

NetApp Takes-on Oracle Databases

A Study of Oracle User Experiences with a Cross-functional, High Performance Storage Platform

By John Webster

July 2020



Evaluator Group

Enabling you to make the best technology decisions

Executive Summary

Despite the broad diversification of database technologies available today that have been driven by open source software development, most enterprises—from no matter what industry segment—continue to use Oracle database solutions to support their core production systems. We believe this is because Oracle continues to offer, advance and support one of the most trusted and widely used relational database engines available for enterprise computing. The proof: Oracle databases are now used by approximately 310,000 customers worldwide.

As Oracle-based applications continue to evolve in the presence of hybrid cloud and application containerization initiatives, enterprise IT still needs a stable, high performance and cost-efficient storage environment to deliver Oracle applications to business users. One of these is NetApp with its portfolio of data-centric solutions that include ONTAP-based, fabric-attached storage systems, HCI platforms and cloud storage offerings.

Here we review NetApp's storage systems as a storage platform for Oracle applications. We do so by first noting and addressing storage requirements. We then look at the data management aspects of NetApp's ONTAP storage operating system to automate data protection and data distribution to applications manager, database administrators and data scientists. Finally, we present our findings from end-user interviews of an Oracle database environment supported by NetApp FAS storage systems as seen in real-world applications.

NetApp ONTAP Systems

Performance Fundamentals for Oracle Database

NetApp offers all-SSD (AFF) and Fabric-Attached (FAS) storage systems for use with critical production Oracle database applications. NetApp refers to fabric-attached storage systems as storage connected to hosts via any network – FC SAN for block or IP NAS for file. All Flash FAS (AFF) systems have embedded software to optimize use of flash. The A800 model group, released in May 2018, includes support for NVMe over Fabric (Fibre Channel with Gen 6 at 32Gb/s initially with Ethernet RDMA at 100Gb/s to follow) and NVMe devices, including support for 30TB SSDs.

NetApp AFF and FAS systems offer excellent performance for block and file storage and retrieval for mixed types of files¹. The most important storage system performance consideration for most database applications is the delay, or latency incurred in the I/O path rather than the number of I/O operations a storage system can deliver. To minimize storage path latency, ONTAP systems include all-flash storage architectures. New NVMe-based systems with NVMe over fabric connectivity are also available to deliver the low latencies required by critical database workloads.

¹ SPEC benchmarks are available that detail block and file performance.

Additionally, ONTAP systems support multiple protocols simultaneously, including Oracle's direct NFS client for multiple network ONTAP connections, increasing I/O performance, and enabling multipathing for fault tolerance, similar to FC multipathing. NetApp has tested multiple systems with Oracle and found performance for all-flash AFF A800 of more than 2M IOPs at less than 200 micro-seconds of latency.²

ONTAP in an Oracle Environment

ONTAP is NetApp's operating system for FAS storage. ONTAP implements a global namespace across all storage nodes in a scale-out FAS cluster which simplifies access to files or LUNs. Performance and capacity can be scaled simultaneously with the addition of nodes to the cluster. This allows customers to maintain performance as the demands for increased capacity are met.

ONTAP systems feature an operational environment called Cluster Mode. In Cluster Mode, non-disruptive operations are supported. This allows IT operations staff to perform critical management functions from a single management interface while maintaining continual Oracle database application availability. These include:

- Automated load balancing via redistributing active data across different nodes
- Automated capacity balancing when new nodes are added to the cluster
- Transparent data migration when nodes are removed from the cluster
- Non-disruptive software updating and patching
- Non-disruptive servicing
- Space saving features including compression/compaction and deduplication

Leveraging Public Cloud

NetApp offers two options for cloud storage usage to improve Oracle database performance, data protection and disaster recovery capabilities vs. providing these services locally. NetApp Cloud Volumes ONTAP, a user-managed software appliance built on NetApp ONTAP, delivers advanced capabilities that allow IT administrators to scale database capacity, improve resource utilization, and increase DBA productivity. Cloud Volumes can be deployed in a hybrid cloud architecture for augmented data protection and disaster recovery capabilities leveraging public cloud storage. For added simplicity and performance, Azure NetApp Files, a native Azure storage service, and Cloud Volume Service for AWS, a marketplace fully managed offering, accelerate Oracle database applications running in the cloud.

² NetApp AFF A800 with Oracle RAC Database - <https://www.netapp.com/us/media/tr-4767.pdf>

High-value Oracle Data Management Capabilities

While the hardware aspects of NetApp storage systems deliver storage performance for Oracle database applications, it's the data management tools and features built into the storage system's embedded ONTAP software that users have found to be highly advantageous as well. Evaluator Group has seen that the wealth of these advanced features is greater than any other competing storage system, particularly when used for storage efficiency and data protection/business continuance practices. These include:

SnapShot

ONTAP SnapShots are a copy of pointers to a data set rather than the actual data itself. As such, they are highly space efficient and created quickly. Typically, a point-in-time SnapShot copy of an Oracle database is created in less than 5 seconds. Only data that has changed from the original SnapShot is used to update subsequent copies. Administrators can utilize up to 255 separate SnapShot instances, which may be taken at pre-scheduled intervals or as needed at other times. Access to the original data set continues while the SnapShot is being created or updated.

SnapMirror

SnapMirror is used to create full data set copies. Once a mirror copy is made, it is kept current by updating it with only new or changed data blocks. Like the base Snapshot feature, access to the original data set continues. SnapMirror allows administrators to control update frequency down to as short as one-minute intervals.

SnapCenter

SnapCenter provides application-consistent management GUI for ONTAP SnapShot and FlexClone data copies. It provides monitoring, administrative notification, logging, reporting, automated backup scheduling and clone management. In an Oracle environment it can be used to distribute data copies for development/test and DevOps teams as well as disaster recovery implementations and AI use cases. In addition, IT operations administrators can enable self-service SnapCenter access for application and database administrators such that they can initiate FlexClone volumes for example for development and testing independent of storage administrators.

FlexClone

A FlexClone volume is a writable, point-in-time image of a data set which can be created in a few seconds, and initially is based on the underlying volume. Similar to snapshots, FlexClone volumes do not alter the underlying data and do not interrupt access to the volume being cloned. Use cases in an Oracle environment include:

- Testing and development using production data but without endangering the integrity of the production volume
- Distributing data (to DBAs for example) in changeable form but without endangering the integrity of the original

- Initiating parallel independent operations using identical temporary data

Using ONTAP to Maximize Oracle Storage and Operational Efficiency

We heard from NetApp users we interviewed for this study that the aforementioned ONTAP data management features can be used by other departmental entities within IT such as developers, database administrators and data scientists to deliver Oracle data sets in the course of their normal activities. In a NetApp ONTAP storage environment, one way to do this while preserving performance for primary Oracle applications is via the use of ONTAP SnapMirror and FlexClone in combination. In this scenario (shown in Figure 1 below), a storage administrator or an authorized user (DBA, developer, data scientist) makes a physical copy using of the Oracle production data using SnapMirror. This physical “Parent” copy is then used as a source for the distribution of virtual copies. SnapMirror creates the Parent copy; FlexClone creates and manages the distribution of the virtual copies that are taken from the Parent.

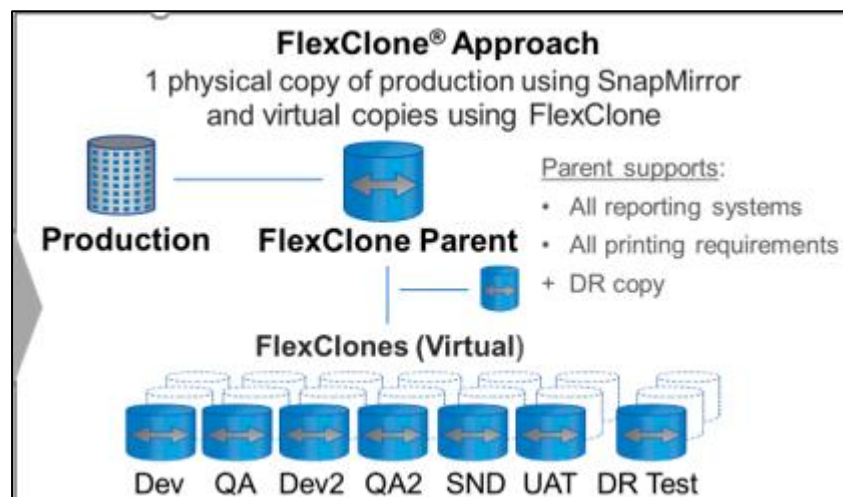


Figure 1. Using ONTAP SnapMirror and FlexClone to create and distribute multiple copies of Oracle production data sets

Using NetApp’s data copy technologies offers a number of advantages in an Oracle database application environment:

1. Performance as seen by the primary application users is maintained without having to add storage capacity in order to make and deliver up to 100 or more clones of the same production database. One user has reported an 81% savings in storage capacity using this process vs. using physical copies to accomplish the same end. Data center space, power and cooling is also saved.
2. Usable database copies are typically created in 5 minutes meaning that developers and DBAs can shorten development and test cycles.

3. Disruption to the production operations environment is minimized. Availability of 99% or better has been achieved by one user we spoke to.

Storage System-Internal Data Protection for Oracle

We have also reviewed NetApp storage in the context data protection for Oracle environments. We have seen that NetApp ONTAP systems have been utilized by IT administrators and DBA's to deliver high-performance storage to Oracle database applications along with efficient and flexible data protection options. NetApp SnapCenter is one of the most important tools for coordinating data protection for ONTAP systems. With a specific plug-in for Oracle, SnapCenter is able to catalog the backup data so that an administrator can insert and remove data protection points. Additional features include the capability to create space-efficient point-in-time copies of production databases using NetApp FlexClone for rapid, non-disruptive cloning of databases.

Snapshots are the basis for data protection. With ONTAP, recovery using snapshots has been simplified with Oracle version 12c and newer, as there is no need to place a database in hot backup mode. ONTAP snapshots may be scheduled while still preserving the ability to perform complete or point-in-time recovery, due to the new Oracle ability in RMAN to replay archive logs on datafile backups that were not in hot backup mode.

The basis for data protection with ONTAP includes:

- Creation of PIT Snapshots
- Cataloging of protection points (coordinated with Oracle RMAN via SnapCenter)
- Restoration via SnapRestore (also utilizing Oracle RMAN via SnapCenter)
- Replication options including SnapMirror, SyncMirror and MetroCluster
- Support for Oracle Data Guard for offline and online protection

Data protection in ONTAP also includes support for the use of consistency groups in both Oracle standalone and RAC environments with multiple storage connectivity options on virtual and physical deployments. Replication support for Oracle includes synchronous, asynchronous with SnapMirror and NetApp MetroCluster. MetroCluster delivers zero data-loss for mission-critical workloads and can replace multiple data protection strategies with one system that provides backup, recovery, DR and high availability (HA) using a clustered ONTAP system spanning two geographical regions.

User Experience

We interviewed two NetApp users to better understand real-world implementations of NetApp storage based on ONTAP data services in critical, production-ready Oracle database environments. One is a

large, multi-region healthcare provider with over 300 offices across four states. The other is a large developer and supplier of IT infrastructure.

Large Healthcare Provider

It is common to find healthcare providers using Epic Electronic Health Record (HER) solutions. This is the case with the healthcare provider we spoke to. In this case, NetApp storage was used to support a 100TB instance of Epic Clarity running on Oracle. We spoke to the lead storage administrator as well as the head of the database administration team about this application as well as other processes enabled by NetApp hardware and software technology.

The Clarity database is a large subset of data that is generated by Epic systems and used as a data warehouse. Data is transferred by an Export/Transform/Load (ETL) process to the Clarity database. Using Clarity, this customer runs complex, data intensive reports daily. Epic Clarity plays a critical role in the overall functioning of medical services delivery to the communities served.

Data Protection and Database Copy Distribution

Prior to NetApp, the IT organization was using a large-scale, high-performance system from a leading vendor of storage arrays to support Clarity on Oracle. What led them to look seriously at replacing this array was the fact that Oracle RMAN backups ran “forever.” As a result, the organization was at risk of losing data when backup processes went unfulfilled. The replacement of this array with NetApp allowed them to replace the inadequate data streaming backup process with snapshots generated by ONTAP. This allowed storage administration to assure not only the completeness of backup processes by rapid restores when called for.

This customer was also using RMAN to create database copies for use by DBAs. Again, this process became highly inefficient and time consuming, reducing the productivity of the organization’s DBAs. FlexClone is now used to create and distribute database copies in a matter of “minutes.” DBAs now provision their own clone copies without continually requesting them from storage administration, allowing them to focus on other more pressing matters. Both the NetApp Snapshot backup and the FlexClone copy processes now complete within the required time windows.

Compression and Encryption

Two other capabilities of the NetApp ONTAP storage platform are highly valued by this healthcare provider. They appreciate and fully exploit the data reduction and compression services resident in ONTAP. These are seen as cost-avoidance factors, allowing them to upwardly scale the capacity of the Clarity storage environment in a cost-efficient way. And, because healthcare systems are highly regulated, encryption is also enabled for this application environment. All of this is handled concurrently by ONTAP while delivering to their required performance levels.

Technology Infrastructure Supplier

The second interview was with a storage architect for a large technology infrastructure supplier of hardware, software and cloud computing resources. Internal IT for the infrastructure supplier is tasked with developing and deploying new business features while maintaining reliable access to existing IT services for business user groups that also include outside business partners. The company's Oracle database environment must support IT's mission without compromising application stability. In addition, IT tests and implements new products, promotions or business workflows that require development and integration work. To meet a six-week release cycle, IT runs parallel development and iterative test cycles requiring multiple application development and test environments.

The company operates five data centers with 100 PB of raw storage supporting 375 enterprise applications and 418 databases, including production and non-production. Oracle applications are used for ERP, HR, SAP, Oracle BI, Oracle SOA, and Oracle ECM applications. All the applications use Oracle databases Single Instance or RAC mode.

NetApp ONTAP software-based features and functions have become critical to supporting their Oracle databases. These include SnapShot, SnapMirror and FlexClone. As noted above, SnapMirror is used to make a copy of the production databases. FlexClone then uses this copy as the "Parent FlexClone" to 1) create numerous virtual copies of the database, 2) create a DR copy, and 3) support all reporting and printing needs. Use of the "Parent" removes potential risk of human error on the production databases, resulting in a more stable operational environment. Creating a dedicated virtual copy of the database also provides them with the ability to isolate and test new business functionalities that might negatively impact the application ecosystem.

Data Protection

The ONTAP software used for data protection and disaster recovery includes:

SnapShot – Local snapshots of an Oracle database are used for operational backup and recovery. Snapshots are taken every 4 hours and retained for seven days

Remote SnapMirror – A copy of operating systems and applications are created every 24 hours; databases every four hours; transaction and archive logs every 10 minutes

SnapVault – For frequent audit and discovery requests against critical application data, SnapVault is used for short term archive. Copies are created on a daily basis and retained for 60 days.

The company also uses NetApp StorageGrid object storage as a private S3 cloud for three to seven-year retention.

Using the data protection processes described above, the company sees the following completion times (based on a 10 TB Oracle database):

- Database backup – 5 min.
- Database restore – 5 min.
- System refresh – 15 min.
- DR synchronization – 10 min.

SnapMirror is also used for on-site storage to storage device migrations as well as remote data center migration when required. In these cases, completion time for the on-site migration (10 TB database) is approximately 1 hour; two hours for a data center migration effort.

Final Thoughts

NetApp has a long history of working closely with Oracle to deliver the storage and data management features that IT administrators and DBAs need for critical applications. This has resulted in the development of a storage environment that caters to Oracle database applications. We also note that Oracle database certification is an important factor among users who are deciding which vendors to include in their Oracle environments.

We have seen that hardware performance is more than adequate to support Oracle's critical, high transaction volume workloads. And while supporting these workloads, administrators are also able to provision and distribute data copies, apply data reduction and encryption features, and meet data protection targets. Indeed, the first user we interviewed was able to replace an inadequate data streaming backup process with snapshots generated by ONTAP. This allowed storage administration to assure not only the completeness of backup processes by rapid restores when called for. In our customer interviews we also saw that the data management features in ONTAP can be used to streamline operational support by storage administrators and accelerate the productivity of application managers and developers.

While every enterprise data center is different, we believe that consolidating Oracle database storage to NetApp storage platforms will save administrative time while satisfying the needs of Oracle database users, developers and DBAs. This will equate to cost savings that will justify the effort in replacing an inadequate storage platform with NetApp AFF or FAS.

This paper was sponsored by NetApp, Inc.

About Evaluator Group

Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. www.evaluatorgroup.com @evaluator_group

Copyright 2020 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.