



LUXOUG - COMMUNITY

Oracle Sharding Technical Deep Dive

Y V Ravi Kumar

Oracle Certified Master (OCM) & Oracle ACE Director (ACED)






LuxOUG Virtual Tech Days
22nd - 25th June 2020



Venkata Ravi Kumar Yenugula (YVR)

- **Oracle** Certified Master (OCM)
- **Oracle** ACE Director
- Co-author - **Oracle** books
 - **Oracle** Database Upgrade and Migration Methods (Includes 12cR2)
 - **Oracle** High Availability, Disaster Recovery, and Cloud Services
 - **Oracle** GoldenGate with MicroServices
- Co-Author - 100+ **Oracle** Technology Network (OTN) - English, Portuguese & Spanish
- Speaker @**Oracle** Open World | Independent Oracle User Group (IOUG), USA
- Profile published in Spotlight on Success in **Oracle** Certification Program
- Profile published in **Oracle** Magazine in July/Aug 2017 Edition



ORACLE
Certified Master

A | **ORACLE**
ACE Director

500+ Technical Experts Helping Peers Globally



ORACLE
ACE Director



ORACLE
ACE




ORACLE
ACE Associate

bit.ly/OracleACEProgram

Nominate yourself or someone you know: acenomination.oracle.com

ORACLE SHARDING

Oracle Database Sharding Technical Deep Dive



1 Oracle Database Sharding and Architecture

2 Benefits and Deployment Models

3 Oracle Exadata With Sharding

4 Methods of Oracle Database Sharding

5 Components of Oracle Database Sharding

6 Deployment of Oracle Database Sharding



Oracle Database Sharding

It's a Database tier architecture in which data is **horizontally** partitioned across different databases

Basically Splitting a Databases into Smaller Databases

Smaller Databases runs in **dedicated storage** and **dedicated server**

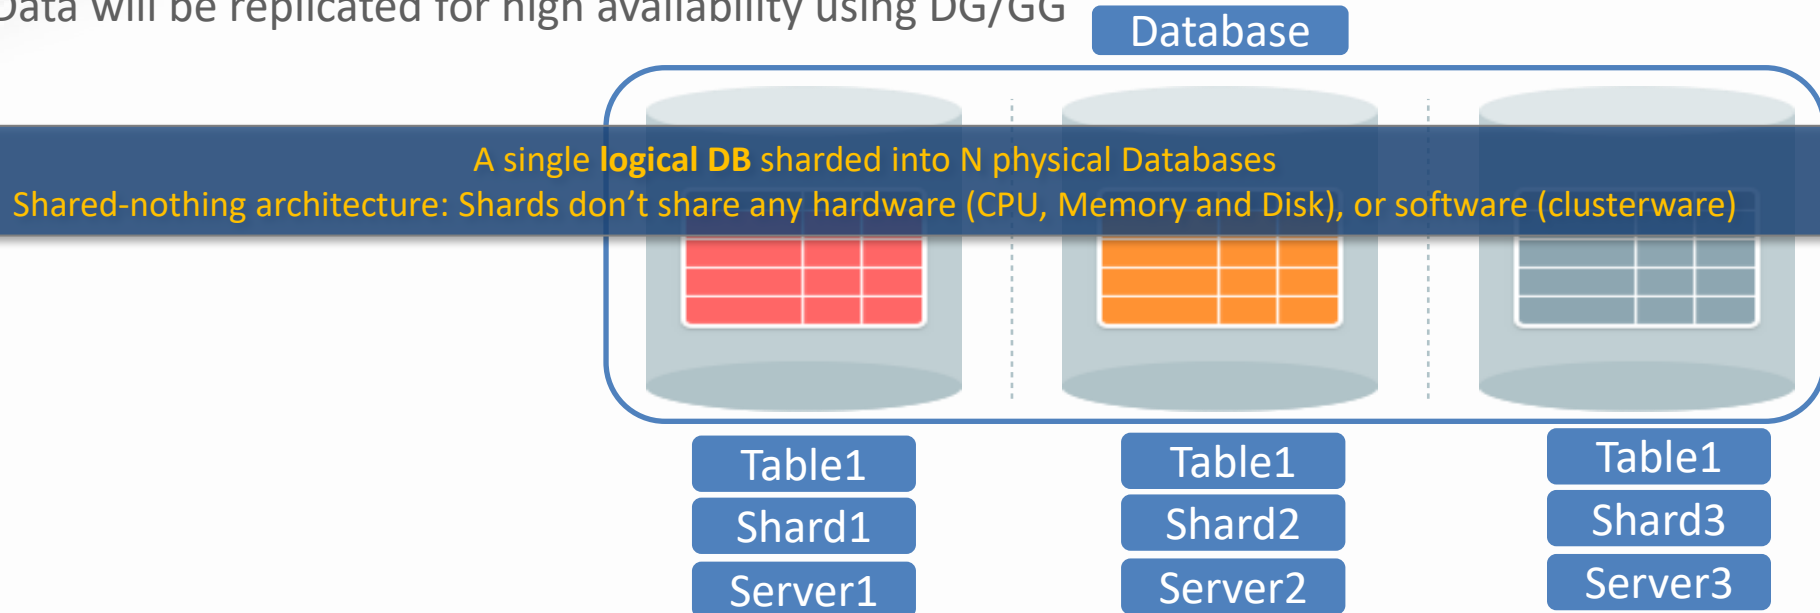
Each Database in sharding configuration is referred as **Shards**

All shards together makeup the single logical group called **Sharded Database (SDB)**

Each Database in sharding configuration is referred as **Shards**

Each Shards contains a table with same column but different subsets of rows called **Chunks**

- **Horizontal partitioning** of data across independent databases (shards)
- Each shard holds a subset of the data
- Can be single-node or RAC or Pluggable Database (PDB)
- Data will be replicated for high availability using DG/GG



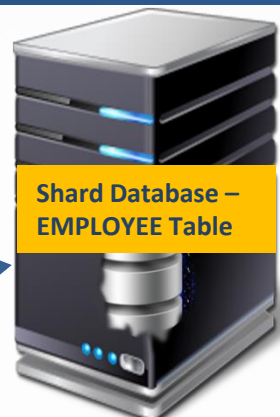
Oracle Database with Sharding Architecture



EMPLOYEE		
EMP_NO	EMP_SCALE_DT	EMP_SCALE
..	..	6
..	..	4
..	..	-
..	..	3
..	..	2
..	..	1
..	..	8



EMPLOYEE		
EMP_NO	EMP_SCALE_DT	EMP_SCALE

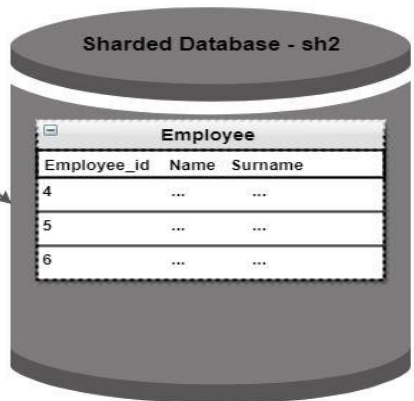
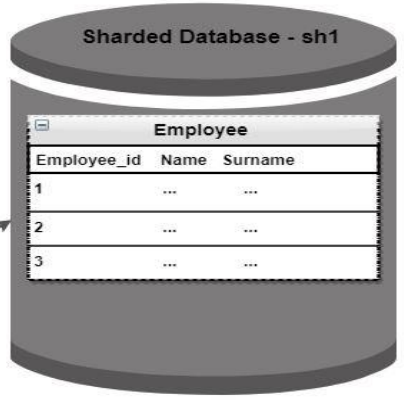




EMPLOYEE		
EMP_NO	EMP_SCALE_DT	EMP_SCALE
..	..	6
..	..	4



EMPLOYEE		
EMP_NO	EMP_SCALE_DT	EMP_SCALE
..	..	-
..	..	3
..	..	2
..	..	1
..	..	8

Transformation to Sharding





Sharding Architecture

Horizontal partitioning of data across up to 1000 independent Oracle Databases (shards)

Horizontal partitioning splits a database table across shards so that each shard contains the table with the same columns but a different subset of rows.

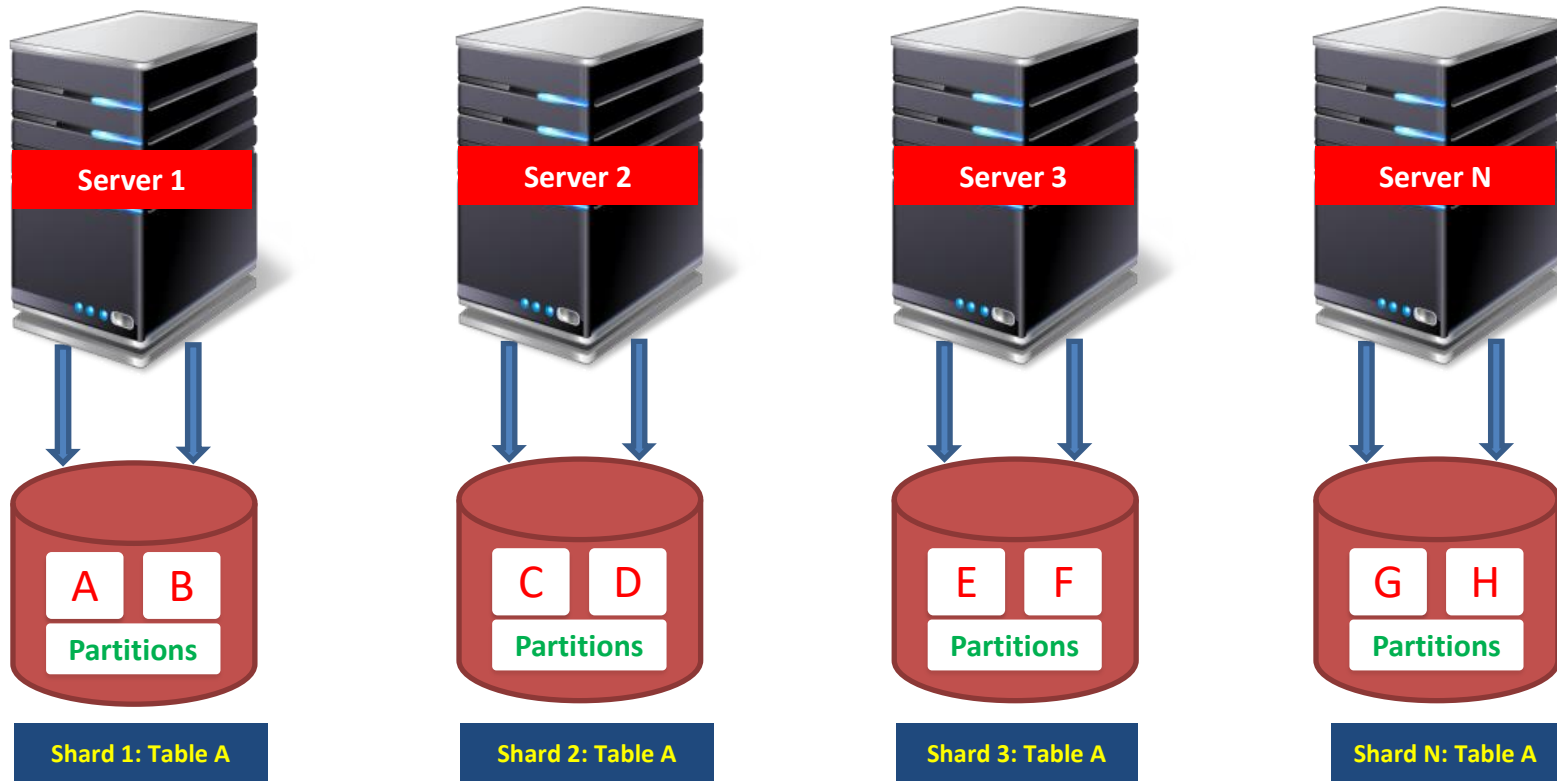
Shared-nothing hardware architecture

- Each shard runs on commodity server
- No shared storage components
- No clusterware components

Data is partitioned using a *sharding key* (i.e. *account_id* in *Account Table*).

NoSQL databases made it easy to deploy Sharding, **Oracle Database Native Sharding** makes it easy for full-featured RDBMS.

- Table that is partitioned into smaller and more manageable pieces among multiple databases





Each shard is an independent Oracle Database.



Oracle Sharding requires a minimum release of Oracle DB 12.2.0 and Oracle Client 12.2.0.



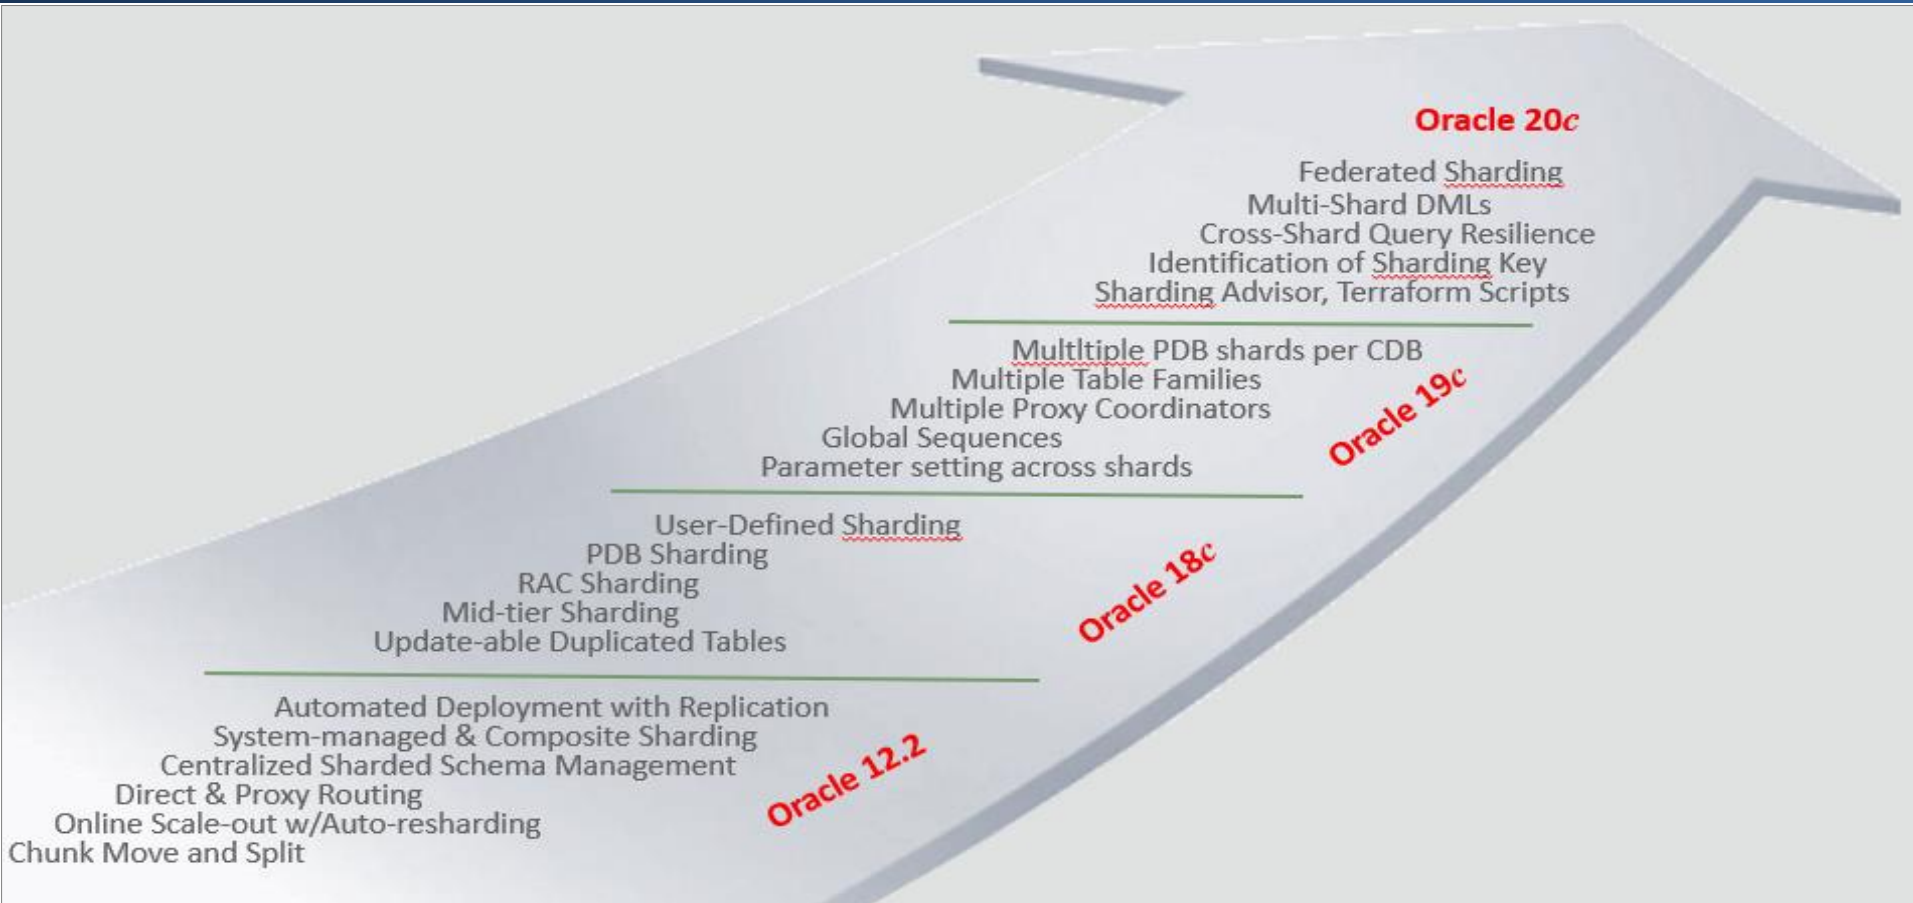
The initial release of Oracle Sharding (12.2.0) does not support Oracle Multitenant.



Started supporting Oracle Multitenant from the release of Oracle Sharding (18c).

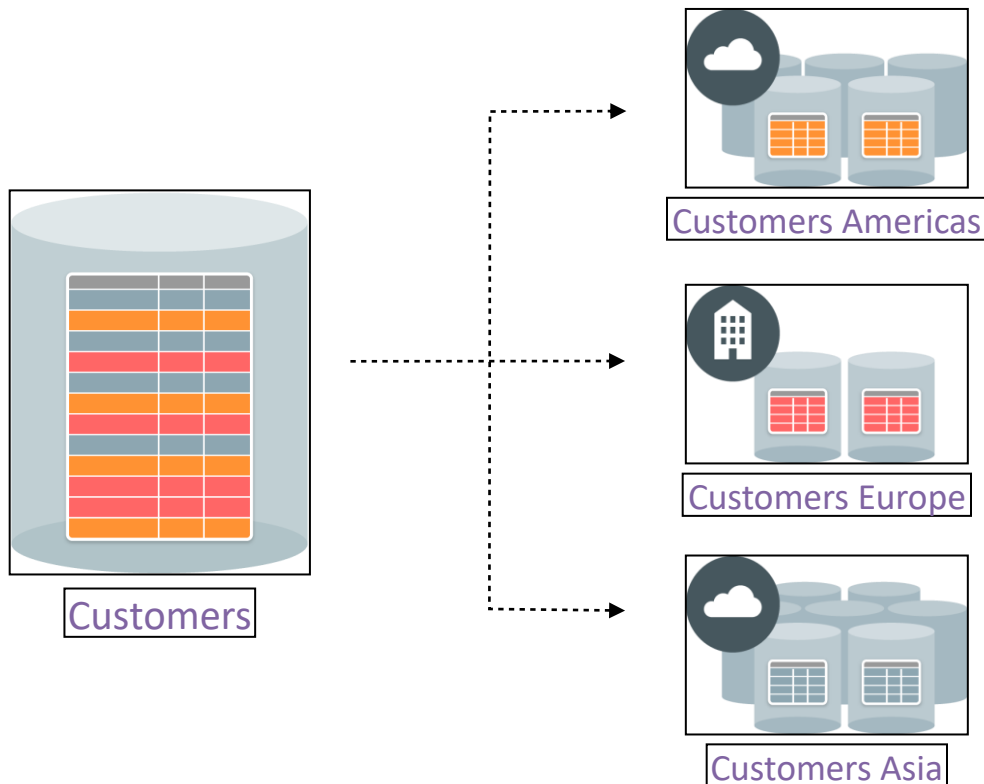


Started supporting Multiple PDB shards per CDB from the release of Oracle Sharding (19c).



Oracle Sharding – Benefits

One mission critical database partitioned into many small databases (shards)



- Extreme scalability by adding shards (independent databases)
- Rolling upgrades with independent availability of shards.
- Global-Scale OLTP applications prefer to shard massive databases into a farm of smaller databases
- Native SQL for sharding tables across up to 1000 Shards
 - Routing of SQL based on shard key, and cross shard queries
 - Online addition and reorganization of shards.

Oracle Database Sharding – Benefits



Linear Scalability



- Add shards online to increase database size and throughput.
- Online split and rebalance.



Extreme Availability



- Shared-nothing hardware architecture.
- Fault of one shard has no impact on others.



Distribution of Data globally





- User defined data placement for performance, availability, DR or to meet regulatory requirements.

Online Addition and Rebalancing of Shards



Online Addition and Rebalancing of Shards





Online Addition and Rebalancing of Shards

- A chunk is lowest level of granularity
- Group of related partitions of a sharded table family
 - *Ex: Chunk#1 contains Customers_P1, Orders_P1, Lineltems_P1*
- Move is initiated automatically or manually (by DBA)
- Automated - Uses RMAN Incremental Backup & Transportable Tablespace

Chunk #1

Sharded Tables →

Books_P1 (1-5000000)

Orders_P1

Sales_P1

Oracle Database Sharding – Flexible Deployment Models



On-Premises




Complete Cloud



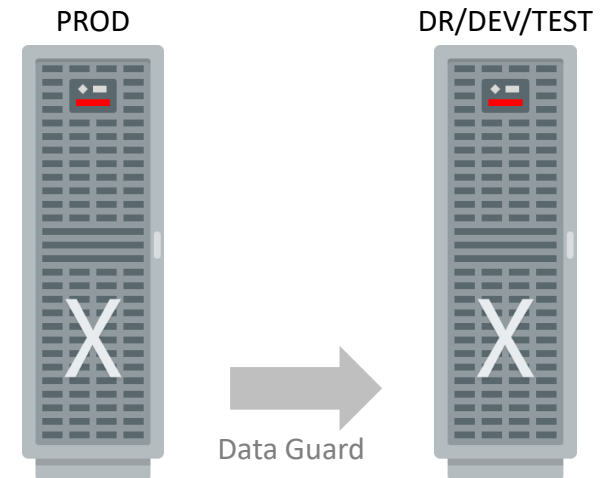
Shards support all DB deployment models like Single Instance, Pluggable Database (PDB), Oracle Exadata & RAC



Hybrid Cloud



- **Best platform for deploying Oracle database**
 - All benefits of Exadata available with Sharding
 - Versatility of Oracle Database PLUS the power of the Exadata infrastructure
 - Simplified lifecycle operations
 - Elastically scale enabled processor capacity with no downtime
- **All Sharded workloads on a single platform**
 - Mission-critical OLTP, analytics, unstructured data (Text/JSON)
 - Disaster Recovery, Reporting
- **Oracle Exadata Storage cells with query offloading provide order of magnitude performance boost**
- **RAC nodes provide instant protection from node failure**



-  Auto deployment & Replication used for shard-level HA
-  Multiple sharding methods
-  Centralized schema maintenance
-  A single global service accesses any shard
-  Direct routing and proxy routing
-  Elastic scaling with automatic rebalance

Oracle Sharding – Advantages



High Availability (HA) features



Compression and Advanced Security



Backup and Recovery at Enterprise-Scale



Database Partitioning applicable



Online schema changes



Guarantees ACID properties and read consistency



Rolling upgrades with independent availability of shards



Custom OLTP Applications

- Large billing systems
- Airline ticketing systems
- Online financial services
- Media companies
- Online information services
- Social media companies



Characteristics / Implementation

- Application must specify a sharding key for optimal performance
- Sharding is **not application transparent** method



System Managed Sharding (*Consistent Hash*)

- System Managed Sharding does not require the user to specify mapping of data to shards.
- Data is automatically distributed across shards using partitioning by consistent hash.



User Defined Sharding (*Range, List*)

- User has control over database means user specifies the mapping of data to individual Shards.
- It is useful in cases where application decides certain data need to be kept in a particular Shard and user have control on moving data between Shards.



Composite Sharding (*Range- Consistent Hash, List- Consistent Hash*)

- Combination of system managed and user defined Sharding.
- Data first partitioned by list or range across multiple shardspaces, and then further partitioned in to consistent hash across multiple shards in each shardspace.

System Managed Sharding	User Defined Sharding	Composite Sharding
Based on partitioning by CONSISTENT HASH	Based on partitioning by RANGE or LIST	Provides two-levels of data organization
Range of hash values assigned to each chunk	List or Range of sharding key values assigned to each chunk by the user	Partitions are distributed among buckets by LIST or RANGE (specified by the user)
Data is uniformly sharded / re-sharded automatically	Need to manually maintain balanced data distribution	Within a bucket, a range of hash values is automatically assigned to each chunk
User has no control on location of data	Full control on location of data provides	Requires two sharding keys (super_sharding_key and sharding_key)



```
SQL> CREATE SHARDED TABLE customers(  
Customer_No          NUMBER NOT NULL,  
Customer_Name        VARCHAR2(50),  
Customer_Address     VARCHAR2(250),  
CONSTRAINT Cust_PK PRIMARY KEY(Customer_No))  
PARTITION BY CONSISTENT HASH (Customer_No)  
PARTITIONS AUTO TABLESPACE SET tbs1;
```

```
SQL> create tablespace tbs1;  
SQL> create tablespace tbs2;  
SQL> CREATE TABLE orders (  
id                   NUMBER,  
country_code        VARCHAR2(5),  
customer_id         NUMBER,  
order_date          DATE,  
order_total         NUMBER(8,2),  
CONSTRAINT orders_pk PRIMARY KEY (id))  
PARTITION BY LIST (country_code)  
(  
PARTITION part_usa VALUES ('USA') tablespace tbs1,  
PARTITION part_uk_and_ireland VALUES ('GBR', 'IRL') tablespace tbs2);
```


```
SQL> create tablespace tsp_set_1;  
SQL> create tablespace tsp_set_2;  
SQL> alter session enable shard ddl;  
SQL> CREATE SHARDED TABLE Customers (  
CustId              VARCHAR2(60),  
FirstName           VARCHAR2(60),  
LastName            VARCHAR2(60),  
Class               VARCHAR2(10),  
Geo                  VARCHAR2(8),  
CustProfile         VARCHAR2(4000),  
Passwd              RAW(60),  
CONSTRAINT pk_customers PRIMARY KEY (CustId),  
CONSTRAINT json_customers CHECK (CustProfile IS JSON))  
partitionset by list(GEO)  
partition by consistent hash(CustId)  
partitions auto (  
partitionset america values ('AMERICA') tablespace set tsp_set_1,  
partitionset europe values ('EUROPE') tablespace set tsp_set_2);
```




*Set
partitionsets
and
tablespace sets
for composite
partitioning*



Active Data Guard with Fast-Start Failover



GoldenGate 'chunk-level' active-active replication with automatic conflict detection/resolution



Optionally – complement replication with Oracle RAC for server HA



Sharded Database (SDB)



Shards in Sharded Database (SDB) - Independent physical Oracle databases that host a subset of the sharded database



Global Service – Database services that provide access to data in an SDB



Shard Directors – Network listeners that enable high performance connection routing based on a sharding key



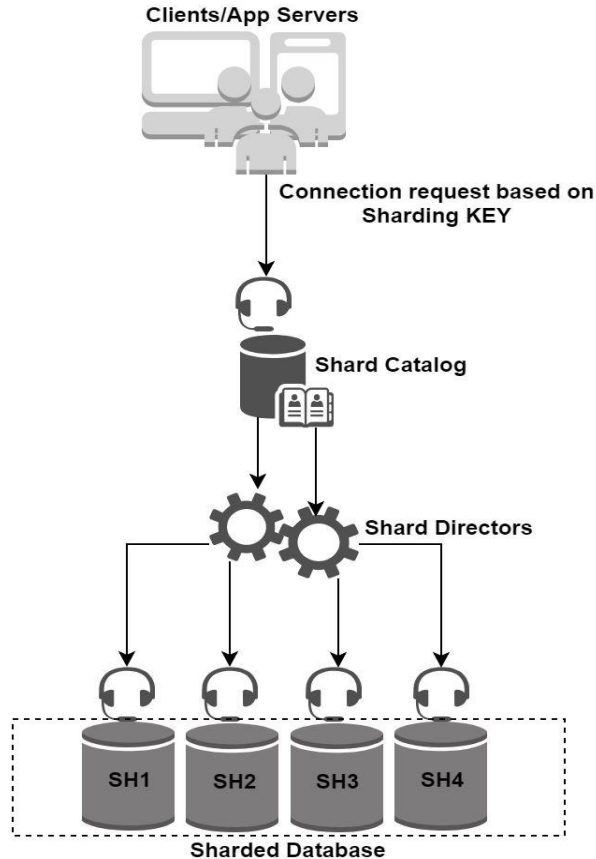
Management Tools – OEM 13C and Global Database Services (GDSCTL)



Connection Pools – At runtime, act as shard directors by routing DB requests across pooled connections



Shard Catalog - Oracle db supports automated shard deployment, centralized mgmt. of a sharded db & multi-shard queries



- Clients pass sharding key (e.g. Customer ID) to Connection pool, connection is routed to the right shard.
- Fast: caching key ranges on client ensures that most accesses go directly to the shard

Shard Catalog

- Stores SDB metadata
- Acts as a query coordinator for multi-shard queries

Shard Director

- Global service manager for direct routing of connection requests to shards
- Publishes run-time SDB topology map, load balancing advisory, FAN events via ONS
- Acting as a regional listener for clients to connect to an SDB.

Sharded Database

- Set of all shards
- Independent physical Oracle databases that host a subset of the sharded database

Global Service

- Single service to access any shard in the SDB



Direct Routing (Based on Sharding_key)

- OLTP workloads must specify sharding_key (e.g. customer_id) during connect
- JDBC/UCP, OCI, and ODP.NET recognize sharding keys



Proxy Routing via a Coordinator(shard catalog)

- For workloads that cannot specify sharding_key (as part of connection).
 - Reporting, Batch jobs workloads
 - Queries spanning one or more or all shards
 - Applications unable to specify a sharding_key



Database Connections



- Must use Oracle integrated connection pools (UCP, OCI, ODP.NET, JDBC)
- Must be able to separate workloads that use Direct Routing from those that use Proxy Routing
- Each uses separate connection pools

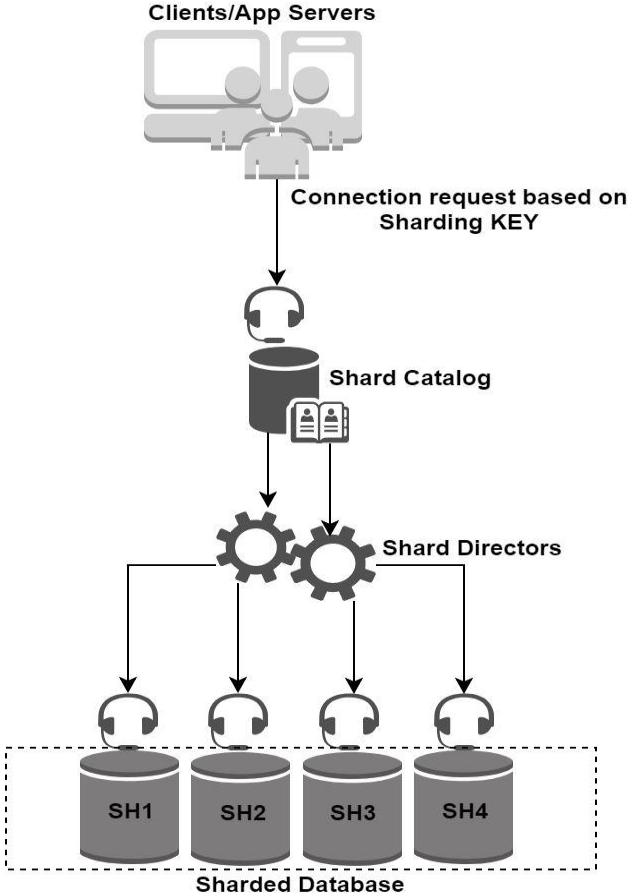
Shard Catalog




Host:shrcat

Install Oracle Database in all five nodes with the following options:
➢ Install Database software only
➢ Oracle Enterprise Edition

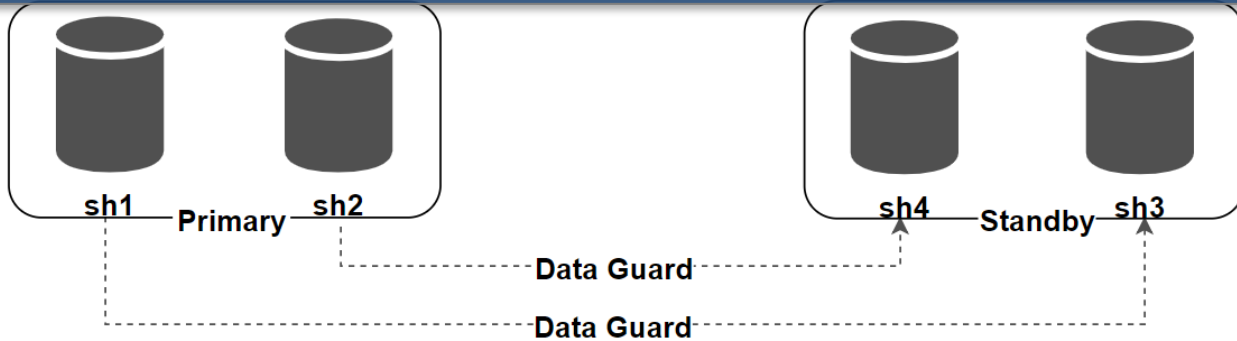




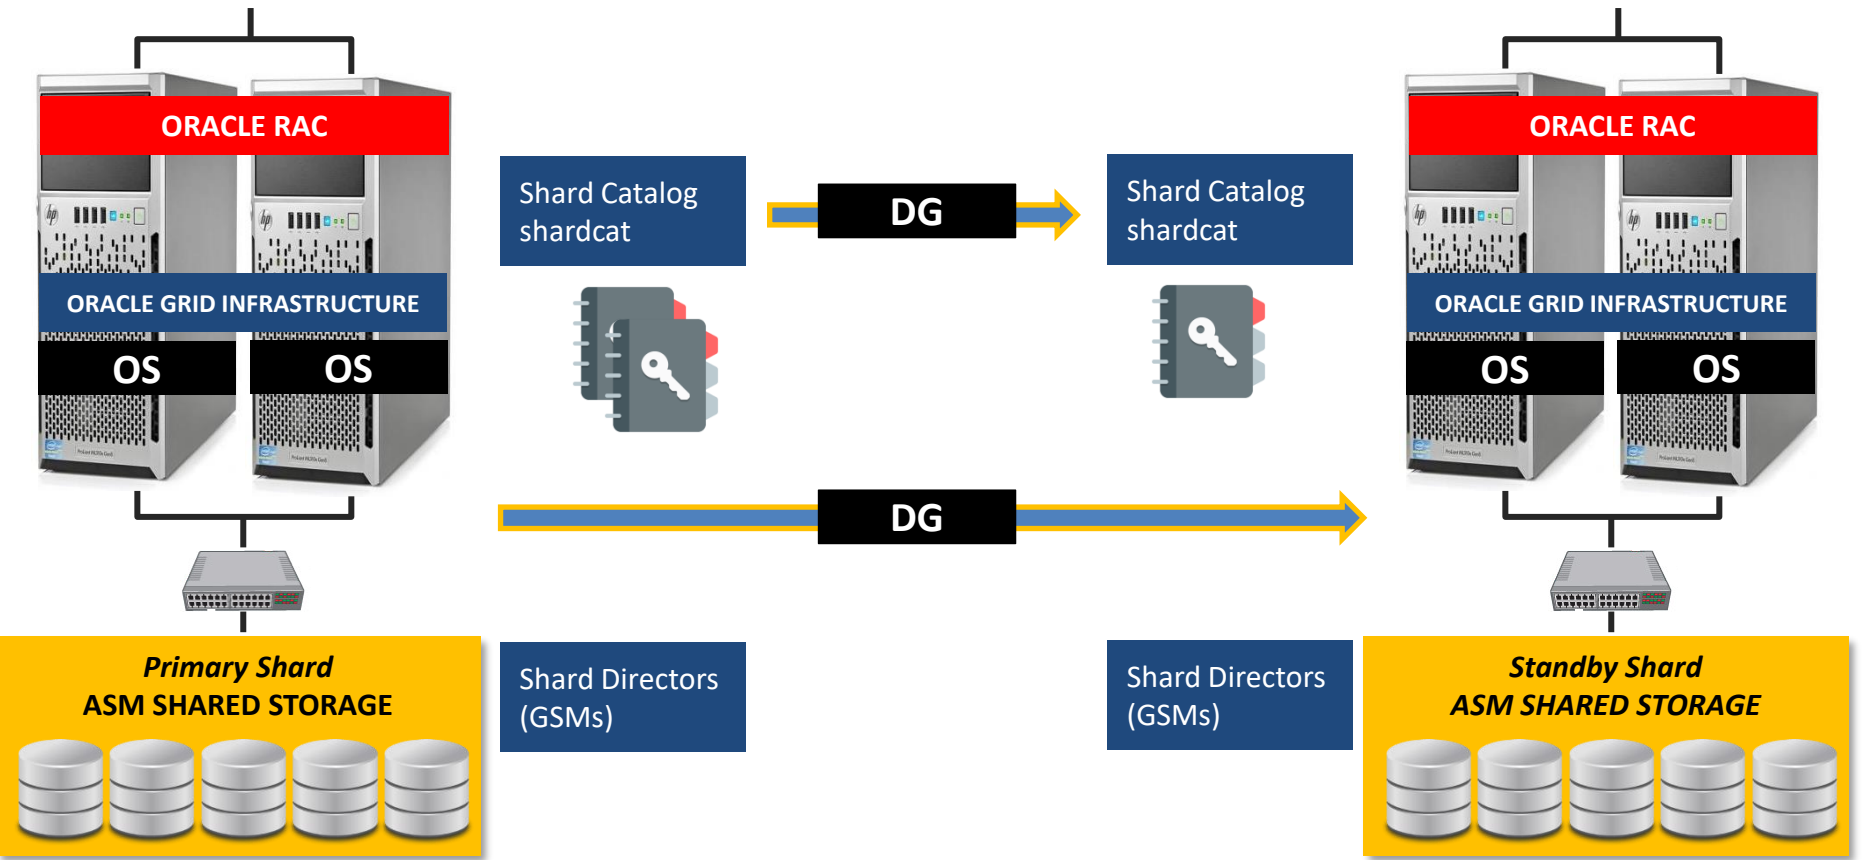


GSM

High Availability for Shard Catalog



Oracle Sharding using Oracle RAC



Setting up Oracle Sharding

Installation of GSM

Specify Installation Location

Specify Installation Location

Specify an Oracle base path to place all Oracle software and configuration-related files. This location is the Oracle base directory.

Oracle base

Specify a location for storing Oracle software files. This location is the Oracle home directory.

ORACLE
DATABASE 12c

Install Product

Specify Installation Location

Prerequisite Checks

Summary

Install Product

Finish

Progress

94%

Successfully executed utility tool: ADR Setup Utility

Status

Task	Status
Oracle Database Global Service Manager Installation	Succeeded
• Prepare	Succeeded
• Copy files	Succeeded
• Link binaries	Succeeded
• Seton	Succeeded
• Seton	In Progress

The following script needs to be executed as the "root" user.

Scripts

```
/u01/app/gsm/12.2.0/root.sh
```

To execute the configuration scripts:

1. Open a terminal window
2. Login as "root"
3. Run the scripts
4. Return to this window and click "OK" to continue

ORACLE
DATABASE

Install the software in a separate home `/u01/app/gsm/12.2.0`
`./runInstaller -silent -responseFile /u01/install/gsm/response/gsm_install.rsp`

Summary

Specify Installation Location

Prerequisite Checks

Summary

Install Product

Finish

Oracle Database 12c Release 2 GSM Installer

Global Settings

- Disk Space: required 1.5 GB available 9.47 GB [\[Edit\]](#)
- Source Location: /u01/install/gsm/install/./stage/products.xml
- Oracle Base: /u01/app/gsm [\[Edit\]](#)
- Oracle Home location: /u01/app/gsm/12.2.0 [\[Edit\]](#)

ORACLE
DATABASE 12c

Set catalog database environment and start listener

```
ORACLE_SID = [oracle] ? SHRCAT
The Oracle base has been set to /u01/app/oracle
$ lsnrctl start
```

open_links and open_links_per_instance are set to 16 (Optional)

```
$ sqlplus / as sysdba
SQL> alter system set open_links=16 scope=spfile;
SQL> alter system set open_links_per_instance=16 scope=spfile;
SQL> shutdown immediate
SQL> startup
```

Grant roles and privileges on the database

```
SQL> alter user gsmcatuser account unlock;
SQL> alter user gsmcatuser identified by oracle;
SQL> create user gsmadmin identified by oracle;
SQL> grant connect, create session, gsmadmin_role to gsmadmin;
SQL> grant inherit privileges on user SYS to GSMADMIN_INTERNAL;
```



Connect to shard director host and start GDSCTL

```
$ export ORACLE_HOME=/u01/app/gsm/12.2.0/  
$ export ORACLE_BASE=/u01/app/gsm  
$ export PATH=$ORACLE_HOME/bin:$PATH  
[oracle@sourcedb1 ~]$ gdsctl  
Current GSM is set to SHARDDIRECTOR
```



Create the shard catalog and configure the remote scheduler agent on the shard catalog

```
GDSCTL> create shardcatalog -database shcat:1521:SHRCAT -user gsmadmin/oracle -chunks 12 -sdb  
SHRCAT -region us-west,us-east,us-north  
Catalog is created
```



Create and start the shard director, set the operating system credentials

```
GDSCTL>add gsm -gsm mysharddirector -listener 1522 -pwd oracle -catalog shcat:1521:SHRCAT -region us-north  
GSM successfully added  
GDSCTL>start gsm -gsm mysharddirector  
GSM is started successfully  
Check the status of gsm:  
GDSCTL> status gsm  
GDSCTL>add credential -credential os_oracle_cred -osaccount oracle -ospassword oracle  
The operation completed successfully
```



Connect to catalog database, to set the scheduler port and password

```
$ sqlplus / as sysdba
SQL> exec DBMS_XDB.sethttpport(8080);
SQL> commit;
SQL> exec DBMS_SCHEDULER.SET_AGENT_REGISTRATION_PASS('welcome1');
SQL> alter system register;
```



Connect to each of the shard hosts, register remote scheduler agents on them, and create directories for Oradata and FRA

```
$ schagent -registerdatabase shcat 8080
[oracle@shdb1 ~]$ schagent -start
[oracle@shdb1 ~]$ schagent -status
[oracle@shdb1 ~]$ mkdir /u01/app/oracle/oradata
[oracle@shdb1 ~]$ mkdir /u01/app/oracle/fast_recovery_area
[oracle@shdb1 ~]$ exit
$ schagent -registerdatabase shcat 8080
[oracle@shdb2 ~]$ schagent -start
[oracle@shdb2 ~]$ schagent -status
[oracle@shdb2 ~]$ mkdir /u01/app/oracle/oradata
[oracle@shdb2 ~]$ mkdir /u01/app/oracle/fast_recovery_area
```

```
GDSCTL>create shardcatalog -database shcat:1521:SHRCAT -user gsmadmin/oracle -chunks 12 -sdb SHRCAT -region us-  
west,us-east,us-north  
  
GDSCTL>add gsm -gsm mysharddirector -listener 1522 -pwd oracle -catalog shcat:1521:SHRCAT -region us-north  
  
GDSCTL>start gsm -gsm mysharddirector  
GDSCTL>status gsm  
  
GDSCTL>set gsm -gsm mysharddirector  
GDSCTL>connect gsmadmin/oracle  
  
GDSCTL>add shardgroup -shardgroup shgroup1 -deploy_as primary -region us-west  
GDSCTL>add shardgroup -shardgroup shgroup2 -deploy_as active_standby -region us-east  
  
GDSCTL>add invitednode shdb1 -group shgroup1  
GDSCTL>create shard -shardgroup shgroup1 -destination shdb1 -credential os_oracle_cred -dbparamfile /tmp/dbcaparams -  
sys_password oracle  
  
GDSCTL>add invitednode shdb2 -group shgroup1  
GDSCTL>create shard -shardgroup shgroup1 -destination shdb2 -credential os_oracle_cred -dbparamfile /tmp/dbcaparams -  
sys_password oracle  
  
GDSCTL>add invitednode shdb3 -group shgroup2  
GDSCTL>create shard -shardgroup shgroup2 -destination shdb3 -credential os_oracle_cred -dbparamfile /tmp/dbcaparams -  
sys_password oracle  
  
GDSCTL>add invitednode shdb4 -group shgroup2  
GDSCTL>create shard -shardgroup shgroup2 -destination shdb4 -credential os_oracle_cred -dbparamfile /tmp/dbcaparams -  
sys_password oracle
```

```
GDSCCTL>deploy
deploy: examining configuration...
deploy: deploying primary shard 'sh1' ...
deploy: network listener configuration successful at destination 'shdb1'
deploy: starting DBCA at destination 'shdb1' to create primary shard 'sh1' ...
deploy: deploying primary shard 'sh2' ...
deploy: network listener configuration successful at destination 'shdb2'
deploy: starting DBCA at destination 'shdb2' to create primary shard 'sh2' ...
deploy: waiting for 2 DBCA primary creation job(s) to complete...
deploy: waiting for 2 DBCA primary creation job(s) to complete...
deploy: waiting for 2 DBCA primary creation job(s) to complete...
deploy: waiting for 2 DBCA primary creation job(s) to complete...
deploy: DBCA primary creation job succeeded at destination 'shdb1' for shard 'sh1'
deploy: deploying standby shard 'sh3' ...
deploy: network listener configuration successful at destination 'shdb3'
deploy: starting DBCA at destination 'shdb3' to create standby shard 'sh3' ...
deploy: DBCA primary creation job succeeded at destination 'shdb2' for shard 'sh2'
deploy: deploying standby shard 'sh4' ...
deploy: network listener configuration successful at destination 'shdb4'
deploy: starting DBCA at destination 'shdb4' to create standby shard 'sh4' ...
deploy: waiting for 2 DBCA standby creation job(s) to complete...
deploy: waiting for 2 DBCA standby creation job(s) to complete...
deploy: waiting for 2 DBCA standby creation job(s) to complete...
deploy: DBCA standby creation job succeeded at destination 'shdb3' for shard 'sh3'
deploy: DBCA standby creation job succeeded at destination 'shdb4' for shard 'sh4'
deploy: requesting Data Guard configuration on shards via GSM
deploy: shards configured successfully
The operation completed successfully
```

- DBMS_SCHEDULER (executed on Shard Catalog) communicates with Scheduler Agents on remote hosts
 - Agents run DBCA and NETCA to create shards and listeners
- Oracle Active Data Guard replication
 - Primary databases are created first
 - DBCA uses RMAN duplicate to create corresponding standbys
 - Redo transport and Broker are configured
 - Observers are started on Shard Director hosts and Fast Start Failover is enabled
- Oracle GoldenGate replication
 - Replication pipelines are configured and replication is started

```
$ ls /u01/app/oracle/cfgtoollogs/dbca/silent.log*
$ tail -f /u01/app/oracle/cfgtoollogs/dbca/silent.log
Copying database files
DBCA_PROGRESS : 1%
DBCA_PROGRESS : 2%
DBCA_PROGRESS : 16%
DBCA_PROGRESS : 25%
DBCA_PROGRESS : 45%
DBCA_PROGRESS : 78%
DBCA_PROGRESS : 86%
```

Query the ***dba_scheduler_running_jobs*** on catalog database for monitoring or diagnosing during the deploying.

The exact location of a given GSM's log and trace files can be obtained using the status gsm command.

- ***GDSCTL> status gsm***

GDSCTL>config shard

Name	Shard Group	Status	State	Region	Availability
----	-----	-----	-----	-----	-----
sh1	shgroup1	Ok	Deployed	us-west	ONLINE
sh2	shgroup1	Ok	Deployed	us-west	ONLINE
sh3	shgroup2	Ok	Deployed	us-east	READ ONLY
sh4	shgroup2	Ok	Deployed	us-east	READ ONLY

GDSCTL>databases

Database: "sh1" Registered: Y State: Ok ONS: N. Role: PRIMARY Instances: 1 Region: us-west
Alert: Data Guard observer is not running.

Registered instances:

shrcat%1

Database: "sh2" Registered: Y State: Ok ONS: N. Role: PRIMARY Instances: 1 Region: us-west
Alert: Data Guard observer is not running.

Registered instances:

shrcat%11

Database: "sh3" Registered: Y State: Ok ONS: N. Role: PH_STNDBY Instances: 1 Region: us-east

Registered instances:

shrcat%21

Database: "sh4" Registered: Y State: Ok ONS: N. Role: PH_STNDBY Instances: 1 Region: us-east

Registered instances:

shrcat%31

Configuration of Oracle Sharding

Tools

Upgrade VMs

Primary-Standby

19c_Security

ogg

RAC 12cR1 Athena

19cr7-oe18

Exadata-Quarter Rack

CentOS7

11g_RAC

Sharding_DG

SHCAT Running

SHDB1 Running

SHDB2 Running

SHDB3 Running

SHDB4 Running

19c-Sharding

General

Name: SHCAT
Operating System: Oracle (64-bit)
Groups: Sharding_DG

System

Base Memory: 4096 MB
Boot Order: Floppy, Optical, Hard Disk
Acceleration: VT-x/AMD-V, Nested Paging, PAE/NX, KVM Paravirtualization

General

Name: SHDB1
Operating System: Oracle (64-bit)
Groups: Sharding_DG

System

Base Memory: 4096 MB
Boot Order: Floppy, Optical, Hard Disk
Acceleration: VT-x/AMD-V, Nested Paging, PAE/NX, KVM Paravirtualization

General

Name: SHDB2
Operating System: Oracle (64-bit)
Groups: Sharding_DG

System

Base Memory: 4096 MB
Boot Order: Floppy, Optical, Hard Disk
Acceleration: VT-x/AMD-V, Nested Paging, PAE/NX, KVM Paravirtualization

Shard Catalog

Primary Shard Database (SH1)
Physical Standby Database (SH3)

Primary Shard Database (SH2)
Physical Standby Database (SH4)

```
Copyright (c) 2011, 2016, Oracle. All rights reserved.

Welcome to GDSCTL, type "help" for information.

Current GSM is set to MYSHARDDIRECTOR
GDSCTL>status database
Database: "sh1" Registered: Y State: Errors ONS: N. Role: PRIMARY Instances: 1 Region: us-west
  Service: "primdb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Service: "standb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Registered instances:
    shrcat/1
Database: "sh2" Registered: Y State: Errors ONS: N. Role: PRIMARY Instances: 1 Region: us-west
  Service: "primdb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Service: "standb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Registered instances:
    shrcat/11
Database: "sh3" Registered: Y State: Errors ONS: N. Role: PH_STNDBY Instances: 1 Region: us-east
  Service: "primdb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Service: "standb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Registered instances:
    shrcat/21
Database: "sh4" Registered: Y State: Errors ONS: N. Role: PH_STNDBY Instances: 1 Region: us-east
  Service: "primdb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Service: "standb" Globally started: Y Started: N
    Scan: Y Enabled: Y Preferred: Y
  Registered instances:
    shrcat/31

GDSCTL>
```

```
oracle@shcat ~1$ . oraenv
ORACLE_SID = [SHRCAT] ? SHRCAT
The Oracle base remains unchanged with value /u01/app/oracle
oracle@shcat ~1$ sqlplus appuser/oracle123

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 18:58:49 2020

Copyright (c) 1982, 2016, Oracle. All rights reserved.

Last Successful login time: Wed Jun 17 2020 18:53:56 -04:00

Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a10
SQL> col auth_lname format a10
SQL>
SQL> insert into books values (0625,'Expert Oracle Exadata','Tanel','Poder');

1 row created.

SQL> commit;

Commit complete.

SQL> select * from books where bookid=0625;

BOOKID TITLE                                AUTH_FNAME AUTH_LNAME
-----
   625 Expert Oracle Exadata                Tanel      Poder

SQL> _
```

Transactions appeared in sh2 and sh4

```
oracle@shdb2 ~1$ . oraenv
ORACLE_SID = [sh2] ? sh2
The Oracle base remains unchanged with value /u01/app/oracle
oracle@shdb2 ~1$ sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 19:03:28 2020

Copyright (c) 1982, 2016, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a10
SQL> col auth_lname format a10
SQL>
SQL> select open_mode,database_role from v$database;

OPEN_MODE          DATABASE_ROLE
-----
READ WRITE         PRIMARY

SQL> select * from appuser.books where bookid=0625;

BOOKID TITLE          AUTH_FNAME AUTH_LNAME
-----
   625 Expert Oracle Exadata      Tanel     Poder

SQL>
```

**Primary Shard
SH2**

```
oracle@shdb4 ~1$ . oraenv
ORACLE_SID = [oracle] ? sh4
The Oracle base has been set to /u01/app/oracle
oracle@shdb4 ~1$ sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 18:51:42 2020

Copyright (c) 1982, 2016, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a20
SQL> col auth_lname format a20
SQL>
SQL> select open_mode,database_role from v$database;

OPEN_MODE          DATABASE_ROLE
-----
READ ONLY WITH APPLY PHYSICAL STANDBY

SQL>
SQL> select * from appuser.books where bookid=0625;

BOOKID TITLE          AUTH_FNAME AUTH_LNAME
-----
   625 Expert Oracle Exadata      Tanel     Poder

SQL>
```

**Physical Standby
SH4**

```
loracle@shcat ~1$ . oraenv
ORACLE_SID = [SHRCAT] ? SHRCAT
The Oracle base remains unchanged with value /u01/app/oracle
loracle@shcat ~1$ sqlplus appuser/oracle123

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 19:09:43 2020
Copyright (c) 1982, 2016, Oracle. All rights reserved.

Last Successful login time: Wed Jun 17 2020 18:58:49 -04:00

Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a10
SQL> col auth_lname format a10
SQL>
SQL> insert into books values (0626,'Effective Oracle By Design','Thomas','Kyte');

1 row created.

SQL> commit;

Commit complete.

SQL> select * from books where bookid=0626;

BOOKID TITLE                                AUTH_FNAME AUTH_LNAME
-----
   626 Effective Oracle By Design    Thomas      Kyte

SQL>
```

```
loracle@shdb1 ~]# . oraenv
ORACLE_SID = [sh1] ? sh1
The Oracle base remains unchanged with value /u01/app/oracle
loracle@shdb1 ~]# sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 19:13:12 2020

Copyright (c) 1982, 2016, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a10
SQL> col auth_lname format a10
SQL>
SQL> select open_mode,database_role from v$database;
```

**Primary Shard
SH1**

OPEN_MODE	DATABASE_ROLE
READ WRITE	PRIMARY

```
SQL>
SQL> select * from appuser.books where bookid=0626;
```

BOOKID	TITLE	AUTH_FNAME	AUTH_LNAME
626	Effective Oracle By Design	Thomas	Kyte

```
SQL> _
```

```
loracle@shdb3 ~]# . oraenv
ORACLE_SID = [sh3] ? sh3
The Oracle base remains unchanged with value /u01/app/oracle
loracle@shdb3 ~]# sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Wed Jun 17 19:16:48 2020

Copyright (c) 1982, 2016, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production

SQL> set lines 100 pages 100
SQL> col bookid format 9999
SQL> col title format a28
SQL> col auth_fname format a10
SQL> col auth_lname format a10
SQL>
SQL>
SQL> select open_mode,database_role from v$database;
```

**Physical Standby
SH3**

OPEN_MODE	DATABASE_ROLE
READ ONLY WITH APPLY PHYSICAL STANDBY	

```
SQL>
SQL> select * from appuser.books where bookid=0626;
```

BOOKID	TITLE	AUTH_FNAME	AUTH_LNAME
626	Effective Oracle By Design	Thomas	Kyte

```
SQL>
```

Oracle Sharding – Patching for SDB



OPatchauto supports

- All sharding schemes and replication methods
- Single instance and clustered databases (also handles Grid Infrastructure)



Patch a sharded database

- `<CATALOG_DB_HOME>/OPatch/opatchauto apply <patch loc> -sdb -wallet <wallet file loc> -sid <sid of shardcat> -port <shardcat port>`



Patch a Oracle Data Guard - OPatchauto supports rolling mode (default: parallel mode)

- Data Guard configurations are patched one after the other
- For a given configuration, standbys are patched first followed by primary



Patch a Oracle GoldenGate

- All shardspaces are patched in parallel
- Within a shardspace, shards are patched in rolling mode



Existing MAA best practices for backup apply to SDB

- Best practices for Disk, Tape, or Oracle Secure Backup (OSB)
- Determine frequency and retention period
- Use Recovery Manager (RMAN) catalog
- Enable Block Change Tracking (BCT)
- Enable auto backup for control file and server parameter file
- Offload backups to physical standby

Thanks for your TIME



yvrk1973@gmail.com



@yvrk1973



<https://yvrk1973.blogspot.com/>



<https://www.linkedin.com/in/yv-ravikumar/>

