



Confidence in the connected world.

Veritas Storage Foundation Cluster File System 5.0 from Symantec

Performance and Scalability
Analysis: Extended Write
Workload

Veritas Storage Foundation Cluster File System 5.0 from Symantec Performance and Scalability Analysis

Contents

Introduction 4

About the Veritas Storage Foundation product line from Symantec..... 5

Extended write workload introduction..... 7

Extended write workload hardware and system specifications..... 8

Extended write workload tests and conclusion 9

Summary 10

Introduction

This white paper focuses on the performance and scalability of Veritas Cluster File System 5.0 in two workloads – a sequential write workload and an NFS file serving workload – in a configuration running from 1 to 16 nodes. The white paper will demonstrate that Veritas Storage Foundation Cluster File System 5.0 provides the most scalable infrastructure of any cluster file system in the Linux operating system. The workloads will be discussed separately and greater attention will be paid to the NFS file serving workload.

NFS is the most popular protocol on the market to share data between UNIX and Linux servers; it was first introduced in 1984 by Sun Microsystems and later was made an official standard for transferring files between computers over the network.

The general findings are that Cluster File System achieves a scalability factor of approximately 99% at the 16th node in the sequential write workload and a 91% factor in the NFS file serving workload.

This paper has been written using a specific hardware and storage configuration (to be discussed later in the paper). Although hardware is inevitably a dependency in any product analysis, lengths were taken to ensure that hardware specifics did not result in any bottlenecks. (Note: for these reasons, the paper focuses on scalability factors rather than the number of transacted operations as the latter can be more easily controlled based upon hardware specifications.)

For a list of supported storage arrays, please consult the Symantec Hardware Compatibility List, this applies to all products in this white paper.

Hardware Compatibility List for Storage Foundation 5.0

<http://seer.entsupport.symantec.com/docs/283161.htm>

About the Veritas Storage Foundation product line from Symantec

Veritas Storage Foundation

Storage Foundation provides easy-to-use online storage management, enables high availability of data, optimized I/O performance, and allows freedom of choice in storage hardware investments. Veritas Storage Foundation is the base storage management offering from Symantec. It includes Veritas File System and Veritas Volume Manager. Both Veritas File System and Volume Manager include advanced features such as journaling file system, storage checkpoints, Dynamic Multi-Pathing, off-host processing, volume snapshots and tiered storage. Storage Foundation comes in three editions, Basic, Standard and Enterprise, each targeting different environments:

Storage Foundation Basic is the freeware version of Storage Foundation. Available as a free download, it is limited to a maximum of 2 CPU and 4 volumes and 4 file systems. For more information, please visit: <http://www.symantec.com/business/theme.jsp?themeid=sfbasic>

Storage Foundation Standard is intended for SAN connected servers with high performance requirements and availability features, such as multiple paths to storage. This product is a minimum requirement for High Availability solutions.

Storage Foundation Enterprise includes the entire feature set of both File System and Volume Manager. It is designed for servers with large SAN connectivity, where high performance, off-host processing and storage tiering are desired.

http://www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=203_1

Veritas Storage Foundation Cluster File System

Veritas Storage Foundation Cluster File System provides an integrated solution for shared file environments. The solution includes Veritas Cluster File System, Cluster Volume Manager and Cluster Server to help implement robust, manageable, and scalable shared file solutions. With Veritas Storage Foundation Cluster File System, cluster-wide volume and file system configuration allows for simplified management; and extending clusters is simplified as new servers adopt cluster-wide configurations.

http://www.symantec.com/business/products/overview.jsp?pcid=2247&pvid=209_1

Veritas Cluster Server

Veritas Cluster Server is the industry's leading cross-platform clustering solution for minimizing application downtime. Through central management tools, automated failover, features to test disaster recovery plans without disruption, and advanced failover management based on server capacity, Cluster Server allows IT managers to maximize resources by moving beyond reactive recovery to proactive management of application availability in heterogeneous environments.

http://www.symantec.com/business/products/overview.jsp?pcid=2247&pvid=20_1

Veritas Storage Foundation for Databases

Veritas Storage Foundation for Databases is an integrated suite of industry-leading Symantec products that delivers easier manageability, superior performance and continuous access to DB2, Oracle and Sybase databases. This suite is built on Veritas Storage Foundation, a storage infrastructure layer that enables database storage optimization with online storage virtualization and RAID. Storage Foundation for Databases also delivers the manageability of file systems with the performance of raw devices for a database environment.

http://www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=208_1

Veritas Storage Foundation for Oracle RAC

Veritas Storage Foundation for Oracle RAC includes storage-management and high availability technologies allowing you to implement robust, manageable, and scalable Oracle Real Application Clusters. The solution delivers the industry's first heterogeneous cluster file system supporting Oracle RAC.

http://www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=145_1

Extended write workload introduction

The goal of this project was to create an I/O test that produced scalability factors that provide relative comparisons with other cluster file system solutions.

In order to arrive at relevant and comparable data points, the following assumptions were made:

- The basic test would use an I/O driver creating sequential writes that were extending a file. Thus, the writes would create metadata activity. (The alternative that was rejected was to pre-allocate the files and simply overwrite each file.)
- Tests with multiple (N) nodes will use multiple (N) processes. Each process will write to its own file, but all files will be in a single file system.
- Tests would be done over NFS and local.

Extended write workload hardware and system specifications

Below we describe the configuration used in the 16-node testing. For tests with fewer than 16 server nodes, the number of data arrays matched the number of server nodes.

Hardware

- Servers:
 - 16x Sun v20z . Each has
 - Two 2.2 GHz AMD Opteron x86_64 CPUs
 - 8 GB RAM
 - HBA: dual-port 2-Gbit FC
 - Four 1-Gbit Ethernet
- Clients: 16x SuperMicro. Each has
 - One 3.0 GHz Intel Pentium4 CPU
 - 1 GB RAM
 - One 1-Gbit Ethernet
- Storage:
 - 17x Sun StorEdge 3510 arrays. Each has:
 - One controller with two 2-Gbit FC ports
 - Each controller has 1GB cache, write back, optimized for random I/O
 - Two expansion units
 - 36x 15k rpm 36GB disks total
- Ethernet and FC Switches:
 - Two Ethernet switches (Dell 2724 PowerConnect) for GAB/LLT traffic
 - Ethernet switch (Cisco Catalyst 4948) for the SPECsfs private network
 - Ethernet switch (Dell PowerConnect 5324) for the public network (this switch is not used in the benchmark tests).
 - Two Brocade FC switches (Silkworm 4100) each having 32 ports; these are used to connect the 16 server nodes to the 17 arrays

Software and Tuning

- RHEL AS4 Update 4 (kernel 2.6.9-42.ELsmp)
- SFCFS 5.0MP1

Other Configuration Information

- The hardware is set up for the full 16node test. When doing a test with $N < 16$ nodes, only those N nodes were powered up. The remaining nodes were shutdown.
- The cluster only consisted of the N nodes.
- Prior to each run:
 - All server nodes and clients were rebooted; all tuning was reapplied
 - The volumes and volume sets were created for a run with N nodes.
 - The file systems were created and mounted.

Extended write workload tests and conclusion

Benchmark Tests

Two sets of tests were performed; the difference between the tests was where the I/O driver program ran in each test.

Test 1: Server nodes run the I/O driver script (no NFS).

Test 2: Client nodes run the I/O driver script accessing the shared clustered file system (using NFS).

Table 1-1 Results of test 1

Results with Server Nodes Running the IO driver (no NFS)

Number of Nodes	Average MB/sec per Node	Total MB/sec to Cluster File System	Ratio of total N-node CFS MB/sec to 1-node MB/sec	Percentage of average N-node CFS MB/sec to 1-node MB/sec
1	146.9	146.9	1.00	100%
2	146.2	292.5	1.99	100%
4	145.8	583.3	3.97	99%
8	145.9	1166.8	7.94	99%
12	145.6	1747.1	11.89	99%
16	145.4	2326.0	15.83	99%

As the table shows, the scalability was linear from 1 to 16 nodes, and the total throughput with N-nodes was very close to N times the throughput with 1-node. For example, with 16-nodes, the throughput of the cluster was 15.83 times that of the 1-node test.

Table 1-2 Results of test 2

Results with Client Nodes Running the IO driver accessing the NFS-mounted Cluster File System

Number of Nodes	Average MB/sec per Node	Total MB/sec to Cluster File System	Ratio of total N-node CFS MB/sec to 1-node MB/sec	Percentage of average N-node CFS MB/sec to 1-node MB/sec
1	102.9	102.9	1.00	100%
2	102.9	205.8	2.00	100%
4	102.5	409.9	3.98	100%
8	101.4	811.3	7.88	99%
12	101.7	1220.4	11.86	99%
16	102.3	1636.3	15.90	99%

As the table shows, the scalability was linear from 1 to 16 nodes, and the total throughput with N-nodes was very close to N times the throughput with 1-node. For example, with 16-nodes, the throughput of the cluster was 15.90 times that of the 1-node test.

Summary

Veritas Cluster File System from Symantec provides near linear performance scalability in a single file system schema that can span up to 32 servers in a cluster. With Storage Foundation Cluster File System, cluster-wide volume and file system configuration allows for simplified management—and extending clusters is simplified as new servers adopt cluster-wide configurations. Because all files can be accessed by all servers, applications can be allocated to servers to balance load or meet other operational requirements.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

12/07 13579374