
KALISTA IO

Get ready for a data storage revolution

Phalanx

Hadoop TestDFSIO

benchmarks

Test system

Test system configuration

Processor Intel(R) Core(TM) i7-4771 CPU @ 3.50GHz

Memory 16GiB DDR3 Synchronous 2400 MHz

Storage interface SATA 3.2, 6.0 Gb/s

OS device 128GB Samsung 840 PRO Series SSD (MZ-7PD128)

Metadata device 480GB Samsung SM843T (MZ-7WD4800/003)


Storage devices

Device info




Model number	HSH721414ALN6M0 (Hs14)	WUH721414ALE6L4 (He14)
Firmware revision	L4GMT240	LDGNW07G
Drive type	Host managed SMR	CMR

Disk throughput at different LBA offsets Hs14 vs He14



4KB random read IOPS Hs14 vs He14 (0 to 128GB from OD)





Apache Hadoop is a collection of open-source software utilities that facilitate using a network of many computers to solve problems involving massive amounts of data and computation. It provides a software framework for distributed storage and processing of big data using the MapReduce programming model.

https://en.wikipedia.org/wiki/Apache_Hadoop

Methodology and SW versions

Methodology

Each test executed 3 times to capture average and standard deviation values

XFS and ext4 initialized and benchmarked with He14

Phalanx initialized and benchmarked with Hs14 (data) and Samsung SM843T (metadata)

Single device configuration with disk read & write cache enabled

Write test: 32x 16GB files written to HDFS (512GB dataset)

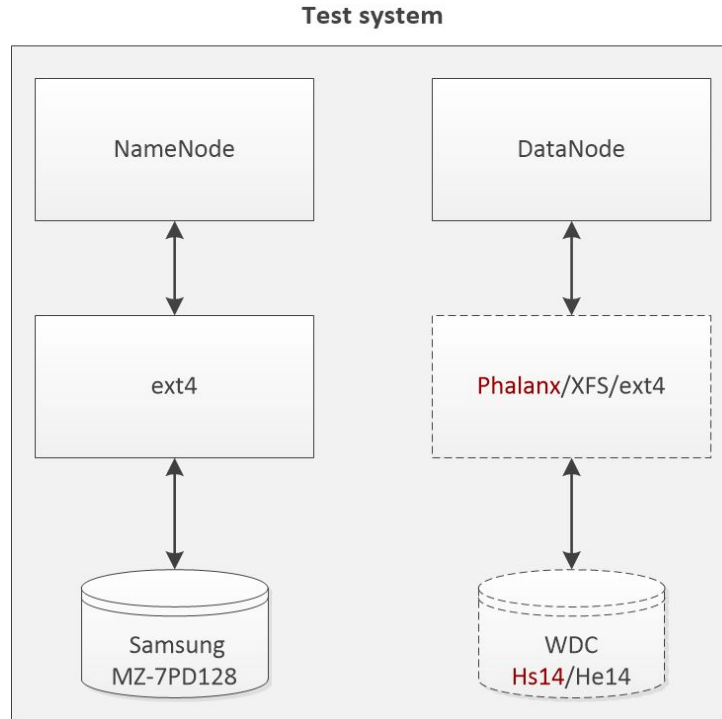
Read test: 32x 16GB files read from HDFS (512GB dataset - from write test)

Software

	Version
Apache Hadoop	3.2.0
TestDFSIO	1.8
OpenJDK	1.8.0_222
Operating system	Ubuntu 18.04.2 LTS (5.0.0-25-generic)

Topology and configuration

Hadoop cluster topology (single node pseudo-distributed)



Hadoop configuration

	Value	Description
io.file.buffer.size	131072	The size of buffer for use in sequence files. Determines how much data is buffered during read and write operations.
dfs.replication	1	Number of block replicas.
dfs.client-write-packet-size	1048576	Packet size for clients to write.
dfs.blocksize	268435456	The default block size for new files, in bytes.

Hadoop configuration (continued)

	Value	Description
dfs.namenode.name.dir	ext4 mount with Samsung MZ-7PD128	Determines where on the local filesystem the DFS NameNode should store the name table(fsimage).
dfs.namenode.checkpoint.dir	ext4 mount with Samsung MZ-7PD128	Determines where on the local filesystem the DFS secondary name node should store the temporary images to merge.
dfs.datanode.data.dir	Phalanx/XFS/ext4 mount with WDC Hs14/He14 depending on test configuration	Determines where on the local filesystem an DFS data node should store its blocks.



Configuration references

Links


<https://hadoop.apache.org/docs/r3.2.0/index.html>

<https://hadoop.apache.org/docs/r3.2.0/hadoop-project-dist/hadoop-hdfs/hdfs-default.xml>

<https://hadoop.apache.org/docs/r3.2.0/hadoop-project-dist/hadoop-common/core-default.xml>


Reads

TestDFSIO read results (32x 16GB files - 512GB dataset)



Writes

TestDFSIO write results (32x 16GB files - 512GB dataset)



Contact

<http://www.kalista.io>

@kalista.io

info@kalista.io