

Oracle Healthcheck Data collection Readme

June 2021

Copyright © 2021, Oracle and/or its affiliates
Public

Purpose Statement

This document provides an overview of the data collection process required for the Oracle Healthcheck service. It is intended solely to help you have a better understanding of the diagnostic tools used.

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation of the practices described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

Due to the nature of the products architecture and complexity, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

DATA COLLECTION README

AWR MINER

What is AWR Miner?

The AWR Miner script queries the AWR repository and gathers key database metrics from each AWR snapshot over the last n days. These metrics can assist in sizing and tuning exercises. An AWR report provides much more detailed metrics but only between 2 snapshots. This script and the subsequent graphs produced from its data provide higher level aggregate data over a longer period of time.

Will AWR Miner Impact My System?

It should have little to no impact on your system. It generally takes less than 45 seconds to run on anything from a laptop to an Exadata Full Rack and incurs less load than generating 1 AWR report. All queries are serial so at most this will occupy 1 CPU core for under a minute. Absolutely no objects are created in the Database.

Can I View the Output of an AWR Miner?

Absolutely. It's just a plain text file with standard reports composed of rows and columns. Ultimately this file is parsed and plotted with an R script by the Oracle team responsible with your Healthcheck report.

How do i Run the Script?

Unzip the contents of the AWR Miner archive in a preferred location on the database host.

If this is a RAC cluster you only need to run this from 1 node as the DBA_HIST_ views contain data from all nodes.

To avoid any SQLPATH environment variables which might cause problems (in case of login.sql files), you can optionally unset the SQLPATH shell environment variable before logging in to sqlplus:

```
[bash] unset SQLPATH
[windows] set SQLPATH=
```

Connect to the database you want to collect the AWR data from with Sqlplus as SYS, SYSTEM or a DBA user account.

Run from sqlplus command-line:

```
SQL> @awr_miner.sql
```



Is Diagnostic Pack License Required for AWR Miner and Statspack?

As long as the bellow process is followed (detailed in the MOS note), an exception is granted for running the scripts provided.

Please consult the following Knowledge articles for more details:

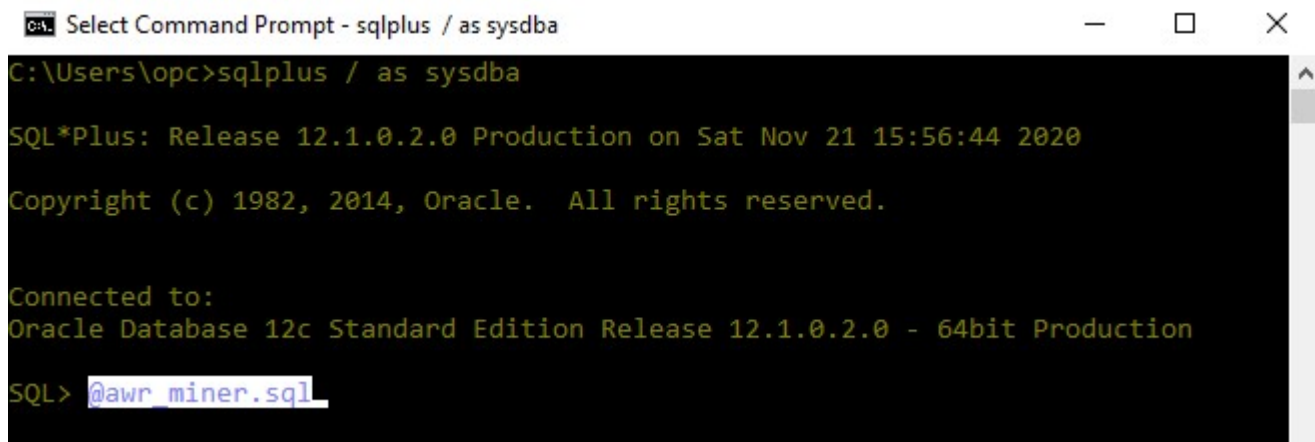
[Doc ID 2771065.1 OEM Metrics License Exception for Analysis by Oracle Pre-Sales or Consulting](#)

[Note: 1490798.1 – AWR Reporting – Licensing Requirements Clarification](#)

Will the AWR Miner and Statspack access my production data?

No, user data is not queried at all.

The AWR Miner script collects data from the Automated Workload Repository(AWR) for diagnostic purposes. It does not collect any data from user schema objects. Statspack uses the PERFSTAT schema which is created before running the snapshots and reports.



```

C:\Users\opc>sqlplus / as sysdba

SQL*Plus: Release 12.1.0.2.0 Production on Sat Nov 21 15:56:44 2020

Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Standard Edition Release 12.1.0.2.0 - 64bit Production

SQL> @awr_miner.sql

```

Wait for the script to finish. It could take longer over slower network connections. The last line displayed on the screen will be `~~END-TOP-SQL-BY-SNAPID~~`, then standard SQL prompt will return.

Do NOT SPOOL the output, it will automatically spool it to a file with a specific name (`awr-hist*.out`).

Do NOT CHANGE the file name as the code to parse this file is looking for files with a specific pattern

Send the output to your Oracle account manager and we will analyze and transform it into a Powerpoint presentation.

STATSPACK

What is statspack?

Oracle introduced Statspack as a diagnostic data collection tool. Along with additional reporting, statspack can store snapshots of system statistics over time, allowing greater accuracy and flexibility. Statspack comes packaged inside Oracle Database installation kit and does not require extra licensing.

How can i set up statspack?

All Statspack related information can be found in the following document, which can be found in each Oracle Database home:

`$ORACLE_HOME/rdbms/admin/spdoc.txt`

Installation of statspack involves the creation of the PERFSTAT user along with the necessary schema objects and the STATSPACK package. This is achieved by running the following script as SYS:

```
$ORACLE_HOME/rdbms/admin/spcreate.sql
```

Once installed you can connect to the PERFSTAT user and take a snapshot of the system statistics using the following procedure call:

```
SQL> EXEC STATSPACK.snap;
```

At a later time you can take another system snapshot, giving you a potential start and end point for your analysis. You can take multiple snapshots and use any for your start and end point.

The collection of system snapshots can be automated with the DBMS_JOB package. The spauto.sql script can be used to schedule system snapshot collections on the hour, every hour.

```
$ORACLE_HOME/rdbms/admin/spauto.sql
```

If you are automating snapshot collection you will need to delete snapshots from time to time. This can be done by running the sppurge.sql file as the PERFSTAT user. This script deletes a range of snapshots by prompting for the start and end points.

```
$ORACLE_HOME/rdbms/admin/sppurge.sql
```

Once you have at least two snapshots you can run the statspack report and find out the change in the statistics over the analysis period. The script prompts you for the start and end snapshots along with a filename for the output report.

```
$ORACLE_HOME/rdbms/admin/spreport.sql
```

ExaChk and OraChk (For Oracle hardware)

What are Exachk and Orachk?

Oracle Exachk and Orachk are diagnostics and reporting tools engineered specifically for Oracle Hardware.

On which systems i can use exachk and orachk?

You can use Oracle Exachk for all Oracle engineered systems except Oracle Database Appliance.

For Oracle Database Appliance, use Oracle Orachk.

You have access to Oracle ORAchk and Oracle EXAchk as a value add-on to your existing support contract. There is no additional fee or license required to run Oracle ORAchk and Oracle EXAchk.

How do I run these reports?

To run ExaChk reports, please follow the steps in the link below:

```
https://docs.oracle.com/cd/E68491\_01/OEXUG/oracle-exadata-and-zero-data-loss-recovery-appliance.htm#OEXUG-GUID-CDCA6A7B-1CCE-453E-9B08-784F2D497F89
```

To run OraChk reports, please follow the steps in the link below:

```
https://docs.oracle.com/cd/E68491\_01/OEXUG/automated-daemon-mode-operation.htm#OEXUG-GUID-2C32CB18-C837-40E8-AB6E-1FC6672A83E8
```

Troubleshooting AWR Miner

Below are some common issues encountered in gathering this data and how to potentially avoid them:

Invalid File Names

These errors occur when the name of the .out file is changed. **Do not change** the file name or extension. It should be in the form:

awr-hist-dbid-dbname-snapidmin-snapidmax.out.

If compressed using gzip, it should have a trailing .gz suffix.

Unparseable content

This almost always occurs when the .out file has been edited to search and replace for schema names or host names. The file layout is columnar so if you replace the schema name SCOTT with XX, that will shift the columns 3 positions to the left for those rows, making it impossible to parse.

Missing Information

This occurs when a collection does not have critical sections; `~~BEGIN-MAIN-METRICS~~` being one critical section.

Unexpected Format

This occurs when the collection script is different from the one on the website.

Truncated output

This has been a minor but persistent problem. The reason for this is unknown at this time, but could be due to a very slow connection to the server which causes longer than expected run times for the script. If possible, run the script directly from the database host. Make sure to wait for the script to end.

The last line will be `~~END-TOP-SQL-BY-SNAPID~~`.

There should be 9 or more sections which you can verify with grep (the number at the end of the file name is the count of sections).

Example:

```
grep -c "~~END" *.out
```

```
awr-hist-148658377-TWELVE1-1225-1320.out:12
```

```
awr-hist-3107787169-TENTWO-227-307.out:12
```

```
awr-hist-3107787169-TENTWO-251-330.out:12
```

```
awr-hist-3846920754-RAC81P-289-1209.out:12awr-hist-389926331-ORCL-775-985.out:10
```

Connect with us

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

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2021, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120