The science behind the report:

Make business decisions faster with value SAS and NVMe mainstream SSDs from KIOXIA

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report Make business decisions faster with value SAS and NVMe mainstream SSDs from KIOXIA.

We concluded our hands-on testing on May 23, 2019. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on May 21, 2019 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

The table below presents our findings in detail.

We used list pricing information obtained from HPE on June 10, 2019. These numbers, while recent at the time of testing, are subject to change. To calculate the cost-per-iteration metric, we divided the list price for each configuration by the number of data analytics workloads the configuration could complete in two years, running 20 hours a week.

	Enterprise SATA SSDs	KIOXIA RM5 Series value SAS SSDs	KIOXIA CD5 Series NVMe™ mainstream SSDs
Total time to complete a 22-query set (mm:ss)	36:53	27:37	20:00
Avg. data read latency (ms)	160.34	111.38	22.05
Avg. CPU utilization (%)	13	16	37
Price	$38,954	$40,474	$43,642
Cost per iteration (USD)	$11.51	$8.96	$6.99

Storage throughput

As the system runs each query, it pulls data from the disk as quickly as possible for the CPU to process. With the NVMe SSDs' higher transfer speeds, the bursts of throughput for the configuration were much larger in MB/s and shorter in seconds, which translates to a quicker result for each query. As the graph shows, the NVMe mainstream SSD configuration had higher throughput and completed the 22 queries in less time.

Storage throughput

— CD5 Series NVMe mainstream SSDs   — RM5 Series value SAS SSDs   — Enterprise SATA SSDs

MB/s

Seconds
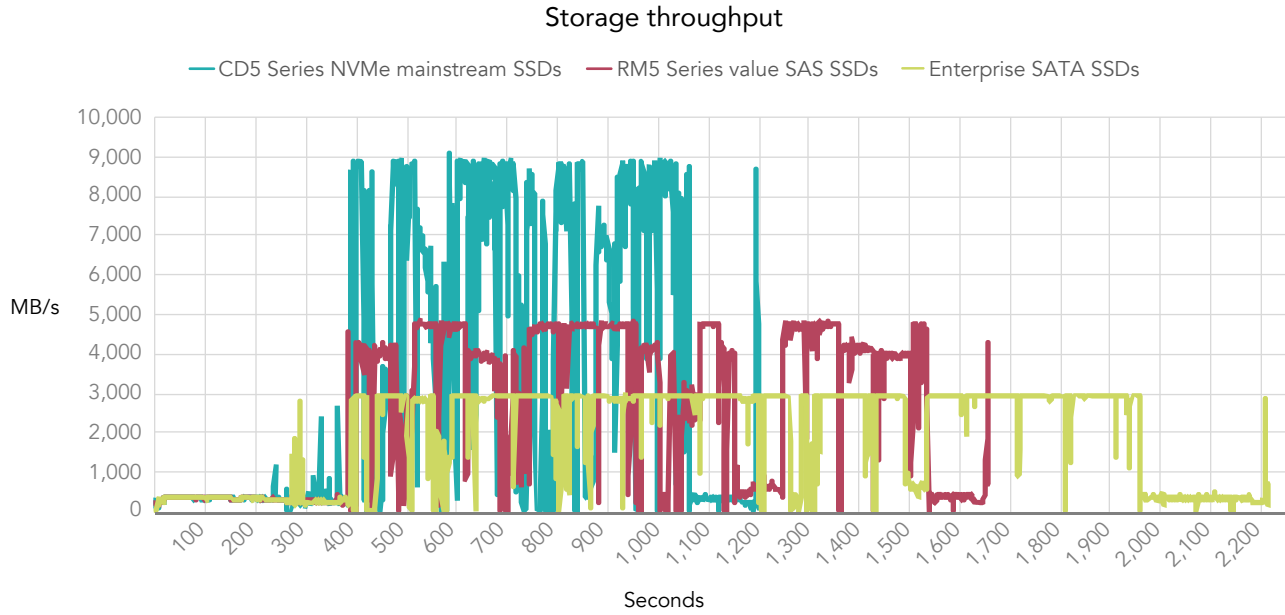
System configuration information

The table below presents detailed information on the systems we tested.

Server configuration information	HPE ProLiant DL385 Gen10
BIOS name and version	A40 v1.40
Operating system name and version/build number	Microsoft Windows Server 2016
Date of last OS updates/patches applied	05/20/18
Power management policy	Balanced
Processor	
Number of processors	2
Vendor and model	AMD EPYC™ 7601
Core count (per processor)	32
Core frequency (GHz)	2.20
Stepping	ZP-B2
Memory module(s)	
Total memory in system (GB)	128
Number of memory modules	16
Vendor and model	Micron MTA18ASF2G72PZ-2G6D1QK
Size (GB)	16
Type	PC4-2666V
Speed (MHz)	2,666
Speed running in the server (MHz)	2,666
Storage controller	
Vendor and model	HPE Smart Array E208i-a SR Gen10
Firmware version	1.98
Driver version	B
Local storage - SATA	
Number of drives	8
Drive vendor and model	HPE VK000960GWCFF
Drive size (GB)	960
Drive information (speed, interface, type)	SATA SSD
Local storage - NVMe	
Number of drives	8
Drive vendor and model	HPE VO000960KWVBR (KIOXIA CD5 Series NVMe mainstream)
Drive size (GB)	960
Drive information (speed, interface, type)	NVMe SSD

Server configuration information	HPE ProLiant DL385 Gen10
Local storage - SAS	
Number of drives	8
Drive vendor and model	HPE VO000960RWUEV (KIOXIA RM5 Series value SAS)
Drive size (GB)	960
Drive information (speed, interface, type)	SAS SSD
Local storage - OS	
Number of drives	1
Drive vendor and model	KIOXIA THNSF8960CCSE
Drive size (GB)	960
Drive information (speed, interface, type)	SATA SSD
Network adapter	
Vendor and model	HPE Ethernet 1Gb 4-port 331i Adapter
Number and type of ports	4 x 1GbE
Driver version	20.14.54
Cooling fans	
Vendor and model	HPE 867810-B21
Number of cooling fans	6
Power supplies	
Vendor and model	HPE 865414-B21
Number of power supplies	2
Wattage of each (W)	800

How we tested

Configuring the RAID volumes

Our testing used three sets of SSDs—eight enterprise SATA, eight value SAS, and eight NVMe mainstream. We used HPE Intelligent Provisioning to create a six-SSD RAID 10 volume for data and a two-SSD RAID 1 volume for logs. We installed the OS on a single SSD in the server backplane.

Installing Microsoft Windows Server 2016

1. Insert the Microsoft Windows Server 2016 installation DVD into the optical drive.
2. Power on the server, and boot into the Windows installer.
3. Once the installer has copied the installation files into the memory, select the appropriate language, time and currency format, and keyboard layout, and click Next.
4. Click Install Now.
5. Enter the product key, and click Next.
6. Choose Windows Server 2016 Datacenter (Server with a GUI), and click Next.
7. Check I accept the license terms, and click Next.
8. Choose Custom: Install Windows only (advanced).
9. From the list, select the OS hard drive, and click Next. This will begin the installation.

Configuring Windows Update on the HPE ProLiant DL385 Gen10 server

1. In the left pane of the Server Manager window, click Local Server.
2. In the main frame, next to Windows Update, click Not configured.
3. In the Windows Update window, in the main pane, click Let me choose my settings.
4. Under Important updates, select Never check for updates (not recommended), and click OK.
5. In the left pane, click Check for updates, and install all available updates.
6. Close the Windows Update window.

Installing Microsoft SQL Server 2017

1. Attach the installation media ISO for SQL Server 2017 to the VM.
2. Click Run SETUP.EXE. If Autoplay does not begin the installation, navigate to the SQL Server 2017 DVD, and double-click it.
3. In the left pane, click Installation.
4. Click New SQL Server stand-alone installation or add features to an existing installation.
5. To accept the license terms, click the checkbox, and click Next.
6. Click Use Microsoft Update to check for updates, and click Next.
7. To install the setup support files, click Install.
8. If the installer doesn't display any failures, click Next.
9. At the Setup Role screen, choose SQL Server Feature Installation, and click Next.
10. At the Feature Selection screen, select Database Engine Services, Full-Text and Semantic Extractions for Search, Client Tools Connectivity, and Client Tools Backwards Compatibility.
11. Click Next.
12. At the Instance Configuration screen, leave the instance default selection, and click Next.
13. At the Database Engine Configuration screen, select the authentication method you prefer. For our testing purposes, we selected Mixed Mode.
14. For the system administrator account, enter and confirm a password.
15. Click Add Current user. This may take several seconds.
16. Click the Data Directories tab.
17. Change the database directory, the database log directory, and the backup directory storage locations to the data, log, and backup volumes.
18. Click the TempDB tab.
19. Change the number of files to 8.
20. Change the initial file size to 1,024 MB.
21. Change the Data Directory to the data volume.
22. Change the initial size of the log to 1,024 MB.
23. Change the Log Directory to the log volume.

24. Click Next.
25. At the Error and usage reporting screen, click Next.
26. At the Installation Configuration Rules screen, check that there are no failures or relevant warnings, and click Next.
27. At the Ready to Install screen, click Install.
28. Once the install finishes, navigate to the Installation tab in the Installation Center, and click Install SQL Server Management Tools.
29. In the browser that pops up, click Download SQL Server Management Studio 17.X.
30. To open the installer, click the download.
31. Click Run.
32. Click Install.
33. Once the installation completes, close the installation window.

Creating a database

1. Open SQL Server Management Studio.
2. Right-click Databases→New Database.
3. Name the database. We named ours `tpch`.
4. To add seven more database files, click Add.
5. Name the database files, and click OK.

Installing HammerDB

1. Download the latest version of HammerDB from www.hammerdb.com/download.html.
2. Double-click the .exe file, choose English, and click OK.
3. Click Yes.
4. Click Next.
5. Choose a destination location, and click Next.
6. Click Next.
7. Click Finish.

Populating the database

1. Open HammerDB, and click Options→Benchmark.
2. Choose MSSQL Server and TPC-H.
3. Expand SQL Server→TPC-H→Schema Build.
4. Double-click Options.
5. Choose 300 scale and 1 virtual user.
6. Check the Clustered Columnstore box.
7. Click OK.
8. Double-click Build. This build could take several hours.

Backing up the database

1. Open SQL Server Management Studio.
2. Right-click the TPC-H database, and click Tasks→Back up….
3. Choose a location to store the backup, and click OK.

Performing the test on Windows Server 2016

1. Start HammerDB.
2. Set the database server to SQL Server, and set the workload to TPC-H.
3. Open the Options panel for the Schema Build: SQL Server→TPC-H→Schema Build→Options.
4. Change the ODBC Driver to ODBC Driver 13 for SQL Server.
5. Change the MAXDOP to 96.
6. For the Scale Factor, select 300.
7. Click OK.
8. Open the Options panel for the Virtual Users: SQL Server→TPC-H→Virtual User→Options.
9. Use 1 Virtual Users.
10. Select the following: Show Output, Log Output to Temp, Use Unique Log Name.
11. Click OK.

12. Open the Options panel for the Driver: SQL Server→TPC-H→Driver Script→Options.
13. Set Total Query Sets per User to 1.
14. Click OK.
15. Load the Driver Script: SQL Server→TPC-H→Driver Script→Load.
16. Start a custom Perfmon data collector set to capture CPU, RAM, and disk performance counters.
17. Click the green arrow.
18. Note the name of the log file. To start the run, click OK.
19. When the run finishes, stop Perfmon.
20. Save the HammerDB results text file and Perfmon output.
21. Open SQL Manager, delete the TPC-H database, and restore the database from the backup file.
22. Reboot the server.

Read the report at http://facts.pt/2h8emuf ▶

This project was commissioned by KIOXIA.