

# Architecture of Oracle Parallel Server

R.Bamford, D.Butler, B.Klots, N.Macnaughton

rbamford@us.oracle.com dobutler@us.oracle.com bklots@us.oracle.com nmacnaug@us.oracle.com

Oracle Corp.

500 Oracle Parkway

Redwood Shores CA 94065

## Abstract

*Oracle Parallel Server (OPS) is a shared disk RDBMS. We present a high level overview of the main architectural issues of OPS and their evolution throughout the releases of the Oracle database.*

The Oracle relational database can be used either in exclusive or shared mode. In exclusive (the most traditional) mode only one Oracle instance can mount or open the database. The Oracle Parallel Server (OPS) option allows a database to be opened in shared mode. In shared mode multiple instances of Oracle can mount the same database. This allows concurrent direct read and write access by multiple users from multiple instances to all the data in the database. This means that multiple computers can work with the same database. It is this approach which ensures that Oracle holds world records in the TPC benchmarks (more than 100,000 tpmC).

---

*Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.*

**Proceedings of the 24th VLDB Conference  
New York, USA, 1998**

Oracle Parallel Server runs on a shared disk cluster. The shared disk architecture employed by Oracle has a number of distinct features:

- high availability: when one node fails, other instances can proceed with all the data still available to them.
- dynamic partitioning/load balancing
- aggregate CPU and memory resources of the cluster are utilized
- enhanced throughput (see the TPCC results above).

These advantages of OPS do not come for free. The challenge of this architecture (as of any clustered or distributed architecture) is to provide data coherency for the independent users of the system. Oracle does that using sophisticated locking schemes. Oracle uses multiple level locking: row locks on transaction levels, instance locks within instances, and global locks among the instances. The latter are specific to Oracle Parallel Server. The basic problems associated with the locking architecture are:

- lock modes and their compatibility,
- locking granularity,
- binding data to locks,
- interaction between data access and locking protocols.

Another important OPS architecture issue is maintaining a global (cluster wide) logical clock.

The global coherency protocol coupled with shared disk access, provides high availability feature of Oracle Parallel Server: if an instance dies the data changed by the dead instance are recovered and after that other instances can continue their work. We will give a brief

overview of OPS failure detection and subsequent recovery.

Oracle Parallel Server works on both clusters and MPP's. It supports both on line transaction processing and data warehouse workloads.

Oracle's Parallel Query feature is expanded to run across multiple CPUs and multiple nodes. Oracle Parallel Query architecture in conjunction with OPS provides a very successful hybrid of shared disk and shared nothing software functionality.

The latest version of Oracle Parallel Server (8.0) contains a significant number of important optimizations. The Distributed Lock Manager was integrated and the OPS interfaces were considerably optimized and changed. The new OPS employs the latest developments in the inter process communication and system area network technology. We will present an overview of those changes and the underlying motivations for them.