



The science behind the report:

Get more from your Amazon EC2 investment by migrating your workloads to next-generation EC2 instances

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 [Get more from your Amazon EC2 investment by migrating your workloads to next-generation EC2 instances](#).

We concluded our hands-on testing on January 15, 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 December 11, 2018 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

Performance testing

The table below presents our findings in detail.

Instance	Processor	Instructions	Peak measured Gflops
General purpose instances			
m4.large	Intel® Xeon® processor E5-2686 v4	AVX2	38.4
m5.large	Intel Xeon Platinum 8175M	AVX-512	69.0
m5.large	Intel Xeon Platinum 8175M	AVX2	39.7
Compute-optimized instances			
c4.large	Intel Xeon processor E5-2666 v3	AVX2	44.6
c5.large	Intel Xeon Platinum 8124M	AVX-512	76.9
c5.large	Intel Xeon Platinum 8124M	AVX2	46.2
Memory-optimized instances			
r4.large	Intel Xeon processor E5-2686 v4	AVX2	38.6
r5.large	Intel Xeon Platinum 8175M	AVX-512	68.7
r5.large	Intel Xeon Platinum 8175M	AVX2	40.1

Pricing for Amazon EC2 On-Demand instances (US East region)

The table below presents pricing for the instances we tested at the time of testing. Source: Amazon Web Services, "Amazon EC2 Pricing," accessed January 15, 2019, <https://aws.amazon.com/ec2/pricing/on-demand/>.

Instance	Price per hour
General purpose instances	
m4.large	\$0.100
m5.large	\$0.096
Compute-optimized instances	
c4.large	\$0.100
c5.large	\$0.085
Memory-optimized instances	
r4.large	\$0.133
r5.large	\$0.126

Price/performance

The table below presents the peak measured Gflops for each EC2 instance divided by its price per hour.

Instance	Processor	Instructions	Peak measured Gflops/ Price per hour
General purpose instances			
m4.large	Intel Xeon processor E5-2686 v4	AVX2	290
m5.large	Intel Xeon Platinum 8175M	AVX-512	545
m5.large	Intel Xeon Platinum 8175M	AVX2	318
Compute-optimized instances			
c4.large	Intel Xeon processor E5-2666 v3	AVX2	446
c5.large	Intel Xeon Platinum 8124M	AVX-512	905
c5.large	Intel Xeon Platinum 8124M	AVX2	544
Memory-optimized instances			
r4.large	Intel Xeon processor E5-2686 v4	AVX2	384
r5.large	Intel Xeon Platinum 8175M	AVX-512	719
r5.large	Intel Xeon Platinum 8175M	AVX2	414

System configuration information

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

Instance	Processor	# of cores	Memory (GB)
m4.large	Intel Xeon processor E5-2686 v4	1	8.00
m5.large	Intel Xeon Platinum 8175M	1	8.00
c4.large	Intel Xeon processor E5-2666 v3	1	3.75
c5.large	Intel Xeon Platinum 8124M	1	4.00
r4.large	Intel Xeon processor E5-2686 v4	1	15.25
r5.large	Intel Xeon Platinum 8175M	1	16.00

How we tested

1. In the AWS account console, click Launch a Virtual Machine.
2. Select Amazon Linux 2 AMI (HVM), SSD Volume Type and 64-bit (x86).
3. Select General purpose m4.large instance type, and click Review and Launch.
4. Click Launch.
5. Once the instance is up and running, click Connect, and follow the instructions to SSH to the VM.
6. Install LINPACK:
 - a. Log into the Linux VM and download the latest LINPACK benchmark suite from Intel:

```
> wget http://registrationcenter-download.intel.com/akdlm/irc_nas/9752/l_mklb_p_2018.3.011.tgz
```
 - b. Unzip the newly downloaded tarball:

```
> tar -zxvf l_mklb_p_2018.3.011.tgz
```
7. Install the numactl utility:

```
> sudo yum install numactl
```
8. cd to the directory:

```
> cd /home/ec2-user/l_mklb_p_2018.3.011/benchmarks_2018/linux/mkl/benchmarks/linpack
```
9. Edit the problem definition file lininput_xeon64, and customize the problem sizes. The most significant part is to determine the largest problem size that fits into the memory. Based on the available memory size in this instance, start with a problem size interactively and fine-tune it until you find the largest problem size:
10. Run the benchmark:

```
> runme_xeon64
```
11. The benchmark will run one thread per core by default. Change the number of threads to match the number of vCPUs, and rerun the benchmark to see if it makes any difference in performance:

```
> export OMP_NUM_THREADS=<# of threads>
> runme_xeon64
```
12. While the test is in progress, measure the actual CPU frequency:

```
> sudo watch -n1 "cat /proc/cpuinfo | grep 'MHz'"
```
13. At the end of the run, LINPACK reports the best observed performance in number of GFlops. We used the best observed number (peak measured GFlops) as the final score for this test.
14. Run AVX2 instructions on the M5/C5/R5 instances. By default, the M5/C5/R5 instances will run AVX512 instructions. To run AVX2 instructions, set the environment variable and repeat steps 10 through 13 for the m5.large, c5.large and r5.large instances:

```
> export MKL_ENABLE_INSTRUCTIONS=AVX2
```

Read the report at <http://facts.pt/o3zyqcn> ►

This project was commissioned by Amazon Web Services.



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.