

**IBM Information**

>>> On Demand

**2007**



# IBM Data Servers Application Development and Administration Trends and Directions

*Curt Cotner, IBM Fellow, [cotner@us.ibm.com](mailto:cotner@us.ibm.com)*

*Session 1298A*



***Act.Right.Now.***

**IBM INFORMATION ON DEMAND 2007**

**October 14 - 19, 2007**

**Mandalay Bay**

**Las Vegas, Nevada**

# IBM Data Servers

## Reduce cost of deployment and management of data

- *Innovation to reduce the cost of infrastructure*
- *Innovation to manage the lifecycle of data - from modeling and design through change management and sunsetting*

## Enable rapid use of data throughout the enterprise

- *Innovation that accelerates SOA and XML initiatives*
- *Innovation that leverages Web 2.0 and situational applications*



---

# *IBM Data Studio*

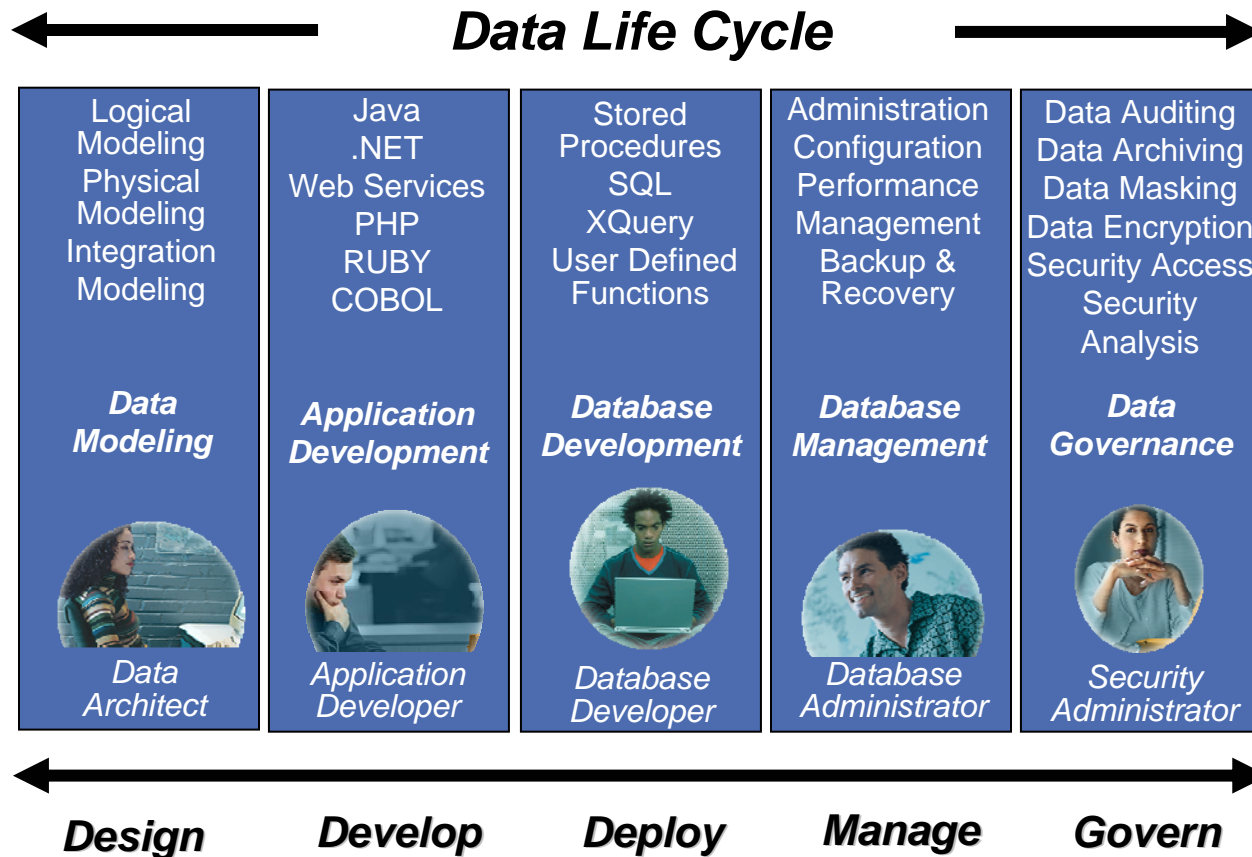
## *What is it ?*

*IBM Data Studio is a comprehensive data management solution that empowers you to effectively design, develop, deploy and manage your data, databases and database applications throughout the entire application development life cycle utilizing a consistent and integrated user interface*



# IBM Data Studio

## Who will use it ?



# IBM Data Studio

A Consistent, Integrated Solution



Application Developer



Database Developer

- Develop
- Coding
  - Debugging
  - Teaming
  - Testing
  - Tuning

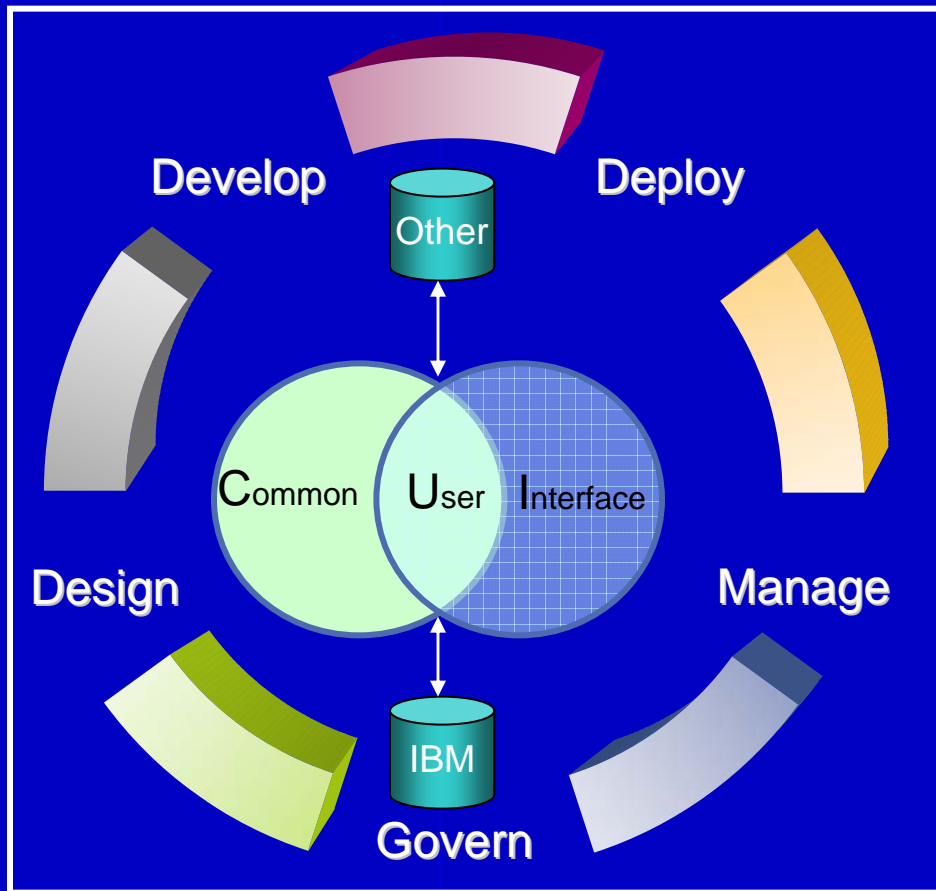
- Design
- Logical Modeling
  - Physical Modeling
  - Integration Modeling



Business Analyst



Database Architect



Database Administrator

- Manage
- Database Administration
  - Data Management
  - Change Management
  - Recovery Management
  - Storage Management
  - Performance Management

- Govern
- Security Access
  - Security Analysis
  - Data Auditing
  - Data Archiving
  - Data Masking
  - Data Encryption



Security Administrator



# IBM Data Studio v1.1

- Empowering developers and database administrators
- Complimentary and available in October of 2007
- Support for DB2 on all platforms and IDS

DB2 for LUW	DB2 for z/OS	DB2 for i5/OS	IDS
<ul style="list-style-type: none"> <li>▪ Physical Data Modeling</li> <li>▪ Data Distribution Viewer</li> <li>▪ Integrated Query Editor</li> <li>▪ SQL Builder</li> <li>▪ SQL Routine Debugger</li> <li>▪ Java Routine Debugger</li> <li>▪ XML Editor</li> <li>▪ XML Schema Editor</li> <li>▪ pureQuery for Java</li> <li>▪ Data Web Services</li> <li>▪ Object Management</li> <li>▪ Data Management</li> <li>▪ Update Statistics</li> <li>▪ Health Monitoring *</li> <li>▪ Visual Explain</li> <li>▪ Security Access Controls</li> <li>▪ Project Management</li> </ul> <p>* Technical Preview</p>	<ul style="list-style-type: none"> <li>▪ Physical Data Modeling</li> <li>▪ Data Distribution Viewer</li> <li>▪ Integrated Query Editor</li> <li>▪ SQL Builder</li> <li>▪ SQL Routine Debugger</li> <li>▪ Java Routine Debugger</li> <li>▪ XML Editor</li> <li>▪ XML Schema Editor</li> <li>▪ pureQuery for Java</li> <li>▪ Data Web Services</li> <li>▪ Object Management</li> <li>▪ Data Management</li> <li>▪ Update Statistics</li> <li>▪ Visual Explain</li> <li>▪ Security Access Controls</li> <li>▪ Project Management</li> </ul>	<ul style="list-style-type: none"> <li>▪ Physical Data Modeling</li> <li>▪ Data Distribution Viewer</li> <li>▪ Integrated Query Editor</li> <li>▪ SQL Builder</li> <li>▪ SQL Routine Debugger</li> <li>▪ Java Routine Debugger</li> <li>▪ XML Editor</li> <li>▪ XML Schema Editor</li> <li>▪ pureQuery for Java</li> <li>▪ Data Web Services</li> <li>▪ Object Management</li> <li>▪ Data Management</li> <li>▪ Security Access Controls</li> <li>▪ Project Management</li> </ul>	<ul style="list-style-type: none"> <li>▪ Physical Data Modeling</li> <li>▪ Data Distribution Viewer</li> <li>▪ Integrated Query Editor</li> <li>▪ SQL Builder</li> <li>▪ XML Editor</li> <li>▪ XML Schema Editor</li> <li>▪ pureQuery for Java</li> <li>▪ Data Web Services</li> <li>▪ Object Management</li> <li>▪ Data Management</li> <li>▪ Security Access Controls</li> <li>▪ Project Management</li> </ul>





# Why IBM Data Studio ?

## **Increase productivity for all roles throughout the data life cycle**

- Slash development time up to 50% with an integrated data management environment
- Promote collaboration across roles to optimize data server and application performance
- Accelerate Java development productivity with new pureQuery data access
- Simplify development of applications implementing industry specific XML standards
- Monitor data server operation & performance anywhere, anytime from a Web browser

## **Simplify and speed development of new skills**

- Learn once, use with all supported data servers
- Easy-to-use and integrated user interface, compatible with Rational Software Development Platform
- Extensible with Eclipse plug-ins to customize the environment for each team member

## **Accelerate data as a service for Service Oriented Architecture**

- Develop and publish data as a Web service without programming
- Info 2.0 Ready - support for Web 2.0 protocols and format





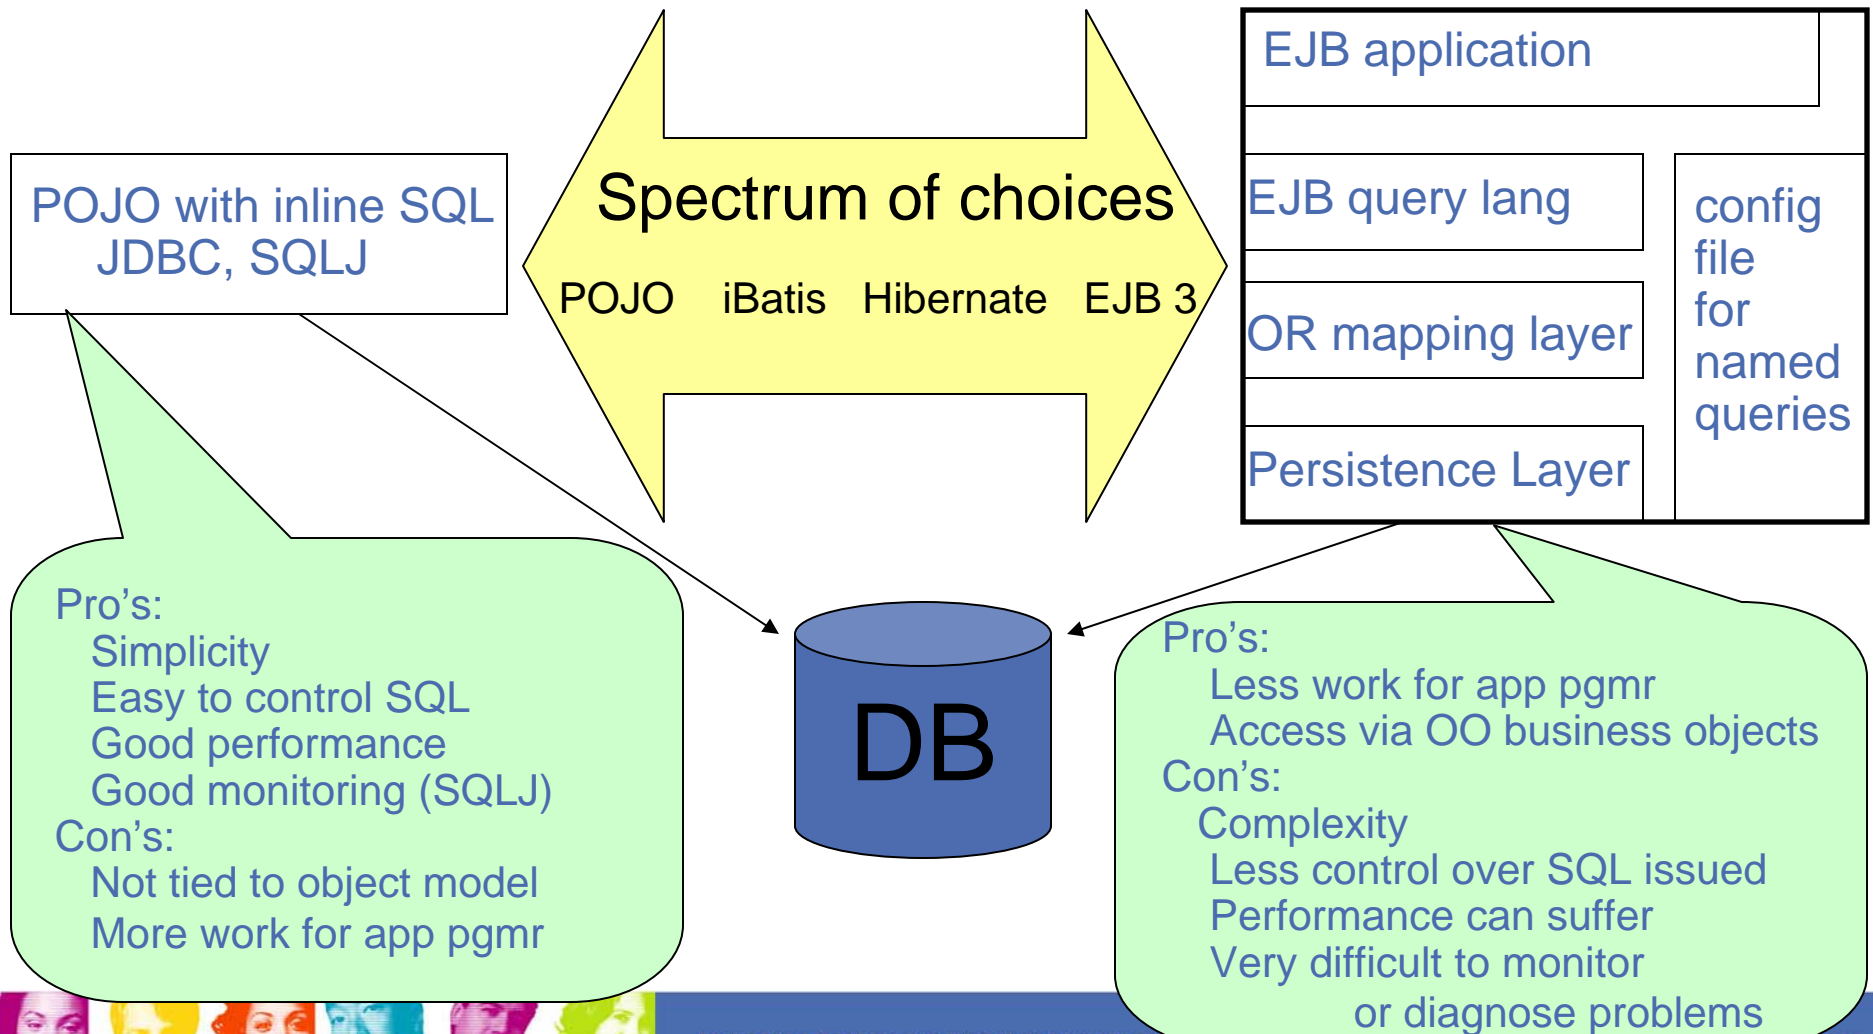
# IBM Data Studio -- Java Database Application Development, Monitoring, Management, Problem Determination

*Act.Right.Now.*

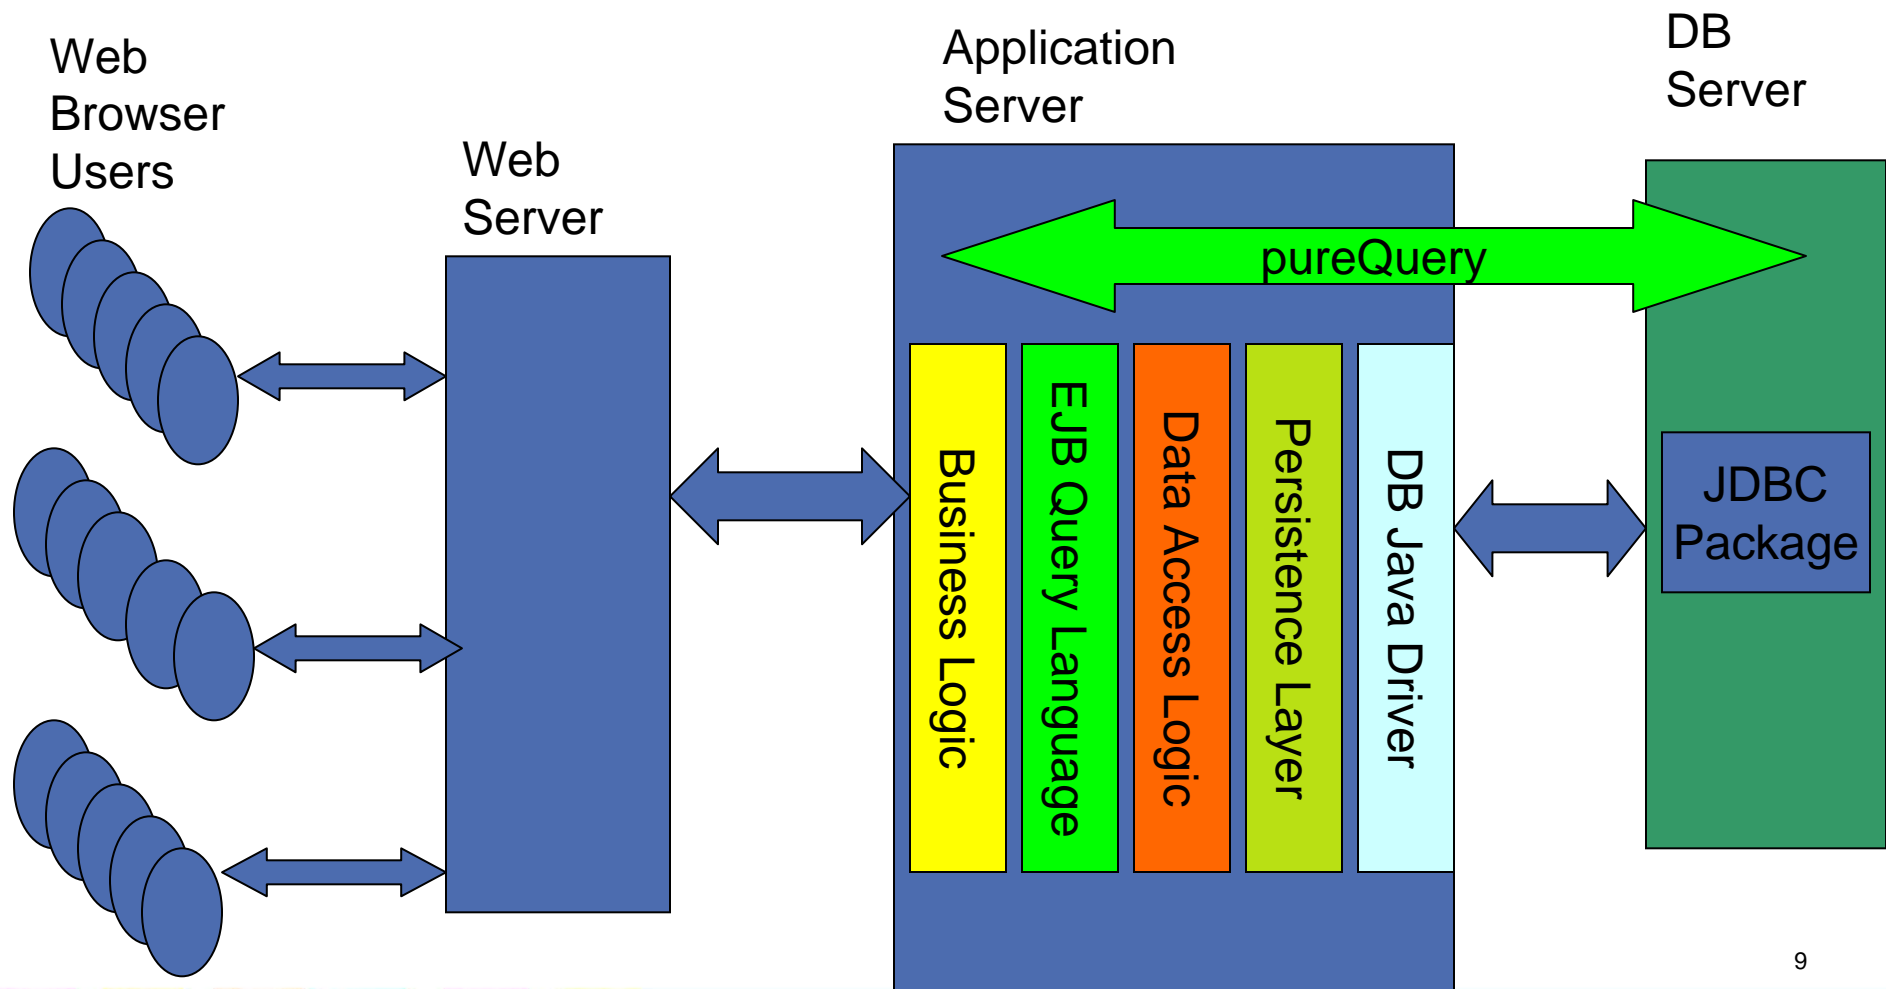




# Java Data Access – many forms



# Toughest issue for Web applications – Problem diagnosis and resolution



# What performance/diagnosis challenges?

## EJB Query Language:

```
SELECT object(e) FROM Employee e
WHERE e.dept=?1 AND e.salary>=?2
```


Query language is a subset of SQL. Doesn't have all the SQL features you want.

App query syntax is different from SQL query. How do you track problem SQL queries back to the app that issued the original query???

In most cases, queries map to JDBC. No ability to lock in access path at program deployment. No ability to search catalog to see which queries are issued by a given program.

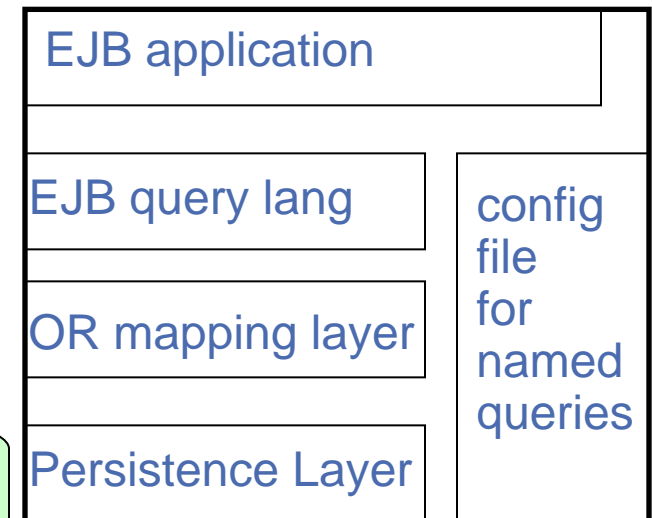
Often, app query is intercepted by persistence layer, and the resulting SQL query looks nothing like the app query.

- Resulting query might perform badly.
- Changing app query might not result in a similar change in the SQL query...



## SQL issued to database:

```
SELECT * FROM PROD.EMP
WHERE DEPT=? AND SALARY>?
```



# pureQuery – Beyond Function

- Development of applications
  - Tools to assist SQL development in .java source file
  - Simple SQL APIs, easy to write to and extend
  - Multiple API “styles” to align with popular Java frameworks
- Query important data sources simply
  - Database, Cache, Collections, XML
- Problem Determination
  - When problems occur, find source quickly.
- Governance / Management
  - Track SQL back to individual apps, lock in access paths with static SQL packages, align with customer change control processes
- Provide high performance/scalability
  - Application: short path length, coding over metadata, optional code gen, JDBC and static SQL runtime optimizations
  - Database: static SQL, batching, pass app SQL directly to database

11



# pureQuery API “Styles”

- Support several API styles to fit well into all of the popular Java programming models/frameworks
  - Inline style (familiar JDBC and SQLJ approach)
  - Method style (similar to JDBC 4 ease of use enhancements)
  - Named query style (similar to iBatis/JDO/Hibernate/JPA)



# Retrieve a single row from Database

**pureQuery:**

*Automatically Optimizes for 1 row*

```
addr = db.queryFirst("SELECT ADDRESS FROM EMP  
WHERE NAME=:name", String.class, name);
```

*-or-*

```
addr = getAddress(name);
```

XML file or Java annotation  
SELECT ADDRESS FROM EMP  
WHERE NAME=?1;

**SQLJ:**

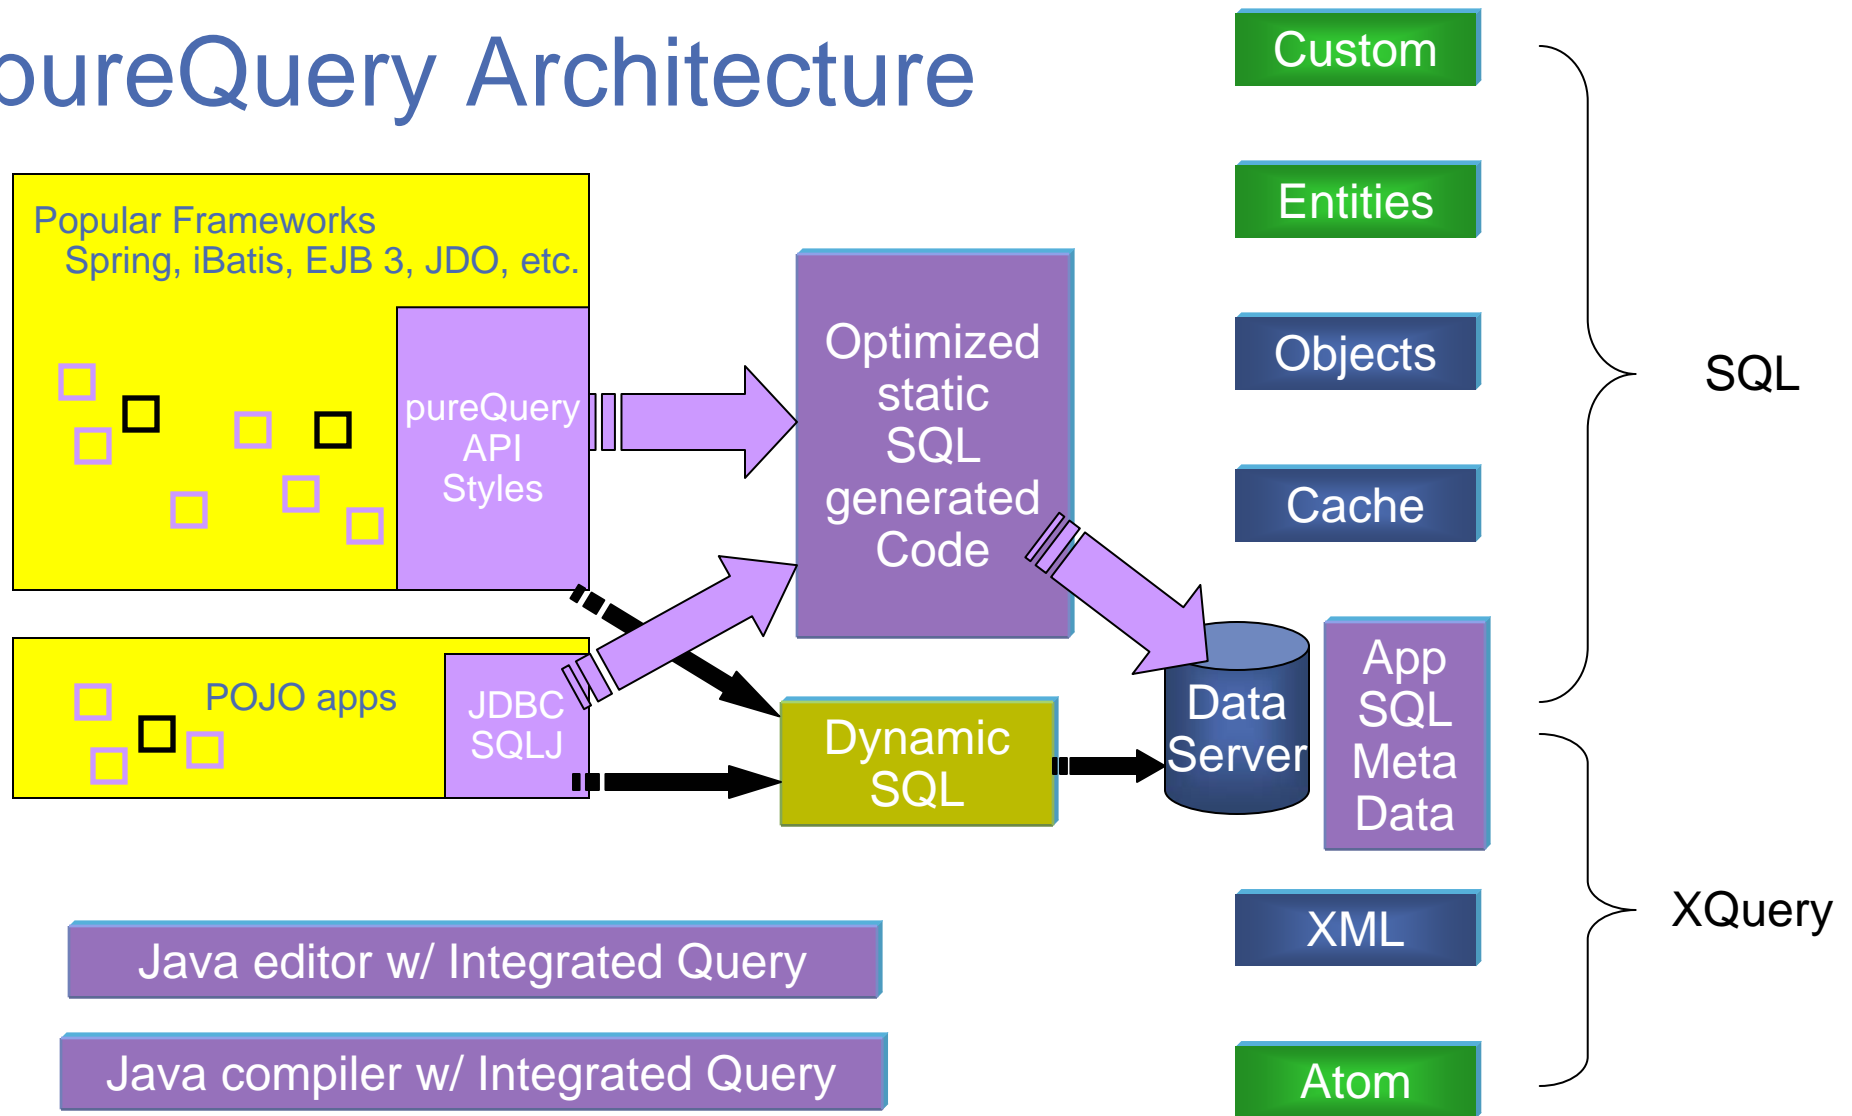
```
#sql [con] { SELECT ADDRESS INTO :addr FROM EMP  
WHERE NAME=:name };
```

**JDBC:**

```
java.sql.PreparedStatement ps = con.prepareStatement(  
    "SELECT ADDRESS FROM EMP WHERE NAME=?");  
ps.setString(1, name);  
java.sql.ResultSet names = ps.executeQuery();  
names.next();  
addr = names.getString(1);  
names.close();
```



# pureQuery Architecture

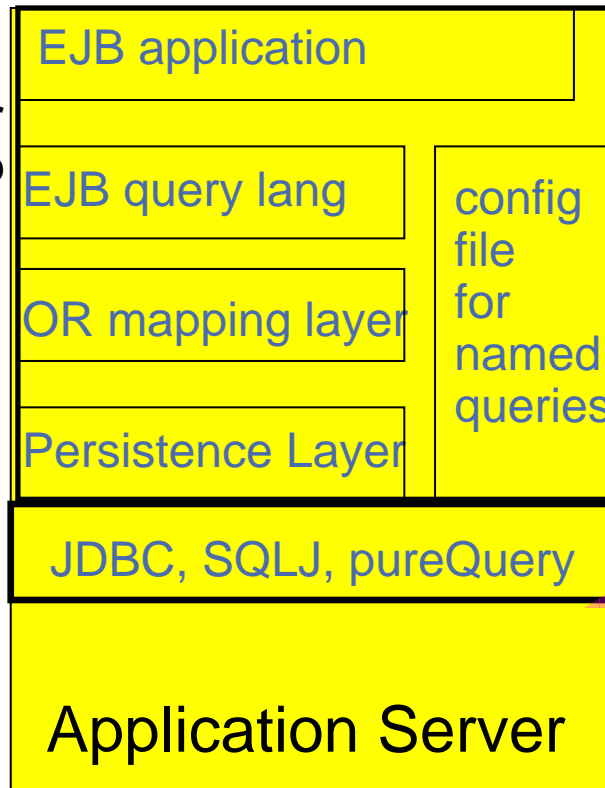


# Problem Determination and Monitoring

Original query issued by app  
 Java class and line number  
 XML filename and line number  
 Last compile date/time for app  
 Full history of all SQL issued  
 by each application

Generated static SQL  
 -- SQL after O/R mapping and persistence  
 -- plan lockdown  
 -- package versioning

Set Client Information APIs  
 -- end user's ID  
 -- end user's IP address  
 -- application name  
 -- accounting chargeback data



JDBC Static SQL Profiling  
 -- SQL after O/R mapping and persistence  
 -- Java stack trace for SQL call  
 -- plan lockdown  
 -- package versioning

Set Client Information APIs  
 -- end user's ID  
 -- end user's IP address  
 -- application name  
 -- accounting chargeback data

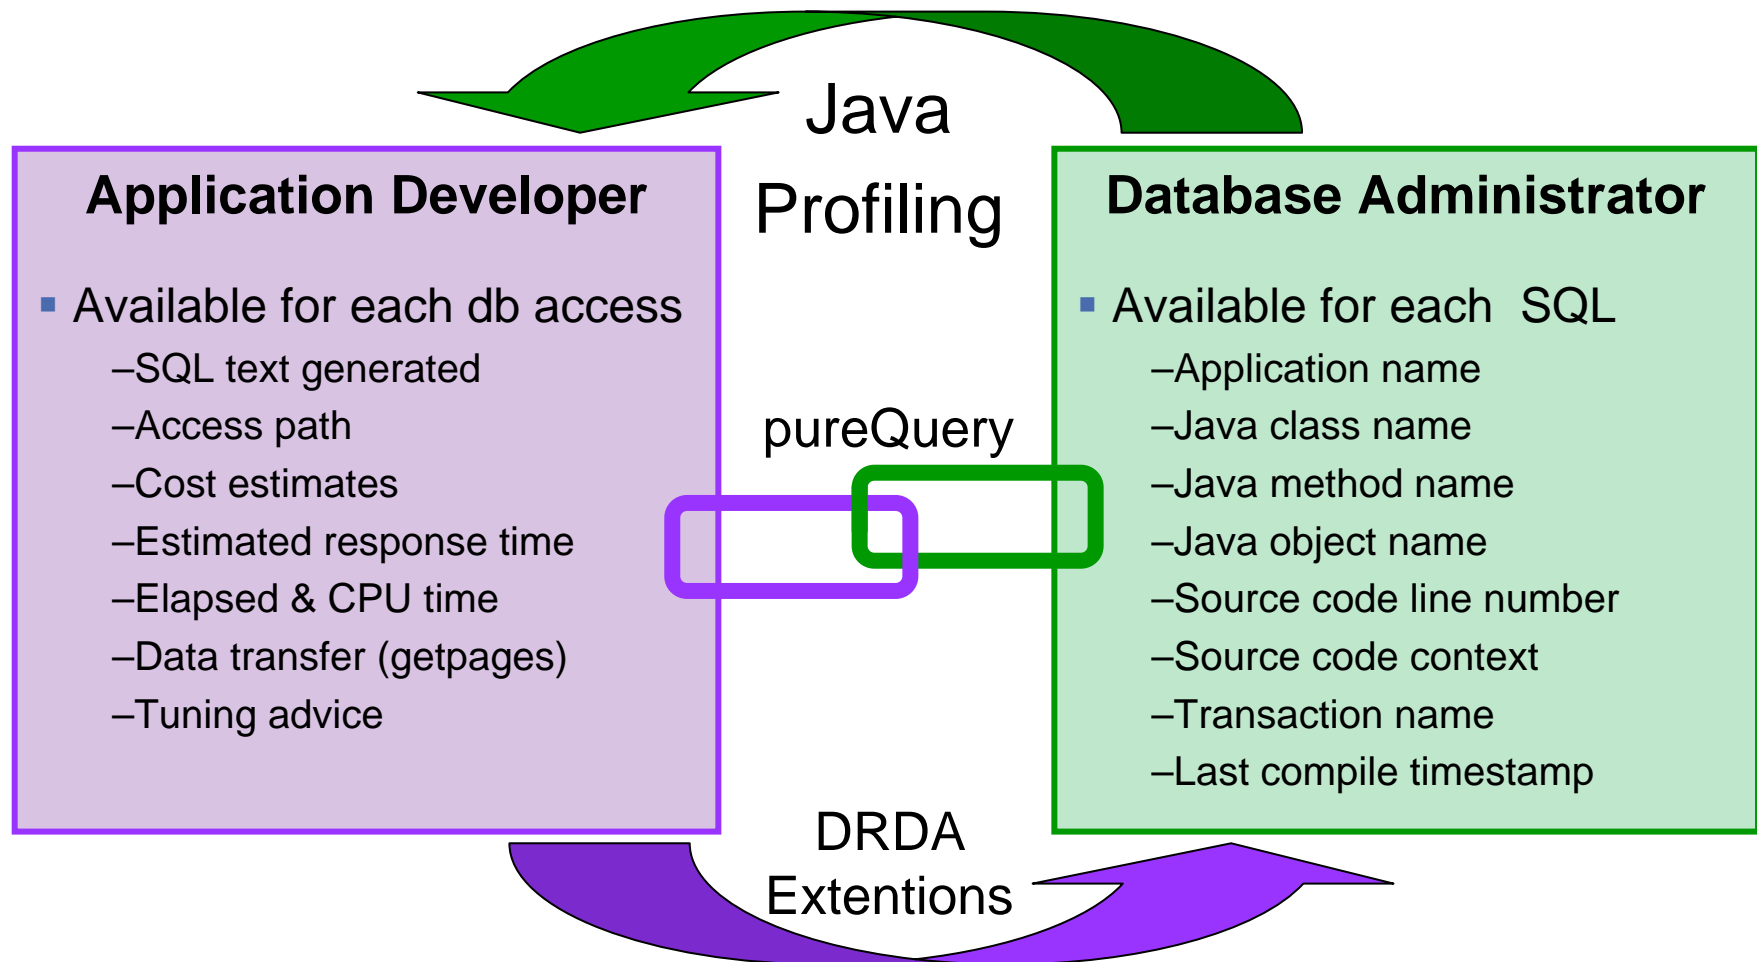
App server IP address  
 App server connection pool userid  
 JDBC driver package name

pureQuery

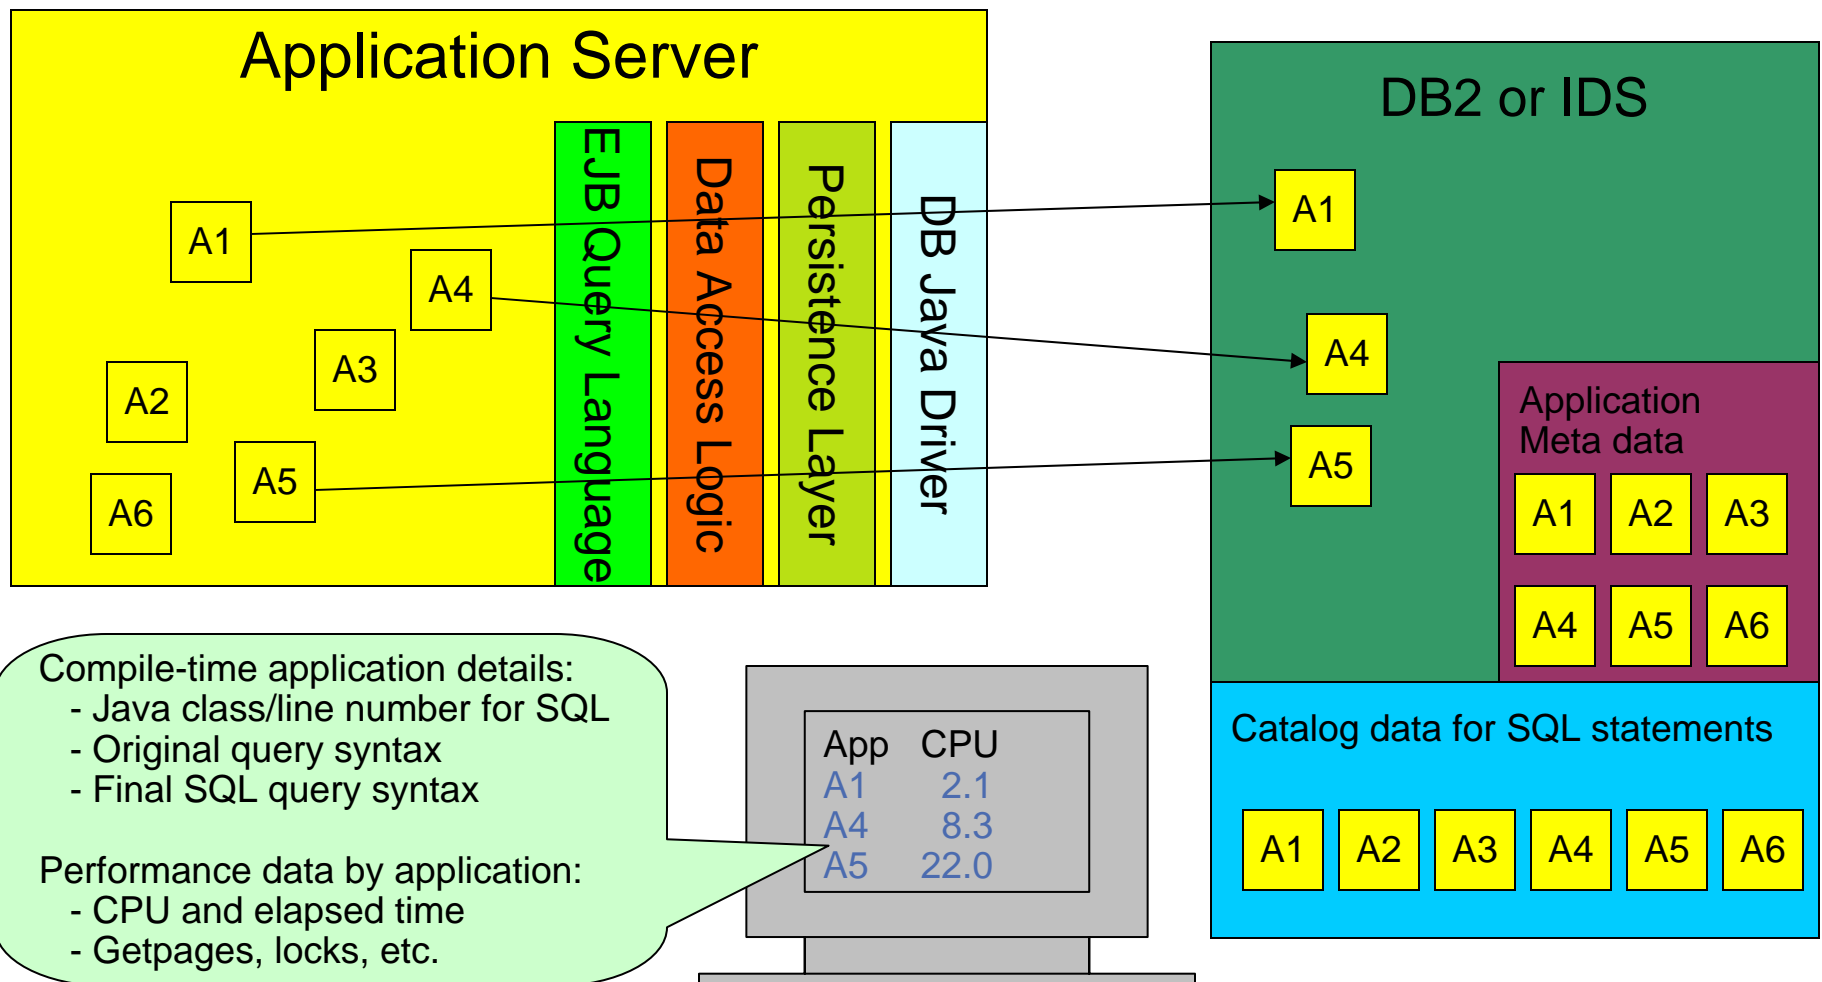
JDBC



# Simplifying Problem Determination Scenario



# pureQuery with IBM Runtime/Tooling





# IBM Data Studio in 2007

*Act.Right.Now.*



# IBM Data Studio – Workbench

A single productive work environment

The screenshot displays the IBM Data Studio interface. The main window is titled "Data - EMPLOYEE - IBM Viper Studio". The interface includes a menu bar (File, Edit, Navigate, Search, Project, Data, Run, Window, Help) and a toolbar. On the left, a "Database Explorer" pane shows a tree view of database objects, including connections like "SAMPLE [DB2 for Linux, UNIX, and Windows v9.5]" and "STLECI1". The central "Data Object Editor" pane is focused on the "EMPLOYEE" table, showing its properties: Name (EMPLOYEE), Label, Schema (RBUGLIO), and Data capture (NONE). Below this, there are sections for "Impacted Objects" and "DDL". At the bottom, a "Sample Contents" table displays data for the EMPLOYEE table.

EMPNO	FIRSTNME	MIDINIT	LASTNAME	WORKDI
000010	CHRISTINE	I	HAAS	A00
000020	MICHAEL	L	THOMPSON	B01
000030	SALLY	A	KWAN	C01
000050	JOHN	B	GEYER	E01
000060	IRVING	F	STERN	D11
000070	EVA	D	PULASKI	D21
000090	EILEEN	W	HENDERSON	E11
000100	THEODORE	Q	SPENSER	E21
000110	VINCENZO	G	LUCCHESI	A00
000120	SEAN		O'CONNELL	A00





# IBM Data Studio – Workbench

An integrated query editor for SQL and XQuery

The screenshot displays the IBM Data Studio Workbench interface. At the top, the title bar reads "Data - statement.sql - IBM Viper Studio". Below the title bar is a menu bar (File, Edit, Navigate, Search, Project, Data, Run, SQL, Window, Help) and a toolbar. The main area is divided into three sections:

- SQL Editor:** Contains a SQL query: 

```
SELECT "RBUGLIO ".EMPLOYEE.EMPNO, "RBUGLIO ".EMPLOYEE.FIRSTNAME,
"RBUGLIO ".EMPLOYEE.MIDINIT, "RBUGLIO ".EMPLOYEE.LASTNAME, "RBUGLIO ".EMPLOYEE.WORKDEPT,
"RBUGLIO ".EMPLOYEE.PHONENO, "RBUGLIO ".EMPLOYEE.HIREDATE, "RBUGLIO ".EMPLOYEE.JOB,
"RBUGLIO ".EMPLOYEE.EDLEVEL, "RBUGLIO ".EMPLOYEE.SEX, "RBUGLIO ".EMPLOYEE.BIRTHDATE,
"RBUGLIO ".EMPLOYEE.SALARY, "RBUGLIO ".EMPLOYEE.BONUS, "RBUGLIO ".EMPLOYEE.COMM,
RBUGLIO.EMP_RESUME.RESUME_FORMAT, RBUGLIO.EMP_RESUME.RESUME
FROM
"RBUGLIO ".EMPLOYEE JOIN RBUGLIO.EMP_RESUME ON "RBUGLIO ".EMPLOYEE.EMPNO = RBUGLIO.EMP_RESUME.EMPNO
WHERE "RBUGLIO ".EMPLOYEE.LASTNAME LIKE "B%"
```
- Schema Diagram:** Shows two tables: "EMPLOYEE" and "EMP\_RESUME". The "EMPLOYEE" table has columns EMPNO, FIRSTNAME, MIDINIT, LASTNAME, and WORKDEPT. The "EMP\_RESUME" table has columns EMPNO, RESUME\_FORMAT, and RESUME. A line with a diamond at the end indicates a one-to-one relationship between the EMPNO columns of the two tables.
- Column List:** A table showing the columns selected in the query. The "Statement" field is set to "statement.sql" and the "DISTINCT" checkbox is unchecked.

Column	Alias	Output	Sort Type	Sort Order
"RBUGLIO ".EMPLOYEE.EMPNO		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.FIRSTNAME		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.MIDINIT		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.LASTNAME		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.WORKDEPT		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.PHONENO		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.HIREDATE		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.JOB		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.EDLEVEL		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.SEX		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.BIRTHDATE		<input checked="" type="checkbox"/>		
"RBUGLIO ".EMPLOYEE.SALARY		<input checked="" type="checkbox"/>		

The bottom status bar shows the current database connection: "SAMPLE (SAMPLE: jdbc:db2://localhost:50000/SAMPLE)".



# IBM Data Studio – Workbench

## Fully integrated with the Administration Console

The screenshot displays the IBM Data Studio Workbench interface, which is fully integrated with the Administration Console. The main window is titled "Data - Integrated Solutions Console - IBM Viper Studio" and shows a "Viper Studio Server Console" with a "Health Monitor" tab selected. The Health Monitor displays a "Heat Chart" and an "Alert History" table. The Alert History table shows the following data:

Name	Monitor Status		Alert		System				
	Monitor Status	Data Server Status	Warning	Critical	CPU Usage	Database Storage	Memory Usage	Locking	Recovery
SAMPLE	Green	Green	0	0	Green	Green	Green	Green	Green
TOOLSDB	Green	Green	0	0	Green	Green	Green	Green	Green

The interface also includes a "Database Explorer" on the left, a "Database Project Explorer" at the bottom, and a "Messages" section at the bottom right. The overall layout is clean and professional, with a blue and white color scheme.



# IBM Data Studio - 2007

## The Administration Console



- What is it ?
  - A web based interface to perform operational database management tasks
  - Immediate access to critical data server information and functions from anywhere, anytime
- What's available in the first release
  - Health and availability monitoring
    - DB2 for LUW V9.5
      - Problem determination
      - Problem recommendations
      - 72 Hours of History
      - Ability to monitor up to 100 data servers
  - Q Replication monitoring and administration
- What will be available in the future?
  - Support for all IBM data servers – DB2 for z/OS, IDS and DB2 for i5/OS
  - The future home for all operational database management tasks
    - Performance Monitoring
    - Performance Analysis
    - Configuration Management
    - Utility Management
    - Storage Management
    - Security Management
    - Advanced Deployment Capabilities



# Quick & Easy Problem Determination

## Heatchart – Overall Health Status



Where are the most important hotspots that need my attention?

## Dashboard – Adhoc Investigation



Something doesn't seem quite right. I wonder what's happening?



Database Administrator

## Alert List – Historical Investigation

Severity	Alert Type	Timestamp	Database
High	CPU Utilization of LPAB/System	05:31am	Support(O/S)
High	CPU Utilization of LPAB/System	04:37am	Accounts(O/S)
High	CPU Utilization of LPAB/System	08:07am	Accounts(O/S)
Medium	Application timeout	07:12am	Support(O/S)
Medium	Application timeout	05:44am	Accounts(O/S)
Medium	Application timeout	04:37am	Marketing(O/S)
Medium	Application timeout	11:18am	Marketing(O/S)
Medium	Application timeout	11:07am	Marketing(O/S)
Medium	Application timeout	2007/06/04	Accounts(O/S)
Medium	CPU Utilization of LPAB/System	2007/06/04	Accounts(O/S)
Medium	CPU Utilization of LPAB/System	2007/06/05	Accounts(O/S)

What happened when I was out for lunch? ... Away for weekend?

## Recommendations – Root Cause Analysis

**Table space TS1 in the TA database is off-line**

At the time of the alert, the TS1 tablespace is off-line, and, as a result, is inaccessible.

**Symptoms**  
The TS1 tablespace is inaccessible.

**Causes**  
A table space is in this state if there is a problem preventing access to one or more of its containers. This is often caused by media problems that are either permanent (for instance a bad disk) or temporary (for instance an offline disk or unmounted file system). After the problem has been remediated and the containers are all accessible, the table space can be brought backonline.

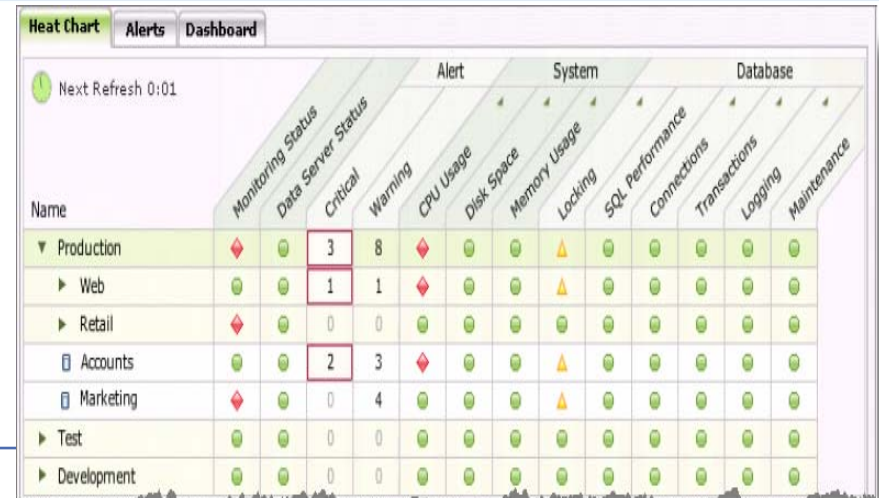
**Diagnosing the problem**

- Table space container is missing**  
If one or more containers of a table space cannot be found by the database management system, the table space will be taken offline and put in an inaccessible state.
- Table space container is damaged**  
If one or more containers of a table space are found to be damaged by the data manager, the table space will be taken offline and put in an inaccessible state.

Guide me to the root cause and help me fix it properly; I need to know all the relevant info to make best decision.



# The Heatchart Overall Health Status



Locking	<ul style="list-style-type: none"> <li>Application causing lock escalation</li> <li>Table space inaccessible due to quiesced state</li> <li>Application waiting indefinitely for lock held by indoubt transaction</li> </ul>
Database Storage	<ul style="list-style-type: none"> <li>Table space inaccessible because it is offline</li> <li>Running out of database storage</li> <li>Storage I/O problem</li> </ul>
Recovery	<ul style="list-style-type: none"> <li>Database inaccessible because recovery is incomplete or failed</li> <li>Table space cannot be updated b/c it is in backup pending state</li> <li>Table space inaccessible b/c it is in drop pending state</li> </ul>
System	<ul style="list-style-type: none"> <li>Bad response times due to high CPU usage</li> <li>Bad response times due to high memory usage</li> </ul>
Logging	<ul style="list-style-type: none"> <li>Long running app is using too much log space</li> <li>Transaction log access problems</li> <li>Full log held by indoubt transaction</li> <li>Reduced application performance due to logging problems</li> </ul>
Monitor Status	<ul style="list-style-type: none"> <li>Monitor agent offline</li> <li>Monitoring turned off by user</li> </ul>
Data Server Status	<ul style="list-style-type: none"> <li>Data server not responding</li> </ul>





# The Dashboard Adhoc Investigation



Locking	<ul style="list-style-type: none"> <li>Number of locks held</li> <li>Deadlock rate</li> <li>Lock escalation rate</li> <li>Number of lock waits</li> </ul>
Transaction Activity	<ul style="list-style-type: none"> <li>Transaction rate</li> <li>Failed transaction rate</li> <li>Number of indoubt transactions</li> </ul>
I/O Activity	<ul style="list-style-type: none"> <li>Data Volume</li> <li>I/O Volume</li> </ul>
Connections	<ul style="list-style-type: none"> <li>Number of connections</li> <li>High-water-mark of connection</li> </ul>
System	<ul style="list-style-type: none"> <li>CPU load %</li> <li>Used/Free real memory</li> <li>Used/Free virtual memory</li> <li>Used/Free swap memory</li> </ul>
Logging	<ul style="list-style-type: none"> <li>Logging volume</li> </ul>
Recovery	<ul style="list-style-type: none"> <li>Days since last backup</li> </ul>





# The Alert List Historical Analysis

Severity	Alert Type	Timestamp	Database
■	▶ CPU Utilization of LPAR/System	05:31pm	Support(z/OS)
■	▶ CPU Utilization of LPAR/System	04:37pm	Account(z/OS)
■	▶ CPU Utilization of LPAR/System	08:07am	Account(z/OS)
▲	▶ Application timeout	07:12pm	Support(z/OS)
▲	▶ Application timeout	05:44pm	Account(z/OS)
▲	▶ Application timeout	04:37pm	Marketing(z/OS)
▲	▶ Application timeout	11:18am	Marketing(z/OS)
▲	▶ Application timeout	11:07am	Marketing(z/OS)
▲	▶ Application timeout	2007/05/06	Marketing(z/OS)
▲	▶ CPU Utilization of LPAR/System	2007/05/06	Account(z/OS)
▲	▶ CPU Utilization of LPAR/System	2007/05/05	Account(z/OS)

Locking	<ul style="list-style-type: none"> <li>Application causing lock escalation</li> <li>Table space inaccessible b/c in quiesced state</li> <li>Application waiting indefinitely for lock held by indoubt transaction</li> </ul>
Database Storage	<ul style="list-style-type: none"> <li>Table space inaccessible b/c it is offline</li> <li>Running out of database storage</li> <li>Storage I/O problem</li> </ul>
Recovery	<ul style="list-style-type: none"> <li>Database inaccessible b/c recovery is incomplete/failed</li> <li>Table space cannot be updated b/c it is in backup pending state</li> <li>Table space inaccessible b/c it is in drop pending state</li> </ul>
System	<ul style="list-style-type: none"> <li>Bad response times due to high CPU usage</li> <li>Bad response times due to high memory usage</li> </ul>
Logging	<ul style="list-style-type: none"> <li>Long running app is using too much log space</li> <li>Transaction log access problems</li> <li>Full log held by indoubt transaction</li> <li>Reduced application performance due to logging problems</li> </ul>
Monitor Status	<ul style="list-style-type: none"> <li>Monitor agent offline</li> <li>Monitoring turned off by user</li> </ul>
Data Server Status	<ul style="list-style-type: none"> <li>Data server not responding</li> </ul>



# Recommendations

## Detailed Root Cause Analysis

TA ▶ Alert Detail : Table space TA.TS1 is offline @Jun 25, 2007 12:48:06 PM

Alert Recommendation Community

Table space **TS1** in the **TA** database is off-line

- Table space container is missing
  - A table space container file for **TS1** was renamed or moved
  - The physical disk is offline
  - The file system is unmounted
  - The physical disk is corrupted
  - A table space container file for **TS1** was deleted
- Table space container is damaged
  - The disk sector is damaged
  - The container file is tampered

### Table space **TS1** in the **TA** database is off-line

At the time of the alert, the **TS1** tablespace is off-line, and, as a result, is inaccessible.

#### Symptoms

The **TS1** tablespace is inaccessible.

#### Causes

A table space is in this state if there is a problem preventing access to one or more of its containers. This is often caused by media problems that are either permanent (for instance a bad disk) or temporary (for instance, an offline disk or unmounted file system). After the problem has been remedied and the containers are all accessible, the table space can be brought back online.

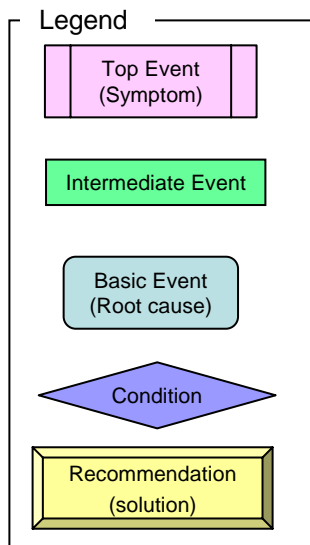
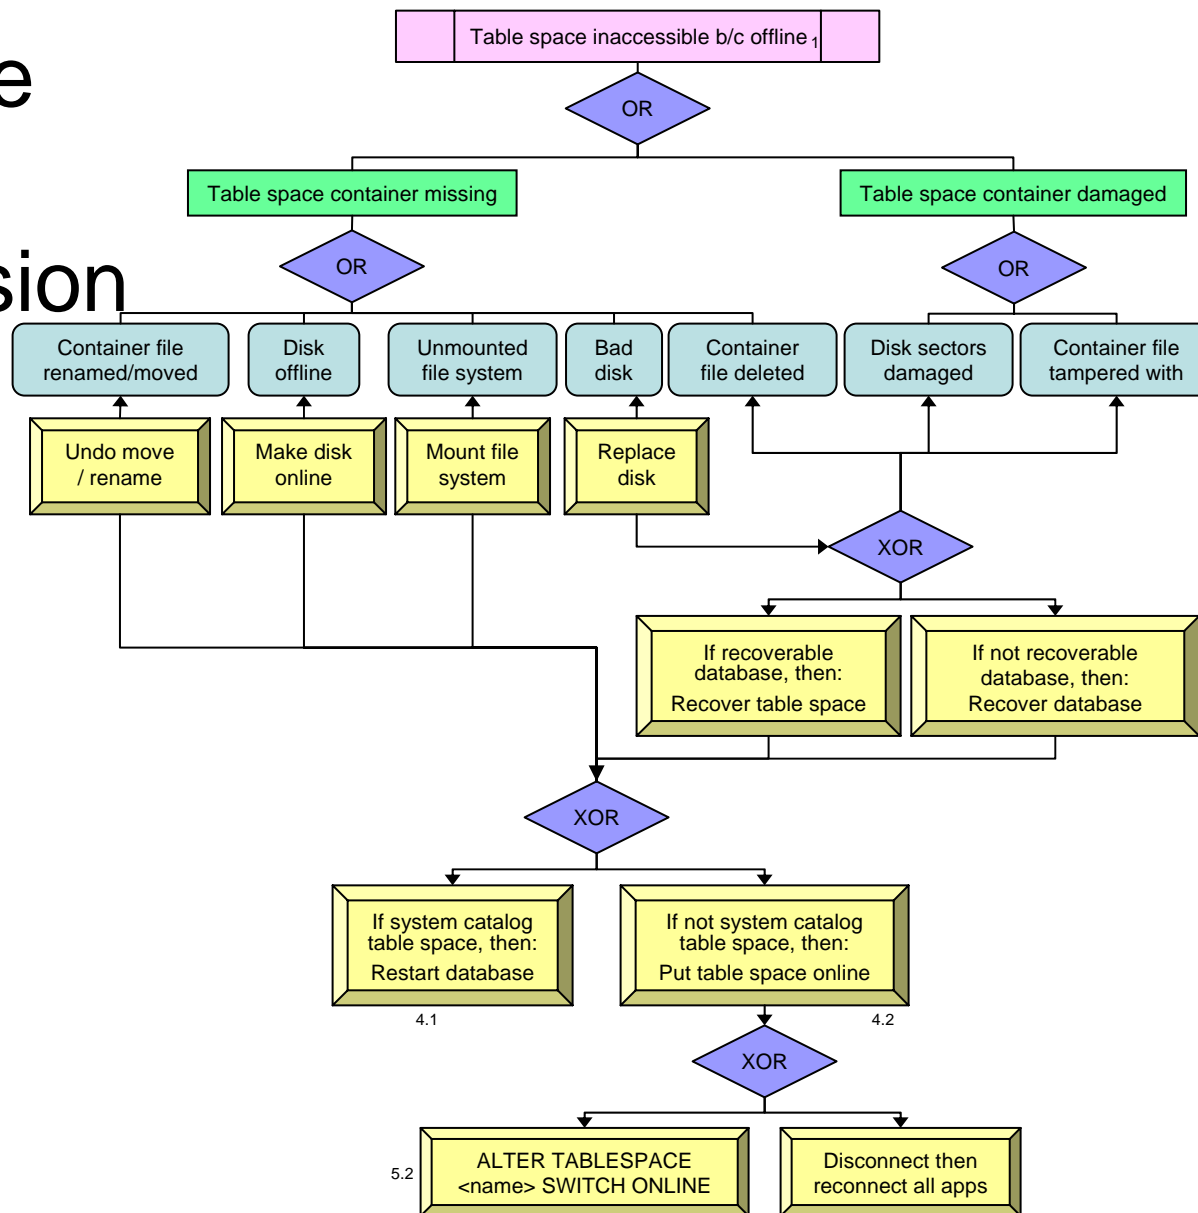
#### Diagnosing the problem

[Table space container is missing](#)  
If one or more containers of a table space cannot be found by the database management system, the table space will be taken offline and put in an inaccessible state.

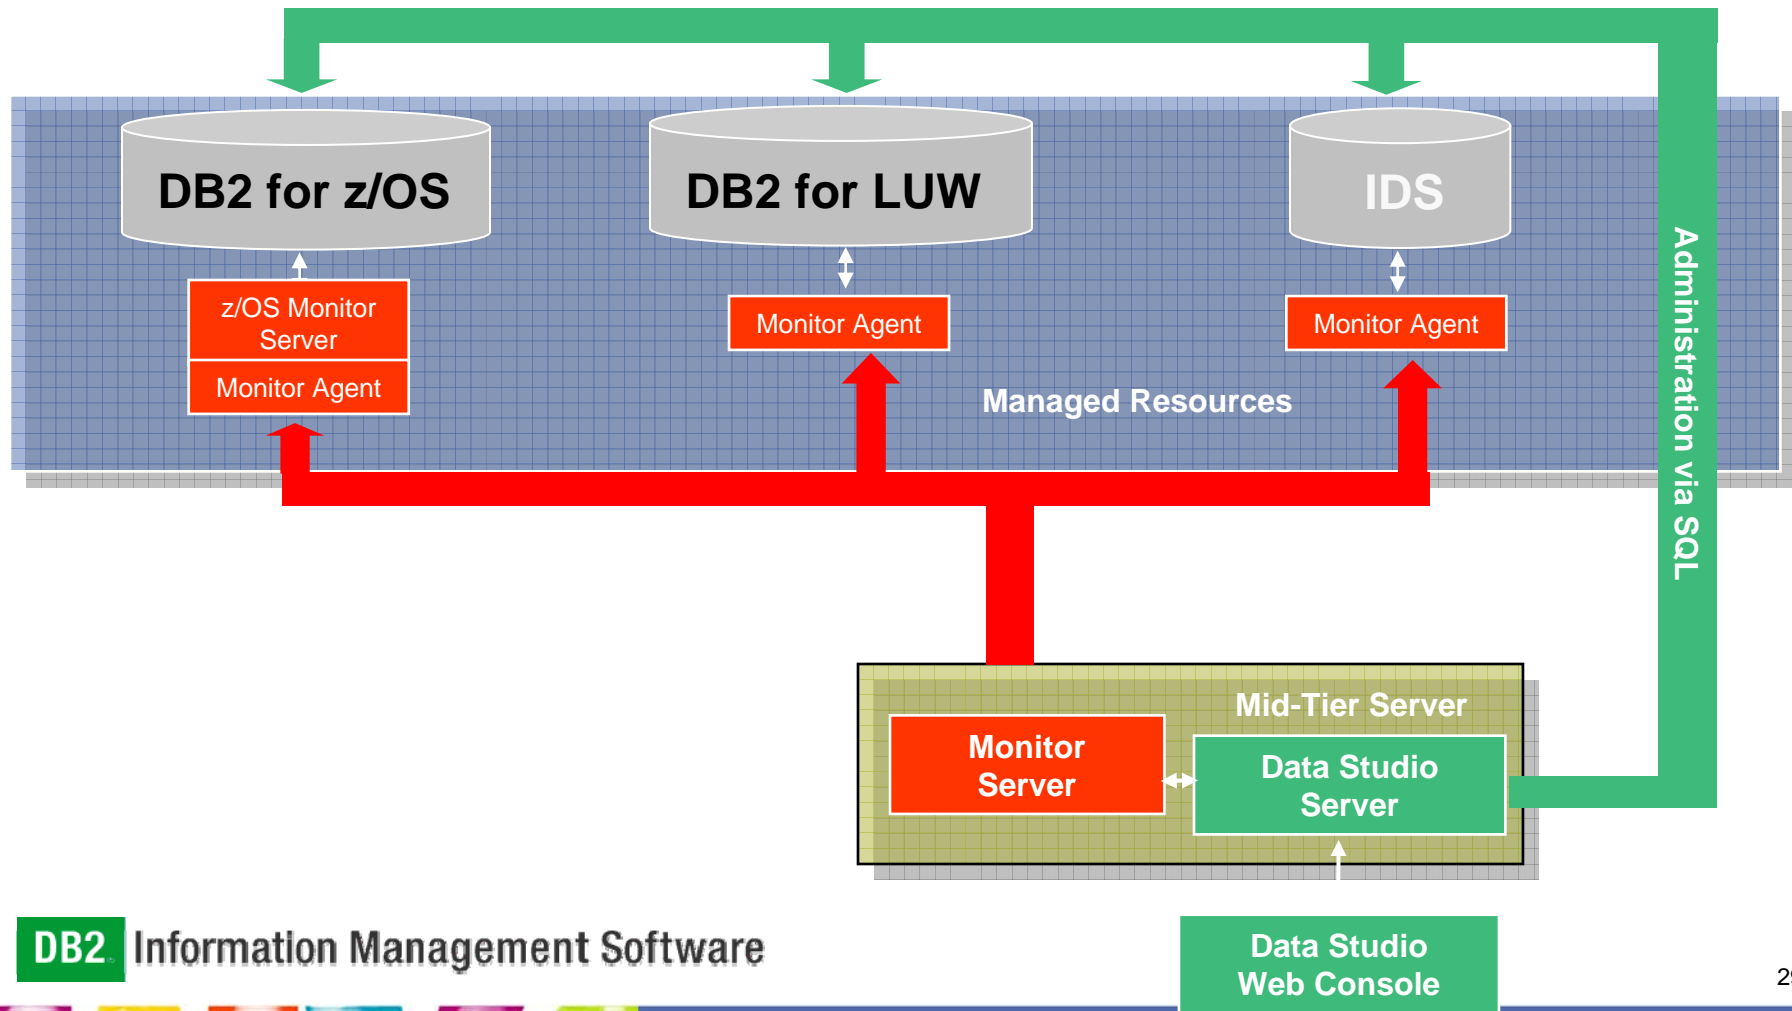
[Table space container is damaged](#)  
If one or more containers of a table space are found to be damaged by the data manager, the table space will be taken offline and put in an inaccessible state.



# Root Cause Analysis Sample Decision Tree



# Data Studio Strategic Monitoring Architecture





# Future Direction

*Act. Right. Now.*



## Top priorities for 2008/2009

- Reduce footprint of the client offerings:
  - Eclipse RCP version of Data Studio Developer for the Application DBA
    - SQL stored procedure create/debug, schema admin, etc.
    - No Java development tooling
  - Eclipse plug-in version for the Java developer
    - Plugs into Rational products (RAD, etc.)
    - Plugs into non-IBM Eclipse IDE solutions
  - Greatly reduce Tivoli agent footprint





## Top priorities for 2008/2009 -- continued

- Make DB2 and IDS the best choices for WebSphere:
  - pureQuery technology (tooling and runtime)
  - JDBC “capture” to enable pureQuery for any Java program
  - Significant performance monitoring and problem determination aids (added to DSAC and DB2 PE)
  - openJPA support for pureQuery
  - Spring, iBatis, and possibly Hibernate support for pureQuery
  - Tight integration with Object Grid (SQL syntax to cache, persistence to DB2 or IDS, replication from DB2 or IDS to Object Grid)



## Top priorities for 2008/2009 -- continued

- Upgrade DB2 for z/OS and IDS support to current functionality levels
- Separate client offerings from the LUW server deliverable
- Integrate key product offerings into Data Studio
  - DB2 Change Management Expert
  - DB2 Performance Expert
  - DWE
  - Princeton SoftTech
  - High Performance Unload



## Top priorities for 2008/2009 -- continued

- Deliver high-value Data Studio Performance Manager
  - SQL statement-level performance detail + historical trends
  - Ability to report database resource use in many ways
    - By SQL statement
    - By package or collection
    - By application
    - By app server
    - By Java class name
- Deliver initial Data Studio Query Tuner
  - Visual Explain of cached statements
  - Visual Plan Hints
  - Technology harvested from OSC and OE



## Top priorities for 2008/2009 -- continued

- Expand functionality in Data Studio Administrator
  - Alter anything (DB2 Change Management Expert)
  - View/change config parameters
  - FTAs for deadlock and timeout events
  - CREATE/START/STOP DATABASE and other CC items that improve up and running experience



# Where to get IBM Data Studio ?

## ■ IBM Data Studio

- [www.ibm.com/software/data/studio](http://www.ibm.com/software/data/studio)
  - FAQs / Tutorials
  - Downloads
  - [Forum](#) / Blogs
  - Join the IBM Data Studio user community

## ■ Meet the IBM Data Studio team at IOD 2007

- TLU-1298A
  - Monday | 3:45 PM | Mandalay Bay South Convention Center - South Seas F | Curt Cotner
- TDN-2627A - IBM Data Servers Application Development : The Details
  - Tuesday | 1:30 PM | Mandalay Bay South Convention Center - Mandalay Bay B | Rick Buglio
- HOL-1963A Developer Workbench for Application DBAs
  - Tuesday | 1:15 PM - 04:15 PM | Mandalay Bay South Convention Center - Breakers K
- HOL-2490A DB2 Viper 2 - Data Server Administration - The Next Generation
  - Tuesday | 9:30 AM - 12:30 PM | Mandalay Bay South Convention Center - Breakers C
- Developer Den and Expo Area



---

THANK  
YOU

