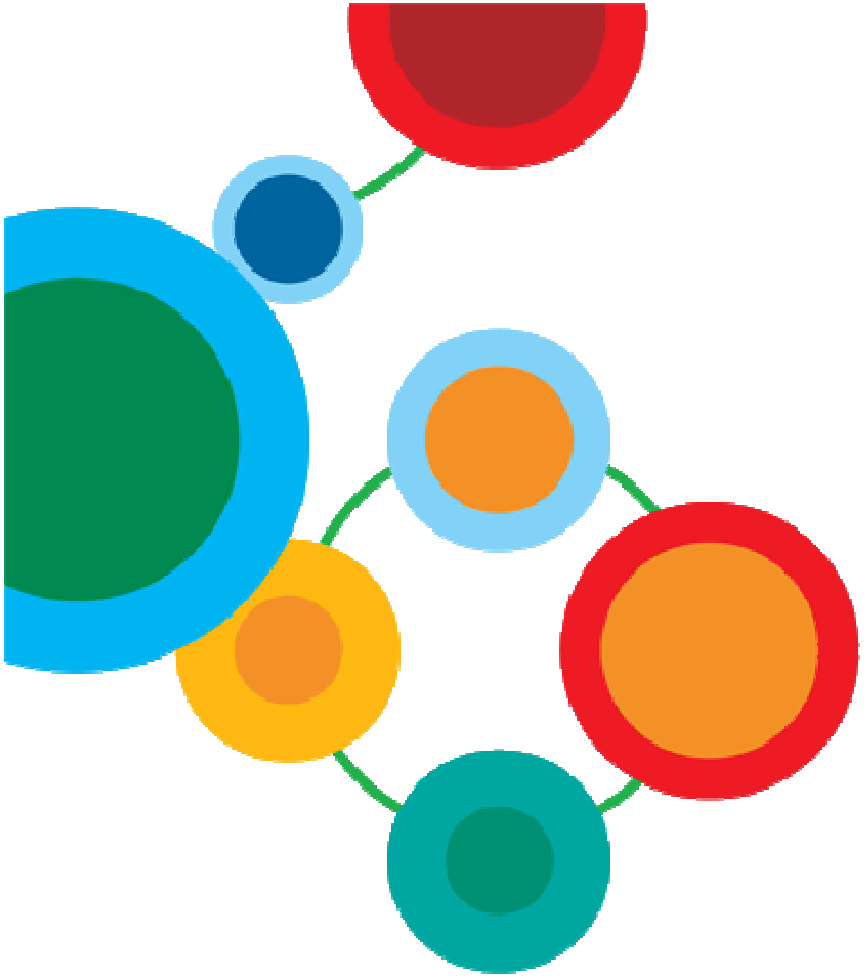

Managing application performance using pureQuery

Session Number 2222

Leif Pedersen, IBM
Patrick Titzler, IBM

IBM Software
Information On Demand 2011

Please Note:

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Why customers have chosen pureQuery

- Better performance
  - Stable/predictable query performance
  - Reduced CPU consumption or improved throughput
- Greater control
  - Improved insights result in more efficient processes
  - Tighter security
- Simplified problem determination and impact analysis
  - Trace SQL to source code
  - Identify dependencies

Managing data access performance using pureQuery

Capture

- Capture performance and application metadata

Review and tune SQL

- Analyze coded/generated SQL

Enable static SQL execution

- Lock-in access plan

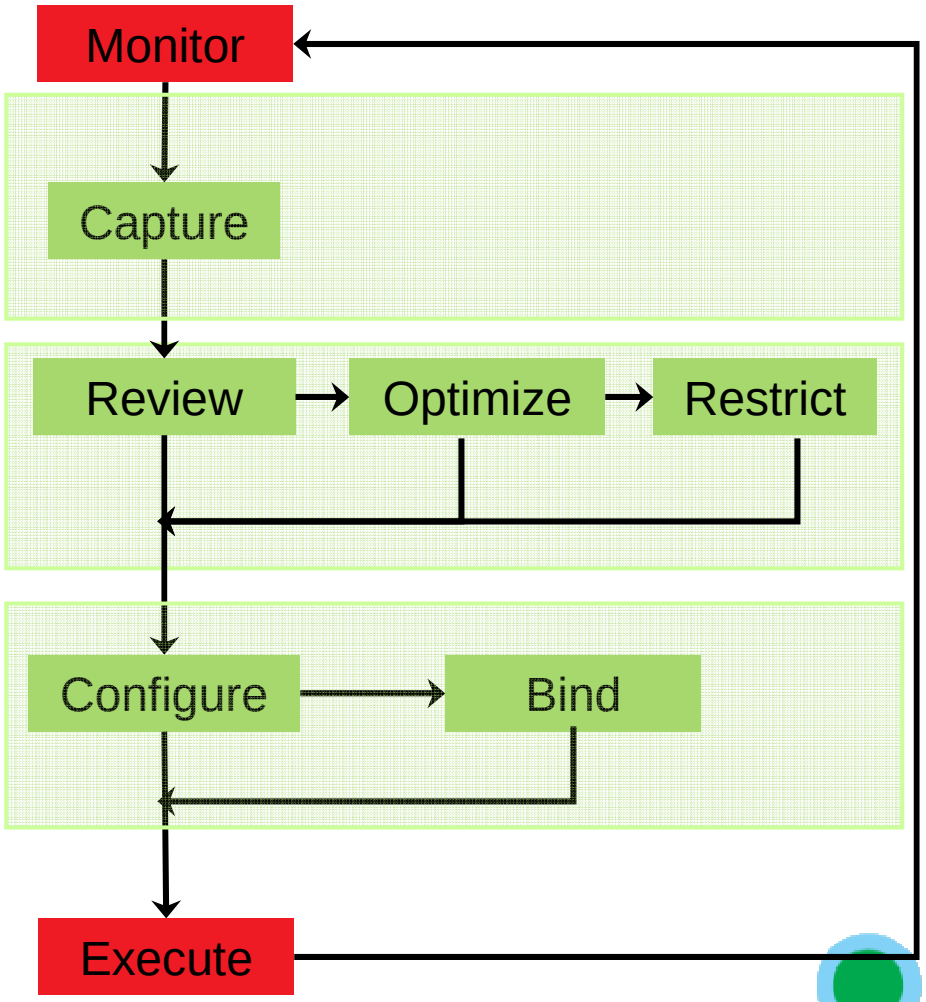

Recap: How does pureQuery work for existing applications?

Capture → Configure → Bind → Execute

SQL | SQL/package mapping | package/DBRM

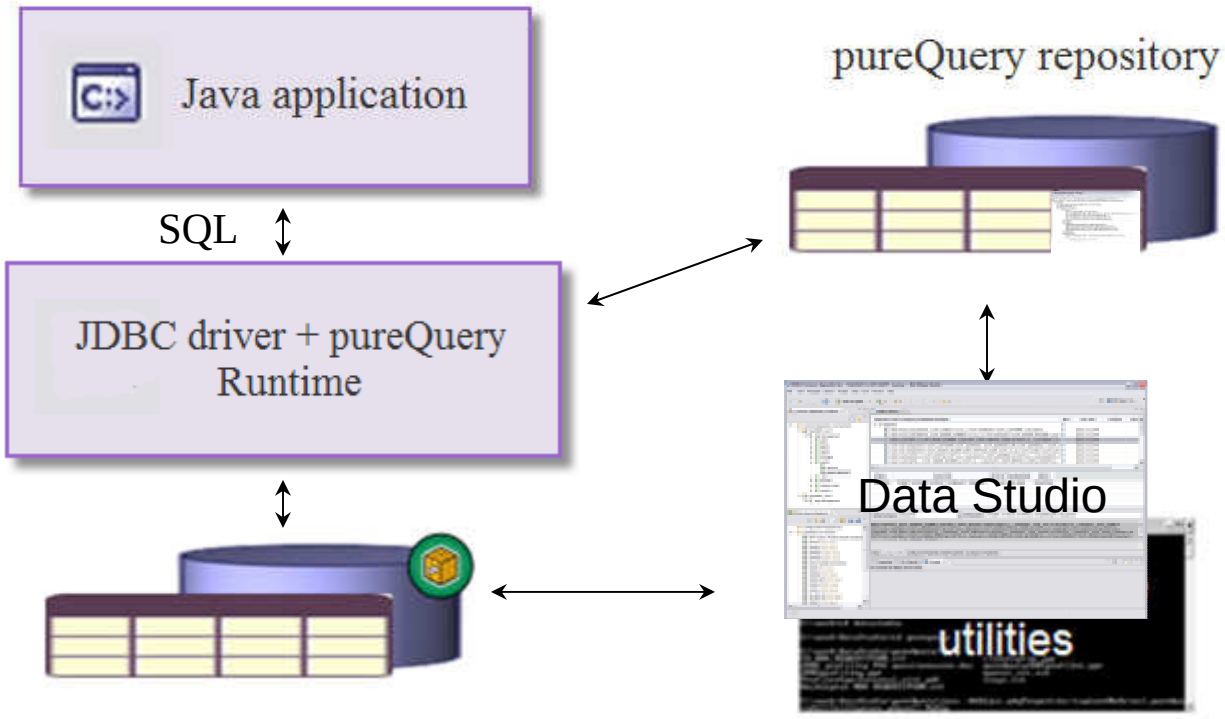

Approaches to capturing pureQuery metadata

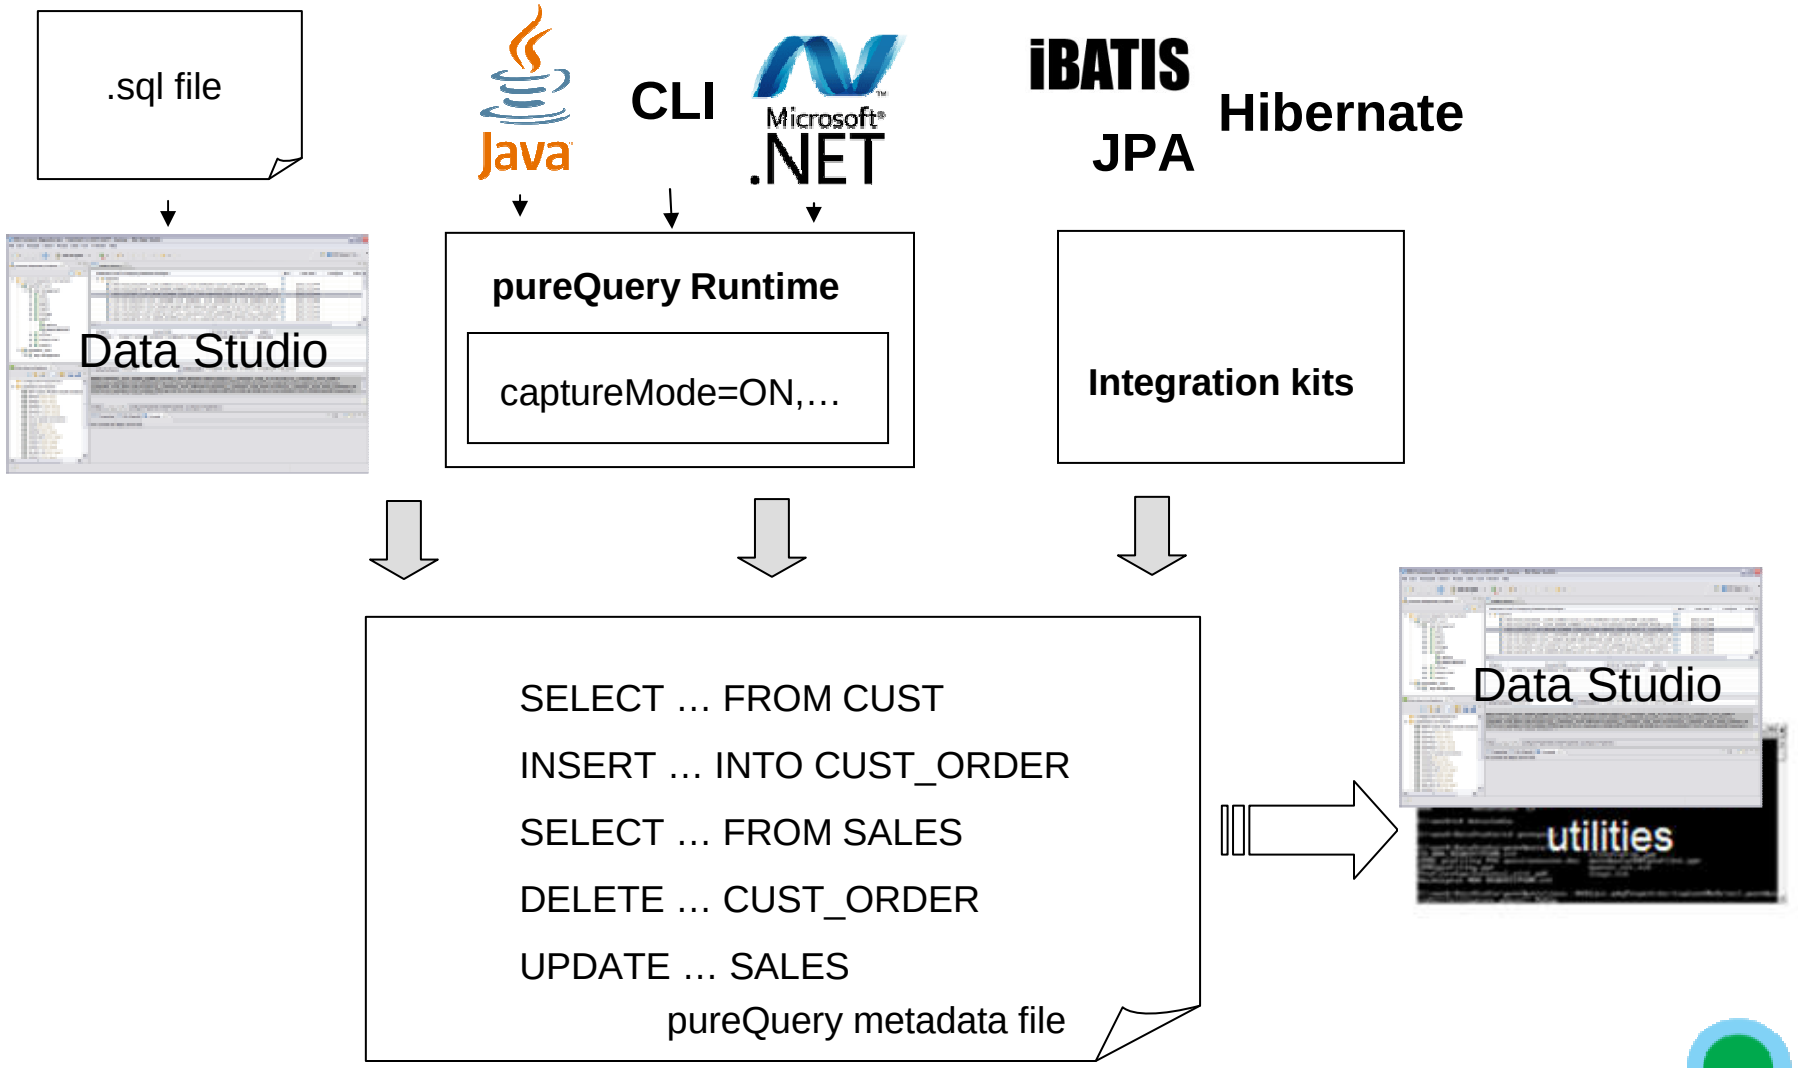

A typical pureQuery deployment topology - before

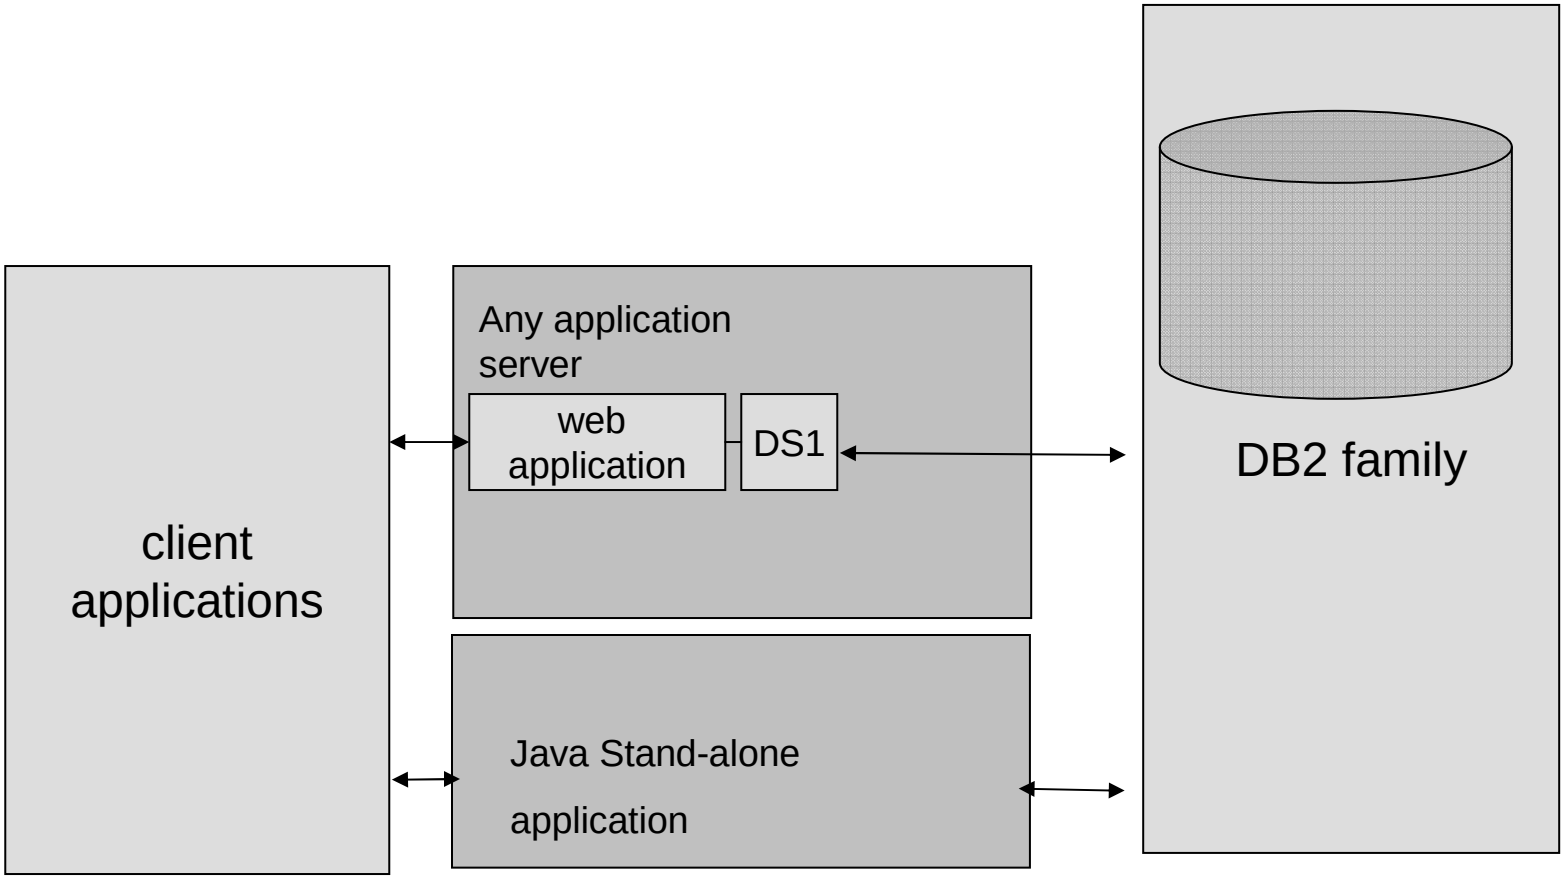

A typical pureQuery deployment topology - after

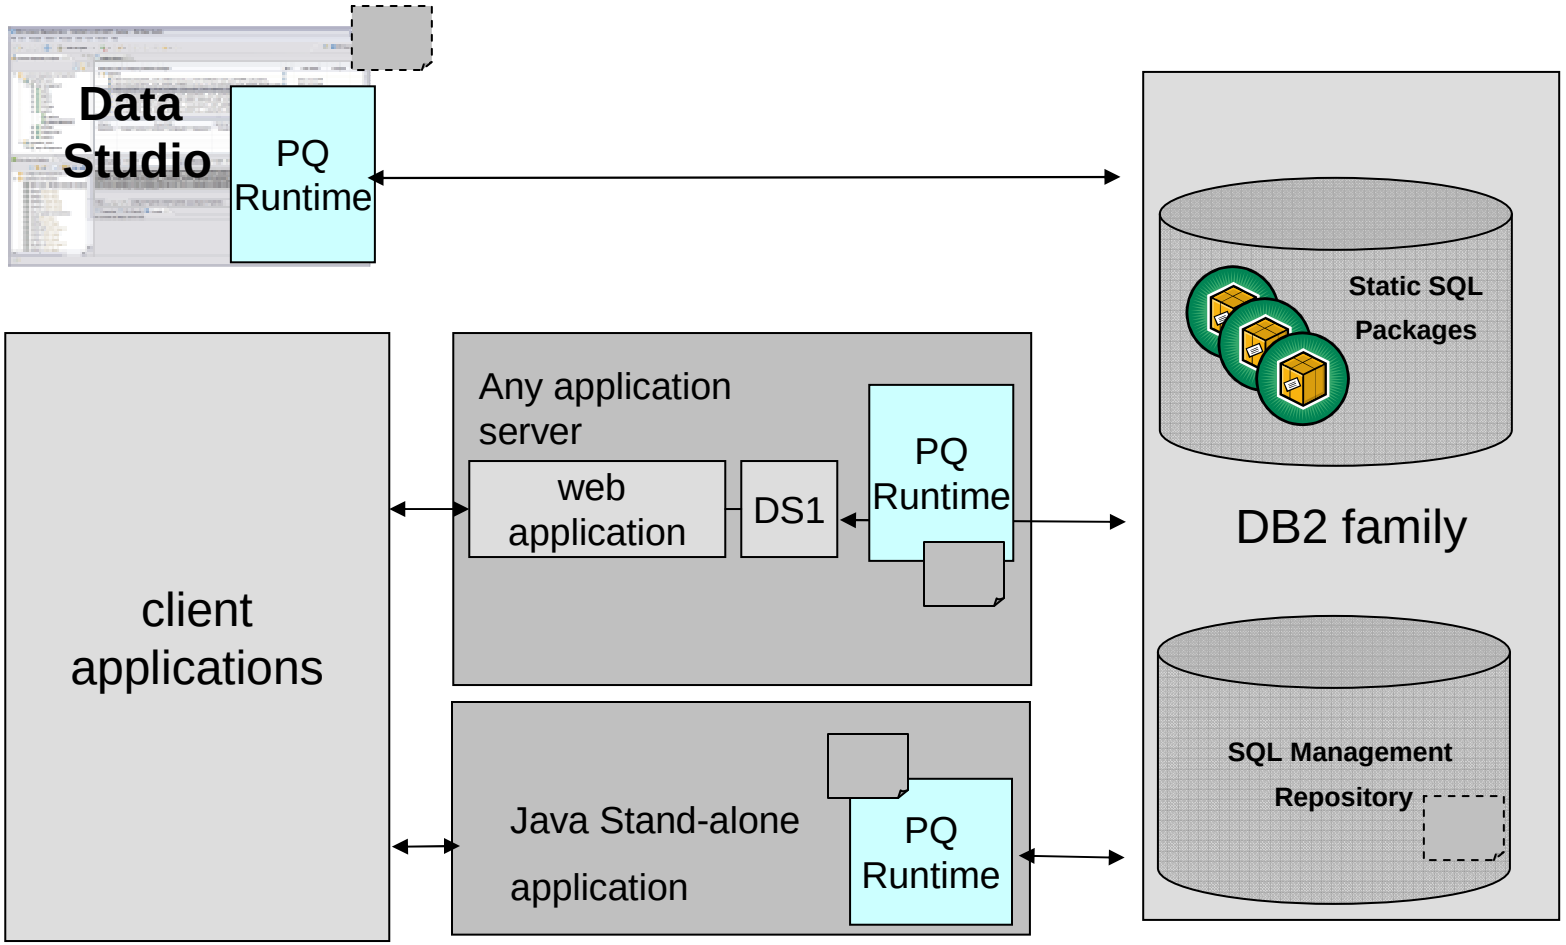

pureQuery deployment – installation and setup

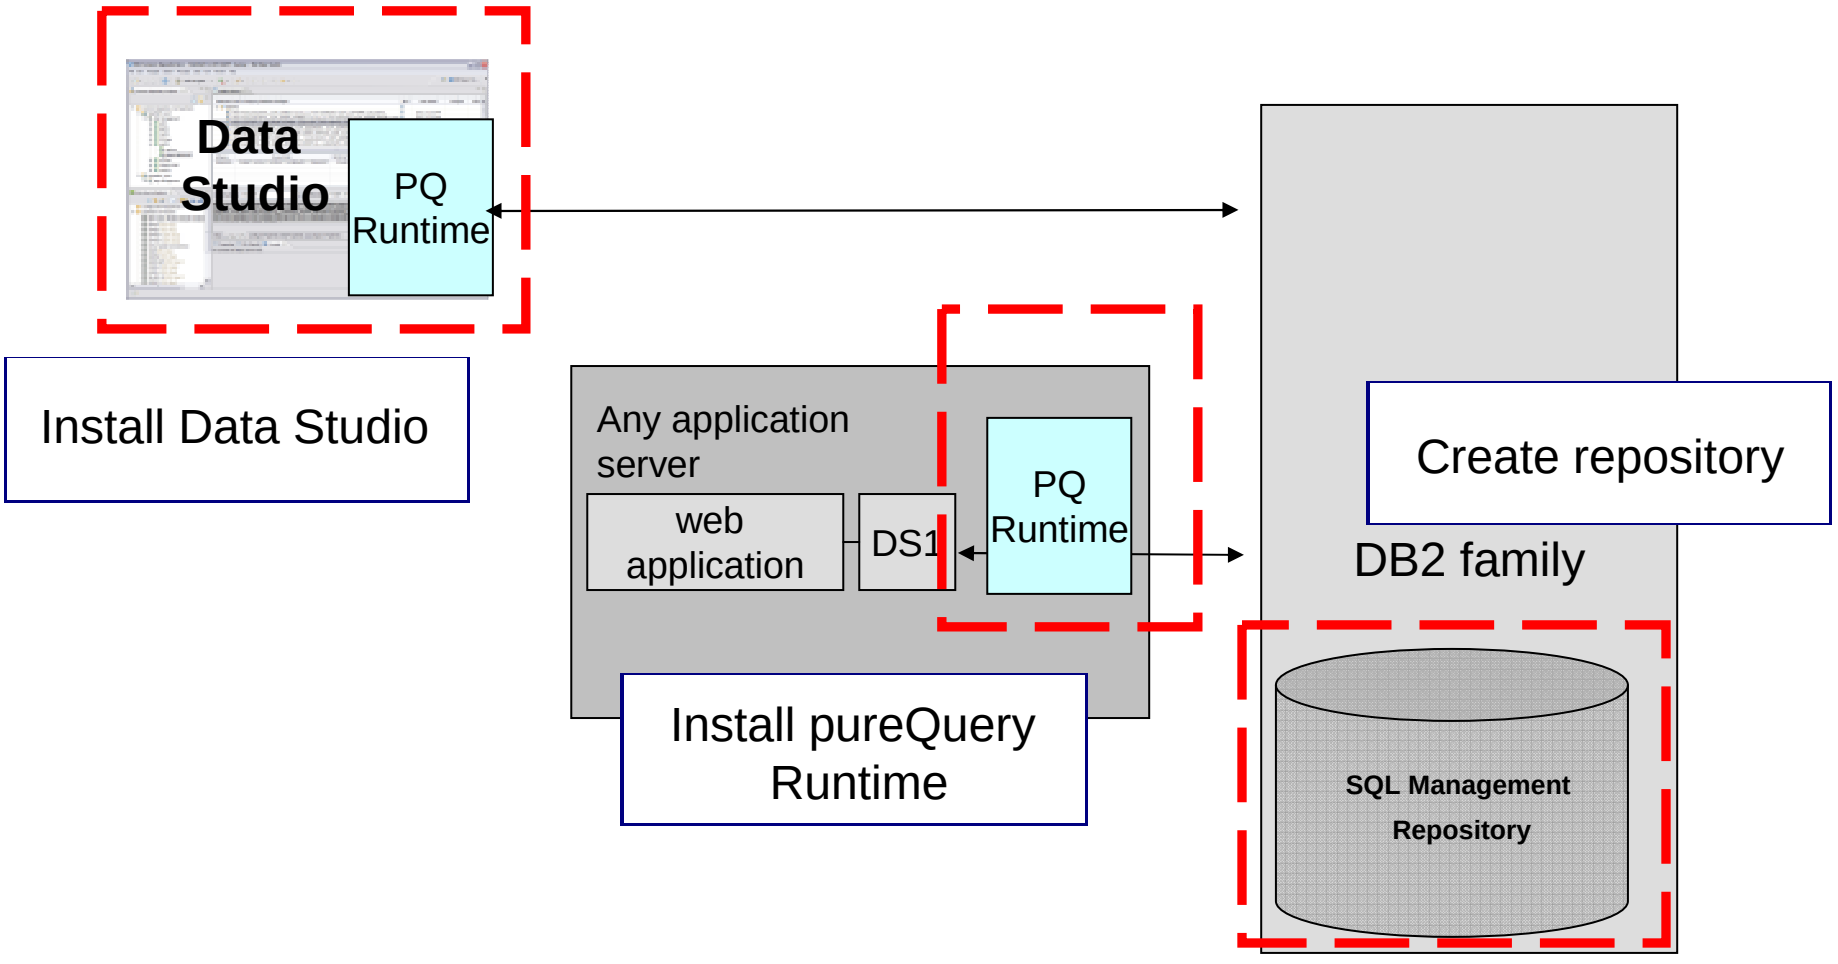

pureQuery deployment – application setup

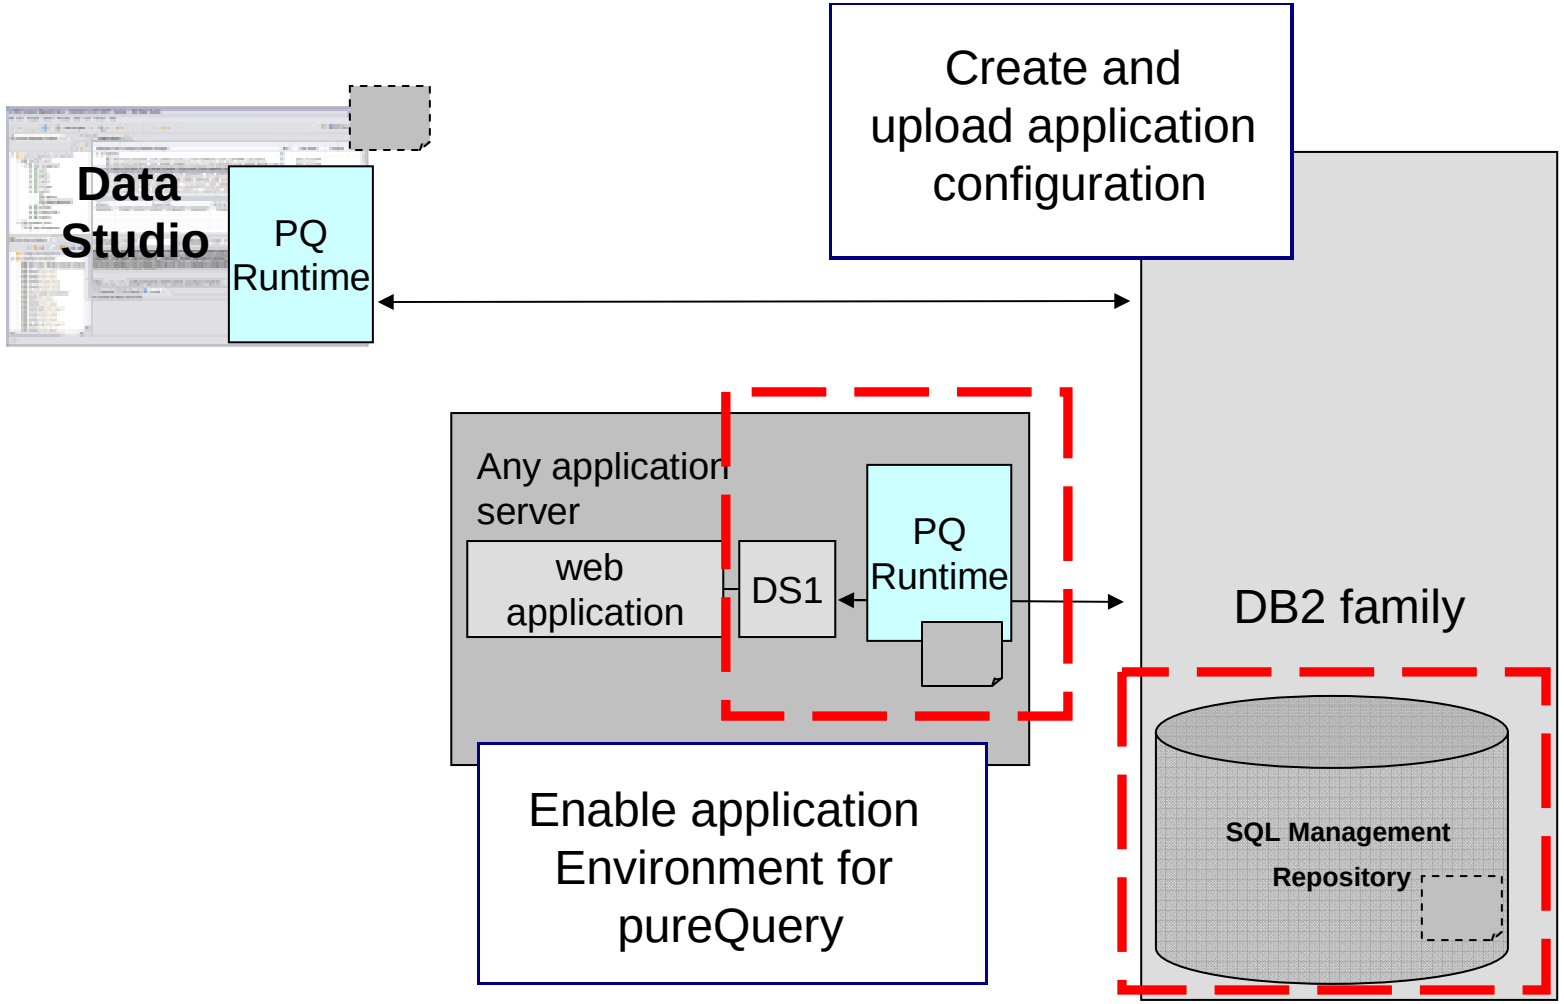

Client tier

Application tier

Database tier

pureQuery deployment – defining the application configuration

- An application configuration (“runtime group”) consists of
  - application metadata (SQL, source code correlation information, …)
  - configuration files
- Application configurations are stored in the SQL Management repository
- Data Studio provides the repository administration interface

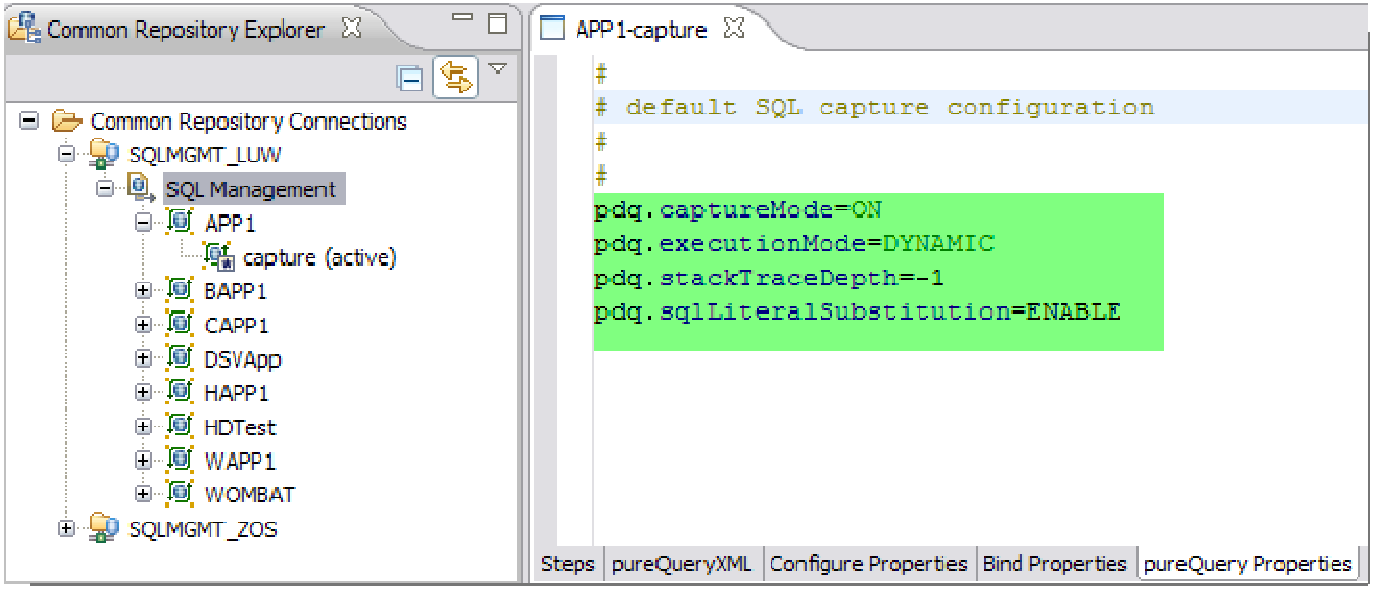

pureQuery deployment - configuring the application environment

- Most common approach:
  - Create pureQuery-enabled JDBC provider
  - Define SQL Management repository connection properties

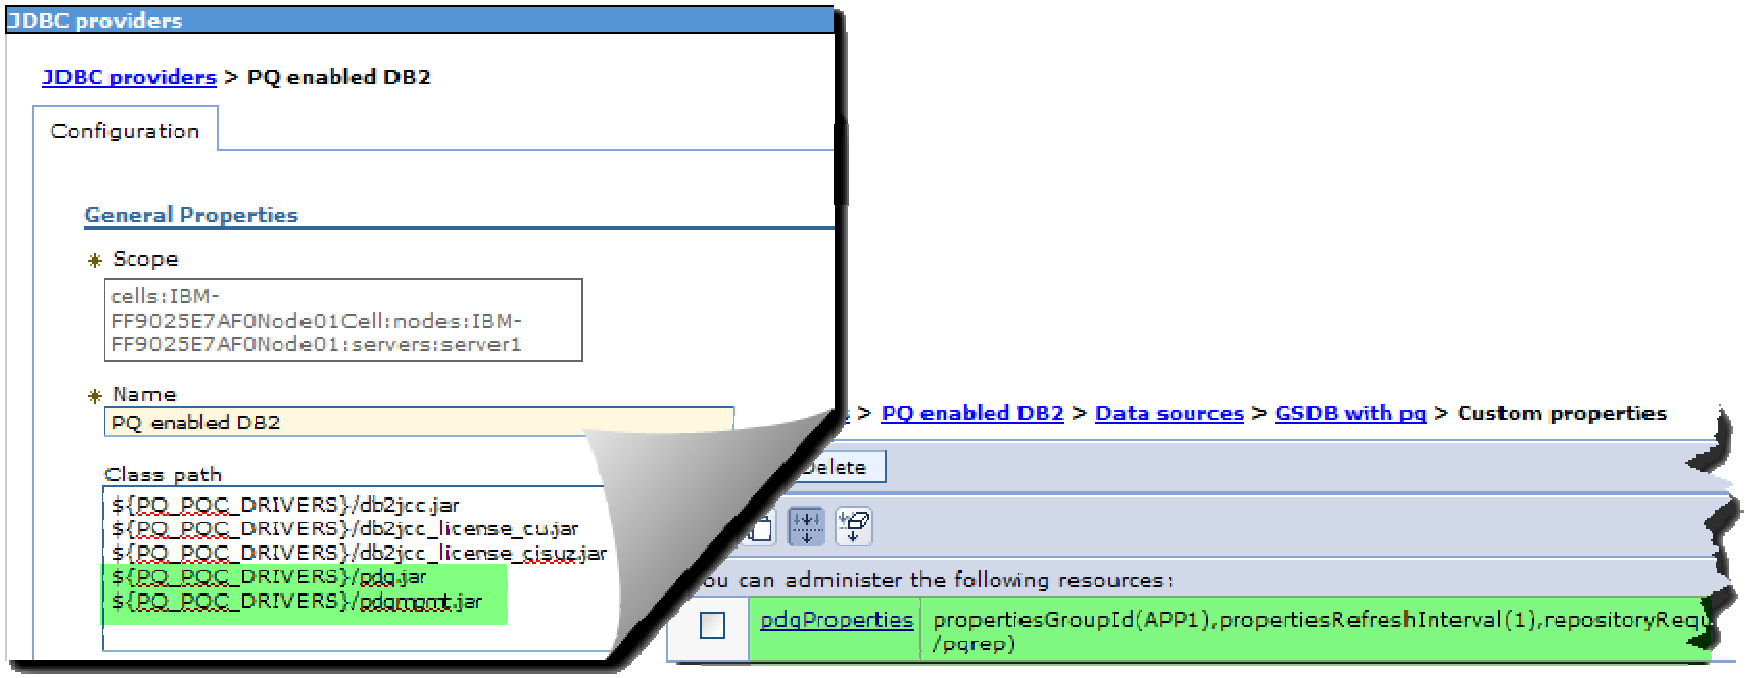

Completing the pureQuery client-optimization process

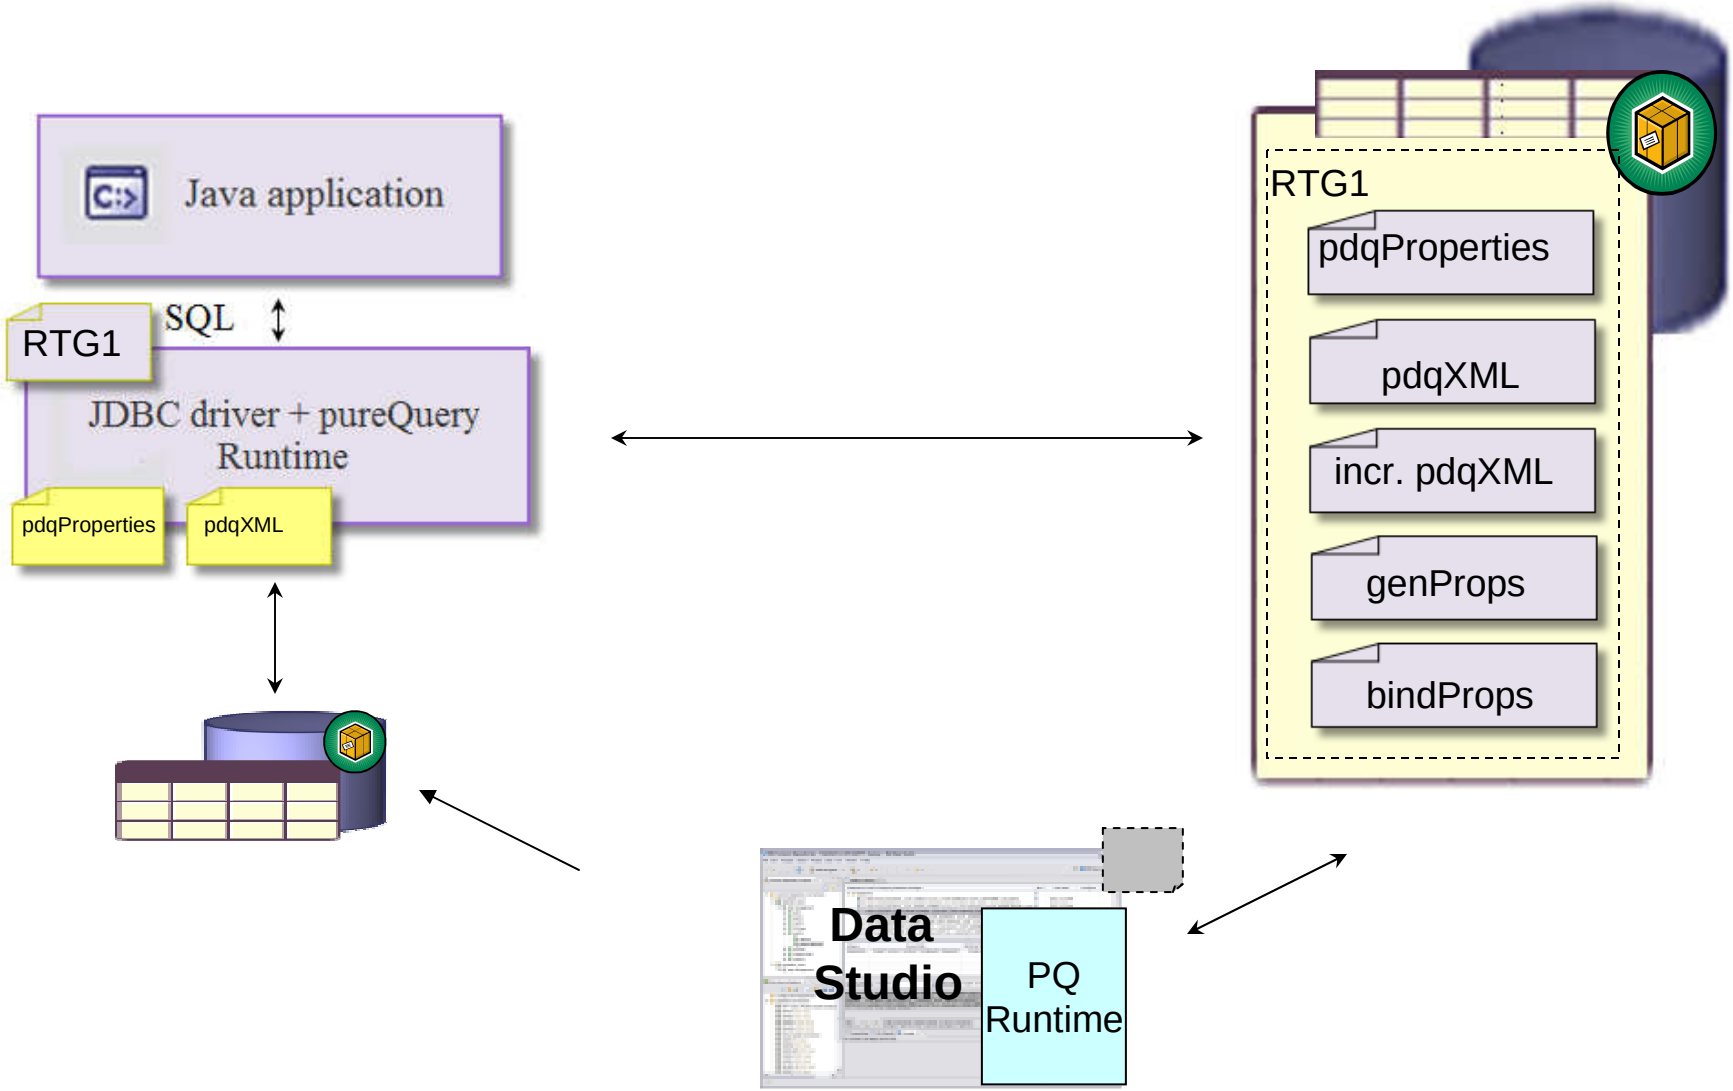

It's the insights!

- Application insights streamline common development tasks

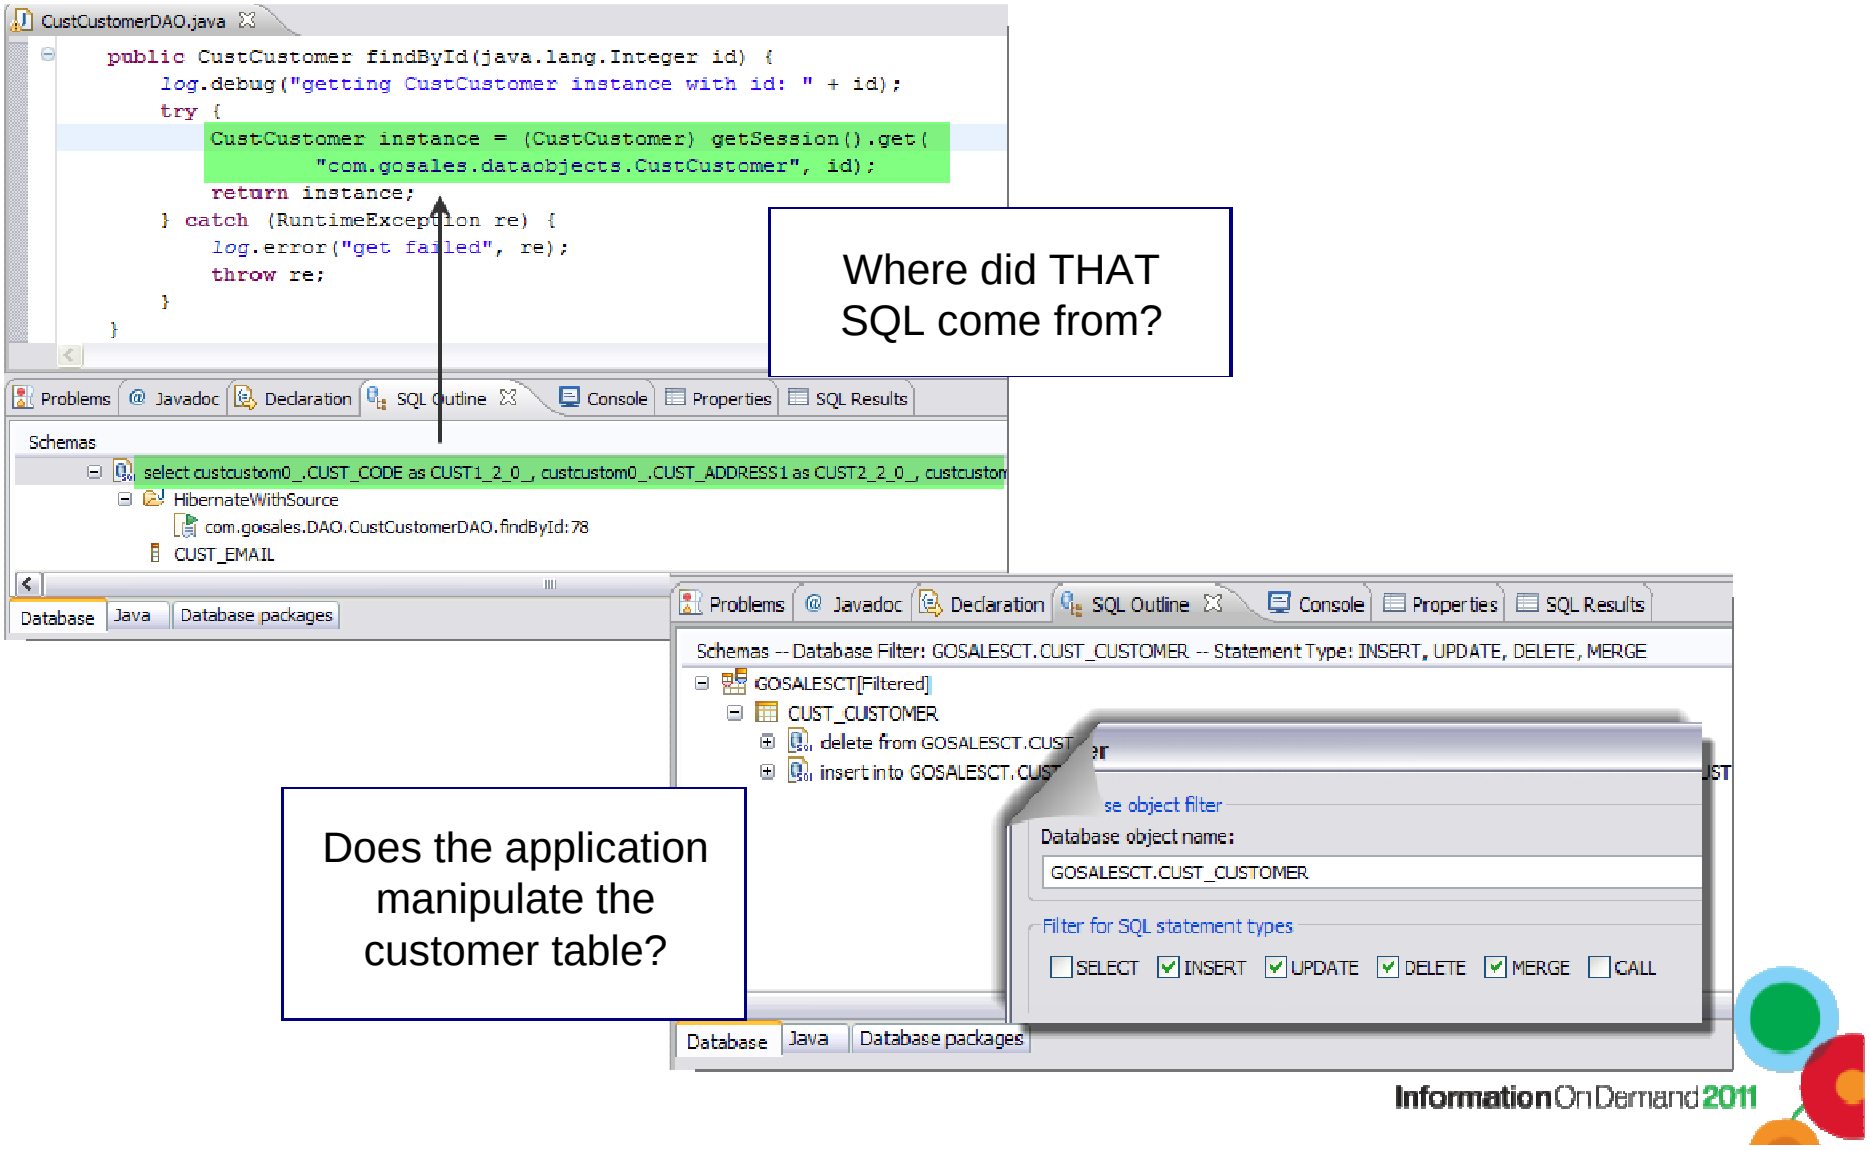

Managing the pureQuery application life-cycle

- Applications are typically deployed in multiple environments:

  development → QA → pre-production → production

- Use multiple SQL Management repositories

- Most common approach: one repository per environment

- Migrate between repositories using export/import

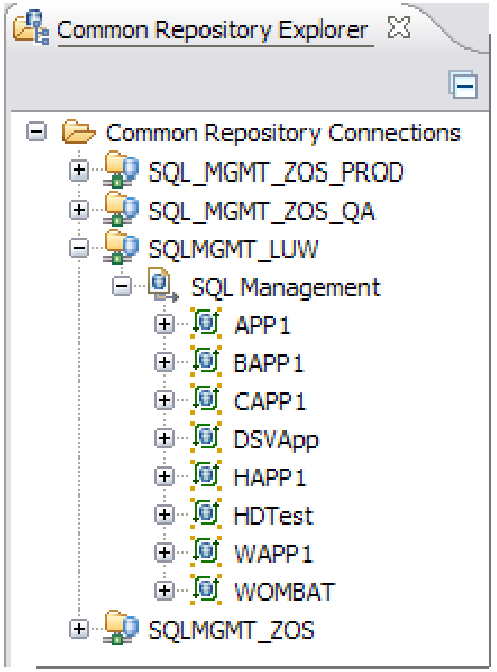

Managing the pureQuery application life-cycle

- Applications change over time: pureQuery metadata and packages map to application versions

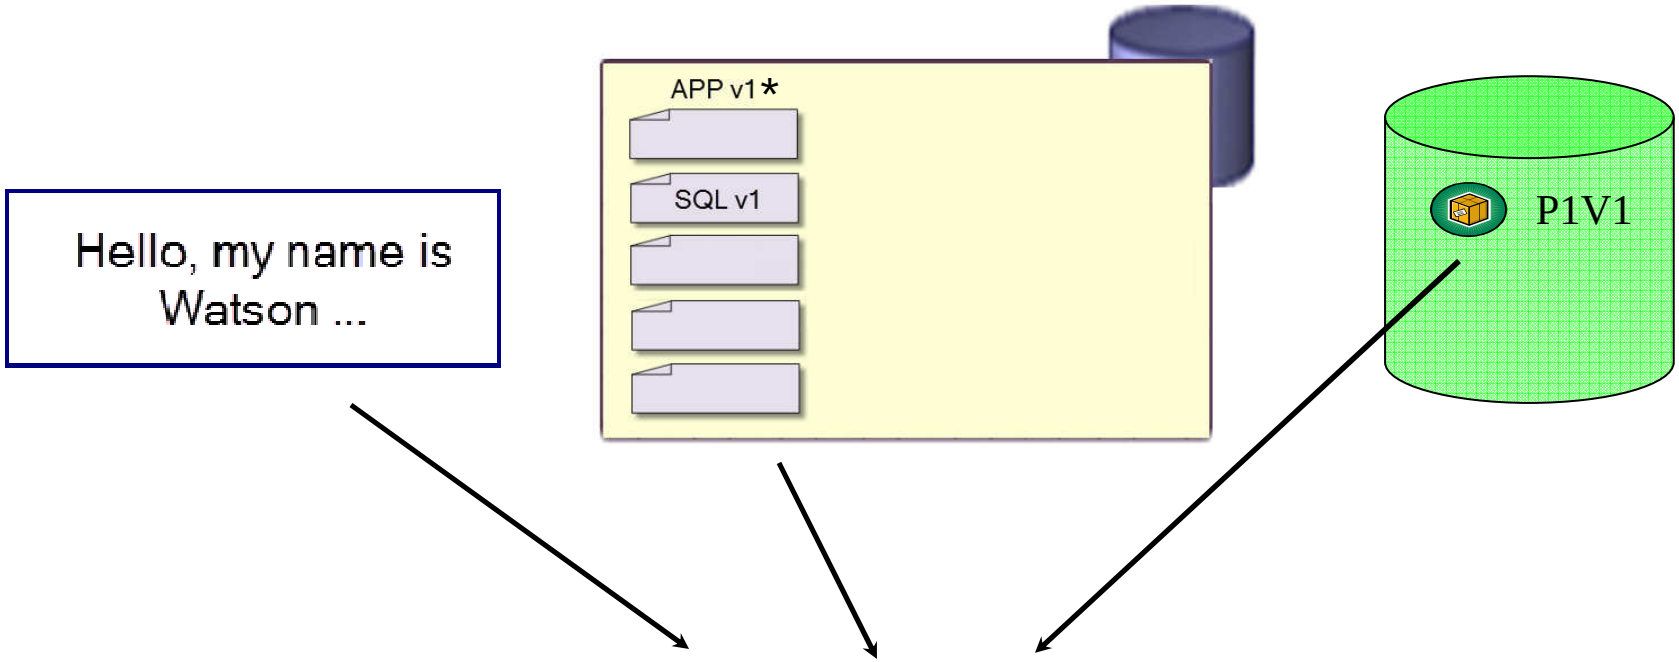

Managing the pureQuery application life-cycle

- Switch to another version by activating it

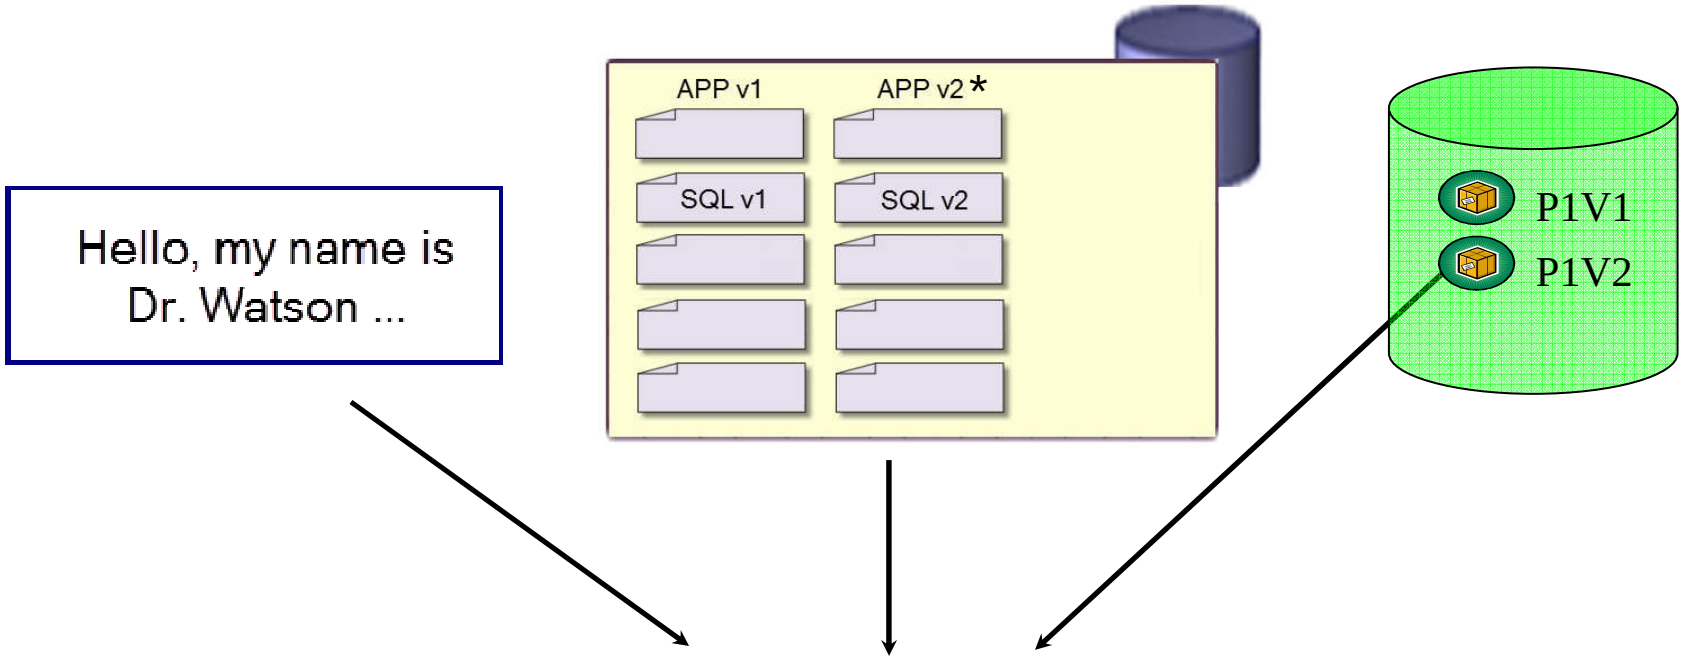

What's new in version 3.1?

- New efficient capture mode: NEW_STMTS
- Enhanced SQL literal replacement support
- Easier portability if SQL contains hard-coded schema qualifiers are used
- Remove invalid SQL statements from capture file (e.g. referenced table was dropped)
- Metadata generation:
  - Perform Literal Substitution during generation
  - Use Optim Query Workload Tuner's exported XML as input
- Native pureQuery implementation: nested bean support
- … and much, much more

Thank You!
Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
  - Your smart phone or web browser at: iodsmartsite.com
  - Any SmartSite kiosk onsite
  - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech

Communities

- **On-line communities, User Groups, Technical Forums, Blogs, Social networks, and more**
  - Find the community that interests you…
    - **Information Management** **ibm.com**/software/data/community
    - **Business Analytics** **ibm.com**/software/analytics/community
    - **Enterprise Content Management** **ibm.com**/software/data/content-management/usernet.html

- **IBM Champions**
  - Recognizing individuals who have made the most outstanding contributions to Information Management, Business Analytics, and Enterprise Content Management communities
    - **ibm.com**/champion

Acknowledgements and Disclaimers:

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. 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. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

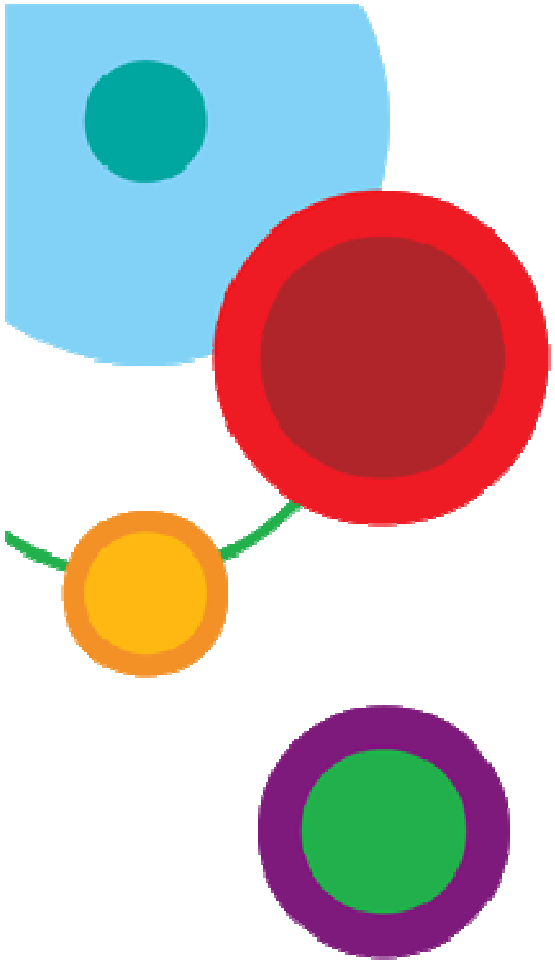

Backup

What information is stored in the database repository?