Micron® and Hortonworks® Power Advanced Big Data Solutions

Flash Energizes Your Analytics

Overview

Competitive businesses rely on the big data analytics provided by platforms like open-source Apache® Hadoop® software to get the answers they need faster.

The volume, velocity and variety of big data demand a lot of storage. For years, legacy HDDs have been the primary storage technology for big data solutions, with price taking precedence over performance (this has been reinforced by the predominance of batch-processed analytics). However, companies are now seeing the value of real-time analytics and faster time to insights.

Historically, one of the major challenges for data scientists has been providing CPUs or GPUs with data fast enough to fully utilize these expensive resources, thereby reducing idle time. CPU/GPU idle time is inefficient and prevents us from attaining real-time results. This is where judicious, cost-effective additions of flash SSDs to existing big-data deployments can help.

Micron's leadership and expertise in developing enterprise SSDs enable us to harness the benefits of advanced memory and flash storage to help reduce the time to actionable answers. Attaining the benefits of real-time analytics requires faster storage, which can be provided by our NVMe™ SSDs.

In this technical brief, we will discuss how we added an NVMe SSD to an all-HDD Hortonworks Data Platform® (HDP) — a leading Hadoop platform implementation for managing large data repositories — as a cost-efficient way to achieve results faster.

Fast Facts

Adding NVMe SSDs to an all-HDD data platform resulted in:

- TCP-DS® benchmarks completing 1.7X faster
- Complex queries experiencing up to 2.5X better response times
- Large clusters experiencing reductions in node count requirements of up to 40% resulting in:
  - Reduced power consumption
  - Reduced cooling
  - Improved reliability
  - Reduced data center space

Micron® 9200 MAX SSD With NVMe

Faster Time to Insights

Adding a single NVMe SSD to an existing Hadoop cluster can improve time to insights by more than 2X depending on the types of queries.

Background

We used a standard benchmark based on the Transaction Processing Council's TPC-DS query set to show that introducing SSDs into existing big data clusters that currently leverage typical "high-performance" 15K RPM SAS HDDs brings real value. This benchmark simulates a typical data analytics workload.

Our goal was to determine if CPU wait times could be reduced by adding a flash SSD as a YARN cache to the Hortonworks cluster. This method of adding flash to an existing cluster gives quick access to the benefits of faster storage without incurring the higher costs of an all-flash replacement.

Test Configuration

For this analysis, we used an early beta version of HDP 3.0. The Hadoop cluster software consisted of a Hortonworks HDP 3.0 Hive database on HDFS/YARN deployed on two separate four-node clusters configured as shown in Table 1.

The two clusters differed only in that one cluster used a group of 15K SAS HDDs and the second cluster used the same HDD configuration plus a single Micron 9200 MAX SSD with NVMe added to each node as a YARN cache.

To ensure true measurement of the storage I/O, the database size-to-memory ratio was targeted at about 2-to-1 (2TB of data with an aggregate cluster memory of 822GB available after operating system overhead).

	Cluster 1 Node Configuration	Cluster 2 Node Configuration
Software		
Operating System	Red Hat® Enterprise Server 7.5 (7.5.1804)	
Hortonworks® HDP	HDP 3.0.0 (beta)	
Hive	3.0.0	
HDFS/YARN	3.1.0	
Hardware		
Server	Supermicro® SYS- 2028U-TNRT+	
CPU	2x Intel® E5-2690v4 at 2.6 GHz	
Memory	256GB (8x 32GB) Micron DDR4-2666	
Network	2x Broadcom® BCM957304 dual-port NIC	
Data Storage	15x 300GB 15K SAS HDDs	
YARN Cache Storage	N/A	1x 3.2TB Micron 9200 MAX SSD with NVMe

Table 1: Test Cluster Configurations

The test performed 94 of the 99 queries used within the TPC-DS benchmark and measured the total completion time for each query on both clusters as well as the individual query completion times. Of the 99 benchmark queries queued, 94 completed with enough confidence to publish the results. Each query was processed three times and a mean result for each query was calculated for this brief.

The Results[1]

Our testing showed that the introduction of a single NVMe SSD to the all-HDD cluster resulted in an average overall 1.7X improvement in benchmark completion times as shown in the figure below.

Figure 1: Completion Time Is Faster With NVMe SSD

The primary contributor to this improvement is the near elimination of CPU wait time. **Error! Reference source not found.** and 3 show system CPU I/O wait time for each cluster. As you can see in Figure 3, the SSD-enhanced cluster CPU wait time is nearly zero, meaning it was not a limiting factor to cluster performance.

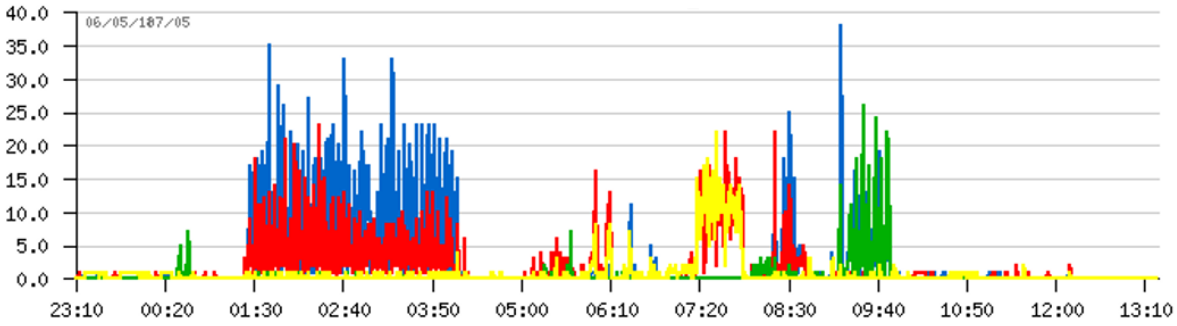

Figure 2: HDD-Only Cluster CPU Wait Time

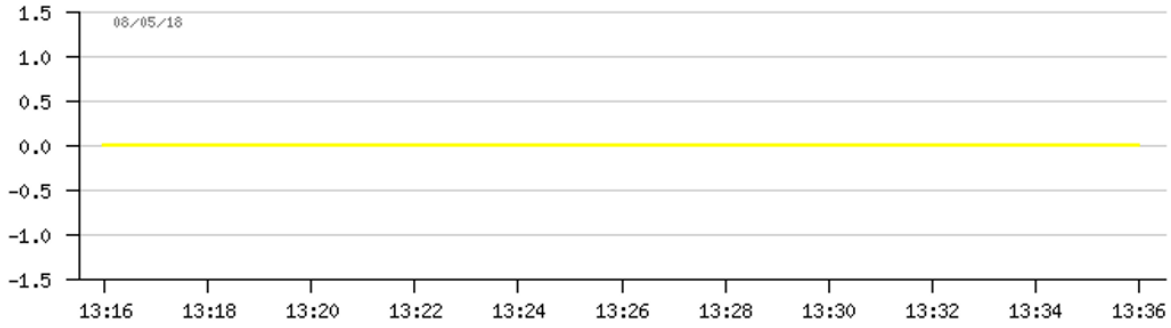

Figure 3: HDDs + NVMe Cluster CPU Wait Time

1. All test results quoted in this brief are based on benchmark tests as described and are for illustrative purposes only. Your actual experience may differ from the results discussed in this brief.

Individual queries for the HDD + NVMe cluster can experience a more than 2X performance improvement over the HDD-only cluster. The majority (55%) of queries resulted in improvement between 1.25X and 1.75X. The queries that showed more than 2X performance improvement are shown in Figure 4. SSD performance is affected by many factors such as query complexity, distribution of data across nodes and read percentage; therefore, actual improvement of real-world workloads may differ.

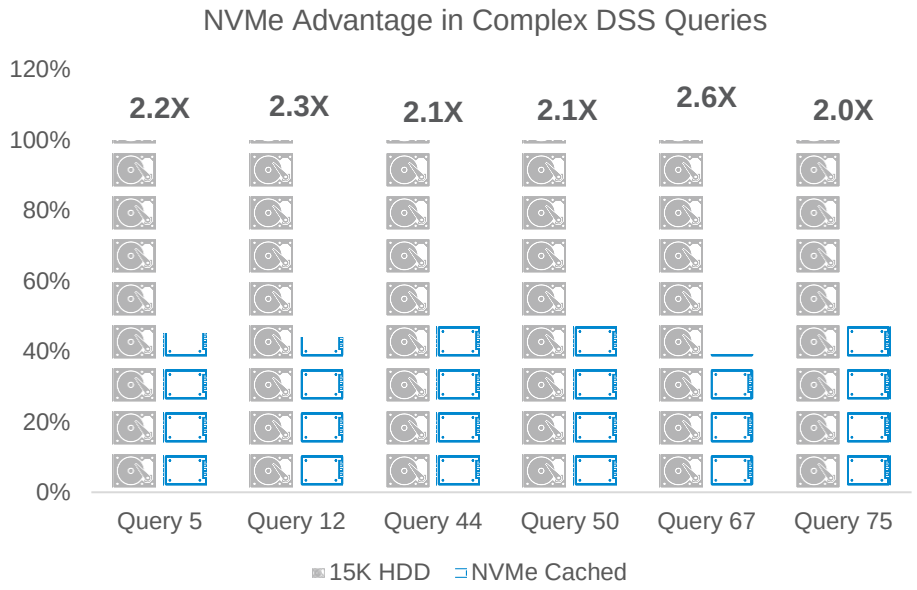

Figure 4: DSS Queries With Largest Improvements Using SSDs

This improvement can directly affect real-world big data solutions. As shown in the figure below, it was determined that an existing 200-node big data cluster used for manufacturing efficiency analytics[2] would require an additional 80 HDD-based nodes to reach the same performance as adding an NVMe SSD to each existing 200 nodes.[3]

Figure 5: Adding NVMe SSD Cache Is Equivalent to Adding 80 HDD-Based Nodes

2. See Dataworks Summit "How to use flash drives with Apache Hadoop 3.x: Real world use cases and proof points" presentation for details.
3. Based on real-world analysis of existing production Hadoop cluster using Apache Hive for data management. Your individual results may vary from those mentioned here.

What could this mean in real dollars? Using the advertised prices of the server configuration described earlier in this brief, 80 nodes would cost approximately $1.8 million. Adding a 3.2TB Micron 9200 MAX SSD to each sever in the existing 200-node cluster would cost only $372,000.[4] This is an 80% potential cost reduction.

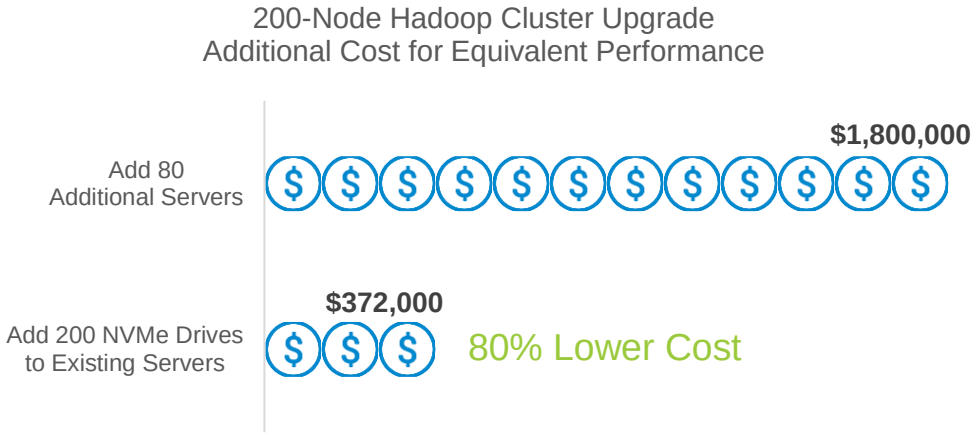

Figure 6: Massive Cost Advantage With NVMe SSDs

The Bottom Line

Adding SSDs can be a cost-effective way to scale existing Hadoop analytics deployments like Hortonworks HDP. Consider adding SSDs if you want faster time to insights than the current all-HDD Hadoop solution can provide. Choosing to add additional servers to your compute cluster can be cost-prohibitive. There are better ways to improve results while staying within budget.

For more information about Micron's SSDs and how they can potentially improve the performance of your existing open-source solutions like Hadoop, visit micron.com.

4. Based on advertised server cost for a server configured as describe in table for Cluster 1, "all-HDD" cluster versus advertised cost of 3.2TB Micron 9200 MAX SSD as of the date of publishing of this brief.

micron.com