

Oracle Database 21c Technical Architecture



Copyright © 2021, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

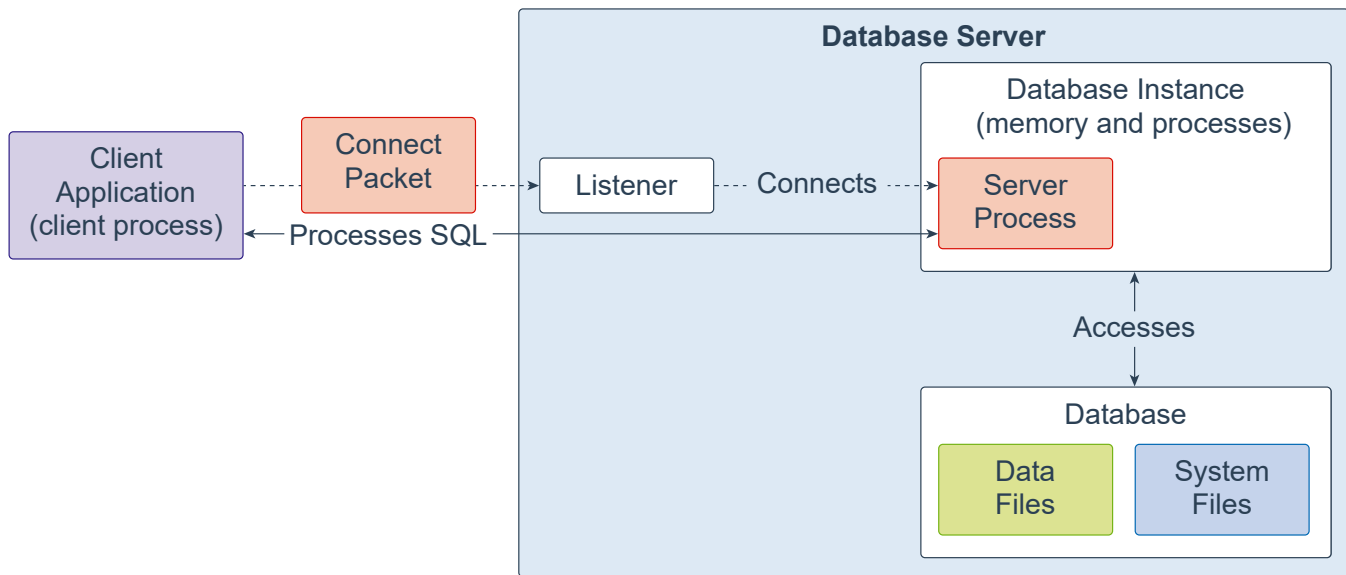
This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Database Server



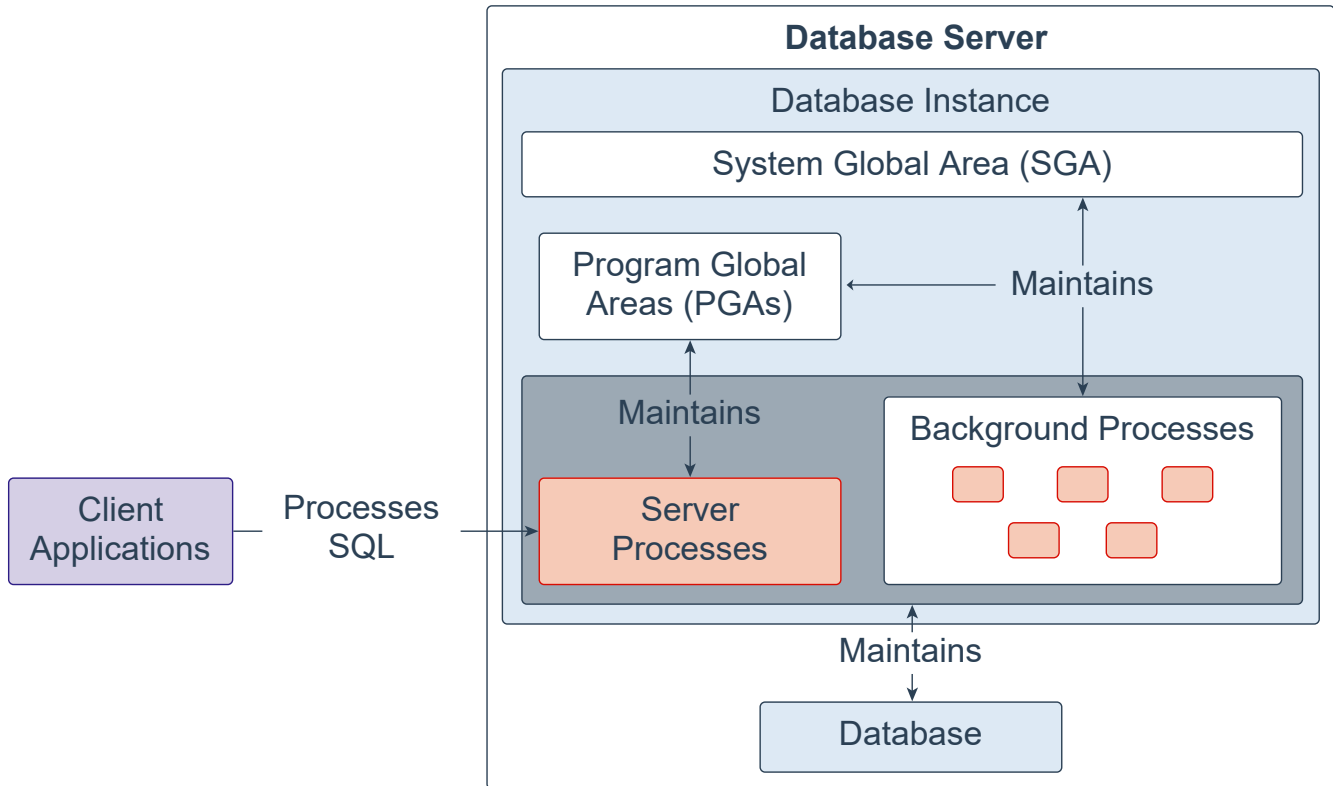
An Oracle Database consists of at least one database instance and one database. The database instance handles memory and processes. The multitenant container database consists of physical files called data files. An Oracle Database also uses several database system files during its operation.

A single-instance database architecture consists of one database instance and one database. A one-to-one relationship exists between the database and the database instance. Multiple single-instance databases can be installed on the same server machine. There are separate database instances for each database. This configuration is useful to run different versions of Oracle Database on the same machine.

An Oracle Real Application Clusters (Oracle RAC) database architecture consists of multiple instances that run on separate server machines. All of them share the same database. The cluster of server machines appear as a single server on one end, and end users and applications on the other end. This configuration is designed for high availability, scalability, and high-end performance.

The listener is a database server process. It receives client requests, establishes a connection to the database instance, and then hands over the client connection to the server process. The listener can run locally on the database server or run remotely. Typical Oracle RAC environments are run remotely.

Database Instance

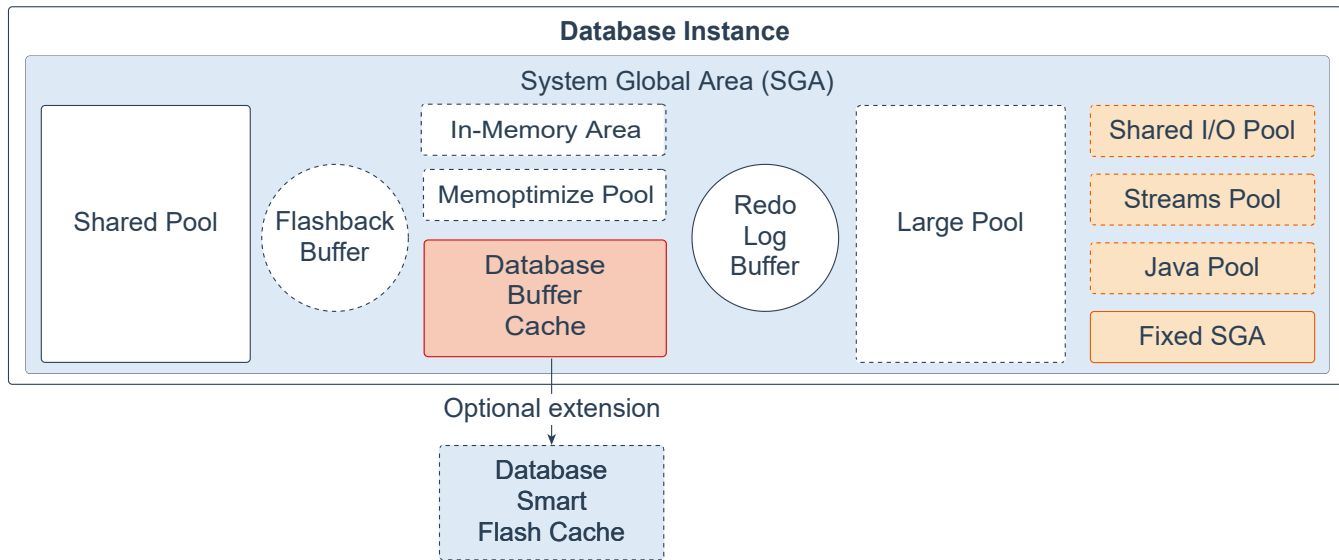


A database instance contains a set of Oracle Database background processes and memory structures. The main memory structures are the System Global Area (SGA) and the Program Global Areas (PGAs). The background processes operate on the stored data (data files) in the database and use the memory structures to do their work. A database instance exists only in memory.

Oracle Database also creates server processes to handle the connections to the database on behalf of client programs, and to perform the work for the client programs; for example, parsing and running SQL statements, and retrieving and returning results to the client programs. These types of server processes are also referred to as foreground processes.

For more information, see [Oracle Database Instance](#).

System Global Area



The System Global Area (SGA) is the memory area that contains data and control information for one Oracle Database instance. All server and background processes share the SGA. When you start a database instance, the amount of memory allocated for the SGA is displayed. The SGA includes the following data structures:

- **Shared pool:** Caches various constructs that can be shared among users; for example, the shared pool stores parsed SQL, PL/SQL code, system parameters, and data dictionary information. The shared pool is involved in almost every operation that occurs in the database. For example, if a user executes a SQL statement, then Oracle Database accesses the shared pool.
- **Flashback buffer:** Is an optional component in the SGA. When Flashback Database is enabled, the background process called Recovery Writer Process (RVWR) is started. RVWR periodically copies modified blocks from the buffer cache to the flashback buffer, and sequentially writes Flashback Database data from the flashback buffer to the Flashback Database logs, which are circularly reused.
- **Database buffer cache:** Is the memory area that stores copies of data blocks read from data files. A buffer is a main memory address in which the buffer manager temporarily caches a currently or recently used data block. All users concurrently connected to a database instance share access to the buffer cache.
- **Database Smart Flash cache:** Is an optional memory extension of the database buffer cache for databases running on Solaris or Oracle Linux. It provides a level 2 cache for database blocks. It can improve response time and overall throughput for both read-intensive online transaction processing (OLTP) workloads and ad-hoc queries and bulk data modifications in a data warehouse (DW) environment. Database Smart Flash Cache resides on one or more flash disk devices, which are solid state storage devices that use flash memory. Database Smart Flash Cache is typically more economical than additional main memory, and is an order of magnitude faster than disk drives.
- **Redo log buffer:** Is a circular buffer in the SGA that holds information about changes made to the database. This information is stored in redo entries. Redo entries contain the information necessary to reconstruct (or redo) changes that are made to the database by data manipulation language (DML), data definition language (DDL), or internal operations. Redo entries are used for database recovery if necessary.
- **Large pool:** Is an optional memory area intended for memory allocations that are larger than is appropriate for the shared pool. The large pool can provide large memory allocations for the User Global Area (UGA)

for the shared server and the Oracle XA interface (used where transactions interact with multiple databases), message buffers used in the parallel execution of statements, buffers for Recovery Manager (RMAN) I/O workers, and deferred inserts.

- **In-Memory Area:** Is an optional component that enables objects (tables, partitions, and other types) to be stored in memory in a new format known as the columnar format. This format enables scans, joins, and aggregates to perform much faster than the traditional on-disk format, thus providing fast reporting and DML performance for both OLTP and DW environments. This feature is particularly useful for analytic applications that operate on a few columns returning many rows rather than for OLTP, which operates on a few rows returning many columns.
- **Memoptimize Pool:** Is an optional component that provides high performance and scalability for key-based queries. The Memoptimize Pool contains two parts, the memoptimize buffer area and the hash index. Fast lookup uses the hash index structure in the memoptimize pool providing fast access to the blocks of a given table (enabled for `MEMOPTIMIZE FOR READ`) permanently pinned in the buffer cache to avoid disk I/O. The buffers in the memoptimize pool are completely separate from the database buffer cache. The hash index is created when the Memoptimized Rowstore is configured, and is maintained automatically by Oracle Database.
- **Shared I/O pool (SecureFiles):** Is used for large I/O operations on SecureFile Large Objects (LOBs). LOBs are a set of data types that are designed to hold large amounts of data. SecureFile is an LOB storage parameter that allows deduplication, encryption, and compression.
- **Streams pool:** Is used by Oracle Streams, Data Pump, and GoldenGate integrated capture and apply processes. The Streams pool stores buffered queue messages, and it provides memory for Oracle Streams capture processes and apply processes. Unless you specifically configure it, the size of the Streams pool starts at zero. The pool size grows dynamically as needed when Oracle Streams is used.
- **Java pool:** Is used for all session-specific Java code and data in the Java Virtual Machine (JVM). Java pool memory is used in different ways, depending on the mode in which Oracle Database is running.
- **Fixed SGA:** Is an internal housekeeping area containing general information about the state of the database and database instance, and information communicated between processes.

For more information, see [Overview of the System Global Area \(SGA\)](#).