How To Simplify and Shrink Your High-Performance Cassandra Cluster

8TB 5100 ECO SSD[2]

Simpler, Smaller Clusters

Simplify and Shrink	Upgrading your Cassandra storage? Few nodes and fewer racks means less space, less hardware to manage
Reduce License Costs	License software by node? 5100 ECO cuts these costs too!

Micron® 5100 ECO Delivers Massive Scale and Blazing Speed While Cutting Licensing Costs[1]

Apache Cassandra supports massive scale, making it a great choice as data pools have turned into lakes, then oceans. Applications' demands have grown as well. They need real results, and they need them faster. Cassandra can keep up, but legacy storage can't.

This technical brief highlights how the Micron® 8TB 5100 ECO SSD can help tame data growth, reduce cluster complexity, improve delivery and cut licensing costs for your Cassandra cluster.

We used the Yahoo! Cloud Serving Benchmark (YCSB) Workloads A–D and F[3] to compare two four-node Cassandra test clusters — one with our 5100 ECO enterprise SSD, the other built with legacy hard disk drive (HDD) storage. We used these results to calculate relative cluster sizes for similar performance[4] and similar capacity, as well as to compare license costs.

Note: *Due to the broad range of Cassandra deployments, we tested multiple thread counts from 48 to 240. See the How We Tested section below for details.*

5100 Family

Scale from 240GB to 8TB per SSD.[5] Discover more at micron.com.

Massive Scale Meets Massive Performance

When legacy storage HDDs dominated Cassandra deployments, we either scaled out our clusters by adding nodes or we grew our storage inside each node by replacing smaller drives with larger drives.

Adding more nodes was effective, but unwieldy — increasing complexity, cost and rack space requirements. Replacing drives worked to an extent. We got more capacity per node and more capacity per cluster — but that was all we got. Performance stagnated.

The 8TB Micron 5100 ECO SSD is an attractive alternative. A single 8TB 2.5-inch SSD is a great building block to tame data growth and boost Cassandra performance at the same time.

The Capacity For More

If you are looking to increase your existing HDD-based Cassandra cluster's capacity, you may be considering higher-capacity HDDs with plans to pull what you have, replace and start afresh, knowing that you'll get capacity, but that's about all.

With the 8TB 5100 ECO, you can get more. This high-capacity, high-performance SSD can dramatically increase per-node capacity and database operations per second.

We investigated the results of replacing our test cluster's legacy storage (two 1.6TB HDDs) with Micron's 5100 ECO SSDs (two 8TB 5100 ECO SSDs in each node). We measured legacy cluster's YCSB performance for each configuration. Figure 1 shows these results:

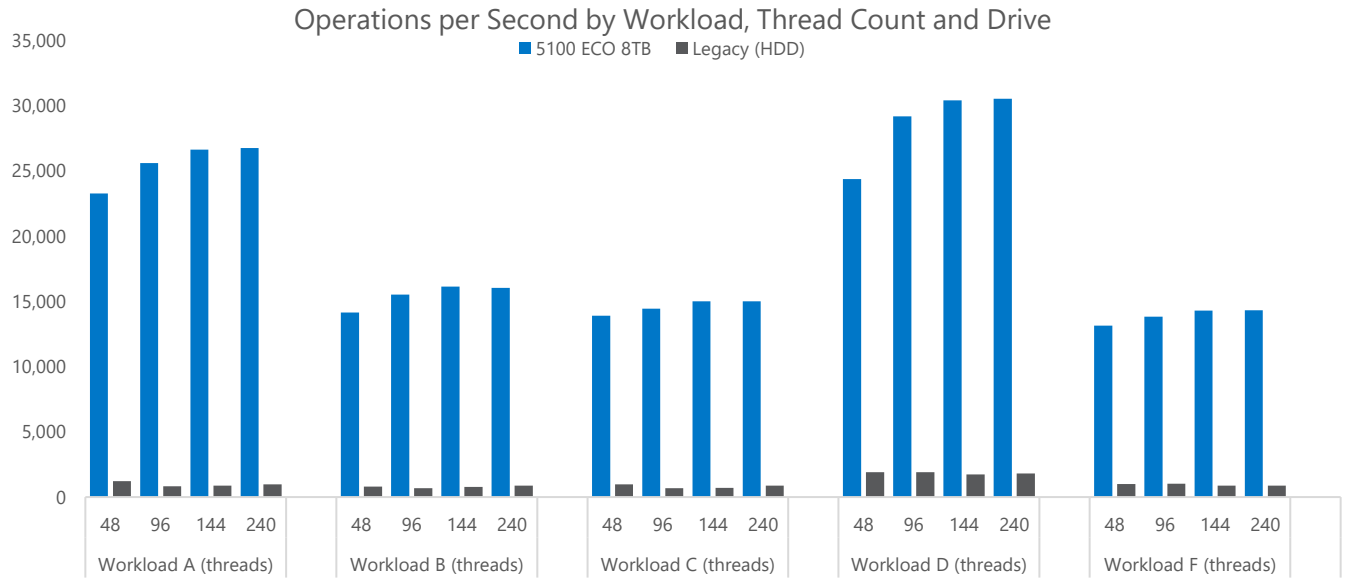

Figure 1: Relative Performance

The results are compelling. The HDD test cluster could store 3.2TB per node while our Micron 5100 ECO test cluster stored 16TB per node, which is a 5X capacity increase. We also saw a tremendous increase in performance over all the workloads and thread counts tested, ranging from a low of 13X to a high of 31X.

Smaller And Simpler With Lower Software License Costs

The Micron 5100 ECO SSD can also help reduce cluster sprawl, complexity and license costs. To illustrate this, we started with our 8TB test cluster performance, as shown in Figure 1, and then calculated the size of a legacy cluster (in racks) that would deliver about the same performance as our four-node 5100 ECO test cluster[5].

In this comparison, we averaged the results shown in Figure 1 across each workload's thread counts for a clearer overall view (see How We Tested for calculation details). Figure 2 shows the results:

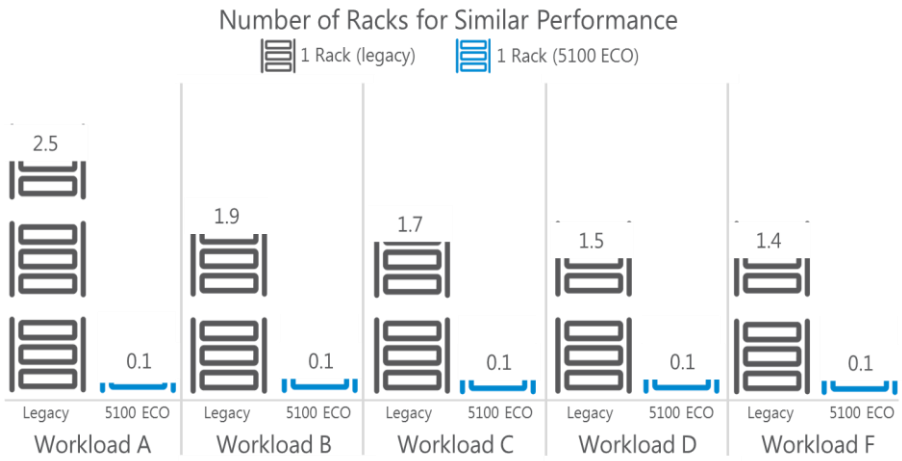

Figure 2: Number of Racks for Similar Performance

The 5100 ECO SSD test cluster (blue) required about one-tenth of a 42U rack (four 1U servers occupied 4U total). The number of racks of legacy nodes we would need for similar performance (gray) varies, but is consistently much greater.

We would need between a rack and a half to as much as two and a half racks of legacy nodes to approximate the performance of our four-node 5100 ECO cluster. The 5100 ECO clusters are much smaller and far less complex for each workload. Smaller clusters also reduce software license cost when licensing software per node (depending on the per-node license cost, cluster level savings may vary).

We can also estimate the software license cost for Figure 2 (clusters of similar performance). Figure 3 shows these calculated results assuming a fixed, per-node software license cost (actual licensing may vary). Our 5100 ECO SSD cluster is always built from four nodes, so its license cost is the same for all workloads (four times the single-node license cost).

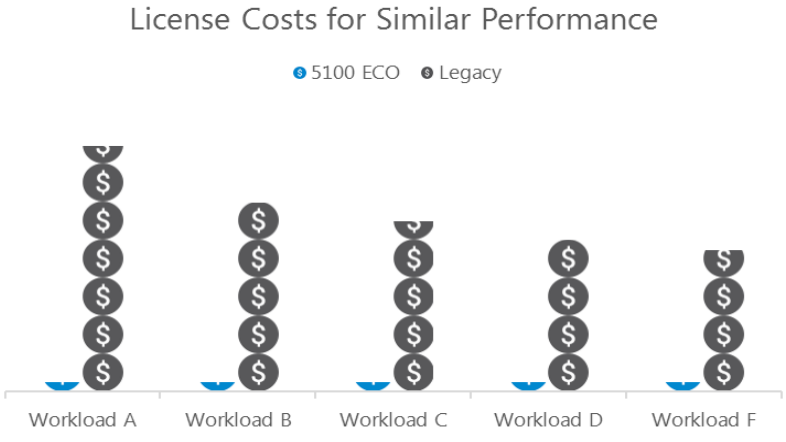

Figure 3: Software License Costs

The legacy cluster license cost is calculated using single-node software license cost multiplied by the number of legacy nodes needed for similar performance. Since the 5100 ECO cluster uses fewer nodes, its overall license cost is much lower.

Naturally, actual license cost reduction will vary depend on licensing model, licensed features, per node pricing, renewal and support terms and other factors, but the potential is compelling.

The Future of Cassandra Is High-Capacity, High-Performance SSDs

Micron's 8TB 5100 ECO SSD brings both massive performance and capacity, reducing cluster complexity, improving delivery and helping cut licensing costs for your Cassandra cluster.

How We Tested

Cluster Performance by Configuration, Workload and Thread Count

We used YCSB to measure performance (database operations per second) for Workloads A–D and F with thread counts ranging over 48 to 240 threads for our two test clusters. Aside from storage, the node configurations and networking were identical:

Cluster	Drive Type	Drive Qty	Platform	CPU	DRAM	Network
5100 ECO	SSD	8	2 CPU	Xeon E5-2690 v4 (14 Core)	256GB	10 GbE
Legacy	HDD					

Table 1: Configurations

Average Cluster Performance by Workload Across Threads

To simplify comparisons and calculations while preserving data fidelity, we calculated the average performance for each test cluster across all tested thread counts, as shown in Figure 4.

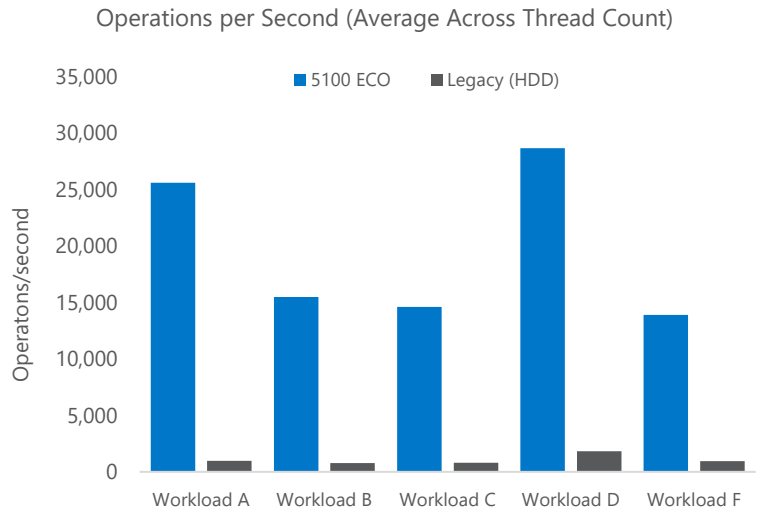

For example, the four-node 5100 ECO cluster's average performance for Workload A is:

$$\text{Workload A Average Cluster Performance} = \frac{\text{Sum of Workload A Performance (All Thread Counts)}}{\text{Number of Tested Thread Counts}}$$

$$= \frac{(23{,}279 + 25{,}613 + 26{,}651 + 26{,}764)}{4}$$

$$= 25{,}777 \text{ Operations per Second}$$

Cluster Performance Ratios by Workload

We divided the average 5100 ECO cluster performance data in Figure 4 by the legacy cluster performance also in Figure 4 to calculate performance ratios of the two test clusters for each workload:

$$\text{Workload A Performance Ratio} = \frac{\text{5100 ECO Cluster Workload A Performance}}{\text{Legacy Cluster Workload A Performance}} = \frac{\text{25,577 Operations/second}}{\text{981 Operations/second}} = 26X$$

This gives a Workload A performance ratio of about 26. Table 2 summarizes average operations per second and calculated ratios[5]:

	Average Operations/sec		
Workload	5100 ECO	Legacy (HDD)	Performance Ratio
A	25,577	981	26
B	15,469	784	20
C	14,599	817	18
D	28,642	1,838	16
F	13,904	939	15

Table 2: Average Operations/second Ratios

Estimated Cluster Size (for Similar Performance) by Workload

To put Table 2 data into a practical perspective, we estimated the number of legacy nodes (and hence rack space) needed to approximate the performance of the four-node 5100 ECO test cluster by multiplying the performance ratio for each workload by four:

		Nodes and Racks (similar performance)			
Workload	Performance Ratio	5100 ECO Nodes	Legacy (HDD) Nodes Needed	5100 ECO 42U Racks	Legacy (HDD) 42U Racks
A	26	4	104	0.1	2.5
B	20	4	80	0.1	1.9
C	18	4	72	0.1	1.7
D	16	4	64	0.1	1.5
F	15	4	60	0.1	1.4

Table 3: Nodes and Racks

micron.com

1. When software is licensed per node. Reduced cost due to 5100 ECO performance (database operations/second) and/or capacity per node resulting in fewer nodes for similar performance and/or capacity and commensurate total software license cost reduction.
2. 5100 ECO 8TB tested. Other capacities and classes available. See 5100 SATA SSD on micron.com for additional details.
3. YCSB Workload E not tested because it is not universally supported.
4. Performance defined as database operations per second.
5. As of the release of this document.
6. Datastax Academy Tutorial, Table 1