



# DB Sizing Toolkits

NetApp Solutions

NetApp  
July 31, 2024

# Table of Contents

- DB Sizing Toolkits ..... 1
- Oracle Sizing Guidance for Azure NetApp Files ..... 1

# DB Sizing Toolkits

## Oracle Sizing Guidance for Azure NetApp Files

Allen Cao, Niyaz Mohamed, NetApp

This solution provides an useful toolkit for sizing compute and storage for Oracle deployment on ANF in Azure cloud.

### Purpose

Moving existing Oracle workload from one platform to another, such as from on-prem to public cloud, needs sizing compute and storage in the target platform to meet performance and service level requirements. This documentation demonstrates a simple toolkit to accomplish that goal.

Unlike a new database application, which may grow over time, an existing Oracle workload has established workload patterns in compute and storage requirements, which are recorded in an Oracle Workload Repository or AWR. This toolkit utilizes an HTML parser to retrieve relevant information from Oracle AWR. The results are supplemented by additional sizing information obtained via SQL scripts against the database to provide meaningful compute and storage guidance when relocating the Oracle database.

This solution addresses the following use cases:

- Provide sizing guidance for Oracle database server compute when relocating database from on-prem to Microsoft Azure cloud.
- Provide sizing guidance for Oracle database server storage when relocating database from on-prem to Microsoft Azure NetApp Files.

### Audience

This solution is intended for the following people:

- A DBA who manages Oracle databases in on-prem private data center or Microsoft Azure cloud environment.
- A storage administrator who manages on-prem storage or Microsoft Azure NetApp Files storage that supports Oracle databases.
- An application owner who likes to migrate Oracle database from on-prem to Microsoft Azure cloud.

### License

By accessing, downloading, installing or using the content in this toolkit repository, you agree the terms of the License laid out in [License file](#).



There are certain restrictions around producing and/or sharing any derivative works with the content in this toolkit repository. Please make sure you read the terms of the License before using the content. If you do not agree to all of the terms, do not access, download or use the content in this repository.

## Solution deployment

### Prerequisites for deployment

Deployment requires the following prerequisites.

- Oracle AWR reports that capture the snapshots of database activities during peak application workload.
- Access to Oracle database to execute SQL scripts with DBA privilege.

### Download the toolkit

Retrieve the toolkit from repository [Oracle Sizing Guidance for ANF](#)

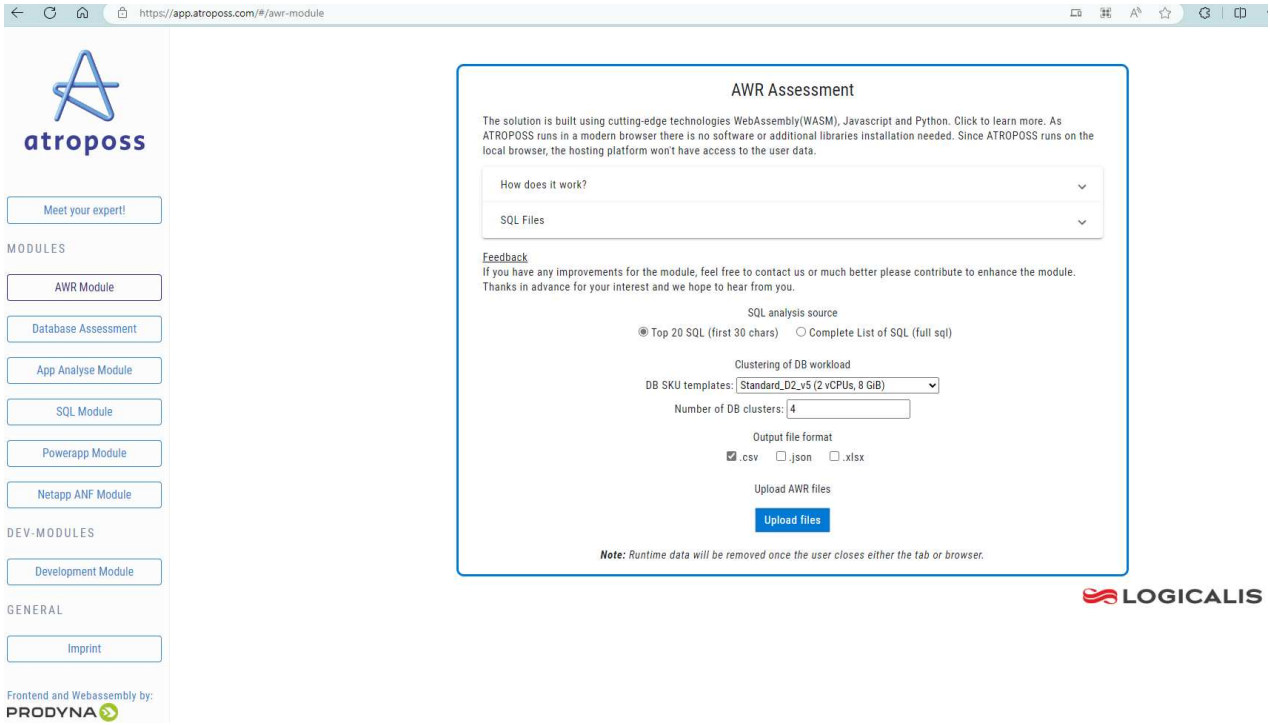
### How to use the toolkit?

The toolkit consists of a web-based HTML parser and two SQL scripts to gather Oracle database information. The output is then input into an Excel template to generate sizing guidance of computing and storage for the Oracle database server.

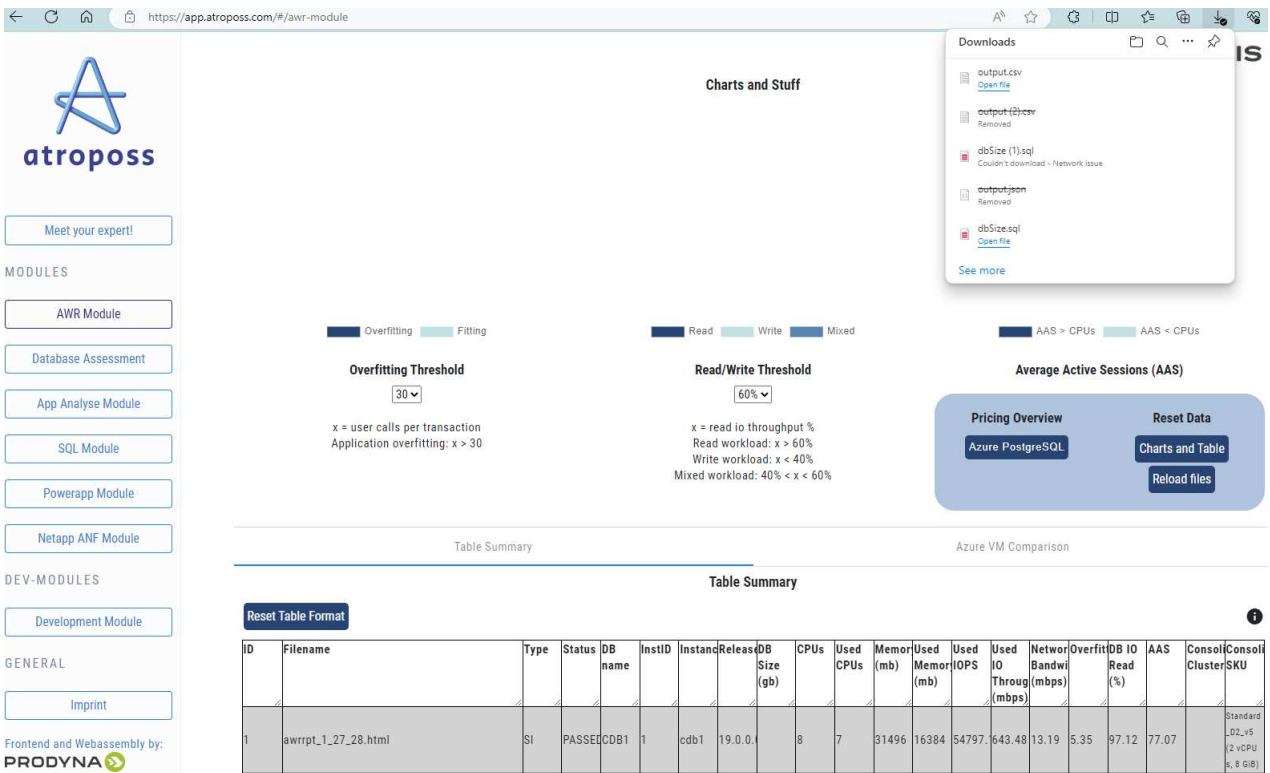
- Use an [HTML parser](#) AWR module to retrieve sizing information of a current Oracle database from an AWR report.
- Execute `ora_db_data_size.sql` as a DBA to retrieve physical Oracle data file size from database.
- Execute `ora_db_logs_size.sql` as a DBA to retrieve Oracle archived logs size with desired archive logs retention window (days).
- Input sizing information obtained above into excel template file `oracle_db_sizing_template_anf.xlsx` to create a sizing guidance on compute and storage for Oracle DB server.

### Toolkit usage demonstration

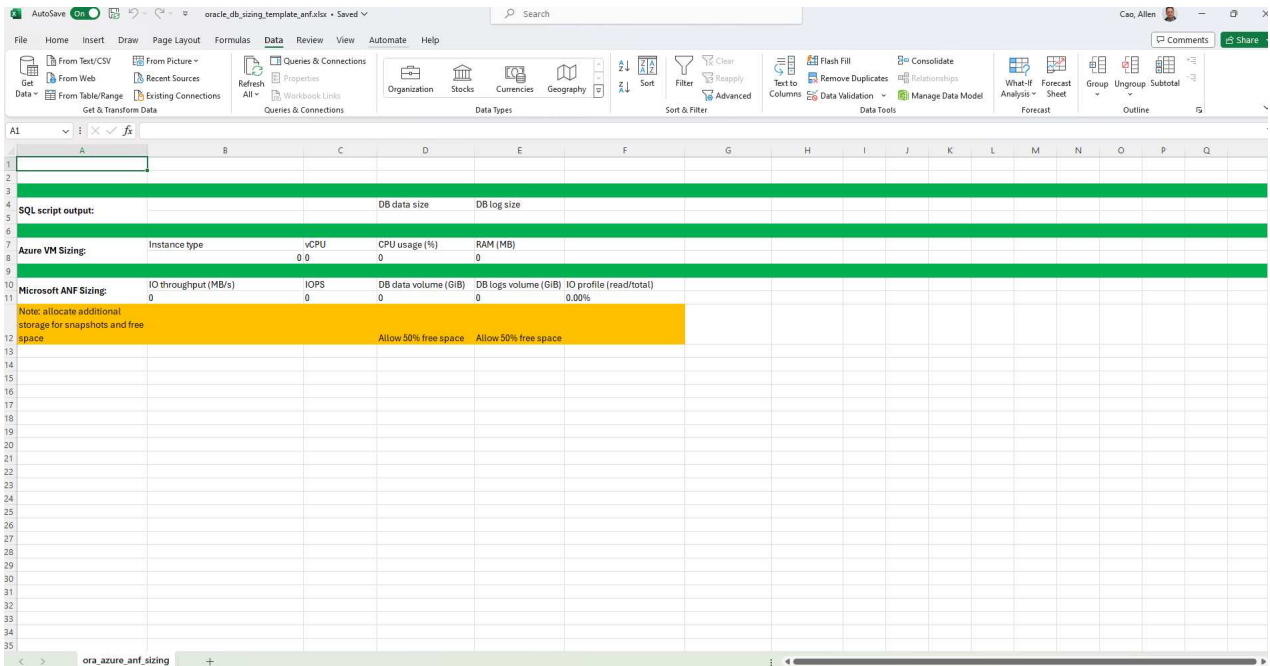
1. Open HTML parser AWR module.




2. Check output format as .csv and click Upload files to upload awr report. The parser returns results in a HTML page with a table summary as well as an output.csv file in Download folder.



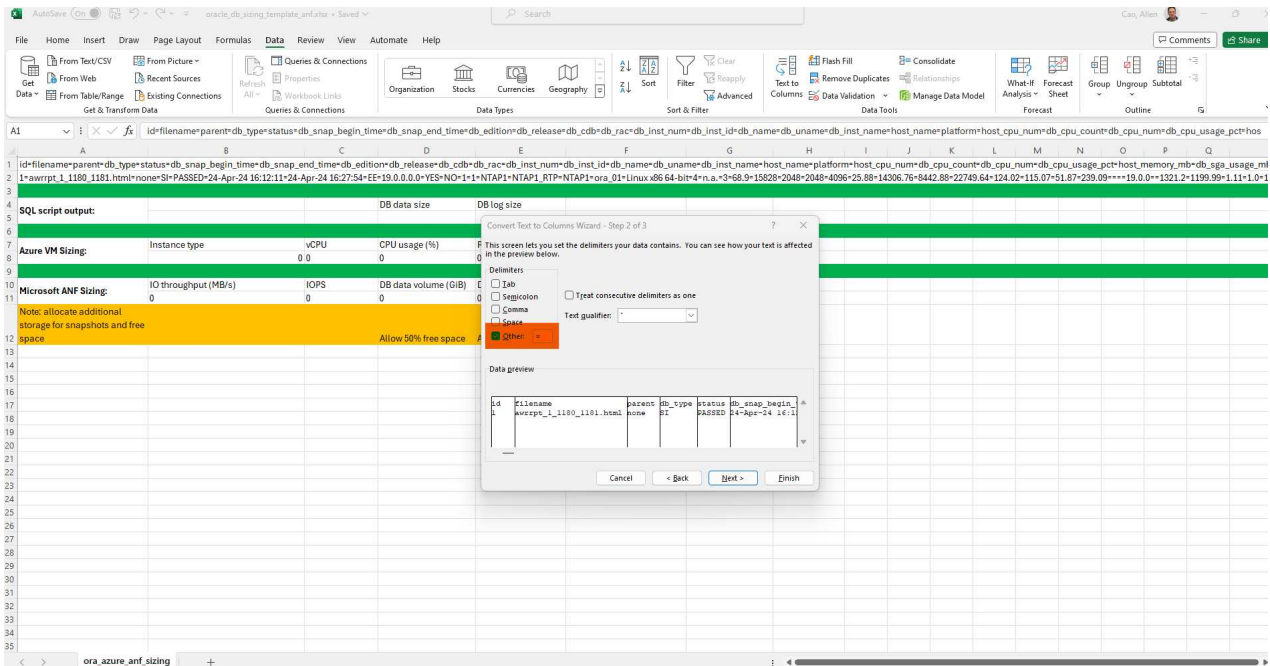
3. Open the excel template file and copy paste the csv content into column A and cell 1 to generate the DB server sizing information.




4. Highlight column A and fields 1 and 2, click on Data, then Text to Columns to open the Text Wizard. Choose Delimited, then Next to next screen.



5. Check Other, then enter '=' as Delimiters. Click on Next to next screen.



6. Click on **Finish** to complete the string conversion into readable column format. Note the VM and ANF sizing fields have been populated with data retrieved from the Oracle AWR report.



| A  | B  | C                         | D                    | E                    | F                       | G                  | H                | I          | J                | K      | L          | M          | N       | O       | P         | Q        |          |           |
|----|--|---------------------------|----------------------|----------------------|-------------------------|--------------------|------------------|------------|------------------|--------|------------|------------|---------|---------|-----------|----------|----------|-----------|
| 1  | id   | filename                  | parent               | db_type              | status                  | db_snap_begin_time | db_snap_end_time | db_edition | db_releas_db_cdb | db_rac | db_inst_ni | db_inst_ic | db_name | db_unam | db_inst_n | host_nam | platform |           |
| 2  |  | 1 awrrpt_1_1180_1181.html | none                 | SI                   | PASSED                  | 4/24/2024 16:12    | 4/24/2024 16:27  | EE         | 19.0.0.0.0       | YES    | NO         | 1          | 1       | NTAP1   | NTAP1_RT  | NTAP1    | ora_01   | Linux x86 |
| 4  | SQL script output:   |                           |                      | DB data size         | DB log size             |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
| 7  | Azure VM Sizing:   |                           |                      |                      |                         |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
|    | Instance type  | vCPU                      | CPU usage (%)        | RAM (MB)             |                         |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
|    | SI   | 4                         | 68.9                 | 15828                |                         |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
| 10 | Microsoft ANF Sizing:  |                           |                      |                      |                         |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
|    | IO throughput (MB/s)   | IOPS                      | DB data volume (GiB) | DB logs volume (GiB) | IO profile (read/total) |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
|    | 239.09   | 22749.64                  | 0                    | 0                    | 62.89%                  |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |
| 12 | Note: allocate additional storage for snapshots and free space |                           |                      | Allow 50% free space | Allow 50% free space    |                    |                  |            |                  |        |            |            |         |         |           |          |          |           |

7. Execute script ora\_db\_data\_size.sql, ora\_db\_logs\_size.sql as a DBA in sqlplus to retrieve existing Oracle database data size and archived logs size with the number of days of retention window.



```

[oracle@ora_01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Mar 5 15:25:27 2024
Version 19.18.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version 19.18.0.0.0

SQL> @/home/oracle/ora_db_data_size.sql;

Aggregate DB File Size, GiB Aggregate DB File RW, GiB Aggregate DB
File RO, GiB
-----
-----
                                159.05                                159.05
0

```

```

SQL> @/home/oracle/ora_db_logs_size.sql;
Enter value for archivelog_retention_days: 14
old 6:      where first_time >= sysdate -
&archivelog_retention_days
new 6:      where first_time >= sysdate - 14

Log Size, GiB
-----
          93.83

SQL>

```



The database sizing information retrieved using above scripts is the sum of actual size of all physical database data files or log files. It does not factor into the free space that may be available inside each data file.

8. Input the result into excel file to complete the sizing guidance output.

| id   | filename                | parent | db_type              | status      | db_snap_begin_time   | db_snap_end_time     | db_edition              | db_releas  | db_cdb | db_rac | db_inst_ni | db_inst_ic | db_name | db_uname | db_inst_ni | host_nam | platform  |
|--|-------------------------|--------|----------------------|-------------|----------------------|----------------------|-------------------------|------------|--------|--------|------------|------------|---------|----------|------------|----------|-----------|
| 1  | awrrpt_1_1180_1181.html | none   | SI                   | PASSED      | 4/24/2024 16:12      | 4/24/2024 16:27      | EE                      | 19.0.0.0.0 | YES    | NO     | 1          |            | 1.NTAP1 | NTAP1_RT | NTAP1      | ora_01   | Linux x86 |
| SQL script output:   |                         |        | DB data size         | DB log size |                      |                      |                         |            |        |        |            |            |         |          |            |          |           |
|  |                         |        | 159.05               | 93.83       |                      |                      |                         |            |        |        |            |            |         |          |            |          |           |
| Azure VM Sizing:   |                         |        | Instance type        | vCPU        | CPU usage (%)        | RAM (MB)             |                         |            |        |        |            |            |         |          |            |          |           |
|  |                         |        | SI                   | 4           | 66.9                 | 15828                |                         |            |        |        |            |            |         |          |            |          |           |
| Microsoft ANF Sizing:  |                         |        | IO throughput (MB/s) | IOPS        | DB data volume (GiB) | DB logs volume (GiB) | IO profile (read/total) |            |        |        |            |            |         |          |            |          |           |
|  |                         |        | 239.09               | 22749.64    | 318.1                | 187.66               | 62.89%                  |            |        |        |            |            |         |          |            |          |           |
| Note: allocate additional storage for snapshots and free space |                         |        | Allow 50% free space |             | Allow 50% free space |                      |                         |            |        |        |            |            |         |          |            |          |           |

9. ANF uses a three-tier service level (Standard, Premium, Ultra) to manage database volume throughput limit. Refer to [Service levels for Azure NetApp Files](#) for details. Based on sizing guidance output, choose an ANF service level that provides throughput that meet the requirement for the database.

## Where to find additional information

To learn more about the NetApp database solutions, review the following website [NetApp Enterprise Database Solutions](#)

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.