DB2 9.5 What's New

weberg@de.ibm.com
IM Techsales
Günter Weber

ON DEMAND BUSINESS™

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Agenda

- ***Administration Enhancements***
  - Installation
  - Manageability
  - Compression
- ***Workload Management***
- ***Performance Enhancements***
- ***Security Enhancements***

The ultimate data server – the innovation continues

DB2 Viper 1

- ***Industry's first hybrid data server***
- ***Innovative row compression tools***
- ***Automation tools***

DB2 Viper 2

- *Reflex response.* ***Increases agility through automation capabilities***
- *Confident compliance.* ***Enhances data governance and security***
- *Transactional XML.* ***Extracts new value for operational processes***

Reflex Response – react immediately

- ***Ensure availability, optimal performance and lower cost***

- ***Introducing reflex response features such as…***
  - Integrated automated failover and backups
  - Automatic deep compression
  - Expanded memory management

"DB2 9 is the future. It's key in helping us **reduce the size of our databases by up to 83 percent.** This ultimately helps us minimize storage costs and increase performance."

—Jean Holley, CIO, Tellabs

Confident compliance – manage business risks

- ***Staying ahead of threats, regulations and policies***

- ***Simplified and enhanced governance throughout the entire data lifecycle with…***
  - Expanded audit capabilities benefiting auditors and administrators
  - Continued evolution of control, security and accountability
  - Expanded use of encryption throughout the data lifecycle

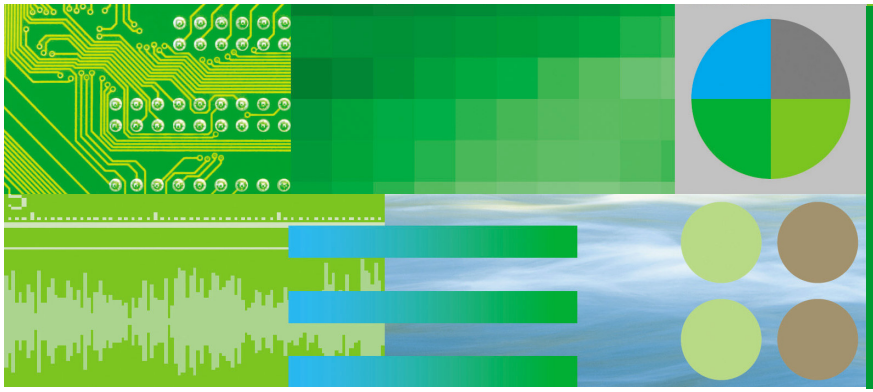

"The top challenge for 43% of CFOs is improving governance, controls and risk management."

—CFO Survey: Current state and future direction, IBM Business Consulting Services

Adapt to changing standards

- ***Extract greater value from data to improve operational processes***

- ***The industry's first hybrid data server—now with support for transactional XML—so you can…***
  - Improve transactional performance
  - Quickly integrate XML into core processes and applications
  - Reduce storage requirements

"Our customers are seeing 5 to 10 times performance improvements. We're seeing particular gains in the insurance industry, where Viper helps agents turn around quotes to customers in minutes rather than days."

—Keith Feingold, CEO, Skytide

DB2 9 for SAP – delivering unique value

- *Ease of use*
- *High performance and availability*
- *Easier troubleshooting*

"Our planned system response improvement was around 20%, whereas in reality we have observed *a 40% cut in response times* ***with DB2."***

—Peter Boegler, Solution Architect, SAP IT

IBM DB2 wins SAP Pinnacle Award for Best Partner Technology Solution — May 2006

DB2 9.5

Administration Enhancements

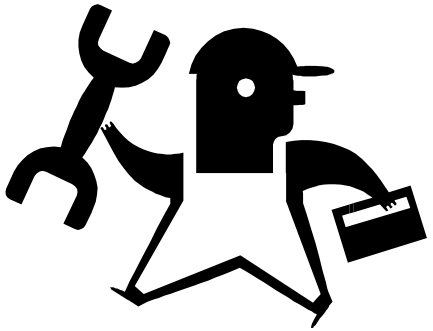

Installation Enhancements

- ***Non-root Installation (for non partitioned DBs)***
  - Capability to install and service a non-root DB2 on Unix or Linux
  - May still require that some root-based features (e.g. OS authentication or HADR) be enabled post-install by a root user (i.e. with `db2rfe -f <filename>`)
- ***Fix Pack Installation***
  - Eliminate the instance update phase
  - Two high-maintenance activities that need to be eliminated
    - Automatic execution of db2iupdt and dasupdt after a fix pack has been deployed
    - Automatic binding of packages against the database using auto rebind for all utilities

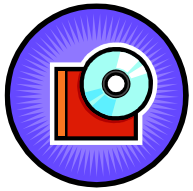

- ***Fix Pack Only***
  - No migration from one point release to another via a fix pack
  - Eliminate “side-effect” of moving to a new release just to get product fixes
  - Continue a separate maintenance stream for each Version and Release level
  - Note: A catalog migration is required from DB2 9 to this release
- ***Enhanced Unicode***
  - Create Unicode databases by default
  - Provide collation compatibility for customers moving from non-Unicode databases to Unicode databases

Manageability Enhancements

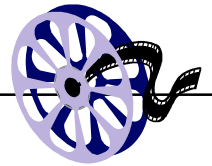

- ***Single system view for Database Partition Facility***
  - Various commands need to “view” the database as one entity rather than multiple partitions
    - Backup database, Update configuration parameters

```
BACKUP DB  mydb ON ALL DBPARTITIONNUMS
BACKUP DB mydb ON DBPARTITIONNUMS (0,1)
```

- ***Simplify Backup Administration***
  - Automatically delete expired backup images and log files
  - Add additional options for automated backup
    - Compression, incremental and delta backups, log files in backup
- ***Support Fully integrated flash copy support***
  - Automate the manual steps currently required for backup and restore with flash copy
- ***Automatic Storage Enhancements***
  - Allow user to free unused space at the end of a tablespace

Enhanced Row Compression

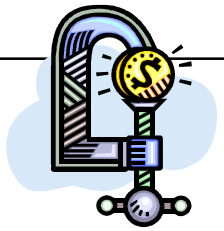

- ***Enhance row compression***
  - Automatically Compress a Table without DBA intervention
    - Once a predetermined table size is reached then automatically creates a dictionary based on sampling and data
    - All data that is inserted or updated after dictionary creation will be compressed
    - Table must be defined with COMPRESS YES to be eligible
  - Removes need for DBA to manually run INSPECT or reorganize all tables

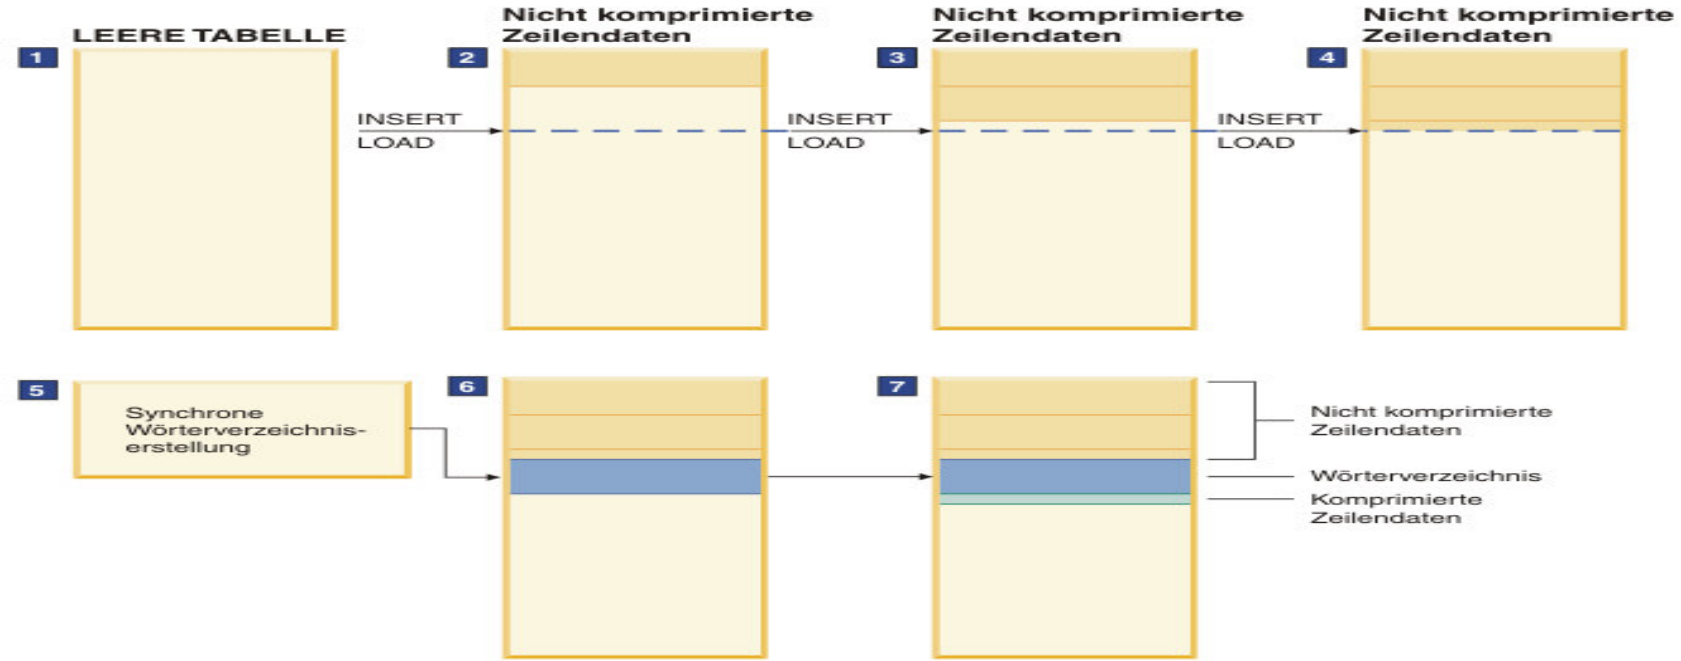

DB2 9.5

Workload Management

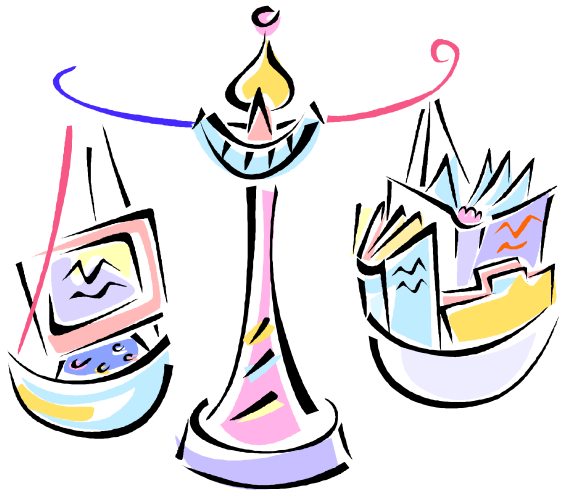

Workload Management Objectives

- ***A stable, predictable execution environment***
- ***A light-weight, granular way to monitor active work***
- ***Better resource management***
  - Be able to explicitly allocate resources amongst work (CPU priority)
  - Be able to limit excessive, unexpected resource consumption (controlling rogue queries)
- ***Better request management***
  - Be able to manage work based on its business priority
  - Be able to track performance of work
- ***End-to-end workload management solutions***

Existing DB2 9 Workload Management Solution

- ***Query Patroller (QP)***
  - A predictive governing tool for DML SQL statements based on SQL Compiler cost estimate (i.e. timerons)
  - Primary control mechanism is the establishment of different concurrency rates for different query classes
    - Each query class is based on cost estimate ranges
    - Ranges and rates defined by user

- ***DB2 Governor***
  - A reactive governing tool for applications based on system monitor information
  - User provides limit values for predefined counters and what action is to be taken when limit is reached

Existing DB2 9 Execution Environment (with QP)

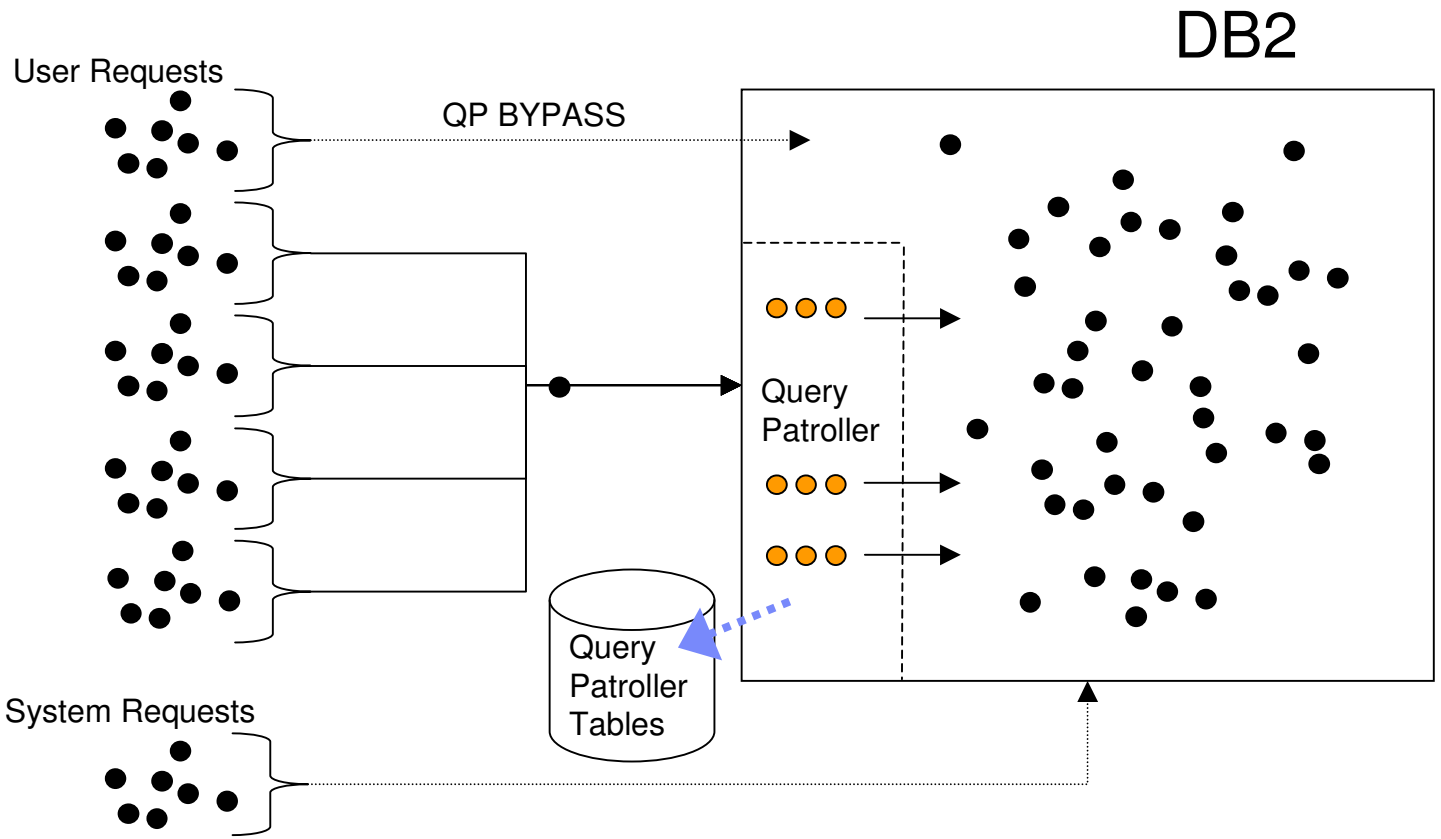

New WLM Concepts

- ***DB2 Service Class***
  - This is where all work for a database is executed
  - Serves as the primary point of resource assignment, control, and monitoring for all work executing within DB2
  - Supports a two-tier hierarchy consisting of super and sub classes
    - A super class is a logical entity providing common attributes across sub classes

- ***DB2 Workload***
  - This how incoming work is routed to a DB2 Service Class
  - Serves as the primary point of identity and control for submitters of work to the system
  - Connections are mapped to a specific workload when established
    - Re-evaluated at unit of work boundaries as required
  - Default: user and admin workload

Customization: DB2 Workloads

- ***A DB2 workload can be created to uniquely identify any connections of interest***
  - Provides the ability to independently monitor and control them (and their activities)

- ***New workloads are defined by providing mapping values for a set of connection attributes***
  - Application name
  - SYSTEM_USER
  - SESSION_USER
  - Any group of SESSION_USER
  - Any role of SESSION_USER
  - CLIENT USERID
  - CLIENT APPLNAME
  - CLIENT WRKSTNNAME
  - CLIENT ACCTNG

- ***Can grant or revoke USAGE privilege on a defined workload***
- ***Can modify the evaluation order of defined workloads***

Customization: DB2 Service Classes

- ***A DB2 service class can be created to act as a unique execution environment for any grouping of work***
  - Provides the ability to independently monitor and control this group of work
  - Can assign different resource priorities to each service class
- ***Can create two-tier hierarchy of service classes using super and sub classes***
  - Allows for more complex division of execution environment and better emulation of real world model
- ***Can modify the resources available to a service class***
  - CPU Priority
    - Controlled either through integration with AIX WLM or agent priority setting of DB2 Service Class
  - Prefetch I/O Priority

Controlling the Execution Environment

- ***Introducing new DB2 Thresholds as an automated way to enforce rules or establish limits for activities running on DB2***
  - "Trigger"-like mechanism based on predefined elements
  - Can be defined at one or more levels (e.g. Database, service class, workload, etc)
  - Activities covered include SQL statements and LOAD utility

- ***Enhanced support provided via DB2 Thresholds for common scenarios:***
  - Controlling "rogue" queries
    - Based on both predictive (prior to execution) and reactive (during execution) elements
      - Example of a Predictive Threshold: Maximum Estimated Cost
      - Example of a Reactive Threshold: Maximum Execution Time
  - Concurrency control
    - Can limit both number allowed to execute concurrently as well as maximum numbers allowed to wait for their turn

Example of Workload Management

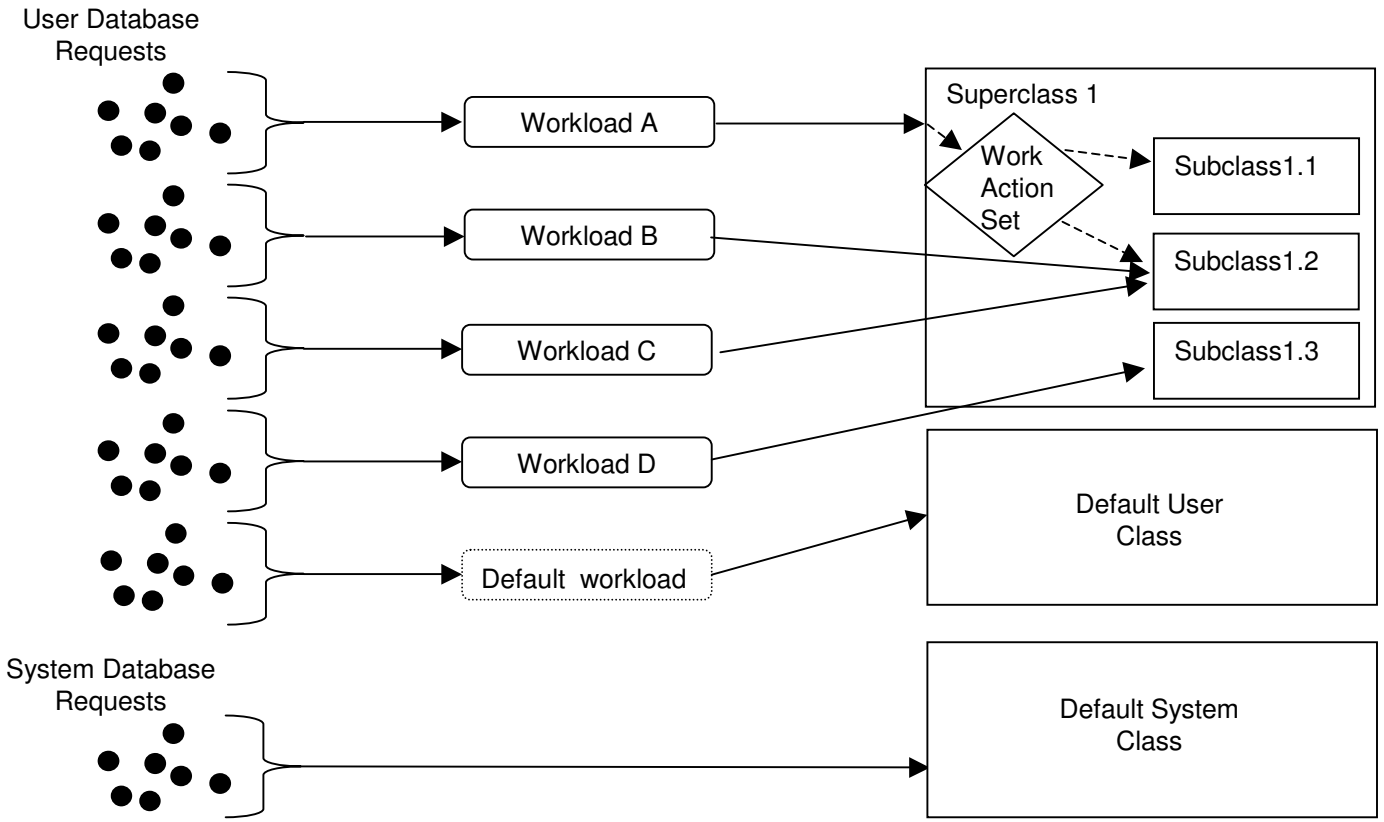

DB2 Viper 2 Execution Environment (Coexistence)

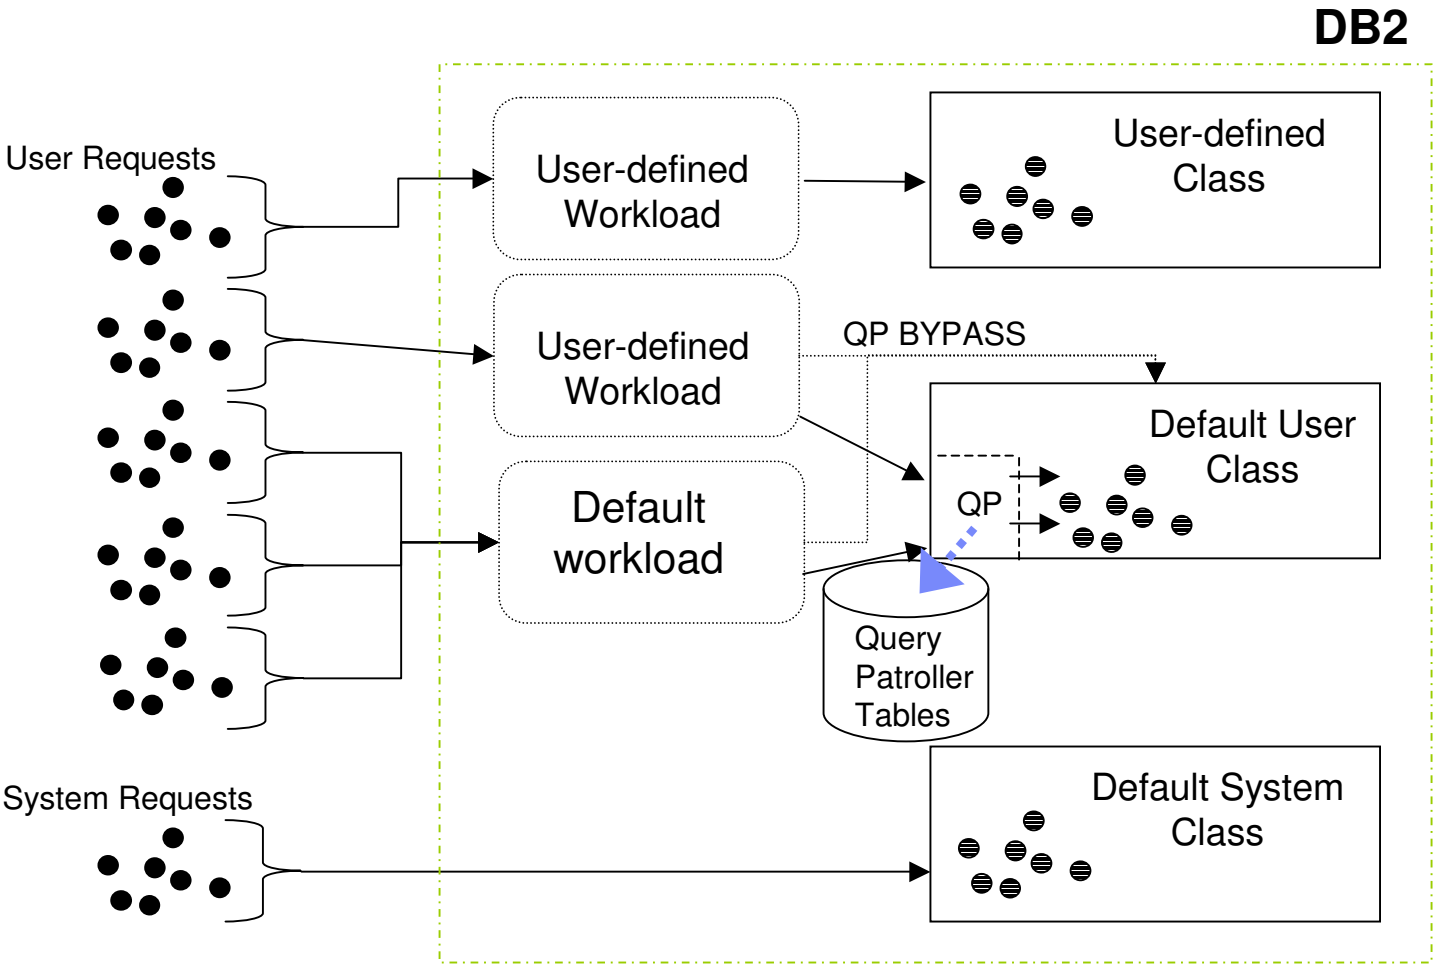

Monitoring Enhancements

- ***New table functions to allow quick ad hoc access to internal information via SQL from service classes and workload occurrences***
  - Can provide information from one or more database partitions with one invocation
  - Type of information: statistics and current activities and thresholds
- ***New event monitors***
  - allow capture of detailed information of SQL statements from database or service classes
  - allow capture of detailed information of SQL statements from activity limits
  - allow capture of service class statistical information at regular intervals
- ***New stored procedures***
  - To cancel a database activity
  - To capture detailed information on a database activity (i.e. SQL statement)

DB2 9.5

Performance Enhancements

Thread-based Process Model

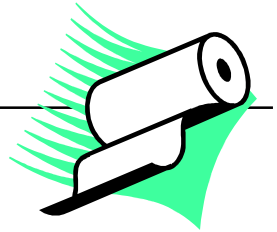

- ***Remove Process-Model Limitations***
- ***Move to a threaded engine architecture***
  - Provide basic execution parallelism independent of server type
  - Provide a single memory parameter for an entire DB2 node
    - control total memory allocated in DBMS, DB, and private memory
- ***Eliminate most agent-level configuration parameters and automate the remainder (i.e., maxagents, numpoolagents, numinitagents)***
  - Simplifies tuning
- ***Reduce agent memory consumption***

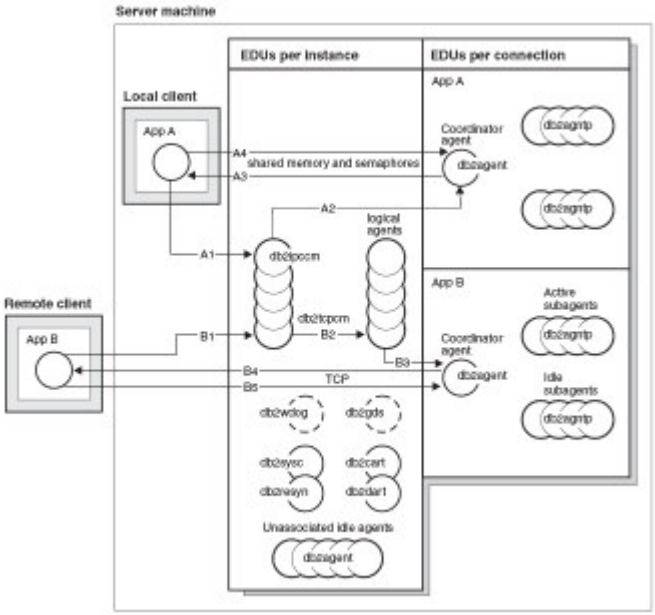

Thread-based Execution Motivation

- ***General Resource Savings***
  - significantly less system file descriptors used
    - All threads in a process can share the same file descriptors
    - No need to have each agent maintain it's own file descriptor table
- ***Performance Improvements***
  - Context switching between threads is generally faster than between processes
    - No need to switch address space
    - Less cache pollution
- ***Memory Savings***
  - Operating system threads require less context than processes
    - Share address space, context information (such as uid, file handle table, etc)
- ***Enables Key Usability / TCO Improvements***
  - Expected benefits with threading includes:
    - Automatic memory growing and shrinking support on Linux/HP/Sun
    - Single memory knob to control memory consumption of an entire instance
    - No need to worry about 'application groups'
    - More automatic and dynamic configuration parameters

DB2 9.5

Business Critical Reliability

Enhanced Audit Facility

- ***New audit records for the actual execution of dynamic and static SQL statements***
  - Includes SQL text with compilation environment (where appropriate) and, optionally, input data values
- ***Improved audit log file infrastructure to support configuration of location and maximum file size***
- ***Implement Fine grained audit control based on***
  - inclusion lists of user, group, or role authorization IDs
  - table object based audit indicators
  - trusted context audit indicator
- ***Improve audit performance***
- ***Increase family compatibility with DB2 for zOS***

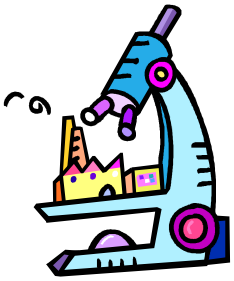

Why do we need more Security enhancements ?

We don‘t want to allow any application server to access our DB2 production system

We cannot identify the individual user, if they come through our Web Application Server

Current problems with 3-tier architectures

Common application server userid used for all communication with DB2

End-user not propagated to DB2

Not possible to audit actions performed by different end-users

Application server userid must hold all authorizations needed to perform all parts of the application

Application server must make sure that end-users are only allowed to use what they are authorized for

Privileges held by application server userid are valid from any “location”

If application server userid is compromised, the impact is large

Solution: Trusted context and roles

A trusted context is a database object that identifies a certain “location”

Connections created using a trusted context are trusted connections

Within a trusted connection you can
- Switch to another authid (with or without authentication)
  allows the reuse of connections !
- Exercise the use of roles

A role provides context dependent privileges
- The role is available only within the trusted context
- The privileges granted to the role can be exercised only through the trusted connection

Role

New database entity
- Not a userid
- Holds privileges like a userid (or group)

Can be associated with a userid within a trusted context via the “Default Role” clause

What is the intention?
- Limiting privileges to a particular context

Creating a Role

Creating a Role:

```
CREATE ROLE My_ROLE;
```

Grant privleges to a ROLE:

```
GRANT SYSADM TO My_ROLE;
```

Use in a trusted context, as default or by user:

```
CREATE TRUSTED CONTEXT My_TC
BASED UPON CONNECTION USING SYSTEM AUTHID aUserID
ATTRIBUTES (ADDRESS `1.2.3.4`)
DEFAULT ROLE My_Role
ENABLE;
```

A sample Trusted Context DDL

```
CREATE TRUSTED CONTEXT CTX1
BASED UPON CONNECTION USING SYSTEM AUTHID WASADM1
ATTRIBUTES (ADDRESS '1.2.3.4',
ADDRESS '1.2.3.5')
ENCRYPTION LOW
ENABLE;
```

Trusted Context and Role Definition are reflected in the DB2 Catalog

A sample Trusted Context DDL

```
CREATE TRUSTED CONTEXT CTX1
BASED UPON CONNECTION USING SYSTEM AUTHID WASADM1
ATTRIBUTES (ADDRESS '1.2.3.4',
ADDRESS '1.2.3.5')
ENCRYPTION LOW
ENABLE;
```

The possible Connection trust attributes are:

PROTOCOL: The communication protocol trust attribute. This is used to control which network communication protocols can use the trusted context.

ADDRESS: The network address trust attribute. This is used in conjunction with the PROTOCOL attribute to control which addresses the trusted context can be used with. This is the actual client's IP address or domain name, used by the connection to communicate with the database manager.

ENCRYPTION: The network encryption trust attribute. This specifies the minimum level of encryption of the data stream ("network encryption") for the connection.

AUTHENTICATION: The authentication trust attribute. The attribute specifies the level of authentication required to be performed on the system authorization ID during the establishment of the connection.

Connection Types

Explicit Trusted connection

You request that the connection be trusted AND the connection meets the server's criteria for being trusted

Implicit Trusted connection

You do not request that the connection be trusted AND the connection meets the server's criteria for being trusted

can be any kind of connection

-> a regular connection with add. Role privleges

Regular Connection

The connection does not meet the server's criteria for being trusted – if an explicit trusted connection was requested, it results in a regular connection and warning SQL20360W (SQLSTATE 01679) is returned

User Switching with explicit trusted connections

An explicit trusted connection is created via:

1.	**CLI/ODBC**	SQLConnect, SQLSetConnectAttr
2.	**XA CLI/ODBC**	XA_open
3.	**JAVA**	getDB2TrustedPooledConnection, getDB2TrustedXAConnection

Switching to a different user is done via:

1.	**CLI/ODBC**	SQLSetConnectAttr
2.	**XA CLI/ODBC**	SQLSetConnectAttr
3.	**JAVA**	getDB2Connection, reuseDB2Connection

Simple and Robust HA

- ***Integrated High Availability and Disaster Recovery solution (automated takeover)***
  - Support integrated install, setup, maintenance and uninstall of Tivoli System Automation (TSA) with DB2
  - DB2 will maintain the cluster configuration of TSA for both HADR and non-HADR failover scenarios

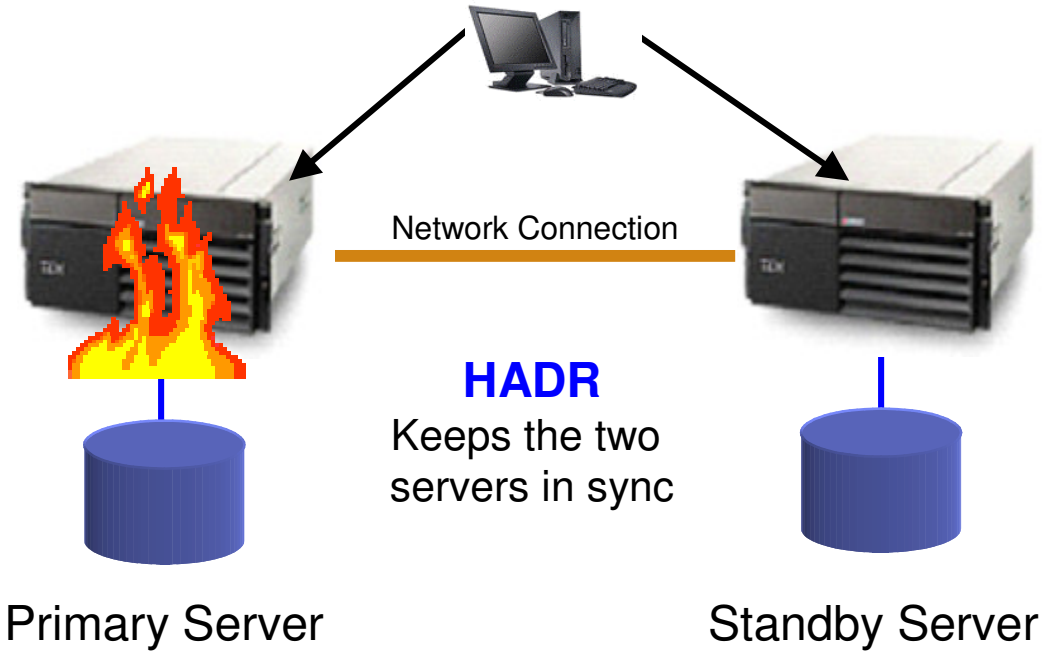

DB2 9.5 High Availability Enhancements

- *HA Cluster Manager Integration*
  - Coupling of DB2 and TSA on Linux and AIX, other platforms coming later
  - DB2 interface to configure cluster
  - DB2 to maintain cluster configuration, add node, add tablespace, …
  - Exploitation of new vendor independent layering (VIL), providing support for any cluster manager
- *NO SCRIPTING REQUIRED!*
  - one set of embedded scripts that are used by all cluster managers
- ***Automates HADR failover***
  - **Exploit HA cluster manager integration previously described**

DB2 – Cluster Manager Integration

- ***New utility db2haicu***
  - DB2 High Availability Instance Configuration Utility
  - Sets up DB2 with the Cluster Manager
    - For HADR or Shared Storage cluster
  - Interactive or XML file driven interface

- ***DB2 communicates with the Cluster Manager***
  - Keeps the cluster manager in sync with DB2 changes
    - CM needs to be aware of new tablespaces, containers etc
    - Avoids failover problems due to missing resources
  - Keeps the XML cluster configuration file up to date
  - Automates HADR Failover

Cluster Configuration XML file

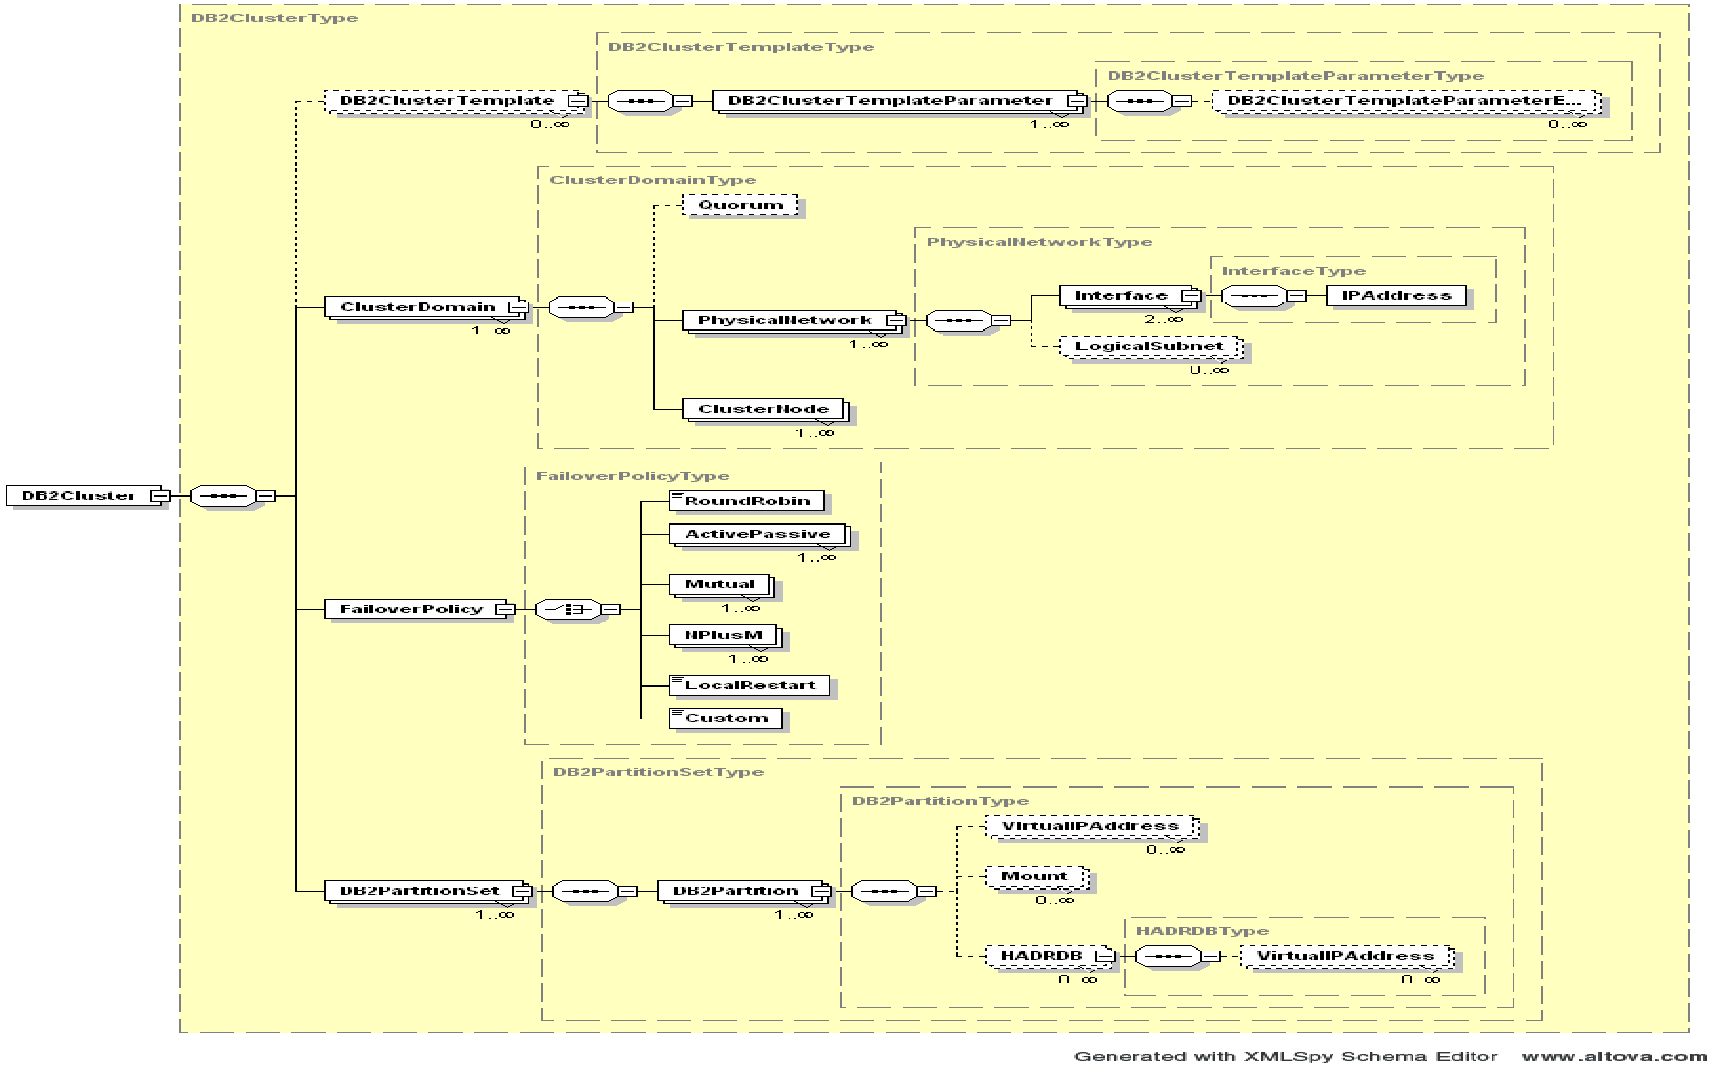

DB2 9.5 HADR “Peer Window”

- ***Preventing data loss during failover***
  - Some users prefer to hold transactions when the primary loses connection to the standby. For these users, preserving data (having a 2nd copy on standby) is more important than continuing transaction processing on the primary.
  - Example: financial customers.
  - New feature “peer window” in DB2 9.5
    - User configurable amount of time to hold transactions when the primary loses connection to the standby.
    - The time to “stay in peer” even if you lose connection is called “peer window”
    - Peer window size can be tuned from a few seconds to decades
    - Not applicable to async mode.

The bottom line – why IBM DB2 9.5 ?

- ***A leading-edge data server from an application-neutral vendor***
- ***The right architecture to support all your applications***
- ***The right performance, at the right cost***
- ***Easier risk management***
- ***Agility to adapt faster to change***

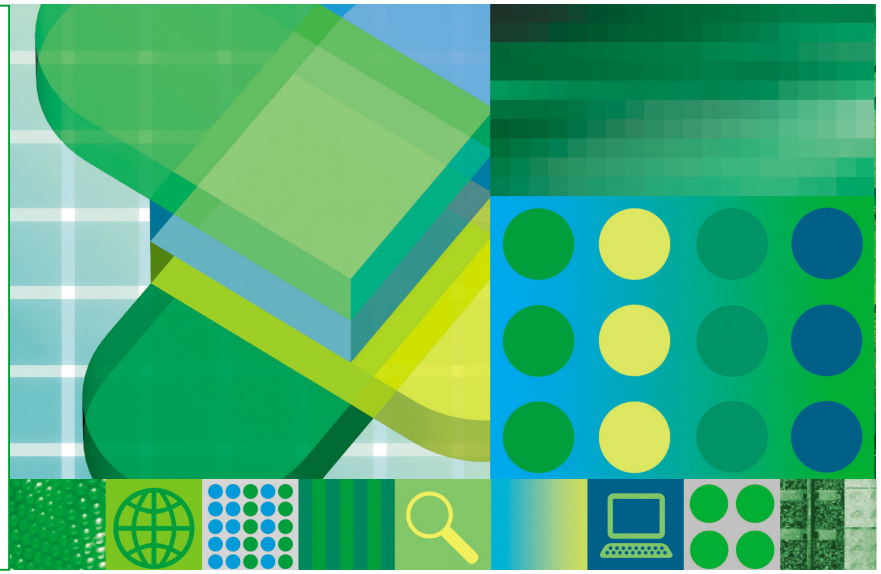

Questions ?