Performance & Price/Performance Benchmark of IaaS Providers 2021

This report is based on the benchmark testing and report of independent analyst company Cloud Mercato.

Introduction

Infrastructure as a Service (IaaS) is at the heartbeat of digitalization. It is the basis of business developing services for the future. Usually, IaaS is considered an easily available resource that is exchangeable to a great extent. But if you take a second look, there are not only differences in the public cloud providers' offers beyond IaaS like artificial intelligence, IoT etc., but as well on the very basic level of IaaS.

Constantly evolving new CPUs, network, storage technology and varying architectures lead to considerable differences in the performance of the offered services. And on that basis, users receive a rather different amount of services for their money. Only a price/performance analysis can reveal which hyperscalers deliver best "bang for the buck".

In 2021, T-Systems asked Cloud Mercato to perform such a benchmark to gain transparency about the performance and the price/performance of the leading hyperscalers Amazon Web Services (AWS), Microsoft Azure and Google Cloud. Cloud Mercato also compared the results with respective data from its analysis of the Open Telekom Cloud (OTC). The benchmark took place in 2021 referring to the prices of August 1st (see appendix for prices).

Like last year, it covers compute performance, RAM performance, storage IOPs and network bandwidth. 2021, two new categories made it to the benchmark: OpenSSL encryption performance and the performance of relational database offerings.

Product specifications

For the comparison, a test panel of mainly 16 vCPU machines was created containing different series categories such as general purpose, compute- or memory-optimized (the so called "basic test panel"). Additionally, AMD and Intel processors were analyzed. In 2021, the report also features a broad variety of Hyperscaler VMs with capabilities for specific use cases. Additionally, the report highlights RDS and encryption speed.

AWS is present with 12 VMs. c5, c5a and c5n classes are designed for compute-intensive workloads delivering cost-effective high performance. m5, m5a, m5n, m6i and m5zn (with 12 vCPUs) offer balance of compute, memory, and networking resources for a broad range of workloads. r5, r5a, r5n and r5b target to accelerate performance for workloads that process large data sets in memory. AWS VMs were used from the region EU (Frankfurt).

Azure VMs used for the benchmark were Standard F16s, D16s, E16s-4 v4 and as-v4 variants. Fs v2 features a sustained all-core turbo clock speed of 3.4 GHz and a maximum single-core turbo frequency of 3.7 GHz. D class VMs are a general-purpose type. They offer a combination of vCPU, memory and remote storage options for most production workloads. E class VMs are memory-optimized. They accelerate performance for workloads that process large data sets in memory. Azure VMs used were delivered from region West Europe.

Google Cloud VMs in the benchmark were N2, N2d, n2, n2d and c2. N types are general-purpose with balanced price/performance across a wide range of VM shapes while c2 offers ultra-high performance for compute-intensive workloads. Google resources were taken from region Netherlands.

Open Telekom Cloud appears with c4, m4 and s3 machines. c4 offers compute-optimized and dedicated CPU, m4 extended memory and dedicated CPU, s3 is considered a general-purpose machine. Open Telekom Cloud resources were provided from Magdeburg and Biere (Germany East).

Source: IaaS Benchmark study by Cloud Mercato 2021

The test panels in detail

Provider	VM	vCPU	RAM
Amazon Web Services	c5.4xlarge	16	32768
Amazon Web Services	m5.4xlarge	16	65536
Amazon Web Services	r5.4xlarge	16	131072
Google Cloud	Custom N2 16 vCPU 32GB	16	32768
Google Cloud	c2-standard-16	16	65536
Google Cloud	n2-highmem-16	16	131072
Microsoft Azure	Standard_F16s_v2	16	32768
Microsoft Azure	Standard_D16s_v4	16	65536
Microsoft Azure	Standard_E16s-4s_v4	16	131072
Open Telekom Cloud	c4.4xlarge.2	16	32768
Open Telekom Cloud	s3.4xlarge.4	16	65536
Open Telekom Cloud	m4.4xlarge.8	16	131072

Table 1: Basic test panel

Provider	VM	vCPU	RAM
Amazon Web Services	m5zn.3xlarge	12	49152
Amazon Web Services	m6i.4xlarge	16	8192
Amazon Web Services	c5a.4xlarge	16	32768
Amazon Web Services	c5n.4xlarge	16	43008
Amazon Web Services	m5a.4xlarge	16	65536
Amazon Web Services	m5n.4xlarge	16	65536
Amazon Web Services	r5b.4xlarge	16	131072
Amazon Web Services	r5a.4xlarge	16	131072
Amazon Web Services	r5n.4xlarge	16	131072
Google Cloud	Custom N2D 16 vCPU 32GB	16	32768
Google Cloud	n2d-standard-16	16	65536
Google Cloud	n2-standard-16	16	65536
Google Cloud	n2d-highmem-16	16	131072
Microsoft Azure	Standard_D16as_v4	16	65536
Microsoft Azure	Standard_E16as_v4	16	131072

Table 2: Test panel with additional VMs for specific use cases

Provider	VM
Amazon Web Services	mysql db.m5.4xlarge
Google Cloud	db-n1-standard-16 MySQL
Microsoft Azure	MySQL GeneralPurpose 4096GB max
Open Telekom Cloud	rds.mysql.c2.4xlarge 5.7

Table 3: RDS VMs tested

To run a VM, an allocation of block storage volumes to the VM is necessary. To apply a fair benchmark methodology, the necessary block storage was sized with a rule of thumb: 50 GB per CPU. This rule mimics a traditional scaling up with growing computing power. It allows to go through the different performance tiers applied to some storage classes.
The following block storage was used for the volumes attached to VMs, but as well for the specific block storage benchmark.

Provider	Storage
Amazon Web Services	General purpose SSD
Amazon Web Services	GP3 16000 IOPS
Google Cloud	Persistent Disk SSD
Microsoft Azure	Premium LRS
Open Telekom Cloud	Ultra-High I/O

Table 4: Block storage used

To minimize errors, duplicate VMs were deployed during testing.

Source: IaaS Benchmark study by Cloud Mercato 2021

Performance testing and data

CPU Performance

The VM performance is mainly qualified by computing power delivered by CPU model. This value derives from a lot of other characteristics bound to virtualization and CPU specifications such as frequency, built-in instructions and more. Facing a real-life workload in the cloud, the simple MIPS (million instructions per second) or FLOPS (floating-point operations per second) declared by CPU vendors aren't enough to determine if one machine performs better than another.

To collect a synthetic performance value, the benchmark suite Geekbench 5 is used. This software runs workloads with integer, floating point and cryptographic domains. The set of tests includes various kinds such as compression, machine learning or compute vision and each test is performed with single and multi-thread modes.

In this test, most flavors deliver a Geekbench performance value of about 8k for multi-threading (an average of 8,032 exactly omitting Azure's E16–4vs) (fig. 1). AWS's m6i is the exception delivering more than 10k – about 25 percent additional performance. Nonetheless, all the respective Open Telekom Cloud VMs outperform the hyperscaler VMs. With values of more than 12 k (exactly 12,681 on average) they even excel the leading m6i (fig. 2). The performance plus vs. the average value is about 58 percent.

Bars in the graphic show the single-thread (dark) and multi-thread scores (light). The higher the score, the better the performance.

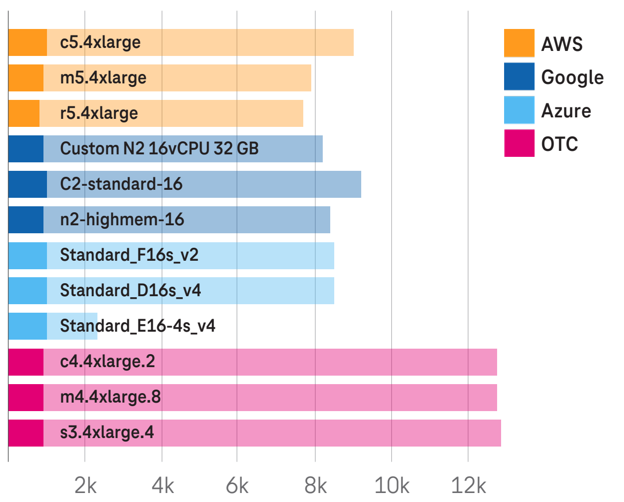

Fig. 1: Performance comparison of basic test panel (Geekbench Score)

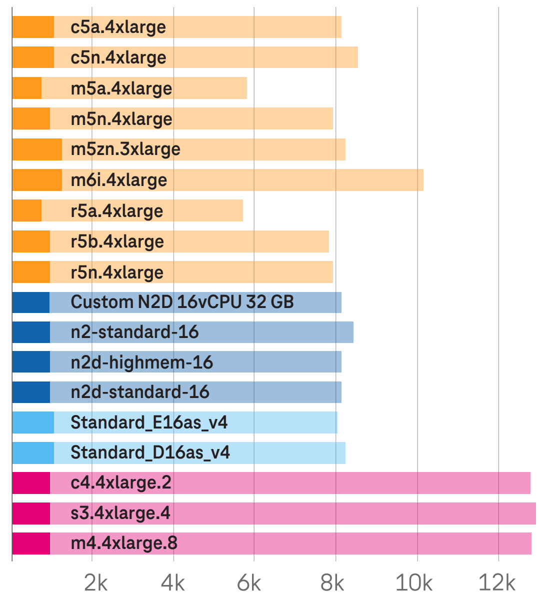

Fig. 2: Geekbench performance of additional VMs (Geekbench Score)

CPU Steal

CPU steal is real phenomenon that occurs from time to time in cloud environments, especially if VMs are not accurately separated. By nature, VMs mandatorily share resources with other tenants or at least with the hypervisor. That could allow CPU steal, but overcommitting could be an additional reason for CPU steal. As a rule of thumb: The older the hypervisor and the more occupied the data center, the higher the risk for CPU steal. CPU sharing can be measured by collecting a Linux kernel counter called "CPU steal".

Expressed in percentage, this number represents the amount of time that a task was not able to be done by CPU because of someone else's usage. When measuring CPU steal CPUs are loaded with Prime95. This software stresses all CPU up to 100 percent and allows capturing the amount of tick stolen during the load.

Results show that CPU steal is a very rare phenomenon that only occurs to a very minor extent at AWS. Open Telekom Cloud is no exception and thus sticks to the common cloud standards of minimum CPU steal.

Source: IaaS Benchmark study by Cloud Mercato 2021

RAM Performance

For some use cases the availability of efficient RAM is even more important than CPU power, e.g. for memory-intensive workloads caused by in-memory databases like SAP's HANA®. The volatile memory is the fastest one present on a system. To evaluate RAM performance the bandwidth between CPU and memory is tested using Sysbench. Its test scenarios are based on random access to memory with 1 k blocks.

For this test, read and write access have been used. By nature, write mode is slower as it suffers from latency occurred by data storing, whereas read only retrieves cached data. RAM performance is time related and may decrease while the hypervisor is filling.

AWS flavors show a broad range of write performance with c5 providing the best write performance of the test panel at all (fig. 3). Azure's e16–4s has the lowest score for writing with only about 40 percent of AWS's c5 flavor. AWS's c5 also provides a good read performance with more than 71 k units. Only Google's C2-standard and Open Telekom Cloud's m4 outstrip AWS's c5 with 72 k (Google) resp. 75 k (Open Telekom Cloud). None of the VMs of the additional set can reach that values (fig. 4).

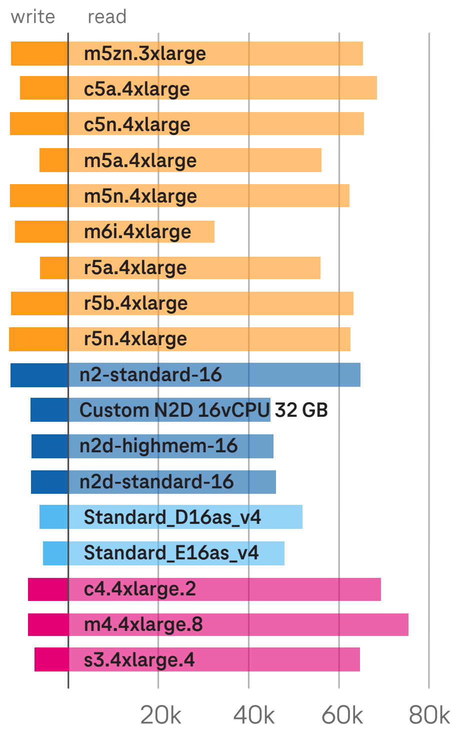

Fig. 4: RAM performance additional VMs (KB/sec)

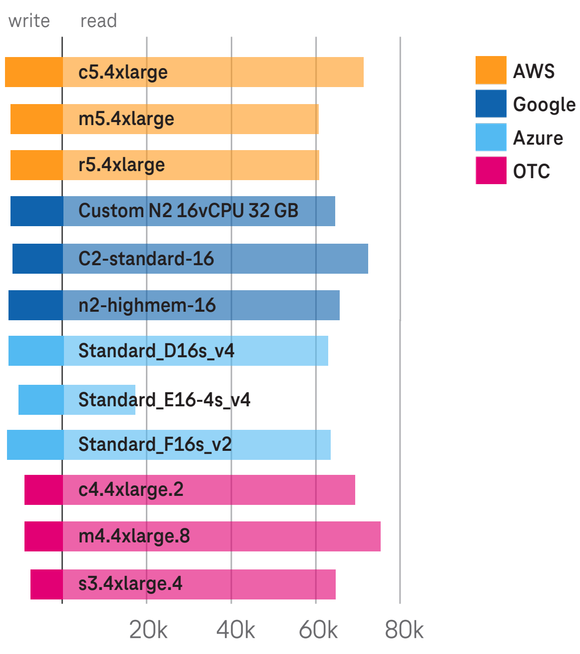

Fig. 3: RAM performance basic test panel (KB/sec)

Network performance

Cloud providers generally have a high-performance internal network throttled in consumer usage to guarantee a certain level of services for all tenants. The maximum performance is completely virtual and is defined by vendor in the VM's network specification. For the evaluation of the maximum bandwidth iPerf3 was used. In the testing setup 2 identical VMs from the same region and availability zone were used. These were loaded with a number of threads equal to CPU to generate the maximum throughput.

Cloud Mercato measured local bandwidth between two identical machines in the same datacenter. The tool used was Iperf 3 with TCP mode and a number of threads equal to CPU.

Google Cloud displays the best performance with more than 29 Gbps, but with a high variation, especially compared to the theoretical maximum bandwidth. AWS's n-series adds a real gain to network by more than doubling performance (from about 10 to more than 24 Gbps). Thus, AWS nearly reaches the standard level of the equivalent Open Telekom Cloud c4 flavor. Azure VMs showed the lowest values for network performance (fig. 5 and 6).

Source: IaaS Benchmark study by Cloud Mercato 2021

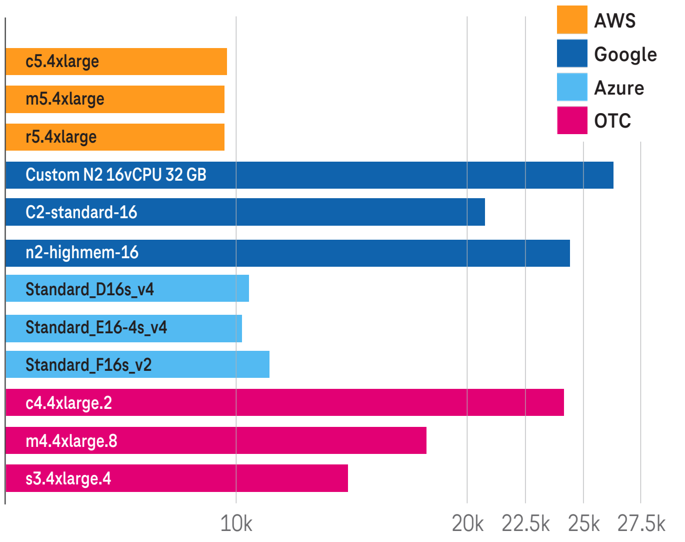

Fig. 5: Network performance for basic test panel (in Mbps)

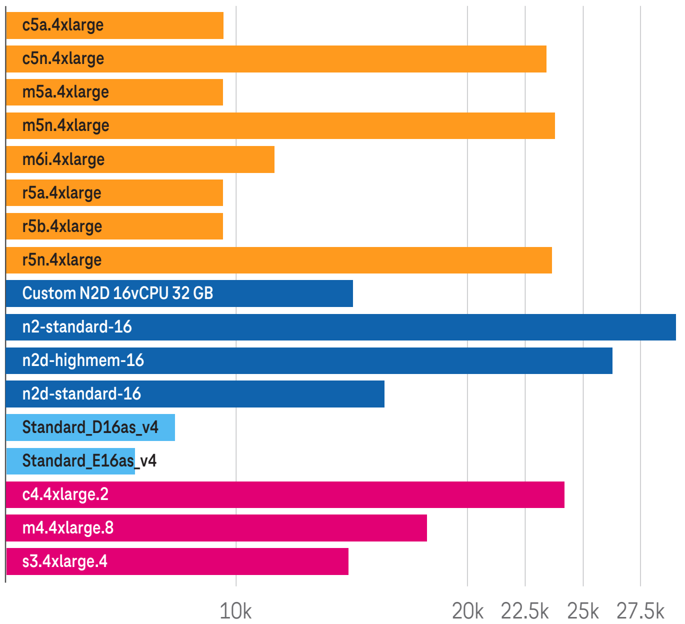

Fig. 6: Network performance for additional VMs (in Mbps)

Block storage

The persistent block storage is mainly qualified by its maximum IOPs (input/output per seconds). This value helps to determine the maximum number of transactions that the virtual device is able to provide and at this game, no storage system is equal to another. Like for network bandwidth, storage is a shared resource, that is throttled by vendors. Each provider has defined its own rules about IOPS and the performance factors are: definition of base rules via storage class, maximum performance via volume size and – in the case of Azure – total volume size (performance depends on size tiers calculated on total ownership).

I/O per seconds (IOPS) is one of the most used values to evaluate the performance of a volume, raw SSD, magnetic or block storage. The average and maximum IOPS gives a good preview of what will be the speed of a database or any disk-intensive task. To measure IOPS, Flexible I/O Tester (FIO) was used with the following configuration: 4 kB blocks, random access, read then write, direct access to device without filesystems, libaio engine.

Google stands out from the crowd with more than 50k IOPS. What's more, Google's write performance excels the read performance (!) by about 10k IOPS. Open Telekom Cloud is second with 20k IOPS (fig. 7).

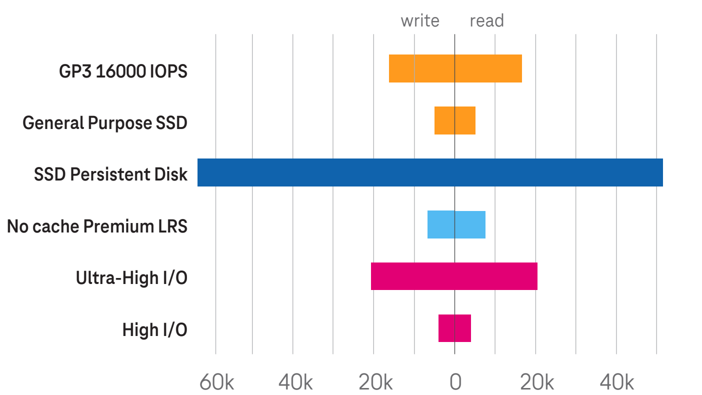

Fig. 7: Block storage performance (in IOPS)

Source: IaaS Benchmark study by Cloud Mercato 2021

OpenSSL encryption performance

With OpenSSL encryption a new test was added to the benchmark this year.

OpenSSL speed is a lightweight opensource tool provided with the OpenSSL command-line tool. It times the encryption/decryption through different algorithms and block sizes. Its output delivers a good picture of how much a system is able to perform with the different encryption systems available in the industry.

In this test, Open Telekom Cloud is clearly ahead of the three hyperscalers in both, aes-256 cnc and sha512 encryption. It delivers about 60 percent additional performance – independent of algorithm and block size. Figure 8 shows values for aes-256 cbc encryption with block size 16.384 Byte.

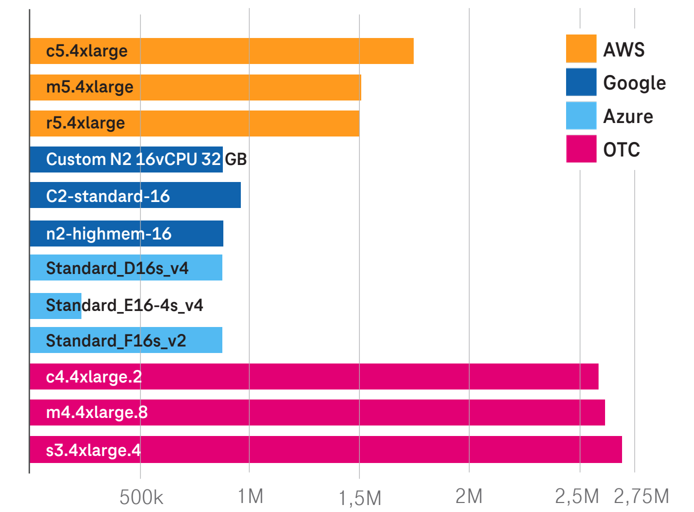

Fig. 8: OpenSSL encryption/decryption speed (in Bytes/s)

As the hyperscaler VMs of the basic panel do not provide great performance for encryption, users might choose flavors from the specific VMs panel (s. fig. 9). Compared to the basic VMs, AMD VMs provide a largely increased encryption capability – at least for a block size of 16.384 Byte.

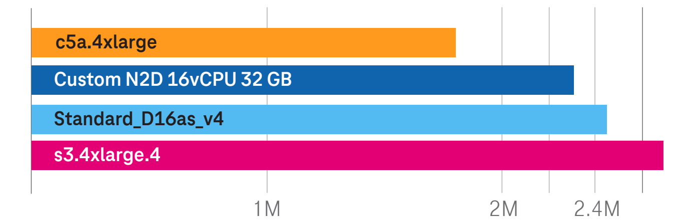

Fig. 9: OpenSSL encryption/decryption speed – best VMs compared (in Bytes/s)

Using a bigger block size (65.535 Byte) for aes-256 cbc encryption or using an alternative encryption method like sha512 doesn't change the picture: Open Telekom Cloud's VMs stay on top. However, there is one interesting finding when using aes-256 with bigger block size. Five VMs – all of them delivering improved capabilities for aes-256 cbc encryption with 16.384 block size – show considerably lower performance: All AMD-Rome-based Google VMs and Azure's Standard D16as and E16as. Columns in fig. 10 display the testing data. Grey columns show the performance decrease from about 2.3 resp. 2.4 million to values of around 900 k encrypted bytes per second at the respective flavors. Open Telekom Cloud's c4 VM, on the other hand, serves as an example for stable encryption speed with bigger block size.

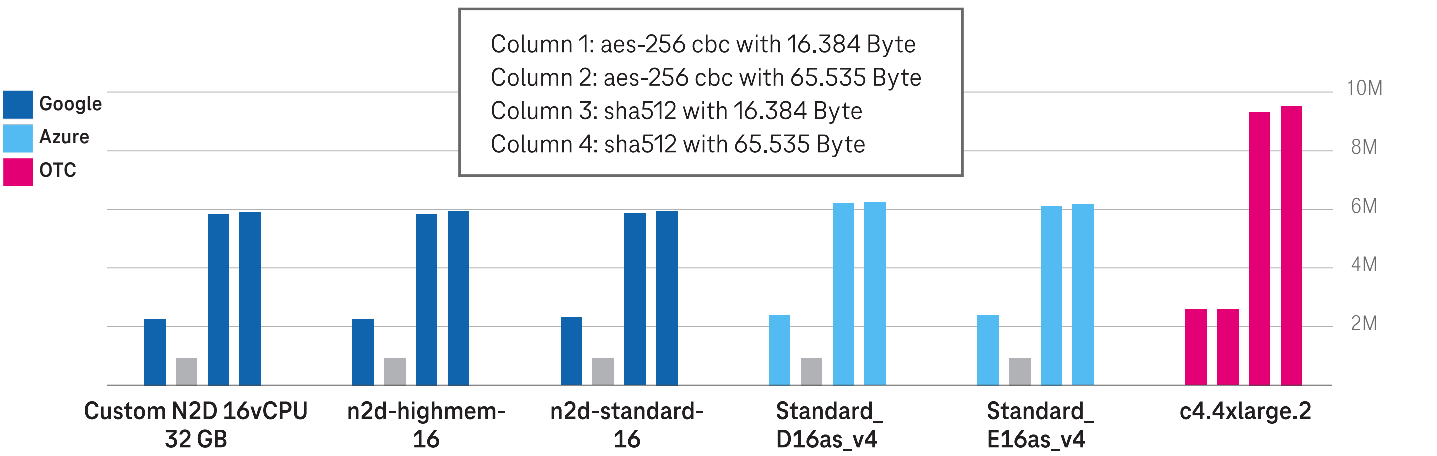

Fig. 10: Block size change – some flavors show reduced encryption performance (grey) (in Bytes/s)

Source: IaaS Benchmark study by Cloud Mercato 2021

Performance of relational database service

Cloud Mercato tested the performance of managed database services with sysbench OLTP, a benchmark tool evaluating the capacity of relational database such as MySQL and PostgreSQL. In the benchmark database types with a single server and 16 vCPUs were selected. They were hit with 32 vCPUs virtual machines in the same region/availability zone.

No special configuration has been injected to the RDBMS. To match with different usages and scenarios, Cloud Mercato scaled up from 1 client to 128 using read-only, write-only and read-write modes bringing additional workload to the database. The test showed that Google has a good ability to scale up write requests. Unfortunately, a bottleneck seems to collapse performance at high level (fig. 11). Open Telekom Cloud shows the best read scaling especially with high client numbers – followed by AWS (fig. 12).

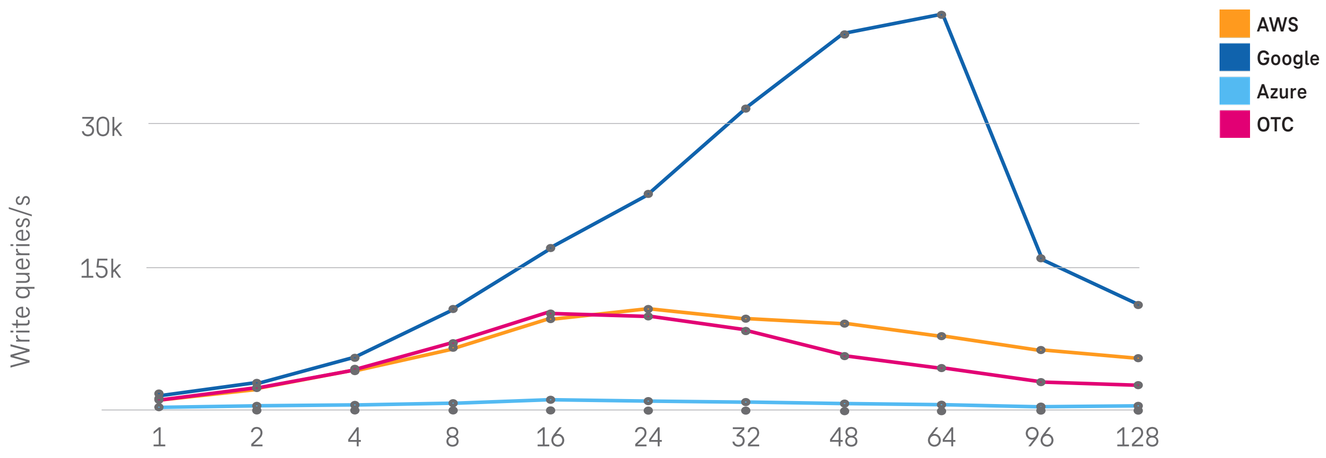

Fig. 11: Write performance while scaling up clients (databases see table 3)

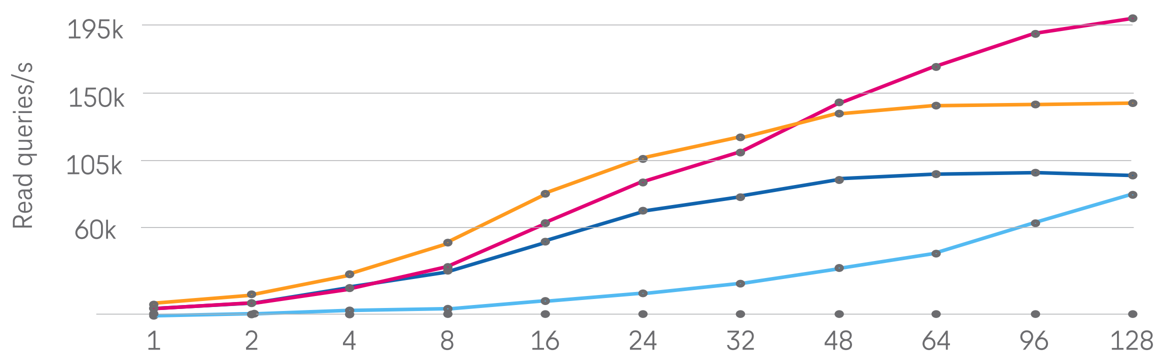

Fig. 12: Read performance of databases while scaling up clients

Source: IaaS Benchmark study by Cloud Mercato 2021

Price/performance value

Compute

A user can quickly lose track of which choice to make among the many cloud provider offerings. The real value of an offer is not apparent at first glance: Performance is one side of the coin, the other side is price. To evaluate the real value a user gains from a cloud provider, a consolidated view containing performance and price is the best approach. Price/performance offers a universal metric for comparing service value. This section shows the price/performance calculation based on the prices listed and the results from performance tests. As Open Telekom Cloud is billed in €, the following exchange rate has been used: 1 US$ = 0.858 € (see appendix for prices).

Calculation of price/performance values

For Geekbench and Sysbench OLTP price/performance is simply calculated by dividing the determined performance value by the price. For IOPS the formula takes into account the read and write performance (value normalization to hours): (write performance x 2 + read performance)/(price x 730) x 1000

Fair and comparable benchmark methodology

In this benchmark, the basic panel of best-comparable VMs of the leading hyperscalers and the Open Telekom Cloud are used. The flavors are similar in CPUs, network and storage technology.

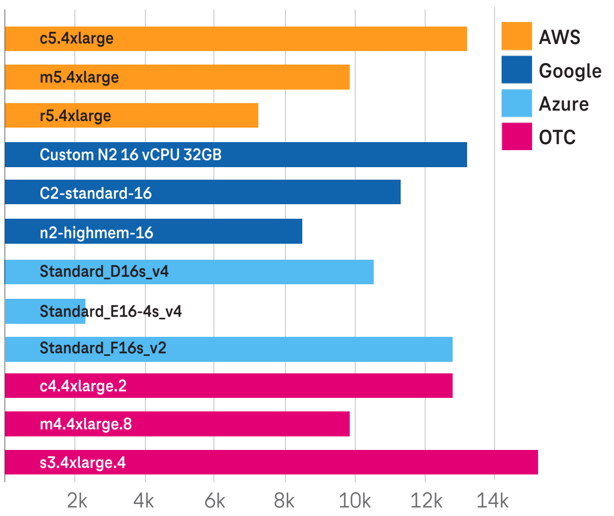

Fig. 13: Price/Performance comparison of basic panel

Open Telekom Cloud's s3.4xlarge offers the best price/performance in the test (fig. 14). Compared to the next best offer from a hyperscaler, AWS's c5a.4xlarge, s3 provides 15 percent of additional price/performance.

When using AWS, the basic c5.4xlarge and c5a.4xlarge from the additional VM set provide best price/performance. The value is on par with Google Cloud's Custom N2D from the additional VM set and the Custom N2 from the basic set. Open Telekom Cloud's c4 is following with about four percent less price/performance.

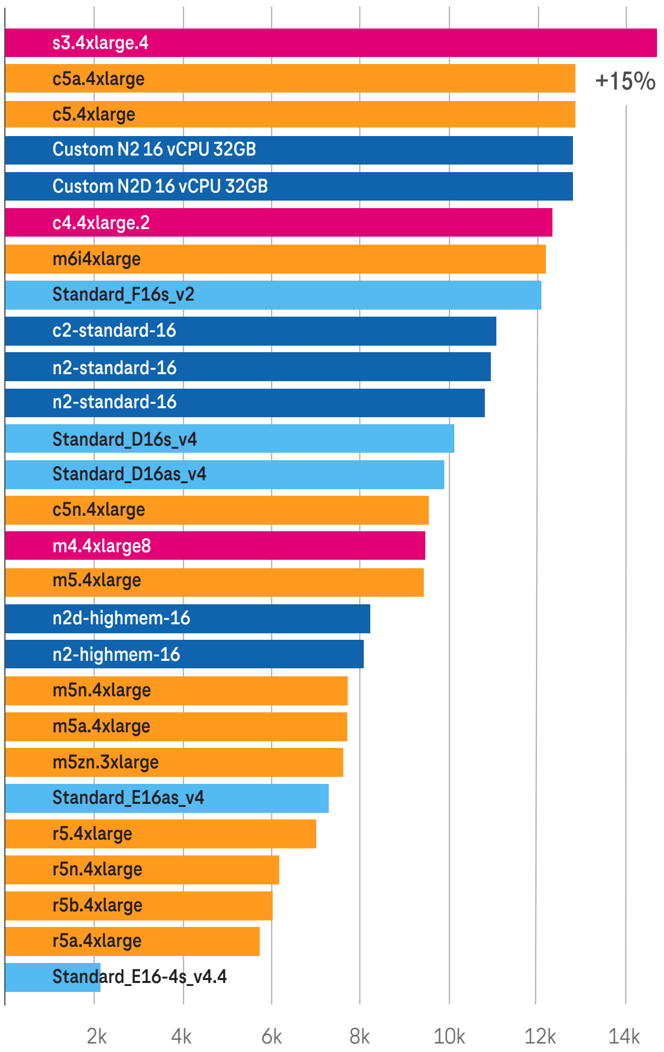

Fig. 14: Price/Performance comparison: all VMs at a glance

Source: IaaS Benchmark study by Cloud Mercato 2021

A look at the rest of the benchmark's test field reveals: Although the variety of AWS VMs optimized for specific use cases is impressive, the balanced solutions of the Open Telekom Cloud are on par, and in some cases even better: Their performance can also cover the specific use cases for which AWS flavors are designed.

Storage

Similarly to the price/performance calculation of the CPUs, the IOPS performance testing data was divided by the price. Long-term and short-term usage don't change the results. The formula to calculate the storage price/performance puts an emphasis on writing, which is valued double. Google's extraordinary performance makes it the best choice also in price/performance. Another interesting finding is that AWS's new GP3 is far better than the General Purpose SSD (GP2), but it does not reach the price/performance of Open Telekom Cloud's Ultra-I/O ranking second (fig. 15).

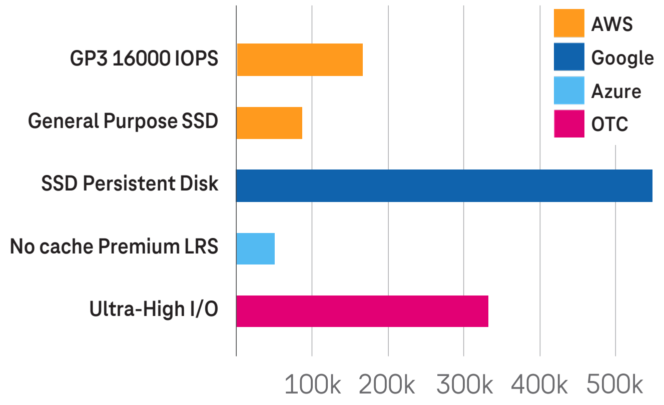

Fig. 15: Price/Performance comparison for block storage

Relational database service

Following the methodology, Cloud Mercato also calculated the price/performance values of the relational database offerings. While Google excels price/performance for writing, Open Telekom Cloud yields outstanding results for read operations. The more clients the database access, the better the price/performance. Second best choice for read is AWS with a rather stable value from 48 clients on.

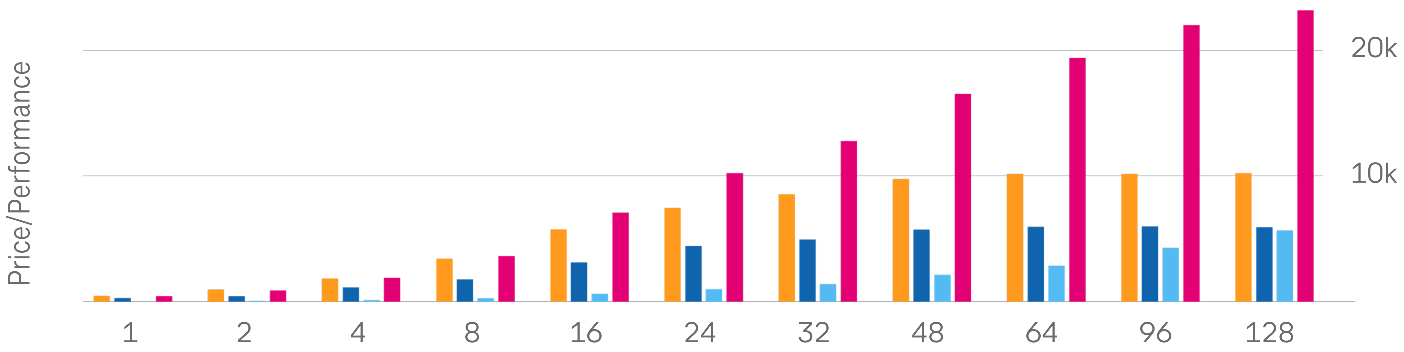

Fig. 16: Price/performance for read operations

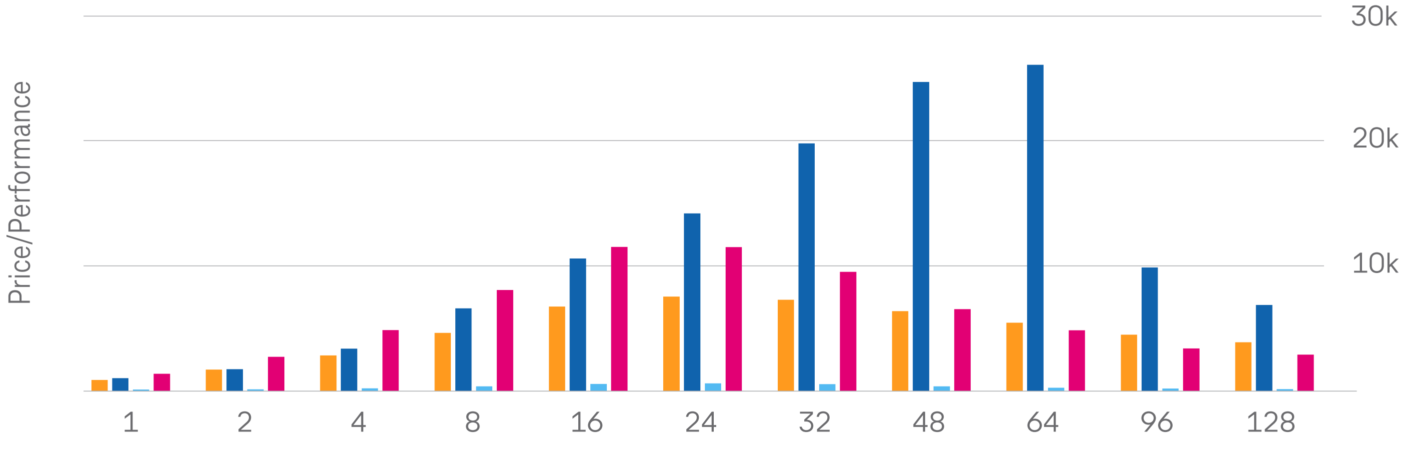

Fig. 17: Price/performance for write operations

Source: IaaS Benchmark study by Cloud Mercato 2021

Conclusion

Cloud usage is getting more and more elaborate and sophisticated as various flavors for different use cases prove. But these numerous specialized flavors and the evolving technology might cause confusion selecting the right VMs. Furthermore, governance and FinOps discussions around cloud usage arise.
One component of this discussion is also the question which providers offer best price/performance to sustainably limit infrastructure costs.

This year's benchmark from independent analyst company Cloud Mercato once again confirms: As in 2020, the offerings of the Open Telekom Cloud, which are well-balanced for all use cases, show at par performance and price/performance with the leading hyperscalers – and even exceed competitors.
This is especially valid in the area of encryption where hyperscalers only deliver good performance when using specific VMs (that still don't reach Open Telekom Cloud's performance).

Users don't need to worry about picking the right flavor for their use case – Open Telekom Cloud's s3 and c4 are a good and consistent choice for rather any case with s3 leading the field regarding price/performance with an additional 15 percent compared to the next best offer. Beyond that, Google Cloud delivers outstanding results for IOPS, IOPS price/performance and network bandwidth only followed by Open Telekom Cloud. With its GP3 storage and n class flavors AWS made good progress, but is still behind the performance of Open Telekom Cloud's Ultra-High I/O and s3 or c4 or m4 flavor.

Contact:
open-telekom-cloud.com/en/contact

Internet:
open-telekom-cloud.com/en

Publisher:
T-Systems International GmbH
Hahnstraße 43d
60528 Frankfurt am Main
Germany

Source: IaaS Benchmark study by Cloud Mercato 2021

Appendix – Prices

Provider	VM	Price (in US$)
Amazon Web Services	c5.4xlarge	0.776
Amazon Web Services	m5.4xlarge	0.920
Amazon Web Services	r5.4xlarge	1.216
Google Cloud	Custom N2 16 vCPU 32GB	0.706
Google Cloud	c2-standard-16	0.919
Google Cloud	n2-highmem-16	1.154
Microsoft Azure	Standard_F16s_v2	0.776
Microsoft Azure	Standard_D16s_v4	0.920
Microsoft Azure	Standard_E16s-4s_v4	1.216
Open Telekom Cloud	c4.4xlarge.2	1.153
Open Telekom Cloud	s3.4xlarge.4	0.927
Open Telekom Cloud	m4.4xlarge.8	1.492

Table 5: Prices for basic test panel

Provider	VM	Price (in US$)
Amazon Web Services	m5zn.3xlarge	1.187
Amazon Web Services	m6i.4xlarge	0.920
Amazon Web Services	c5a.4xlarge	0.696
Amazon Web Services	c5n.4xlarge	0.984
Amazon Web Services	m5a.4xlarge	0.832
Amazon Web Services	m5n.4xlarge	1.128
Amazon Web Services	r5b.4xlarge	1.424
Amazon Web Services	r5a.4xlarge	1.096
Amazon Web Services	r5n.4xlarge	1.424
Google Cloud	Custom N2D 16 vCPU 32GB	0.693
Google Cloud	n2d-standard-16	0.817
Google Cloud	n2-standard-16	0.855
Google Cloud	n2d-highmem-16	1.076
Microsoft Azure	Standard_D16as_v4	0.920
Microsoft Azure	Standard_E16as_v4	1.216

Table 6: Prices for test panel with additional VMs for specific use cases

Provider	VM	Hourly Price (in US$)
Amazon Web Services	GP3 16000 IOPS	0.09520
Amazon Web Services	General Purpose SSD	0.11676
Google Cloud	SSD Persistent Disk	0.18700
Microsoft Azure	No cache Premium LRS	0.29200
Open Telekom Cloud	Ultra-High I/O	0.11655

Table 7: Prices for Storage

Provider	VM	Hourly Price (in US$)
Amazon Web Services	mysql db.m5.4xlarge	1.62
Google Cloud	db-n1-standard-16 MySQL	1.85
Microsoft Azure	MySQL GeneralPurpose 4096GB max	3.33
Open Telekom Cloud	rds.mysql.c2.4xlarge 5.7	0.98

Table 8: Prices for RDS

Note: This report is based on the benchmark testing and report of the independent analyst company Cloud Mercato. T-Systems commissioned Cloud Mercato to do this report.