Turn Massive SQL Data Into Real Insights 13X Faster

Learn More From Your Data With Micron® 8TB ECO SATA SSDs

Overview

The amount of data we're generating, collecting and analyzing is growing at previously unimaginable rates.

Finding valuable, actionable data in the lakes and oceans of data has never been more challenging and never more important.

Enterprises are global. Competition is global. Success relies on a company's ability to make the right decision at the right time — quickly and accurately.

Traditional, rotating storage has long been the default choice for Enterprise Business Intellignece (BI) database platforms, but should it be the default choice today?

SSD capacities of up to 8TB each and a simple, cost effective SATA interface are making businesses rethink how they architect their business-intelligent (BI) systems.

Large-scale BI systems are critical. Storing immense data sets and supporting fast, precise queries on them — with maximum power efficienc — overtaxes traditional storage.

In this technical brief we compare three BI platforms:

1. Built with four Micron 8TB ECO SATA SSDs
2. Built with eight of the same SSDs
3. Built with eight 10K RPM HDDs (legacy configuration)

We found that the SSD configurations eclipse the traditional BI platform, completing 13X to 24X more queries per hour, consuming less energy and supporting more stream counts—illustrating how Micron ECO SSDs provide more value from your data, faster.

8TB ECO SATA SSDs[1]

Vast Data, Keen Insights

Immense Data Sets	Capacities up to 8TB in 2.5" form factor
Faster Insights	Up to 24X more queries per hour than legacy storage
Energy Efficient	Energy-stingy processing

Large-scale BI systems are more critical now than ever before. These systems need to support faster, more precise decisions from oceans of data. Large-capacity, simple deployments using 8TB ECO SSDs drive better systems and better results.

Get Results 13X to 24X Faster

We tested maximum degree of parallelism (Max DoP, a tuning parameter)[2] values from 2 to 28 for both SSD configurations to determine the setting that gives the maximum queries per hour (QPH)[3] for each configuration. We set the traditional configuration's Max DoP=7[4] (the HDD configuration's optimal value).

Because users would typically adjust the stream count to provide the greatest number of QPH, we selected the ECO SSD stream count and DoP providing the greatest QPH and compared with the legacy configuration. Figure 1 shows the results.

The legacy configuration completed two QPH (one stream, the only tested stream count[5]).

The 4x 8TB ECO configuration achieved 26.6 QPH (4 streams) and the 8x 8TB ECO configuration achieved 48.0 QPH (8 streams) for a difference of 13X and 24X.

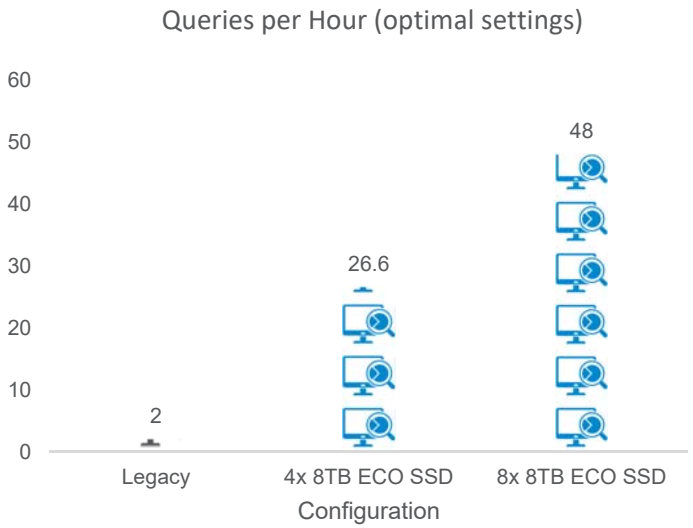

Figure 1: Queries per Hour

	Queries Per Hour		
Configuration	8x HDD	4x 8TB ECO SSD	8x 8TB ECO SSD
Streams	1	4	8
DoP	7	14	28
QPH	2	26.6	48.0
Multiplier	**1X**	**13X**	**24X**

Table 1: Queries per Hour Details

SSDs Give Energy Efficient Insights

We recorded the power draw (kW) for each configuration and how much time each took to process the test query set at DoP=7 (the legacy configuration's optimum), and then calculated total energy consumed (kWh) for each configuration to complete the test query set.

While users would not typically adjust the stream count for DoP to provide the greatest energy efficiency (minimizing completion time is more important), energy efficiency gives a direct comparison of how high-capacity SATA SSDs like the 8TB ECO drives we tested bring more than raw performance to BI platforms.

Figure 2 on the next page shows each configuration's energy consumption to complete the test query sent (with DoP=7). Note that while Figure 2 shows power consumed at different stream counts for the ECO SSD configurations, it shows one stream for the legacy configuration. Table 2 shows additional details (note: the entry "---" indicates that the stream count was not completed by that configuration).

Figure 2 shows that the legacy configuration consumed far more energy (3.24 kWh) to complete the test query set compared to either of the ECO SSD configurations at any number of streams.

In the single stream tests (the far left of each configuration), the 4x ECO configuration used just 10.6% of the energy used by the legacy configuration to complete the same query set.

The 8x ECO configuration showed even lower consumption, just 9.6% of the energy of the legacy configuration (also to complete the same query set).

Despite limiting the Max DOP to 7, Figure 2 shows that we can increase the number of streams for both the ECO SSD configurations while still consuming less energy than the legacy configuration.

The 4x and 8x ECO configurations processed 8X the number of queries while using 59% and 40% of the legacy configuration, respectively.

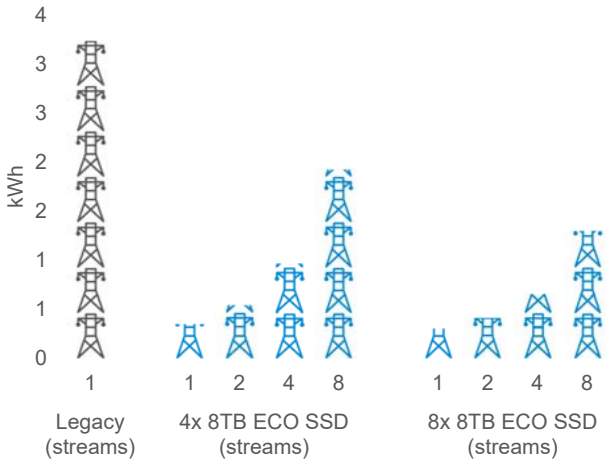

Figure 2: Energy Consumed: Test Query Set (DoP=7)

Summary

Should rotating drives still be the default choice for fast, large-capacity BI platforms? That depends on your particular needs — but high-capacity SSDs are a compelling choice for obtaining real insights from massive data.

	Energy Consumption by Streams (kWh)			
	1	2	4	8
Legacy	3.24	---	---	---
4x 8TB ECO SSD	0.35	0.54	0.97	1.99
8x 8TB ECO SSD	0.31	0.41	0.66	1.30

Table 2: Energy Consumption

Data is everywhere, and more is created every day. Data silos have grown into pools, and pools into lakes. With this unparalleled growth comes incredible complexity—finding the information we need is becoming much harder. Finding it in a timely manner, even more so. Competition on a global scale is helping drive the need for better, more informed decisions faster. Fast, responsive Business intelligence systems (BI) are vital to success.

We tested two 8TB ECO SSD BI platforms (with four and eight drives) against a legacy platform (with eight 10K RPM HDDs). We tested each configuration's BI capability (QPH) when we set test parameters to get the most out of the HDD configuration (as we'd expect users would do when deploying).

We found that the SSDs configuration completed up to 24X the queries per hour of the legacy platform. We also found that when we limited the DoP to 7, the SSD configurations were far more energy efficient—even when processing eight times the data of the legacy platform.

High-capacity, Enterprise SSDs like the 8TB ECO drive better BI systems for better results.

Learn more at micron.com.

How We Tested

We used HammerDB's OLAP implementation for all query performance tests[6]. This free, standardized measurement tool is based on the TPC-H benchmark. It uses a series of 22 business-oriented, ad-hoc queries and concurrent data modifications to gauge platform capability.

To ensure test results matched real-world use, we set two test run exit conditions:

- Run time: The test run time exceeded 12 hours (longer run times decrease the usefulness of the results)
- Loading: Additional loading resulted in lower queries completed per hour (QPH)[3]

The first condition ensures results are available in a reasonable timeframe. The second reflects the common practice of adding load until reaching a maximal queries per hour. When the test met either condition, we stopped the test.

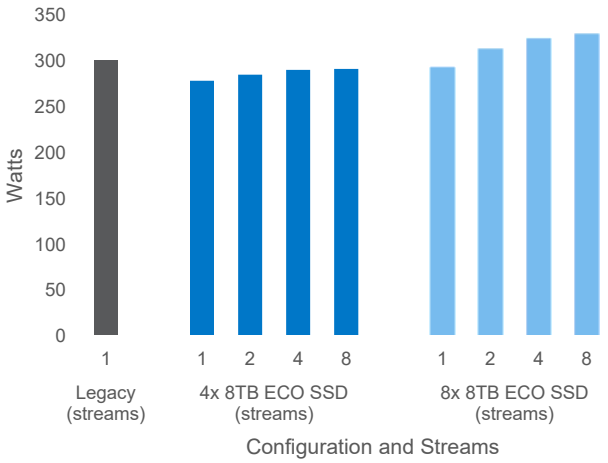

Figure 3: Power Consumption (DoP=7)

We were forced to limit the HDD stream count to one; if we added more streams, the test took longer than 12 hours (making its results less useful in the real world). A single stream test using Max DoP=7 took about 12 hours to complete. Becau se total run time and number of streams have a roughly linear relationship, we fixed the HDD configuration stream count to one and did not test additional stream counts on the HDD configuration. The 8TB ECO SSD configurations completed all tested stream counts.

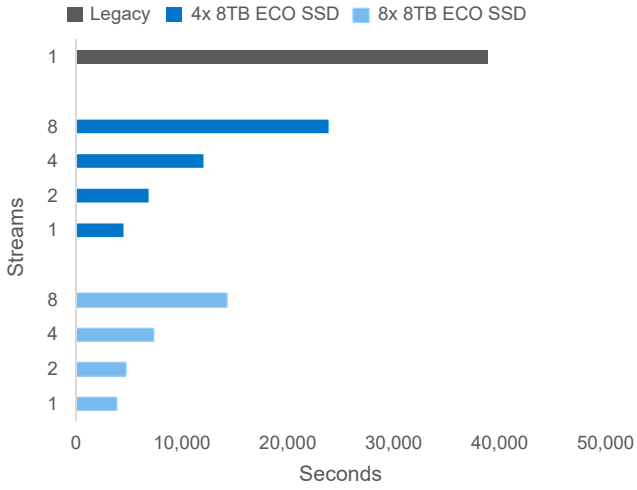

Figure 4: Completion Time

Energy Efficiency

A platform is more energy efficient when it consumes less energy to complete a task (like processing a test query set). Two factors affect energy efficiency: power consumption (in kilowatts) and the time to complete a set of queries. When two test platforms draw similar power, but one completes the test query set faster, that platform is more energy efficient because it consumes less energy to complete the task.

Figure 3 shows the power consumption (in watts) for each configuration with DoP=7 (values do not change much). All configurations at all tested streams consume between about 280 and 330 watts.

The 8TB ECO SSD configurations consume slightly less at lower numbers of streams and slightly more at higher numbers of streams. The legacy configuration is near the middle of this range, consuming about 300 watts.

Figure 4 shows how long (in seconds) it took each configuration to complete the test query set with DoP=7.

[1] 8TB 5100 ECO SSDs tested; other capacities and models available. See micron.com for details. Different capacities and models may show different results.
[2] Max DoP is an adjustable parameter that tells the SQL Server Planner how many parallel operations it can use for a given query.
[3] Queries per hour = $\frac{(\text{Streams})*(22\ \text{queries})}{\text{Total Run Time}}$
[4] Max DoP=7 is the optimal value for the legacy configuration.
[5] With larger stream counts, the HDD configuration took longer than 12 hours to complete the test run.
[6] For additional details on HammerDB and Decision Support testing, see: http://www.hammerdb.com/hammerdb_dssintro.pdf