

Taming the PDB

Resource Management & Lockdown Profiles in Oracle Database 18c & 19c

Markus Flechtner

 markusdba

 markusdba.de|.net

Tuesday, December 1st
1:00 PM – 1:45 PM GMT

BASEL | BERN | BRUGG | BUCHAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.BR.
GENEVA | HAMBURG | LAUSANNE | MANNHEIM | MUNICH | STUTTGART | VIENNA | ZURICH

trivadis

Markus Flechtner

- Principal Consultant / Trivadis Germany GmbH
- Studied Mathematics a long time ago
- Focus
 - Oracle High Availability
 - Database Upgrade + Migration
- Teacher:
RAC, New Features, Multitenant, PostgreSQL



@markusdba



www.markusdba.net|.de



ORACLE®
ACE

DOAG

BASEL | BERN | BRUGG | BUKAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.B.R. | GENÈVE
HAMBURG | KOPENHAGEN | LAUSANNE | MANNHEIM | MÜNCHEN | STUTTGART | WIEN | ZÜRICH

trivadis


FOUNDED IN
1994

250 SLA's
(SERVICE LEVEL AGREEMENTS)

 **700**
EMPLOYEES

 **16** TRIVADIS
WORKSPACES
SWITZERLAND, GERMANY,
AUSTRIA, DENMARK,
ROMANIA

4000 
TRAINING PARTICIPANTS PER YEAR

5 MILLION
CHF 
BUDGET FOR SCIENCE
AND DEVELOPMENT PER YEAR

118 MILLION
CHF 
TURNOVER

800 
CUSTOMERS

1900 PER
YEAR
PROJECTS

trivadis

Motivation

- In a Multitenant Database, ..
 - ... access to common resources like OS and network should be restricted
 - ... the system resources like CPU, memory and I/O should be distributed among the PDBs in a controlled way
 - ... the DBaaS customer should be hindered from lifting these restrictions
 - ... charging the customers by resource usage (storage, I/O, CPU, memory) should be possible

Agenda

- **Resource Management in a Container Database**
 - CPU
 - Memory
 - I/O
 - Other Resources
- **Lockdown Profiles**

Resource Management in a Container Database

Resource Management for CDBs (1)

- Adequate resource management is essential for Multitenant Databases
- Very often, with DBaaS Service Level Agreements (SLAs) guarantee a certain amount of resources
- Oracle can manage the following resources
 - CPU
 - Number of parallel server processes
 - Memory (since Oracle 12.2)
 - I/O (since Oracle 12.2)
- On Exadata and Oracle Super Cluster I/O-management is available since Oracle 12.1 (not covered in this presentation)

Resource Management for CDBs (2)

- For some of the resource limits, the following procedure applies
 1. Resource limits are defined in CDB\$ROOT
 2. Resource limits are activated in the PDB ("ALTER SYSTEM ..")
- To prevent a PDB administrator from disabling these resource limits, PDB resource management should be combined with lockdown profiles
 - Restrict "ALTER SYSTEM .." on PDB level

Evolution of Resource Management for CDBs

- **Oracle 12c Release 1:**
 - Resources are managed using Resource Manager (CPU share, CPU Limit (%), Parallel Server Limit (%))
- **Oracle 12 Release 2:**
 - Parameter CPU_COUNT can be set on PDB level (CPU threads)
 - Oracle Recommendation: do not use Resource Manager for CPU resources (CPU share, CPU limit), use CPU_COUNT instead

Evolution of Resource Management for CDBs

- **Oracle 18c:**

- Parameter `PARALLEL_SERVERS_TARGET` can be set on PDB level
Default:
 $\text{PARALLEL_THREADS_PER_CPU} * \text{CPU_COUNT} * \text{concurrent_parallel_users} * 2$
- Oracle Recommendation: do not use Resource Manager (Parallel Server Limit)

- **Oracle 19c:**

- Parameter `CPU_MIN_COUNT`: specifies the minimum number of CPU (threads) for a PDB

CPU

trivadis

Two ways to limit CPU resources

- Initialization parameters
 - CPU_COUNT (on PDB level) – since Oracle 12.2
 - CPU_MIN_COUNT (19c)
- Resource Manager (12.1)
 - Takes precedence over parameter settings (CPU_COUNT, CPU_MIN_COUNT)

CPU_COUNT and CPU_MIN_COUNT (19c)

- Dynamic PDB level parameters for instance caging on PDB level
- Resource Manager must be enabled
- Preferred method for limiting CPU resources (threads) on PDB level
- Advantages
 - Values are transferred when the PDB is unplugged and plugged into another CDB
 - Values are "static" and do not depend on the CPU shares of the other PDB

```
SQL> ALTER SYSTEM SET cpu_count = 4;
```

Managing CPU Resources with Resource Manager

EE

- Only method in Oracle 12.1, not recommended for Oracle 12.2 and higher
- CDB Level – between PDBs
 - Define resource allocation to specific PDBs
 - Limit resource utilizations for specific PDBs
 - CDB DBA can give more resources to more important PDBs
 - System resource shares and limits can be configured
- PDB Level – within PDBs
 - Define resource allocations (consumer groups) within specific PDBs
 - As we know it from Non-CDBs

Example for CDB Resource Plan

	Shares	Utilization Limit %	Parallel Server Limit %
Default per PDB	1	100	100
PDB1	1	50	20
PDB2	2	75	20
PDB3	3	100	100

- PDB1 gets guaranteed 1 share of total 6, so 17% of system resources (CPU, Exadata I/O Bandwidth, queued parallel statements) – PDB2 33% – PDB3 50%
- PDB1 can utilize max. 50% of system resources
- PDB1 can utilize max. 20% of parallel server processes specified by the instance parameter `PARALLEL_SERVERS_TARGET`
- One row in the table is defined as a default ("CDB plan directive")

CDB Resource Plan (1)

- Create a pending area and a new CDB plan

```
SQL> execute dbms_resource_manager.create_pending_area;
SQL> execute dbms_resource_manager.create_cdb_plan(
  2 plan                => 'MY_PLAN',
  3 comment              => 'My CDB plan'
  4 );
```

- Create a CDB plan directive for each PDB

```
SQL> execute dbms_resource_manager.create_cdb_plan_directive(
  2 plan                => 'MY_PLAN',
  3 pluggable_database  => 'PDB1',
  4 comment              => 'give 1 share',
  5 shares              => 1,
  6 utilization_limit   => 50,
  7 parallel_server_limit => 20
  8 );
```


CDB Resource Plan (2)

- Update the Default Directive

```
SQL> execute dbms_resource_manager.update_cdb_default_directive (  
  plan                => 'MY_PLAN' ,  
  new_shares          => 1 ,  
  new_utilization_limit => 100 ,  
  new_parallel_server_limit => 100 );
```

- Update the directive for the automated maintenance tasks in CDB\$ROOT

```
SQL> execute dbms_resource_manager.update_cdb_autotask_directive (  
  2 plan                => 'MY_PLAN' ,  
  3 new_shares          => 1 ,  
  4 new_utilization_limit => 20 ,  
  5 new_parallel_server_limit => 10);
```

- Validate and submit

```
SQL> execute dbms_resource_manager.validate_pending_area;  
SQL> execute dbms_resource_manager.submit_pending_area;
```

Query CDB Resource Plans

- Query the assigned directives

```
SQL> SELECT pluggable_database pdb,  
2         shares,  
3         utilization_limit util,  
4         parallel_server_limit par  
5 FROM dba_cdb_rsrc_plan_directives  
6 WHERE plan='MY_PLAN';
```

PDB	SHARES	UTIL	PAR
ORA\$AUTOTASK	1	20	10
PDB1	1	50	20
PDB2	2	75	20
PDB3	3	100	100
ORA\$DEFAULT_PDB_DIRECTIVE	1	100	100

Activate a CDB Resource Plan

- Activate instantly

```
SQL> ALTER SYSTEM SET resource_manager_plan = MY_PLAN
```

- Regularly via Scheduler Window

```
SQL> execute dbms_scheduler.create_window(  
2 window_name      => 'MY_WINDOW',  
3 resource_plan    => 'MY_PLAN',  
4 start_date       => to_timestamp_tz('24.07.2013 20:00:00',  
5                                     'DD.MM.YYYY HH24:MI:SS'),  
6 repeat_interval  => 'freq=daily',  
7 duration         => INTERVAL '2' HOUR  
8 );
```

Example for Combined CDB and PDB Resource Plan Usage

- CDB resource plans and PDB resource plans can be combined:

CDB		
PDB	Shares	Util Limit
PDB1	1	20 %
PDB2	2	100 %
PDB3	3	50 %

PDB3		
Consumer Group	Shares	Util Limit
OLTP	1	30 %
BATCH	2	20 %
BOARD	2	50 %

- How much CPU resources gets consumer group BOARD in PDB3 ?
 - Guaranteed is: $3/6 * 2/5 = 6/30 = 20\%$
 - Limited to: $50 \% * 50 \% = 25 \%$

Performance Profiles

- A performance profile is a collection of resource manager settings for PDBs
 - E.g. SLA level (gold, silver, bronze)
- Easy way to modify the resource limits for a group of PDBs

```
DBMS_RESOURCE_MANAGER.CREATE_CDB_PROFILE_DIRECTIVE(  
  plan => 'newcdb_plan',  
  profile => 'gold',  
  shares => 3,  
  utilization_limit => 100,  
  parallel_server_limit => 100);
```

- Activation on PDB level

```
ALTER SYSTEM SET DB_PERFORMANCE_PROFILE=gold SCOPE=SPFILE;
```

- The parameter is static; the PDB has to be bounced to activate the setting

Management of Parallel Server Processes on PDB level

- **Oracle 12.1 and Oracle 12.2**

- Resource Manager
- Resource "PARALLEL_SERVER_LIMIT" (Percentage of parallel server processes specified by the instance parameter PARALLEL_SERVERS_TARGET)
- Not recommended in Oracle 18c and higher

- **Oracle 18c**

- Set parameter **PARALLEL_SERVERS_TARGET** on PDB level
- Default:
 $\text{PARALLEL_THREADS_PER_CPU} * \text{CPU_COUNT} * \textit{concurrent_parallel_users} * 2$

I/O

I/O Rate Limits for PDBs (1)

- New parameters on container level (CDB\$ROOT or PDB):
 - Values set in CDB\$ROOT are the default for PDBs

```
ALTER SYSTEM SET MAX_IOPS = 1000 SCOPE = BOTH  
ALTER SYSTEM SET MAX_MBPS = 5 SCOPE = BOTH
```

- To disable a limit, set the parameter to 0 (Default)
- Event "resmgr: I/O rate limit" (V\$SYSTEM_EVENT, V\$SESSION_EVENT) is raised when the limit is hit
- DBWR-I/O, Controlfile-I/O and Password file are exempted
- Not supported on Exadata or on Non-CDB databases (IORM is used)

I/O Rate Limits for PDBs (2) - Example

```
SQL> select * from dba_tab_columns;  
.. 9 seconds
```

```
SQL> alter system set max_iops=50 scope=both;
```

```
SQL> alter system set max_mbps=12 scope=both;
```

```
SQL> select * from dba_tab_columns;  
.. 19 seconds
```

```
SQL> select con_id,event,total_waits,time_waited from v$session_event  
2 where event='resmgr: I/O rate limit';
```

CON_ID	EVENT	TOTAL_WAITS	TIME_WAITED
7	resmgr: I/O rate limit	32	1559

Other resources

Memory Management for PDBs (1)

- The following memory related parameters can be set on PDB level:

Parameter	Meaning
DB_CACHE_SIZE	Minimum guaranteed buffer cache for the PDB
SHARED_POOL_SIZE	Minimum guaranteed shared pool for the PDB
PGA_AGGREGATE_LIMIT	Maximum PGA size for the PDB
SGA_MIN_SIZE	Minimum SGA size for the PDB
SGA_TARGET	Maximum SGA size for the PDB
INMEMORY_SIZE	Maximum size of the In-Memory-Column-Store (IMCS, available and modifiable since 12.1.0.2, In-Memory-Option required (see notes))

Memory Management for PDBs (2)

- Requirements
 - Parameter NONCDB_COMPATIBLE=FALSE in CDB\$ROOT
 - MEMORY_TARGET not set in CDB\$ROOT (or MEMORY_TARGET=0)
- Restrictions for SGA related parameters (if SGA_TARGET=0)
 - Sum of all minimal values for SGA (DB_CACHE_SIZE, SGA_MIN_SIZE, SHARED_POOL_SIZE) for all PDBS must no be higher than 50% of the corresponding value for CDB\$ROOT
- Restrictions (PGA)
 - PGA_AGGREGATE_LIMIT: not more than PGA_AGGREGATE_LIMIT in CDB\$ROOT
- INMEMORY_SIZE
 - Overprovisioning allowed

Other resources

Other resource-relevant Parameters on PDB-level

Parameter	Default	Meaning
MAX_DATAPUMP_JOBS_PER_PDB	100	Maximum number of concurrent DataPump Jobs in the PDB (max. 252 or "AUTO") (Value in CDB\$ROOT is the default for the PDBs) (available since Oracle 12.2)
AWR_PDB_MAX_PARALLEL_SLAVES	10	Amount of resources dedicated to AWR snapshot flushing in PDBs (number of MMON slave processes) (available since 18c)
SESSIONS	from CDB\$ROOT	Maximum of concurrent sessions in the PDB (available since Oracle 12.1)

- Not limiting the number of sessions in a PDB can lead to the problem that one PDB takes all available sessions (instance parameter SESSIONS)
 - ➔ no logins to the other PDBs and to CDB\$ROOT possible

Restricting directory access

```
SQL> create pluggable database PDBTEST
      2 admin user PDBADMIN identified by manager roles=(DBA)
      3 path_prefix='/u02/' create_file_dest='/u02/';
```

- **PATH_PREFIX**

- Is added to relative paths e.g. for directory objects, absolute paths result in "ORA-65254: invalid path specified for the directory"

- **CREATE_FILE_DEST**

- Restricts the path for database files

```
SQL> create tablespace T2 datafile '/u01/oradata/t2.dbf' size 10M;
create tablespace T2 datafile '/u01/oradata/t2.dbf' size 10M
ERROR at line 1:
ORA-65250: invalid path specified for file - /u01/oradata/t2.dbf
```

Database Maintenance Jobs in a Multitenant Database

- Since Oracle 10g, Oracle runs various maintenance jobs (e.g. Statistics Gathering, Tuning Advisor etc.) in a defined maintenance window
- In a Multitenant database this window is the default window for all PDBs!
 - → in a CDB with a huge number of PDBs this can lead to high load during the maintenance window
- Workarounds:

```
ALTER SYSTEM SET AUTOTASK_MAX_ACTIVE_PDBS = 4 SCOPE = BOTH;
```

- Deactivate automatic maintenance jobs on PDB level (Default: TRUE)

```
ALTER SYSTEM SET ENABLE_AUTOMATIC_MAINTENANCE_PDB = FALSE SCOPE = BOTH;
```


Resource Monitoring for PDBs (1)

- The view V\$RSRCPDBMETRIC contains data for the last minute (available since 12.2)
- V\$RSRCPDBMETRIC_HISTORY (with the same structure) contains data for every minute of the last hour
- The AWR-View **CDB_HIST_RSRC_PDB_METRIC** (Diagnostic Pack required!) contains persistent snapshots of the view V\$RSRCPDBMETRIC
 - Can be used for PDB charging (e.g. by used memory, I/O etc.)

Resource Monitoring for PDBs (2)

```
SQL> SELECT p.pdb_name,h.begin_time,h.end_time,h.iops,h.iombps, h.cpu_consumed_time cpu_time,
 2          trunc(h.sga_bytes/1024/1024)+1 sga_mb, trunc(h.pga_bytes/1024/1024)+1 pga_mb
 3 FROM cdb_hist_rsrc_pdb_metric h, cdb_pdbns p
 4 WHERE p.dbid = h.dbid
 5        AND begin_time > sysdate-12/24
 6        AND p.pdb_name='PDB01'
 7 ORDER BY p.pdb_name, h.begin_time;
```

PDB_NAME	BEGIN_TIME	END_TIME	IOPS	IOMBPS	CPU_TIME	SGA_MB	PGA_MB
PDB01	13.06.2019 13:00	13.06.2019 13:38	.68	.00	10	50	7
PDB01	13.06.2019 13:38	13.06.2019 14:00	2.66	.04	1067	121	7
PDB01	13.06.2019 14:00	13.06.2019 15:00	1.66	.03	16940	127	9
PDB01	13.06.2019 15:00	13.06.2019 16:00	.99	.02	438	139	14
PDB01	13.06.2019 19:59	13.06.2019 20:03	.07	.00	10	116	7
PDB01	13.06.2019 20:03	13.06.2019 20:59	1.08	.02	151	177	9
PDB01	13.06.2019 20:59	13.06.2019 21:59	.55	.01	490	203	22

trivadis

Lockdown Profiles

trivadis

Lockdown Profiles (1)

- Restrict feature usage on PDB level
- Areas
 - Network access
 - Common user or object access
 - Administrative features
 - XML database access
 - Database options (e.g. Partitioning)
- Use cases
 - DBaaS
 - Other variations of role separation between CDB administrator and PDB administrator

Lockdown Profiles (2)

- Create a lockdown profile in CDB\$ROOT and assign the restrictions

```
SQL> CREATE LOCKDOWN PROFILE demo_lckdprf;
SQL> ALTER LOCKDOWN PROFILE demo_lckdprf DISABLE STATEMENT = ('ALTER SYSTEM');
SQL> ALTER LOCKDOWN PROFILE demo_lckdprf ENABLE STATEMENT = ('ALTER SYSTEM')
  2 CLAUSE = ('flush shared_pool');
SQL> ALTER LOCKDOWN PROFILE demo_lckdprf DISABLE STATEMENT = ('ALTER SYSTEM')
  2 CLAUSE = ('SET') OPTION=('OPTIMIZER_INDEX_COST_ADJ')
  3 MINVALUE='80' MAXVALUE='120';
SQL> ALTER LOCKDOWN PROFILE demo_lckdprf DISABLE FEATURE = ('NETWORK_ACCESS');
SQL> ALTER LOCKDOWN PROFILE demo_lckdprf DISABLE OPTION = ('Partitioning');
```

- In the example, the lockdown profile restricts the values for "OPTIMIZER_INDEX_COST_ADJ" from 80 to 120. Trying to set the parameter to a value out of this range will result in an error ("ORA-01031: insufficient privileges")

Lockdown Profiles (3)

- Activate the lockdown profile

```
SQL> ALTER SESSION SET CONTAINER=PDB1;  
SQL> ALTER SYSTEM SET PDB_LOCKDOWN = demo_lckdprf SCOPE = BOTH;
```

- What happens in the PDB?

```
SQL> ALTER SYSTEM FLUSH BUFFER_CACHE;  
Error at line 1:  
ORA-01031: insufficient privileges  
  
SQL> ALTER SYSTEM FLUSH SHARED_POOL;  
System altered.  
  
SQL> CREATE TABLE .. PARTITION BY ..  
ERROR at line 1:  
ORA-00439: feature not enabled: Partitioning
```

Lockdown Profiles - Data Dictionary

```
SQL> select
  2     PROFILE_NAME,
  3     RULE_TYPE,
  4     RULE,
  5     clause,
  6     status
  7 from dba_lockdown_profiles
  8 order by profile_name;
```

PROFILE_NAME	RULE_TYPE	RULE	CLAUSE	STATUS
DEMO_LCKDPRF	STATEMENT	ALTER SYSTEM		DISABLE
DEMO_LCKDPRF	STATEMENT	ALTER SYSTEM	FLUSH SHARED_POOL	ENABLE
DEMO_LCKDPRF	FEATURE	NETWORK_ACCESS		DISABLE
DEMO_LCKDPRF	OPTION	PARTITIONING		DISABLE
PRIVATE_DBAAS				EMPTY
PUBLIC_DBAAS				EMPTY
SAAS				EMPTY

Lockdown Profiles - Enhancements in Oracle 18c

- PDB lockdown profiles in the Application root, as well as in the CDB root
- Create PDB lockdown profile that is based on another PDB lockdown profile

```
SQL> REM Static dependency
SQL> CREATE LOCKDOWN PROFILE <profile> FROM <other profile>;
SQL> REM Dynamic dependency
SQL> CREATE LOCKDOWN PROFILE <profile> INCLUDING <other_profile>;
```

- New view V\$LOCKDOWN_RULES to see the lockdown rules

Lockdown Profiles – Local vs. Common Users

- A lockdown profile can have different restrictions for local and common users

```
SQL> ALTER LOCKDOWN PROFILE demo_lockdown_profile
2  DISABLE STATEMENT = ('ALTER SYSTEM');

SQL> ALTER LOCKDOWN PROFILE demo_lockdown_profile
2  ENABLE STATEMENT = ('ALTER SYSTEM') CLAUSE = ('SET') USERS=COMMON;
```

- Valid Clauses for “USERS=” are "ALL", "LOCAL" and "COMMON"
- Recommendation
 - Disable "ALTER SYSTEM" for local users only, so that they cannot disable the lockdown profile or modify resource restrictions

Summary & Further Information

Summary

- Resource management and resource monitoring is a must for DBaaS
- Oracle 12.2+ can manage all kinds of resources (CPU, I/O, memory)
- Lockdown Profiles allow fine granular access to features and administrative commands
- Use lockdown profile to prevent PDB users from lifting restrictions
- Be careful when defining resource limits!

Further Information

- MOS Note 2171135.1: Managing OS Resources Among PDBs Using PDB Performance Profiles - 12.2 New Feature
- MOS-Note 2170772.1: How to Control and Monitor the Memory Usage (Both SGA and PGA) Among the PDBs in Multitenant Database- 12.2 New Feature
- MOS-Note 2326708.1: How to Provision PDBs, based on CPU_COUNT
- Whitepaper PDB Isolation:
<http://www.oracle.com/technetwork/database/multitenant/learn-more/isolation-wp-12c-3614475.pdf>
- Presentation "PDB Isolation and Security" by Stefan Oehrli
http://www.oradba.ch/download/UKOUG_PDB-Isolation-Security.pdf

Questions & Answers

Markus Flechtner

markus.flechtner@trivadis.com

Phone +49 211 5866 64725



@markusdba



markusdba.de|.net

BASEL | BERN | BRUGG | BUCHAREST | DÜSSELDORF | FRANKFURT A.M. | FREIBURG I.BR.
GENEVA | HAMBURG | LAUSANNE | MANNHEIM | MUNICH | STUTTGART | VIENNA | ZURICH

trivadis



Making a **WORLD** possible
in which **intelligent IT**
facilitates **LIFE and WORK** as a
matter of course.