





INDEX

- 1. Introduction
- 2. Determining the Right Size for Cloud VMs
- 3. Benchmarking Methodology
- 4. <u>Benchmarking Results</u>
 - a. <u>CPU Benchmark</u>
 - b. <u>CPU Benchmark Command-line</u>
 - c. <u>CPU Benchmark Results: Micro VM</u>
 - d. CPU Benchmark Results: Small VM
 - e. <u>CPU Benchmark Results: Medium VM</u>
 - f. CPU Benchmark Results: Large VM
 - g. <u>CPU Benchmark Results: XL VM</u>
 - h. CPU Benchmark Results: XXL VM
 - i. <u>CPU Benchmark Conclusions</u>
 - j. <u>RAM Benchmark</u>
 - k. RAM Benchmark Command-line
 - I. RAM Benchmark Results: Micro VM
 - m. RAM Benchmark Results: Small VM
 - n. RAM Benchmark Results: Medium

VM RAM Benchmark Results: Large

VM RAM Benchmark Results: XL VM

- o. RAM Benchmark Results: XXL VM
- p. RAM Benchmark Conclusions
- q. Disk Benchmark
- r. Disk Benchmark Command-line
- s. Disk Benchmark Results: MICRO VMs
 - i. <u>IO Bandwidth</u>
 - 5. Disk Write/Read IOPS
- a. Disk Performance Results: SMALL VMs
 - i. <u>IO Bandwidth</u>
 - ii. Disk Write/Read IOPS
- b. Disk Performance Results: MEDIUM VMs
 - i. <u>IO Bandwidth</u>
 - ii. Disk Write/Read IOPS
- c. Disk Performance Results: LARGE VMs
 - i. <u>IO Bandwidth</u>
 - ii. Disk Write/Read IOPS
- d. Disk Performance Results: EXTRA LARGE
- <u>VMs</u>
 - i. <u>IO Bandwidth</u>
 - ii. Disk Write/Read IOPS
- e. Disk Performance Results: EXTRA EXTRA LARGE VMs
 - i. <u>IO Bandwidth</u>
 - ii. Disk Write/Read IOPS
- f. Disk Benchmark Conclusions



- g. <u>Network Benchmark</u>
- h. WAN Network Benchmark Command-line
- i. LAN Network Benchmark Command-line
- j. LAN Network PPS Benchmark Command-

line Network Benchmark Results: Micro VM

Network Benchmark Results: Small VM

k. Network Benchmark Results: Medium VM

Network Benchmark Results: Large VM

- I. Network Benchmark Results: XL VM
- m. Network Benchmark Results: XXL VM
- n. Network Benchmark Conclusions
- 6. Overall Benchmark Conclusion
- 7. Price/Performance Comparison
- 8. <u>Summary</u>

Introduction

Cloud infrastructure services (IaaS) is the fastest growing segment of the public cloud services market. According to <u>Gartner</u>, IaaS revenues worldwide in 2017 were \$30 billion and are expected to grow to \$83.5 billion by 2021, at an average YoY growth of 29%.

Choosing a cloud service provider (CSP) should be a well-researched process that considers many factors including performance, pricing, how well they will support you, do they meet your corporate/industry compliance and security requirements, and so on.

When it comes to performance/price, cloud costs are very much a function of performance SLAs. Without carefully mapping your compute/storage needs to the appropriate resources, you will pay for capacity (bandwidth, RAM, CPU, disk IO, etc.) that you don't need, and those costs can mount quickly. So when doing your due diligence on CSPs, it's important to compare "apples to apples" in terms of performance/price.

This report summarizes benchmarks that were conducted by IOD Cloud Research Technologies on four CSPs: AWS, Azure, GCP, and CaaB—a global CSP operating out of ten global data centers that has been providing high performance, enterprise-grade, highly flexible IaaS services for over a decade.



Determining the Right Size for Cloud VMs

Even though container adoption is growing by the day in modern enterprises, they too still rely on virtual machines for their daily workloads.

From a corporate aspect, many internal enterprise applications, such as CRM / ERP systems, e-mail servers, or file services, cannot be easily containerized. Even if they could, large companies aren't rushing to change what has worked for them for more than a decade, providing them with both stability and required performance. Even in fast paced environments where new compute resources are being created and destroyed every couple of minutes, VMs remain the dominant choice, especially with today's ease of use and customization options that public cloud providers offer.

In your own on premise environment, provisioning VMs is quite a different task than doing so in the cloud. You can go only as high as your server configuration allows you to, and companies tend to overprovision resources until they reach hardware limitations. But the situation in cloud is different, as you pay for every additional CPU core, extra GB of RAM, or IOPS for your storage needs. That's why you need to carefully assess which VMs to provision and how you pick your public cloud provider: not all of the offerings give you the optimal performance at the desired price point.

From a SaaS vendor or start-up perspective, cost-to-performance ratio is a key consideration when choosing a cloud provider. In order to maximize SaaS margins or get the most out of a startup's limited budget, you need a cloud platform that scales easily and has a flexible pricing model that takes into account both predictable baseline traffic as well as demand spikes. In the case of SaaS vendors, high availability, global geographic distribution, robust security and, in some sectors, regulatory compliance are also critical cloud provider requirements. In addition, it is important for these businesses to avoid vendor lock-in; the cloud provider should seamlessly support hybrid and multi-cloud architectures, and even colocation scenarios in which it hosts the customer's equipment and connects it to the cloud.

There are four main categories to keep in mind when you want to provision cloud VMs:

CPU: Besides selecting the number of cores for your VM, be sure to notice which processor type and family the public cloud provider uses in their hypervisor. Also, familiarize yourself with the workload the specific VM will carry: is it an intensive CPU-aware application? Or does the application rely on some other category, for example, network or storage? For legacy apps, be sure it can support multi-core environments.

RAM: Expanding RAM can be quite expensive in the public cloud, so be sure to test your application locally and see how much RAM it requires. Sometimes, if the application will be used behind a load balancer, it's better to scale the application horizontally than to add more RAM to the VM. If your application is RAM-intensive, it won't hurt to check the RAM read and write speeds, for better throughput.

Storage: Second to CPU, the most important category when provisioning cloud VMs is storage. Providers usually offer you the option of choosing the size of the volume and the IOPS capacity. However, both can be bottlenecks for your application if not planned properly. Be sure to research the storage prices of all public providers before you decide to provision your VMs because you may end up having to pay a fortune if you want a volume with 10k IOPS or more.

Networking: Cloud providers usually offer 1 or 10 Gbps network interfaces for your VMs, along with the option of using static IP addresses (both internally and externally) and bridging interfaces if needed. If your VMs will be publicly reachable from the Internet, be sure to position them geographically in the region from where you expect most of the traffic to come. Don't neglect LAN bandwidth, especially if your VMs will communicate a lot among themselves.



Benchmarking Methodology

In order to carry out the benchmarking, an internal network with Internet access was provisioned through each cloud provider. Each network was comprised of 6 general-purpose VMs, ranging from 1 to 8 CPU cores and 1 to 32 GBs of RAM (see the detailed table <u>below</u>).

The VMs were only reachable via the Secure Shell (SSH) network protocol, for the purpose of running benchmark tests. All internal traffic among VMs was permitted because that was required for LAN speed tests.

The following tools were used for the various types of tests:

- sysbench for CPU performance
- **fio** and **hdparm** for disk bandwidth and read/write speed
- **dd** for RAM performance
- ncmeter for LAN speed and ping for packets per second (with one base VM per provider always available to host ncmeter)
- speedtest for WAN speed (download and upload)

All are open-source and free to use.

Sysbench stresses your CPU cores for a limited amount of time with a standard test of verifying prime numbers, which has proven to be accurate for determining CPU performance. This test is customized with the number of CPU cores and simultaneous threads sysbench will use and the maximum number sysbench will verify if it's a prime number.

Fio and **hdparm** are used to calculate read/write speeds and the total number of IOPS per volume, while system **dd** command is used to measure RAM-disk read and write speeds. (dd is present on all Linux distributions, and can be used for both cloning, but also wiping volumes, so be careful with it.)

Another system command useful for benchmarking local network performance is the well-known **ping**, while **ncmeter** will give you the actual speed of your local network compared to 1 Gbps or 10 Gbps per interface promised by your public provider. **Speed test**, used in this benchmark comparison, is a command-line version of a popular test from the <u>www.speedtest.net</u> website.

For the purpose of our benchmark, we divided the most popular VMs from cloud providers into six categories (**Micro**, **Small**, **Medium**, **Large**, **XL** and **XXL**) and for each, we picked the most popular general-purpose VMs from all four tested public cloud providers. These six categories should fit most enterprise workloads, from very light applications or test/dev environments (suitable for **Micro** and **Small** VMs) up to intensive production or internal applications (databases, CRM/ERP, customer-oriented apps) that require highest availability and top-notch performance.



The following table describes the VMs that were provisioned, with every attempt being made to adapt the parameters for each VM so that the comparisons among the providers would be as close to "apples to apples" as possible:

| Category | CSP | VM Туре | CPU Model (all Intel® Xeon®) | Cores | RAM (GB) |
|----------|------|--------------|--|--------------|----------|
| | AWS | t2.micro | CPU E5-2676 v3 @ 2.40GHz | 1 | 1 |
| Micro | AZ | az.B1s | CPU E5-2673 v3 @ 2.40GHz | 1 | 1 |
| Micro | GCP | gcp.MicroVM | CPU @ 2.20GHz | 1 | 1 |
| | СААВ | kt.1p1g | Gold 6150 CPU @ 2.70GHz | 1 | 1 |
| | AWS | t2.small | CPU E5-2676 v3 @ 2.40GHz | 1 | 2 |
| Small | AZ | az.B1ms | CPU E5-2673 v4 @ 2.30GHz | 1 | 2 |
| | GCP | gcp.SmallVM | CPU @ 2.20GHz | 1 | 2 |
| | СААВ | kt.1p2g | Gold 6150 CPU @ 2.70GHz | 1 | 2 |
| | AWS | m5.large | Platinum 8175M CPU @ 2.50GHz | 2 | 8 |
| | AZ | az.F2s | CPU E5-2673 v4 @ 2.30GHz | 2 | 4 |
| Medium | GCP | gcp.MediumVM | CPU @ 2.20GHz | 2 | 4 |
| | СААВ | kt.2p4g | Gold 6150 CPU @ 2.70GHz | 2 | 4 |
| | AWS | There is | no equivalent AWS general-purpose VM i | in this cate | egory |
| Large | AZ | az.F4s | CPU E5-2673 v4 @ 2.30GHz | 4 | 8 |
| | GCP | gcp.LargeVM | CPU @ 2.20GHz | 4 | 8 |



| | СААВ | kt.4p8g | Gold 6150 CPU @ 2.70GHz | 4 | 8 |
|------------------|--|---------------|------------------------------|----|----|
| | AWS m5.xlarge Platinum 8175M CPU @ 2.50GHz | | 4 | 16 | |
| AZ GCF CAA | AZ | az.D4s_v3 | CPU E5-2673 v4 @ 2.30GHz | 4 | 16 |
| | GCP | gcp.XLargeVM | CPU @ 2.20GHz | 4 | 16 |
| | СААВ | kt.4p16g | Gold 6150 CPU @ 2.70GHz | 4 | 16 |
| XXL | AWS | m5.2xlarge | Platinum 8175M CPU @ 2.50GHz | 8 | 32 |
| | AZ | az.F8s | CPU E5-2673 v4 @ 2.30GHz | 8 | 16 |
| | GCP | gcp.XXLargeVM | CPU @ 2.20GHz | 8 | 16 |
| | СААВ | kt.8p16g | Gold 6150 CPU @ 2.70GHz | 8 | 16 |

The benchmarking script, available at <u>https://github.com/cloudwm/cloudbenchmark</u>, was run on each VM separately, with the output results written directly to the VM. Each script command had three arguments: the VM type, which also determined the name of the output CSV file into which the results were written; the IP address of the **ncmeter** receiver for the LAN test; and the block device being used ("xvd", "sd", or "nvme"). The CLIs for the tests that were run are included with the benchmarking results themselves.



Benchmarking Results

CPU BENCHMARK

Using **sysbench**, CPU speed (in seconds) was tested with three different core settings (single, maximum and double) and two concurrent thread loads (10K and 20K) per configuration, i.e., six times per VM:

| Single core (1Core) | 10K threads | 20K threads |
|--------------------------------------|-------------|-------------|
| Maximum number of cores (MaxCore) | 10K threads | 20K threads |
| Double the number of cores (DblCore) | 10K threads | 20K threads |

CPU BENCHMARK COMMAND-LINE

CORES=\$(cat /proc/cpuinfo | grep processor | wc -1)

- # 1Core-10k sysbench --test=cpu --num-threads=1 --cpu-maxprime=10000 run >>/tmp/results/out.cpu1-10k 2>&1
- # MaxCore-10k
- sysbench --test=cpu --num-threads=\$CORES --cpu-max-prime=10000 run >>/tmp/results/out.cpuMax-10k 2>&1
- # DblCore-10k
- sysbench --test=cpu --num-threads=\$((\$CORES*2)) --cpu-max-prime=10000 run
 >/tmp/results/out.cpuDouble-10k 2>&1
- # 1Core-20k
- sysbench --test=cpu --num-threads=1 --cpu-max-prime=20000 run
 >/tmp/results/out.cpu1-20k 2>&1
- # MaxCore-20k
- sysbench --test=cpu --num-threads=\$CORES --cpu-max-prime=20000 run >>/tmp/results/out.cpuMax-20k 2>&1
- # DblCore-20k
- sysbench --test=cpu --num-threads=\$((\$CORES*2)) --cpu-max-prime=20000 run
 >/tmp/results/out.cpuDouble-20k 2>&1



CPU BENCHMARK RESULTS: MICRO VM



| | СааВ | AZ-az.B1s | GCP-gcp.MicroVM | AWS-t2.micro |
|---------------------------|----------|-----------|-----------------|--------------|
| RAM-Write | 1.1 | 1.4 | 0.0873 | 0.182 |
| RAM-Read | 5.9 | 2.9 | 4.1 | 1.2 |
| Actual available RAM (GB) | 0.967541 | 0.969048 | 0.878426 | 0.962578 |



CPU BENCHMARK RESULTS: SMALL VM



| | AWS-t2.small | СааВ | AZ-az.B1ms | GCP-gcp.SmallVM |
|-------------|--------------|---------|------------|-----------------|
| 1Core-10K | 11.6847 | 8.5387 | 11.3769 | 11.3721 |
| MaxCore-10K | 11.6817 | 8.5341 | 11.0362 | 11.2887 |
| DblCore-10K | 11.6978 | 8.5346 | 11.1063 | 11.4414 |
| 1Core-20K | 30.209 | 22.0309 | 29.4407 | 29.0513 |
| MaxCore-20K | 30.0297 | 22.024 | 29.0036 | 29.0066 |
| DblCore-20K | 30.0051 | 22.0383 | 29.194 | 29.09 |



CPU BENCHMARK RESULTS: MEDIUM VM



| | AWS-m5.large | CaaB | AZ-az.F2s | GCP-gcp.MediumVM |
|-------------|--------------|---------|-----------|------------------|
| 1Core-10K | 10.7547 | 8.5966 | 10.6115 | 13.6937 |
| MaxCore-10K | 6.1043 | 4.31 | 5.1818 | 6.1291 |
| DblCore-10K | 6.0942 | 4.311 | 5.1751 | 6.1233 |
| 1Core-20K | 27.8143 | 22.1739 | 25.6734 | 32.129 |
| MaxCore-20K | 15.4872 | 11.1065 | 13.111 | 15.5965 |
| DblCore-20K | 15.4821 | 11.0862 | 13.3114 | 15.625 |



CPU BENCHMARK RESULTS: LARGE VM



| | AWS-m5.large | CaaB | AZ-az.F4s | GCP-gcp.LargeVM |
|-------------|--------------|---------|-----------|-----------------|
| 1Core-10K | 10.7547 | 8.5269 | 12.0556 | 10.9347 |
| MaxCore-10K | 6.1043 | 2.2705 | 2.8962 | 2.8898 |
| DblCore-10K | 6.0942 | 2.2746 | 2.8951 | 3.0538 |
| 1Core-20K | 27.8143 | 22.0369 | 29.6443 | 28.115 |
| MaxCore-20K | 15.4872 | 5.8499 | 7.5768 | 7.5677 |
| DblCore-20K | 5.8254 | 5.825 | 7.9221 | 7.765 |



CPU BENCHMARK RESULTS: XL VM



| | AWS-m5.xlarge | CaaB | AZ-az.D4s_v3 | GCP-gcp.XLargeVM |
|-------------|---------------|---------|--------------|------------------|
| 1Core-10K | 9.378 | 8.5428 | 10.3825 | 11.0567 |
| MaxCore-10K | 2.7081 | 2.2988 | 3.131 | 2.8383 |
| DblCore-10K | 2.6854 | 2.2683 | 3.1702 | 3.0762 |
| 1Core-20K | 24.1955 | 22.1253 | 27.1566 | 28.2983 |
| MaxCore-20K | 6.8462 | 5.8733 | 8.454 | 7.7078 |
| DblCore-20K | 6.814 | 5.8254 | 8.4067 | 7.8204 |



CPU BENCHMARK RESULTS: XXL VM



| | AWS-kt.8p16g | CaaB | AZ-az.F8s | GCP-gcp.XLargeVM |
|-------------|--------------|---------|-----------|------------------|
| 1Core-10K | 10.4386 | 8.5301 | 10.3723 | 11.1414 |
| MaxCore-10K | 1.5009 | 1.079 | 1.7636 | 1.3936 |
| DblCore-10K | 1.4789 | 1.0702 | 1.4468 | 1.4851 |
| 1Core-20K | 26.8972 | 22.0086 | 27.2476 | 28.9378 |
| MaxCore-20K | 3.7907 | 2.7637 | 3.7478 | 3.7076 |
| DblCore-20K | 3.7639 | 2.7567 | 3.7351 | 3.8797 |



CPU BENCHMARK CONCLUSIONS

CaaB's CPUs brings clear benefits in terms of CPU speed:

- In all configurations, CaaB's CPU results were the lowest, i.e., the best performance.
- The lower CaaB CPU results were particularly noticeable in the smaller VMs (micro and small) under the heavier thread load (20K), providing a very advantageous performance/cost ratio.

RAM BENCHMARK

VM memory speed has an important impact on latency. The higher the transfer rate, the better the performance. Using the **dd** test, RAM transfer speed (in GB/sec) was measured during read and write operations for each VM category (micro, small, medium, large, XL and XXL). In the charts below, the actual memory (in GBs) available to the VM during the test is also noted. For each operation the blocksize was defined as 1MiB and the number of blocks to be read or written (count) was limited to 512.

RAM BENCHMARK COMMAND-LINE

```
# RAM-Write (GB/s):
dd if=/dev/zero of=/tmp/TEST_RAM/data_tmp bs=1M count=512 >>
/tmp/results/out.mem-write 2>&1
```

RAM-Read (GB/s):

```
dd if=/tmp/TEST_RAM/data_tmp of=/dev/null bs=1M count=512 >>
/tmp/results/out.mem-read 2>&1
```



RAM BENCHMARK RESULTS: MICRO VM



| | AWS-t2.micro | CaaB | AZ-az.B1s | GCP-gcp.MicroVM |
|------------------------------|--------------|----------|-----------|-----------------|
| RAM-Write | 1.1 | 1.4 | 0.0873 | 0.182 |
| RAM-Read | 5.9 | 2.9 | 4.1 | 1.2 |
| Actual available RAM (GB) | 0.967541 | 0.969048 | 0.878426 | 0.962578 |



RAM BENCHMARK RESULTS: SMALL VM



| | AWS-t2.small | CaaB | AZ-az.B1ms | GCP-gcp.SmallVM |
|---------------------------|--------------|---------|------------|-----------------|
| RAM-Write | 0.818 | 3.4 | 1.4 | 1.8 |
| RAM-Read | 3.3 | 6.2 | 4 | 4.5 |
| Actual available RAM (GB) | 1.95192 | 1.95342 | 1.8628 | 1.94695 |



RAM BENCHMARK RESULTS: MEDIUM VM



| | AWS-m5.large | CaaB | AZ-az.F2s | GCP-gcp.MediumVM |
|---------------------------|--------------|---------|-----------|------------------|
| RAM-Write | 2.7 | 3.2 | 2.3 | 0.868 |
| RAM-Read | 5.8 | 6.1 | 4.9 | 3.8 |
| Actual available RAM (GB) | 7.50777 | 3.85904 | 3.83158 | 3.85263 |



RAM BENCHMARK RESULTS: LARGE VM



| | AWS-m5.large | СааВ | AZ-az.F4s | GCP-gcp.LargeVM |
|---------------------------|--------------|---------|-----------|-----------------|
| RAM-Write | 2.7 | 3.4 | 2.1 | 1.6 |
| RAM-Read | 5.8 | 6.1 | 4.5 | 5.2 |
| Actual available RAM (GB) | 7.50777 | 7.79628 | 7.76908 | 7.78978 |



RAM BENCHMARK RESULTS: XL VM



| | AWS-m5.xlarge | СааВ | AZ-az.D4s_v3 | GCP-gcp.XLargeVM |
|---------------------------|---------------|---------|--------------|------------------|
| RAM-Write | 3 | 2.7 | 2.2 | 1.2 |
| RAM-Read | 5.9 | 6.1 | 4.4 | 3.8 |
| Actual available RAM (GB) | 15.2185 | 15.6713 | 15.6441 | 15.6648 |



RAM BENCHMARK RESULTS: XXL VM



| | AWS-m5.2xlarge | CaaB | AZ-az.F8s | GCP-gcp.XLargeVM |
|---------------------------|----------------|---------|-----------|------------------|
| RAM-Write | 2.8 | 1.6 | 2.2 | 1.5 |
| RAM-Read | 6.1 | 6 | 4.9 | 4.6 |
| Actual available RAM (GB) | 30.6398 | 15.6707 | 15.6441 | 15.6641 |

RAM BENCHMARK CONCLUSIONS

- CaaB's RAM read-write transfer speeds were significantly faster in the Small VM category, and equivalent to the other cloud service providers in all the other categories.
- The one exception to the above is the Read transfer speed in the Micro VM category. Given the consistency of the other results, this outlying result is probably due to a test artefact.
- CaaB read-write memory speed is particularly notable in the Medium and XXL VM categories, this might be related to running VMware hypervisor where



DISK BENCHMARK

Using the **FIO** (Flexible I/O) random-write tester, disk performance was measured on six mounted volumes per VM: the base volume (50 GB) plus five additional volumes as follows:

| Vol. # | AWS | Azure ¹ | GCP | CaaB ² |
|--------|--------------------|--------------------|--------------------|-------------------|
| 1 | 100GB (300 IOPS) | 64GB (300 IOPS) | 10GB (300 IOPS) | 100GB |
| 2 | 100GB (1000 IOPS) | 256GB (1000 IOPS) | 34GB (1000 IOPS) | 100GB |
| 3 | 100GB (5000 IOPS) | 1000GB (5000 IOPS) | 167GB (5000 IOPS) | 100GB |
| 4 | 200GB (10000 IOPS) | 2048GB (7500 IOPS) | 334GB (10000 IOPS) | 200GB |
| 5 | 400GB (20000 IOPS) | | 500GB (15000 IOPS) | 400GB |

Disk performance was benchmarked according to the following parameters:

- IO bandwidth, in MB/sec
- R/W IO, in MB
- R/W BW, in KB/sec
- R/W IOPS

NOTE: The IO bandwidth and the IOPS read-write results shown in this section are for the base volume as well as the volume disks provisioned for 1K and 10K IOPS. These parameter and configuration benchmarks are fully representative of the results recorded for all of the tested disk parameters and configurations.

The maximum IOPS for Azure disks is 7500 IOPS, so the Azure disk tests were done for the base¹ volume plus the four volumes shown in the table.

CaaB provides unlimited IOPS per volume.²



DISK BENCHMARK COMMAND-LINE

fio --directory=/OS-GP2 --name fio_test_file --direct=1 --rw=randwrite --bs=4k --size=1G
--numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-os-gp2-fio-write

fio --directory=/OS-GP2 --name fio_test_file --direct=1 --rw=randread --bs=4k --size=1G
--numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-os-gp2-fio-read

fio --directory=/IO300 --name fio_test_file --direct=1 --rw=randwrite --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io300-fio-write fio --directory=/IO300 --name fio_test_file -direct=1 --rw=randread --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io300-fio-read

fio --directory=/IO1K --name fio_test_file --direct=1 --rw=randwrite --bs=4k -size=1G numjobs=16 --time_based --runtime=180 --group_reporting --norandommap >>
/tmp/results/out.disk-io1k-fio-write

fio --directory=/IO1K --name fio_test_file --direct=1 --rw=randread --bs=4k -size=1G -numjobs=16 --time_based --runtime=180 --group_reporting --norandommap >>
/tmp/results/out.disk-io1k-fio-read

fio --directory=/IO5K --name fio_test_file --direct=1 --rw=randwrite --bs=4k -size=1G -numjobs=16 --time_based --runtime=180 --group_reporting --norandommap >>
/tmp/results/out.disk-io5k-fio-write

fio --directory=/I05K --name fio_test_file --direct=1 --rw=randread --bs=4k -size=1G -numjobs=16 --time_based --runtime=180 --group_reporting --norandommap >>
/tmp/results/out.disk-io5k-fio-read

fio --directory=/IO10K --name fio_test_file --direct=1 --rw=randwrite --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io10k-fio-write

fio --directory=/IO10K --name fio_test_file --direct=1 --rw=randread --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io10k-fio-read

fio --directory=/IO20K --name fio_test_file --direct=1 --rw=randwrite --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io20k-fio-write

fio --directory=/IO20K --name fio_test_file --direct=1 --rw=randread --bs=4k -size=1G --numjobs=16 --time_based --runtime=180 --group_reporting -norandommap >>
/tmp/results/out.disk-io20k-fio-read



DISK BENCHMARK RESULTS: MICRO VMS

IO BANDWIDTH



| | AWS-t2.micro | CaaB | AZ-az.B1s | GCP-gcp.MicroVM |
|---------------|--------------|---------|-----------|-----------------|
| OS_GP2 (Base) | 34.66 | 1181.1 | 4.92 | 32.09 |
| Ю1К | 81.21 | 1592.24 | 8.44 | 142.31 |
| Ю10К | 81.35 | 1533.29 | 7.47 | 235.28 |



DISK WRITE/READ IOPS



| | AWS-t2.micro | CaaB | AZ-az.B1s | GCP-gcp.MicroVM |
|----------------|--------------|-------|-----------|-----------------|
| OS-GP2-WR-IOPS | 3058 | 17205 | 193 | 1485 |
| OS-GP2-R-IOPS | 3052 | 62693 | 24385 | 749 |
| 1K.WR-IOPS | 1022 | 16806 | 213 | 6283 |
| 1K-R-IOPS | 1021 | 64819 | 18382 | 6680 |
| 10K-WR-IOPS | 1022 | 16815 | 213 | 6283 |
| 10K-R-IOPS | 1021 | 64857 | 18382 | 6680 |



DISK PERFORMANCE RESULTS: SMALL VMS

IO BANDWIDTH



| | AWS-t2.small | CaaB | AZ-az.B1ms | GCP-gcp.SmallVM |
|---------------|--------------|---------|------------|-----------------|
| OS-GP2 (Base) | 34.777 | 1232.08 | 9.22 | 62.16 |
| IO1K | 81.25 | 1572.26 | 8.44 | 247.02 |
| Ю10К | 81.37 | 1694.24 | 7.76 | 246.03 |





DISK WRITE/READ IOPS



| | AWS-t2.small | CaaB | AZ-az.F2s | GCP-gcp.SmalloVM |
|----------------|--------------|-------|-----------|------------------|
| OS-GP2-WR-IOPS | 3060 | 12753 | 520 | 1485 |
| OS-GP2-R-IOPS | 3051 | 90552 | 131000 | 749 |
| 1K.WR-IOPS | 1022 | 17314 | 607 | 6295 |
| 1K-R-IOPS | 1021 | 66533 | 789 | 6482 |
| 10K-WR-IOPS | 1022 | 17173 | 607 | 6295 |
| 10K-R-IOPS | 1021 | 66169 | 789 | 6482 |



DISK PERFORMANCE RESULTS: MEDIUM VMS





| AWS-m5.large | CaaB | AZ-az.F2s | GCP-gcp.MediumVM |
|--------------|--|---|---|
| 55314 | 71682 | 4474.9 | 49845 |
| 74188 | 408026 | 4679.8 | 73857 |
| 55314 | 72928 | 4474.9 | 49845 |
| 74188 | 405792 | 4679.8 | 73857 |
| | AWS-m5.large 55314 74188 55314 74188 | AWS-m5.large CaaB 55314 71682 74188 408026 55314 72928 74188 405792 | AWS-m5.largeCaaBAZ-az.F2s55314716824474.9741884080264679.855314729284474.9741884057924679.8 |



DISK WRITE/READ IOPS



| | AWS-m5.large | CaaB | AZ-az.F2s | GCPgcp.MediumVM |
|-------------|--------------|--------|-----------|-----------------|
| 1K-WR-IOPS | 13828 | 17920 | 1118 | 12461 |
| 1K-R-IOPS | 18547 | 102006 | 1169 | 18464 |
| 10K-WR-IOPS | 13828 | 18231 | 1118 | 12461 |
| 10K-R-IOPS | 18547 | 101448 | 1169 | 18464 |



DISK PERFORMANCE RESULTS: LARGE VMS

IO BANDWIDTH



| | AWS-m5.large | CaaB | AZ-az.F4s | GCP-gcp.LargeVM |
|--------------|--------------|---------|-----------|-----------------|
| OS-GP2(Base) | 105.03 | 1448.54 | 66.64 | 113.32 |
| IO1K | 272 | 1588.38 | 121.44 | 246.68 |
| Ю10К | 334.97 | 1732.77 | 121.25 | 247.07 |





DISK WRITE/READ IOPS



| | AWS-m8.large | CaaB | AZ-az.F4s | GCP-gcp.LargeVM |
|----------------|--------------|---------|-----------|-----------------|
| OS-GP2-WR-IOPS | 2,813 | 18,214 | 518 | 1,485 |
| OS-GP2-R-IOPS | 4,175 | 163,037 | 289,046 | 749 |
| 1K-WR-IOPS | 13,828 | 18,368 | 1,114 | 14,865 |
| 1K-R-IOPS | 18,547 | 162,238 | 1,238 | 15,217 |
| 10K-WR-IOPS | 13,828 | 18,260 | 1,114 | 14,865 |
| 10K-R-IOPS | 18,547 | 160,981 | 1,238 | 15,217 |

DISK PERFORMANCE RESULTS: EXTRA LARGE VMS

IO BANDWIDTH





| | AWS-m5.xlarge | CaaB | AZ-az.D4s_v3 | GCP-gcp.XLargeVM |
|---------------|---------------|-------|--------------|------------------|
| OS-GP2 (Base) | 110 | 1,518 | 57.57 | 116 |
| IO1K | 278 | 1,592 | 64 | 245.44 |
| IO10K | 336 | 1,742 | 61 | 245 |





DISK WRITE/READ IOPS



| | AWS-m5.xlarge | CaaB | AZ-az.D4s_v3 | GCP-gcp.XLargeVM |
|----------------|---------------|---------|--------------|------------------|
| OS-GP2-WR-IOPS | 2,826 | 13,185 | 363 | 1,485 |
| OS-GP2-R-IOPS | 3,975 | 242,908 | 310,758 | 749 |
| 1K-WR-IOPS | 13,825 | 18,564 | 14,387 | 14,387 |
| 1K-R-IOPS | 18,539 | 161,750 | 15878 | 15,878 |
| 10K-WR-IOPS | 13,825 | 18,357 | 14,387 | 14387 |
| 10K-R-IOPS | 18,539 | 162,260 | 15,878 | 15,878 |



DISK PERFORMANCE RESULTS: EXTRA EXTRA LARGE VMS

IO BANDWIDTH



| | AWS- m5.xlarge | CaaB | AZ-az.F8s | GCP-gcp.XLargeVM |
|---------------|-------------------|-------|-----------|------------------|
| OS-GP2 (Base) | 147 | 1,550 | 76.62 | 115 |
| IO1K | 271 | 1,581 | 103 | 245.83 |
| IO10K | 337 | 1,719 | 45 | 246 |





DISK WRITE/READ IOPS



| | AWS-m5.xlarge | CaaB | AZ-az.F8s | GCP-gcp.XLargeVM |
|----------------|---------------|---------|-----------|------------------|
| OS-GP2-WR-IOPS | 2,808 | 19,247 | 375 | 1,484 |
| OS-GP2-R-IOPS | 4,205 | 201,191 | 328,061 | 749 |
| 1K-WR-IOPS | 14,149 | 17,928 | 1,109 | 14,149 |
| 1K-R-IOPS | 15,703 | 192,880 | 592386 | 15,703 |
| 10K-WR-IOPS | 14,149 | 18,108 | 1,109 | 14149 |
| 10K-R-IOPS | 15,703 | 191,551 | 592,386 | 15,703 |

DISK BENCHMARK CONCLUSIONS

In contrast to the other CSPs that require the user to provision bandwidth and IOPS per volume, CaaB places no limits on bandwidth or IOPS for its volumes. Thus all CaaB users can achieve optimal latency and performance for their workloads at no extra cost:

- IO bandwidth: For all VM categories, CaaB's IO bandwidth is far superior (higher) than the other CSPs.
- Read IOPS: In general CaaB's IOPS rate during read operations was far higher than the other CSPs, with the exception of the XXL category, where AWS' IOPS rate was higher. In addition it should be noted that AWS' IOPS read rate was generally higher than CaaB's for the base volume.
- Write IOPS: Although CaaB's IOPS rate during write operations was generally slower than for read operations, it out-performed all the other CSPs for all volume types and disk provisioning configurations.



NETWORK BENCHMARK

Network performance was assessed based on the following parameters:

- WAN Download and Upload speeds, in Mbits/second.
- LAN speed, Gbits/second
- Packet speed, in packets per second (pps)

WAN NETWORK BENCHMARK COMMAND-LINE

```
curl -s
```

```
https://raw.githubusercontent.com/sivel/speedtestcli/master/speedtest.py |
python - >> /tmp/results/out.speedtest 2>&1
```

LAN NETWORK BENCHMARK COMMAND-LINE

ncmeter \$2 256M >> /tmp/results/out.lan 2>&1

LAN NETWORK PPS BENCHMARK COMMAND-LINE

ping -c 20000 -q -s 1 -f \$2 >> /tmp/results/out.packets

NETWORK BENCHMARK RESULTS: MICRO VM

| | AWS-t2.micro | CaaB | AZ-az.B1s | GCP-gcp.MicroVM |
|------------------------|--------------|---------|-----------|-----------------|
| WAN-Download (Mbits/s) | 418 | 3,655 | 1725.24 | 442 |
| WAN-Upload (Mbits/s) | 728 | 1,274 | 139 | 131.44 |
| LAN (Gbits/s) | 1 | 9 | 0.969 | 2 |
| Packets (pps) | 63,314 | 519,379 | 218776 | 240,902 |











NETWORK BENCHMARK RESULTS: SMALL VM

| | AWS-t2.small | CaaB | AZ-az.B1ms | GCP-gcp.SmallVM |
|------------------------|--------------|---------|------------|-----------------|
| WAN-Download (Mbits/s) | 437 | 3,655 | 1725.24 | 442 |
| WAN-Upload (Mbits/s) | 157 | 1,274 | 139 | 131.44 |
| LAN (Gbits/s) | 1 | 9 0.963 | | 2 |
| Packets (pps) | 65,015 | 475,801 | 219396 | 233,601 |













NETWORK BENCHMARK RESULTS: MEDIUM VM

| | AWS-t2.large | CaaB | AZ-az.F2s | GCP-gcp.MediumVM |
|------------------------|--------------|---------|-----------|------------------|
| WAN-Download (Mbits/s) | 3,278 | 2,009 | 2649.08 | 845 |
| WAN-Upload (Mbits/s) | 1,346 | 1,716 | 469 | 609.89 |
| LAN (Gbits/s) | 10 | 9 | 1.43 | 4 |
| Packets (pps) | 289,408 | 546,627 | 201297 | 312,138 |











NETWORK BENCHMARK RESULTS: LARGE VM

| | AWS-m5.large CaaB | | AZ-az.D4s_v3 | GCP-gcp.LargeVM |
|------------------------|-------------------|---------|--------------|-----------------|
| WAN-Download (Mbits/s) | 3,200 | 3,928 | 800 | 500 |
| WAN-Upload (Mbits/s) | 1,200 | 1,557 | 450 | 100 |
| LAN (Gbits/s) | 10 | 9 | 1.91 | 8 |
| Packets (pps) | 276,841 | 615,219 | 227800 | 365,079 |











NETWORK BENCHMARK RESULTS: XL VM

| | AWS-m5.xlarge | CaaB | AZ-az.D4s_v3 | GCP-gcp.XLargeVM | |
|------------------------|---------------|---------|--------------|------------------|--|
| WAN-Download (Mbits/s) | 832 | 3,928 | 2372.2 | 423 | |
| WAN-Upload (Mbits/s) | 168 | 1,557 | 495 | 106.04 | |
| LAN (Gbits/s) | 10 | 9 | 1.91 | 8 | |
| Packets (pps) | 276,841 | 615,219 | 227800 | 365,079 | |











NETWORK BENCHMARK RESULTS: XXL VM

| | AWS-m5.2xlarge | CaaB | AZ-az.F8s | GCP-gcp.XLargeVM |
|------------------------|----------------|---------|-----------|------------------|
| WAN-Download (Mbits/s) | 4,552 | 3,632 | 2833.84 | 146 |
| WAN-Upload (Mbits/s) | 1,273 | 1,477 | 191 | 77.47 |
| LAN (Gbits/s) | 10 | 9 | 3.17 | 12 |
| Packets (pps) | 405,874 | 593,973 | 268913 | 372,653 |











NETWORK BENCHMARK CONCLUSIONS

CaaB's network performance was equivalent to or superior to the other CSPs across all the network parameters that have the greatest impact on performance. Although the AWS WAN download was faster for the Medium and XXL VMs, this is most likely due to differences in the regions on which each CSP was tested.



OVERALL BENCHMARK CONCLUSION

PRICE/PERFORMANCE COMPARISON

As can be seen in the table below, we list all VMs and sizes per provider, per hourly or monthly pricing.

laaS prices will vary widely for the same CSP depending on the instance category, network traffic volumes, data storage tiering, availability zone location(s), level of upfront payment, usage of reserved instances, spot pricing, and more.

For example, in the largest VM category (XXL), CaaB's price per hour is 60% less than the second lowest price (Azure). Similarly, CaaB's price per month in the XXL VM category is 45% less than the second lowest price (GCP). Pricing flexibility is also important. CaaB does allow month-by-month pricing, while some CSPs' monthly price requires a 12month commitment.

| VM Size | Price per Hour (Base traffic and IOPS) | | | Price per Month (base traffic and IOPS) | | | | |
|----------------------------|--|---------|---------|---|----------|----------|----------|----------|
| | AWS | СААВ | AZ | GCP | AWS | CAAB | AZ | GCP |
| Micro (1 vCPU / 1 GB RAM) | \$0.0134 | \$0.018 | \$0.013 | \$0.046 | \$9.78 | \$13.14 | \$9.49 | \$23.58 |
| Small (1 vCPU / 2 GB RAM) | \$0.0268 | \$0.026 | \$0.025 | \$0.052 | \$19.56 | \$18.98 | \$18.25 | \$26.37 |
| Medium (2 vCPU / 4 GB RAM) | \$0.115 | \$0.059 | \$0.103 | \$0.103 | \$83.95 | \$43.07 | \$75.19 | \$52.74 |
| Large (4 vCPU / 8 GB RAM) | \$0.115 | \$0.121 | \$0.204 | \$0.206 | \$83.95 | \$88.33 | \$148.92 | \$117.48 |
| XL (4 vCPU / 16 GB RAM) | \$0.23 | \$0.186 | \$0.24 | \$0.250 | \$167.90 | \$135.78 | \$175.20 | \$187.80 |
| XXL (8 vCPU / 16 GB RAM) | \$0.46 | \$0.255 | \$0.409 | \$0.413 | \$335.80 | \$186.15 | \$298.57 | \$270.96 |



In addition to price per hour or per month, most CSPs add charges per volume and IOPS provisioned, with clear ratios of maximum IOPS permitted per volume size. For example, AWS enforces an IOPS/volume size ratio of 50:1. Thus, if a workload requires 10000 IOPS, on AWS the smallest volume that could be provisioned would be 200GB. The following table compares prices across the four CSPs for both base volumes and volumes provisioned with 1K IOPS:

| | AWS | СААВ | AZ | GCP |
|------------------------------------|---------|-------------------------------------|---------|--------|
| Base volume | \$5.95 | Included in price per hour/month | \$3.01 | \$4.80 |
| Volume provisioned with 1K IOPS | \$92.90 | Included in price per hour/month | \$41.82 | \$6.94 |



Summary

The CaaB public cloud is comprised of 12 data centers hosted in top tier facilities around the globe including: the UK, Germany, The Netherlands, Israel, Hong Kong, Santa Clara, Texas and New York. With servers operating on VMWare, CaaB is a trusted provider to tens of thousands of customers worldwide.

As can be seen in this Benchmark Report, CaaB's infrastructure delivers highly competitive performance across all critical cloud parameters such as vCPU speeds, WAN and LAN network speeds, disk IO / bandwidth / IOPS, RAM speed, and more.

CaaB offers advanced cloud infrastructure services and products such as block storage, load balancer, firewall, VPN and more. CaaB's management console provides an intuitive single-pane control panel to manage, maintain and supervise all elements of the customer's infrastructure. CaaB can also provide fully managed cloud services such as cloud migration planning and execution, disaster recovery, web servers, and email servers.

Service and flexibility are core values at CaaB. Live technical support is available 24/7/365 and CaaB's cloud experts can help customers find the optimal solutions for their cloud storage and computing needs. Customers can pay either per hour or a fixed price on a month-by-month basis that includes unlimited bandwidth and IOPS.

CaaB works closely with affiliates and also makes its platform available to partners on a white-label basis through which they can provide a wide range of branded services including VM/laaS, Cloud LAN, DRP, CDN, WAF, and a full featured marketplace.

Visit our website at <u>https://www.caab.io/</u> to learn more about our customers and partner programs, or to begin your 30-day free trial.