Unlock SQL Data Value with Fast SSDs

Micron® 5200 ECO SSDs: Faster Queries

Overview

It is well known that the most expensive NVMe™ SSDs are amazingly fast, but how far can cost-efficient SATA SSDs drive critical SQL workloads?

Traditional rotating storage has long been the default choice for enterprise business intellignece (BI) database platforms, but should it be the default choice today?

Affordable SATA SSD capacities have grown to 7.68TB each, and their simple, cost-effective SATA interface is making businesses rethink how they architect their BI systems.

With their high capacity, easy support and nearly ubiquitous SATA interface, SATA SSDs store immense data sets and support fast, precise queries with maximum energy efficiency. Traditional storage can't satisfy these demands.

In this technical brief we compare three BI configurations:

- Four 3.84TB Micron® 5200 ECO SATA SSDs
- Eight of the same Micron SSDs
- Eight 10K RPM 1.2TB HDDs (legacy)

We found that the SSD configurations eclipse the legacy BI platform, completing 5X to 6X more queries per hour with greater energy efficiency—illustrating how Micron 5200 ECO SSDs provide more value from your data, faster.

5200 ECO SATA SSDs

Vast Data, Keen Insights

Immense Data Sets	Capacities up to 7.8TB per SSD[1] (3.84TB capacity tested)
Faster Insights	5X to 6X more queries per hour vs legacy storage
Energy Efficient	Energy-stingy processing

Global competition makes large-scale BI systems more critical now than ever before. We need faster, more precise results to support lightning-quick, accurate decisions from oceans of data. Large-capacity, simple BI systems using high-capacity, affordable SATA SSDs drive these results.

Get the Results You Need, When You Need Them

BI tuning enables DBAs and administrators to optimize platform capability to maximize results (measured in queries per hour[2] [QPH]). Two such tuning parameters are Maximum Degree of Parallelism (Max DoP) and the number of streams (stream count)[3].

Tuning these parameters helps enable the greatest number of QPH. We tested several values for Max DoP and stream count for all three configurations.

We measured each configuration's QPH using the stream count and DoP values providing the greatest QPH for that configuration. (These values may be different for different configurations. See How We Tested for additional details.)

Figure 1 shows the results and Table 1 shows how we calculated the 5X and 6X QPH improvements of the SSD configurations over the legacy configuration[4] (see How We Tested for all measured QPH values).

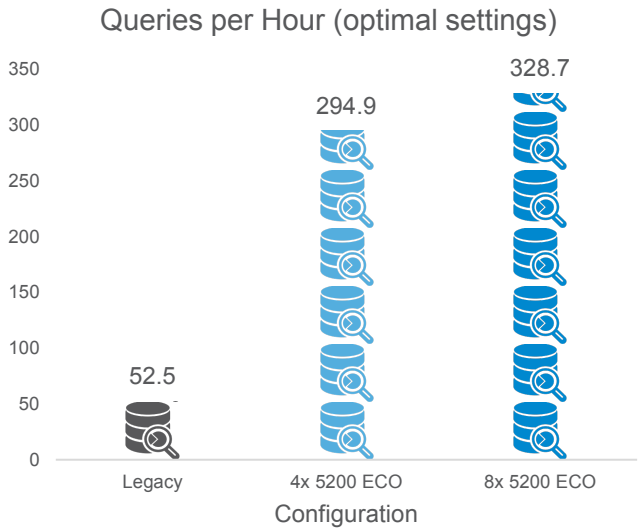

Figure 1: Query Completion Rates

Settings, Values and Ratios			
Configuration	Legacy	4x 5200 ECO	8x 5200 ECO
Streams	1	8	8
DoP	24	24	24
QPH	52.5	294.9	328.7
QPH Ratio	**1X**	**5.6X**	**6.3X**

Table 1: Query Completion Rates (details)

SSDs Provide Energy-Efficient Insights

SSDs drive more than fast results, they drive greater energy efficiency. We calculated energy efficiency—the amount of energy needed to complete a task (such as processing a query set)—by recording the power each configuration consumed (watts) and the time each took to complete the query set. To ensure an unbiased comparison, we did this for one stream (the legacy platform maximum QPH stream count) and eight streams (the SSD configurations' maximum QPH stream count). We used Max DoP = 24 for all measurements (the best Max DoP value for all three configurations).

While users would not typically adjust the stream count or DoP to provide the greatest energy efficiency (minimizing completion time is more important), energy efficiency at each configuration's optimal stream count shows how SSDs like the 3.84TB 5200 ECO bring more than raw performance to BI platforms.

Figures 2a and 2b on the next page show each configuration's energy consumed to complete the test query set. Note that Figure 2a shows energy consumed using the legacy configuration's optimimal QPH stream count (one); Figure 2b shows energy cosumed using the 5200 ECO SSD configurations' optimal QPH stream count (eight). Note too that the legacy configuration consumed significantly more energy with eight streams, requiring us to use a different scale in Figure 2b. Also note that the values are expressed in watt-hours.

These figures show that the legacy configuration consumed far more energy compared to either 5200 ECO SSD configuration at each stream count.

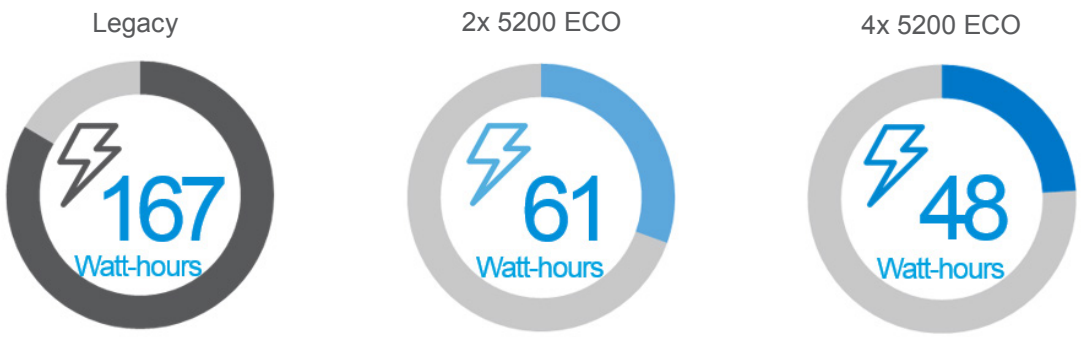

Figure 2a: Energy Consumed to Complete Query Set by Configuration (Stream Count = 1)

Figure 2b: Energy Consumed to Complete Query Set by Configuration (Stream Count = 8; Scale Different from Figure 2a)

Summary

Rotating drives have been the standard storage for BI platforms, but should they continue? Although that utlimately depends on your particular needs, high-capacity SSDs are a compelling choice for obtaining real insights from massive data.

Data is everywhere, and more is created each 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 BI systems are vital to success.

We tested two 3.84TB Micron 5200 ECO SSD BI platforms (with four and eight SSDs) and 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 each configuration (as we'd expect users would do when deploying).

We found that the SSD configurations completed more queries per hour, and did so with much greater power efficiency.

High-capacity enterprise SSDs like the 3.84TB Micron 5200 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[5]. 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 QPH. When the test met either condition, we stopped the test.

Neither condition was met during testing, enabling us to explore several combinations of stream count and Max DoP for each configuration and determine which combination provided best results (measured in QPH).

Table 2 shows these results with the highest QPH for each configuration shown in blue. Note that the legacy configuration showed highest QPH with stream count = 1 while the SSD configurations showed higher QPH with stream count = 8. Note too that setting DoP = 24 yielded maximum QPH for all three configurations.

Configuration	Stream Count	Max DoP					
		4	8	14	24	48	96
Legacy	1	12.2	19.7	36.5	**52.5**	48.8	43.0
	2	34.8	33.3	34.0	33.2	22.7	25.2
	4	37.9	34.9	30.8	32.6	27.8	19.1
	8	34.9	30.8	30.0	22.9	20.9	19.2
4x 5200 ECO	1	36.7	105.2	151.1	150.3	225.6	211.2
	2	70.9	194.2	254.5	269.2	252.0	179.0
	4	155.0	248.7	281.0	283.3	79.4	89.6
	8	170.4	196.3	237.3	**294.9**	151.3	143.6
8x 5200 ECO	1	33.2	106.0	163.0	206.8	246.7	133.6
	2	81.8	200.0	185.0	285.1	311.5	136.3
	4	163.5	300.7	247.7	303.5	185.5	290.7
	8	174.8	295.0	185.5	**328.7**	232.0	172.1

Table 2: Measured QPH by Configuration, Stream Count and Max DoP

Configuration Details

Table 3 summarizes the hardware and software configurations. Note that with columnstore indexes, the total database size is 2190 GiB, far exceeding available memory (see below). This ensures a storage-centric workload.

Item	Configuration Details
RAID	10 (all configurations)
Controller	Dell H740P
HDD Storage	10K RPM 1.2TB HDD (x8)
SSD Storage	Micron 5200 ECO 3.84TB SSD (x4 and x8)

Item	Configuration Details
CPU	Intel® Xeon® Platinum 8168 (x2)
DRAM	256GB, DDR4
SQL	Microsoft SQL Server® 2017 Enterprise Core Edition
OS	Windows Server® 2016 Datacenter Edition

Table 3: Hardware and Software Configuration

1. 5200 ECO 3.84TB SSDs tested; other capacities and models available. See micron.com for details. Different capacities and models may show different results.
2. Queries per hour = $\frac{(\text{Streams})*(22 \text{ queries})}{\text{Total Run Time}}$
3. Max DoP is an adjustable parameter that tells the SQL Server Planner how many parallel operations it can use for a given query.
4. With larger stream counts, the HDD configuration took longer than 12 hours to complete the query set.
5. For additional details on HammerDB and Decision Support testing, see: http://www.hammerdb.com/hammerdb_dssintro.pdf