peakmarks

peakmarks® Runs

January 2025

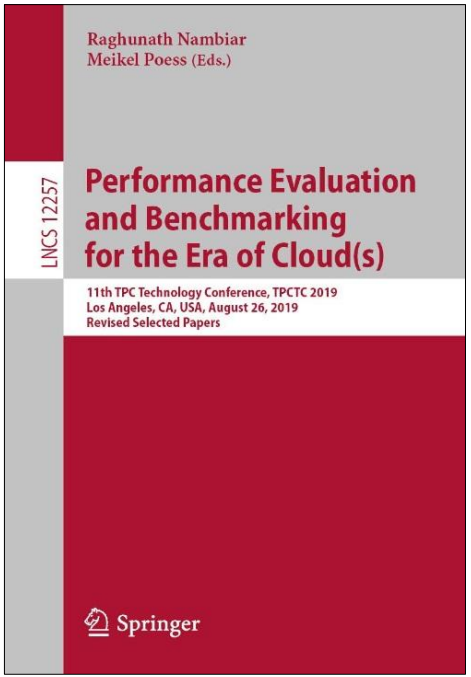

peakmarks® presented its software at the 11th Technology Conference of the Transaction Processing Performance Council (TPC) 2019 in Los Angeles.

Copyright Information

Naming Conventions

Database name	ORA19C / ORA21C / ORA23AI
Instance names	ORA19C / ORA21C / ORA23AI for a single instance
	ORA19C1 / ORA21C1 / ORA23AI1 for RAC instance 1
	ORA19C2 / ORA21C2 / ORA23AI2 for RAC instance 2
peakmarks® PDB	PMK
Connect string SYSTEM user	system/manager@SYSAWR
Connect string peakmarks user	bench/bench@PMK
peakmarks® base directory	../pmk

Agenda

1 Introduction – Runs, Tests and Jobs

2 Automatic Configuration

3 Smart Configuration

4 Sample Configuration

5 Manual Configuration

6 Summary of Scripts and Commands

peakmarks

Performance is not everything.
But without performance, everything is worth nothing.

Introduction – Runs, Tests and Jobs

Overview

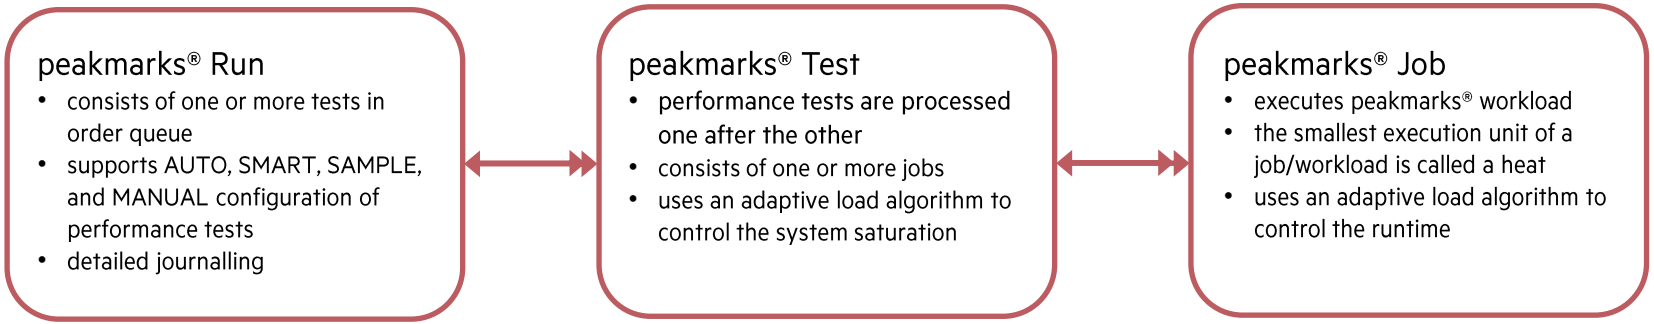

peakmarks® Run Configuration

There are different types of peakmarks® runs dependent on their configuration

- Automatic configuration — fully automated; for administrative tasks
- Manual configuration — to achieve the greatest flexibility in configuring workloads
- Sample configuration — fully automated; for quick sample tests during maintenance windows
- Smart configuration — fully automated; the most convenient and fastest way to get a complete overview of the performance of an Oracle database service

peakmarks® Runs

A unique run ID identifies each peakmarks® run

A peakmarks® run is controlled by following commands

```
SQL> exec pmk.start_manual;
SQL> exec pmk.start_sample;
SQL> exec pmk.start_smart (...);
SQL> exec pmk.stop_run (...);
```

Each peakmarks® run has a status

EXE	–	run executing
ABORT	–	run aborted with pmk.stop_run command
FAIL	–	run has failures
OK	–	run successfully processed without failures

peakmarks® Performance Tests

Each performance test is described by a set of parameters

- ID — unique test number within peakmarks run
- Workload — workload of this performance test
- Parameter — optional workload parameter
- ALC — adaptive load control
  - if > 0: number of heats to be executed
  - if 0: execute the test until the runtime target is reached
- Nodes — number of RAC nodes used for this performance test
- Jobs — number of workload jobs used for this performance test, distributed over all nodes (parameter Nodes)
- DOP — Oracle degree of parallelism for this performance test
- Runtime — runtime target in minutes for this specific performance test

peakmarks

Swiss precision in performance measurement.

Automatic Configuration

Automatic Run Configuration

The Automatic Configuration is fully automated.

It is used for administrative tasks like loading and dropping the peakmarks data and populating the database caches.

Examples

- SQL> exec pmk.load_pdb;
- SQL> exec pmk.purge_pdb;
- SQL> exec pmk.warmup_pdb;

peakmarks

Simple. Representative. Fast.

Smart Configuration

Smart Configuration

The Smart Configuration is fully automated.

It is the most convenient and fastest way to get a complete overview of the performance of an Oracle database service.

It generates a sequence of max 32 tests for each workload with an increasing number of processes. The increase in the number of processes is defined by the parameter P_INCREMENT when calling the Smart Configuration. peakmarks® uses a default value if P_INCREMENT is not specified.

peakmarks® automatically terminates the test sequence when this specific workload's most important performance metric stops increasing but runs at least some tests in any case.

Smart Configuration

To run the Smart Configuration, use the procedure `pmk.start_smart` with the following parameters

- P_WKLGROUP — workload group
- P_INCREMENT — optional step size for some workload groups when increasing processes; default value based on peakmarks® configuration parameter CPUCOUNT
- P_PARAMETER — optional workload parameter; default value based on P_WKLGROUP
- P_REMARK — optional comment; default value based on P_WKLGROUP, P_INCREMENT and P_PARAMETER

Following workload group are supported; each workload group includes a list of workloads:

SRV	SRV-QUERY1, SRV-QUERY25, SRV-REPORT, SRV-SCAN
SRV-MIXED2	SRV-MIXED2
STO	STO-READ, STO-OFFLOAD, STO-RANDOM, STO-SCATTER
LGWR	LGWR-LAT, LGWR-THR
DBWR	DBWR-THR
DL	DL-BUFFER, DL-DIRECT
DA	DA-STORAGE, DA-OFFLOAD, DA-ROWSTORE, DA-COLSTORE
TP	TP-REPORT, TP-LOOKUP, TP-LIGHT, TP-MEDIUM, TP-HEAVY
TP-COMPLEX	TP-MIXED1, TPMIXED2
PLS	PLS-ADD, PLS-BUILTIN
PLS-MIXED2	PLS-MIXED2, PLS-FIBO, PLS-PRIME

Smart Configuration

Examples

- SQL> exec pmk.start_smart ('TP');
- SQL> exec pmk.start_smart ('TP', 4);
- SQL> exec pmk.start_smart ('TP', NULL, 20)
- SQL> exec pmk.start_smart ('TP', NULL, 20, 'TP workloads on test system');

Determination of the JOBS sequence

All job sequences start with 1 job (single thread performance)

For all CPU-bound workloads, peakmarks® generates 5 additional measuring points:

- Utilize 25% of all threads JOBS = 0.25 * CPU_COUNT
- Utilize 50% of all threads JOBS = 0.50 * CPU_COUNT
- Utilize 75% of all threads JOBS = 0.75 * CPU_COUNT
- Utilize all threads JOBS = CPU_COUNT
- Overload system JOBS = CPU_COUNT * 1.25

For all other workloads, peakmarks® generates a maximum of 31 additional measuring points based on parameter P_INCREMENT

The maximum number of JOBS is limited by 3 * CPU_COUNT

Auto Stop of Test Series

peakmarks® stops the configured series of performance tests if the performance does no longer increase

```
BENCH@PMK SQL> @show_teststats

Fri 31-Jan-2025 09:06:35

peakmarks Test Statistics
-------------------------

Run.........:
Test........:
Workload....:

Database....:  PMK               Oracle......:  19.25.0
Instance....:  ORA19C1           Build.......:  250201
RAC nodes...:  2                 Platform....:  peakmarks Ref System


                              Para        KPM value        KPM value KPM value    peakmarks    peakmarks     database     database
 Run Test Status Workload     meter        curr test       prev test   inc [%]       snap 1       snap 2       snap 1       snap 2
---- ---- ------ ------------ ------ ---------------- ---------------- --------- ------------ ------------ ------------ ------------
   5  129     OK STO-SCATTER    100           14,156                0      0.00          125          126        5,438        5,439
      130     OK STO-SCATTER    100           61,535           14,156    334.69          127          128        5,440        5,441
      131     OK STO-SCATTER    100           75,578           61,535     22.82          129          130        5,442        5,443
      132     OK STO-SCATTER    100           87,990           75,578     16.42          131          132        5,444        5,445
      133     OK STO-SCATTER    100           87,081           87,990     -1.03          133          134        5,446        5,447

5 rows selected.

BENCH@PMK SQL>
```

Starting all workload groups

Script to call separate runs for each workload group

- SQL> @run_all

Call for separate run of each workload group (example)

- SQL> exec pmk.start_smart ('SRV');
- SQL> exec pmk.start_smart ('SRV-MIXED2');
- SQL> exec pmk.start_smart ('STO', NULL, 20);
- SQL> exec pmk.start_smart ('LGWR', NULL, 125, NULL);
- SQL> exec pmk.start_smart ('DBWR');
- SQL> exec pmk.start_smart ('DL', Null, 5);
- SQL> exec pmk.start_smart ('DA');
- SQL> exec pmk.start_smart ('TP', NULL, 20);
- SQL> exec pmk.start_smart ('TP-MIXED');
- SQL> exec pmk.start_smart ('PLS');
- SQL> exec pmk.start_smart ('PLS-MIXED2');

Some comments

Server (SRV) workloads

- Parameters P_PARAMETER and P_INCREMENT ignored

Storage (STO) workloads

- Parameter P_PARAMETER only used for an optional second set of STO_RANDOM performance tests (the first set use 100% I/O reads and 0% I/O writes)
- Workload STO-OFFLOAD on Exadata only

Log Writer (LGWR) workloads

- Parameter P_PARAMETER only used for an optional second set of LGWR-LAT performance tests (the first set uses 1 kByte REDO data per transaction)

Some comments

Database Writer (DBWR) workload

- Parameter P_PARAMETER ignored

Data Load (DL) workloads

- Parameter P_PARAMETER only used for an optional second set of DL-BUFFER performance tests (the first set uses 5 rpt)

Data Analytics (DA) workloads

- Parameter P_PARAMETER ignored
- Parameter P_INCREMENT ignored for workloads DA-ROWSTORE
- Workload DA-COLSTORE only when column-store is populated
- Workload DA-OFFLOAD on Exadata only

Some comments

Transaction processing (TP) workloads

- Parameter P_PARAMETER uses default 20 if no value is specified
- If the memory-optimized pool is configured, the workload TP-LOOKUP accesses rows via hash key, otherwise the workload TP-LOOKUP accesses rows via b-tree index

PL/SQL Application (PLS) workloads

- Parameter P_PARAMETER and P_INCREMENT ignored

Reporting peakmarks® Runs after using SQL> @run_all Script for single instance and 2-node RAC cluster

```
BENCH@PMK SQL> @show_runs

Fri 31-Jan-2025 09:07:27

peakmarks Run(s)
----------------

Database....:  PMK               Oracle......:  19.25.0
Instance....:  ORA19C1           Build.......:  250201
RAC nodes...:  2                 Platform....:  peakmarks Ref System
```

Run	Type	Status	Tests	Run begin	Run end	Elapsed time [min]	Database errors	peakmarks errors	Remark
1	Auto	OK	1	29-JAN 20:48:07	29-JAN 22:35:41	107.57	0	0	pmk.load_pdb on 1 node, per instance 8192 GByte dbsize and 36 loader
2	Auto	OK	1	29-JAN 22:53:59	29-JAN 22:55:06	1.12	0	0	pmk.warmup_pdb on 1 node, per instance 8192 GByte dbsize and 36 loader
3	Smart	OK	25	29-JAN 22:55:07	30-JAN 00:23:15	88.14	0	0	pmk.start_smart, workload group SRV
4	Smart	OK	7	30-JAN 00:23:16	30-JAN 00:47:59	24.72	0	0	pmk.start_smart, workload group SRV-MIXED2
5	Smart	OK	160	30-JAN 00:48:03	30-JAN 02:55:03	126.99	0	0	pmk.start_smart, workload group STO, parameter 20, inc 8
6	Smart	OK	96	30-JAN 02:55:05	30-JAN 04:56:23	121.29	0	0	pmk.start_smart, workload group LGWR, parameter 25, inc 8
7	Smart	OK	32	30-JAN 04:56:24	30-JAN 05:14:06	17.69	0	0	pmk.start_smart, workload group DBWR, inc 8
8	Smart	OK	96	30-JAN 05:14:08	30-JAN 07:14:34	120.42	0	0	pmk.start_smart, workload group DL, parameter 5, inc 8
9	Smart	OK	102	30-JAN 07:14:36	30-JAN 08:36:01	81.41	0	0	pmk.start_smart, workload group DA, inc 8
10	Smart	OK	161	30-JAN 08:36:05	30-JAN 13:43:31	307.44	0	0	pmk.start_smart, workload group TP, parameter 20, inc 8
11	Smart	OK	65	30-JAN 13:43:33	30-JAN 15:19:51	96.30	0	0	pmk.start_smart, workload group TP-COMPLEX, inc 8
12	Smart	OK	8	30-JAN 15:19:53	30-JAN 15:44:50	24.95	0	0	pmk.start_smart, workload group PLS-MIXED2
13	Auto	OK	1	30-JAN 17:10:01	30-JAN 18:13:15	63.23	0	0	pmk.load_pdb on 2 nodes, per instance 4096 GByte dbsize and 24 loader
14	Auto	OK	1	30-JAN 18:31:43	30-JAN 18:32:53	1.16	0	0	pmk.warmup_pdb on 2 nodes, per instance 4096 GByte dbsize and 24 loader
15	Smart	OK	25	30-JAN 18:32:59	30-JAN 20:02:42	89.71	0	0	pmk.start_smart, workload group SRV
16	Smart	OK	7	30-JAN 20:02:45	30-JAN 20:27:50	25.08	0	0	pmk.start_smart, workload group SRV-MIXED2
17	Smart	OK	160	30-JAN 20:28:46	30-JAN 22:21:24	112.63	0	0	pmk.start_smart, workload group STO, parameter 20, inc 8
18	Smart	OK	96	30-JAN 22:22:10	31-JAN 00:03:13	101.05	0	0	pmk.start_smart, workload group LGWR, parameter 25, inc 8
19	Smart	OK	32	31-JAN 00:03:28	31-JAN 00:21:27	17.98	0	0	pmk.start_smart, workload group DBWR, inc 8
20	Smart	OK	96	31-JAN 00:22:06	31-JAN 02:02:26	100.32	0	0	pmk.start_smart, workload group DL, parameter 5, inc 8
21	Smart	OK	102	31-JAN 02:03:07	31-JAN 03:28:42	85.60	0	0	pmk.start_smart, workload group DA, inc 8
22	Smart	OK	161	31-JAN 03:29:56	31-JAN 06:52:52	202.93	0	0	pmk.start_smart, workload group TP, parameter 20, inc 8
23	Smart	OK	65	31-JAN 06:53:19	31-JAN 07:50:48	57.48	0	0	pmk.start_smart, workload group TP-COMPLEX, inc 8
24	Smart	OK	8	31-JAN 07:50:51	31-JAN 08:16:11	25.34	0	0	pmk.start_smart, workload group PLS-MIXED2

```
24 rows selected.
```

peakmarks

Simple. Representative. Fast.

Sample Configuration

Sample Run Configuration

The Sample Configuration is a preconfigured and fully automated peakmarks® run.

The peakmarks® user can customize this configuration by changing table PMK_SAMPLE.

The default Sample Configuration is used to take a performance sample of only a few key performance metrics within a short period (< 1 hour) during a maintenance window.

It is used for quality assurance tasks to ensure that a database service (on-premises or in the cloud) meets the agreed performance metrics over an extended period.

Sample Run Configuration

Repeatedly taking a short performance sample within the maintenance windows (e.g., every month or quarter).

- One sample peakmarks® run takes less than 1 hour
- Determine key performance metrics for selected workloads

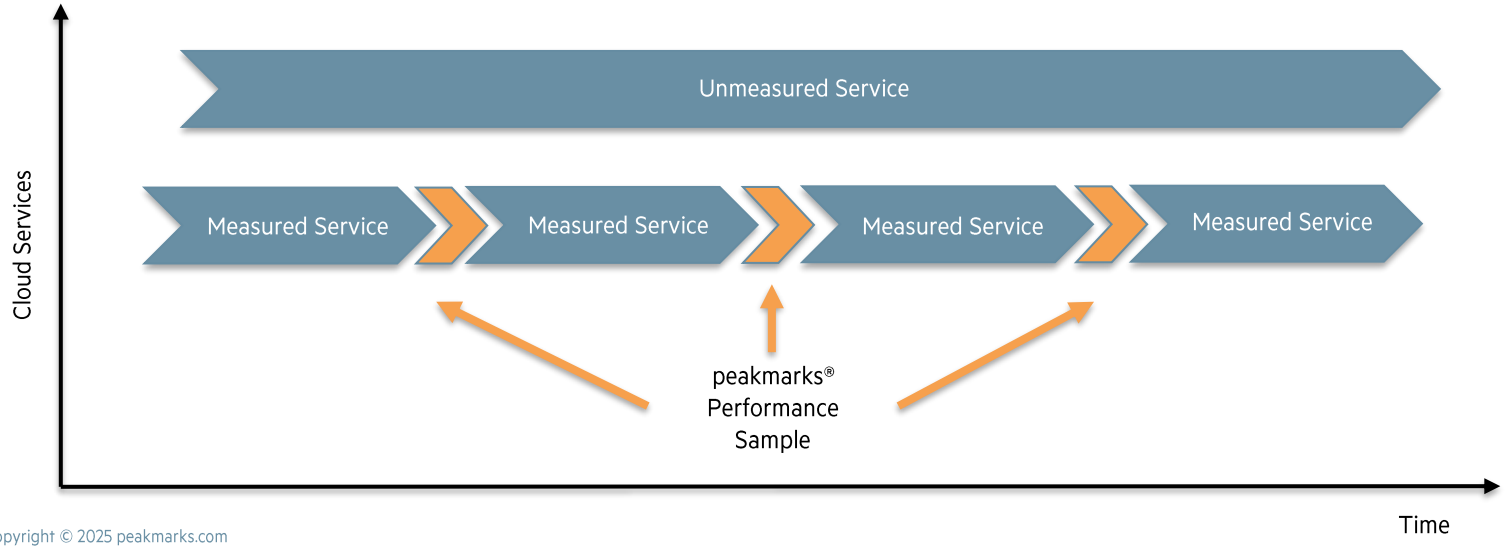

Pre-configured Sample Configuration

```
BENCH@PMK SQL> @show_sample

Fri 31-Jan-2025 09:07:22

Sample peakmarks Configuration
------------------------------

Database....:  PMK               Oracle......:  19.25.0
Instance....:  ORA19C1           Build.......:  250201
RAC nodes...:  2                 Platform....:  peakmarks Ref System

                                              Runtime
                    Para                       target
Test Workload      meter  ALC Nodes  Jobs  DOP   [min] Remark
---- ------------ ------ --- ----- ----- ---- ------- ------------------------------------------------------------------
   1 SRV-MIXED2      N/A    0     1    19    1       3 perform mix of queries and scans on row store
   2 STO-READ        N/A    0     1    48    4       3 perform sequential read using conventional storage system
   3 STO-OFFLOAD     N/A    0     1    48    1       3 perform sequential read using offload technology
   4 STO-RANDOM        0    0     1    48    1       3 perform random single block read, 100% read
   5 STO-RANDOM       20    0     1    48    1       3 perform random single block read, 80% read, 20% write
   6 LGWR-LAT          2    0     1    48    1       3 test LGWR latency with 1 kbyte REDO transactions
   7 LGWR-THR        N/A    0     1    48    1       3 test LGWR throughput
   8 DL-BUFFER         5    0     1    48    1       3 perform buffered data load, 5 rpt
   9 DL-DIRECT       N/A    0     1    48    1       3 perform direct data load
  10 TP-MIXED1       N/A    0     1    19    1       5 perform read-intensive OLTP transactions
  11 TP-MIXED2       N/A    0     1    19    1       5 perform write-intensive OLTP transactionss
  12 PLS-MIXED2      N/A    0     1    19    1       3 perform PL/SQL built-in operations on different data type
                                                -------
Sum                                                  40

12 rows selected.

BENCH@PMK SQL>
```

Sample Configuration

Example

- SQL> exec pmk.start_sample;

Use KPM reports to report the results.

peakmarks

Swiss precision in performance measurement.

Manual Configuration

Manual Run Configuration

The manual configuration supports a maximum of configuration options

It is used by engineers to analyze specific load situations or to find configurations for maximum performance numbers

The manual configuration requires the complete input of all test parameters from an external file in CSV format called manual.csv located in the directory ../pmk/cfg

Manual Configuration

Implemented as external table manual.csv

- Location ../pmk/cfg
- Each row describes one performance test

The ../pmk/cfg directory contains some templates for

- Adapt the template and copy it to manual.csv

Check the current configuration file manual.csv with

- SQL> @show_manual

Manual Configuration

Configuration file manual.csv is managed as an external table

- Column A Workload — workload name
- Column B Parameter — workload parameter
- Column C ALC — adaptive load control, if 0, use peakmarks® configuration parameter RUNTIME otherwise, the number of heats
- Column D Nodes — number of RAC nodes used for this test
- Column E Jobs — number of jobs used for this test, distributed over all nodes configured in column D
- Column F DOP — Oracle degree of parallelism for this test
- Column G Runtime — overrules configuration parameter RUNTIME for this specific test
- Column H Comment — additional information

Monitoring the manual.csv configuration file

```
BENCH@PMK SQL> @show_manual

Fri 31-Jan-2025 09:08:07

Manual peakmarks Configuration
------------------------------

Database....:  PMK                Oracle......:  19.25.0
Instance....:  ORA19C1            Build.......:  250201
RAC nodes...:  2                  Platform....:  peakmarks Ref System

                                                 Runtime
                    Para                          target
Test Workload      meter  ALC Nodes  Jobs  DOP    [min] Remark
---- ------------ ------ ---- ----- ----- ---- ------- ------------------------------------------------
   1 PLS-ADD          SI    0     1     1    1       1 datatype SIMPLE_INTEGER

1 row selected.

BENCH@PMK SQL>
```

Some comments on Column F DOP

When DOP is set to 0 or 1

- No parallel query enabled
- No direct I/O enabled

When DOP is set > 1

- Parallel query and direct I/O enabled, even if DOP = 1 for following workloads
  - STO-READ
  - DA-STORAGE, DA-OFFLOAD
- Parallel query enabled
  - SRV-SCAN (no direct I/O because data is in row cache)

Manual Configuration

Manual peakmarks configurations support all kind of parallelism

```
BENCH@PMK SQL> @show_manual

Fri 31-Jan-2025 09:08:07

Manual peakmarks Configuration
------------------------------

Database....:  PMK               Oracle......:  19.25.0
Instance....:  ORA19C1           Build.......:  250201
RAC nodes...:  2                 Platform....:  peakmarks Ref System

                                                 Runtime
                    Para                          target
Test Workload      meter  ALC Nodes  Jobs  DOP    [min] Remark
---- ------------ ------ --- ----- ----- ---- ------- --------------------------------------------------
   1 STO-READ        N/A    0     1     1    1       3 sequential read          -> Intra-SQL  Parallelism
   2 STO-READ        N/A    0     1     1    2       3 sequential read
   3 STO-READ        N/A    0     1     1    4       3 sequential read          -> Inter-SQL Parallelism + Intra-SQL Parallelism
   4 STO-READ        N/A    0     1     2    4       3 sequential read
   5 STO-READ        N/A    0     1     4    4       3 sequential read
   6 STO-READ        N/A    0     2     4    4       3 sequential read          -> Inter-SQL Parallelism + Intra-SQL Parallelism + Cluster Parallelism
   7 STO-READ        N/A    0     2     8    4       3 sequential read

7 rows selected.

BENCH@PMK SQL>
```

Templates

The ../pmk/cfg directory contains some templates for all workload groups

Just adapt (according to your platform) one of these files copy it to manual.csv

Manual Configuration

Examples

- SQL> exec pmk.start_manual;

Use KPM reports to report the results.

peakmarks

Simple. Representative. Fast.

Summary of Scripts and Commands

Summary of Scripts and Commands

Manual peakmarks® configurations

SQL> @show_manual

SQL> exec pmk.start_manual;

Smart peakmarks® configurations

SQL> exec pmk.start_smart (p_wklgroup,
p_increment, p_parameter, p_comment)

Sample peakmarks® configurations

SQL> @show_sample

SQL> exec pmk.start_sample;

Monitoring queue of peakmarks® tests to process

SQL> @show_orders

peakmarks Mission

Identify Key Performance Metrics for Oracle Database Platforms.

On-Premises and in the Cloud.

For Quality Assurance, Evaluations, and Capacity Planning.