
Open-E DSS V7

Active-Active vs. Active-Passive

Performance comparison of failover solutions

Contents

INTRODUCTION	2
HIGH AVAILABILITY OVERVIEW	2
ACTIVE-PASSIVE	3
ACTIVE-ACTIVE	3
TESTING METHODOLOGY AND FINDINGS	3
ACTIVE-ACTIVE VS. ACTIVE-PASSIVE - DATA THROUGHPUT PERFORMANCE COMPARISON	4
ACTIVE-ACTIVE VS. ACTIVE-PASSIVE - FAILOVER SWITCHING TIME COMPARISON	5
SUMMARY	6
APPENDIX A – TEST ENVIRONMENT DETAILS	7
APPENDIX B – MEASURED VALUES	8

INTRODUCTION

The purpose of this document is to compare two high availability failover configurations, namely Active-Active and Active-Passive, both being possible to configure using Open-E DSS V7 software*.

The comparison was entrusted with two tests conducted by the Open-E Quality Assurance Team, measuring respectively: data throughput performance and failover switching time.

Apart from presenting tests results and conclusions, this document provides brief descriptions of both configurations.

****To take advantage of Active-Active functionality, an additional Feature Pack - Active-Active iSCSI Failover of the Open-E DSS V7 software is required.***

HIGH AVAILABILITY OVERVIEW

Last year (2011), Science magazine published an interesting study presenting measurements of world's total technological capacity to communicate, store, and compute information.

According to the publication: "in 2007, humankind was able to store 2.9×10^{20} optimally compressed bytes and carry out 6.4×10^{18} instructions per second on general-purpose computers"[1]. At the same time: "general-purpose computing capacity grew at an annual rate of 58%" [1].

With amount of data increment at this rate, the need to effectively store data becomes a priority.

Thinking of "effective" ways of storing data – the key element is continuous availability of storage solutions.

Availability of resources is one of the concerns of Business Continuity. Planning involved in that field is about ensuring that companies will be able to operate in case of a serious disruption.

IT components availability has become increasingly important since businesses are more and more dependent on digital information and technology. Considering industries as candidates for implementing high availability solutions, we can name healthcare, banking, online retail and many more. All of them require continuous, uninterrupted access to their resources.

A standard for availability of the products or systems is known as "class of 9s".

It refers to the percentage figures used for the representation of high availability (e.g. 99.999% is five nines). For example, in case of availability value of 98 %, the downtime will equal 2 %. Downtime easily translates to **lost revenue**.

High availability is a goal for storage solutions. The priority is for data or services to be accessed at all times, even if malfunction occurs. This can be achieved in a number of ways involving clustering, redundancy and software that handles the failover configurations.

High Availability clusters, with priority to decrease downtime to a minimum, are one of the implementations of these solutions. Such clusters are usually supported by storage management software and **Open-E DSS V7** can be an example of those.

Typically, a high availability cluster consists of two nodes. In case of a node failure, the client network connections and file systems are moved or mounted to the other node, so the downtime is minimized or eliminated completely. An operation of switching to a redundant or synchronized node is called **failover**. In contrast, **failback** is the process of restoring a system, component or service to a state before failure.

There are two major configurations for implementation of high availability storage solutions – **Active-Passive** and **Active-Active**. Let us briefly describe them, before presenting the results.

Downtime Costs \$26.5 Billion in Lost Revenue

"...more than \$26.5 billion in revenue is lost each year from IT downtime, which translates to roughly \$150,000 is lost annually for each business (...). Of the 200 companies surveyed, small enterprises lost, on average, more than \$55,000 in revenue due to IT failures each year, while midsize companies lost more than \$91,000 and large companies lost more than \$1,000,000."

Source: Information Week, 2011 [2].

ACTIVE-PASSIVE

An **Active-Passive** storage concerns the configuration of two nodes, where one of them plays the role of the primary node (active) and the other is set as standby (passive)—waiting to take over the control if necessary.

The active node processes all the I/O operations and continuously synchronizes its configuration and session information with the passive device, so that in case of a failure, it will be ready to take its place.

The main disadvantage of the Active-Passive configuration is that the secondary node is not operational; therefore the **hardware resources are not fully utilized**.

ACTIVE-ACTIVE

In an **Active-Active** storage configuration, both nodes process I/O's providing balanced access to the logical devices.

In contrast, with Active-Passive configuration, neither instance is designated as primary or secondary. Both nodes are kept synchronized, so any changes to data in one node will be propagated to the other node. The synchronization is made through synchronous replication.

Active-Active configuration provides disaster tolerance, so if one node fails, the other one takes over automatically and all services continue to run without interruption.

The main difference in comparison with Active-Passive configuration is that both nodes are in the "operational mode", balancing read, write and replication traffic, which **significantly improves overall cluster performance**.

With Open-E DSS V7

- » *an Active-Active or Active-Passive High Availability cluster can be easily setup with reduced number of configuration steps.*
- » *the failover configuration can be made on any cluster node, without a need to setup every node separately.*
- » *failover and failback functions are simplified and re-designed.*
- » *reboot of cluster nodes is much more user friendly, allowing node to join the cluster automatically, if no error has been found.*

TESTING METHODOLOGY AND FINDINGS

This paragraph presents tests that were conducted to compare the Active-Active and the Active-Passive failover configurations of Open-E DSS V7 software.

The following measures were tested:

1. Data throughput in both configurations.
2. Failover switching time in both configurations.

The purpose of the tests was to find out **which configuration performs better** in terms of managing data and handling failover processes.

In all cases the test environments were configured using the Open-E DSS V7 software.

- > Each test section consists of a brief background description and findings.
- > A detailed list of hardware components and topology used in tests can be found in Appendix A.
- > Details of measures from both tests can be found in Appendix B.

ACTIVE-ACTIVE VS. ACTIVE-PASSIVE - DATA THROUGHPUT PERFORMANCE COMPARISON

Test Background

- > To enhance the throughput performance a multipath IO between clients and storage servers (nodes) has been implemented in both configurations.
- > In order to achieve Zero-Single-Point-of-Failure with 2 network switches, bonding (balance-alb) of the network interfaces was applied on each node.
- > The performance tests with variable block sizes were performed using Iometer tool, version 2006.07.27.
- > Block sizes of 4 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, 1 MB and 4 MB were used for the tests.
- > Read and write time onto raw iSCSI device for each block size was measured over a period of 90 seconds.
- > The results in MB/s are sums of measures received for both clients (WIN1 and WIN2) taking part in the test.

Test Findings

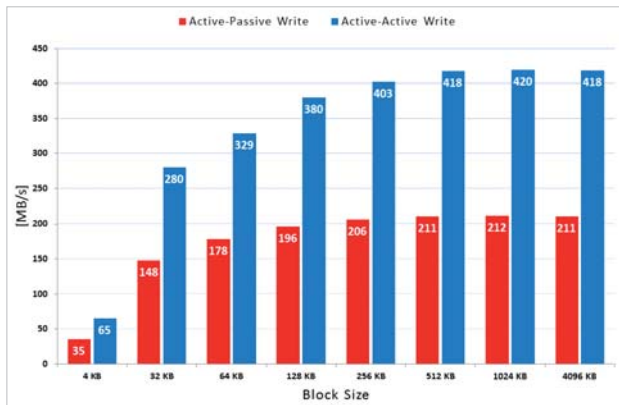


Figure 1. Write performance [MB/s] per block size (KB).

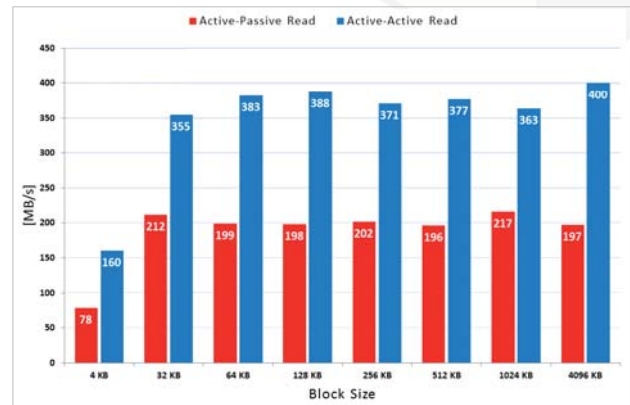


Figure 2. Read performance [MB/s] per block size (KB).

As indicated in figures 1 and 2, the comparison of read and write performance in both failover configurations proved an advantage of the Active-Active set up.

Open-E DSS V7 in the Active-Active iSCSI Failover setup achieved up to 98 % better results in write performance than Active-Passive.

Similar results were calculated for read performance, the advantage in favor of Active-Active configuration reached the maximum of 105 %.

The average values for data write performance ranged from 175 MB/s, in case of the Active-Passive configuration, to 339 MB/s in case of the Active-Active configuration.

For data read performance average values ranged from 187 MB/s in the Active-Passive configuration, to 350 MB/s in the Active-Active configuration.

The above results clearly correspond to the fact that Open-E DSS V7 in the Active-Active failover configuration **better utilizes storage resources and allows for provision of data at higher speed** than the Active-Passive set up.

Detailed results are presented in Appendix B.

ACTIVE-ACTIVE VS. ACTIVE-PASSIVE- FAILOVER SWITCHING TIME COMPARISON

Test Background

The purpose of this test was to calculate failover switching time in both configurations, defined as time needed for data to be accessible again on the second node in case of failover.

To entrust the results, different numbers of logical volumes per node were used for evaluation - 10, 20, 30, 40 and 50 respectively.

Test Findings

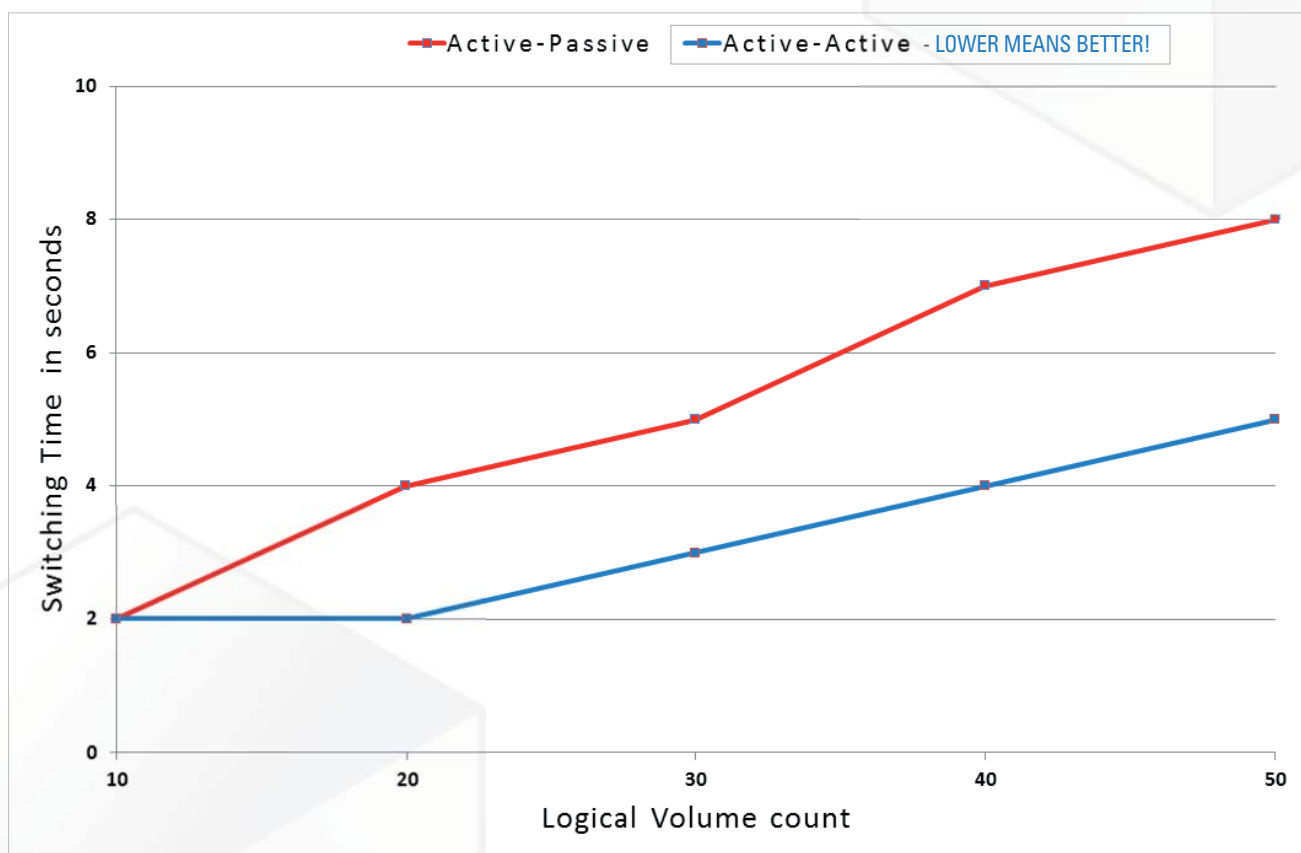


Figure 3. Failover switching time in seconds per LV count.

In the comparison of the Active-Active and the Active-Passive failover switching time, the Active-Active set up delivered better results proportionally to increasing the Logical Volumes count in a node.

While for 10 Logical Volumes the switching times were on the same level for both configurations, with 20-30 Logical Volumes the Active-Active failover solution performed 2 seconds faster than Active-Passive failover set up and with 40-50 this difference has increased to 3 seconds.

Detailed results are presented in Appendix B.

SUMMARY

Tests Key Findings


- > The Open-E DSS V7 iSCSI Active-Active Failover configuration achieved better results in data throughput performance. The difference between tested configurations was up to 105% in Active-Active failover favor. This translates to **improved access to storage devices** for both writing and reading data, allowing for **more operations to be performed** at specific period of time in case of an Active-Active configuration.
- > The Open-E DSS V7 iSCSI Active-Active failover configuration performs better in terms of resources switching time. The time needed to switch operation of the failing nodes was shorter than in case of Active-Passive configuration. This effectively results in **minimized disruptions in case of a failure**.

To summarize, Open-E DSS V7 with the Feature Pack - Active-Active iSCSI Failover helps to deliver much better value than the basic Active-Passive configuration.

Minimized disruptions in case of a failure, better utilized storage resources and strong performance in both storing and accessing data, confirmed in our tests, strongly speak for considering Active-Active failover configuration as an ideal solution for High Availability storages.

What Open-E DSS V7 with Active-Active iSCSI Failover has to offer?

- » It is simple to use – while Active-Active set up is considered as complex to configure - with Open-E DSS V7 it is no longer an issue.
- » Eliminates wasting of hardware, all resources are in use - working towards better system performance.
- » Provides self-validation of the system. When starting a cluster, Open-E DSS V7 checks all the critical settings on each node. This way, clusters will not be started if they were configured wrong.
- » Increases sensibility for network failures, thanks to the possibility of configuring Ping Nodes.
- » Speeds up networking connectivity, since I/O traffic is equally balanced on two nodes.
- » Fully utilizes all processing power on both cluster nodes.



APPENDIX A – TEST ENVIRONMENT DETAILS

1. Active-Active vs. Active-Passive - Data Throughput Performance Comparison

Hardware components details

Nodes specification:

CPU:	Dual-core 3.10GHz processor
MEM:	8GB DDR3 RAM
NICs:	5x 1GbE Ethernet network adapter 1x 10GbE Ethernet network adapter
H/W RAID:	1x PCI-e 6GB SAS with 1GB Cache
HDD:	16x 3TB SATAIII 7200RPM 64MB Cache

Hardware topology

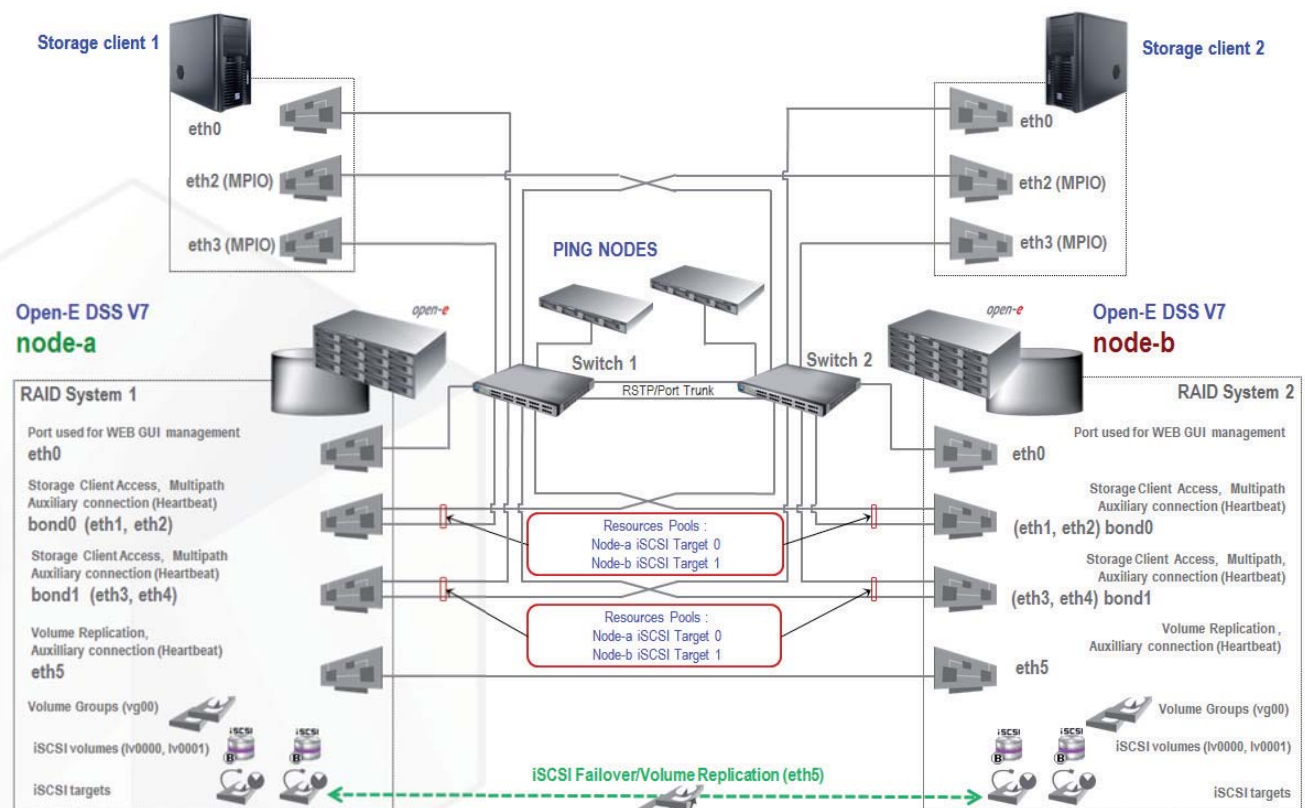


Figure 4. Hardware topology in test

2. Active-Passive vs. Active-Active - Failover Switching Time Comparison

Hardware components details

Nodes specification:

CPU:	Dual-core 3.10GHz processor
MEM:	8GB DDR3 RAM
1NIC:	2x 1GbE Ethernet network adapter
2NIC:	2x 10GbE Ethernet network adapter
H/W RAID:	1x PCI-e 6GB SAS with 1GB Cache
HDD:	16x 3TB SATAIII 7200RPM 64MB Cache

APPENDIX B – MEASURED VALUES

1. Active-Active vs. Active-Passive - Data Throughput Performance Comparison

Block size	Active-Active Write (MB/s)	Active-Active Read (MB/s)
4 K	65,21	160,24
32 K	280,15	354,55
64 K	328,99	382,74
128 K	379,97	387,73
256 K	402,59	371,47
512 K	418,13	377,49
1024 K	419,71	363,46
4096 K	418,36	400,39

Table 1. Data throughput in the Active-Active configuration.

Block size	Active-Passive Write (MB/s)	Active-Passive Read (MB/s)
4 K	35,05	78,22
32 K	148,15	211,54
64 K	178,36	198,82
128 K	196,05	197,84
256 K	206,32	201,58
512 K	211,00	196,20
1024 K	212,02	216,57
4096 K	210,98	197,10

Table 2. Data throughput in the Active-Passive configuration.

2. Active-Passive vs. Active-Active - Failover Switching Time Comparison

Open-E DSS V7 1GbE Active-Active	
Logical Volume Count	Switching Resources Time (s)
10	2
20	2
30	3
40	4
50	5

Table 3. Failover switching time in the Active-Active configuration.

Open-E DSS V7 1GbE Active-Passive	
Logical Volume Count	Switching Resources Time (s)
10	2
20	4
30	5
40	7
50	8

Table 4. Failover switching time in the Active-Passive configuration.



References:

[1] Hilbert, M., Lopez, P. (2011). The World's Technological Capacity to Store, Communicate, and Compute Information. Science 1 April 2011: Vol. 332 no. 6025 pp. 60-65. Retrieved from: <http://www.sciencemag.org/content/332/6025/60.full>

[2] Harris C., (2011). IT Downtime Costs \$26.5 Billion In Lost Revenue, Information Week, Retrieved from <http://www.informationweek.com/storage/disaster-recovery/it-downtime-costs-265-billion-in-lost-re/229625441>