

Oracle Sharding 19c

New Features

White Paper / June 24, 2019

Introduction	3
Multiple Shards in a Single Multitenant Database	4
Multiple Table Families	4
Scalable Cross Shard Query Coordinators	5
Middle Tier Sharding.....	6
Summary	6

INTRODUCTION

Oracle Database Sharding enables globally distributed, linearly scalable, multi-model databases. It allows you to support applications that require low latency, high availability, and meet data sovereignty requirements. Oracle Sharding distributes your data across multiple databases either within the same data center or across multiple data centers in a way that provides locality for applications either on-premises or in the cloud. The pool of databases is presented to the application as a single logical database. Applications can elastically scale by adding more databases(shards) to the pool. It is built on shared-nothing architecture, in which data is horizontally partitioned across databases that share no hardware or software. Oracle Sharding does all this while rendering the strong consistency, full power of SQL and the Oracle Database ecosystem.



Figure 1 – Sharded Databases

Oracle Database Sharding has been available since Oracle Database 12c Release 2.

This white paper discusses new features of Oracle Sharding 19c. It is intended for Enterprise Architects, Database Architects, Database Administrators, Application Architects and those who are involved in the design and architecture of distributed database systems.

MULTIPLE SHARDS IN A SINGLE MULTITENANT DATABASE

Sharding a pluggable database (PDB) provides fault isolation from single CDB failures. For example, by sharding a database using PDB across 100 CDBs, the failure or slowdown of a single CDB can affect at most 1% of the data within that PDB. This can also be used to honor data sovereignty.

A shard can be a PDB. This is referred to as PDB-shards. A single container database (CDB) can host multiple PDB-shards from different sharded database deployments. This enables consolidation of sharded databases.

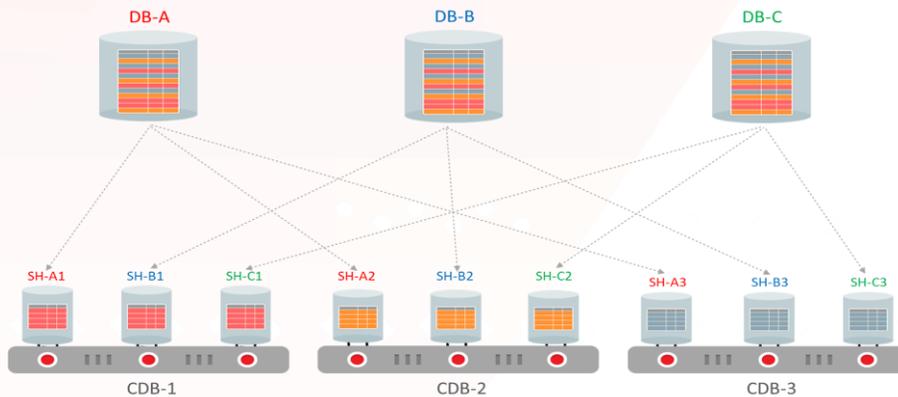


Figure 2 – Sharded Pluggable Databases

In the above example in Figure 2, three logical databases (DB-A, DB-B and DB-C) have been sharded across nine PDB-shards and consolidated in three multitenant databases (CDB-1, CDB-2 and CDB-3).

MULTIPLE TABLE FAMILIES

Using multiple table families on the same sharded database improves system resource utilization and reduces operational cost. A single Oracle Sharding environment can now have multiple table families. A table family is a set of related (through referential integrity) tables sharing the same sharding key. A single sharded database can host tables from different table families which could potentially be used by different applications. However, joins executing joins across table families may take longer as compared to those executed within the same table family. Multiple table families are supported for system-managed sharding (by consistent hash) only.

An example of multiple table families is in Figure 3 below.

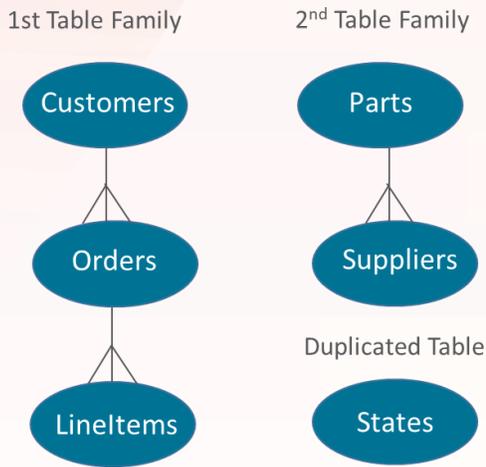


Figure 3 – Multiple Table Families

SCALABLE CROSS SHARD QUERY COORDINATORS

Horizontally scalable cross-shard query coordinators can improve performance and availability of read-intensive cross-shard queries. A Shard Catalog can be protected by one or more Active Data Guard standby databases. The primary and all the read-only standby Shard Catalogs can be used as cross shard query coordinator.

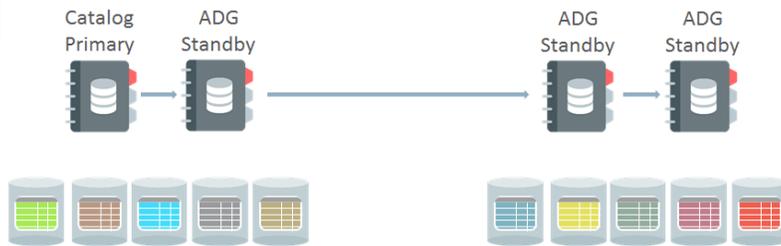


Figure 4 – Scalable cross shard query coordinators

MIDDLE TIER SHARDING

The application middle tier can also be sharded to provide affinity to database sharding. Affinitive grouping of middle tiers with database shards is at times referred to as swim lanes. In such deployments, the application's front end routing tier can call a REST API (provided by a sharded database Middle Tier Routing Service), by-passing the sharding key, to retrieve the swim lane details. The swim lane details include the shard name to help route the request to the appropriate shard. This provides better fault isolation, cache locality, scalability and reduction in database connections used by the middle tier. For geographically distributed shards, it eliminates chatty midtier-to-database connections across datacenters.

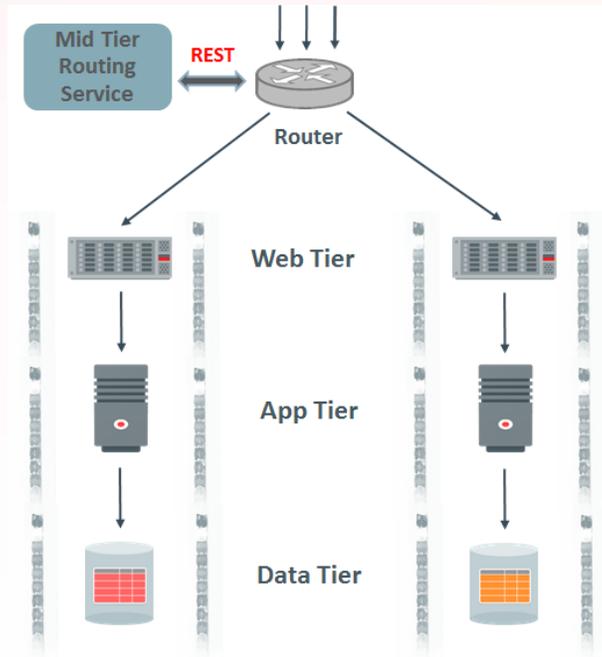


Figure 5 – Middle Tier Sharding with Middle Tier routing Services

SUMMARY

Oracle Databas Sharding enables globally distributed multimodel (relational and document) cloud-native (and on-premises) Database Management System (DBMS). Oracle Database 19c provides additional deployment and scalability options.

References

Oracle Sharding OTN Page – <http://www.oracle.com/goto/oraclesharding>

ORACLE CORPORATION

Worldwide Headquarters

500 Oracle Parkway, Redwood Shores, CA 94065 USA

Worldwide Inquiries

TELE + 1.650.506.7000 + 1.800.ORACLE1

FAX + 1.650.506.7200

oracle.com

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com. Outside North America, find your local office at oracle.com/contact.

 blogs.oracle.com/oracle

 facebook.com/oracle

 twitter.com/oracle

Integrated Cloud Applications & Platform Services

Copyright © 2019, 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. 0619

Oracle Sharding 19c— New Features

June 2019

Author: Shailesh Dwivedi