

Handle more traffic on your WordPress-based websites with Google Cloud Platform VM instances featuring 2nd Generation Intel Xeon Scalable processors – Cascade Lake

Newer N2 and C2 standard VM instances offered up to 1.60 times the WordPress performance as older N1 VM instances with processors from a previous generation

Companies that host websites on the cloud must pay close attention to the cloud instances that power their work. Newer VM instances with current processors could support enough web traffic to improve the experience for your customers and internal users who rely on consistent access to your company's websites.

At Principled Technologies, we compared the WordPress website performance of three series of Google Cloud Platform VM instances: general-purpose N2 standard VM instances featuring Intel® Cascade Lake processors, older N1 standard VM instances featuring older Intel processors, and compute-optimized C2 standard VM instances featuring Intel Cascade Lake processors. In our tests, both the N2 and C2 VM instances supported more website requests per second than the N1 VM instances.

When your website can comfortably support more traffic, it can mean supporting more users and avoiding frustrating downtime for the users who depend on you.



With N2 standard VM instances:

Handle up to **1.42x** the website requests per second*



With C2 standard VM instances:

Get even greater performance of up to **1.60x** the website requests per second

*compared to older N1 VM instances

How we tested

Google Cloud Platform

We compared the following series of VM instances:

- **Newer N2 standard**
VM instances featuring Cascade Lake processors
- **Older N1 standard**
VM instances featuring older Intel processors
- **Newer C2 standard**
VM instances featuring Cascade Lake processors

Note that Google Cloud Platform does not display the specific model number for the processors in each VM instance. To control for processor model, we ensured that the processor speed information in each VM instance series was consistent from test to test, but we cannot say for certain which specific processors were involved in each test scenario. For additional information on hardware and software configuration, see the [Science behind this report](#).

VM instance sizing

We tested three sizes of VM instance for each series (for a total of 9 VM instance sizes spread out across the N2, N1, and C2 standard series):

- **Small instances:** 4 vCPUs and up to 16 GB of RAM
- **Medium instances:** 8 vCPUs and up to 32 GB of RAM
- **Large instances:** 16 vCPUs and up to 64 GB of RAM

We chose these sizes to demonstrate performance at various levels of work, reflecting a range of real-world use cases. Figure 1 illustrates key specifications for each VM instance. Note that we tested all VM instances in the us-east1-b region.

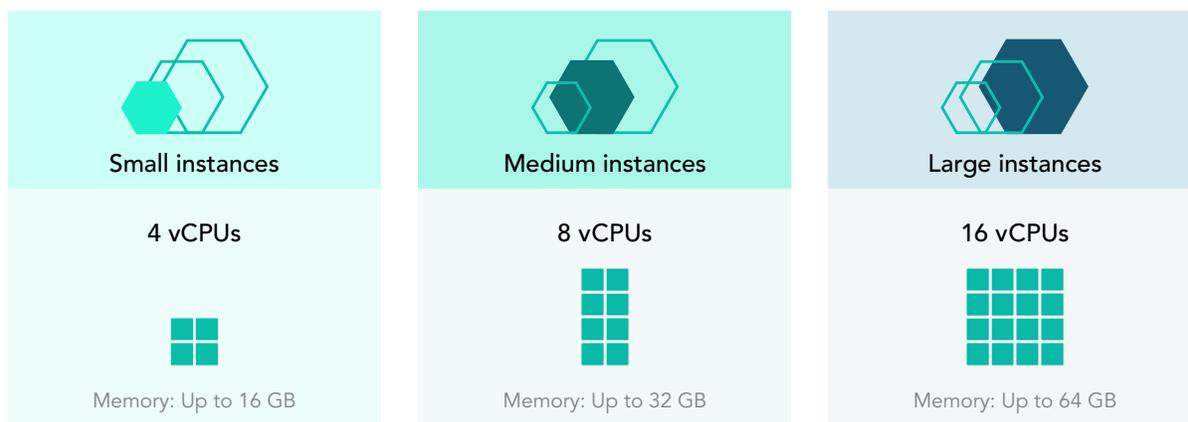


Figure 1: Key specifications for each VM instance. Source: Principled Technologies.

WordPress workload

To run WordPress workloads on the Google Cloud Platform VM instances, we used an open-source website transaction benchmark suite called oss-performance. Intel provided us with modifications to the oss-performance benchmark that enabled it to run using PHP instead of via HHVM. For more information, see the [Science behind this report](#).

Our results: comparing N2 and N1 standard VM instances

One perceived advantage of older instances is their lower cost. After all, why pay more for newer technology when you don't know if the performance boost will be large enough to justify the added cost? Our testing, combined with publicly available cost information, tells a different story.

N2 series VM instances cost 1.02 times as much as the N1 standard series, but they delivered up to 1.42 times the performance. In other words, the N2 standard series VM instances' performance advantage was up to 20 times larger than the cost difference. In the next few sections, we'll detail the performance value that N2 series VM instances can bring your company.

Small VM instances (4 vCPUs)

Application developers who don't need to support as many users may choose smaller VM instances such as those with 4 vCPUs. Figure 2 shows the results of our small VM instance testing. In our testing, the 4vCPU N2 VM instance with Cascade Lake processors fulfilled 1.36 times the number of requests per second (RPS) compared to the N1 series VM instance with older Intel processors.

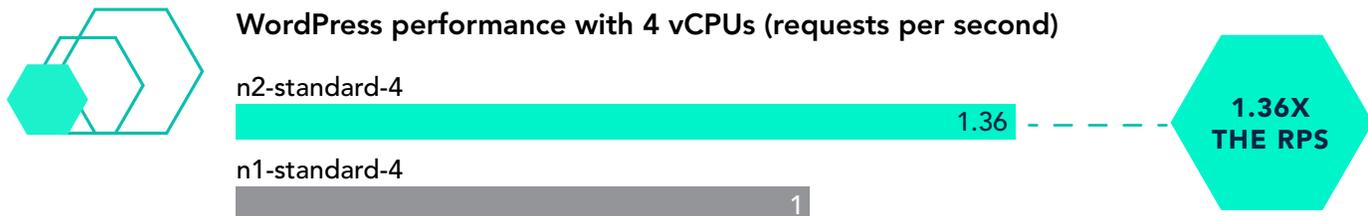


Figure 2: Average requests per second achieved by 4-vCPU VM instances, normalized to the N1 standard series results. Higher is better. Source: Principled Technologies.

Medium VM instances (8 vCPUs)

Developers with large company portals or modest customer-facing websites may benefit from VM instances with a bit more compute power, such as 8 vCPUs each. As Figure 3 shows, the newer N2 VM instances featuring Cascade Lake processors continued to outperform the older N1 VM instances featuring older processors when we moved to 8 vCPUs, handling an average of 1.30 times the number of requests per second.

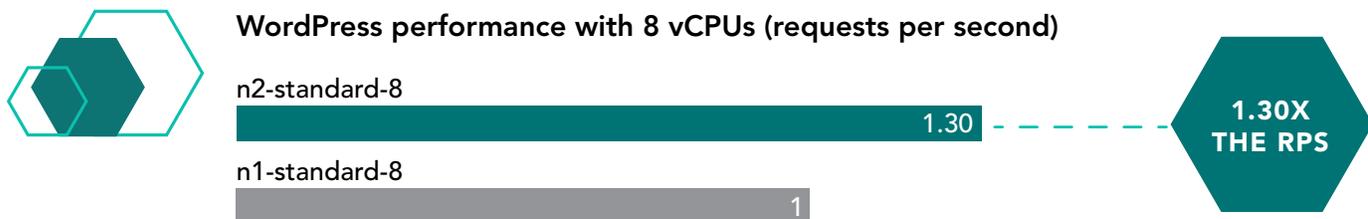


Figure 3: Average requests per second achieved by 8-vCPU VM instances, normalized to the N1 standard series results. Higher is better. Source: Principled Technologies.

How N2 series VM instances help your business

The N2 series offers a few advantages that may have contributed to the improved performance we saw compared to the N1, including:

Higher processor frequency
(2.8GHz vs. 2.2GHz)

Increased memory
at each size

Increased egress network
bandwidth limits

Large VM instances (16 vCPUs)

Developers tasked with supporting large websites may choose to go with larger VM instances of 16 vCPUs or greater. With 16 vCPUs, large N2 VM instances with Cascade Lake processors supported 1.42 times the number of requests per second as the N1 VM instances with previous-generation Intel processors.

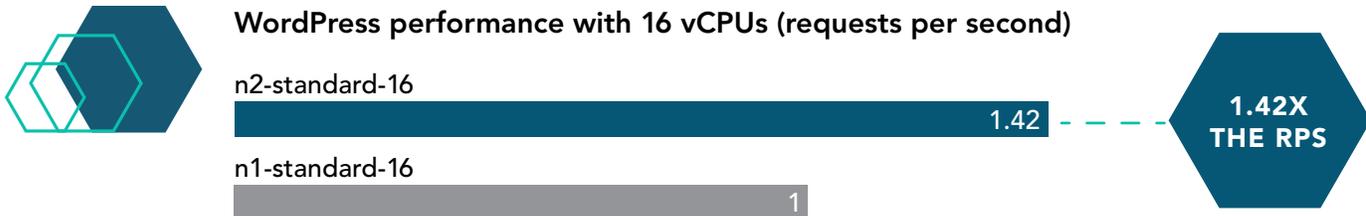


Figure 4: Average requests per second achieved by 16-vCPU VM instances, normalized to the N1 standard series results. Higher is better. Source: Principled Technologies.

Take your work further with C2 VM instances

Application developers seeking even greater website transaction performance than we saw with the N2 standard series might find their match with compute-optimized C2 standard series VM instances. Like the N2 standard series, the C2 series uses Intel Cascade Lake processors and offers more memory than N1 standard series VM instances. C2 also offers a higher base frequency for its CPUs than the N2 standard series, making it more powerful than the general purpose N2 VM instances at each size. So, if your website faces especially busy peak times, the additional processing power and memory of the compute-optimized C2 series could help to cover those performance peaks.

Figure 5 shows a comparison of WordPress workload performance between C2 and N1 standard series VM instances. The C2 series VM instances yielded 1.37–1.60 times the performance, but cost just 1.09 times as much as the N1 VM instances. In other words, the performance gain was up to 6.6 times the cost difference.

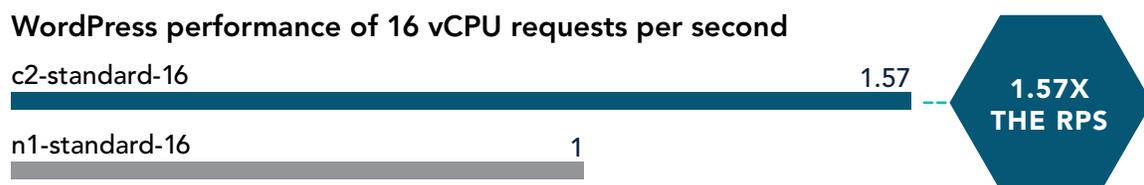
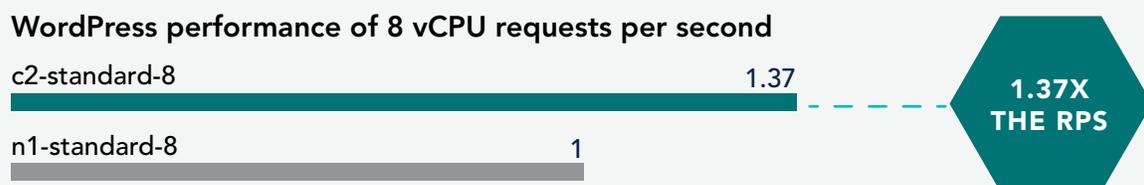
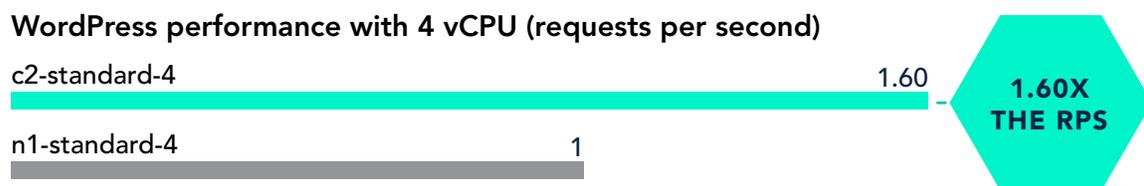


Figure 5: Average requests per second achieved by 4, 8, and 16-vCPU C2 standard series VM instances. Source: Principled Technologies.



Conclusion

In our WordPress web traffic tests on Google Cloud Platform, newer N2 standard VM instances featuring Intel Xeon Scalable Cascade Lake processors handled up to 1.42 times the average number of WordPress website requests per second compared to older N1 standard VM instances that featured Intel Xeon Broadwell processors. For businesses seeking even higher performance, compute-optimized C2 standard VM instances handled up to 1.60 times the average number of WordPress website requests as N1 standard VM instances.

Better-performing WordPress VM instances can mean being able to handle unexpected traffic surges, support a higher baseline of daily users, and give your business room to grow.

Read the science behind this report at <http://facts.pt/dfoB1GE> ►



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.
For additional information, review the science behind this report.

This project was commissioned by Intel.