

Upgrade Your Data to Solid State Performance

Micron® 5300: World's First Data Center SATA SSD With 96-Layer NAND

Overview

Many performance-focused Cassandra cluster deployments are built with NVMe™ SSDs. But building with NVMe means a new storage interface, new support tactics and tools as well as (potentially) new hardware.

While there's little doubt that NVMe is the interface of choice for long-term data center build outs, if you need to extend the life of your existing infrastructure, SATA SSDs are often a great way to optimize performance at low cost.

Micron's 5300 series of data center SATA SSDs helps you:

- Modernize, economize and maximize the value of the resources you already own
- Realize the economic and performance benefits of next-generation 96-layer NAND

This technical brief highlights an example of the benefits of such an upgrade using a NoSQL database demonstrating increased performance and decreased average read latency.

The terms "performance" and "database operations per second (or OPS)" are used interchangeably.



Figure 1: Micron 5300 SATA SSD Family

Fast Facts

Upgrading from hard disk drive (HDD) legacy storage to the Micron® 5300 PRO for your most demanding Apache Cassandra™ NoSQL workloads brings real benefits and real results with ease:

19X Higher Performance

Up to 19X higher database operations per second compared to twice as many hybrid HDDs.

52% Lower Read Latency

Access your data faster – an upgrade to the Micron 5300 PRO yields up to 52% lower average read latency.



An Easy Upgrade

With a proven, standard SATA interface, it is easy to upgrade to the latest 96-layer, Micron data center SATA SSD.

Optimizing Storage to Unleash Data-Hungry Workloads

When upgrading the storage in any platform, it is important to understand that optimizations may change and that finding the new optimizations may yield significant benefit. We built a test cluster (four Cassandra nodes, each with two Micron 5300 PRO 3.84TB SSDs) and then increased the test workload (by increasing the number of threads) to determine optimal loading.

With a 1TB Cassandra database and the replication factor for the Cassandra database to 3 (this means that there are three copies of the data and the cluster can sustain the loss of two data nodes while continuing operation) we found the results in Figure 1. (Note, Figure 1 shows only 5300 PRO results. Storage upgrade comparisons appear later in this document.)

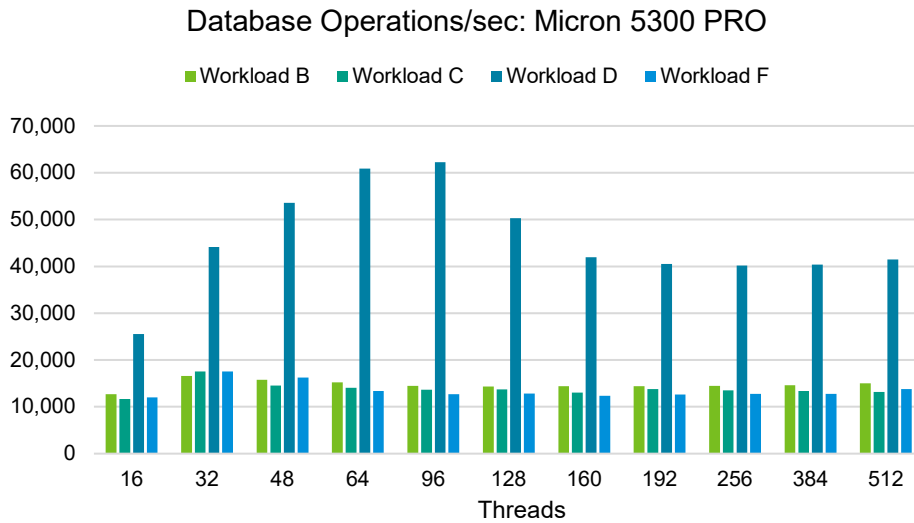


Figure 2: Operations per Second, Micron 5300 PRO Configuration

Optimal thread count for Workload D is easy to see in Figure 2 (96 threads) while others are not. Table 1 shows 5300 PRO test cluster optimal thread counts for each workload:

Workload	Optimal Thread Count ¹
B	32
C	32
D	96
F	32

Table 1: Micron 5300 PRO Test Cluster: Optimal Thread Count by Workload

The remainder of this document uses these thread count results for Micron 5300 PRO cluster analysis combined with legacy platform performance from earlier work².

1. We assume that Administrators would seek to optimize all results
 2. https://www.micron.com/-/media/client/global/documents/products/technical-marketing-brief/qlc_5210_with_cassandra_tech_brief.pdf

The Real Value of a 5300 PRO Upgrade

Once we know how to best load a Micron 5300 PRO cluster, we can analyze the primary question asked of most upgrades: is the upgrade worth it? If we extend the life of an existing infrastructure by simply switching from legacy storage to the 5300 PRO, what do we get in terms of greater database operations per second and lower average response time? As shown below, this simple swap gets us up to a 9X performance increase and a 52% latency improvement.

More Operations per Second: Up to 19X

When we built legacy Cassandra clusters, we scaled out by adding more and more nodes to the cluster. We scaled up by upgrading to larger drives. Sometimes we did both. Adding more legacy nodes was effective (to a point), but it quickly became unwieldy.

Upgrading to larger HDDs can be somewhat effective, but as HDD capacity increases (beyond 2.4TB) spindle speed decreases (from 10K RPM down to 7,200 RPM or fewer), so we could get additional RPM per node, but at the expense of node performance.

More Responsive Clusters: Up to 52%

Many Cassandra deployments rely heavily on fast responses to queries (lower read latency). We compared the average read latency for each configuration by workload.

We used the Yahoo! Cloud Serving Benchmark (YCSB) workloads B–D and F to compare two 4-node Cassandra test clusters: one built with Micron 5300 PRO SSDs and the other built with multiple legacy HDDs.

Figures 3 and 4 show the performance (database operations per second) and average read latency results:

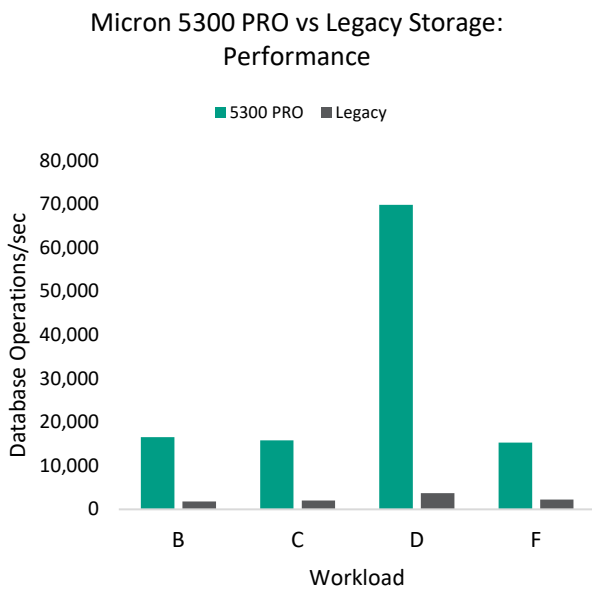


Figure 3: Performance

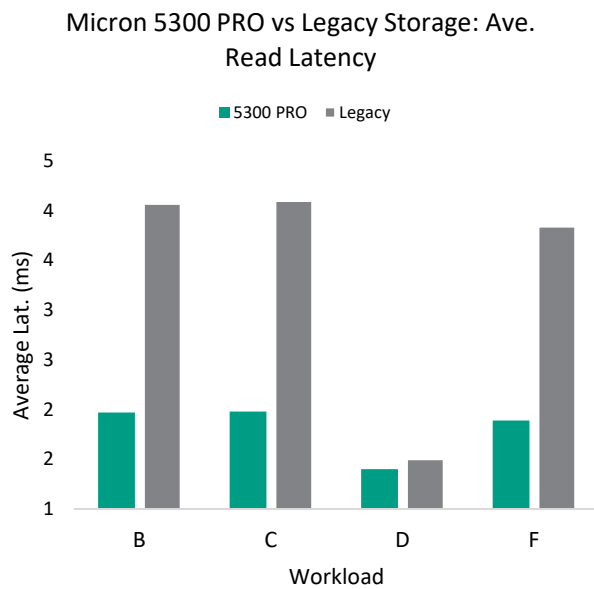


Figure 4: Latency

Figures 3 and 4 show how a simple upgrade from existing legacy storage to Micron 5300 PRO SSDs can drive performance higher while decreasing read latency, as shown in Table 2 (with calculated differences):

Workload	5300 PRO Perf.	Legacy Perf.	5300 PRO Advantage	5300 PRO Aver. Read Lat.	Legacy Aver. Read Lat.	5300 PRO % Lower
B	16,590.82	1,776	9X	1.97	4.07	52%
C	15,844.65	2,018	8X	1.98	4.10	52%
D	69,908.45	3,711	19X	1.40	1.50	7%
F	15,337.61	2,229	7X	1.89	3.84	51%

Table 2: Micron 5300 PRO, Legacy Storage Cluster Comparisons

Table 2 shows the substantial performance and average read latency improvements of the 5300 PRO cluster, with up to 19X higher performance and 52% lower average read latency.

How You Can Use These Results

A simple storage refresh in your existing data center can produce amazing results. We showed YCSB performance increases up to 19X with latency reductions up to 52%, making a 5300 PRO upgrade a cost-effective option to get the results you need. Your needs, systems and infrastructure should determine if a solid upgrade path using the Micron 5300 PRO in your existing, compatible infrastructure is the right choice.

Refresh Your Data Center Today

Extend the life of your infrastructure investment with an easy upgrade to the Micron 5300 PRO SSD. Micron’s 5300 series leverages innovations in NAND technology to help you build a firm foundation on the way to an all-flash future. Maximize the assets you already have (racks, servers, systems) by refreshing your data center with the performance, consistency and expanded data security of the Micron 5300 SSD.

Invigorate Your Infrastructure Now

With cost-effective, 96-layer, 3D triple-layer cell (TLC) NAND, the Micron 5300 rewards your modernization efforts with high performance and fast response times on a proven architecture. Enable strong performance for read-intensive and mixed-use segments: media streaming, BI/DSS as well as OLTP, block and object storage.

With a compelling price-performance ratio, upgrades with the Micron 5300 SSD are a great path forward. With the proven design of the market-leading Micron 5100 and 5200 data center SATA SSDs, the Micron 5300 brings the economies of the latest NAND technology and the ease of a SATA interface upgrade.

Micron’s 5300 features the industry’s broadest product portfolio, best-in-class MTTf reliability ratings and a long-term commitment in the data center SATA space. The Micron 5300 series of SSDs is a key to an easy upgrade that brings real value.

Learn more about Micron’s entire 5300 series line today at <http://www.micron.com/5300>.

How We Tested

The testing methodology is designed to represent a real-world deployment and usage scenario for a Cassandra database. We used four nodes to host a 1TB database with replication factor set to three (three copies of the data and the cluster can sustain the loss of two data nodes while continuing operations).

The database is initially created using YCSB workload A's load parameter. This generated a data set of approximately 3TB (including replication). Table 3 shows the percentage of data owned by each of the four nodes:

Node	Capacity	Tokens	%Owned
Node01	779.18GB	256	74.9
Node02	832.53GB	256	80.1
Node03	777.45GB	256	74.8
Node04	725.03GB	256	70.1

Table 3: Database Ownership and Tokens

Table 4 shows the parameters used when testing each workload: the various workloads.

Parameter	Value	Description
Number of Threads	Load – 240, Execute - Various	Load on the database
Fieldcount	10	Standard 1KB record size
Recordcount	1 Billion	Number of records in the database
Operationcount	1 Billion	Data set size within database
ExecutionTime	60 minutes	Duration of the test

Table 4: Testing Parameters

Testing Configurations

Table 5 shows the main elements of the test platforms used.

Component	5300 PRO Configuration	Legacy Configuration
Server platform	Supermicro® 2028U-TNRT+	Supermicro 2029U-HC0NR
Processor	Xeon™ E5-2690v4	Xeon SP Gold 6142
Drive Count	8 total (2 per node)	16 total (4 per node)
RAID	0	0
Capacity/drive	3.84TB	2.4TB
Type	Micron 5300 PRO	2.5-inch, 10K RPM hybrid

Table 5: Test Platform Configuration

micron.com

©2020 Micron Technology, Inc. All rights reserved. All information herein is provided on an “AS IS” basis without warranties of any kind, including any implied warranties, warranties of merchantability or warranties of fitness for a particular purpose. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners. Products are warranted only to meet Micron’s production data sheet specifications. Products, programs and specifications are subject to change without notice. Rev. B 07/2020 CCM004-676576390-11376