

HardwareBenchmarks

Cluster benchmark

How does different hardware perform with hadoop? Hopefully this page can help us answer that question and help new users buy their own based on our experience.

Please add your own configurations and sort benchmarks below. Information on how to run sort benchmark at: <http://wiki.apache.org/hadoop/Sort>

It basically generates 10gb of random data per node and sorts it.

Hardware

Cluster name	CPU model	CPU freq	Cores	RAM	Disk size	Disk interface	Disk rpm	Disks	Network type	Number of machines	Number of racks
Herd1	Intel Xeon LV	2.0ghz	4	4gb	0.25tb	SATA	7200rpm	4	GigE	35	2
Herd2	Intel Xeon 5320	1.86ghz	8	8gb	0.75tb	SATA2	7200rpm	4	GigE	20	2

Benchmark

All benchmarks run with the default randomwriter and sort parameters.

Cluster name	Version	Sort time s	Mappers	Reducers	Max map tasks / node	Max reduce tasks / node	Map speculative ex	Reduce speculative ex	Parallel copies	Sort mb	Sort factor
Herd1	0.14.3	3977 s	5600	175	?	?	Yes	Yes	20	200	10
Herd2	0.18.3	1715s	1520	136	7	8	No	Yes	20	100	50