Cloud Providers in the Netherlands

Price and Price/Performance Benchmark

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

Source – IaaS NL 2022 Study by Cloud Mercato

September 2022

Introduction

The Cloud is gaining more and more traction as numerous European enterprises set sail for digitalization. Companies excited about the future and its possibilities seek to find the best providers for their cloud usage – particularly as the cloud will soon be the standard infrastructure for future business. Cloud discussions often center on the potential for business innovation and risk-minimized business exploration. But anybody who works in the Cloud will soon discover (or already knows) that the cost aspect is relevant too. You might argue that, in this context, all clouds are the same.

Are they really?

This Cloud Mercato benchmark will shed some light on this prejudice, comparing hyperscaler cloud services for the Dutch market.

Amsterdam is one of the cloud centers of the world. Its position as one of the most important Internet hubs makes it a preferred location for providing cloud services for Europe and the Netherlands benefit from this. Google Cloud and Microsoft Azure Cloud are present here, as is the Open Telekom Cloud. AWS is also providing cloud services for the Dutch market from Frankfurt and Paris.

In the second quarter of 2022, T-Systems commissioned independent cloud analysts from Cloud Mercato to run a benchmark on the performance of hyperscalers in the Dutch market. Based on the list prices and the performance results achieved, the analysts calculated the price/performance for the offerings. The price/performance ratio is a standardized approach for comparing the services and finding an answer to one of the most relevant cloud questions:

What is the Cloud's value for money?
How much do users get for their euro?

It is often debated that cheaper hardware (e.g. for CPUs) influences the price and thus the price/performance of cloud services. Hyperscalers therefore offer a variety of CPUs for different purposes, e.g. for price-sensitive customers. They add hardware vendors like AMD, ARM or Neoverse to their CPU pool.

This benchmark focuses on 16 vCPU virtual machines (flavors). The analysts designed three test panels, one for general purpose machines (with a CPU/GB RAM ratio of 1:4, i.e. 64 GB), one for compute-optimized machines (1:2 – 32 GB) and one for memory-optimized machines (1:8 – 128 GB). For this report, Cloud Mercato categorized the flavors as per Open Telekom Cloud's offering and similar specifications, though. Categorization and nomenclature may differ from provider to provider. These machines are used for different workloads and use cases. They also reflect different CPU vendors (Intel, AMD) in order to provide insights into the influence of different CPUs on performance and price/performance. An additional panel consists of 8vCPU machines from Intel and AMD.

AWS and Google Cloud have eight VMs, and Azure and Open Telekom Cloud six each. While the international hyperscalers offer Intel and AMD processors, Open Telekom Cloud only uses Intel processors. Open Telekom Cloud has c4, m4, and s3 machines. c4 offers a compute-optimized and dedicated CPU, while m4 offers extended memory and dedicated CPU, and s3 is considered a general-purpose machine.

The following machines were chosen for the respective panels:

GENERAL PURPOSE - 16 VCPU, 64 GB RAM

AWS (Paris)	m5a.4xlarge (AMD) m6i.4xlarge (Intel)
Google Cloud (Amsterdam)	n2-standard-16 Intel Cascade Lake n2d-standard-16 AMD Rome
Azure (Amsterdam)	Standard_D16as_v5 (AMD) Standard_D16s_v5 (Intel)
Open Telekom Cloud (Amsterdam)	c4.4xlarge.4 (Intel) s3.4xlarge.4 (Intel)

COMPUTE-OPTIMIZED - 16 VCPU, 32 GB RAM

AWS (Paris)	c5a.4xlarge (AMD) c6i.4xlarge (Intel)
Google Cloud (Amsterdam)	Custom N2 16 vCPU Intel Cascade Lake Custom N2D 16 vCPU AMD Rome
Azure (Amsterdam)	Standard_F16s_v2 (Intel)
Open Telekom Cloud (Amsterdam)	c4.4xlarge.2 (Intel)

MEMORY-OPTIMIZED - 16 VCPU, 128 GB RAM

AWS (Paris)	r5a.4xlarge (AMD) r6i.4xlarge (Intel)
Google Cloud (Amsterdam)	n2-highmem-16 Intel Cascade Lake n2d-highmem-16 AMD Rome
Azure (Amsterdam)	Standard_E16as_v5 (AMD) Standard_E16s_v5 (Intel)
Open Telekom Cloud (Amsterdam)	m4.4xlarge.8 (Intel)

GENERAL PURPOSE - 8 VCPU, 32 GB RAM

AWS (Paris)	t3.2xlarge (Intel) t3a.2xlarge (AMD)
Google Cloud (Amsterdam)	Custom N2D 8 vCPU AMD Rome Custom N2 8 vCPU Intel Cascade Lake
Azure (Amsterdam)	Standard_A8_v2 (Intel)
Open Telekom Cloud (Amsterdam)	s3.2xlarge.2 (Intel) s3.2xlarge.4 (Intel)

Compute: Geekbench computing performance benchmark

Cloud Mercato used the established Geekbench performance test to evaluate the compute power of the machines offered. Geekbench 5 is a standard benchmark suite aiming to create a score based on different application performances. Cloud Mercato used its outputs to capture a single score (for CPU performance) and multi score (for whole system performance), the latter being the more realistic scenario for cloud usage. We will thus focus on the multi score results. For comparison, single score values are shown in the figures as gray bars.

General purpose machines

In the first test panel, Open Telekom Cloud's s3 clearly leads the field for the multi score. The s3 machine achieves a Geekbench score of more than 12,600. Azure's Standard_D16 (Intel processor) takes second place with about 10,400 while AWS's m6i is third with about 10,100. In other words: s3 delivers about 22 percent more computing power than the next best machine of Azure.

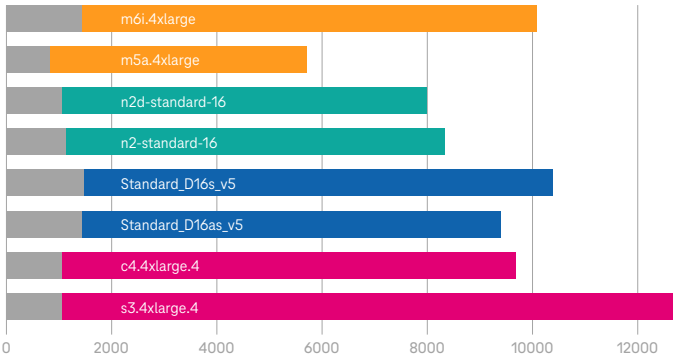

Fig 1: Geekbench performance results for general-purpose machines (single-score in grey, multi-score gray plus color) – the higher, the better

Throughout the panel, AMD's CPUs cannot keep up with Intel's performance, though the performance does differ greatly. For Google Cloud, the respective Geekbench score is only about 300 below the Intel score (8,003 vs. 8,323), whereas for Azure the difference is about 1,000 (9,394 vs. 10,384) and for AWS, AMD only provides about 57 percent of Intel's performance (5,701 vs. 10,071).

Compute-optimized machines

In the second test panel, Open Telekom Cloud once again delivers the best performance with its c4 machine. Its Geekbench score is 12,212. AWS's c6i takes second place with 10,162. With values of 8,000 to 8,500 the other machines are below the 10,000 mark. AMD-based machines again rank behind Intel-based ones. As previously seen, the difference in Google Cloud is only about 2.4 percent, while the difference in AWS is about 25 percent.

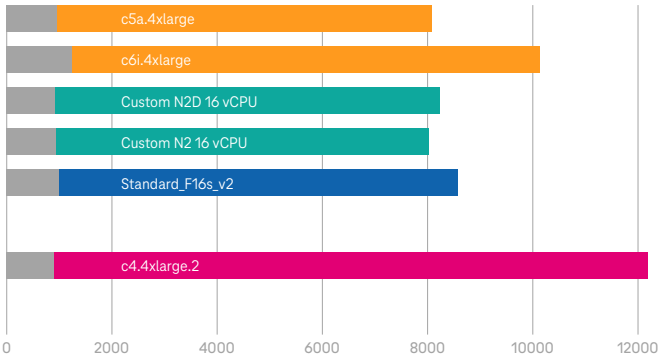

Fig 2: Geekbench performance results for compute-optimized machines (single-score in gray, multi-score gray plus color) – the higher, the better

Memory-optimized machines

The picture for memory-optimized machines is similar to the previous ones. Open Telekom Cloud's m4 reaches a value of 12,194 and leads the field. Azure's Standard E16 follows with 10,416, and AWS's Intel CPU (r6i) is slightly behind with 10,132. Azure's AMD machine performs far better than those of Google Cloud and AWS.

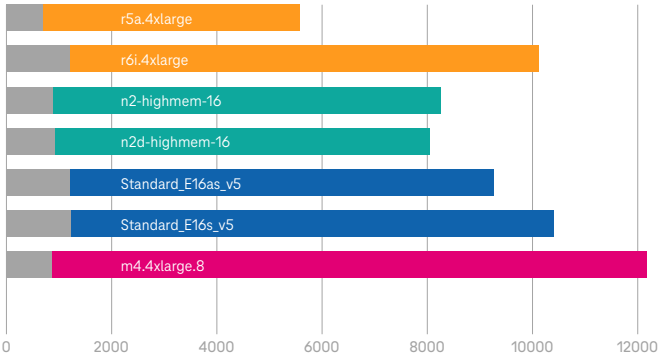

Fig 3: Geekbench performance results for memory-optimized machines (single-score in gray, multi-score gray plus color) – the higher, the better

8 vCPU machines

Due to machine availability, this panel contains 8 vCPU machines with 16 and with 32 GB RAM. AWS's t3 and t3a feature 32 GB RAM, while Azure's A8 and Google Cloud's N2 and N2D offer 16 GB. Open Telekom Cloud offers machines with 16 and 32 GB RAM (s3.2xlarge.2 resp. s3.2xlarge.4). Open Telekom Cloud's s3 takes first place, regardless of RAM. Both machines achieve similar values with 6,863 and 6,939.

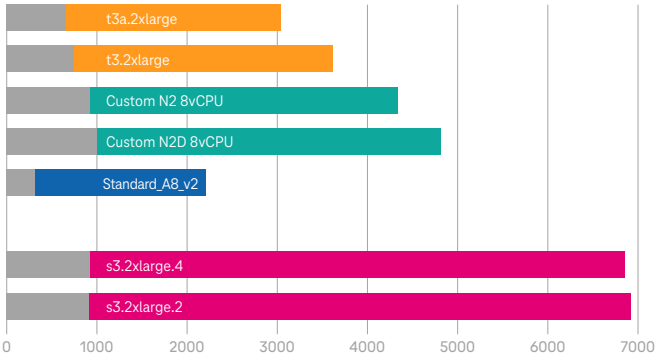

Fig 4: Geekbench performance results for 8vCPU machines (single-score in gray, multi-score gray plus color) – the higher, the better

Cloud Mercato assertion: Despite the AWS upgrades and the latest AMD EPYC and Intel Ice Lake, Open Telekom Cloud virtual machines still top the Geekbench ranking. Open Telekom Cloud's compute- and memory-optimized VMs clearly deliver better performance than their competitors.

A word about cloud steal

CPU steal is a real phenomenon that occasionally occurs in cloud environments, especially if VMs are not accurately separated. By nature, VMs have to share resources with other tenants or at least with the hypervisor. While this could make CPU steal possible, overcommitting could also be a factor. As a rule of thumb: The older the hypervisor and the more occupied the data center, the higher the risk of CPU steal. CPU sharing can be measured by collecting a Linux kernel counter called "CPU steal".

As part of the test, analysts only observed the phenomenon of "cloud steal" in one Cloud – AWS. This effect is neglectable in most machines. However, there is an exception. The aforementioned 8 vCPU machines t3 and t3a show an average "steal" of 0.3 and 2.5 per-cent respectively. Cloud steal may also influence machine performance.

Compute: Price comparison

In any case, it is to be expected that machines with good or great performance will be expensive. On the other hand, however, cheap machines may be the better choice when it comes to optimum value for money.

Let's first have a look at the prices collected on July 14th, 2022 (on demand usage model with hourly rates). Cloud Mercato used the Paris prices for AWS, and the Netherland prices for Open Telekom Cloud the prices in the Netherlands. Unsurprisingly, prices increase with the resources delivered – the more vCPUs, the more RAM, the higher the price. With identical resources (e.g. two machines with 16 vCPU and 32 GB RAM), the AMD processor offer is cheaper than the respective Intel-based offer.

The report uses euro as the currency (exchange rate USD 1 = EUR 0.98). Original benchmark values were calculated based on USD.

With an absolute price of € 0.670, Open Telekom Cloud's s3.4xlarge.4 offers the lowest price in the 16 vCPU panel, a general-purpose machine with 64 GB RAM. Google Cloud's compute-optimized N2D AMD with 16 GB RAM follows in second place (€ 0.679). Google Cloud takes third place (€ 0.692) with the respective Intel flavor. The most expensive machines are in the memory-optimized panel: Azure's Standard E16 (Intel) is € 1.192.

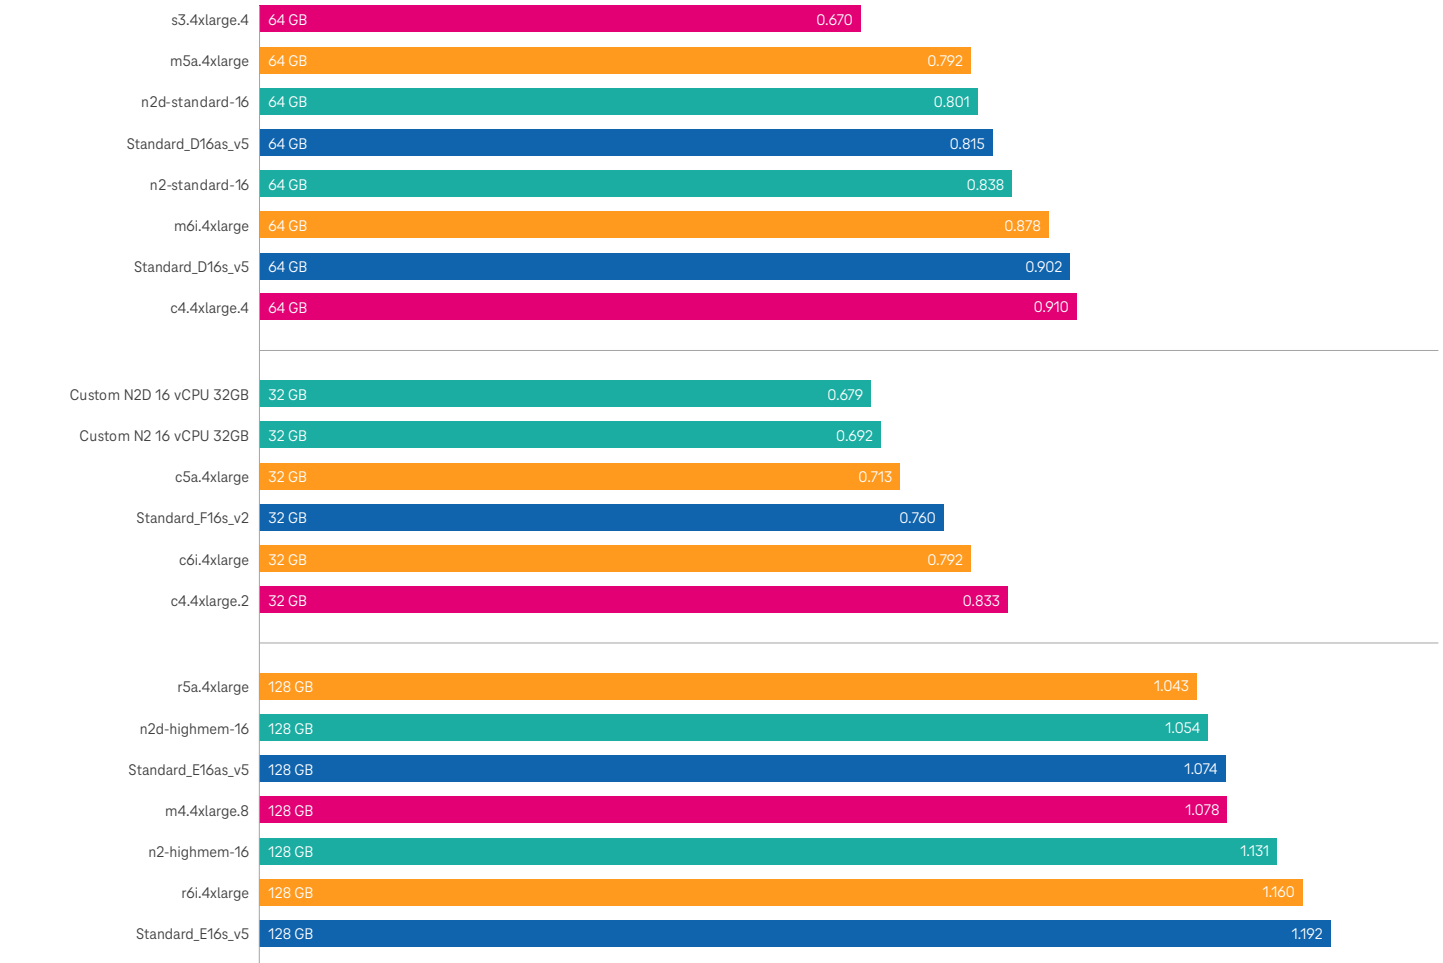

Fig. 5: 16vCPU machines – General purpose (64 GB), compute-optimized (32 GB) and memory-optimized (128 GB). Prices in €

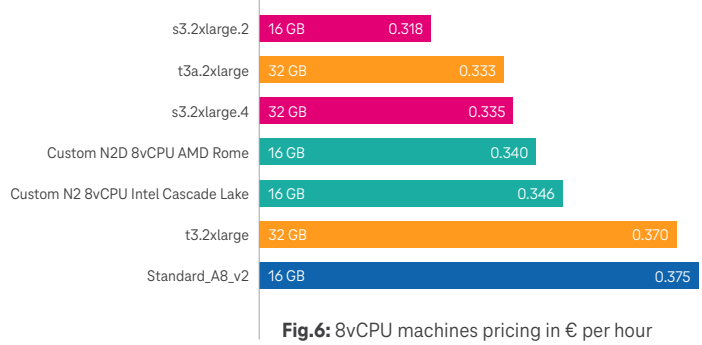

Fig.6: 8vCPU machines pricing in € per hour

Compute: Price/performance evaluation

Performance is one side of the coin, and price the other. Only by looking at everything as a whole can there be solid "value for money" discussions and appropriate comparison of offerings. Cloud Mercato thus calculated the price/performance values from the benchmark results and official prices.

$$\text{Price/Performance-ratio} = \frac{\text{Geekbench performance}}{\text{Price}}$$

Irrespective of the panel, an Open Telekom Cloud machine delivers the best price/performance. This is especially true for the general-purpose panel and the 8vCPUs. In the general-purpose field, s3.4xlarge.4 delivers more than 60 percent additional price/performance, while the 8 vCPU s3.2xlarge.2 offers more than 50 percent additional price/performance compared to the second best machine.

Comparing the best with the worst price/performance, the general-purpose panel shows a striking difference of 262 percent (7,349 vs. 18,906), and the 8 vCPU panel a difference of 370 percent (6,015 vs. 21,853), meaning that the value for money varies starkly.

The panels for memory and compute-optimized machines are far better balanced.

The results for Google Cloud and Azure provide one more interesting insight: From the perspective of price/performance, it doesn't matter if users choose an Intel or an AMD CPU; the price/performance values are nearly the same. E.g. in the general purpose panel, Standard D16as (AMD) yields a price/performance of 11,761, while the respective Intel machine D16s provides 11,756. Differences for Google Cloud are also minimal: 10,140 (Intel) vs. 10,203 (AMD), which is less than one percent.

In the case of AWS, the AMD machines price/performance is far less than that of the Intel processors. The latter provides about 60 per-cent more value for money for general-purpose and memory-optimized machines. For the other panel, the difference is 7 and 13 per-cent respectively. In all cases, Intel CPUs deliver a better price/performance with the exception of 8 vCPUs.

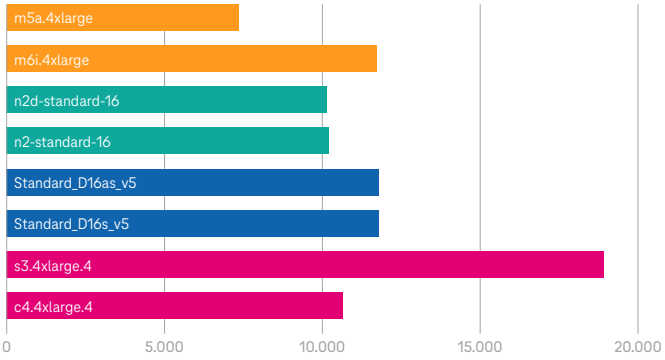

Fig. 7: Price/Performance for general-purpose (16 vCPU/64 GB RAM) panel (based on € prices) – the higher, the better

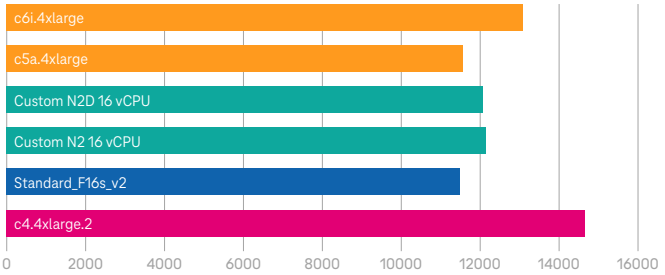

Fig. 8: Price/Performance for compute-optimized (16 vCPU, 32 GB RAM) panel (based on € prices) – the higher, the better

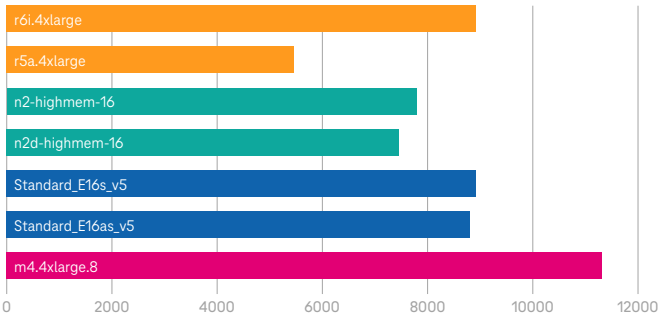

Fig. 9: Price/Performance for memory-optimized (16 vCPU, 128 GB RAM) panel (based on € prices) – the higher, the better

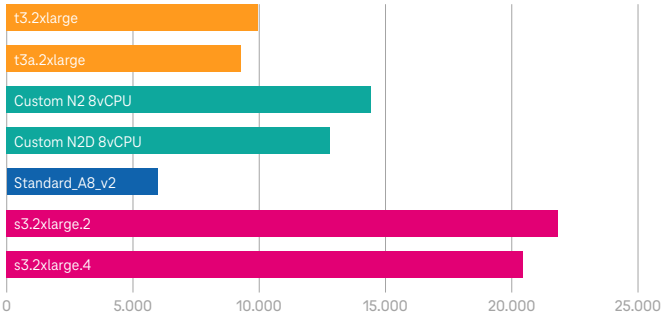

Fig. 10: Price/Performance for 8vCPU/32 GB RAM panel (based on € prices) – the higher, the better

Storage and storage bandwidth

Compute resources need storage resources to perform real-world processing tasks. Looking at storage performance is thus crucial to understanding overall system performance. Input/output per seconds (IOPS) is one of the most commonly used parameters for evaluating the performance of volume, raw SSD, magnetic or block storage. The average and maximum IOPS provides a good preview of the speed of a database or any disk-intensive task.

The persistent block storage is mainly qualified by its maximum IOPS. This value helps determine the maximum number of transactions the virtual device is able to provide, and no storage system is the same here. As was also the case for network bandwidth, storage is a shared resource that is throttled by vendors.

Each provider has defined its own rules for IOPS, and the performance factors are: Definition of base rules using storage category, maximum performance using volume size and sometimes total volume size (performance depends on size tiers calculated based on total ownership). 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, and number of jobs equal to CPU. Regionally balanced persistent disks were used in the test.

VMs used for the test:

AWS	c5a.4xlarge for GP3 16000 IOPS
AWS	c6i.4xlarge for General Purpose 2 SSD
Google Cloud	n2-standard-16 Intel Cascade Lake for Balanced Persistent Disk
Google Cloud	n2-standard-16 Intel Cascade Lake for SSD Persistent Disk
Azure	Standard_F16s_v2 for No cache Premium LRS
Open Telekom Cloud	s3.2xlarge.2 for Ultra-High I/O

The main findings: Google Cloud excels in storage performance (IOPS), while the observed AWS and Azure (no cache Premium LRS and General-Purpose SSD) offerings lag behind. However, the AWS GP3 16000 IOPS delivers as promised with a value of 16,529 IOPS in read and write. But this comes with additional costs: AWS charges users for a guaranteed IOPS. Even higher values are possible. Users should also keep in mind that Azure's IOPS scales with allocated storage. The greater the volumes, the higher the IOPS. Users wanting/needing higher IOPS will therefore need to purchase more storage.

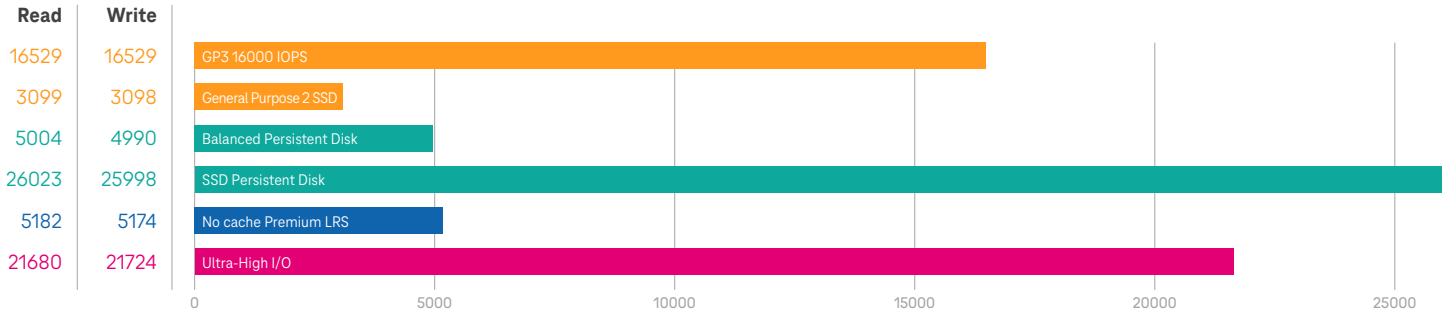

Fig. 11: Read vs. write values for block storage – the higher, the better the performance

Throughput bandwidth

Using FIO, Cloud Mercato also evaluated the maximum number of bytes transferrable by the respective instance (throughput). Findings are similar to the first evaluation: Google Cloud's SSD Persistent Disk achieves 399 MB/s, while Open Telekom Cloud follows with 360 MB/s. Surprisingly, the AWS General-Purpose SSD delivers about 260 MB/s as maximum throughput – putting it in third place. The GP3 ends up last with 130 MB/s.

Storage price comparison

Hyperscalers' prices for block storage are quite similar between € 0.836 and 1.074/h. Open Telekom Cloud charges € 0.365/h. The complete price comprises VM and attached volume. Figures 13 and 14 provide a breakdown of the total storage pricing in the test.

Fig. 12: Total Prices (VM plus storage) for storage offerings (in €)

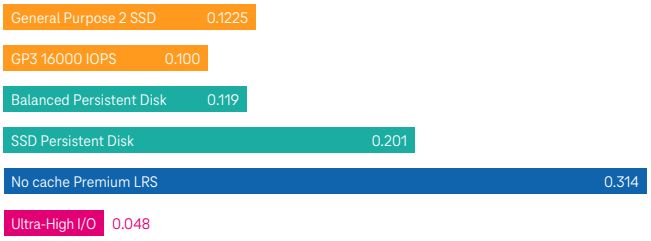

Fig. 13: Portion of pure storage to offering (in €)

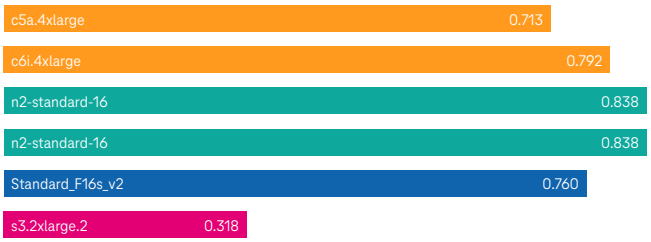

Fig. 14: Portion of VM to offering (in €)

Price/performance for storage

$$\frac{\text{Read value} * 2 + \text{write value}}{\text{Price} * 730} * 1{,}000{,}000$$

Excellent performance is essential for good price/performance. But the pure-performance picture changes when the costs for the block storage are taken into account. Low prices for Open Telekom Cloud NL usage result in an additional 130 percent value for money, which, in this case, reflects an unmatched combination of performance and price (see above for calculation of price/performance value). However, the price/performance approach doesn't change the complete picture: The basic offerings of AWS and Azure lag behind, and AWS's GP3 delivers a far better result than the General-Purpose 2 SSD.

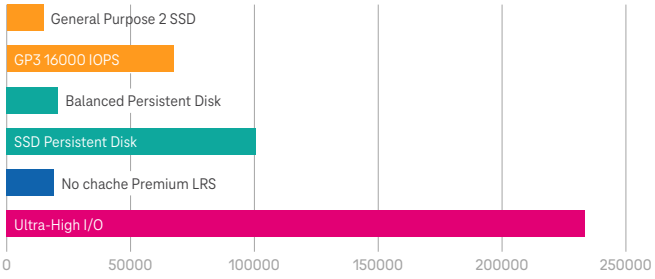

Fig. 15: Price/Performance comparison for storage – the higher, the better

Cloud Mercato assertions: Google PD-SSD and Open Telekom Cloud Ultra-High I/O lead this ranking with >20,000 IOPS. AWS doesn't reach a good level of performance without a monthly IOPS. Azure requires a large amount of storage to deliver fast access. With Google being among the most expensive, Open Telekom Cloud's performance make it the best bang for your buck.

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 the vendor as per the VM's network specification.

Two identical VMs from the same region and availability zone were used in the testing setup. 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 data center. The tool used was Iperf 3 with TCP mode and a number of threads equal to CPU.

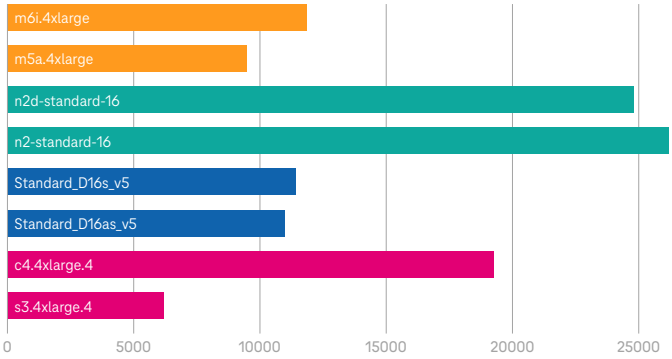

Fig. 16: iPerf for general purpose panel – the higher, the better

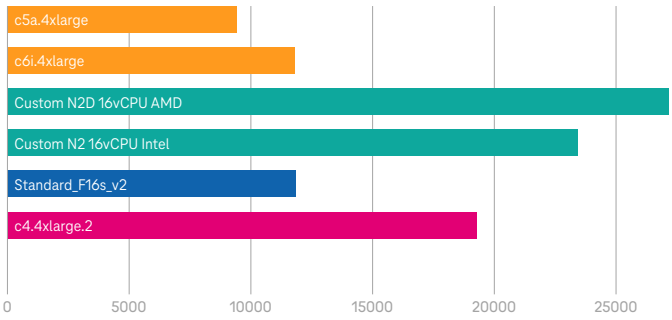

Fig. 17: iPerf for compute-optimized panel – the higher, the better

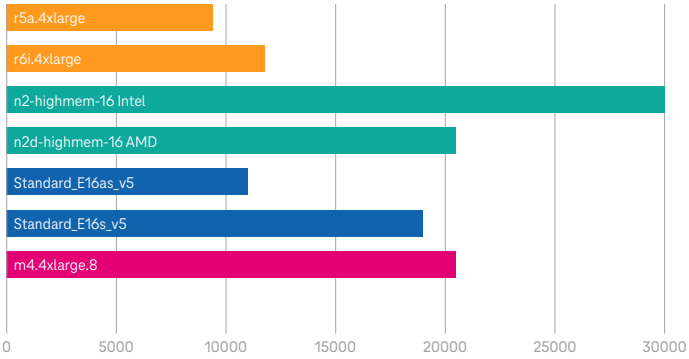

Fig. 18: iPerf for memory-optimized panel – the higher, the better

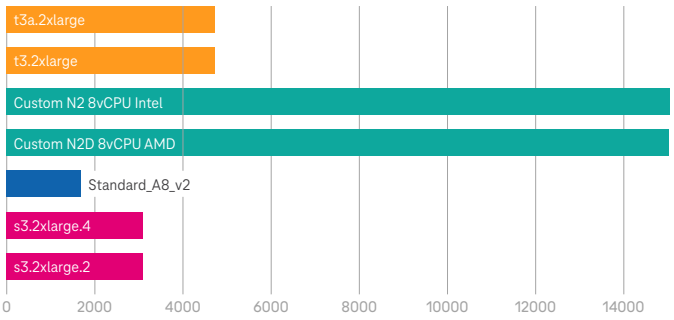

Fig. 19: iPerf for 8 vCPU panel – the higher, the better

Cloud Mercato assertion: Google Cloud provides the best internal network performance of 20 Mbps up to 30 Mbps. network bandwidth varies strongly depending on the different VMs. Open Telekom Cloud at least provides one VM in each panel that delivers around 20 Mbps.

Conclusion

It's only IaaS, right? The Cloud Mercato benchmark reveals the truth, with huge differences being observed in the panels with comparable VMs for general purpose, compute-optimized and memory-optimized workload. These differences cover performance and even more so price/performance. Be it raw compute or storage, the Open Telekom Cloud Netherlands is among the best performing and cost-effective offerings. And the flavors of Open Telekom Cloud deliver consistently reliable performance for any use case.

Customers do not need to search and pick the "right" specialized VM out of an abundance of offers; using a generic Open Telekom Cloud flavor, they can be sure to receive the best price/performance. These findings relate to similar results in the previous report covering the German market: Open Telekom Cloud provides a consistently good performance and price/performance ratio with its standard Intel machines across the regions served.

Cloud users should keep in mind that the choice of provider will have significant impact on the costs associated with operating workloads over the long term. This is essential for all computing requirements, including at a scientific level. Advantages in performance and price/performance will pay off substantially when handling huge and long-lasting workloads.

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