

A Real-World Comparison of the NoPlsql and Thick Database Paradigms

ORACLE WHITEPAPER | November 2016





Abstract

The NoPlsql paradigm starts with an object oriented domain model in the middle tier. This leads naturally to treating the database as a bag of tables, so that primitive SQL statements express its API. The Thick Database paradigm starts with the design of the relational model. This leads naturally to allowing client-side code only to call PL/SQL subprograms, so that all SQL is done from PL/SQL. These two paradigms are mutually incompatible.

In the first part of this paper, we show, based on real-world use cases, how the NoPlsql paradigm brings problems with correctness, security, and performance, and we note that projects that adopt the NoPlsql paradigm use most of their resources on developing plumbing code rather than code that implements actual business functionality. These problems are just the tip of the iceberg and we explain how the NoPlsql paradigm also brings maintenance nightmares.

In the second part, we turn to the Thick Database paradigm, rehearse the reasons for its superiority, and explain that it implies more than just allowing database calls to invoke only PL/SQL subprograms. We formalize a layered code classification scheme which leads to optimal understandability and maintainability of both your PL/SQL and your SQL code—which properties bring the maximum probability of correct behavior. We will convince participants that adopting the Thick Database Paradigm guarantees the avoidance of the awful problems brought by the NoPlsql paradigm and establishes the database tier as a reliable central service provider in the application landscape. We will debunk the myth that “moving business logic out of the data tier and into the middle tier” improves scalability.

The paper draws upon the empirical findings acquired over many years of extensive hands-on experience of building, and tuning, applications that use Oracle Database. We have worked with applications that use the Thick Database paradigm, and with ones that use the NoPlsql paradigm. So we mean what we say by “real-world comparison”. We also discuss the results of a very carefully designed and conducted apples-to-apples experiment to compare the performance characteristics of the two paradigms. We show that the Thick Database paradigm is two-to-three orders of magnitude better than the NoPlsql paradigm.



To be posted on the Oracle Database PL/SQL and EBR blog

We are still working on the whitepaper. As soon as we can, we'll upload it to the blog's resources section and link to it from this post:

» [NoPlsql versus ThickDB¹](#)

When we do, we'll each send out a tweet. So do make sure that you follow us!

Notice that the post already has a link to a YouTube video, recorded at Oracle HQ, where Toon presents a detailed account of the performance experiment that this paper's abstract refers to.

Toon Koppelaars,
Consulting Member of Technical Staff, Real World Performance Group,
Database Server Technologies Division, Oracle Netherlands
toon.koppelaars@oracle.com
Twitter: [@ToonKoppelaars](#)

Bryn Llewellyn,
Distinguished Product Manager,
Database Server Technologies Division, Oracle Headquarters
bryn.llewellyn@oracle.com
Twitter: [@BrynLite](#)

1. blogs.oracle.com/plsql-and-ebr/entry/noplsq_vs_thickdb



Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Hardware and Software, Engineered to Work Together

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 1014

A Real-World Comparison of the NoPsql and Thick Database Paradigms
November 2016
Authors: Toon Koppelaars and Bryn Llewellyn