Small instances:
**Support 1.27x**
the transactions per minute

Medium instances:
**Support 1.20x**
the transactions per minute

Large instances:
**Support 1.41x**
the transactions per minute

# Achieve more PostgreSQL transactions per minute with newer Amazon Web Services instances powered by 2nd Generation Intel Xeon Scalable Processors – Cascade Lake

## M5n instances powered by 2nd Generation Intel Xeon Scalable (8259CL) Cascade Lake processors handled more transactions per second compared to older instances powered by Intel Xeon (E5-2686 v4) Broadwell processors

Ecommerce businesses that are in the market for new cloud instances should seek a solution that delivers strong online transaction processing (OLTP) performance. Purchasing older instances may seem like a good idea due to familiarity or low cost, but newer instances can provide stronger ecommerce performance to help your site keep up with user demand.

At Principled Technologies, we tested an OLTP workload on two series of small, medium, and large general-purpose instances for Amazon Elastic Cloud Compute (Amazon EC2) running PostgreSQL databases:

- Newer M5n instances powered by 2nd Generation Intel® Xeon® Scalable 8259CL processors
- Older M4 instances powered by Intel Xeon (E5-2686 v4) Broadwell processors

In our tests, the newer M5n instances completed more PostgreSQL transactions per minute than the M4 instances, which could help your organization handle more customers and give users a better experience when using your services.

## M5n instances for Amazon EC2

In 2019, Amazon introduced M5n instances to their EC2 offerings, extending 100Gbps networking and adding improved packet-processing performance capabilities to its general-purpose M family.[1] In addition to adding 2nd Generation Intel Xeon Scalable processors to the M-series options, M5n instances support the following:[2]

- Higher CPU core frequency (3.1 GHz vs. 2.3 GHz)
- 25 Gbps of peak bandwidth for small and medium instances, 75Gbps for large instances
- AWS Nitro System, which provides better security and performance compared to older AWS hypervisors[3]
- Options for direct-attached NVMe SSDs instead of EBS volumes
- Higher EBS-dedicated bandwidth

## How we tested

We used an OLTP workload from the HammerDB suite to test the instances for Amazon EC2 upon which we installed PostgreSQL: new M5n instances powered by 2nd Generation Intel Xeon Scalable (8259CL) Cascade Lake processors, and older M4 instances powered by Intel Xeon (E5-2686 v4) Broadwell processors. The M4 instances are available in a variety of different processor configurations, however we chose to test with the Broadwell processors.

The HammerDB developers derived their OLTP workload from TPC-C benchmark specifications. However, because the HammerDB test is not a full implementation of the official TPC-C benchmark, our results are not directly comparable to published TPC-C results.

We compared the M5n and M4 instances across three sizes to demonstrate their value at different workload and performance levels. Like other cloud service providers, Amazon defines an instance's size by its vCPU count. We chose to test instances with 8, 16, and 64 vCPUs.

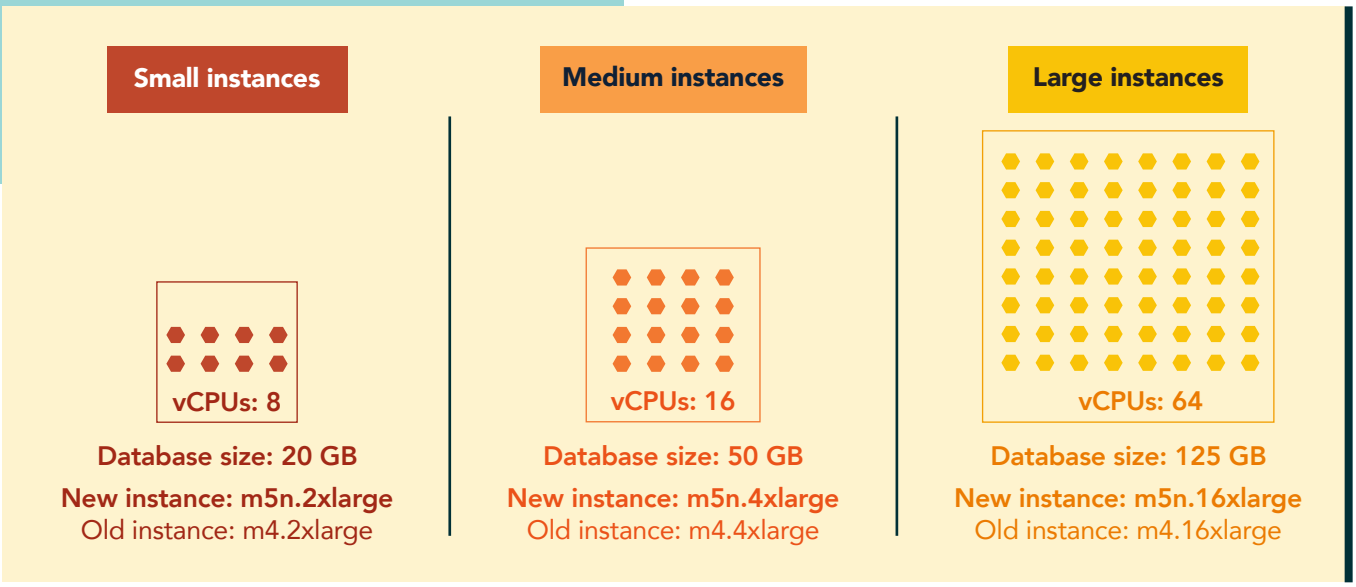


Figure 1: Specifications for the Amazon EC2 VMs we used for testing.
Note that we tested each instance in the us-east-1f region. Source: Principled Technologies.

# Our results

For each instance size we tested, new 2nd Generation Intel Xeon Scalable (8259CL) Cascade Lake processor-powered M5n instances handled more transactions than the Intel Xeon (E5-2686) Broadwell processor-powered M4 instances. With the advantages of the newer instances we saw during testing, organizations would be able to support more online customers and help them have a smoother experience when using online services during peak traffic times.

## Small instances

Small businesses may find that 8-vCPU instances are enough to meet their ecommerce needs and keep customers happy. The newer m5n.2xlarge instances processed 1.27 times the transactions per minute on average compared to the older m4.2xlarge instances.

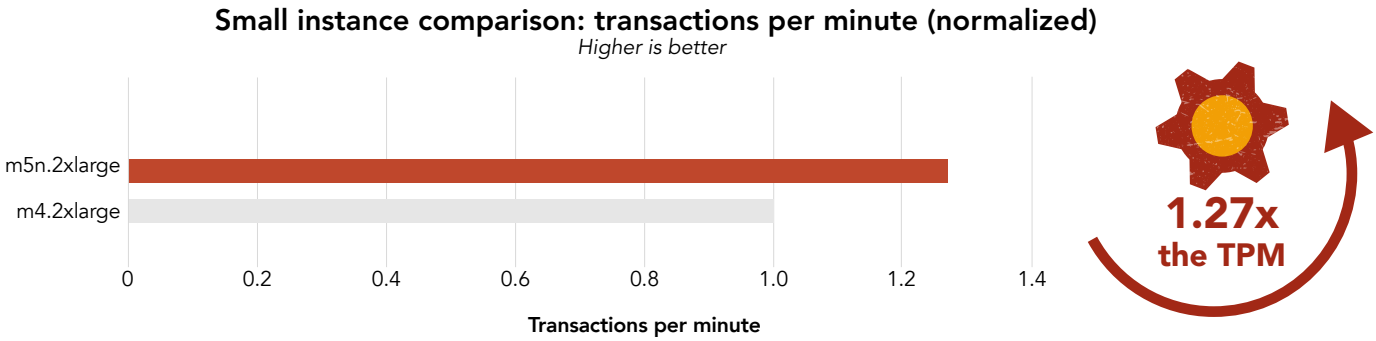


Figure 2: Normalized comparison of the average number of TPC-C-like transactions per minute each small PostgreSQL database instance achieved during the HammerDB workload. A higher rate is better. Source: Principled Technologies.

## Medium instances

Businesses with a larger customer base may require 16-vCPU instances to power their ecommerce work. Figure 3 shows that the new m5n.4xlarge instances again processed more transactions per minute on average, achieving 1.20 times the rate of the older m4.4xlarge instance.

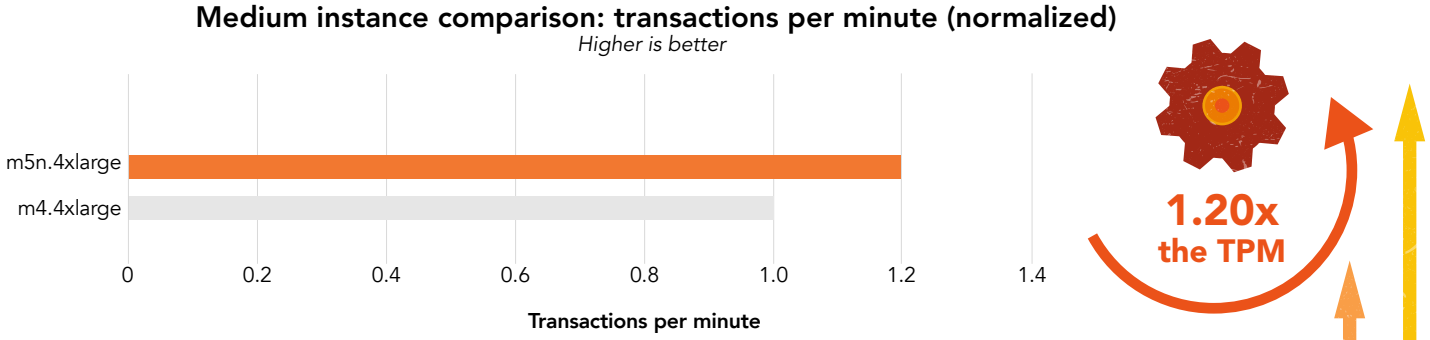


Figure 3: Normalized comparison of the average number of TPC-C-like transactions per minute each medium PostgreSQL database instance achieved during the HammerDB workload. A higher rate is better. Source: Principled Technologies.

## Large instances

Businesses that must meet demand for many thousands of users each day may need 64-vCPU instances or larger to keep up. In our tests, the new m5n.16xlarge instances powered by 2nd Generation Intel Xeon Scalable (8259CL) Cascade Lake processors handled 1.41 times the average transactions per minute as the m4.16xlarge instances.

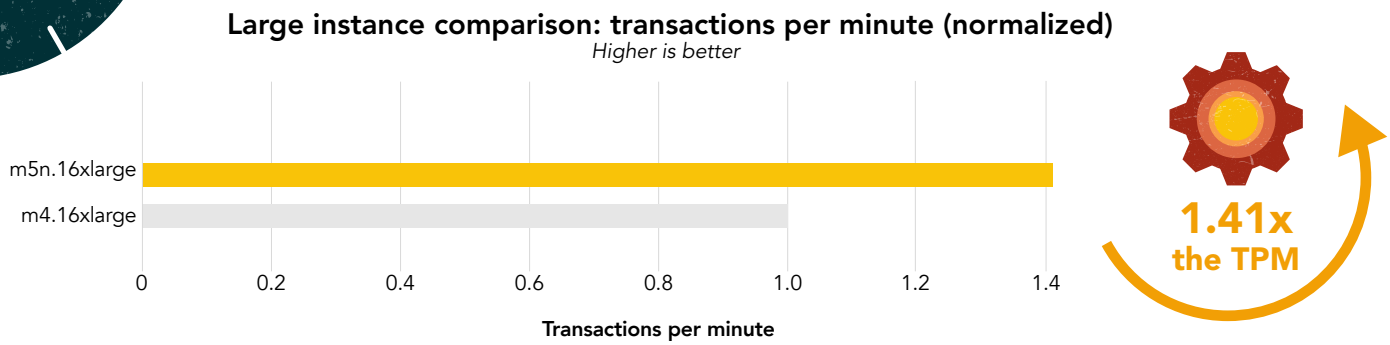


Figure 4: Normalized comparison of the average number of TPC-C-like transactions per minute each large PostgreSQL database instance achieved during the HammerDB workload. A higher rate is better. Source: Principled Technologies.

**1.27 times the TPM**
for small instances

**1.20 times the TPM**
for medium instances

**1.41 times the TPM**
for large instances

## Greater performance at high vCPU count

In our tests, the performance margin between the M5n and M4 series instances grew from 1.20x with 16vCPU instances to 1.41x for large 64vCPU instances.

### Better performance, better investment

Our performance findings suggest the M5n instances can provide anywhere from 1.20 to 1.41 times the rate of transactions per second as the M4 instances. At the time of this writing, however, all sizes and specifications of the M5n instances cost just 1.19 times as much as their M4 counterparts.[4] By investing in M5n instances as opposed to M4 instances, you would be getting better performance for your money.

## Conclusion

No matter what size operation you run, your business' ecommerce site needs to be able to keep up with user demand. Doing so can facilitate customer satisfaction and help to foster repeat patrons.

In our tests, newer 2nd Generation Intel Xeon Scalable (8259CL) Cascade Lake processor-powered M5n-series instances for Amazon EC2 processed more PostgreSQL transactions per second compared to older M4-series instances powered by Intel Xeon (E5-2686 v4) Broadwell processors. Choosing faster instances can help to ensure you meet the needs of your current user base while paving the way for the future.

---

1 Julien Simon, "New M5n and R5n EC2 Instances, with up to 100Gbps networking," accessed November 3, 2020, https://aws.amazon.com/blogs/aws/new-m5n-and-r5n-instances-with-up-to-100-gbps-networking/.

2 "Amazon EC2 Instance Types," accessed November 3, 2020, https://aws.amazon.com/ec2/instance-types/.

3 "AWS Nitro System," accessed November 16, 2020, https://aws.amazon.com/ec2/nitro/.

4 "Amazon EC2 On-Demand Pricing," accessed October 28, 2020, https://aws.amazon.com/ec2/pricing/on-demand/

**Read the science behind this report at http://facts.pt/3YlY74Z ►**

Principled Technologies®

Facts matter.®



This project was commissioned by Intel.