



Technical Report

A Comparison of NetApp and EMC Isilon File Services Performance Using SPECsfs2008

Bhavik Desai, Chris Lemmons, NetApp
November 2011 | TR-3990

ABSTRACT

In November 2011, NetApp delivered the best SPECsfs2008 result ever published using NFS and NetApp® Data ONTAP® 8 running in Cluster-Mode using 24 x FAS6240 storage controller nodes. This technical report compares those test results to competitive SPECsfs2008 results posted by Isilon, a division of EMC Corporation, to demonstrate the superior performance of the NetApp solution.

TABLE OF CONTENTS

1 INTRODUCTION 3

2 COMPETITIVE SUMMARY 3

 2.1 PERFORMANCE AT LARGEST CLUSTER SIZE3

 2.2 PERFORMANCE SCALING AT DIFFERENT CLUSTER SIZES4

3 CONCLUSION..... 6

LIST OF TABLES

Table 1) Performance results for largest cluster size. 3

LIST OF FIGURES

Figure 1) NetApp performance results..... 5

Figure 2) EMC Isilon performance results. 5

1 INTRODUCTION

NetApp has consistently demonstrated leadership in high-performance NAS environments using our FAS series of unified storage systems. In September 2011, NetApp announced the availability of Data ONTAP 8. This newest version offers a storage operating system (OS) platform to help address the data management challenges faced by growing and dynamic businesses. Building on the NetApp leading storage efficiency features, Data ONTAP 8 enhances this industry-leading unified storage platform by providing:

- Enterprise nondisruptive operations that eliminate planned downtime, even during servicing and upgrades
- Massive scalability
- Improved management for large-scale deployments

Data ONTAP 8 extends the storage domain of a high-availability (HA) pair of controllers to multiple pairs of controllers. With Data ONTAP 8, capacity scales from terabytes to tens of petabytes, all of which is transparent to running applications. Storage is virtualized across as many as 12 HA pairs and is managed as a single logical pool of resources and name space.

Successful organizations of tomorrow must evolve their virtualized IT environments to a more flexible and efficient service-oriented model. NetApp Data ONTAP 8 helps companies achieve this goal by providing a massively scalable unified platform that delivers a shared IT infrastructure foundation for nondisruptive operations, operational flexibility and efficiency, and on-demand IT services.

2 COMPETITIVE SUMMARY

This section compares the NetApp SPECsfs2008 NFS test results to those recently published by EMC Isilon. We compare the performance and latencies of the different configurations for both the largest and smaller cluster sizes.

2.1 PERFORMANCE AT LARGEST CLUSTER SIZE

With the arrival of Data ONTAP 8 running in Cluster-Mode, NetApp has delivered the capability to provide scalable storage capacity and performance in one package. Demonstrating the stability and scalability of Data ONTAP 8 running in Cluster-Mode, NetApp delivered the best SPECsfs2008 result ever published with NFS using a cluster consisting of 24 x FAS6240 storage controller nodes.

Table 1 shows the leading NetApp result as well as the next-highest result generated by EMC Isilon in May 2011, on a 140-node cluster of EMC Isilon S200 systems. Compared to the EMC Isilon benchmark, the NetApp FAS6240 system delivers:

- 36% more performance
- 40% faster overall response time (ORT)

Table 1) Performance results for largest cluster size.

	SPECsfs2008_nfs.v3 Ops/sec	ORT (ms)	Total Disk Drives	Ops/sec per Node	Ops/sec per Disk
NetApp FAS6240 at 24 nodes	1,512,784	1.53	1,728	63,033	875
EMC Isilon S200 at 140 nodes	1,112,705	2.54	3,360	7,948	331
Percent better than EMC Isilon	36%	40%	49%	693%	164%

In addition to the throughput and response time data, we found the following items of interest when comparing the specifics of the NetApp and EMC Isilon configurations used to generate these test results.

- The NetApp configuration efficiently required only approximately 50% of the disk drives used in the EMC Isilon configuration: 1,728 for NetApp compared to 3,360 disk drives for EMC Isilon.
- The NetApp configuration used significantly fewer nodes: 24 FAS6240 storage controller nodes compared to 140 EMC Isilon S200 controller nodes.
- The NetApp configuration delivered over 60,000 SPECsfs2008_nfs.v3 ops per second per cluster node compared to less than 8,000 SPECsfs2008_nfs.v3 ops per second per node in the EMC Isilon configuration.
- The NetApp configuration delivered 875 SPECsfs2008_nfs.v3 ops per second per disk drive compared to 331 SPECsfs2008_nfs.v3 ops per second per disk drive in the EMC Isilon configuration.

2.2 PERFORMANCE SCALING AT DIFFERENT CLUSTER SIZES

Like EMC Isilon, NetApp recently published a series of additional SPECsfs2008 test results. These results demonstrate the performance of a cluster of FAS6240 controller nodes as the number of nodes in the cluster increases. For these tests, we ran the SPECsfs2008 NFS benchmark on clusters containing 4, 8, 12, 16, and 20 FAS6240 storage controller nodes to accompany the record 24-node result. The EMC Isilon configurations consisted of 7, 14, 28, and 56 nodes to accompany its 140-node S200 result.

Figure 1 and Figure 2 compare the published test results for both the NetApp and EMC Isilon configurations, respectively, for the different cluster sizes noted in Table 1. Items worth noting are:

- In both the NetApp and EMC Isilon configurations, we observed nearly linear scaling as the number of ops per second generated increased proportionally with the number of nodes in the cluster.
- As the size of the NetApp cluster increased, the ORT reported at each cluster size remained nearly constant at approximately 1.5ms as the number of ops per second generated by the cluster increased. Overall, the ORTs reported by the NetApp configuration were 40% to 50% faster than those reported for the EMC Isilon configuration at all cluster sizes.

Figure 1) NetApp performance results.

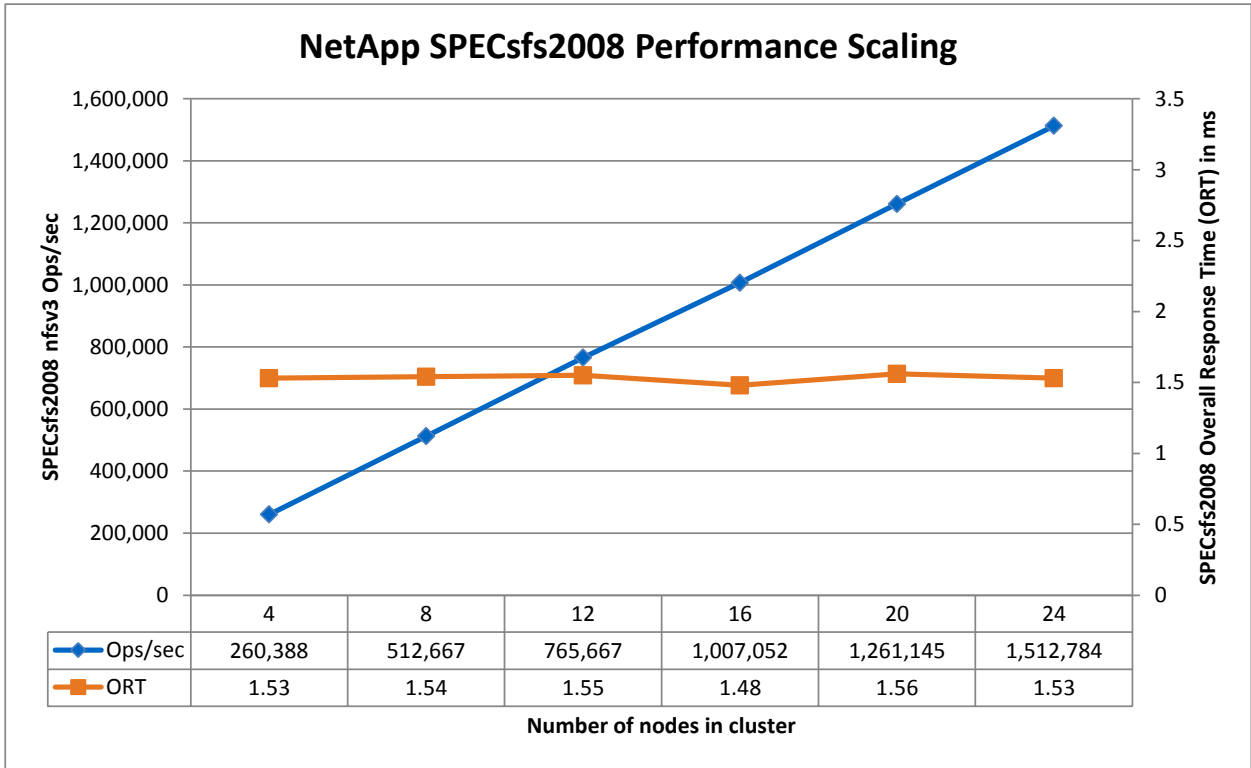
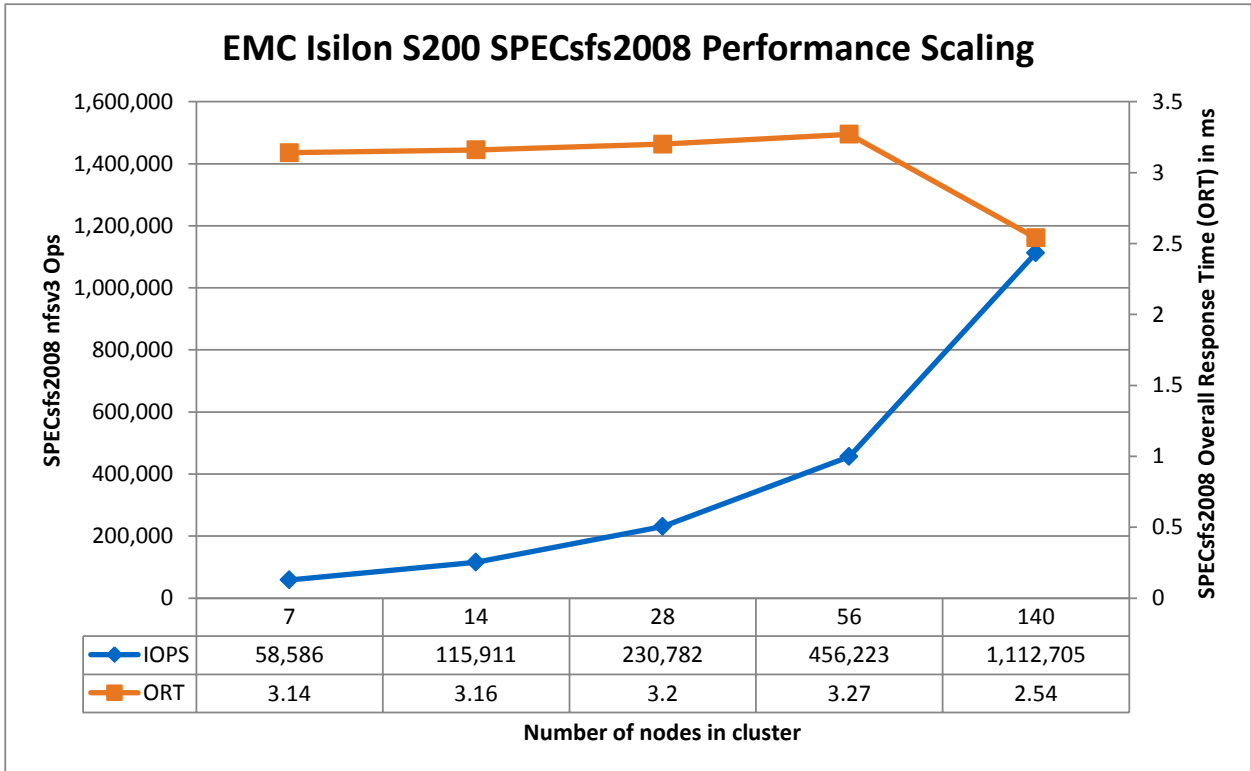


Figure 2) EMC Isilon performance results.



3 CONCLUSION

The SPECsfs2008 NFS test results clearly show that a cluster of FAS6240 systems running Data ONTAP 8 running in Cluster-Mode deliver record-setting NFS-based file services performance while also demonstrating leading capabilities to scale out storage environments. What this means for our customers is that NetApp can meet the challenges of high-performance file-serving environments using a single unified storage cluster while delivering superior responsiveness compared to EMC Isilon. In the real world, this means that applications respond faster and queries complete faster, resulting in better performance and lower overall response times. Finally, these results demonstrate that NetApp storage provides near-linear performance scaling as the size of the storage cluster increases while providing significant efficiency advantages by using only approximately 50% of the disk drives of EMC Isilon's 140-node configuration.

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

© 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, and Data ONTAP are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3990-1111