The Twelve Days of 12.2

Walk through Oracle Database12c Release 2 (12.2) New Features
with Viscosity Experts & Consultants

On the Twelfth Day of 12.2 my DBA gave to me...

ORACLE®

Welcome to the 12 days of 12c | 12.2 – New Features

For 12 days, Viscosity will release a new article about 12.2 new features, written by our Oracle Ace Directors and Consultants. Be sure to check back every day for the newest article!

On the Twelfth day of 12.2, my DBA gave to me... Pluggable Databases (PDBs)

Dec 12, Day 1: RAC and Grid Infrastructure	Dec 13, Day 2: Data Guard	Dec 14, Day 3: Partitioning	Dec 15, Day 4: ASM
Dec 16, Day 5: SQL Performance Tuning	Dec 17, Day 6: ACFS	Dec 18, Day 7: DB Security	Dec 19, Day 8: Index
Dec 20, Day 9: In-Memory	Dec 21, Day 10: RMAN	Dec 22, Day 11: Utilities, PL/SQL, & More	Dec 23, Day 12: Pluggable Databases

Evolution of the DBA

Kind of DBA	Timeline
CLI DBA	Early 90's DBAs
GUI DBA	Late 90's and Dot Com
Google DBA	Dot Com and 2000's
iDBA	Dot Com, IOUG iDBA Master Curriculum
RAC DBAs	2000+ after 9.2 (but major spike with 10.2)
DMA	2010+ Database Machine Administrator
vDBA / vRAC DBA	2010+ Evolving role of a DBA in the virtual world
Cloud DBA	2011+ Database Consolidation with Private Database Cloud Oracle Database 12c Launches June 2013
Public Cloud DBA	2015+ Oracle Public Cloud with Database Cloud Service, Database Backup Cloud Service, Storage Cloud Service, IaaS Cloud Service
PDBAs	2017+ Multi-Tenant with Oracle Database 12c Release 2 GA – March 15, 2017

"It is not the strongest or the most intelligent who will survive but those who can best manage [adapt to] change."

4 Quarters of Tech Deep Dives

Q1 2017

Data Guard 12.2 New Features

January 2017

By Charles Kim, Oracle ACE Director

VISCOSITY NORTH AMERICA

ORACLE® Platinum Partner

My Speaking Events 2017

March TBD	Cloud Experience Houston	Houston, TX
•March TBD	Sun Coast OUG Tech Day - Cloud Experience	Tampa, FL
March 13-14	Utah Oracle Users Group Training Days	Salt Lake City, UT
March 15-16	Cloud Experience @GaOUG Tech Day	Atlanta GA
March TBD	Cloud Experience Oklahoma	Oklahoma
April 2-6	IOUG Collaborate17 at Mandalay Bay	Las Vegas, NV
May 17-18	Great Lakes OUG Conference	Cleveland, OH
June 25-29	ODTUG Kscope17 @ JW Marriot	San Antonio, TX
Nov 8	Michigan Oracle User Group	Detroit, MI

@ViscosityNA

VISCOSITY NORTH AMERICA

Agenda and Topics

- Why Cloud
- Upgrade Considerations

- RAC and Grid Infrastructure
- ASM
- ACFS

- In-Memory

- General Purpose New Features

- Data Guard

- PDB
- RMAN

VISCOSITY NORTH AMERICA

Collaborate – Independent Oracle User Group
Cloud Experience @IOUG 2017

A boot camp design by Oracle User Community Technology Experts for Oracle Technology Professionals to experience uses of the Oracle Public Cloud and Oracle Technology on-premise to accomplish work critical for their businesses and careers.

IOUG Cloud SIG
ioug.org/cloud
@cloud_sig

@ViscosityNA

VISCOSITY NORTH AMERICA

cloud.oracle.com/tryit

Oracle Database Cloud Service (DBaaS - PaaS)

Use Cases to Consider:

- Disaster Recovery (DR) to the Cloud (Backup)
- Business Critical Workloads in the Cloud (Scale)
- **Dev/Test for 12c Database in the Cloud (Test)**
- **Web Application Development Anywhere (Dev)**
- Migration of On-Premise Apps to the Cloud (Migrate)

@ViscosityNA

VISCOSITY NORTH AMERICA

Oracle 12cR2 introduced on the Cloud 9/18/2016 On-Premise 12cR2 released just under 6 months later on 3/1/17

Announcing: Exadata Express Cloud Service
Simple to Use, Lowest Cost Database Cloud Service

- Oracle Enterprise Edition including all DB options
- Runs on Exadata in Oracle Public Cloud
- Fully managed by Oracle
- Low cost, starting at $175 per month
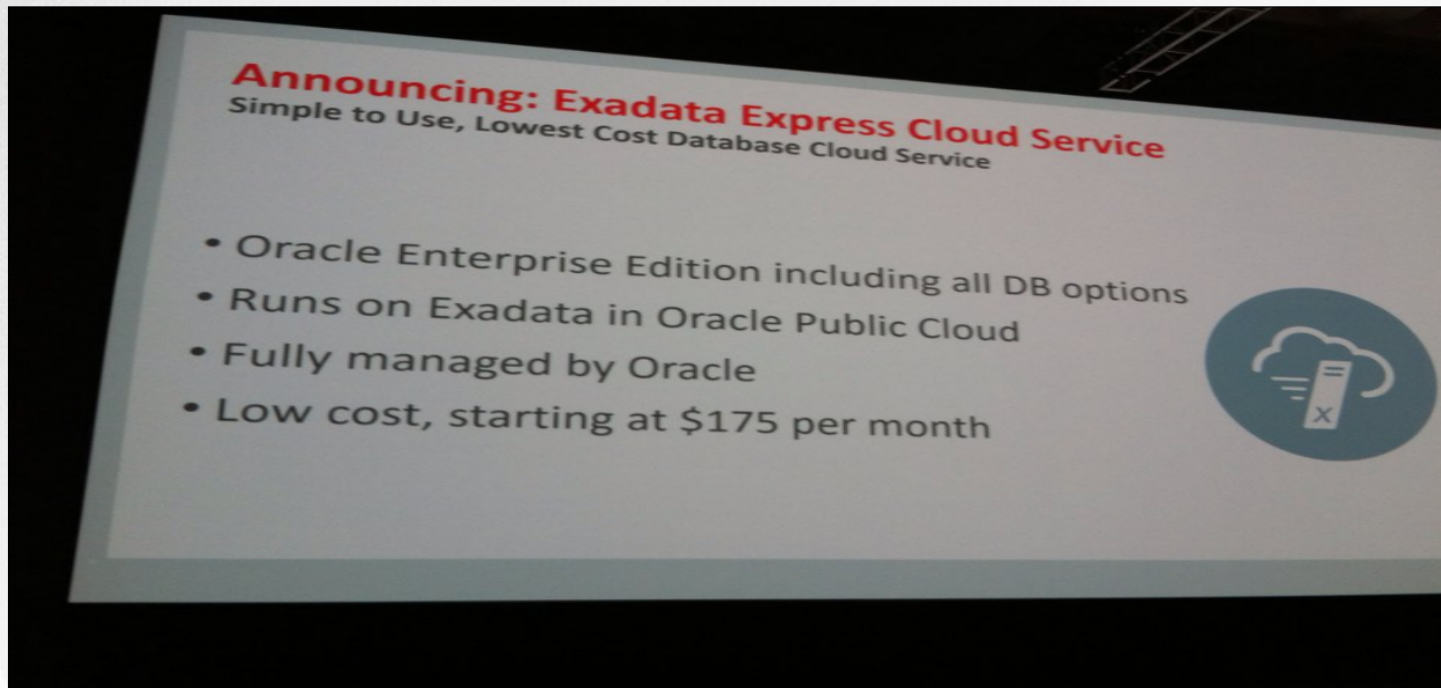
12cR2 Tuning Book allowed to be released 3/10/2017

Rich Niemiec @RichNiemiec · 18 Sep 2016
Larry announces Exadata Express Cloud Service running 12cR2 for $175/month with all features. Available now! pic.twitter.com/3tqJdDuzJe

2   21   23

12cR2 Initialization Parameters

- **Oracle 12c Release 2 now has 4649 initialization parameters**
- **Oracle 12c Release 2 now has 412 documented and 4237 hidden.**

- **Oracle 12c R2 has 746 V$ views and 1,312 X$ tables**

- There is also an **additional column on many V$ views** for Container ID (**CON_ID**) with PDBs (Pluggable Databases) and CDBs (Container Database).
- Make sure that you adjust any scripts that you have to accommodate this new column.

Oracle Database 12c Release 2: Upgrade Paths

Direct Upgrade Path

Source Database	Target Database
10.1.0.5 (or higher)	11.2.0.4 or 12.1.0.2
11.1.0.7 (or higher)	11.2.0.4 or 12.1.0.2
11.2.0.3 (or higher; includes 12.1)	**12.2.x**

In-Direct Upgrade Path

Source Database	Upgrade Path for Target Database	Target Database
7.3.3.0.0 (or lower)	7.3.4.x --> 9.2.0.8	11.2.0.3+
8.0.5.0.0 (or lower)	8.0.6.x --> 9.2.0.8	11.2.0.3+
8.1.7.0.0 (or lower)	8.1.7.4 --> 9.2.0.8	11.2.0.3+
9.0.1.3.0 (or lower)	9.0.1.4 --> 9.2.0.8	11.2.0.3+

14

Database Upgrade Assistant (DBUA)

- DBUA checks before the upgrade:
  - Invalid user accounts or roles
  - Invalid data types or invalid objects
  - De-supported character sets
  - **Adequate resources** (rollback segments, tablespaces, and free disk space)
  - Missing SQL scripts needed for the upgrade
  - Listener running (if Oracle Enterprise Manager Database Control upgrade or configuration is requested)
  - Oracle Database software linked with Database Vault option. If Database Vault is enabled, Disable Database Vault before upgrade (Vault installed by default on).
- Deinstallation Tool integrated with Installation Media

RAC and Grid Infrastructure

- Streamlined GI Installation
  - GI Software available as an image file for download and installation
  - Simply create new GI home with the appropriate user/group permissions, and extract image file into new home
  - Execute gridsetup.sh to invoke the setup wizard
  - Can be used for all RAC and Standalone configurations

- Rapid Home Provisioning- Method of deploying software homes to any number of nodes in a data center
  - Create, store, & manage templates of Oracle home gold images(databases, middleware, & applications)
  - Make working copy of gold image and deploy
  - Provision GI Homes of 11.2.0.3/4, 12.1.0.2, and 12.2
  - Provision RAC clusters, Patch, and Upgrade Oracle Grid Infrastructure

RAC and Grid Infrastructure

- RAC Reader Nodes. - Further leverages Flex Cluster Hub-Leaf architecture
  - Designate Leaf nodes in a Flex Cluster as Reader Nodes
  - Updates from read-write instances are immediately propagated read-only instances on the Leaf Nodes
  - Up to 64 Reader Nodes per Hub Node

- Autonomous Health Framework
  - Cluster Resource Activity Log – contains resource activity information across the cluster. Used to track resource activity against resource behavior in the cluster, and to understand chain reactions triggered by the planned or unplanned relocation of individual resources
  - Cluster Health Advisor - Provides early warning of pending performance issues, root causes,& corrective actions for RAC databases and cluster nodes.   Use chactl command line interface
  - What-If Command Evaluation - crsctl eval delete server racnode2-3 -explain

RAC and Grid Infrastructure

- Load-Aware Resource Placement
  - Define database resources such as CPU and memory to Clusterware

```
$ srvctl modify database -db vnadb -cpucount 8 -memorytarget 32G
```

  - Clusterware places databases instances only on servers with sufficient number of CPUs, Memory or Both.
  - Prevents overloading a server with more database instances than the server is capable of running
  - Downside - Requires Instance Caging/DBRM to be enabled and use AMM ☹

RAC and Grid Infrastructure

- Extended Clusters
  - Build a Extended Cluster from the initial install/config or Convert afterwards
  - ConvertToExtend script converts an existing non-Flex ASM cluster to an extended cluster

  - On first node -> rootcrs.pl -converttoextended -first -sites *racnode1,racnode2* –site racnode1
  - On other nodes -> rootcrs.pl -converttoextended -site *racnode2*

VISCOSITY NORTH AMERICA

ASM
Flex Disk Groups and File Groups

- Flex Disk groups are the foundation to support new features such as File groups

- File group is a group of files that share the same set of properties and characteristics
  - This set of properties includes redundancy, rebalance rebuild priority, rebalance power limit, client compatibility, striping, quota group, and the access control list.

  - A major benefit is the ability to have a different availability specification for each database that shares the same disk group. File groups are also useful for point-in-time database clones.

ASM
Flex Disk Groups and File Groups

VISCOSITY NORTH AMERICA

ASM
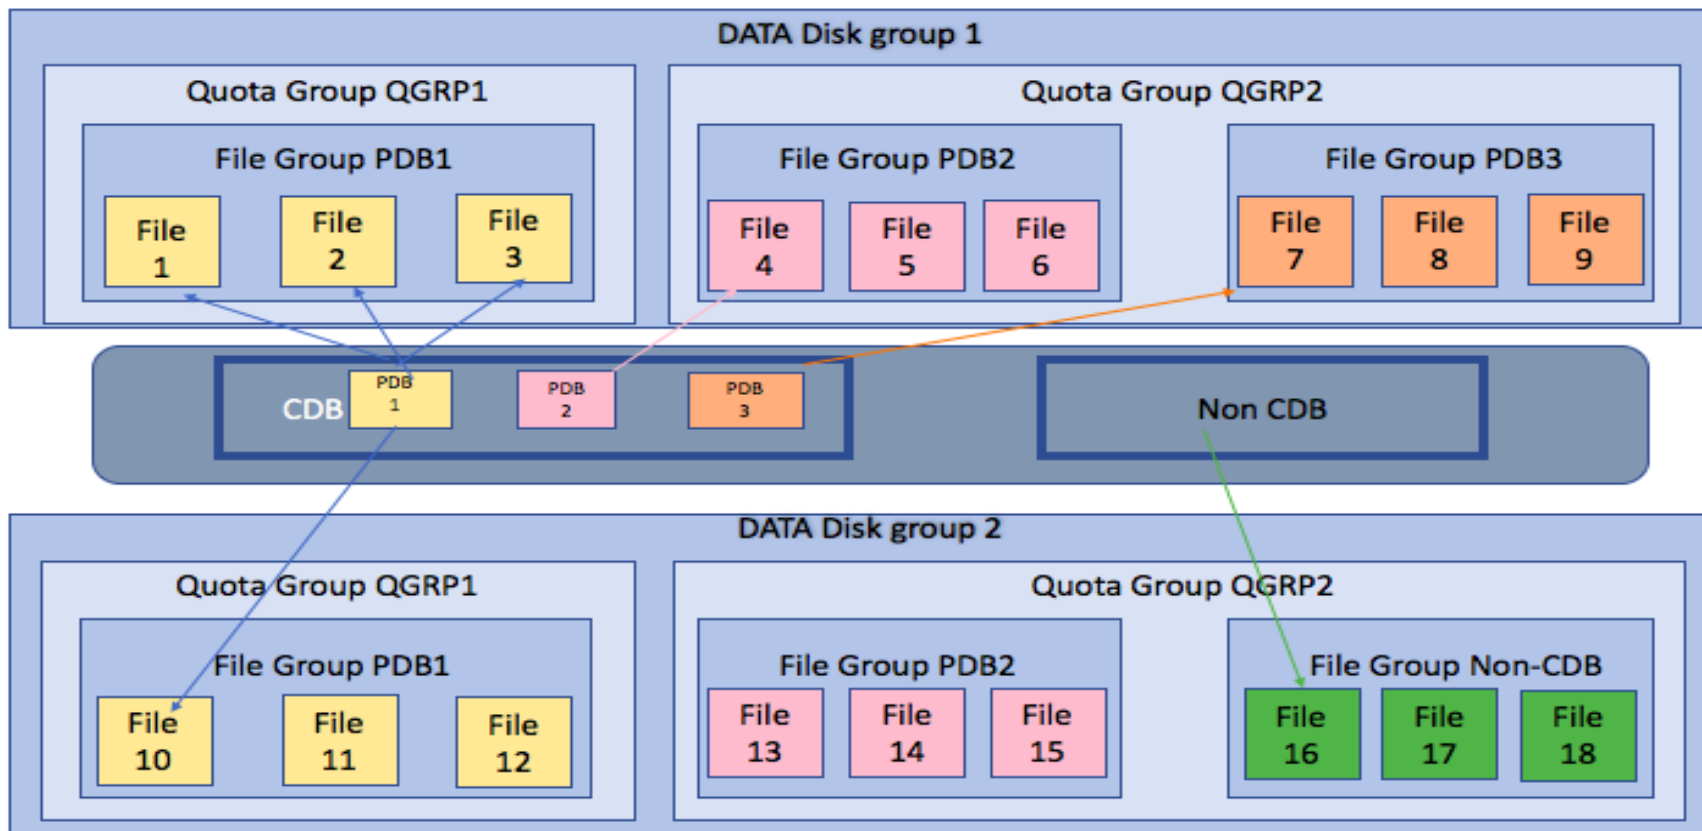
Flex Disk Groups and File Groups

- Disk group contains at least one file group, the default file group.
- Disk group can contain multiple file groups.
- Disk group must have FLEX redundancy to contain a file group.

- File group can belong to only one disk group.
- File group can describe only one database, PDB, CDB, volume, or cluster.
- File group can belong to only one quota group

- Database can span multiple disk groups with multiple file groups in different disk groups.

VISCOSITY NORTH AMERICA

ASM
Flex Disk Groups and File Groups

```
CREATE DISKGROUP data FLEX REDUNDANCY DISK '/dev/asmdata1','/dev/asmdata2','/dev/asmdata3'
ATTRIBUTE
'compatible.asm' = '12.2.0.2.0',
'compatible.rdbms' = '12.2.0.2.0',
'compatible.advm' = '12.2.0.2.0';
OR
SQL> ALTER DISKGROUP DATA MOUNT RESTRICTED;
SQL> ALTER DISKGROUP DATA CONVERT REDUNDANCY TO FLEX;


Created a CDB called NISHAN


$ASMCMD> lsfg
File Group          Disk Group   Quota Group   Used Quota MB   Client Name         Client Type
DEFAULT_FILEGROUP   DATA         GENERIC       0
NISHAN_CDB$ROOT     DATA         GENERIC       4624            NISHAN_CDB$ROOT     DATABASE
NISHAN_PDB$SEED     DATA         GENERIC       1552            NISHAN_PDB$SEED     DATABASE
```

ASM
Flex Disk Groups and File Groups

CREATE A PDB called OBIWAN. ASM automagically builds an OBIWAN file group.

```
ASMCMD> lsfg
File Group          Disk Group    Quota Group   Used Quota MB   Client Name        Client Type
DEFAULT_FILEGROUP   DATA          GENERIC       0
NISHAN_CDB$ROOT     DATA          GENERIC       4624            NISHAN_CDB$ROOT     DATABASE
NISHAN_PDB$SEED     DATA          GENERIC       1552            NISHAN_PDB$SEED     DATABASE
OBIWAN              DATA          GENERIC       1760            OBIWAN             DATABASE
```

Note, There's currently no QuotaGroup For Filegroup OBIWAN.

```
ASMCMD> lsqg
Group_Num   Quotagroup_Num   Quotagroup_Name   Incarnation   Used_Quota_MB   Quota_Limit_MB
2           1                GENERIC           1             7936            0
```

VISCOSITY NORTH AMERICA

ASM
Flex Disk Groups and File Groups

```
ALTER DISKGROUP DATA ADD QUOTAGROUP QGNishan SET quota = 2000m

ALTER DISKGROUP DATA MODIFY FILEGROUP OBIWAN SET 'quota_group' = 'QGNishan';

$ asmcmd lsfg
File Group          Disk Group   Quota Group   Used Quota MB   Client Name        Client Type
DEFAULT_FILEGROUP   VNADATA      GENERIC       0
NISHAN_CDB$ROOT     VNADATA      GENERIC       4624            NISHAN_CDB$ROOT    DATABASE
NISHAN_PDB$SEED     VNADATA      GENERIC       1552            NISHAN_PDB$SEED    DATABASE
OBIWAN              VNADATA      QGNISHAN      1760            OBIWAN             DATABASE

$ asmcmd lsqg
Group_Num   Quotagroup_Num   Quotagroup_Name   Incarnation   Used_Quota_MB   Quota_Limit_MB
2           1                GENERIC           1             6176            0
2           3                QGNISHAN          3             1760            2000
```

ASM
Flex Disk Groups and File Groups

```
ASMCMD> mkqg -G data Clonewars quota 10g

ASMCMD> lsqg
Group_Num  Quotagroup_Num  Quotagroup_Name  Incarnation  Used_Quota_MB  Quota_Limit_MB
2          1               GENERIC          1            11360          0
2          3               QGNISHAN         3            1760           100
2          4               CLONEWARS        1            0              10240
```

Create a new non-CDB called Hansolo

```
ASMCMD> lsfg
```

File Group	Disk Group	Quota Group	Used Quota MB	Client Name	Client Type
DEFAULT_FILEGROUP	DATA	GENERIC	0		
HANSOLO	**DATA**	**GENERIC**	**5384**	**HANSOLO**	**DATABASE**
NISHAN_CDB$ROOT	DATA	GENERIC	4680	NISHAN_CDB$ROOT	DATABASE
NISHAN_PDB$SEED	DATA	GENERIC	1552	NISHAN_PDB$SEED	DATABASE
OBIWAN	DATA	QGNISHAN	1760	OBIWAN	DATABASE

ASM
Flex Disk Groups and File Groups

```
ASMCMD> mvfg -G data --filegroup HANSOLO CLONEWARS
ASMCMD> lsfg
```

File Group	Disk Group	Quota Group	Used Quota MB	Client Name	Client Type
DEFAULT_FILEGROUP	DATA	GENERIC	0		
HANSOLO	DATA	CLONEWARS	5384	HANSOLO	DATABASE
NISHAN_CDB$ROOT	DATA	GENERIC	4680	NISHAN_CDB$ROOT	DATABASE
NISHAN_PDB$SEED	DATA	GENERIC	1552	NISHAN_PDB$SEED	DATABASE
OBIWAN	DATA	QGNISHAN	1760	OBIWAN	DATABASE

```
ASMCMD> lsqg
```

Group_Num	Quotagroup_Num	Quotagroup_Name	Incarnation	Used_Quota_MB	Quota_Limit_MB
1	1	GENERIC	1	6232	0

ASM
Flex Disk Groups and File Groups

```
ASMCMD> lsqg

Group_Num   Quotagroup_Num   Quotagroup_Name   Incarnation   Used_Quota_MB   Quota_Limit_MB
2           1                GENERIC           1             6232            0
2           2                NISHAN            1             0               6144
2           3                QGNISHAN          3             1760            100
2           4                CLONEWARS         1             5384            10240
```

@ViscosityNA

VISCOSITY NORTH AMERICA

ASM
Flex Disk Groups and File Groups

```
ASMCMD> lsfg -G DATA --filegroup obiwan
File Group    Disk Group    Property      Value     File Type
OBIWAN        DATA          PRIORITY      MEDIUM
OBIWAN        DATA          STRIPING      COARSE    CONTAINER
OBIWAN        DATA          STRIPING      FINE      CONTROLFILE
OBIWAN        DATA          REDUNDANCY    MIRROR    DATAFILE
OBIWAN        DATA          STRIPING      COARSE    DATAFILE
OBIWAN        DATA          REDUNDANCY    MIRROR    ONLINELOG
OBIWAN        DATA          STRIPING      COARSE    ONLINELOG
OBIWAN        DATA          REDUNDANCY    MIRROR    ARCHIVELOG
OBIWAN        DATA          STRIPING      COARSE    ARCHIVELOG
OBIWAN        DATA          REDUNDANCY    MIRROR    TEMPFILE
OBIWAN        DATA          STRIPING      COARSE    TEMPFILE
….
…..
z
```

ASM
Flex Disk Groups and File Groups

ASM Configuration Assistant: Disk Groups

Disk Groups

ORACLE 12c
GRID INFRASTRUCTURE

- ASM
  - ASM Instances
    - +ASM
  - Disk Groups
  - Vol
  - ACF
  - Settings

Disk Group Name	Size (GB)	Free (GB)	Usable (GB)	Redundancy	State
DATA	24.00	23.66	11.83	FLEX	MOUNTED

Add Disks...
Drop Disks...

Add Quota Group...
Drop Quota Group...
List Quota Groups
Move File Groups...
Add File Group...
Drop File Group...

Edit Attributes...
Manage Templates...
Create ACFS for Database Use...

Mount
Dismount
Drop
Drop and Clear Labels

Note: Use right click to see more options.

Create... | Mount All | Dismount All | Refresh

Help   Exit   Y NORTH AMERICA

ASM

- **Oracle IOServer (IOS)**
  - DB Member clusters do not have connectivity to ASM disks.
  - IOServer instance provides ASM file access for Oracle Database instances on Member Clusters nodes
  - On client side, database use dNFS to communicate with an IOServer instance.
  - The IOServer instance receives data packets from client & performs appropriate IO to ASM disks similar to any other database client.
  - No client side configuration -> discovery of the Oracle IOS instance occurs automatically

VISCOSITY NORTH AMERICA
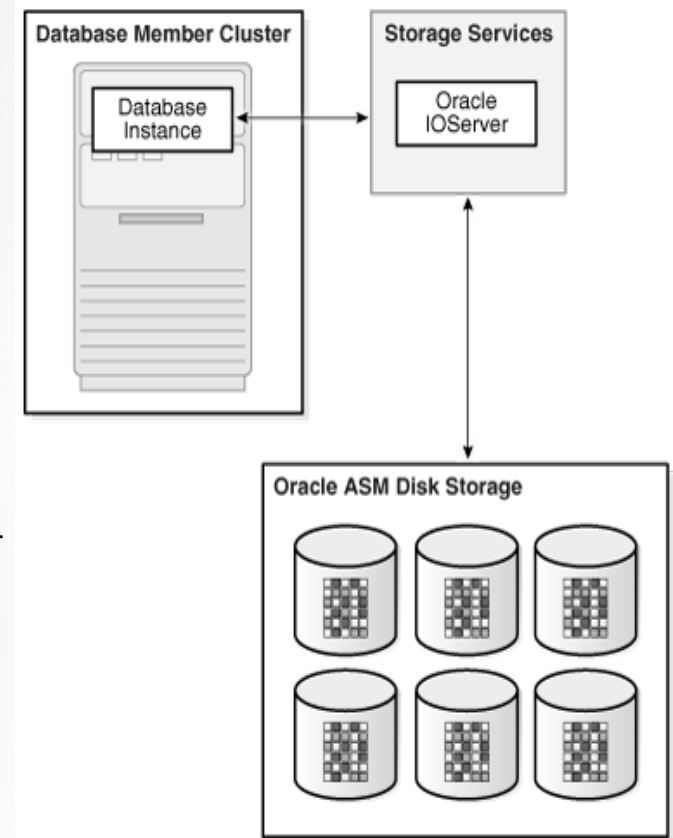
ASM
Misc Features

- **Filter Driver Installation and Configuration**
  - Installation and configuration for ASM Filter Driver (ASMFD) is enabled as part of Oracle Grid Infrastructure installation

- **Extended Support for 4K Sector Size**
  - Logical_sector_size defines the logical sector size (in bytes) of the disk group and specifies the smallest I/O that can be used

- **Support for Preferred Read on Extended Clusters**
  - The preferred read failure groups capability is automatically detected and set in the ASM instance when extended clusters are deployed.

ASM
Filter Driver

ASM

Filter Driver

ACFS #1

- **Snapshot Enhancements**
  - Snapshot based replication
  - Admins can impose quotas to snapshots to limit amount of write operations that can be done on a snapshot
  - Rename an existing ACFS snapshot, for more user-friendly names
  - ACFS snapshot remaster capability allows for a snapshot in the snapshot registry to become primary file system

- **4K Sectors**
  - If the COMPATIBLE.ADVM ASM Diskgroup attribute is set to 12.2 or greater, then the metadata block is 4096 bytes by default.
  - Use mkfs –i
  - If COMPATIBLE.ADVM attribute is set to less than 12.2, then the block size is set to 512 bytes.

VISCOSITY NORTH AMERICA

ACFS

- **Compression Enhancements**
  - ACFS compression can be enabled for specific ACFS file systems for database files, RMAN backup files, archivelogs, data pump extract files, and general purpose files.
    - Oracle does not support redo log/flashback logs/control file compression
  - When enabling ACFS compression for a file system, only new files will be compressed. Likewise, if you decide to uncompress a file system, Oracle will disable compression for new files.
  - Databases with 2k or 4k block sizes are not supported for ACFS compression.
  - ACFS compression is supported on Linux and AIX, as well as ACFS snapshot-based replication.

- **Defragger**
  - "acfsutil defrag dir" or "acfsutil defrag file" commands for on-demand defragmentation.
  - ACFS will perform all defrag operations in the background.
  - With the –r option of the "acfsutil defrag dir"command, you can recursively defrag subdirectories.

VISCOSITY NORTH AMERICA

ACFS

- **Loopback Devices**
  - We can now take OVM images, templates, and virtual disks and present them as a block device
  - Files can be sparse or non-sparse
  - ACFS also supports Direct I/O on sparse images

- **Metadata Accelerator**
  - Dramatically speeds up filesystem operations
  - Enables many critical ACFS metadata structures, including extent metadata, storage bitmaps, volume logs, and some snapshot metadata to be placed on accelerator storage
  - The accelerator volume can be created on Linux environments with the -a option of the mkfs command.
  - Recommended starting accelerator size is minimally 0.6% of the size of the file system

ACFS

- **Auto-Resize Enhancements**
  - The auto-resize feature, allows to "autoextend" a file system if the size of the file system is about to run out of space.
  - Just like an Oracle datafile that has the autoextend option enabled, we can now "autoextend" the ACFS file system to the size of the increment by option.
  - With the –a option to the "acfsutil size" command, we can specify the increment by size.
  - We can also specify the maximum size or quota for the ACFS file system to "autoextend" to guard against a runaway space consumption.
  - To set the maximum size for an ACFS file system, execute the "acfsutil size" command with the –x option.

VISCOSITY NORTH AMERICA

ACFS New Features – Complete List

- Oracle ACFS Snapshot-Based Replication
- Oracle ACFS Snapshot Enhancements
- Oracle ACFS Compression and Defragger
- Oracle ACFS Support for 4K Sectors
- Oracle ACFS Automatic Resize
- Oracle ACFS Metadata Acceleration
- Oracle ACFS NAS Maximum Availability eXtensions
- Oracle ACFS Sparse Files
- Oracle ACFS Scrubbing Functionality
- Oracle ACFS Loopback Functionality
- Oracle ACFS Diagnostic Commands

In-Memory

Basics

- New component in the SGA , in 12.1 was a static pool, and must be sized correctly using inmemory_size parameter
- Populated using 'alter table tab_name inmemory
  - Only for Tables, partitions, subpartitions, Mviews
  - Can Populate specific columns with priority
- Compressed by default using columnar compression – memcompress for query, dml, or capacity
- User_tables list which tables are marked IMMEMORY
- V$IM_SEGMENTS list state of segement
- Storage Indexes allow pruning

In-Memory Virtual Columns – 12cR2

Oracle Database 12c
Release 2 Performance
Tuning Tips and Techniques

To put the table INMEMORY (in the main IM area IMCU):

alter table scott.emp_rich INMEMORY; (**virtual column IM if below parameter set**)

The following initialization parameter must be set (can set when DB running):

INMEMORY_VIRTUAL_COLUMNS=ENABLE (set to DISABLE to turn it off)

To specifically put virtual column INMEMORY (a **separate area of IM – IMEU**):

alter table scott.emp_rich INMEMORY(yearly_sal);

IMCU=*In-Memory Compression Unit*; IMEU=*In-Memory Expression Unit*

In-Memory

Expressions, Virtual Columns, Fast Start, and ADO

- IM column store dynamic resizing - dynamically increase the size of the In-Memory Area without reopening the database

- IM expressions - Speeds queries of large data sets by precomputing computationally intensive expressions. Beneficial for frequently executed table joins, projections, and predicate evaluations.
  - Improved efficient than Mviews
  - Expression Statistics Store (ESS) automatically identifies frequently used ("hot") expressions that are candidates for population in IM column store
  - For example, aliased weekly_sal and ann_comp:
    ```
    SELECT employee_id, last_name, salary, commission_pct, ROUND(salary*12/52,2) as
    "weekly_sal", 12*(salary*NVL(commission_pct,0)+salary) as "ann_comp" FROM
    employees ORDER BY ann_comp;
    ```

  - DBMS_INMEMORY_ADMIN package is the primary interface for managing IM expressions: config, capture and drop IME

In-Memory

Expressions, Virtual Columns, Fast Start, and ADO

- <u>IM virtual columns</u> enable the IM column store to materialize some or all virtual columns in a table
  - IM virtual columns are user-created and exposed, whereas IM expressions are database-created and hidden
  - Like any other column in a table enabled for INMEMORY, except that its value is derived by evaluating an expression
  - `CREATE TABLE hr.admin_emp ( empno NUMBER(5) PRIMARY KEY, ename VARCHAR2(15) NOT NULL, job VARCHAR2(10), sal NUMBER(7,2),` **`hrly_rate NUMBER(7,2) GENERATED ALWAYS AS (sal/2080),`** `deptno NUMBER(3) NOT NULL) INMEMORY;`
  - `ALTER TABLE hr.admin_emp INMEMORY(hrly_rate);`
  - Storing the pre-calculated IM virtual column values in the IM column store improves query performance.

- <u>IM column store as a new ADO data tier.</u>
  - When enabled, the Heat Map feature automatically tracks data access patterns; ADO uses this Heat Map data to implement user-defined policies at the database level.
  - ADO manages the In-Memory Column Store, by moving objects (tables, partitions or subpartitions) in and out of the memory, based on Heat Map statistics.

In-Memory

Expressions, Virtual Columns, Fast Start, and ADO

- IM FastStart optimizes the population of database objects in the IM column store by storing IMCUs directly on disk.
  - FastStart area is a designated tablespace where IM FastStart stores and manages data for INMEMORY objects.
  - Database manages the FastStart tablespace without DBA intervention

- Object-level support for RAC Services DB object, the INMEMORY ... DISTRIBUTE clause has a FOR SERVICE subclause that limits population to the database instance where this service is allowed to run.
  - Configure an INMEMORY object to be populated in the IM column store on instance 1 only, or on instance 2 only, or in both instances.
  - ALTER TABLE sales INMEMORY DISTRIBUTE FOR SERVICE "Salessrv";

- IM column store on a standby database
  - Enable an IM column store in an Oracle Active Data Guard standby database
  - Populate a completely different set of data in the in-memory column store on the primary and standby databases

SQL Performance Tuning

- **AWR at PDB Level**
  - Earlier AWR reports could be created only at CDB level. 12.2 reports can now be generated at PDB level
  - Snapshot interval & retention can be controlled at PDB level
  - To check AWR information, at ROOT & PDB level there are two new views
    - AWR_ROOT%
    - AWR_PDB%
- **AWR for Active Data Guard**
  - Primary database or any other remote database can host snapshot repository
  - AWR tables on catalog database gather information on the snapshot of ADG database using database links

- **AWR for Active Data Guard Cont.**
  - Snapshots can be taken manually or automated

- **Optimizer Statistics Advisor**
  - New Optimizer Statistics Advisor, AUTO_STATS_ADVISOR_TASK, aids in diagnosing problems with existing stats gathering practices
  - Like other advisors, it runs in the same maintenance window & recommend changes.
  - Reports are in a user-friendly format, similar to ADDM reports

VISCOSITY NORTH AMERICA

SQL Performance Tuning #2

- **Per-Process PGA Limits**
  - Ability to restrict the PGA usage at session level, in a particular consumer group
    - SESSION_PGA_LIMIT and DBMS_RESOURCE-MANAGER.CREATE_PLAN_DIRECTIVE

- **Index Monitoring**
  - Turned on by default
  - Tracks usage at execution level rather than parse level
  - V$INDEX_USAGE_INFO keeps track of index usage since last flush
  - DBA_INDEX_USAGE historically tracks and stores when the index was last used, total number of times index was accessed, and much more.

- **Adaptive Plans**
  - The latest update has made Optimizer_Adaptive_Features parameter obsolete
  - Plan adaptive features are controlled by two new parameters:
    - Optimizer_Adaptive_Plans : controls whether the optimizer creates adaptive plans and defaults to TRUE.
    - Optimizer_Adaptive_Statistics : controls whether the optimizer uses adaptive statistics and defaults to FALSE. It includes generations of SQL plan directives, dynamic sampling, automatic re-optimization, and much more.

Security

Zero Downtime Encryption and Decryption

- TDE encryption now available for all tablespaces; including SYSTEM, SYSAUX, and UNDO

- Extend encryption and hashing algorithms to include ARIA, GOST, and SEED Encryption Algorithms for International Standards.

- Perform offline conversion of a tablespace, without additional storage overhead.

- To encrypt an existing tablespace online, you must login to the database with the SYSKM role. To encrypt the SYSTEM or SYSAUX tablespace, you must login with the SYSDBA role.
  - `SQL> ALTER TABLESPACE `**`sysaux`**` ENCRYPTION ONLINE USING 'AES256' ENCRYPT;`

- Decrypt a tablespace online without any downtime
  - `SQL> ALTER TABLESPACE sysaux ENCRYPTION ONLINE DECRYPT;`

Security

- 12.2 RAC introduces an administrative privilege called SYSRAC.
  - The SYSRAC privilege is for RAC, whereas the SYSDG privilege is for Data Guard. The SYSRAC privilege reduces the reliance on SYSDBA on production RAC systems.

- Database Vault:
  - Oracle provides the ability to group multiple Database Vault realms and protections into a Single Policy for simple deployment and integration

- Application security in the database:
  - Enhancement to Real Application Security (RAS Session, Column, Schema, and Integration Privilege), will limit users to execute privileged user commands.
  - The RAS policy can be enforced to each schema, versus the entire database as before

- Privilege Analysis:
  - We can review and compare access across PL/SQL packages and Java Oracle programs.

Oracle Database 12c R2 — New Features for Administrators

SQL*Plus

History

SQL> help hist
HIST[ORY] [N {RUN | EDIT | DEL[ETE]}] | [CLEAR]

SQL> show hist

SQL> hist 2 run
SQL> hist 3 del

SQL> hist clear

```
SQL> hist
  1  select name from v$database;
  2  select instance_name, host_name from v$instance;
  3  help hist

SQL> hist 3 del
SQL> hist 1 run

NAME
---------
MERIT

SQL> hist 1 edit
```

```
SQL> select name from v$database;

NAME
---------
MERIT

SQL> select instance_name, host_name from v$instance;

INSTANCE_NAME
----------------
HOST_NAME
--------------------------------------------------
merit
ika82

SQL> hist
  1  select name from v$database;
  2  select instance_name, host_name from v$instance;
```

- set hist on

- SQL> set hist 1000
- SQL> show hist
  history is ON and set to "1000"

Our History

VISCOSITY NORTH AMERICA

HUGE Improvements – 12cR2

- In 12c, **object names for users, roles, tables, columns, indexes, constraints, etc. have been increased from 30 bytes to 128 bytes** with a few limitations.

- The limit for **tablespace names and pluggable databases is still 30 bytes**, but others all increase to 128 bytes.

- You will **notice this change in the dictionary views** where the **VARCHAR2 columns will shows as 128 bytes** instead of 30 bytes.

- It also **helps in migrations** from non-Oracle systems where the name is longer than 30 characters.

- The best enhancement in **12c R2** is **32K VARCHAR is default**. This allows the <u>extending of the VARCHAR data types without having to enable the extended mode specifically (early 12c)</u>. The size limit for both **VARCHAR2 and NVARCHAR2 is 32K**.

VISCOSITY NORTH AMERICA

Long Name Identifiers

- Table Name to 128 bytes
- Column Name to 128 bytes

Long Table Name

SQL> create table rmoug_oracle_dbas_are_super_cool_in_denver_colorado_in_2017_tab
(oracle_database_12c_release_2_will_be_available_for_general_availability_column_1 integer);

Table created.

Long Column Name

@ViscosityNA

VISCOSITY NORTH AMERICA

Approximate Query – 12cR2

- **Approximate Query Processing**, counts distinct values and adds approximate percentile aggregation.
- This allows **faster processing of large data sets using approximation** instead of exact aggregation.
- Since this is an aggregation it is **not assured to be *completely* accurate**, however, in most cases it is very close and acceptable considering the large performance boost it provides.
- Note that the **results other than approximated value returned are 100% accurate**, it is only how the query is processed and the amounts that are approximated (for instance the departments with approximately $1M in sales will give the correct departments that are **within 97%** of $1M in sales **with 95% accuracy** … but, NOT give a department with only $100 in sales).

VISCOSITY NORTH AMERICA

Approximate Query: 100x+ Faster
(Depending on the Query)

Instead of (100% accurate – 12.1.0.2):
select count(distinct(empno)) from emp;

Use this for speed (97% accurate – 12.1.0.2):
select **approx_count_distinct**(empno) from emp;
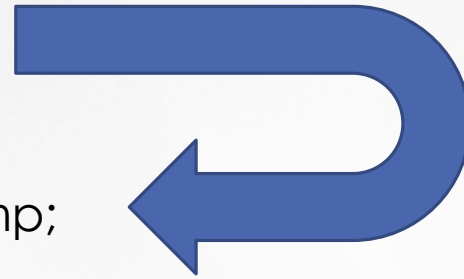
In 12c R2, additional approximate percentile functions & capability to reuse approximate aggregations for multiple queries (via materialized views and query rewrite).

In 12c R2, you can toggle to this mode:
alter session set approx_for_count_distinct = TRUE; (12.2 only – distinct counts **all** approximate)

- Approximate amount within 97% & Explain Plan: SORT GROUP BY to **SORT AGGREGATE APPROX**

Approximate Query – 12cR2 Only

Other initialization parameters:

approx_for_aggregation=TRUE

approx_for_percentile=TRUE

Other Approximate Functions:

- APPROX_COUNT_DISTINCT_DETAIL
- APPROX_COUNT_DISTINCT_AGG
- TO_APPROX_COUNT_DISTINCT
- APPROX_MEDIAN
- APPROX_PERCENTILE
- APPROX_PERCENTILE_DETAIL
- APPROX_PERCENTILE_AGG
- TO_APPROX_PERCENTILE

Also in 12cR2 is support for Materialized Views and Query Rewrite

Scheduler

Anti-Affinity Rules
– Specify that a scheduled job cannot run at the same time with another scheduled job
– Leveraging the dbms_scheduler.create_incompatibility procedure
– Add a job or program to an existing compatibility definition with dbms_scheduler.add_to_incompatibility proc
– Remove a job or program from an existing incompatibility definition with the
   dbms_scheduler.remove_from_incompatibility

Resource Queues
– Define a set of resources for a scheduled job
– A resource, can be anything that is defined with two attributes: name and count with the
   DBMS_SCHEDULER.CREATE_RESOURCE procedure
– Ensure that the scheduled job will not exceed the resources specified for the job definition.

In-Memory Job
– Create two types of in-memory jobs: repeating (in-memory runtime jobs) and one-time jobs (in-memory full jobs) in-memory jobs
– With repeating in-memory jobs, only the job metadata is written to disk

Data Pump

Parallel Export and Import of Metadata

Functionality has been extended to include metadata for export and import operations

Rename Data Files During TTS Import

* Symbol matches multiple characters and question mark (?) matches a single character

- In Oracle 12.2, use wildcards to simplify our TTS process and lower our management overhead.
- TRANSPORT_DATAFILES=vna_data*.dbf
- TRANSPORT_DATAFILES=vna_index?.dbf

Additional Wildcard Substitution Variables

- New options in 12.2 expdp for date and time are:
  - %d or %D to specify Day of Month in DD format
  - %m or %M to specify number of month in MM format
  - %y or %Y to specify year in YYYY format
  - %t or %T to specify full date in YYYYMMDD format

VISCOSITY NORTH AMERICA

Data Pump #2

Additional DATA_OPTIONS Parameters
– DATA_OPTIONS=VALIDATE_TABLE_DATA for import - validates date & number formats of table data
  • Default behavior for Data Pump is to perform no validation
  • Want to use this option when the originator of the data is not from a trusted source.

Direct Load in Network Mode
– Data pump now leverages OCIDirPathUnload to unload and load data from the source database for Long and Long Raw data types
– Specify ACCESS_METHOD=DIRECT_PATH parameter with the NETWORK_LINK parameter

Instant Client
– Data pump is now available in the Instance Client.
– Also added SQL*Loader, expdp, impdp, exp, and imp to the tools for instant client

Interactive Commands
– Enable tracing for a running job; do not need to stop/restart the job for the tracing to take effect
– Can kill individual works that we believe to be hung or stuck, with the STOP_WORKER option

Materialized View

Historical Statistics

- DBA_MVREF_STATS –

Stores basic statistics for a materialized view refresh such as the refresh ID and **basic timing statistics** for the refresh operation.

- DBA_MVREF_RUN_STATS –

Stores **detailed** information about each materialized view refresh operation including the following

- parameters specified when running the refresh operation such as list of materialized views, refresh method, purge option, and so on.
- number of materialized views refreshed in the refresh operation.
- detailed timing statistics for the refresh operation including start time, end time, and elapsed time.

DBA_MVREF_CHANGE_STATS –
Contains **change data load information** for the base tables associated with a materialized view refresh operation.

The details include base table names, materialized view names, number of rows inserted, number of rows updated, number of rows deleted, number of direct-load inserts, PMOPs details, and number of rows at the beginning of the refresh operation.

DBA_MVREF_STMT_STATS –
Contains information related to **each refresh statement** that is part of a single materialized view refresh operation.

This includes information such as materialized view name, refresh ID, the refresh statement, SQLID of the refresh statement, and execution plan of the statement.

@ViscosityNA

VISCOSITY NORTH AMERICA

Sharding

Introductory Primer

- Sharding is a data tier architecture in which data is horizontally partitioned across independent databases.

- Horizontal partitioning involves splitting a database table across shards so that each shard contains the table with the same columns but a different subset of rows. A table split up in this manner is also known as a sharded table.

- Oracle 12.2 supports up to 1000 Shards

Unsharded Table in One Database

Server

Sharded Table in Three Databases

Server A   Server B   Server C
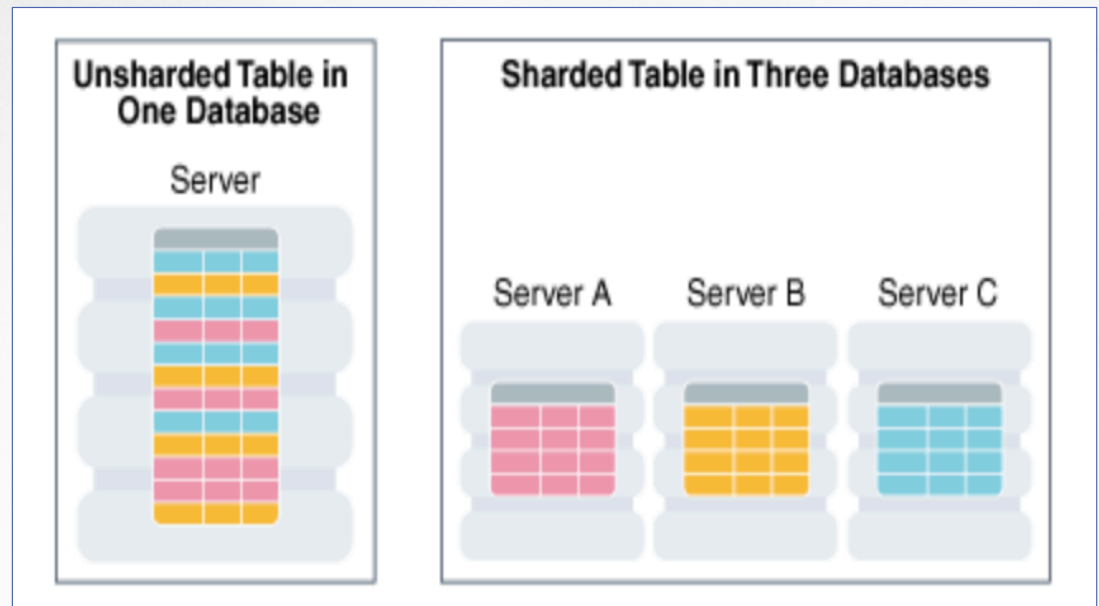
Sharding

Introductory Primer

- Shards (Sharded Table): Table that is partitioned into smaller and more manageable pieces among multiple databases

- Sharded Table Family: A sharded table family is a set of tables that are sharded in the same way.
  - The recommended way to create a sharded table family is to specify parent-child relationships between tables using reference partitioning.

- Duplicated Tables: In addition to sharded tables, an SDB can contain tables that are duplicated on all shards.

Partitioning
Support for VLDB Databases

- Convert Non-Partitioned Table to a Partitioned Table

- Read-Only Partitions

- Multi-Column List Partition

- Split Partition with Online Maintenance

- Create a Partitioned External Table
  - Support to map partitioned Hive tables into the Oracle Database ecosystem as well as providing declarative partitioning on top of any Hadoop Distributed File System (HDFS) based data store.
  - External Tables Can Access Data Stored in Hadoop Data Sources Including HDFS and Hive

VISCOSITY NORTH AMERICA

Partitioning

Convert Non-Partitioned Table To A Partitioned Table Example

- ALTER TABLE mytable MODIFY
  PARTITION BY RANGE (tab_id) INTERVAL (25)
  (PARTITION p1 VALUES LESS THAN (100),
  PARTITION p2 VALUES LESS THAN (1000))
  ONLINE
  UPDATE INDEXES
  (IDX1_tab_id LOCAL (PARTITION ip1 VALUES LESS THAN (MAXVALUE)));

- With **ONLINE** option, zero downtime conversion to Partitioning Strategy

VISCOSITY NORTH AMERICA

Partitioning

Read-Only Partitions Example

- CREATE TABLE mytable (
tab_id NUMBER(10) NOT NULL,
tab_desc VARCHAR2(20))
READ WRITE PARTITION BY RANGE (tab_id)
(PARTITION mytable_p1 VALUES LESS THAN (100),
PARTITION mytable_p2 VALUES LESS THAN (1000) READ ONLY)
);

- Old Method: Make tablespace read-only

- Implications: Simplified Backups, Tiering of Data

Partitioning

Multi-Column List Partitions & Split Partitions Example

- CREATE TABLE mytable (
  tab_id NUMBER(10) NOT NULL,
  tab_desc VARCHAR2(20),
  tab_code VARCHAR2(2),
  tab_channel NUMBER(2))
  PARTITION BY LIST (tab_code, tab_channel)
  (PARTITION mytable_p1 VALUES (('AB',10), ('AC',20)),
  PARTITION mytable_p2 VALUES (DEFAULT));

- ALTER TABLE mytable
  SPLIT PARTITION p1 at (100)
  INTO (PARTITION p1_1,
  PARTITION p1_2);

- Can Split Sub-Partitions too

@ViscosityNA

VISCOSITY NORTH AMERICA

Partitioning

Misc New Features

- **Automatic List Partitioning**
  - The automatic list partitioning method enables list partition creation on demand.

- **Deferred Segment Creation for Automatic List Partitions and Interval Subpartitions**
  - Saving disk space when partition and subpartition creation is in the presence of data.

- **Creating a Table for Exchange with a Partitioned Table**
  - Tables can be created with FOR EXCHANGE WITH clause, to exactly match the definition of the partitioned table for partition exchange command.

Online Table Move

Zero DownTime

- ALTER TABLE ... MOVE ... ONLINE

- ALTER TABLE ... MOVE PARTITION ... ONLINE

- ALTER TABLE ... MOVE SUBPARTITION ... ONLINE

- ALTER TABLE consultant_details
  MOVE
  ONLINE
  COMPRESS
  TABLESPACE data_ts1
    UPDATE INDEXES (idx1 TABLESPACE index_ts1,
                    idx2 TABLESPACE index_ts2);

- Table move operation now also supports automatic index maintenance as part of the move.

- With UPDATE INDEXES clause, the indexes remain usable during the move operation

VISCOSITY NORTH AMERICA

Indexes

- Advanced Index Compression (AIC) has been added to this release, allowing for high compression; versus the low index compression in prior releases.
  - With the introduction to AIC in Oracle 12.2, bitmapped indexes and index-organized tables are not supported
  - At the gain of higher compression, we sacrifice additional CPU penalty.
- Starting with this release, Oracle provides support hash and list partitions for Spatial indexes.

- High Index compression implements new compression algorithms, which allow for greater compression ratios.
- To create a high level advanced compressed index:
  - 1 CREATE INDEX consultant_idx ON consultants(consultant_id)

    2* COMPRESS ADVANCED HIGH

    SQL> /
- We can also rebuild the index and change the compression from LOW to HIGH

Indexes

Index Tracking and XMLIndex Partitioning

- New index tracking has been added that provides information about the use of indexes in the database.
  - The V$INDEX_USAGE_INFO tracks index usage since last flush.
- A flush occurs every 15 minutes and is recorded to the V$INDEX_USAGE_INFO view.
  - The INDEX_STATS_COLLECTION_TYPE column, represents whether the type of collection was sampled or comprehensive (ALL) statistics; these are collected for each access of the index.

- The DBA_INDEX_USAGE view, holds cumulative statistics for each index in the database.
  - Usage data includes total access count, executions involving the index, and when the index was last used.
- Oracle 12.2 now provides partitioning support for XMLIndex, allowing for more scalability during parallel operations.
  - Paralyzing XML operations, allows for increased performance by taking advantage of parallel hardware configurations such as Oracle Exadata.

"Disaster strikes without warning.
> Data loss strikes without warning."

VISCOSITY NORTH AMERICA

Evolution of Data Guard

Oracle Version	Big Feature Added
Oracle 7.3	Hot Standby
Oracle 8i	Read-Only Physical Standby Managed Recovery
Oracle 9i	Logical Standby, Switchover/Failover Automatic Gap Resolution Cascaded Redo
Oracle 10g	Real-time apply Standby Redo Logs FSFO Flashback Database on DG
Oracle 11g	Active Data Guard Redo Compression Snapshot Standby Rolling Transient Upgrades RMAN Duplicate from Active Database
Oracle 12.1.0.1	Far Sync Fast Sync (LOG_ARCHIVE_DEST_n attributes SYNC and NOAFFIRM) Data Guard Broker Switchover Enhancements SYSDG Role Online Standby Database File Rename Significantly Improved Rolling Upgrades (DBMS_ROLLING)

Data Guard

dbca

- Create standby databases with DBCA
  - Must use command line dbca interface
  - Can only create from Non Multi-tenant primary database
  - Cannot do RAC
  - `dbca -createDuplicateDB`
    `-gdbName global_database_name`
    `-primaryDBConnectionString`
    `easy_connect_string_to_primary`
    `-sid database_system_identifier`
    `[-createAsStandby`
    `[-dbUniqueName`
    `db_unique_name_for_standby]]`
    `[-customScripts scripts_list]`

- `-createAsStandby` followed by dbUniqueName For Standby

- `[-customScripts scripts_list]`
  - List of scripts to be executed
  - Oracle will execute the scripts in the order they are listed

- No Need for TNSNAMES.ORA entry
- Can use easy connect string
- `"host[:port][/service_name][:server][/instance_name]"`

Data Guard

dbca

```
$ dbca -silent -createDuplicateDB \
-gdbName viscorcl \
-primaryDBConnectionString rac01:1521/viscorcl \
-sid viscorcl -createAsStandby \
-dbUniqueName sviscorc
Enter SYS user password:
Listener config step
33% complete
Auxiliary instance creation 66% complete
RMAN duplicate 100% complete
Look at the log file "/u01/app/oracle/cfgtoollogs/dbca/sviscorc/viscorc.log" for further
details.
```

Data Guard

Support for Diagnostic Pack and SQL Tuning Advisor

Diagnostic Pack

- ADG in read-only mode can take advantage of diagnostic pack
- Can take advantage of ADDM from the AWR Repository
- AWR Snapshots on ADG are called remote snapshots
- Database node is called a destination
- Work off of database links
- We can create remote snapshots manually or at automatically at a scheduled interval and are started by a destination node.

SQL Tuning Advisor

- DBAs can initiate the SQL Tuning Advisor on the primary database
- SQL tuning process workload is executed on the Active Data Guard
- Can offload the tuning workload of the primary database to the Active Data Guard standby database
- Results of the SQL tuning efforts are written back to the primary database over database links
- SQL profile recommendations on the primary database are applied to the Active Data Guard

Data Guard

NoLogging & MIRA

NoLogging on Primary

- Nologging blocks are recorded in the control file on the physical standby database
- On the Standby
  `RMAN RECOVER DATABASE NONLOGGED BLOCK`
- After switchover, the old primary database must be in a mounted state

Multiple Instance Redo Apply (MIRA)

- Can run Redo Apply in some or all of the standby instances
- Redo Apply scales as wide as standby RAC configuration
- ALTER DATABASE RECOVER MANAGED STANDBY DATABASE command now accepts a new INSTANCES [ALL | integer] clause
- All the instances must be in the same mounted or open mode
- cannot specify which RAC instance(s) will perform the redo apply

Data Guard

Multiple Observers

```
SQL> desc V$FS_FAILOVER_OBSERVERS
Name                                      Null?       Type
----------------------------------------- --------    ------------------------
NAME                                                  VARCHAR2(513)
REGISTERED                                           VARCHAR2(4)
HOST                                                 VARCHAR2(513)
ISMASTER                                             VARCHAR2(4)
TIME_SELECTED                                        TIMESTAMP(9)
PINGING_PRIMARY                                      VARCHAR2(4)
PINGING_TARGET                                       VARCHAR2(4)
CON_ID                                               NUMBER
```

- Up to 3 Observers to monitor and support a single Oracle Data Guard Broker configuration
- Each Observer must be assigned a name
  - Case sensitive
  - Unique in configuration
  - The "start observer" command is enhanced to accept a name
- Concept of Master Observer – has privilege of coordinating the FSFO with the Data Guard broker
- Other registered observers serve the role of backup observers until the master observer is not available
- Considerations
  - Never place the observers on the same server or the same virtual server
  - Place one or more of the observers in separate data centers

VISCOSITY NORTH AMERICA

Data Guard

Subset Standbys

- In 12.1, ENABLED_PDBS_ON_STANDBY initialization parameter only supported two values: all PDBs or none
- The ENABLED_PDBS_ON_STANDBY parameter is only applicable on the physical standby database
  - can accept a list of PDB names or a glob pattern such as "VNAPDB?" or "VNAPDB*a", "VNAPDB2"
  - Glob pattern rules are similar to regular expression rules in common UNIX shells
  - asterisk (*) and question mark (?) wildcard characters are supported.
    - The question mark (?) represents a single unknown character;
      - the asterisk (*) represents matches to any number of unknown characters.
      - This parameter also accepts a minus sign (-) which can be used as the first character in a PDB name to designate that the PDB should be excluded on the standby database

VISCOSITY NORTH AMERICA

Data Guard

Subset Standbys ENABLED_PDBS_ON_STANDBY Examples

- ENABLED_PDBS_ON_STANDBY="*" means that all PDBs will be created on physical standby.
- ENABLED_PDBS_ON_STANDBY="VNAPDB1*" means that VNAPDB1A, VNAPDB1B, and VNAPDB1C will be created on the physical standby.
- ENABLED_PDBS_ON_STANDBY="VNAPDB*A" means that VNAPDB1A, VNAPDB2A, and VNAPDB3A will be created on physical standby.
- ENABLED_PDBS_ON_STANDBY="VNAPDB1*", "-VNAPDB*A" means that VNAPDB1B and VNAPDB1C will be created on physical standby but VNAPDB1A will be excluded

Oracle Multitenant

12ᶜ
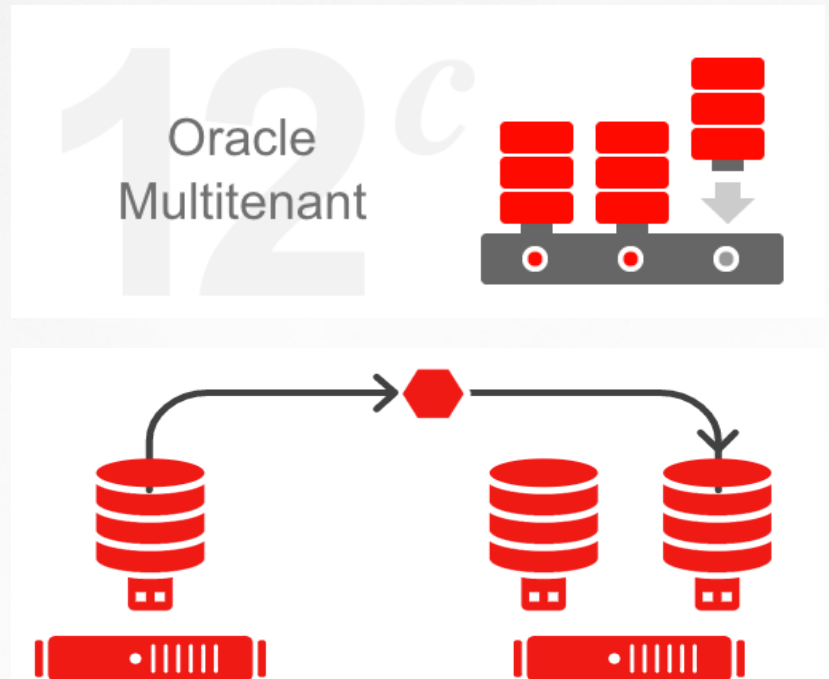
Data Guard

Broker Support for ADG Rollling Upgrade & Alternate Destinations

ADG Rolling Upgrade

- In 12.1, Oracle introduced rolling upgrades for ADG

- In 12.2, The broker now supports rolling upgrades
  - Broker no longer needs to be disabled

- During the rolling upgrade process, the "show configuration" command from DG Broker will report the status of "ROLLING DATABASE MAINTENANCE IN PROGRESS".

@ViscosityNA

VISCOSITY NORTH AMERICA

Data Guard

Database Block Compare

- DBMS_DBCOMP.DBCOMP, is introduced to detect lost writes and to identify inconsistencies between the primary and physical standby database
- Does not require the DB_LOST_WRITE_PROTECT
  - dbverify utility cannot detect for lost write disk errors
  - Monitor the progress by querying the V$SESSION_LONGOPS
- Executed on the primary or on the physical standby with the database in MOUNT or OPEN mode

```
DECLARE
  DataFile VARCHAR2(1000);
  OutputFile VARCHAR2(1000);
BEGIN
  DataFile := 'all' ;

OutputFile:='BlockCompareFULL_';
  SYS.DBMS_COMP.DBCOMP(DataFile,
OutputFile, true);
END;
/
```

Data Guard

Database Block Compare Broker Integration & Output

Output

- The datafile can be a number, name of the datafile or ALL for all the datafiles
- The output file is a prefix in the name of the output file
- All output is stored in the $ORACLE_HOME/dbs directory
  - Can be modified with either relative or absolute path
  - The block dump parameter is a Boolean parameter. By default, this parameter is FALSE
  - We can set this Boolean parameter to TRUE if we want the content of the block to be dumped into the output file when a pair of blocks between the primary and standby databases is not the same.

Broker Integration

VALIDATE DATABASE [database-name | ALL] DATAFILE [datafile-name | datafile-number | ALL] OUTPUT="output-file-name";

- Can validate the database at the datafile level or at the database level
- The output file is generated in the trace directory
- Can specify the datafile name or the datafile number or the ALL option

```
DGMGRL> VALIDATE DATABASE
[database-name  |  ALL] DATAFILE
[datafile-name  | datafile-number |
ALL] OUTPUT="output-file-name";
```

Data Guard

Validate Database

```
DGMGRL> validate database svisorc;

  Database Role:    Primary database

  Ready for Switchover:  Yes

  Flashback Database Status:
    svisorc:  Off

  Managed by Clusterware:
    svisorc:  NO
    Warning: Ensure primary database's StaticConnectIdentifier property
    is configured properly so that the primary database can be restarted
    by DGMGRL after switchover

DGMGRL>
```

VISCOSITY NORTH AMERICA

Data Guard

Validate Database Verbose

```
DGMGRL> validate database verbose svisorc;

  Database Role:    Primary database

  Ready for Switchover:  Yes

  Flashback Database Status:
    svisorc:  Off

  Capacity Information:
    Database    Instances        Threads
    svisorc   1                1

  Managed by Clusterware:
    svisorc:  NO
    Warning: Ensure primary database's StaticConnectIdentifier property
    is configured properly so that the primary database can be restarted
    by DGMGRL after switchover
```

```
Temporary Tablespace File Information:
    svisorc TEMP Files:  1

  Data file Online Move in Progress:
    svisorc:  No

  Transport-Related Information:
    Transport On:  Yes

  Log Files Cleared:
    svisorc Standby Redo Log Files:  Cleared

  Automatic Diagnostic Repository Errors:
    Error                          svisorc
    No logging operation           NO
    Control file corruptions       NO
    System data file missing       NO
    System data file corrupted     NO
    System data file offline       NO
    User data file missing         NO
    User data file corrupted       NO
    User data file offline         NO
    Block Corruptions found        NO
```

VISCOSITY NORTH AMERICA

Data Guard

Password File Synchronization

- When the password for SYS, SYSDBA, SYSOPER and/or SYSDG are modified, password file is replicated to the password files to all the standby database ORACLE_HOME
- The password file is updated on the physical standby database server when redo is applied
- For far sync configurations, the changes to the password file must be manually copied
- When the password file is updated on the far sync server, the password will automatically be propagated to the target physical standby database servers

VISCOSITY NORTH AMERICA

Data Guard

In-Memory with Active Data Guard

In-Memory with Active Data Guard

- IM option can be configured on the primary database, on an ADG standby database, or on both the primary and the ADG standby databases

- The INMEMORY_ADG_ENABLED parameter needs to be enabled
  - By default, the INMEMORY_ADG_ENABLED parameter is set to true.

- For RAC configuration, this parameter must be set to the same value across all the RAC instances

"Queries run 100 times faster and updates double in speed, simply by throwing a switch."

Larry Ellison
CEO, Oracle
#DB12c

ORACLE
DATABASE IN-MEMORY

No Application Changes

@ViscosityNA

VISCOSITY NORTH AMERICA

Data Guard

Minimize Impact to Multiple SYNC Standby Databases
FSFO in Maximum Protection Mode

- DATA_GUARD_SYNC_LATENCY specifies how long the primary database Log Writer (LGWR) should wait for a response from multiple synchronous standby databases during redo transport.

- DATA_GUARD_SYNC_LATENCY parameter defines the number of seconds that the primary database must wait; once one of the synchronous standby databases acknowledges receipts of the redo

FSFO in Maximum Protection Mode

- Oracle Data Guard supports FSFO in Maximum protection mode for zero data loss protection when multiple synchronous destinations exist
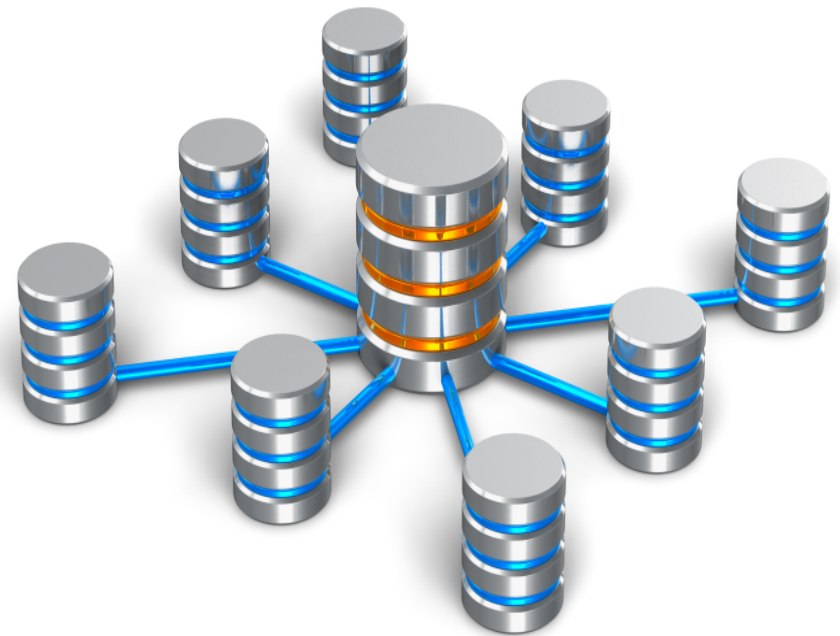
Data Guard

**Duplicate Database Enhancements &
Connection Preservation During Role Change**

Duplicate Database Enhancements

- In Oracle 12.1, we are restricted to create a physical standby database with the DUPLICATE command
- In Oracle 12.2, Oracle extends the DUPLICATE command to create a database from the physical standby database
- Enhanced the DUPLICATE command to create far sync standby instance
  - DUPLICATE TARGET DATABASE FOR FARSYNC FROM ACTIVE DATABASE;
  - DUPLICATE TARGET DATABASE FOR FARSYNC BACKUP LOCATION '+DATA/BACKUP';

Connection Preservation During Role Change

OLD
- All connections established on the **ADG** are terminated, and connections must be re-established

NEW
- Incorporate a database service that is architected to run both primary and the standby databases, users will stay connected as the role transition occurs
- If we use a database service that only connects to the standby database, user sessions will be terminated and forced to re-connect.

- **STANDBY_DB_PRESERVE_STATES** can have following values:
  - NONE — No sessions on the standby are retained during a switchover/failover.
  - SESSION or ALL — User sessions are retained during switchover/failover.

VISCOSITY NORTH AMERICA

Data Guard

Broker DGMGRL Command Scripts

Script Execution

- Execute scripts like we do in SQL*Plus with the @ sign (i.e. @scriptname)
  - Can use the word REM or -- (two dashes followed by a space) to comment a line
  - OS commands such as host or ! can be executed as we do in SQL*Plus
- Each line in the script must have a semi-colon (;) at the end
  - We cannot execute the command "Start Observer" in the script
  - Execute Start Observer In Background;" instead

- **SPOOL** records input and output of the DGMGRL console
- SET TIME turns timestamp printing on and off
- **/ (slash)** repeats the last command input

Data Guard

Multiple Automatic Failover Targets

Multiple Automatic Failover Targets
- Achieve a higher level of success for automatic failover when the need arises
- Configure multiple failover targets in a fast-start failover configuration

- The list of target standby databases are called candidate fast-start failover targets
- The FastStartFailoverTarget property can be set to the list of DB_UNIQUE_NAME,
- Broker will attempt failover in the order they are listed.
  – ANY, then the broker can select any of the candidate targets as the current fast-start failover target.

- `DGMGRL> edit database 'ORCL_SITE' set property FastStartFailoverTarget='DRSITE1,DRSITE2';`
  `Property "faststartfailovertarget" updated`
  – Order of the list matters
- `DGMGRL> edit database 'DRSITE1' set property FastStartFailoverTarget='ANY';`
  `Property "faststartfailovertarget" updated`

VISCOSITY NORTH AMERICA

Data Guard

Broker Support for Redo Transport of Different Endianness to ZDLRA

ZDLRA

- Can manage a remote redo destination that has a different endianess than the primary database

- Consider cross platform migrations with the ZDLRA and the ability to synchronize databases of heterogeneous configurations with different endianness

- Significantly improve the flexibility of migrating databases from one platform to another with the ZDLRA.

ORACLE
ZERO DATA LOSS
RECOVERY APPLIANCE

ZERO DATA LOSS
RECOVERY APPLIANCE

Data Guard

Broker Support for PDB Migration

- New command to the Data Guard Broker, MIGRATE PLUGGABLE DATABASE
  - Movement of a single PDB from one container to another or failover a PDB from the standby database to a new production container database
- Common use cases include:
  - Moving a PDB from one container to another container on the same server
  - Failover a PDB from a physical standby database to a new production database on the same physical standby server

- In either of the scenarios, data files for the PDB being "migrated" must presented and made available to both the current container database and the target container database.
- The source and target CDBs must participate in different Data Guard Broker configurations.

VISCOSITY NORTH AMERICA

Data Guard

Switchover Wait to Drain for Application Continuity

- SWITCHOVER TO database_name [WAIT [timeout_in_seconds]];
  - wait for sessions to drain before proceeding with the switchover

- DGMGRL> switchover to 'svisorc' wait 10;
  Stopping services and waiting up to 10 seconds for sessions to drain...
  Performing switchover NOW, please wait...
  Operation requires a connection to database "svisorc"
  Connecting ...
  Connected to "svisorc"
  Connected as SYSDBA.
  New primary database "svisorc" is opening...
  ...
  Connected to "visorcl"
  Switchover succeeded, new primary is "svisorc"

- This WAIT timeout overrides the drain_timeout value that is in the srvctl utility
  - The drain-timeout value is an option that's part of srvctl add service command.

VISCOSITY NORTH AMERICA
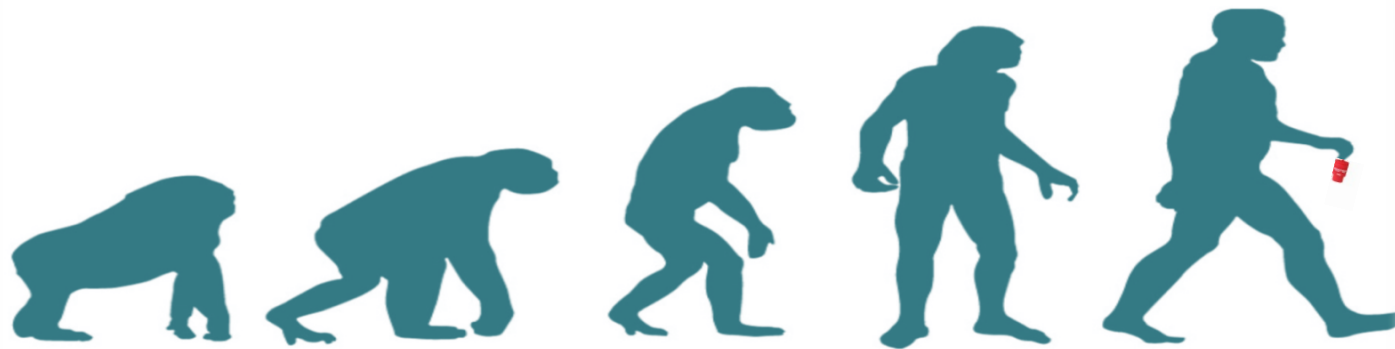
RMAN

Enhancements to Table Recovery

REMAP TABLE clause, to recover tables or table partitions into a schema, that is different from the schema in which these objects originally existed

- ```
  RMAN> RECOVER TABLE VNA.employees,
  VNASALES.channels
  UNTIL SCN 2456
  REMAP TABLE
  VNA.employees:VNAEX.employees,
  VNASALES.channels:VNASCH.channels
  REMAP TABLESPACE
  'SALES_TBS':'NEW_SALES_TBS'
  AUXILIARY DESTINATION
  '/u01/oradata/vnaauxdest';
  ```
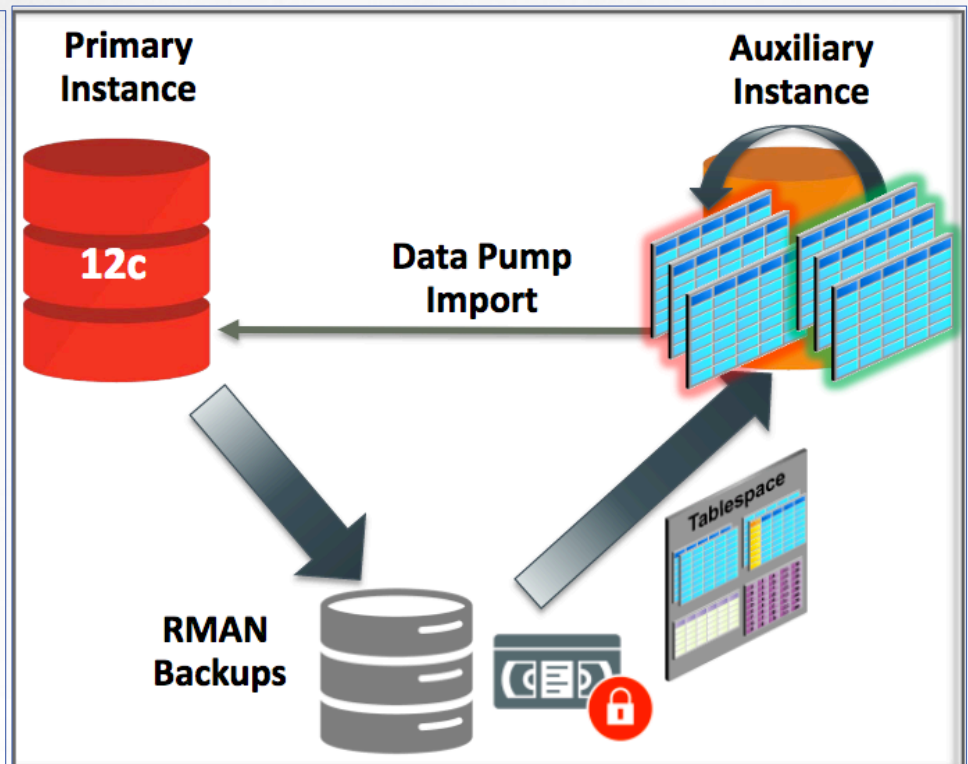
Evolution of Multitenant Cloning

NEW IN
12.2

Read Only Source – **Cold/Remote Clone**	→	Read Write Source – **Hot/Refresh Clone**	→	Online **Relocate**

Containers 0 – 254 (4096 PDBs in 12cR2)

Oracle Database 12c
Release 2 Performance
Tuning Tips and Techniques

- Entire CDB => Container ID = 0
- Root (CDB$ROOT) => Container ID = 1
- Seed (PDB$SEED) => Container ID = 2
- PDBs => Container ID = 3 to 254

```
(While in PDB1):                    (Connect to ROOT):
SQL> SHO CON_ID CON_NAME                SQL> connect / as sysdba
                                    SQL> SHO CON_ID CON_NAME

CON_ID                              CON_ID
------------------------------      ------------------------------
3                                   1


CON_NAME                            CON_NAME
------------------------------      ------------------------------
PDB1                                CDB$ROOT
```

PDB

- **Hot Clones**
  - Oracle 12.2 Multitenant option, fully integrates the concept of "hot clones" with the ability to perform on-line cloning of PDBs.
    - With hot clones, the source database is still open for read-write mode.
    - All PDB clones in Oracle 12.2 are hot clones and will be referred to as clones.
- **Read-Only Refreshable PDBs**
  - On Oracle 12.2, PDBs can be cloned while the source PDB is online.
  - With Refreshable PDBs, we can have the PDB refresh data from the source PDB with delta changes, since the last refresh automatically (define in nnn Minutes) or on demand.

PDB Relocate

- **PDB Relocate**
  - Oracle offers a new RELOCATE clause to the CREATE PLUGGABLE DATABASE, to "relocate" a PDB from one container database (CDB) to another container database.

  - The PDB can be relocated to another CDB on the same server, within the same data center or across data centers.

  - While the relocation process is occurring, database connections still persist on the original PDB. When the statement completes, there will be two transactionally consistent PDBs running.
    - When the PDB is ready, the relocated PDB will be brought online on the new target CDB.
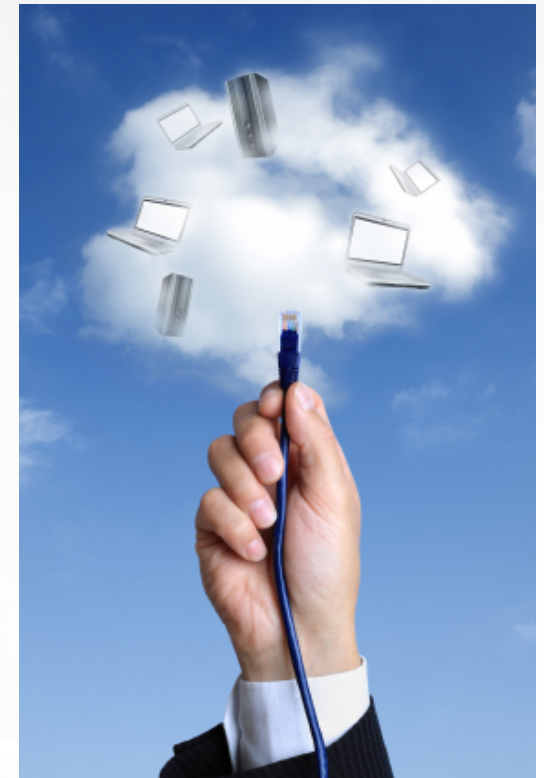
PDB Restore Points

– What is a PDB Restore Point
- A user-defined name associated with an SCN corresponding to the time the restore point was created unique to the PDB

– Normal Restore Point
- A label for an SCN or point in time
- Stored in the control file
- Ages out

– Guaranteed Restore Point
- A restore point for which the database is guaranteed to retain the flashback logs for a flashback operation.
- Stored in the controlfile
- Must be explicitly dropped

– Clean Restore Point
- A PDB restore point that is created when the PDB is closed and when there are no outstanding transactions for that PDB
- Only applies to CDBs that use shared undo

NEW IN 12.2

Isolation with Economies of Scale

- Flashback PDB
- Per-PDB character set
- 4k PDBs per CDB
  (4,096 - increased from 252)
- AWR data at PDB level
- Memory & I/O resource prioritization in addition to CPU
- Configurable isolation via Lockdown Profiles
- PDB-level failover capability with Data Guard Broker

Pricing   Retail

Multitenant Container
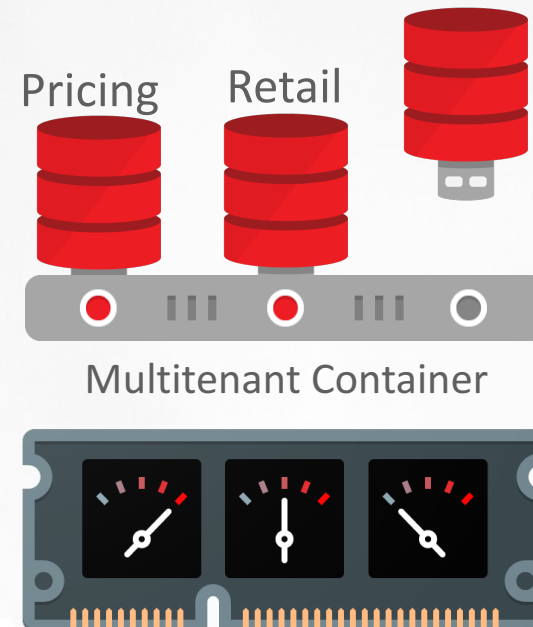
PDB

Localized UNDO and Memory Allocation

- **Localized UNDO**
  - As of Oracle 12.2, we have the option to run UNDO in the local PDBs Snapshot interval and retention can be controlled at PDB level.
  - If the local UNDO option is enabled, each of the PDBs will have it's own UNDO tablespace.
  - There is no mix of partial local UNDO and shared UNDO. You are either 100% local UNDO or 100% shared UNDO.
  - Another powerful feature that is inherited from localized UNDO, is the capability to flashback a PDB.

- **PDB Memory Allocation**
  - Now, with Oracle 12.2, we can set the following initialization parameters at the PDB level:
    - DB_CACHE_SIZE
    - SHARED_POOL_SIZE
    - PGA_AGGREGATE_TARGET
    - SGA_MIN_SIZE
    - SGA_TARGET
  - SGA_MIN_SIZE is the guaranteed SGA size, which includes DB_CACHE_SIZE and SHARED_POOL_SIZE) for PDBs.
  - In general, the total of all SGA_MIN_SIZE should not be greater than 50% of the SGA.

VISCOSITY NORTH AMERICA

Set PDB Level Memory Parameters in 12cR2
(*there are many restrictions not all listed **)

Oracle Database 12*c*
Release 2 Performance
Tuning Tips and Techniques

You can now set at the PDB level (must have MEMORY_TARGET=0 in CDB root):
- DB_CACHE_SIZE** (<=50% of CDB & sum of all PDBs<=50% of CDB level)
- SHARED_POOL_SIZE** (<=50% of CDB & sum of all PDBs<=50% of CDB level)
- PGA_AGGREGATE_TARGET (<= CBD level setting)
- PGA_AGGREGATE_LIMIT* (<= CBD level setting)
- SGA_MIN_SIZE*
- SGA_TARGET*
- INMEMORY_SIZE*

Must have NONCDB_COMPATIBLE=FALSE (in CDB root)
**DB_CACHE_SIZE + SHARED_POOL_SIZE <= 50% CDB level SGA_TARGET

PDB

CPU Count, Max PDBs, Localized AWR, PDBAs & Character Set

- **PDB CPU Count**
  - With Oracle 12.2, we can set the CPU_COUNT initialization parameters at the PDB level.
  - In Oracle 12.1, CPU limits were imposed on the PDB as part of the CDB Resource Plan definition.

- **4096 PDBs per CDB**
  - When PDBs were first introduced in Oracle 12.1, the number of PDBs per CDB was 252
  - Now it has increased 4,096.

- **Localized AWR**
  - Having the AWR at the PDB level, allows for a granular level of tuning and diagnostics.

- **PDBAs**
  - Oracle now introduces the concept of the Pluggable DBAs (PDBAs), so that each PDB can be potentially maintained by a PDB Administrator.

- **PDBS with Different Character Sets**
  - In Oracle 12.2, CDBs can have PDBs with different character sets, as long as the root container's character set is the superset of all of the PDB character sets.

PDB

Mutiple Time Zone Support

- **Heat Map**
  - PDBs are able to benefit from Oracle's Automatic Data Optimization capabilities, because Oracle PDBs now have full Heat Map support.

- **Support for PDBs with Different Time Zone File Versions and Database Time Zones**
  - Now, PDBs in a CDB can be in different time zone file versions and database time zones.

VISCOSITY NORTH AMERICA

Resource Manager Enhancements with Multitenant

NEW IN 12.2

Memory Management

- Much requested capability
  - Not possible in 12.1
  - However, consolidation very successful without this capability!
- Memory RM parameters now at PDB level
- New parameter SGA_Min_Size
  - Partition memory per PDB
  - Use for low density consolidation of platinum / core applications
  - Avoid in other use cases

I/O on Commodity Storage

- 2 new per PDB parameters
  - MAX_IOPS & MAX_MBPS
  - Can be dynamically altered
- Rate limits for PDBs on non-Exadata storage
  - Error if set for CDB$Root
  - Error if set on Exadata system
- IORM only possible with Exadata storage in 12.1
  - Still more flexible on Exadata!
    - Automatic throttling based on shares
    - DBA doesn't need to know IOPS and MBPS of storage & workloads
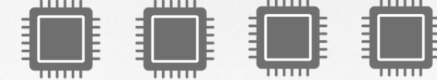
CPU Count per PDB

- Limits CPU consumed by each PDB
  - In 12.1 this is expressed in percentage terms in CDB resource plan
- In 12.2 we have CPU Count as per-PDB parameter
  - No need to re-compute percentage when migrating PDB between boxes of different configurations
  - Percentage still supported for compatibility
  - Lower setting wins

PDB I/O Rate Limits

NEW IN 12.2

Feature

- 2 new PDB parameters
  - MAX_IOPS: maximum I/O requests per second
  - MAX_MBPS: maximum MegaBytes per second of I/O
  - Parameters can be dynamically altered
- Parameters can only be set at the PDB level on non-Exadata systems
  - Setting at the root level results in an error
  - Setting on Exadata results in an error

RMAN

Perform Flashback on a PDB to PIT

- Perform a flashback database operation to rewind an individual PDB to a previous point in time
- In 12.2 users can execute a database flashback on a CDB across PDB, PITR or PDB flashback operations.
- The PDB on which a Flashback Database operation is being performed must be closed
  - SQL> CREATE RESTORE POINT vnapdb_grp_before_changes GUARENTEE FLASHBACK DATABASE;

Flashback a PDB
  - RMAN> SHUTDOWN IMMEDIATE;
  - RMAN> STARTUP MOUNT;
  - RMAN> FLASHBACK PLUGGABLE DATABASE vnapdb TO RESTORE POINT vnapdb_grp_before_changes AUXILIARY DESTINATION '/temp/aux_dest';
  - RMAN> ALTER PLUGGABLE DATABASE vnapdb OPEN RESETLOGS;
  - **Also, in 12cR2, RMAN checks disk space available on Recover Table!**

Lockdown Profiles: Apply Limits to Privileges

PDB_LOCKDOWN initialization parameter to enable the a PDB lockdown profile

NEW IN 12.2

- Lockdown profiles are complementary to grants
- Grants alone are "all or nothing"
- Lockdown profiles add granular control of the capabilities enabled by the grant

```
grant alter system
to pdb_user;
```

```
alter lockdown
profile p1 disable
statement=
('ALTER SYSTEM')
clause=('SET')
option= ALL EXCEPT
('plsql_code_type'
,'plsql_debug'
,'plsql_warnings');
```

- Scope of 'alter system'
  - ~~cursor_sharing~~
  - ~~ddl_lock_timeout~~
  - ~~optimizer_mode~~
  - ~~parallel_degree_limit~~
  - plsql_code_type
  - plsql_debug
  - plsql_warnings
  - ~~resource_manager_plan~~
  - ~~skip_unused_indexes~~
  - ~~...~~

Follow Us Online!

Facebook.com/ViscosityNA

Linkedin.com/company/Viscosity-North-America

@ViscosityNA

Viscosity North America

Facebook.com/ViscosityNA

@Viscosity_NA