



The science behind the report:

Complete data analysis faster with Google Cloud C3 high CPU instances enabled by 4th Gen Intel Xeon Scalable processors

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 [Complete data analysis faster with Google Cloud C3 high CPU instances enabled by 4th Gen Intel Xeon Scalable processors](#).

We concluded our hands-on testing on July 19, 2023. 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 June 8, 2023 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to <http://facts.pt/calculating-and-highlighting-wins>. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our testing. Source: Principled Technologies.

	Storage type	Time to complete (sec)
n2-highcpu-80	Hyperdisk Extreme	266
c3-highcpu-88	Hyperdisk Extreme	141
C3 vs. N2 speed increase	-	1.88

System configuration information

Table 2: Detailed information on the systems we tested.

General information	
Date testing ended	7/19/2023
Cloud service provider (CSP)	Google Cloud™
Region	us-central-1a

Table 3: Detailed information on the systems we tested.

Workload information	
Workload name and version	HammerDB v4.7 TPROC-H
Workload or software specific parameters	SQL Server: MAXDOP 0, Lock Pages in Memory, 90% Reserved SQL Memory
Iterations and result choice	3 test runs, median reported

Table 4: Detailed information on the systems we tested.

System configuration information	Intel c3-highcpu-88	Intel n2-highcpu-80
Number of instances	1	1
Instance size	c3-highcpu-88	n2-highcpu-80
BIOS name and version	Google, 6/9/2023	Google, 5/27/2023
vCPU	88	80
Number of cores/threads	44/88	20/40
Memory (GB)	176	80
Underlying processor model	Intel® Xeon® Platinum 8481C @ 2.70GHz	Intel Xeon Platinum 8373C CPU @ 2.60GHz
vNIC information/underlying NIC speed	Google Ethernet Adapter, 100/100 Gbps	Google VirtIO Ethernet Adapter, 100/100 Gbps
Operating system information		
Image or template name and UUID	Windows Server 2022 Datacenter	Windows Server 2022 Datacenter
Operating system name	Microsoft Windows Server Datacenter Edition 2022	Microsoft Windows Server Datacenter Edition 2022
Operating system build number	20348	20348
Date patches last applied	06/08/2023	06/08/2023
Changes made from CSP image	Installed SQL Server 2022	Installed SQL Server 2022

System configuration information	Intel c3-highcpu-88	Intel n2-highcpu-80
Instance storage (volume type 1)		
Number of volumes	1	1
Volume use in this test	Database data	Database data
CSP volume type	Hyperdisk Extreme	Hyperdisk Extreme
Volume size (GB)	1,000	1,000
IOPS requested	100,000	100,000
Encryption type	Google-managed	Google-managed
Instance storage (volume type 2)		
Number of volumes	1	1
Volume use in this test	Database logs	Database logs
CSP volume type	SSD Persistent	SSD Persistent
Volume size (GB)	60	60
Encryption type	Google-managed	Google-managed

How we tested

Testing overview

We tested Google Cloud instances featuring 4th Gen Intel Xeon Scalable processors vs. 3rd Gen Intel Xeon Scalable processors versions. We ran a TPROC-H workload on Microsoft SQL Server on GCP to show the time to complete queries for analytic databases that customers can expect to see using the newer C3 instance series vs. the older series. These steps assume that you have already built a database backup.

Uploading your database backup to Google Cloud Storage

1. Open a browser, and log into Google Cloud.
2. Click Cloud Storage, and click Buckets.
3. Click Create.
4. Name the bucket, and give it a label.
5. Under Location type, select Region, and choose the region you will be working in.
6. Select a storage class for the data.
7. Click Create.
8. Click the newly created Bucket, and click Upload Files.
9. Navigate to your database backup folder, and select the backup file.

Creating the Microsoft Windows Server 2022 baseline images

Create the baseline client image instance

1. Log into Google Cloud, and click Go to console.
2. Click Compute engine, and click VM instances.
3. Click Create.
4. In the left window, select New VM instance.
5. Add the following information:
 - Name: Name your instance.
 - Labels: Use any appropriate labels.
 - Region: Select your desired region.
 - Zone: Select your desired zone.
 - Machine family: General-purpose
 - Series: E2
 - Machine type: e2-micro
 - CPU platform: Automatic
 - Boot Disk, and click Change.
 - Operating System: Windows Server
 - Version: Windows Server 2022 Datacenter
 - Boot disk type: SSD persistent disk
 - Size: 50GB
 - Click Select.
 - Identity and API access: App Engine default service account.
 - Firewall: Check Allow HTTP traffic and Allow HTTPs traffic.
6. Click Create.

Configuring Windows Server Datacenter Edition 2022

1. Open Server Manager, and click Local Server.
2. Disable IE Enhanced Security Configuration.
3. Disable Windows Defender Firewall for Domain, Private, and Public networks.
4. Change the time zone to your local time zone.
5. Change the name of your server, and reboot when prompted.
6. Open Server Manager again, and click Local Server.
7. Click to run updates.
8. Run updates, rebooting when prompted, until the server shows no new updates to install.

Installing HammerDB 4.7

1. Download HammerDB from <https://hammerdb.com/download.html>.
2. Double-click the .msi file, choose English, and click OK.
3. Click Yes.
4. Click Next.
5. Chose a destination location, and click Next.
6. Click Next.
7. Click Finish.

Installing Microsoft ODBC Driver 18 for SQL Server

1. Download the latest ODBC Driver for SQL Server from <https://learn.microsoft.com/en-us/sql/connect/odbc/download-odbc-driver-for-sql-server?view=sql-server-ver16>.
2. Double-click the .exe file, and click Next.
3. Check the box to accept the terms, and click Next.
4. Click Next.
5. Click Install.
6. Click Finish.

Creating a machine image of your baseline client instance

1. Log into Google Cloud, and click Go to console.
2. Click Compute engine, and click VM instances.
3. Select the baseline client instance, and click More Actions (the three-dot menu on the right).
4. Click Create new machine image.
5. Give the image a name and optionally a description.
6. Under Location, select Regional, and choose the region you are working in.
7. Leave the Encryption setting default.
8. Click Create.

Adding data and volume disks to the baseline instance

1. Log into Google Cloud, and click Go to console.
2. Click Compute engine, and click Disks.
3. Click Create Disk.
4. Name the disk, and optionally give it a description.
5. Under Location, select Regional, and choose the Region you are working in.
6. Leave the Disk source type as Blank disk.
7. Under Disk type: For the data volume, select Hyperdisk Extreme. For log volume, select SSD persistent disk.
8. Under size, choose 1024GB for the data volume, and 60GB for the log volume.
9. Uncheck Enable snapshot schedule.
10. Add an appropriate label, and leave the rest of the options as default.
11. Click Create, repeating the steps to create both a data and a log volume.
12. Once the disks have been created, click Compute Engine→VM instances.
13. Click into the baseline system under test (SUT) instance, and click Edit.
14. Under Storage→Additional disks, click Attach Existing Disk.
15. Choose the appropriate disk, leave Mode as Read/write, and change the Deletion rule to Delete disk.
16. Click Save.
17. Repeat steps 14 through 16 for the other disk.

Downloading the database backup from Google Cloud Storage

1. RDP into your baseline SUT instance.
2. Open a browser, and log into Google Cloud.
3. Click Cloud Storage, and click Buckets.
4. Select your db backup, and click Download (make sure your selected download destination is large enough to host the file).

Installing SQL 2022 Enterprise on the baseline instance

1. Download or copy the SQL 2022 ISO to the baseline instance, and unzip it.
2. Double-click the Setup application.
3. Click Installation→New SQL Server Standalone installation or add features to an existing installation.
4. Choose the evaluation version, and click Next.
5. Check the I accept the license terms and Privacy Statement box, and click Next.
6. Check the Use Microsoft Update to check for updates (recommended) box, and click Next.
7. On the Install Rules page, click Next.
8. Check the boxes for the following features, and click Next:
 - Database Engine Services
 - Full-Test and Semantic Extractions for Search
9. Leave the Default instance, and click Next.
10. Leave the default Service Accounts, and click Next.
11. On the Server Configuration tab, choose Mixed Mode, and enter and confirm a Password for the SQL Server system administrator (sa) account.
12. Click Add Current User to Specify the SQL Server administrators.
13. Click TempDB and change the initial file size from 8MB to 1024M.
14. Click Next.
15. Once you've passed the rule check, click Next.
16. Click Install.
17. When the install is finished, go back to the SQL Server Installation Center, and click Install SQL Server Management Tools.
18. Download the SSMS file, and install with defaults.
19. Reboot the server when prompted.
20. Run Windows Update one more time to ensure there aren't any new updates for SQL (make sure Windows Updates are set to get updates for other Microsoft products).
21. Once you've installed all available updates, disable Windows Update service by clicking the Start button, typing `services` to open the Services list, and disabling the Windows Update service.

Locking pages in memory

1. Click Start, and type `Local Security Policy`. When it pops up in the search, open the program.
2. Expand Local Policies, and click User Rights Assignment.
3. In the right pane, scroll down and double-click Lock pages in memory.
4. Click Add User or Group, type `NT Service\MSSQLSERVER`, and click OK.
5. To close the Properties window, click OK, and close the Local Security Policy window.

Creating a machine image of your baseline SUT instance

1. Log into Google Cloud, and click Go to console.
2. Click Compute engine, and click VM instances.
3. Select the baseline SUT instance, and click More Actions (the three-dot menu on the right).
4. Click Create new machine image.
5. Give the image a name and optionally a description.
6. Under Location, select Regional, and choose the region you are working in.
7. Leave Encryption setting default.
8. Click Create.

Creating the test instance

To create an instance, you must first have a machine image. The steps below walk you through the creation of an instance from a machine image.

1. Log into Google Cloud, and click Go to console.
2. Click Compute engine, and click Machine images.
3. Select your machine image, and click the More Actions menu at the right.
4. Click Create instance.
5. Add the following information:
 - Name: Name your VM instance.
 - Labels: Use any appropriate labels
 - Region: Select your desired region.
 - Zone: Select your desired zone.
 - Machine family: General-purpose
 - Series: <n2,c3>
 - Machine type: highcpu-80, or highcpu-88 for c3
 - CPU platform: For n2 series: Intel Ice Lake or higher. For c3 series: Intel Sapphire Rapids or higher.
 - Click Create.

Configuring SQL Server on the instances under test

In this section, we list the various SQL Server settings that we changed and the steps to do so.

Setting the SQL memory reserve and max degree of parallelism (MAXDOP)

1. Open the SQL Server Management Studio.
2. Right-click the SQL instance, and click Properties.
3. Click Advanced node, scroll down to the Max Degree of Parallelism, and change the value to 0. Click OK.
4. Right-click the SQL instance again, and go to Memory.
5. Set the Max Memory to 90% of the total memory in the system. Click OK, and close the Properties window.
6. Right-click the SQL instance, and restart the service. When prompted, click Yes.

Running the tests

In this section, we list the steps to run the HammerDB TPROC-H workload on the instances under test. For the maximum number of users we ran, we followed HammerDB TPROC-H recommendations for the size database we were testing. Note that for each test we ran a single-stream test first to cache the database into memory before running the second, multi-stream test. The results we report are for the multi-stream test.

1. On the instance under test, restore the database under test so that the database and log files reside on the disks under test.
2. Ensure that your SQL settings are configured properly according to the instructions above and the instance you're running on.
3. Open HammerDB.
4. Select Options→Benchmark.
5. Choose MSSQL Server and TPC-H.
6. Expand SQL Server→TPC-H→Driver Script.
7. Double-click Options, change the driver to ODBC Driver 18 for SQL Server, set the scale to match your database, and set MAXDOP to 0. Click OK.
8. Expand Virtual User, and double-click Options.
9. Choose 1 user.
10. Check the boxes for Show Output, Log Output to Temp, and Use Unique Log Name.
11. Click OK.
12. Double-click Load to load the Driver Script.
13. Double-click Create users.
14. To capture performance metrics on the system, start Performance monitor set to record CPU, Memory, and drive usage information.
15. To begin the test, click Start.
16. When the run finishes, stop Perfmon and save the HammerDB results file and Perfmon output.
17. Stop the HammerDB user.
18. Double-click User options again and set the number of users to the appropriate count for the multi-stream test.

19. Double-click Create users.
20. To capture performance metrics on the system, start Performance monitor set to record CPU, Memory, and drive usage information.
21. To begin the HammerDB run, click Start.
22. When the run finishes, stop Perfmon, and save the HammerDB results file and Perfmon output.
23. Reboot the instance.
24. Repeat the test two more times for a total of three runs at each user count, and record the median run.

Read the report at <https://facts.pt/8sfbBlx>

This project was commissioned by Intel.



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners.

DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.