# Shrink Your MongoDB Cluster

## Our 5100 ECO Brings You Fast and Powerful MongoDB in Smaller, Space Efficient Clusters


### 5100 ECO vs. HDD

**MongoDB Performance[1]**

| YCSB Workload | 5100 ECO Improvement[2] |
|---|---|
| A | 8.7X |
| B | 9.6X |
| C | 8.5X |
| D | 11.6X |
| F | 7.9X |

1. MongoDB performance (database operations per second) measured at 96 threads. Higher thread counts (up to 240) showed no additional performance gain for most workloads in either configuration. See text for details.
2. 5100 ECO 3-node cluster performance improvement calculated by dividing the 5100 ECO 3-node cluster performance by the legacy 3-node cluster performance. See test for details.

Micron's 5100 ECO's SATA interface enables broad deployment. Its I/O capability enhances MongoDB performance vs. HDDs.

## Overview

MongoDB has the flexibility, adaptability and extensibility to embrace widely varying data types and rapid design/deployment cycles.

To meet these growing demands, we've traditionally scaled MongoDB clusters by adding more nodes; however, budget constraints often make that unsustainable.

Combining MongoDB with our 5100 ECO Enterprise SSD brings amazing results in smaller, simpler deployments compared to legacy storage.

In this technical brief, we compare two MongoDB test clusters, each using the Linux Logical Volume Manager (LVM) for RAID configuration:

- 5100 ECO 3-node cluster: Two Micron 5100 ECO (1.92TB) per node configured as a software RAID 0 (LVM)
- Legacy 3-node cluster: Two 10K RPM SAS HDDs (1.6TB) per node configured as a software RAID 0 (LVM).

We compare MongoDB performance (database operations per second) for these configurations and use that to estimate the size of a legacy cluster needed to approximate the performance of the 5100 3-node cluster for various workloads.

Factors other than storage configuration can affect cluster performance, but while also important, they are beyond the scope of this document.

Your results may vary, but the values in this technical brief show how the 5100 ECO could help enable similar performance in smaller, simpler clusters.

## Database Performance

We measured MongoDB performance (operations per second) for the 5100 ECO and the legacy 3-node clusters using [Yahoo Cloud Serving Benchmark] (YCSB). We tested YCSB Workloads A through D and F, but did not test Workload E as it is not universally supported.

During testing we found that scaling either of our test configurations beyond 96 threads provided no additional performance; this brief is focused on 96 thread results.

| Workload | 5100 ECO | Legacy |
|---|---|---|
| A | 7120.1 | 818.3 |
| B | 10404.4 | 1080.4 |
| C | 11455.4 | 1347.2 |
| D | 28671.5 | 2464.6 |
| F | 6315.0 | 794.7 |

*Table 1: Database Operations per Second*

Table 1 summarizes our findings for the 3-node test clusters.
The 5100 ECO configuration shows much greater performance for all tested workloads.

We used these results to estimate the size of legacy cluster needed to closely match the 5100 ECO 3-node-cluster performance for each workload. These calculated values are estimates only; actual results may vary based on a variety of factors such as network bandwidth, congestion and other deployment-specific variables.

## Estimating Cluster Sizes

We used the data shown in Table 1 to estimate how many legacy nodes we would need to approximate the performance of our 5100 ECO 3-node cluster. We divided the 5100 ECO performance by the legacy performance for each workload to calculate a 'performance multiplier.'

We then multiplied the tested legacy cluster node count (always three) by the performance multiplier for an approximate number of legacy nodes needed for performance that would be similar to that of the 5100 ECO 3-node cluster. The results are shown in Figure 1.

The differing number of legacy nodes by workload is due to the different stresses each workload placed on storage, the varying amount of read and write traffic in each workload and the nature of that traffic. The values in Figure 1 are estimates.

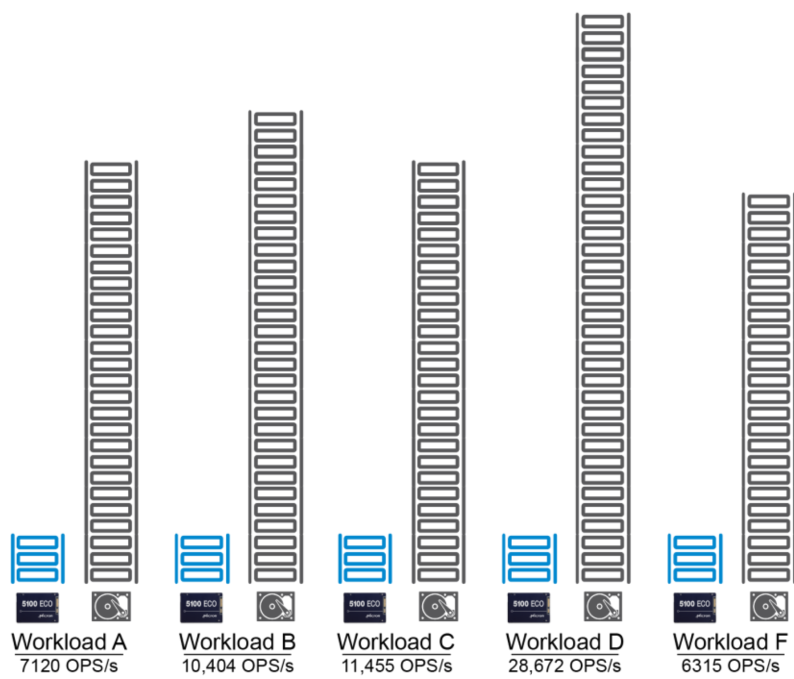


*Figure 1: Estimated Legacy Vs. 5100 ECO Cluster Size for Similar Performance*

For example, Workload A has a performance multiplier of about 8.7. With this value, we can estimate how many legacy nodes are needed to approximate the 5100 ECO 3-node Workload A performance using the following calculation:

Total number of legacy nodes needed = (legacy test cluster node count) x (performance multiplier)

$$= (3) \times (8.7)$$
$$=\sim 26 \text{ nodes}$$

For this example, it would take about 26 legacy nodes to equal the Workload A performance of the 5100 ECO 3-node cluster.

Figure 1 shows the tested workloads along the bottom (A through D and F) and each workload's 3-node 5100 ECO cluster performance. This is the performance level we're approximating with the legacy cluster shown in grey (the approximate number of legacy nodes needed for performance is similar to what we measured for the 5100 ECO 3-node cluster).

We see the greatest difference in Workload D, where about 35 legacy nodes would be needed and the smallest in Workload F, where about 24 nodes would be needed. The remaining workloads range between 26 and 29 legacy nodes to approximate the 5100 ECO 3-node performance.

## About the Workloads We Tested

Workload A is update-heavy with 50% of the total I/Os writing data. At the application level, this workload is similar to recording recent session actions.

Workload B is an update-light, mostly read workload with 5% of the total I/Os writing data. At the application level, this workload is similar to tagging photographs and articles or adding information about videos and music.

Workload C is a read-only workload (100% of the total I/Os read data; there is no write traffic). At the application level, this workload is similar to reading user profiles or other static data where profiles are constructed elsewhere.

Workload D reads the latest entries (most recent records are the most popular). At the application level, this workload is similar to reading user status updates (where users are likely to read the most recent entries). Examples of this workload include social media, frequently changing or updated product literature, or software development repositories.

Workload F is a read/modify/write workload in which records are read, changed and written back. At the application level, this workload is similar to users reading and changing data or tracking user activity.

For further information on YCSB workloads, refer to the Core Workloads page of GitHub.com.

## Why This Matters

We know that many factors affect real-world performance — implementations, infrastructure, tuning and basic design and components. We also know that each of these factors can play an important role in real-world, post-deployment results.

Managing new applications that are hungry for fast storage, developed and deployed quickly, can push storage to its limits, and IT is looking to improve storage efficiency while reducing complexity. Many turn to MongoDB and SSDs like the 5100 ECO.

We used to scale MongoDB by adding nodes to our clusters to meet demand. When demand grew, we added more. But this cycle has a poor long term outlook — we can't reasonably continue with this strategy.

We need better MongoDB systems. We need small, fast and agile. The 5100 ECO lets us reduce our footprint, not grow it, without sacrificing performance.

We can build smaller. We can build simpler. We can build smarter.

micron.com