IBM Information On Demand 2007

IBM

IBM Data Studio and IBM's Java Persistence Strategy: pureQuery

http://www.ibm.com/software/data/studio

Stephen Brodsky, Architect, IBM Data Studio

Act.Right.Now.

December 2007

IBM Data Studio

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

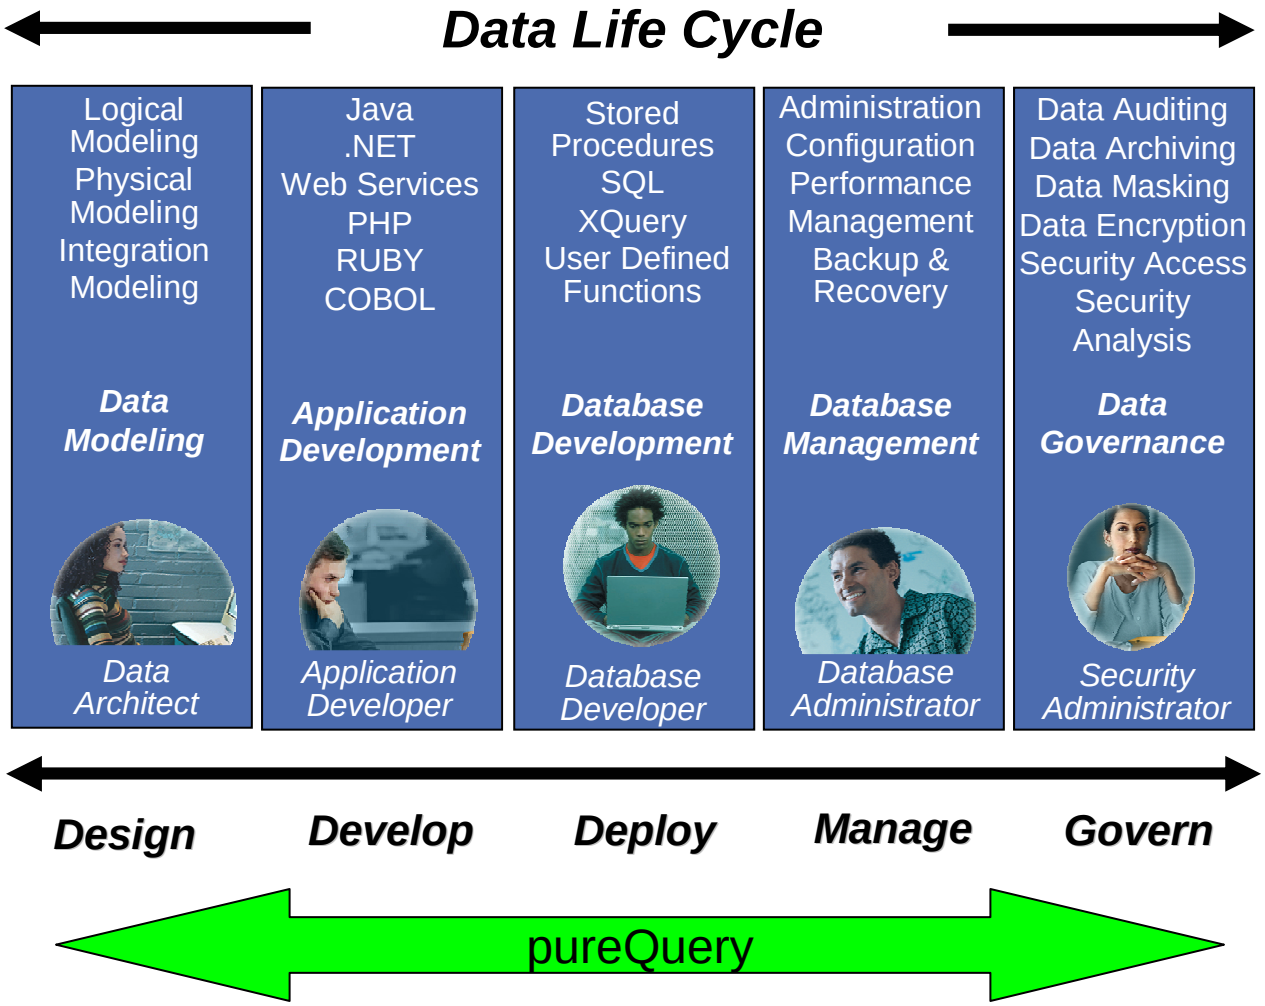

IBM Data Studio

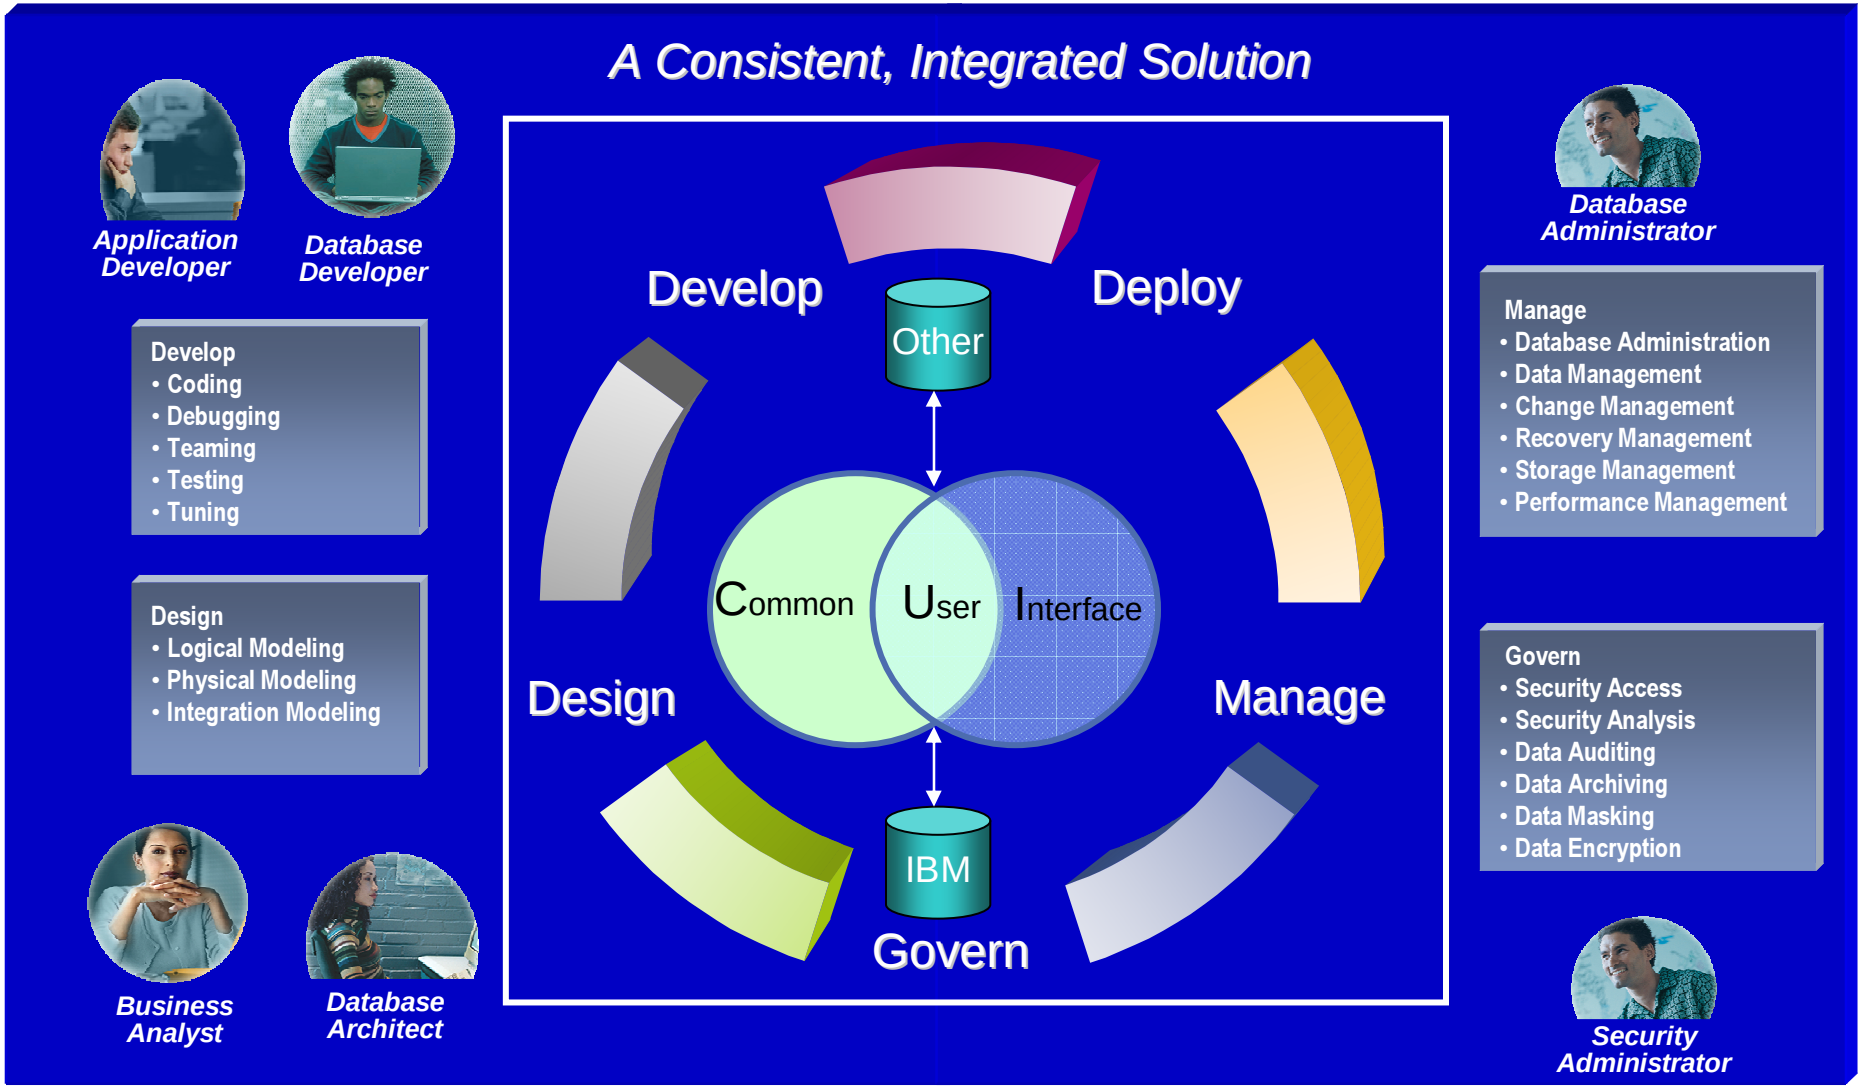

IBM Data Studio v1.1

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

DB2 for LUW	DB2 for z/OS	DB2 for i5/OS	IDS
Physical Data Modeling	Physical Data Modeling	Physical Data Modeling	Physical Data Modeling
Data Distribution Viewer	Data Distribution Viewer	Data Distribution Viewer	Data Distribution Viewer
Integrated Query Editor	Integrated Query Editor	Integrated Query Editor	Integrated Query Editor
SQL Builder	SQL Builder	SQL Builder	SQL Builder
SQL Routine Debugger	SQL Routine Debugger	SQL Routine Debugger	XML Editor
Java Routine Debugger	Java Routine Debugger	Java Routine Debugger	XML Schema Editor
XML Editor	XML Editor	XML Editor	pureQuery for Java
XML Schema Editor	XML Schema Editor	XML Schema Editor	Data Web Services
pureQuery for Java	pureQuery for Java	pureQuery for Java	Object Management
Data Web Services	Data Web Services	Data Web Services	Data Management
Object Management	Object Management	Object Management	Security Access Controls
Data Management	Data Management	Data Management	Project Management
Update Statistics	Update Statistics	Security Access Controls	
Health Monitoring *	Visual Explain	Project Management	
Visual Explain	Security Access Controls		
Security Access Controls	Project Management		
Project Management			

* Technical Preview

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 – Workbench

A single productive work environment

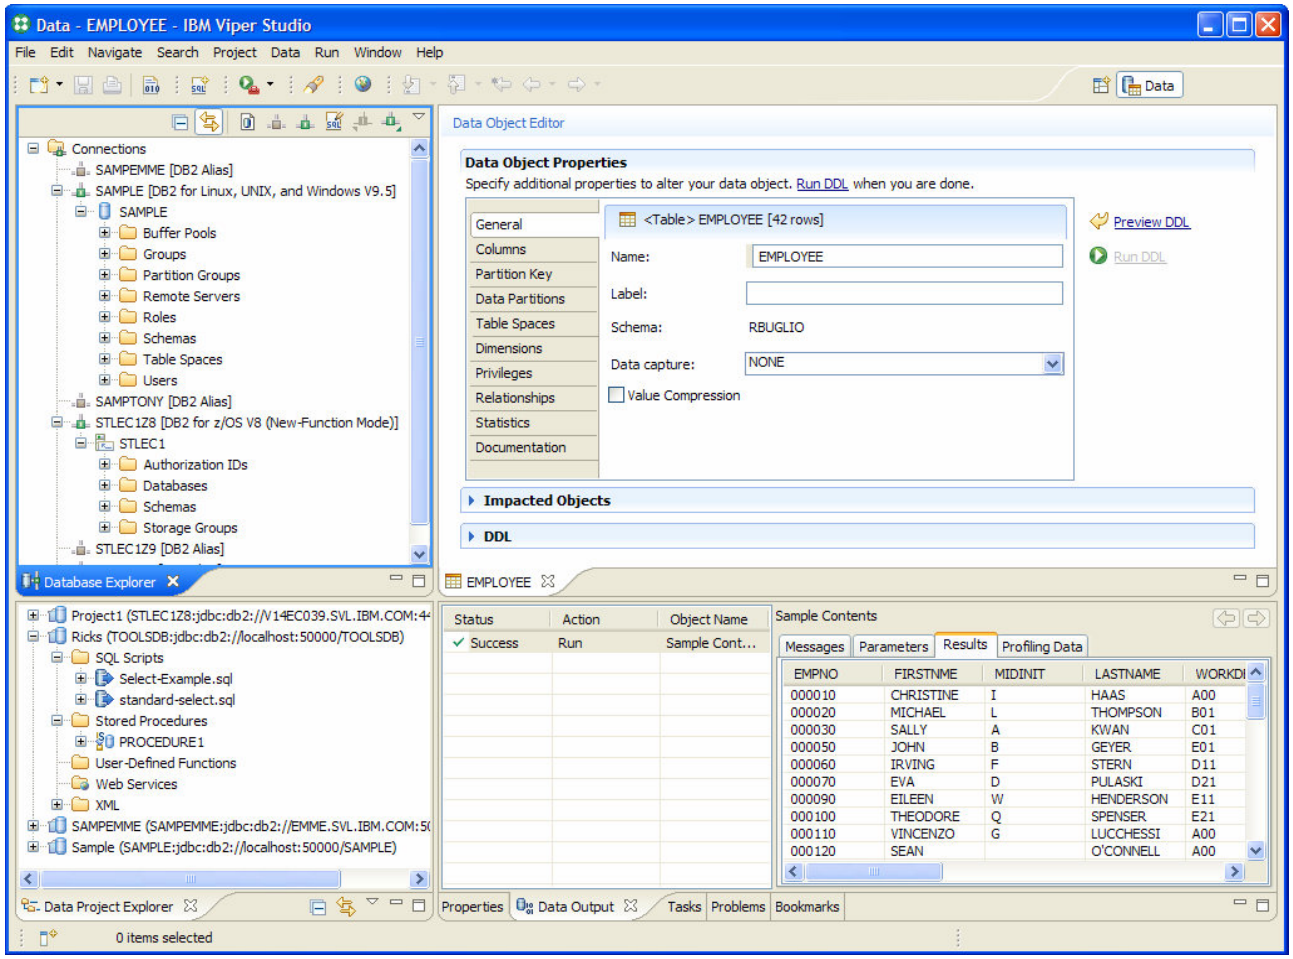

IBM Data Studio – Workbench

An integrated query editor for SQL and XQuery

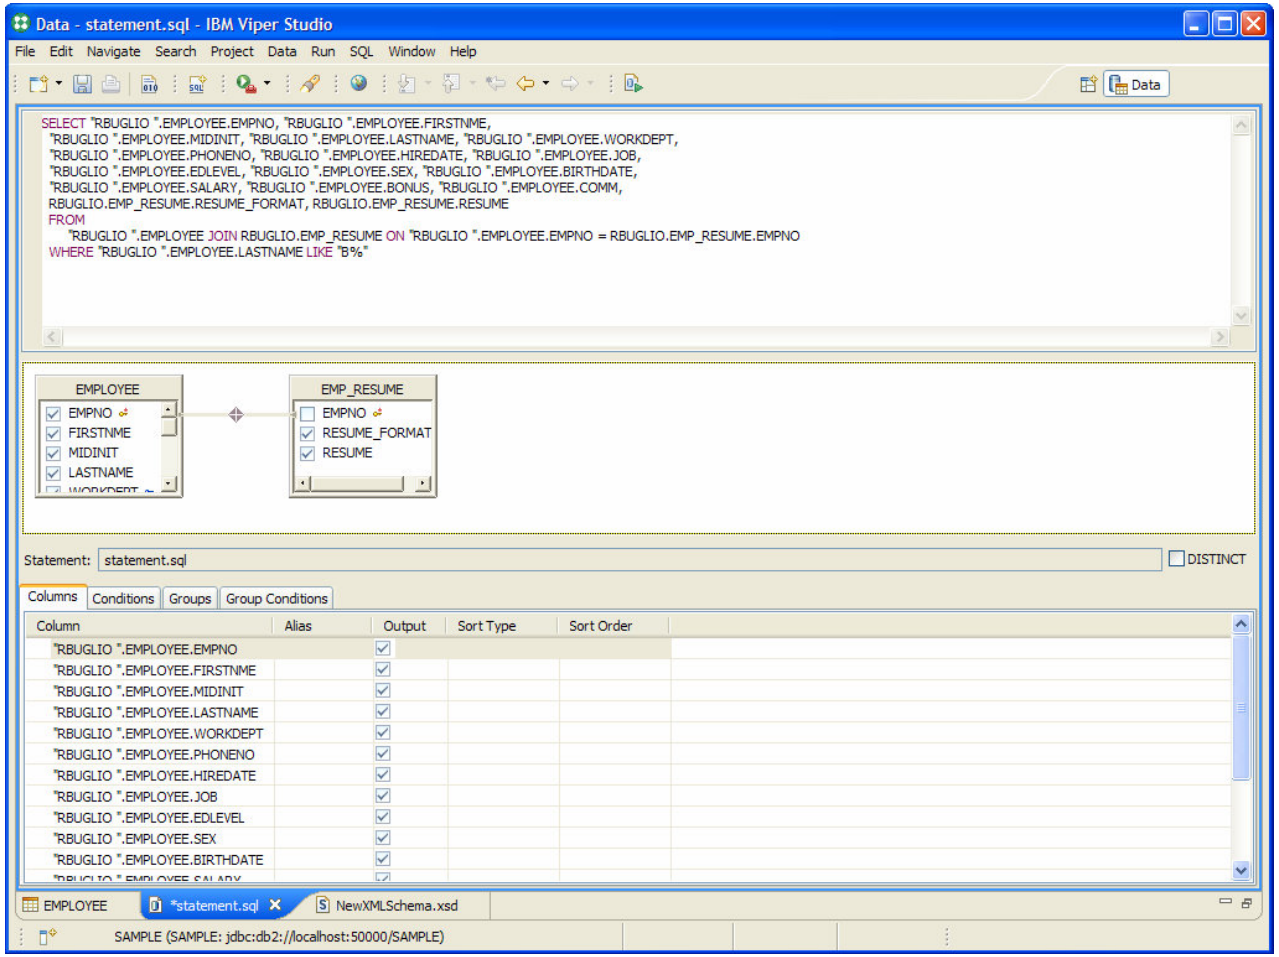

Data Web Services without Programming

Turn frequently used database operations as Web services for reuse

1. Create and Test Queries or Stored Procedures

2. Create Service

3. Drag ‘n Drop Resources

4. Deploy Service

5. Test and Deliver

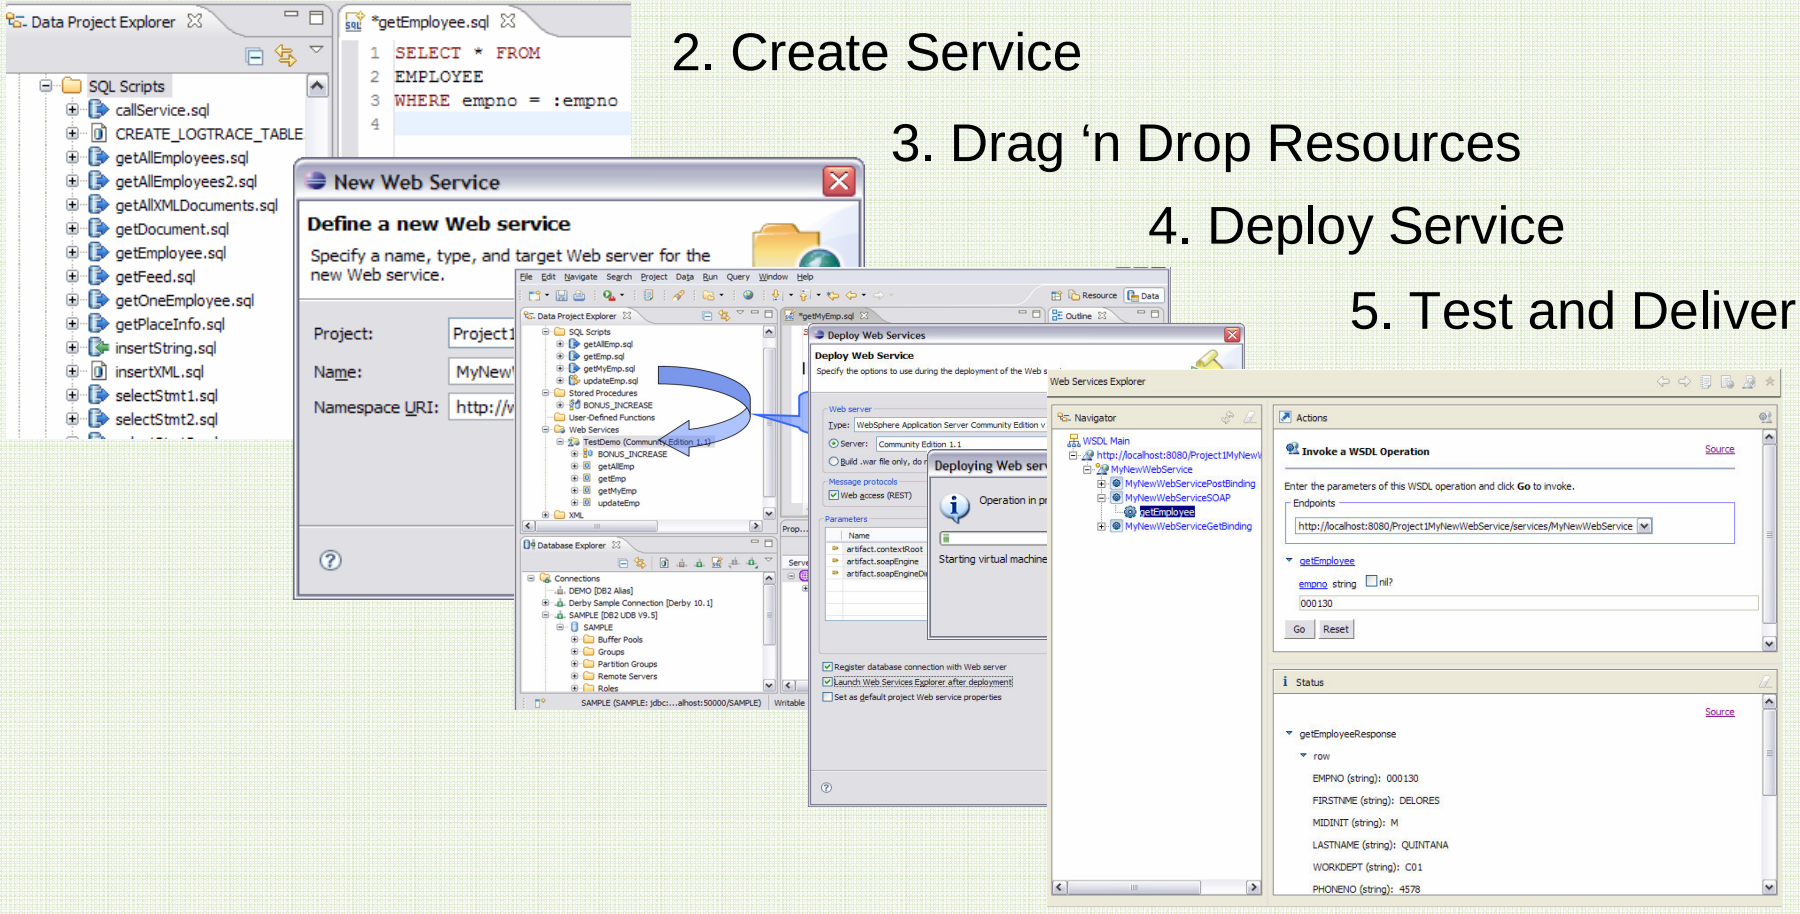

Quick & Easy Problem Determination

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

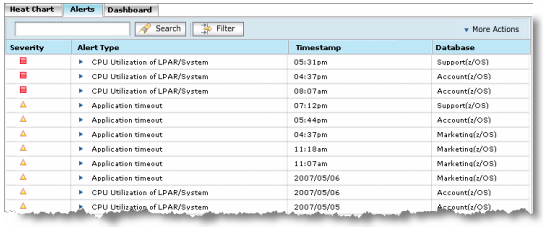

What happened when I was out for lunch? ... Away for weekend?

Recommendations – Root Cause Analysis

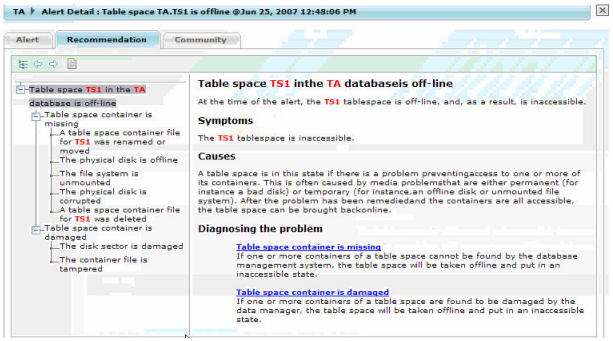

Guide me to the root cause and help me fix it properly; I need to know all the revelant info to make best decision.

Comparison - Developer Workbench vs. IBM Data Studio

2006

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 more

2007

IBM Data Studio

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

- Physical Data Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Servers Support
- pureQuery for Java*

IBM Data Studio – 2008 and Beyond

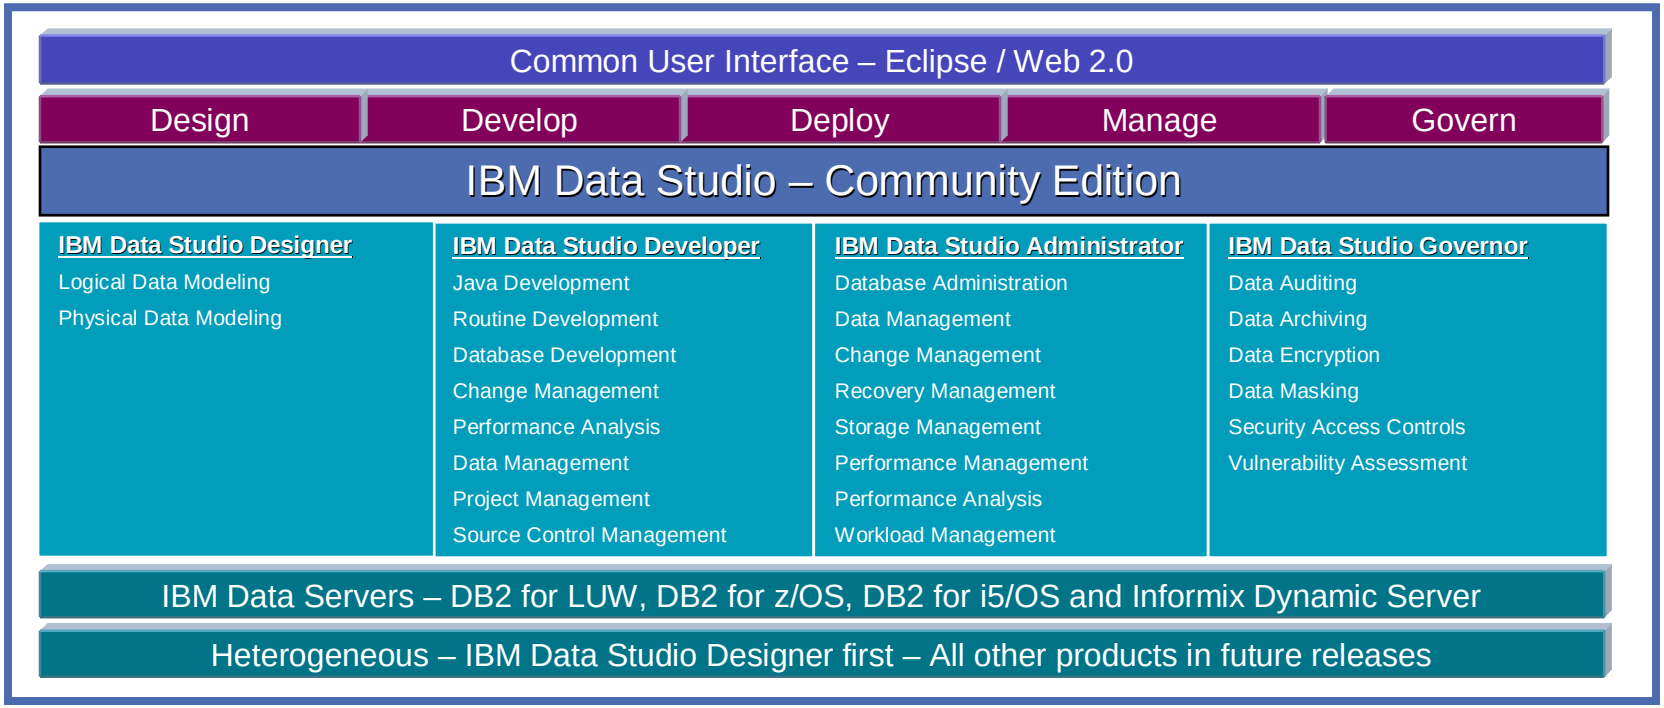

A Complete End to End Integrated Solution

- A single consistent, flexible and intuitive user interface
- Common code based shared by all products
- Products that can be plugged-in as needed to satisfy customer needs and requirements and meet changing market demands
- Products that can be sold individually, packaged as part of an edition or customized in any combination ala carte style
- Products packaged with value added components that satisfy the daily workflow requirements of specific end user roles
- Products that support all IBM Data Servers
- Products positioned to support Non-IBM Data Servers

Java access to relational – pureQuery on ramp for all cases

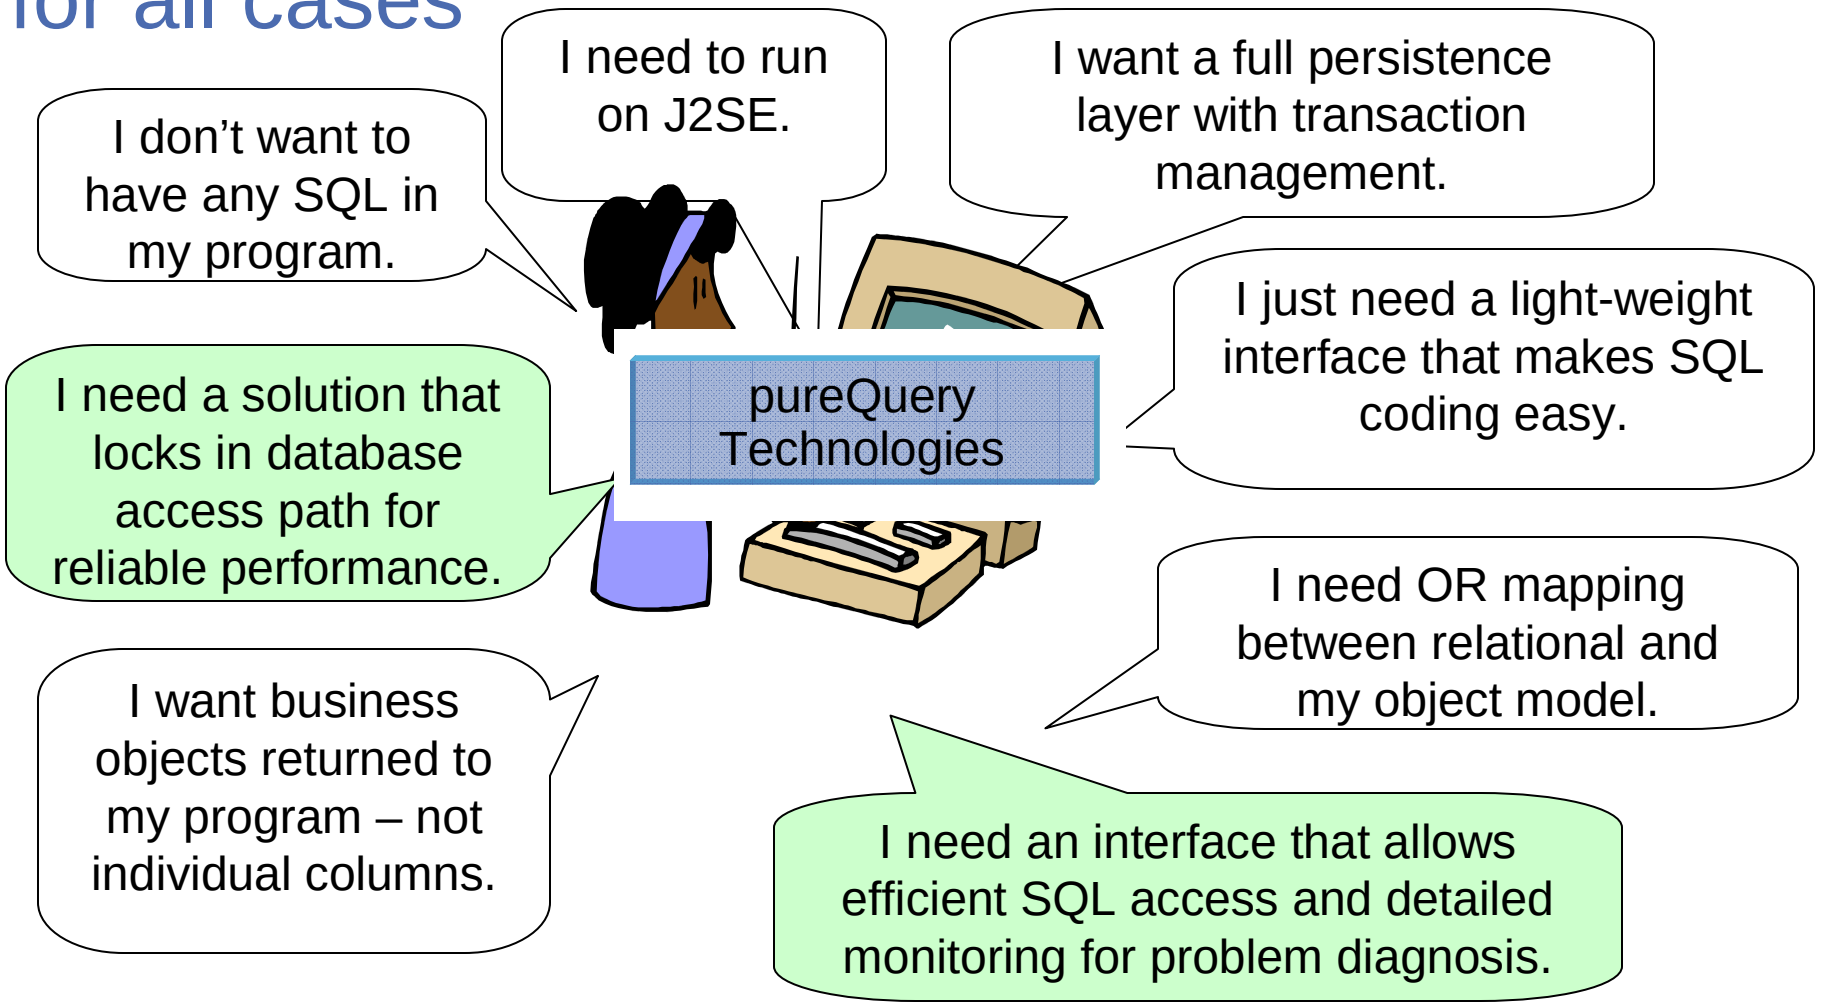

Java Data Access – many forms

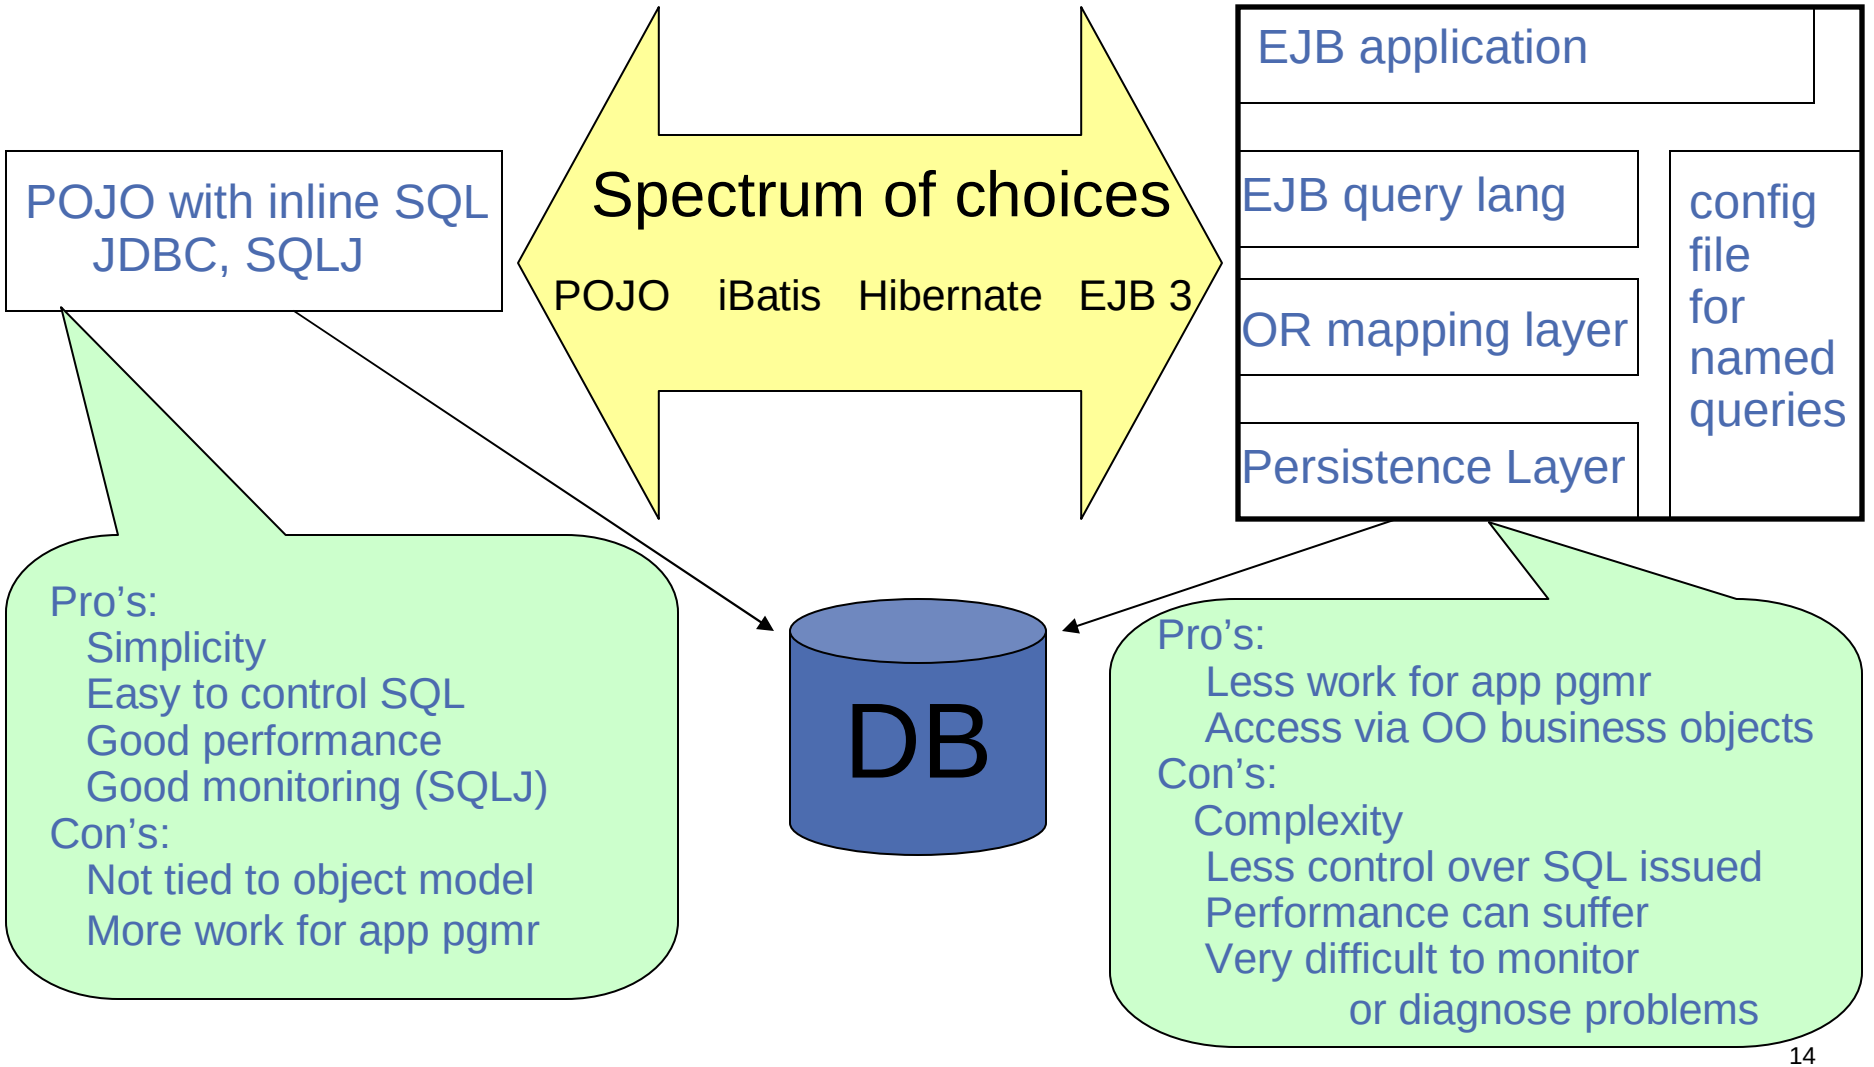

What performance/diagnosis challenges?

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.

OR mapping, Transform to SQL

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.

SQL issued to database:

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

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…

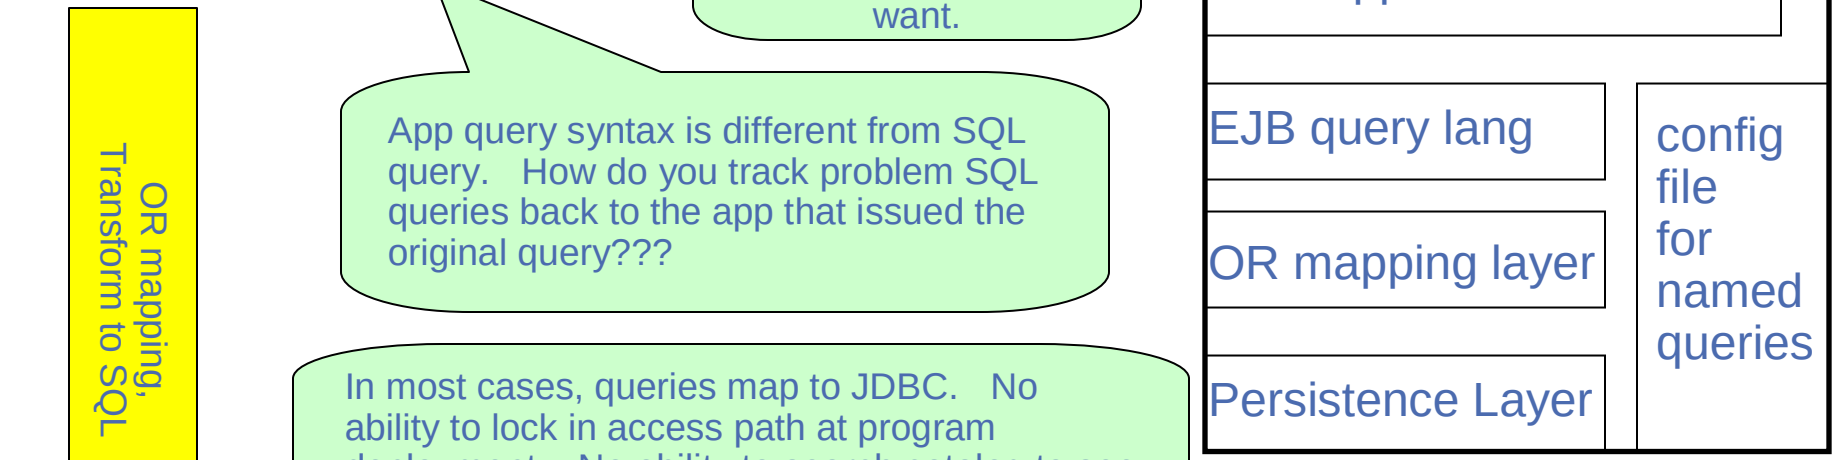

pureQuery

pureQuery is a high-performance Java data access platform focused on simplifying the tasks of developing and managing applications that access data.

Improves the Java data access life cycle

- Tools, APIs, and runtime environment
- Single API to query databases and in-memory Java objects
- Embraces SQL as the common query language
- Simplify SQL Data Access

pureQuery across the Life Cycle

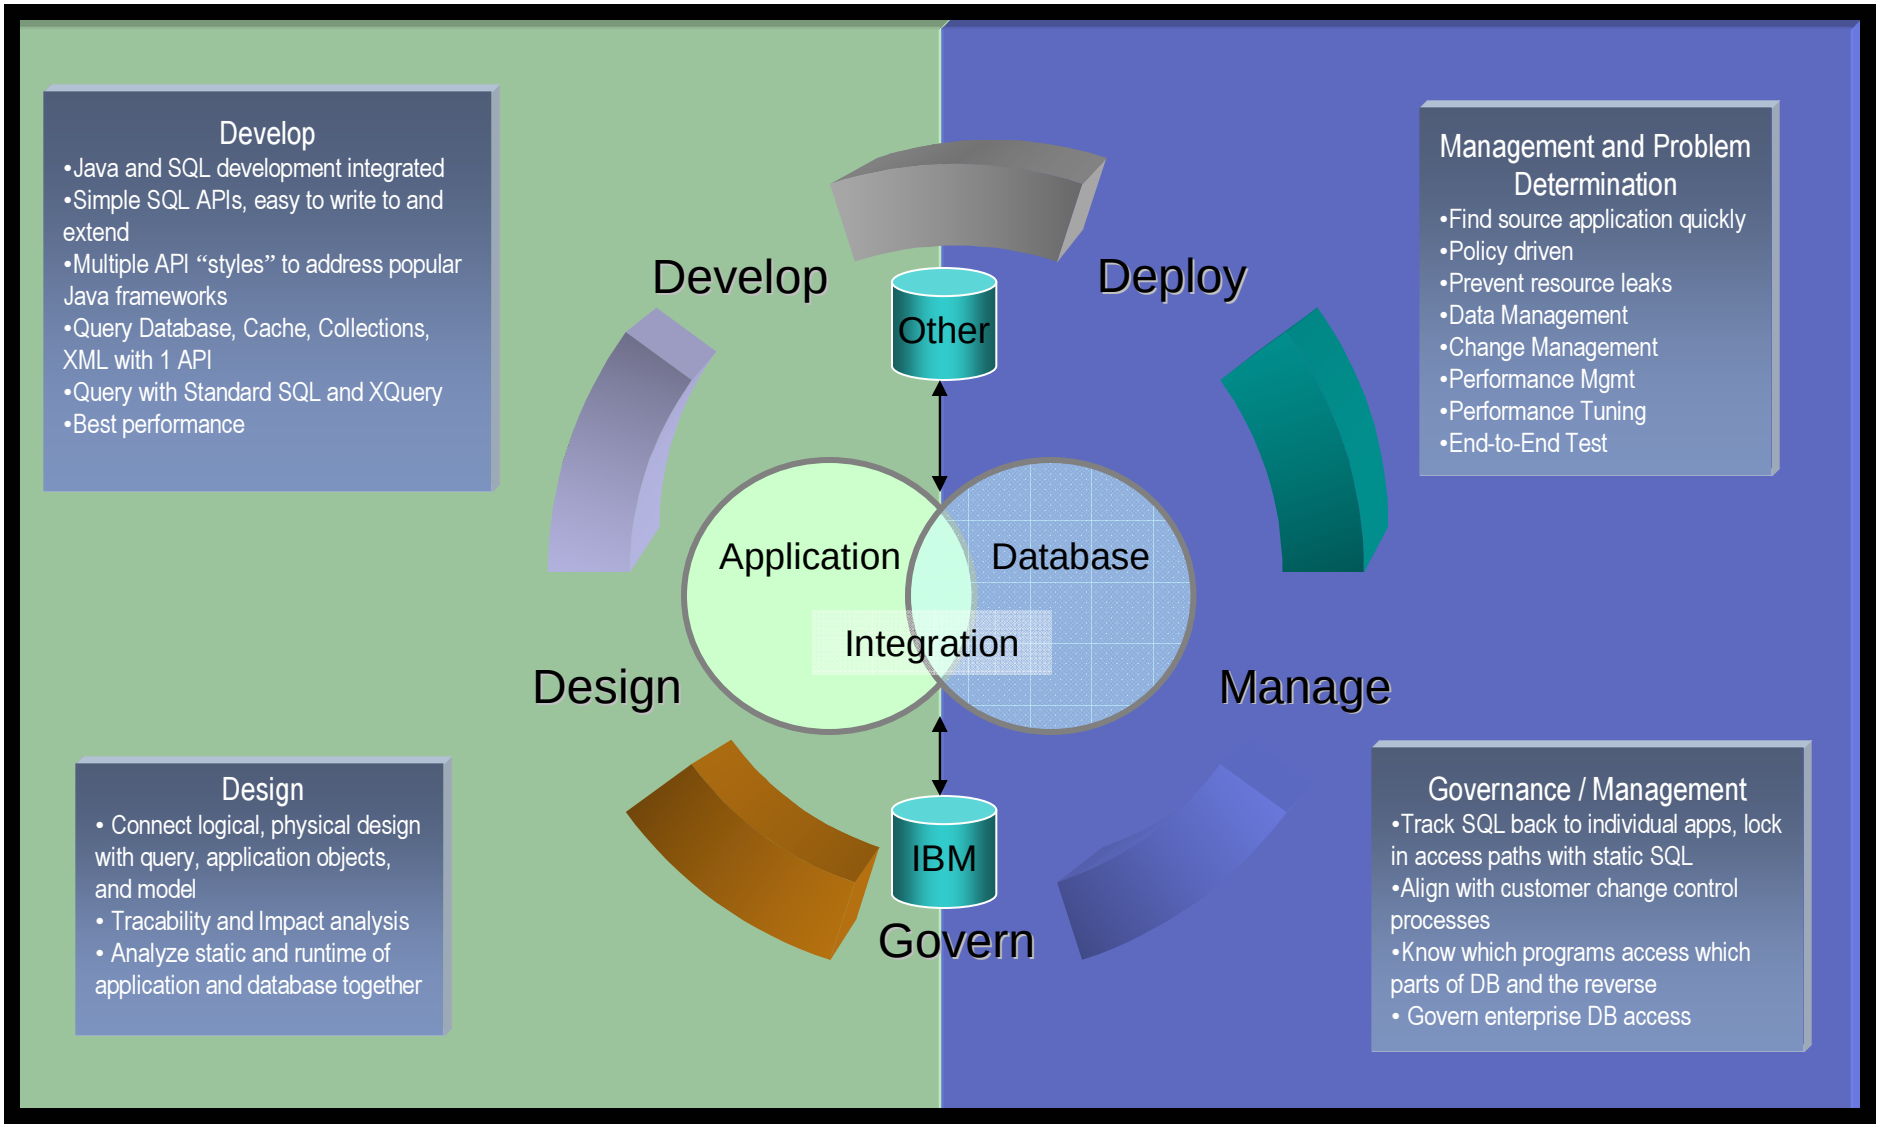

pureQuery – Across the Life Cycle

- 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, including JPA/EJB3
- Query important data sources simply
  - Database, Cache, Collections, XML
- Problem Determination
  - When problems occur, find source quickly across the tiers
- Governance / Management
  - Track SQL back to individual apps, lock in access paths with **static SQL** packages, align with customer change control processes
- 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

pureQuery with IBM Runtime/Tooling

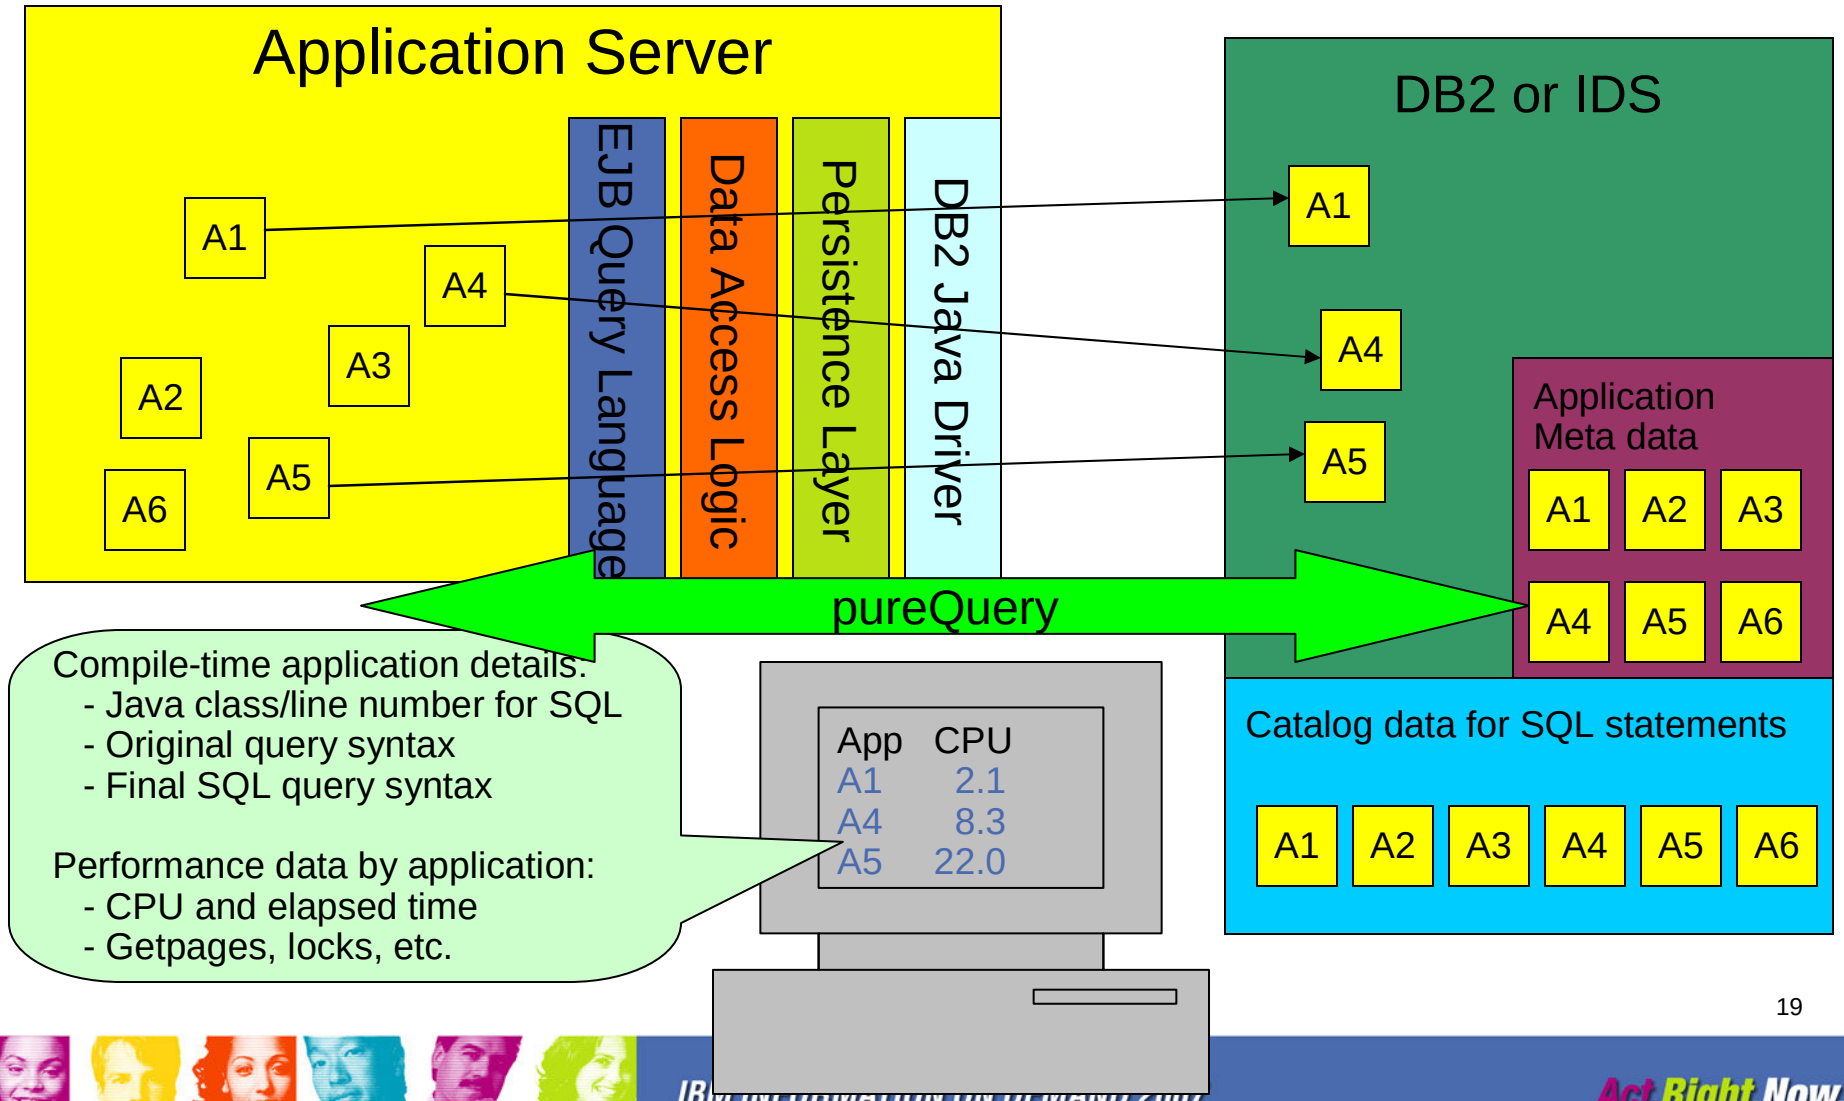

Toughest issue for Web applications – Problem diagnosis and resolution

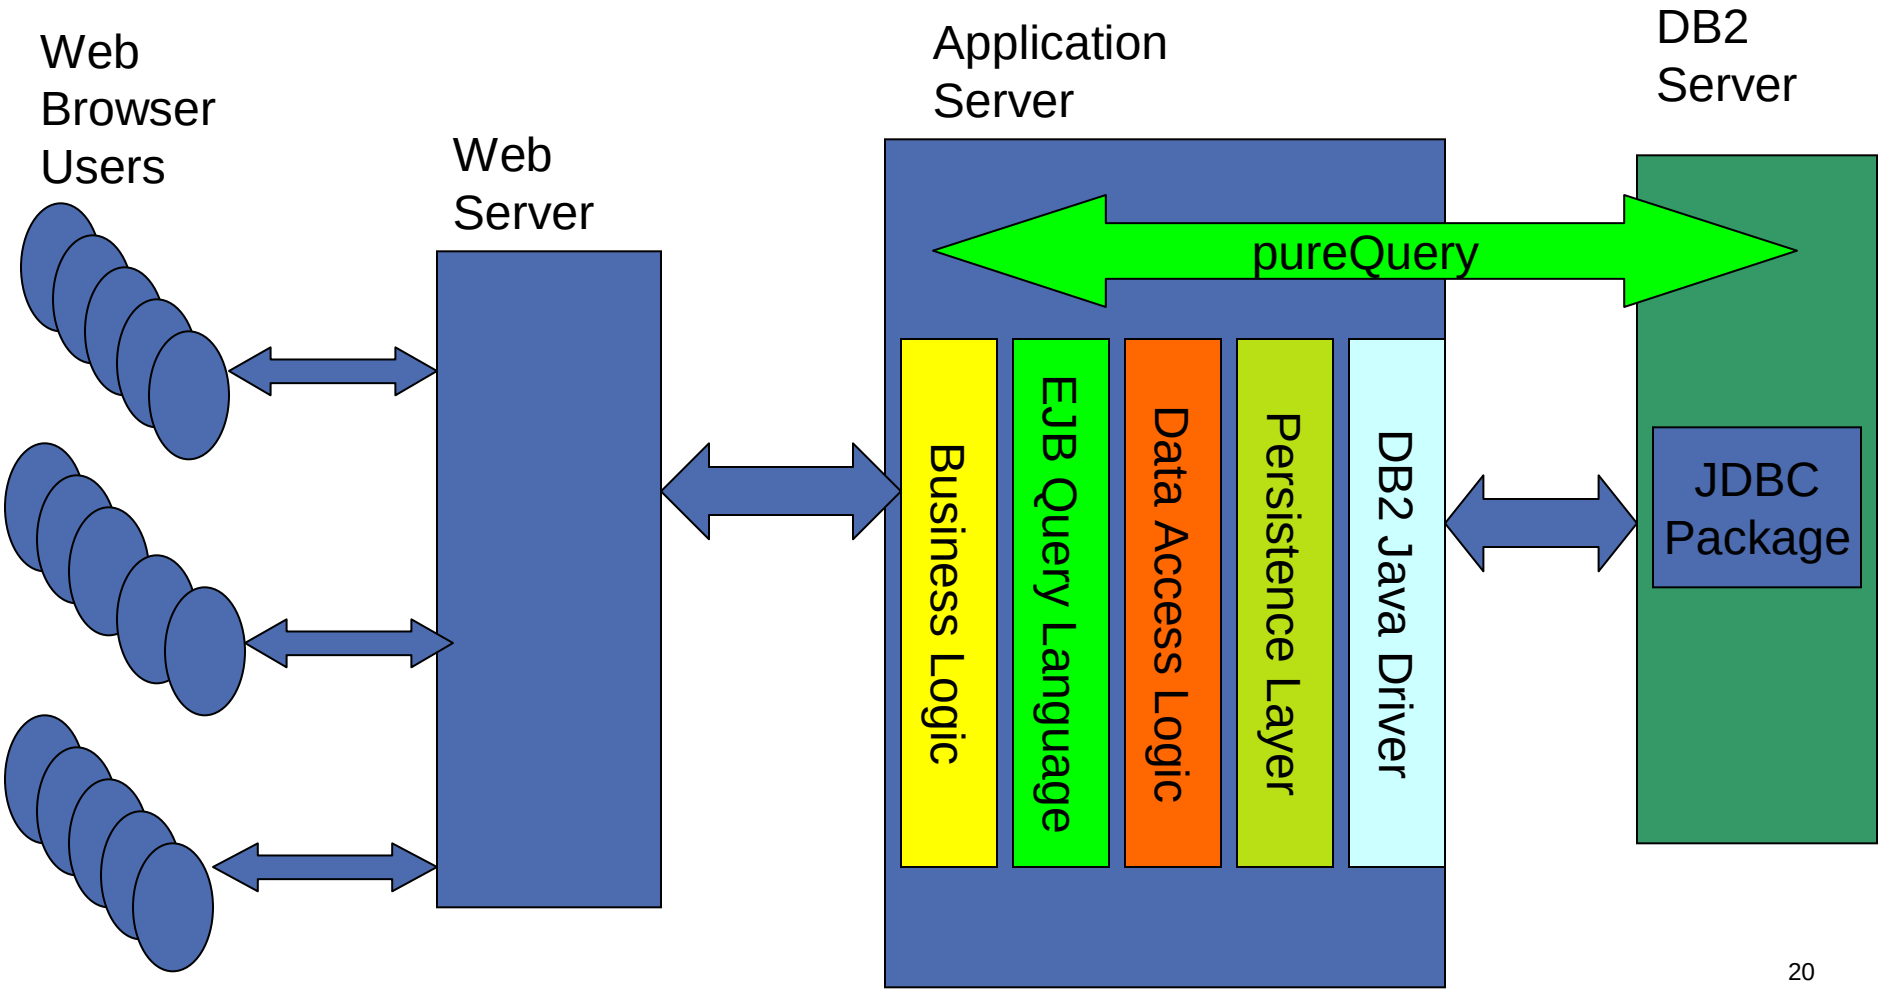

Simplifying Problem Determination Scenario

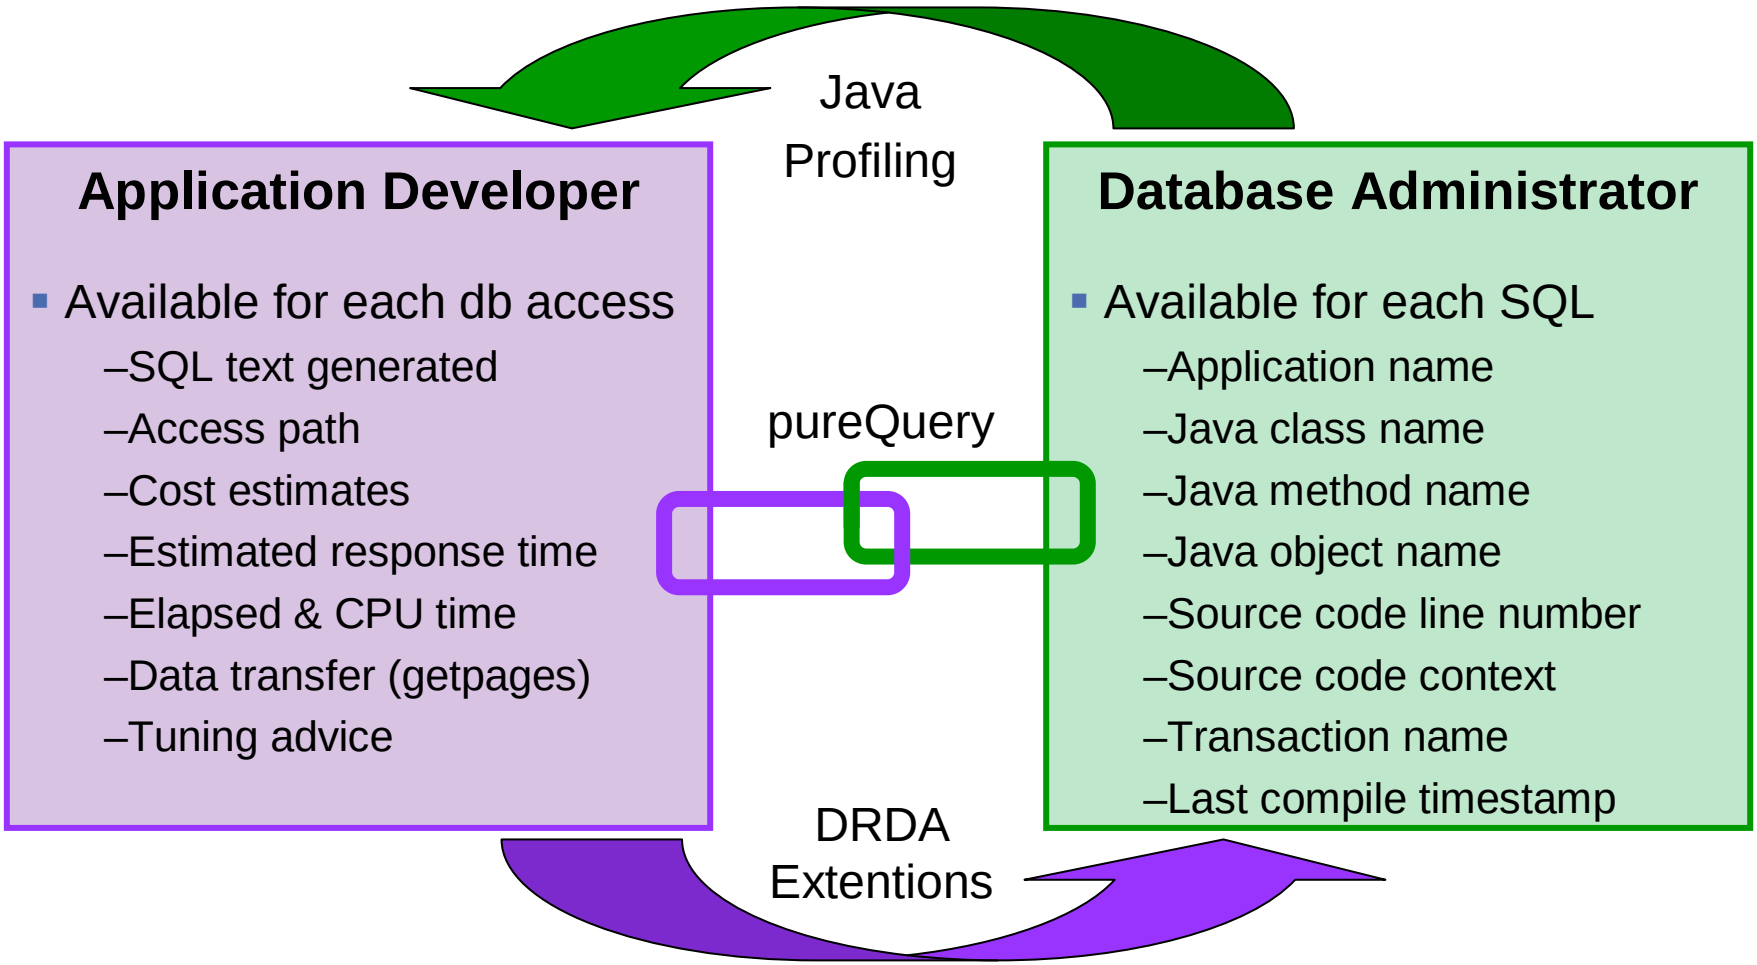

Scenario: Java Data Access in 4 simple steps

1. Select Table

2. Name Bean & Select Styles

3. Map Table to Bean

4. Select template SQL CRUD

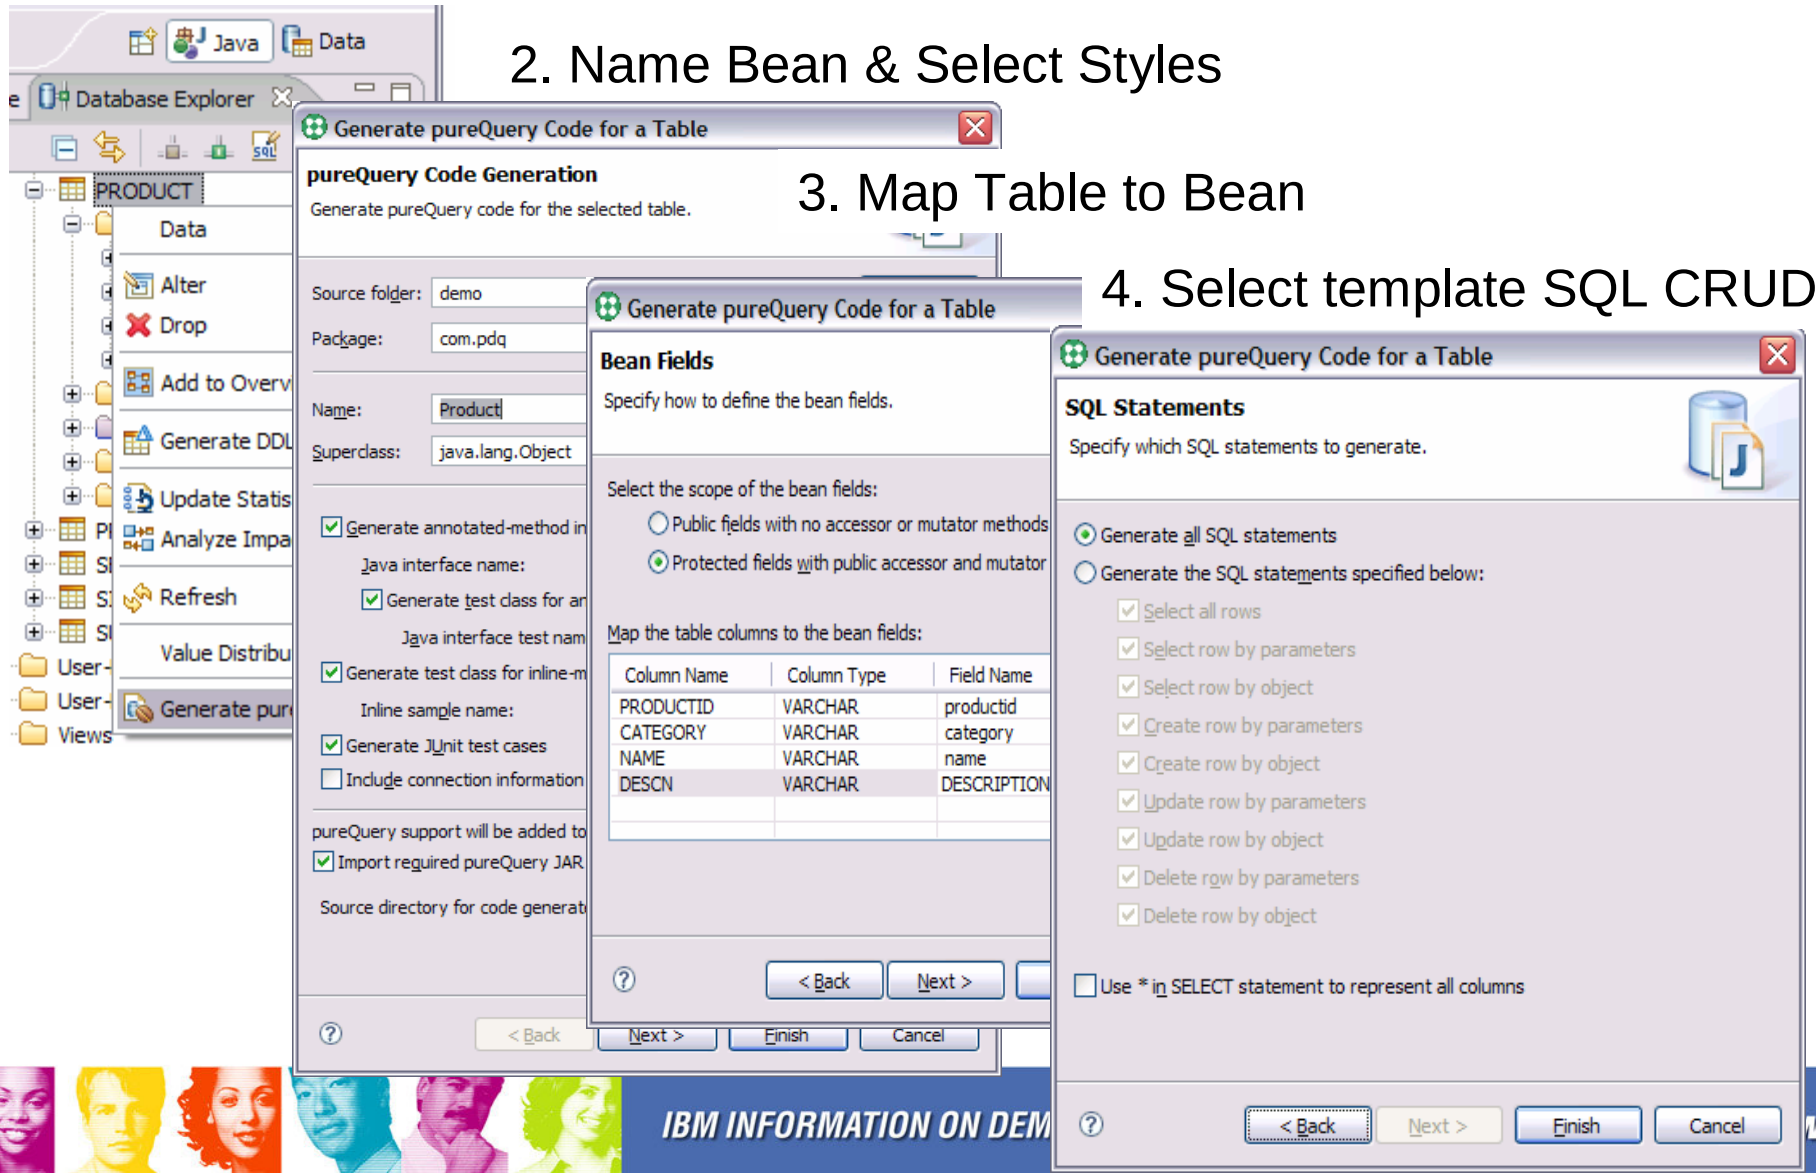

pureQuery Tools

- Generate basic Java data access objects and JUnit

- Integrated Java & SQL editor simplifies development
  - SQL content assist
  - SQL validation
  - SQL execution
  - SQL open definition

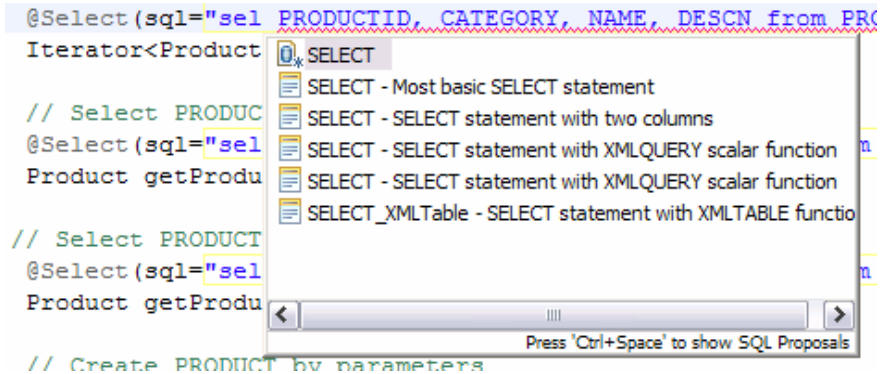

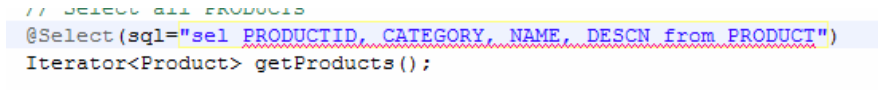

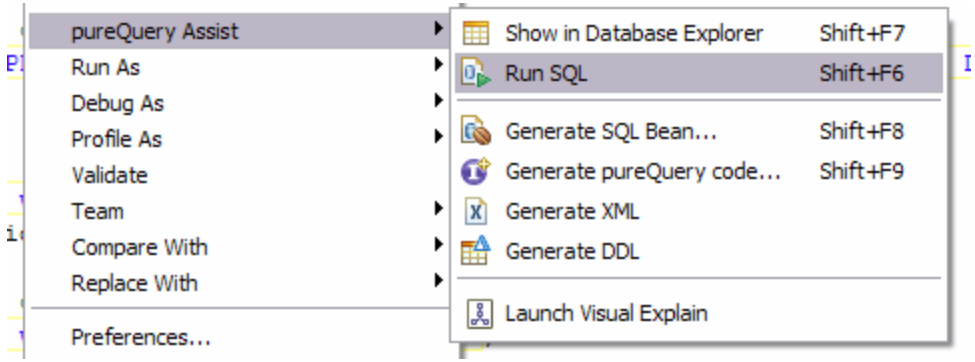

pureQuery Architecture

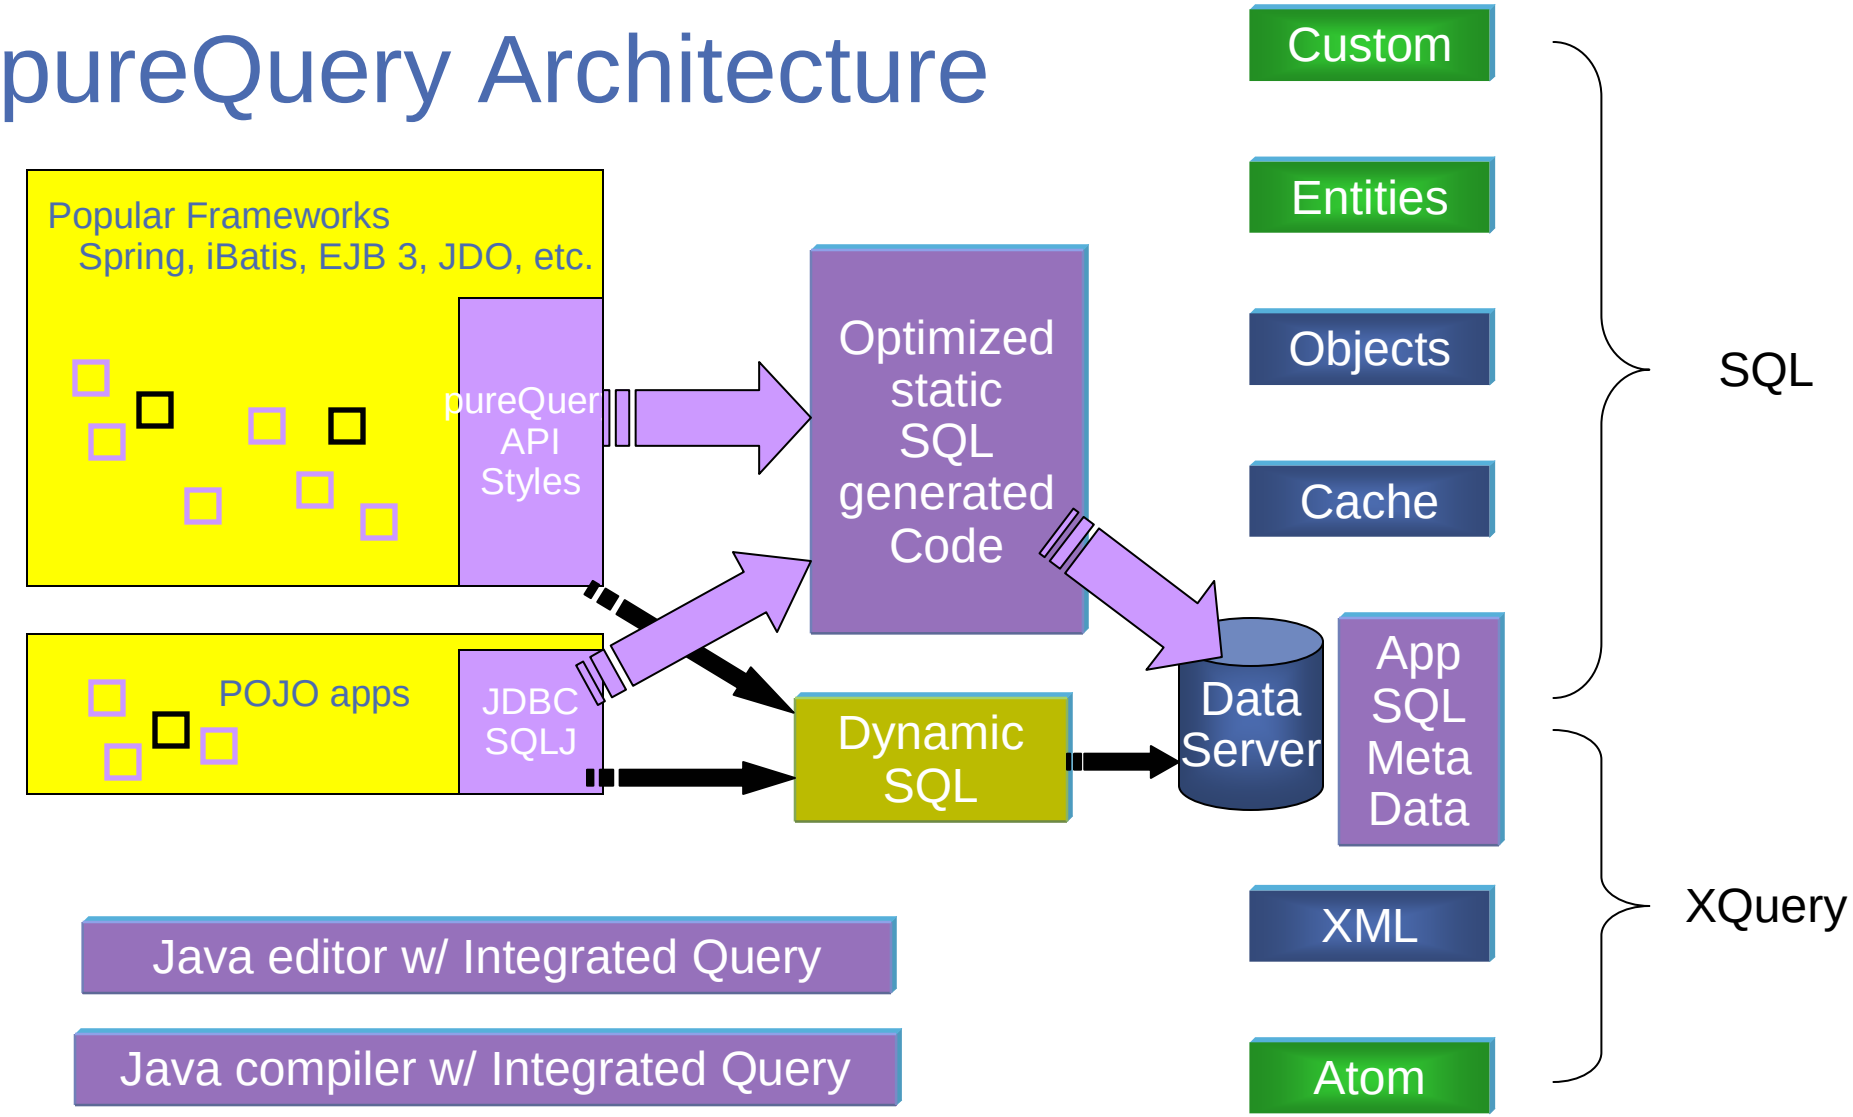

pureQuery enables wide use of Static SQL

- Static SQL
  - Highest speed
  - Greatest reliability
- JDBC is basic access, uses Dynamic SQL
- SQLJ adds Static SQL
- pureQuery supports both Static SQL and Dynamic SQL
  - Code to dynamic SQL, turn on static SQL at deployment

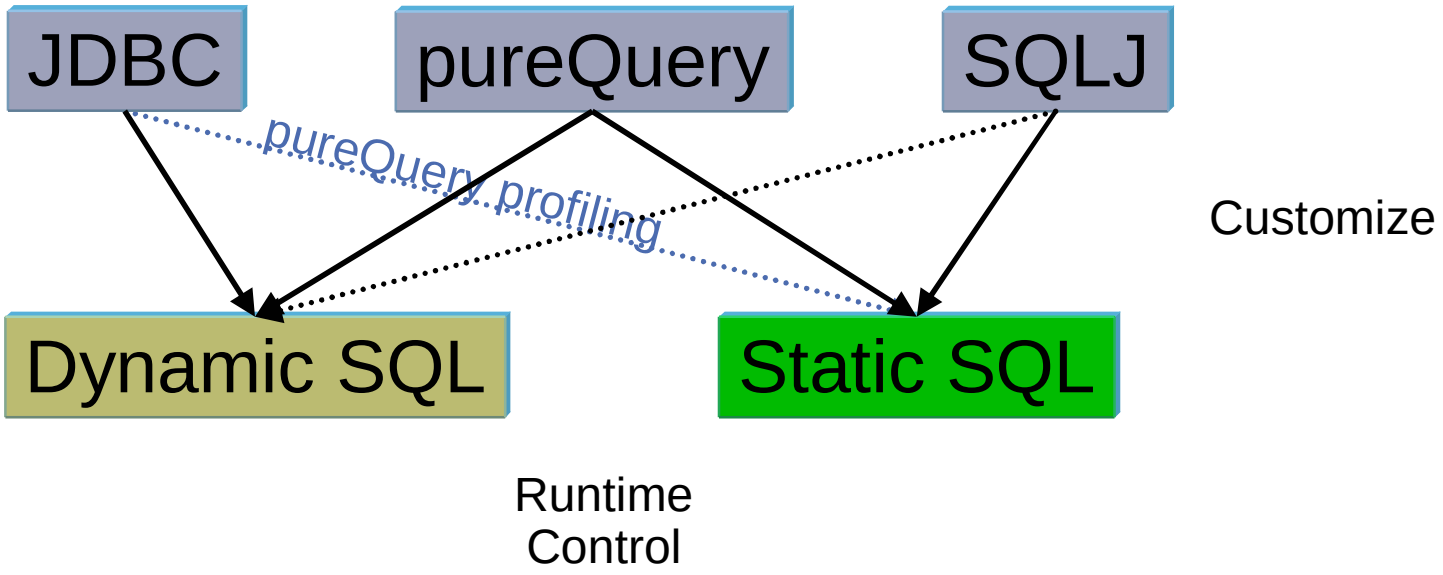

Static SQL Advantages – Comparison with Dynamic SQL

	Dynamic SQL (pureQuery, JDBC)	Static SQL (pureQuery, SQLJ)
Performance	Can approach static SQL performance with help from dynamic SQL caches. Cache misses are costly!	All SQL parsing, catalog access, done at BIND time. Fully optimized during execution.
Access Path Reliability	Unpredictable – Any prepare can get a new access path as statistics or host variables change	Guaranteed – locked in at BIND time All SQL available ahead of time for analysis by EXPLAIN.
Authorization	Pivileges handled at object level. All users or groups must have direct table privileges – Security exposure, and administrative burden	Privileges are package based. Only administrator needs table access. Users/Groups have execute authority. Prevent non-authorized SQL execution.
Montoring, Problem Determiation	Database View is of the JDBC or CLI package – No easy distinction of where any SQL statement came from.	Package View of applications makes it simple to track back to the SQL statement location in the application
Capacity Planning, Forecasting	Difficult to summarize performance data at program level.	Package Level Accounting gives program view of workload to aid accurate forecasting.
Tracking Dependent Objects	No record of which objects are referenced by a compiled SQL statement	Object dependencies registered in database catalog

Java On-Ramps to pureQuery

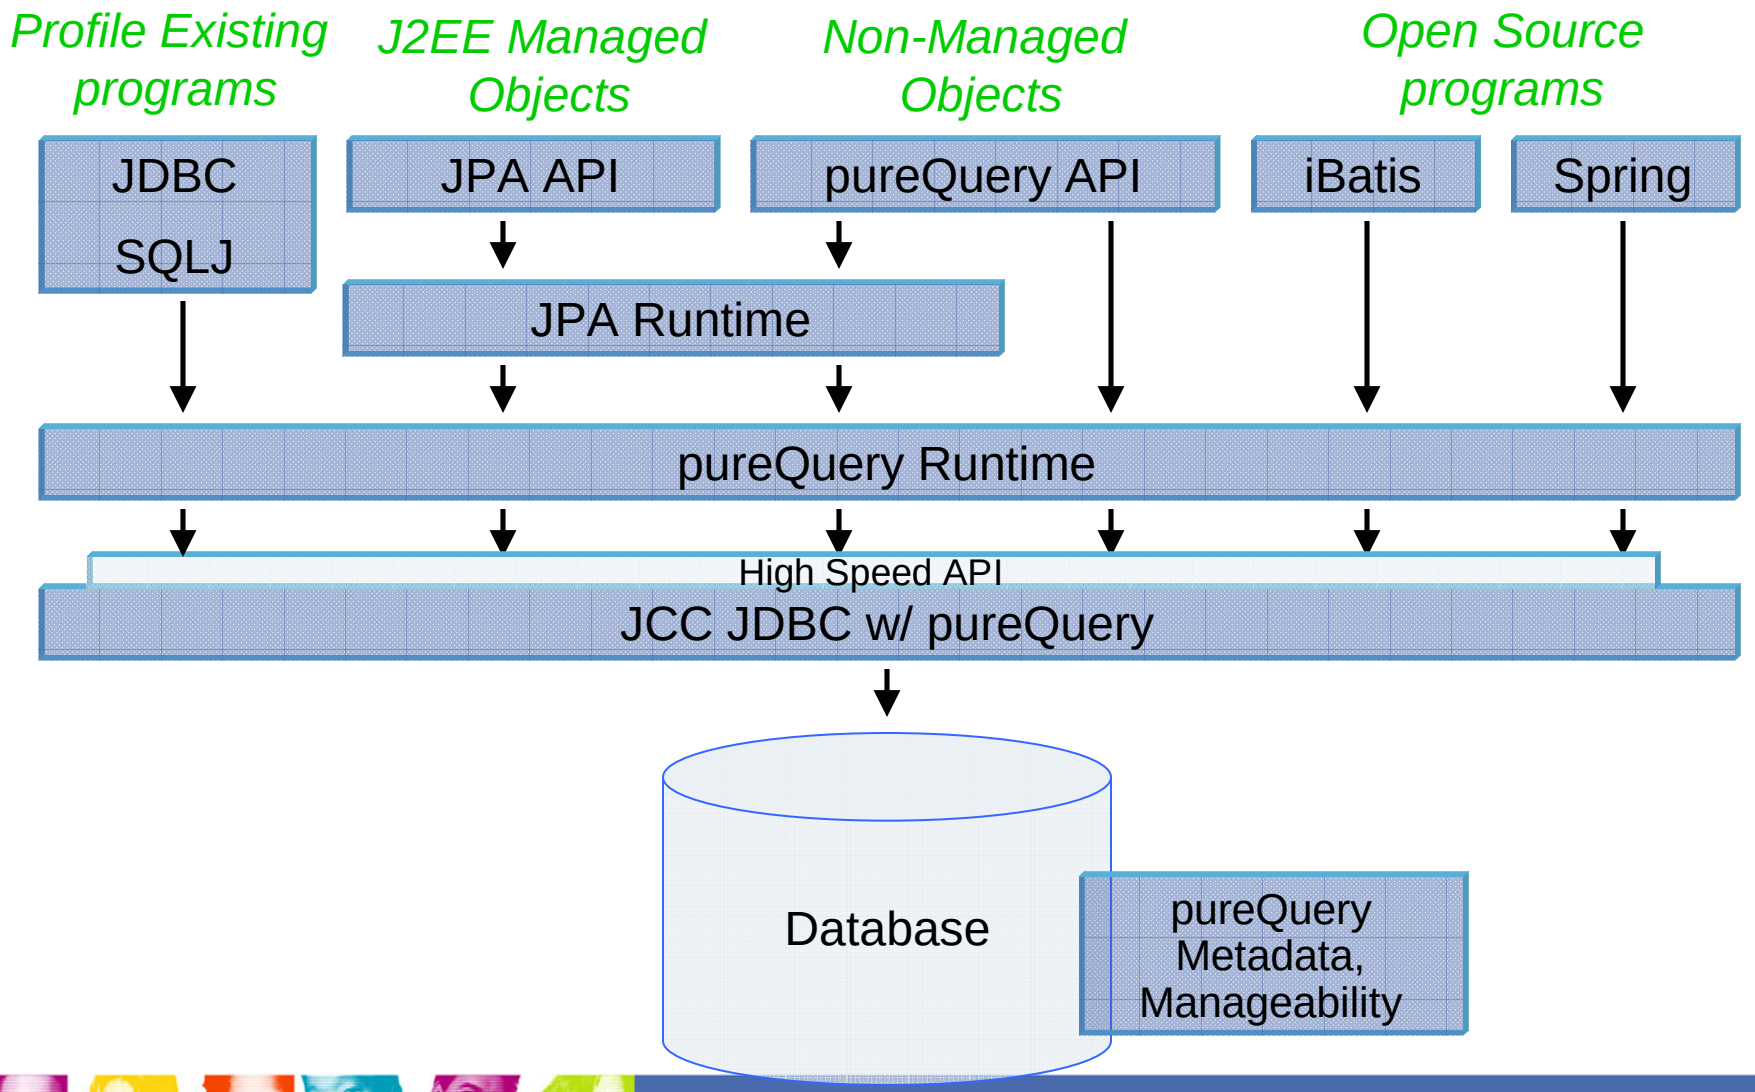

JPA and pureQuery together

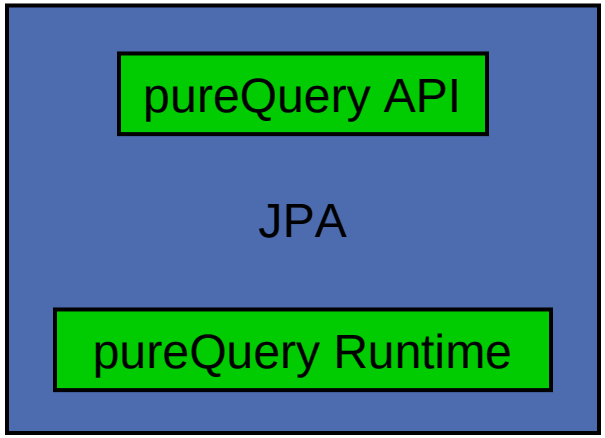

- EJBQL and runtime SQL generation based on object manipulation make the value of pureQuery even more important in the JPA setting
- IBM is enhancing our JPA implementation with both pureQuery APIs and pureQuery runtime lifecycle benefits
- JPA w/pureQuery enables problem determination, optimization, and governance connecting the EJBQL and business logic to the actual SQL and database operation
- JPA / EJB3 is a J2EE5 standard
- WebSphere is delivering JPA
- Apache openJPA is the only JPA implementation supported by more than one major vendor: BEA and IBM
- Hibernate users should use the JPA standard APIs and migrate to the openJPA implementation.

Java Data API Space

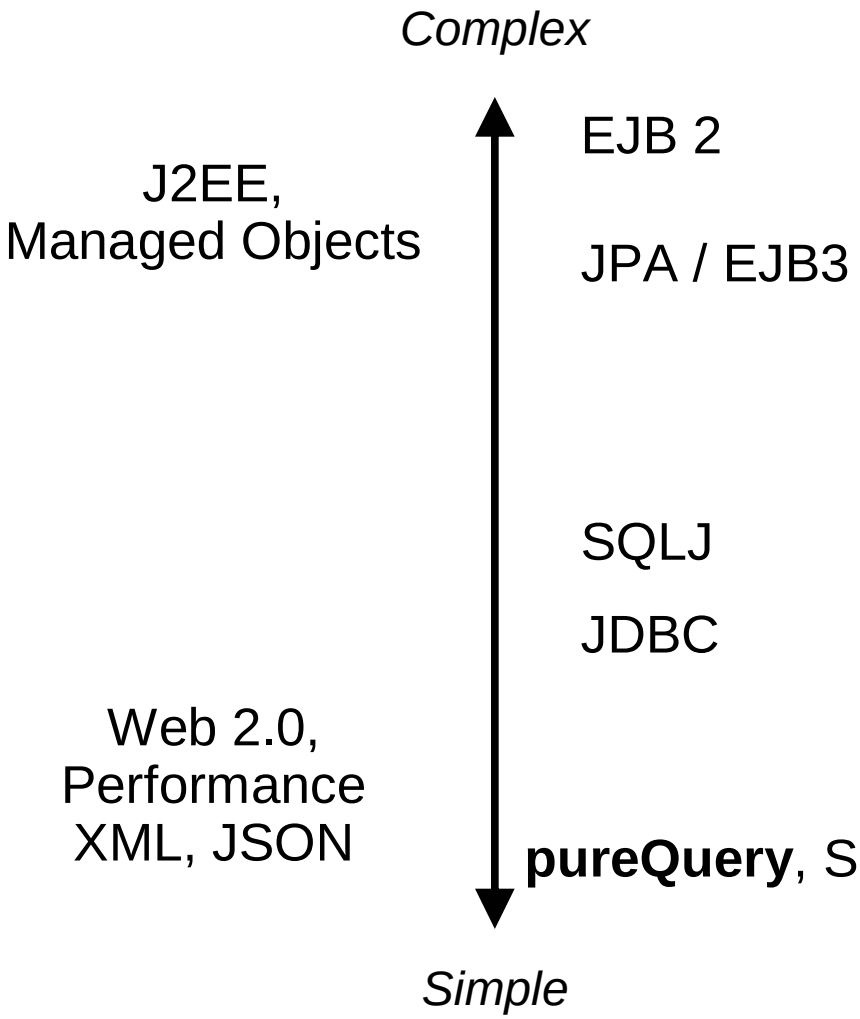

J2EE Managed Objects Complex O/R	JPA + **pureQuery**
All other cases ▪Development speed ▪Performance ▪Web 2.0 / AJAX ▪XML / JSON ▪Dynamic & Static SQL	pureQuery
Already using SQLJ	SQLJ + **pureQuery**

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 Technology SQL Query API

- Simple, straightforward programming model for data access
  - A fairly thin layer on top of JDBC that simplifies the most common tasks
  - Supports DB2, IDS, Oracle, SQL Server, etc. (any JDBC database)
  - Out-of-the-box support for storing/retrieving Beans and Maps to/from the database
- Extensible framework
  - Pluggable custom result processing patterns
    - Use Java to implement the mapping behavior instead of a “mapping language”
    - Instantiate result types other than Beans and Maps
  - Framework itself uses the same extension points to provide the out-of-the-box behavior
    - Library of the most common patterns
- Full expressiveness of SQL available
  - In practice, even simple applications do “sophisticated” SQL
- SQL inlined in data access methods
  - Everything that is needed to understand a data access method is in the method