

36% Faster Hadoop – Without Adding Servers

Add One 9100PRO to HDD-Based Nodes, See 36% Average Faster Run

9100PRO¹ Accelerates Existing Hadoop



Add a 9100PRO to Each Node

- 36% Faster Average Benchmark Completion Time
- More Cost Effective (than adding cluster nodes for similar performance)
- Leverage Existing Resources

Overview

When IT needs better Hadoop performance from their HDD-based nodes, they typically have two options:

They can add more nodes to the existing cluster. This may help meet performance goals, but the incremental cost of the added nodes may be prohibitive.

Alternatively, they can replace the current cluster nodes with new ones. Rebuilding with higher performance nodes may meet performance goals, but the cost is higher still.

Now there is a third choice that enables better performance and economics: add Micron's 9100PRO NVMe SSD to existing HDD-based cluster nodes.

This technical marketing brief shows how adding a single 9100PRO to each existing Hadoop node (10-node cluster) and making a slight change in YARN's resource localization provides 36% average reduction in benchmark runtime (over 10 test runs) and is far more economical than adding more nodes to achieve a similar improvement.

¹2.5" and Card form factors available, 2.5" 3.2TB version tested

Reduce Benchmark Runtime. Keep Your Cluster Investment.

Each Hadoop distribution comes with a set of standardized, built-in benchmarks. These benchmarks enable broad range performance measurements across technologies and deployments.

We used these benchmarks to compare the relative completion times of two 10-node cluster configurations. We ran each benchmark set ten times then used the average across those ten runs for our comparison.

The first configuration used standard HDDs for both the YARN cache and HDFS. The second used the same hardware – adding one 9100PRO and configuring YARN caching to use the 9100PRO instead of the HDDs (we still used the HDDs for HDFS in the second configuration).

The results are in Figure 1.

When we used the 9100PRO as the YARN cache, we saw a 36% reduction in benchmark completion time.



Figure 1: Benchmark Completion Time (Lower is Better)

Add the 9100PRO for Better Economics (Than Expanding the Cluster)

After we saw the above improvement, we also wanted to understand if adding one 9100PRO to the existing cluster nodes was more economical than expanding the cluster (adding more all-HDD nodes) to achieve a similar runtime reduction.

Our goal was to see how far we had to expand the all-HDD cluster to approximate the benchmark completion time for our 9100PRO YARN cache cluster (3518 seconds average across 10 runs), then analyze the cost.

We added 2 more nodes (12 total) and repeated the tests. The 10-run average was faster (4,353 seconds mean), but fell short. We added one more node (13 total) and again ran the same tests. This run was very close.

Table 1 summarizes the results:

Cluster Configuration	Mean Completion Time (10-runs) in Seconds
10 node, all HDD	5,554
12 node, all HDD	4,353
13 node, all HDD	3,398

Table 1: All-HDD Benchmark Completion Time (Lower is Better)

Table 1 also shows that it takes 13 all-HDD nodes to provide similar performance to one 10-node 9100PRO + HDD cluster. Table 2 shows the additional advantage of cost savings of adding one 9100PRO to each node.

Cluster Upgrade Option	Cost
Add one 9100PRO to each node	\$25,000
Expand cluster by three nodes	\$37,000

Table 2: Upgrade Costs for Similar Performance

When we ran these tests, the manufacturer's suggested retail price (MSRP) for the all-HDD cluster nodes (as configured in these tests) is just over \$12,000 each. The price for one 9100PRO (2.5" form factor, 3.2TB capacity) is just under \$2,500.²

Adding the 9100PRO to our existing cluster is more economical than adding cluster nodes to reach similar performance.

Summary

When IT needs more from their Enterprise deployments, they typically weigh several options. They may consider acquisition and recycling costs, performance benefit, and deployment time among a host of others.

Distributed systems like Hadoop typically offer two options for improving performance: Add more of what you already have (cluster expansion) or replace what you have (decommission and build new).

With Micron's 9100PRO NVMe SSD there is another, more attractive option: add one 9100PRO to each existing cluster node and make a small change to YARN resource allocation (to use the 9100PRO as the YARN cache).

The results are compelling.

Adding a single 9100PRO to each node (10-node, all HDD) in a cluster reduced standard benchmark runtime by 36% and costs less than expanding the existing cluster (to reach a similar benchmark runtime reduction).

² Selling price for the 9100PRO 3.2TB current as of the date of this document's first publication (average of prices from www.froogle.com)

How We Measured These Results

We used the standard benchmarks and built it to most Hadoop distributions, running the set 10 times and recording the mean completion time. Table 3 shows the benchmarks we used and their parameters.

Benchmark	Parameters
Wordcount	mapreduce.randomtextwriter.totalbytes=2147483648000 mapreduce.job.maps=100 mapreduce.job.reduces=100 mapreduce.output.fileoutputformat.compress=false
Word Aggregate Histogram	Default (no parameters set)
Sort	-r 1000
Test DFSIO Write	-write -nrFiles 1000 -size 1024MB
Test DFSIO-Read	-read -nrFiles 1000 -size 1024MB
Teragen	mapreduce.job.maps=100 mapreduce.job.reduces=100 21474836480 (bytes)
Terasoft	mapreduce.job.maps=100 mapreduce.job.reduces=100

Table 3: Benchmark Parameters

micron.com

This technical marketing brief is published by Micron and has not been authorized, sponsored, or otherwise approved by The Apache Software Foundation. Products are warranted only to meet Micron's production data sheet specifications. Products and specifications are subject to change without notice.

©2016 Micron Technology, Inc. Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners. All rights reserved. Rev. B 5/18

Source: CCMMMD-676576390-10607, PDF: CCMMMD-676576390-10617