10 Ways to Destroy Your Object-Relational Mapping Project

Craig L. Russell

JAX India 2009

Presenter's "Baggage"

Java™ Data Objects Specification Lead

Apache OpenJPA PMC Chair

Java EE Container Managed Persistence Architect

#10 Ignore Impedance Mismatch

- Objects != Relational DBMS
- Data Representation
  - Primitive Types
  - Multi-valued Types
- Inheritance
- Polymorphism
- Behavior

#9 Stick to Standards

- Standards == Least Common Denominator
- Not All Implementations Are Equal
  - Performance
  - Support
  - Ease of Use
- Some Vendors Have Cool Features
  - that are not standard,
  - but will improve your project

#8 Rely on Industry-Standard Benchmarks

- Benchmarks != Your Application
- Benchmark Results are Not Indicative of:
  - What to expect from Your Own Efforts
  - How Your Application Will Perform

#7 Ignore the Database
Concentrate on the Object Model

- The Database is The Truth
  - For Everyone Else
- The Object Model is Only
  - The transient representation of The Truth

#6 Eager Fetch Nothing

Bring the Data in Only When Needed

- Different Use-Cases Need Different Data
- Consider the 80/20 Rule:
  - If 80% of use-cases need it, eagerly fetch it
  - Or, Use Custom Fetch-Plans

#5 Map All Foreign Keys to Bidirectional Relationships

- Database-to-Class Tools Assume:
  - Foreign Key == Relationship
- But Some Foreign Keys Are Just Constraints
- Know and Understand the Difference

#4 Map Tables and Classes 1:1

- Tables Can Naturally Represent an Entire Class Hierarchy (subclasses)
- Classes Can Be Stored in Multiple Tables (security issues or infrequent usage)

#3 Tell Your DBA

What the Schema Is

- DBA's Are Stakeholders in Your Schema
- Consult, Ask, and Work With Your DBA

#2 Test Performance
Just Before Production

- Performance is a Critical Non-Functional Criterion
- Performance Can Sink Usability
- Measure Performance Early!

Normalized?

Non-Normalized?

#1 Normalize Everything

- Third Normal Form:
  - A Primary Key Identifies Each Row
  - All Related Data is in the Same Row
  - Slots Contain Only Primitive Values
  - No Aggregate Data (e.g. sum, min, max)
- But Variations on Third Normal are NORMAL

Questions?

Answers!

Craig L. Russell
OpenJPA PMC Chair
ASF