Technical Report

NetApp Extreme Analytics Solution for Oracle

Bobby Oommen, Jeff Steiner, NetApp
March 2013 | TR-4145

Executive Summary

This document explains the NetApp® Extreme Analytics solution for Oracle® databases and how NetApp can provide a faster, more efficient, easier managed, and lower-cost engineered solution that meet the needs of an enterprise. This architecture provides the superior performance and storage efficiency that customers have come to expect from NetApp by using a combination of NetApp EF540 flash array and NetApp FAS storage.

TABLE OF CONTENTS

LIST OF FIGURES

1 Introduction

This solution is based on the combination architecture of the NetApp EF540 flash array (referred to as the EF540 from here on) and the NetApp FAS array with SAS disk drives. The EF540 provides ultrahigh application performance and the NetApp FAS array with SAS drives provides Data ONTAP® manageability and efficiency features such as Snapshot™-based backups, recovery, cloning, and replication.

2 Audience

This document is intended for the following audiences:

- People familiar with NetApp FAS storage but who want extreme I/O with submillisecond latency for very specific workloads and are open to the idea of a combination NetApp storage solution based on FAS along with NetApp's new EF540 flash array.
- People familiar with Oracle Database technologies and features.

3 Solution Components

3.1 NetApp EF540 Flash Array

The NetApp EF540 flash array delivers high reliability and performance for enterprises running mission-critical applications with low latency and high IOPS requirement. The EF540 flash array, available with either 9.6TB or 19.2TB of raw capacity, provides performance to meet the requirements of the most demanding applications. The managing and provisioning of storage are done using NetApp SANtricity® storage management software.

3.2 NetApp FAS Array

The NetApp Data ONTAP 8 architecture offers a storage-OS platform that helps address the challenges of growing and dynamic businesses. NetApp's storage efficiency features help IT to easily define end-to-end data protection strategies. The unified storage architecture is one of the key design features in the NetApp FAS array that helps businesses consolidate IT operations, increase efficiency, and magnify data center savings in terms of power, cooling, and space utilization. NetApp Flash Cache™ intelligent caching also helps with savings by requiring fewer disk spindles without compromising application performance.

3.3 NetApp OnCommand Insight

NetApp OnCommand® Insight enables customers to manage their multivendor storage environment as an integrated, end-to-end service. OnCommand Insight is open and integrates with third-party configuration management databases (CMDBs) to extend value across solutions so that customers can better manage and optimize existing resources to drive greater IT efficiency and cost savings. OnCommand Insight lets customers:

- Achieve more accurate capacity forecasting with trending.
- Perform cost reporting for showback and chargeback.
- Plan for consolidation, virtualization, and cloud initiatives.
- Set policies and best practices for operational excellence.
- Gain visibility into and control of heterogeneous environments.

3.4 Snap Creator Framework

NetApp Snap Creator™ Framework is used to manage the backup, recovery, and cloning of the Oracle Database through the Oracle ASM fail group that is configured on the NetApp FAS array. The purpose of Snap Creator is to create a central framework that provides seamless integration with applications such as Oracle Database and NetApp Snapshot technology. The Oracle plug-in that comes with Snap Creator Framework handles application consistency during a Snapshot operation. Snap Creator communicates with NetApp FAS storage and performs various tasks, including policy-based Snapshot management and integration with other NetApp products, especially SnapMirror® and SnapVault® technologies. In the event of data corruption Snap Creator can be used to recover the database to a desired point in time.

4 Architecture

This solution is designed to help customers looking for extreme application performance as well as ultra efficient storage manageability features. NetApp EF540 flash array can deliver high IOPS with submillisecond latency and is suited for customers looking for transactional database-driven applications where responsiveness is critical. The NetApp FAS storage array in combination with the industry-leading Data ONTAP storage operating system lets customers define service layers in the context of data protection through Snapshot, SnapMirror, and SnapVault technologies.

Data from the EF540 to the FAS array can be replicated by using replication methodologies like LVM or ASM mirroring, Data Guard, or other vendors' technologies. For this solution the data replication was performed by using Oracle ASM disk mirroring, which uses two sets of storage arrays: a NetApp EF540 and a NetApp FAS array with SAS drives. The Oracle ASM disk groups verify that the two different arrays are bonded together. This design guarantees 100% synchronous mirroring and mirroring rebuilds in the event of a serious hardware failure. Customers have the flexibility to choose any servers and NetApp storage arrays to use in this solution based on capacity and throughput/IOPS requirements.

The Oracle ASM fail groups are set up in such a way that reads are pinned on the EF540. The layout of the Oracle Database is explained in section 4.5, "Database Layout." The Oracle ASM mirror copy on FAS storage is used for data manageability using NetApp Snapshot features and is explained in section 3.4, "Snap Creator Framework." In the event of database recovery, the Oracle ASM fail groups are rebalanced to keep the mirror copies in sync.

4.1 Solution Advantages

1. Customers can create application-consistent backups instantaneously in conjunction with the NetApp Snap Creator Framework, which creates the Snapshot copies on the NetApp FAS storage array where the second Oracle ASM fail group is configured.
2. Database recovery can be performed quickly through a NetApp SnapRestore® operation. The application can be up and running temporarily on the second Oracle ASM fail group while the rebalancing takes place on the primary fail group, which is on the EF540.
3. The NetApp FAS storage can be configured to replicate data to a remote location by using SnapMirror or SnapVault. Also, FlexClone® volumes can be used to create any number of database copies on the FAS storage array to provision development and test environments.
4. All the read workloads are serviced from the EF540 Oracle ASM disk group fail group through the ASM_PREFERRED_READ_FAILURE_GROUPS setting. The capability of Data ONTAP to coalesce writes and distribute I/O over numerous spindles prevents write performance limitations.
5. The solution is simple and uses standard Oracle ASM mirroring that can be managed easily by DBAs.

Figure 1) Oracle on the NetApp EF540.

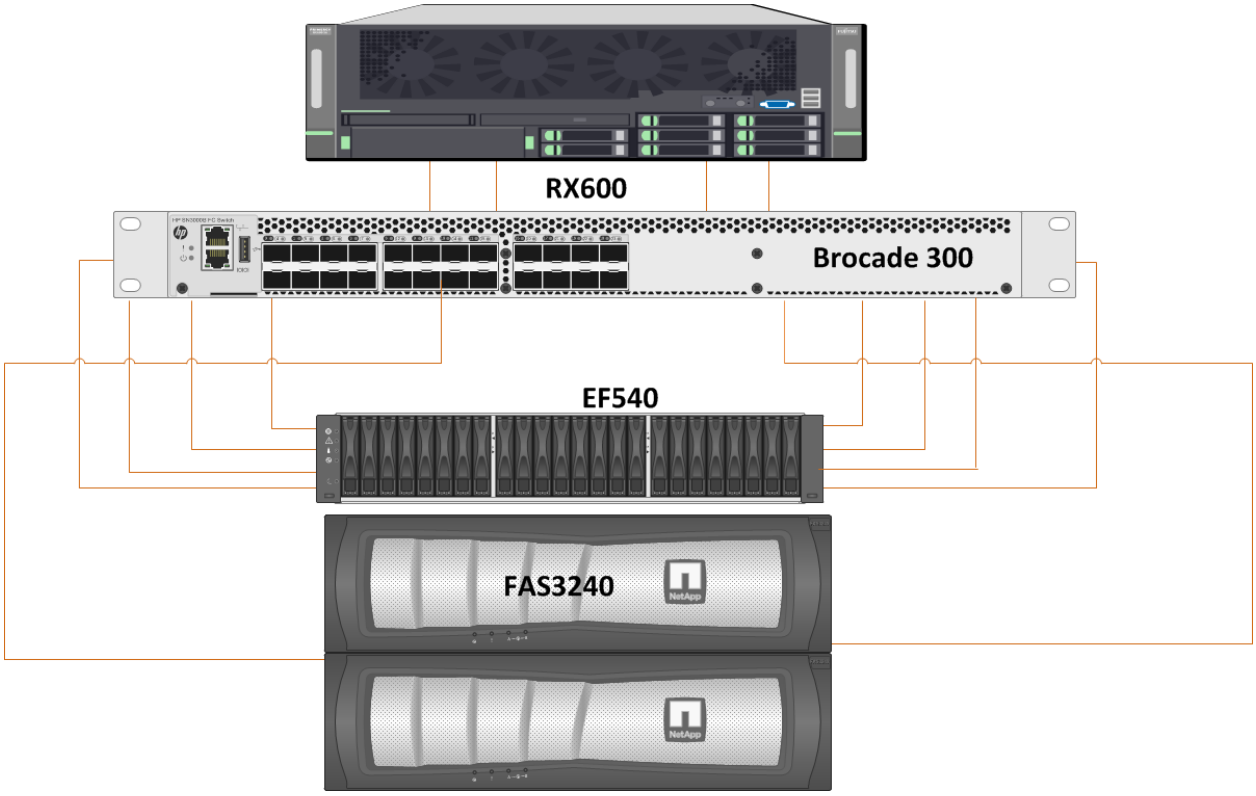

4.2 Hardware Components

- Fujitsu RX600, 512GB RAM, 2 x 8Gb FC dual port
- Brocade 300 8Gb fabric switch
- NetApp EF540 with 24 SSDs
- FAS3240A storage array, 450GB x 48 drives, Flash Cache
- Two Windows Server® 2008 servers for OnCommand Insight systems and Snap Creator Framework

4.3 Software Components

- Oracle Database 11*g* R2 (11.2.0.3.0)
- NetApp OnCommand Insight 2.0
- Data ONTAP 8.1.1 operating in 7-Mode
- SANtricity 10.83
- NetApp Snap Creator Framework 3.6

4.4 Design Considerations

Performance considerations for the workload test were:

- Queue depth. The queue for the disk devices at the Linux® host was changed to 256.
- Scheduler. The Linux I/O scheduler was set to NOOP.
- Multipathing. DM-MP was used to create redundancy and load balancing across all paths.

4.5 Database Layout

When executing any workload test, one of the most important factors is the locality of reference, which is known as the working set. The working set is the amount of data that has been active and been accessed for read and write operations during a synthetic workload.

For an accurate and valid workload test, the locality of reference should have a reasonable size that should be bigger than the memory in the servers and in storage so that all writes and reads go to the disks and the amount of cache used is small. The locality of reference used on this workload was large enough for all reads and writes to go to disk, with less than 3% of a cache hit to storage.

To obtain optimal results, some of the tuning parameters were changed during the database setup and are outlined in section 4.4, “Design Considerations.”

Storage Arrays

- NetApp EF540 flash array. One volume group with 64 volumes (each 30GB) for data and 4 volumes (each 20GB) in the same volume group for logs. One volume (3GB) for the GRID Oracle engine.
- NetApp FAS3240A (Oracle ASM mirror group). One volume with 32 LUNs for data and one volume with 2 LUNS for logs on each controller.

Database Server

The database server included one ASM disk group for each GRID (one LUN). The ORADATA Oracle ASM disk group was created with two fail groups (see Figure 2): one with LUNs from the EF540 and the other with LUNs from the NetApp FAS array. The ORALOG Oracle ASM disk group was also created with two fail groups each, with four LUNs accessing the EF540 and FAS array.

Figure 2) Oracle ASM fail groups on the NetApp EF540 and FAS array.

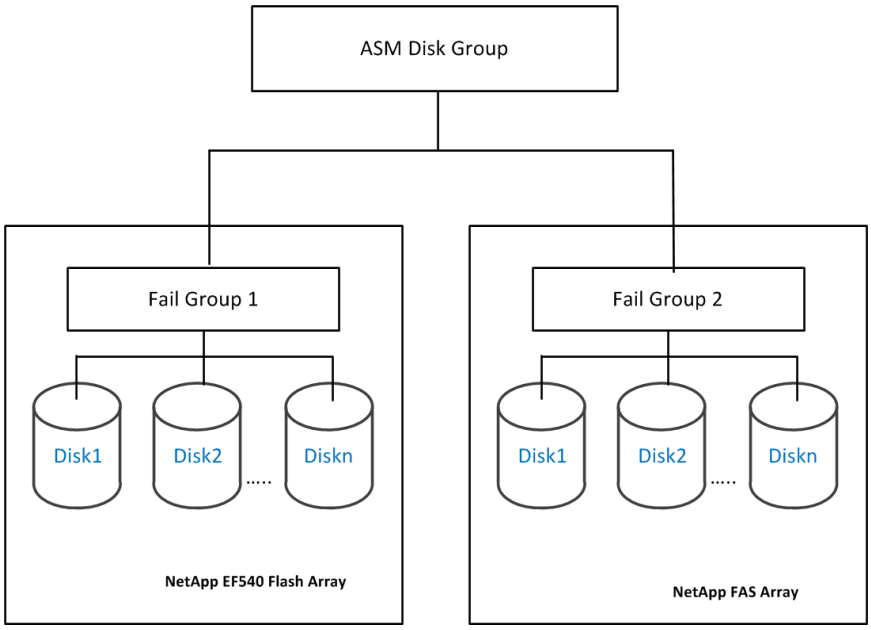

5 Workload Testing

A demanding workload with a random I/O characteristic to simulate an OLTP application was executed against the database. The workload was configured to measure the performance on the storage array by

scaling the IOPS with submillisecond latency, which would be a key requirement for any mission-critical application.

5.1 Read Workload

The solution achieved the expected IOPS under submillisecond latency for random workloads (8k blocks) and is shown in Figure 3. Also, as part of the workload test, a sequential load was generated by performing a reasonable number of direct path reads (for example, table scans) for the queries. The sequential workload showed an I/O performance of 2.6GB/sec and latency of 3ms, as expected.

Figure 3) Oracle random workload on NetApp EF540.

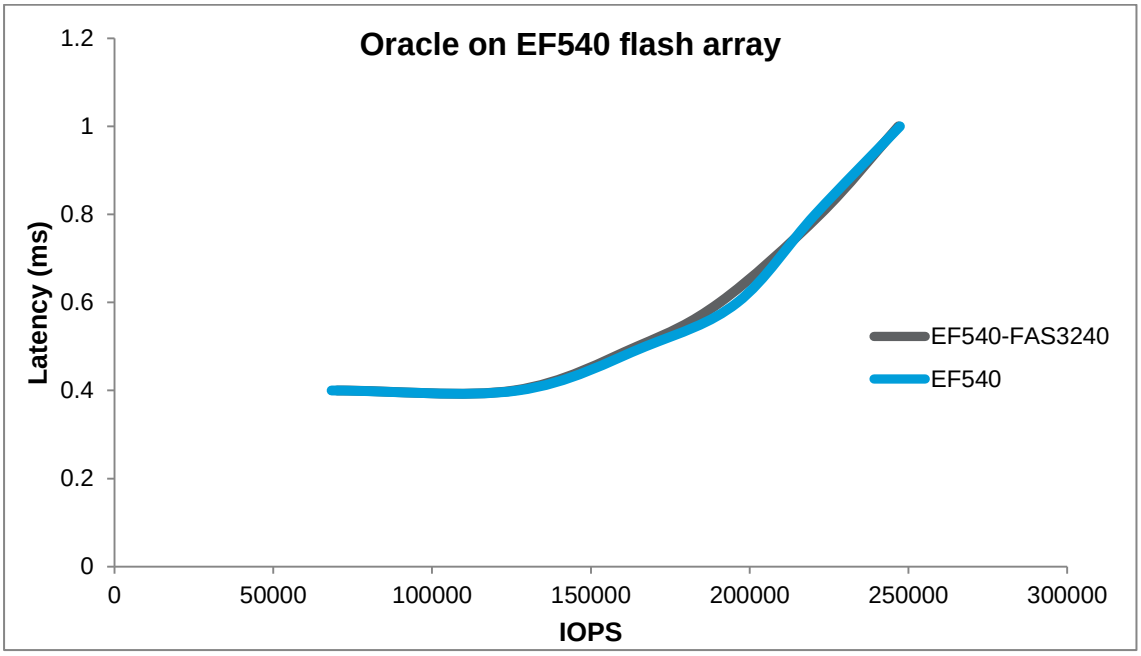

5.2 Workload—Raw I/O

Figure 4 shows the IOPS-to-latency graph during the raw I/O tests on the Oracle ASM fail group, which was configured on the EF540.

Figure 4) 100% random reads on the NetApp EF540.

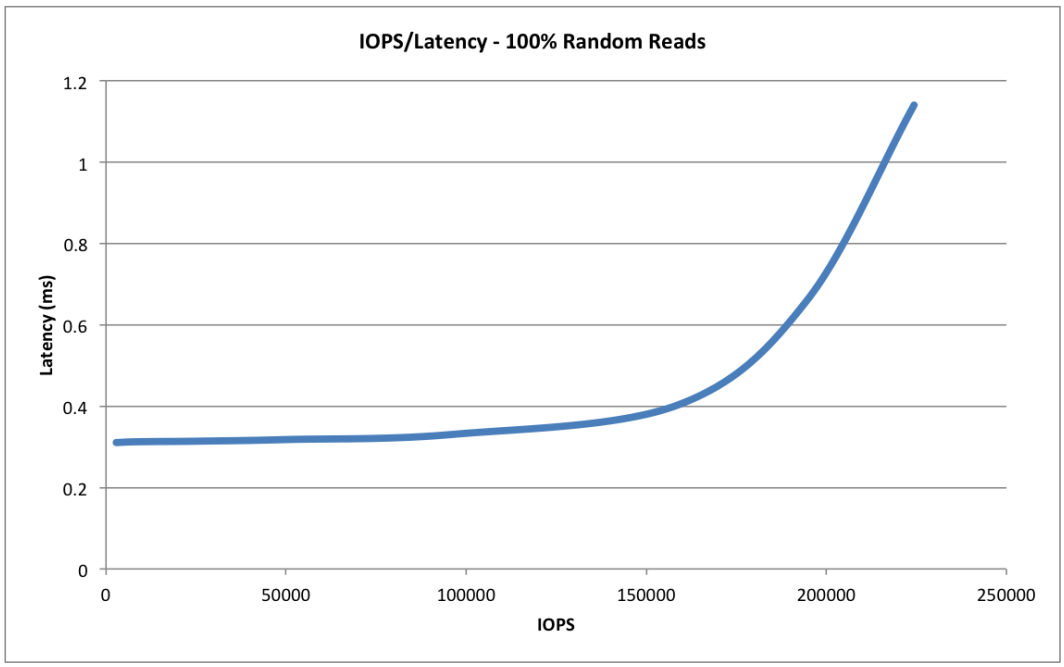

Figure 5 shows the effect on overall latency with the introduction of 30% writes, which is a very high level of write activity. Latency is slightly varied, but the platform still shows an overall IOPS of 100,000 before crossing the 1ms latency mark. The tests conclude that the EF540 can handle high data volume with very low latency, which is the key requirement for a subset of applications found in a customer's environment.

Figure 5) Mixed workload on the NetApp EF540.

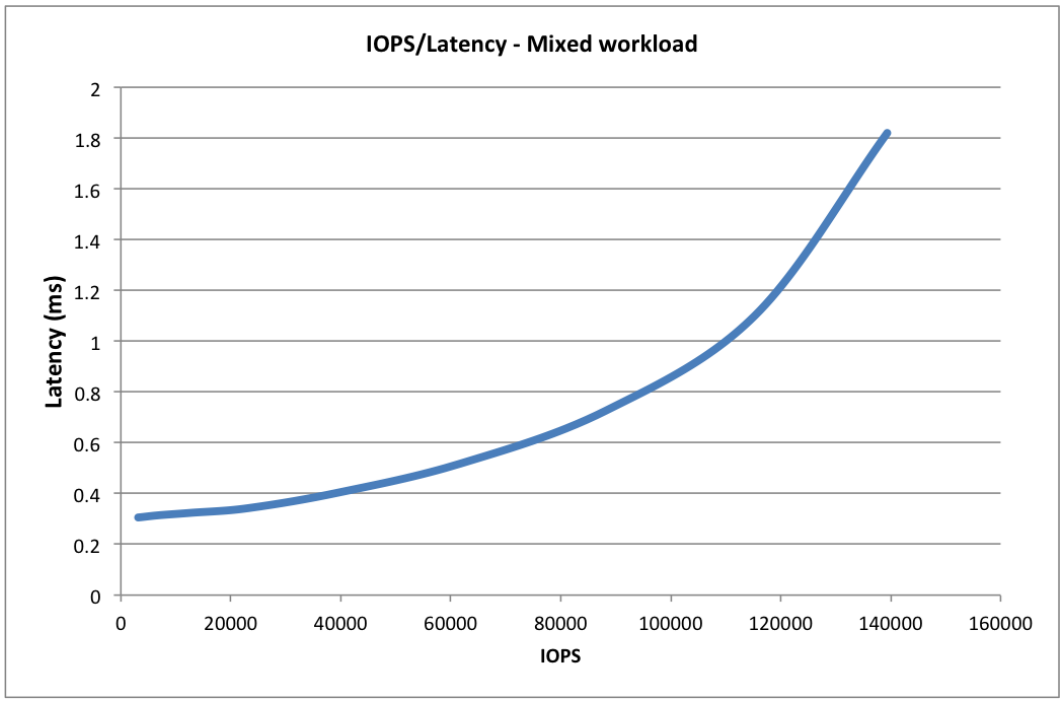

6 Conclusion

This combination NetApp storage solution helps meet the needs of customers with high IOPS and submillisecond latency requirements.

NetApp Snapshot features integrated into the solution enable customers to perform faster backup and recovery operations to meet the needs of the enterprise. Customers can leverage NetApp SnapVault for longer retention of data, thereby meeting regulatory compliance requirements. Finally, NetApp OnCommand Insight and Snap Creator Framework offer customers the ability to monitor and manage storage array performance and automate all aspects of data protection.

7 Acknowledgments

Special thanks to the following people for their valuable input:

- Antonio Jose Rodrigues Neto
- William Heffelfinger
- Jamal Boudi
- Nilesh Patel
- Robert Houser
- Ryan Leonard
- Vibhuti Bhusan

References

The following documents were referenced in this report:

- “NetApp EF540 Flash Array” http://www.netapp.com/us/media/ds-3417.pdf
- “SANtricity Storage Manager” http://www.netapp.com/us/media/ds-3171-66862.pdf
- “Snap Creator Framework” http://www.netapp.com/us/products/management-software/snapcreator-framework.aspx

Version History

Version	Date	Document Version History
Version 1.0	March 2013	Initial release

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com