
Windows



DB2 Audit
Management Expert
Reporting GUI
Windows



Alerts

- Real-time alert monitoring
- Exceptions outside of expected business process
- Immediate triaging & response
- Easy integration with 3rd-party IT ecosystems: SYSLOG, SNMP, Email...

The screenshot displays a software interface for monitoring alerts. The main window is titled 'Alerts (filtered)' and shows a table with one alert: ID 2004, timestamp 08:49:02, and count 11. A detailed view of this alert is shown on the right, titled 'Alert 2004: z/OS Security Policy'. It includes a bar chart showing 11 violations over time. Below the chart is a table 'Alert aggregated by:' with columns 'Distinct value for:' and 'Value'. The table lists 'Custom Rule' (z/OS Security Policy), 'Server Group' (DB2 Mainframe), and 'Source IP'. At the bottom, another 'Violations:' table lists the user 'ajcuser'.

Alert 2004: z/OS Security Policy

Actions: None
Policy: z/OS Security Policy

Aggregated from 08:49:01 (0 hour(s), 2 minute(s)),

Alert aggregated by:

Distinct value for:	Value
Custom Rule	z/OS Security Policy
Server Group	DB2 Mainframe
Source IP	

Violations:

User	OS User
ajcuser	
ajcuser	
ajcuser	

Violations:

User	OS User
ajcuser	

Event 6909825234996570923: Custom Rule Violation !

Key	Value
Violation Description	z/OS Security Policy
Violated Item	Custom Violation

Event Details:

Event Time	Server Group	Service	Application
July 31, 2008 6:24:39 PM	DB2 Mainframe	DB2	Default DB2 Application

Connection	User	DB Application	OS User	OS Host
:0 → :0	ajcuser			

Affected Rows	Response Size	Response Time
0	0 Records	0 msec.

Error Code	Error Message

START TRACE (AUDIT)CLASS (3)RMID (*)DEST (OPX)PLAN (*)AUTHID (*)IFCID (*)BUFSIZE (16)TDAT A (CORRELATION DISTRIBUTED)

Viewing the Audit Logs:

Provides all the details including: date and time , database user name and parsed query

Date/Time	Database Username	Parsed Query	Log Collector
7/21/08 3:45:57 PM	csliivi	display log	z/OS
7/21/08 3:45:58 PM	csliivi	display log	z/OS
7/21/08 3:46:01 PM	csliivi	display log	z/OS
7/21/08 3:46:02 PM	csliivi	display log	z/OS
7/31/08 5:12:09 PM	csliivi	display log	z/OS
7/21/08 3:32:55 PM	csliivi	display utility(*)	z/OS
7/21/08 3:37:41 PM	csliivi	display utility(*)	z/OS
7/21/08 3:38:14 PM	csliivi	display utility(*)	z/OS
7/21/08 3:45:57 PM	csliivi	display utility(*)	z/OS
7/21/08 3:45:58 PM	csliivi	display utility(*)	z/OS
7/21/08 3:46:01 PM	csliivi	display utility(*)	z/OS
7/21/08 3:46:02 PM	csliivi	display utility(*)	z/OS
7/21/08 3:46:02 PM	csliivi	display utility(*)	z/OS
7/31/08 5:12:09 PM	csliivi	display utility(*)	z/OS
7/31/08 5:15:51 PM	csliivi	delete from dsn8710.act where actno=?	z/OS
7/31/08 5:15:51 PM	csliivi	insert into dsn8710.act (actno,actkwd,actdesc) values(?,?,?)	z/OS
7/31/08 5:15:51 PM	csliivi	update dsn8710.act set actdesc=? where actno=?	z/OS
7/21/08 3:28:15 PM	ajcuser	start trace (audit)class (?)rmid (*)dest (opx)plan (*)authid (*)ifcid (*)bufsize (?)tdata (correlation distributed)	z/OS
7/21/08 3:42:23 PM	ajcuser	start trace (audit)class (?)rmid (*)dest (opx)plan (*)authid (*)ifcid (*)bufsize (?)tdata (correlation distributed)	z/OS
7/21/08 3:48:27 PM	ajcuser	start trace (audit)class (?)rmid (*)dest (opx)plan (*)authid (*)ifcid (*)bufsize (?)tdata (correlation distributed)	z/OS


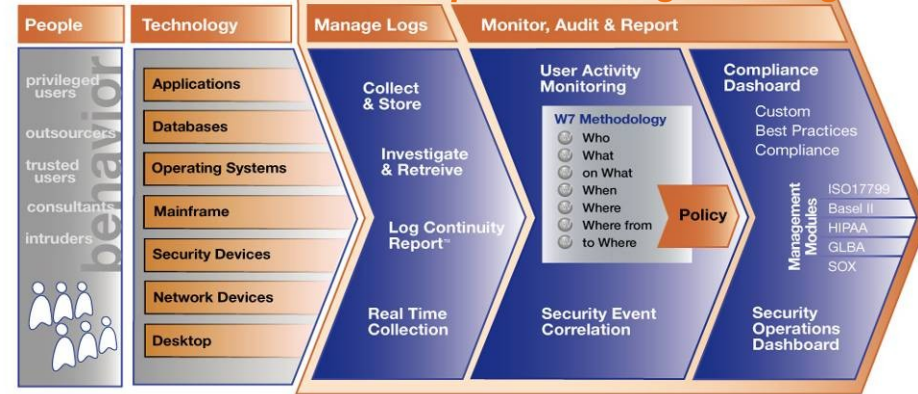
Tivoli Compliance Insight Manager

Tivoli Compliance Insight Manager provides an enterprise security compliance dashboard with in-depth privileged user **monitoring** capabilities, all powered by a comprehensive log and audit trail collection capability

Key Features

- Compliance management modules and regulation-specific reports
- Unique ability to monitor user behavior, including PUMA (Privileged User Monitoring and Audit) reporting
- Broadest, most complete log and audit trail capture capability
- W7 log normalization translates your logs into business terms
- Easy ability to compare behavior to regulatory and company policies – auditors no longer need RACF expertise to monitor activities
- Enabler event source integrates the OS and mainframe database events into TCIM's enterprise compliance dashboard

Tivoli Compliance Insight Manager



TCIM – Representative Screen

Platform History Event List on Platform Z180 (z/OS) - Database GEM on Server CIFDB - Microsoft Internet Explorer

Address: http://9.142.236.76/jview/?expert=platformhistoryevents&GEMCatalog=GEM&plfdett=z%2FOS&plfdetnm=Z180&count=1573&plfdetix=38&EPR1SECatalog=EPR1SEDB&navig=Gem&navname=Gem.GemSummary&stid=1207083708203

IBM. Information Management Lotus Rational Tivoli WebSphere

Dashboard Trends Reports Regulations Policy Groups Distribution Settings

CIFDB > GEM > Platform History Events

Platform History Event List on Platform Z180 (z/OS)

Database GEM on Server CIFDB

Setup:

Start time: Month: June, Day: 20, Year: 2007, Hour: 8, Min: 10
 End time: Month: March, Day: 31, Year: 2008, Hour: 13, Min: 0

Execute Reset

Time zone: Event time zone

Severity	Date / Time	#	What (detail)	Where (detail)	Who (detail)	Where from (detail)	On what (detail)	Where to (detail)
10	Wed Jun 20 2007 06:10:36 GMT+00:00	1	Verify : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.DSN8S91E	Z180 (z/OS)
50	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Modify : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
50	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Modify : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP2	Z180 (z/OS)
50	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Change : Auditlog / Success	Z180 (z/OS)	System	Z180 (z/OS)	SYSTEM : Z180 / SMF	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Read : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Write : Databasespace / Success	Z180 (z/OS)	.SYSOPR	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Write : Dobject / Success	Z180 (z/OS)	.SYSOPR	Z180 (z/OS)	DBOBJECT : - / -	Z180 (z/OS)
10	Wed Jun 20 2007 06:10:28 GMT+00:00	1	Write : Databasespace / Success	Z180 (z/OS)	.CRMBFT1	Z180 (z/OS)	DBTABLESPACE : DB9G / DSN8D91A.XEMP1	Z180 (z/OS)

Regulatory Resource Center

start 10:29:50 - AT&T Ne... Vacation Planner - 2... Session B - [24 x 80] Platform History Eve... IBM Lotus Sametime ... meadowsr@us.ibm.c... 100% 5:02 PM Tuesday 4/1/2008

Redbook on Audit and Encryption on DB2 for z/OS – SG24-7720

IBM

Draft Document for Review March 4, 2009 6:04 pm


SG24-7720-00

Securing and Auditing Data on DB2 for z/OS

Prepare for the threat from within and
without

Comply with IBM Data Server
Security Blueprint

Extend the skills of data
professionals



Paolo Bruni
Felipe Bortoletto
Thomas Hubbard
Ernest Mancill
Hennie Mynhardt
Shuang Yu

Redbooks

ibm.com/redbooks



Summary

→ Take Back Control with IBM Data Governance solutions :

- Transform your information from a Liability into your most strategic, valuable Asset
- Help manage business risk by enforcing security, audit, privacy and policy controls
- Lower operational costs by optimising data management, retention and archiving
- Increase profitability by enabling more accurate business intelligence
- Increase management's confidence in making more informed decisions based on quality and more complete data
- Increase customer satisfaction and retention through targeted advertising and up/cross selling

→ Software, Hardware and Expertise.

- Information Management - the most complete end-to-end Data Governance software solutions
- zSeries the ultimate platform to govern your enterprise data
- IBM Industry Data Models as a fast-start, best practice and help with industry compliance
- GBS - Expertise and skills from DG readiness assessments to solution implementation.