Why OVHcloud Might Be Your PostgreSQL Power Player

Contents

**This benchmark was performed
by BenchANT in August 2025**

PostgreSQL is available as a managed service from nearly every cloud provider, but performance and cost can differ significantly under the hood. In this study, we benchmarked OVHcloud's PostgreSQL DBaaS against AWS RDS PostgreSQL, Azure Flexible PostgreSQL, and DigitalOcean PostgreSQL using widely used OLTP workloads and a strict price-performance methodology. We compared throughput, latency, and efficiency on equivalent instance types to provide an objective view of how these services perform.

Key takeaways

· OVHcloud delivers competitive performance, in some scenarios
  even surpassing established players.

· Its price-performance ratio is strong, also compared
  to hyperscaler services.

For engineering and product teams, these results highlight the real-world impact of choosing one DBaaS provider over another—not just on performance but also on budget efficiency. Explore the full benchmark to see how OVHcloud positions itself in the 2025 PostgreSQL DBaaS market.

Introduction

This article compares the performance and cost performance of OVHcloud's PostgreSQL service against comparable services from other cloud providers using two complementary workloads. In detail, the competitors are the two U.S. hyperscalers Azure (Database for PostgreSQL Flexible Services) and AWS (RDS PostgreSQL) as well as the U.S. cloud provider Digital Ocean (PostgreSQL managed database).

Methodology

Table 1 shows the competitor sizing for all workloads used in this evaluation. The baseline configuration we use is OVH PostgreSQL b3-32. For the competitors, we apply set-ups which are similar in terms of software version, but also hardware configuration. More specifically, for all competitors, we use a two-node cluster configuration where each node runs at 8 cores and provides around 32 GB of RAM. The workloads are issued from a single virtual machine with 16 cores and 64GB of RAM that is deployed in the very same data center / region as the database service. In any case, we verified that the workload machine is not a bottleneck in any dimension.

Table 1: Competitor Sizing for this evaluation

DBaaS	Region	Version	Cluster Size	Instance Size	Hardware	Storage Size
OVH PostgreSQL V2	Paris	16.9	2	b3-32 production	8vCores 32 GB RAM	n/a
AWS RDS PostgreSQL	eu-central-1 (Frankfurt)	16.8	2	db.m7g. 2xlarge	8vCores Graviton 32 GB RAM	GP3 12,000 IOPS
Azure Database for PostgreSQL	northeurope	16	2	D8ds_v5 Intel	8vCores 32 GB RAM	P50 7,500 IOPS
Digital Ocean PostgreSQL	FRA (Frankfurt)	16	2	General Purpose dedicated	8vCores 32 GB RAM	n/a

The first workload uses the well-known Yahoo! Cloud Serving Benchmark (YCSB) as workload generator. In this case, we use the following parameters: (i) the data set size is set to around 100GB; (ii) we use 95% of read operations and 5% of update operations. (iii) we use 25 threads (the YCSB terminology for parallel users). The benchmark runtime is set to 60 minutes and we use a wait time of 15 minutes after the load phase, before starting the run phase.

The second workload we use is a TPC-C like workload issued using BenchBase's TPC-C implementation. In this case, we apply the following configuration parameters: (i) we set the scale factor to 150 which equals around 50GB of data; (ii) we use 75 terminals (the TPC-C terminology for parallel users); the overall runtime of the benchmark is 60 minutes. In addition, as with the first workload, we apply a 15 minute wait time between loading the database and starting the workload.

Cost Analysis

Figure 1 illustrates the monthly costs in € for the configurations shown in Table 1. For this analysis, all monthly costs are computed based on 730 hours usage per month. We do not take into account discounts for reserved capacities or long term contracts. We also do not consider futher operational costs that may occur, for instance, for backup, network gateways, and other types of services.

The plot differentiates between three cost categories: (i) instance costs represent the cost for the compute part of the database cluster. These may come with some default storage; (ii) the category "storage costs" denotes costs that occur when using the amount of storage shown in Table 1; (iii) the category "IOPS cost" captures costs for using a higher I/O bandwidth compared to the default.

Cost / Month [€]
less is better

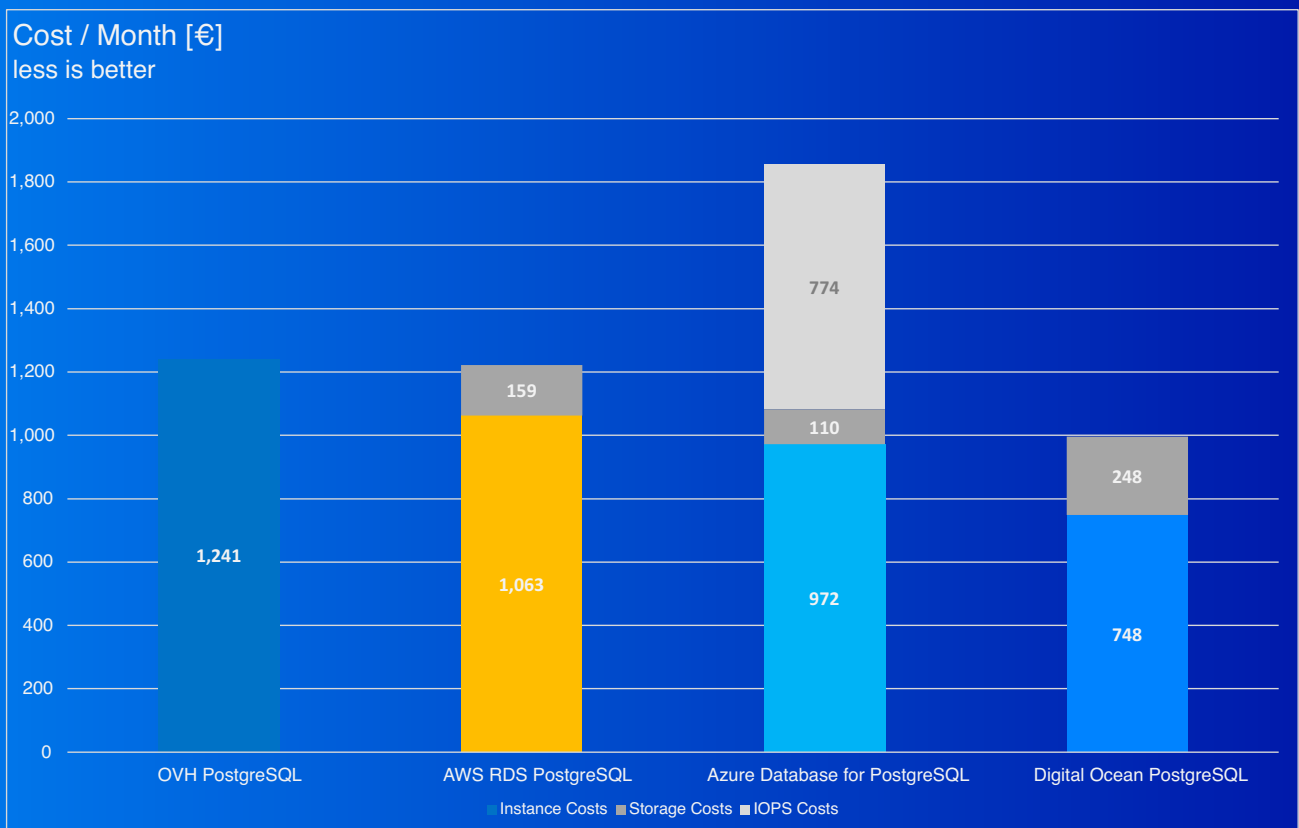|	OVH PostgreSQL	AWS RDS PostgreSQL	Azure Database for PostgreSQL	Digital Ocean PostgreSQL
IOPS Costs			774	
Storage Costs		159	110	248
Instance Costs	1,241	1,063	972	748

Instance Costs ■ Storage Costs ■ IOPS Costs

Figure 1. Cost comparison of all competitors including storage costs

OVH PostgreSQL comes with the largest baseline cost. Running the OVHcloud cluster from Table 1 costs 1,241 € per month. Yet, it is clearly visible that OVHcloud PostgreSQL is the only set-up amongst all competitors which does not cause any additional costs for storage or IOPS. Taking these additional costs into account, the monthly costs for AWS RDS PostgreSQL range at 1,222 € (1,063 € for instances and 159 € for storage), which is 1.5% cheaper than OVHcloud PostgreSQL. For Azure Database for PostgreSQL, the combined cost of instances (972 €) and storage (110 €) are 12.8% cheaper than for OVHcloud PostgreSQL. Yet, in the case of Azure an additional 774 € per month are required to achieve a comparable IOPS level as the competitors leading to a total monthly cost of 1,856€, which is almost 50% more expensive than OVHcloud PostgreSQL. The only offer that is truly cheaper than OVHcloud PostgreSQL is Digital Ocean PostgreSQL. Their instances come at a price of 748 €, while another 248 € are required for storage. Hence, the overall price of 996 € per month is almost 20% cheaper than OVHcloud PostgreSQL.

In the following section, we use this cost analysis to compute the performance-per-€ putting the costs into relation to the overall performance.

Results

The following sections discuss the results we obtained when applying both workloads to the cluster configurations shown in Table 1. For each of the workloads, we discuss the following three metrics: (i) throughput denotes the number of operations (or transactions) per second each database is able to achieve under the applied workload and the number of concurrent users. (ii) While throughput demonstrates the capability of a system to handle many requests in parallel, latency describes its ability to process individual requests as fast as possible. In performance engineering, averages of latencies are usually not considered meaningful as they are very vulnerable to outliers and skew. Hence, in this study, we focus on the 95th percentile of the latencies. The 95th percentile of a certain value means that 95% of all queries are not slower, i.e., have a lower or equal latency. (iii) finally, we use performance per € as the third metric. More precisely, we put the throughput and the costs from Figure 1 in relation to each other. This metric describes how many operations per second one can achieve for each € spent monthly.

Case-A:
Read-heavy YCSB Workload

In this section, we discuss the results for the YCSB workload. The workload as chosen is read-heavy but uses a 5% portion of update operations. That is, despite it having some modifying operations, the overall size of the data set remains widely identical over the benchmark runtime.
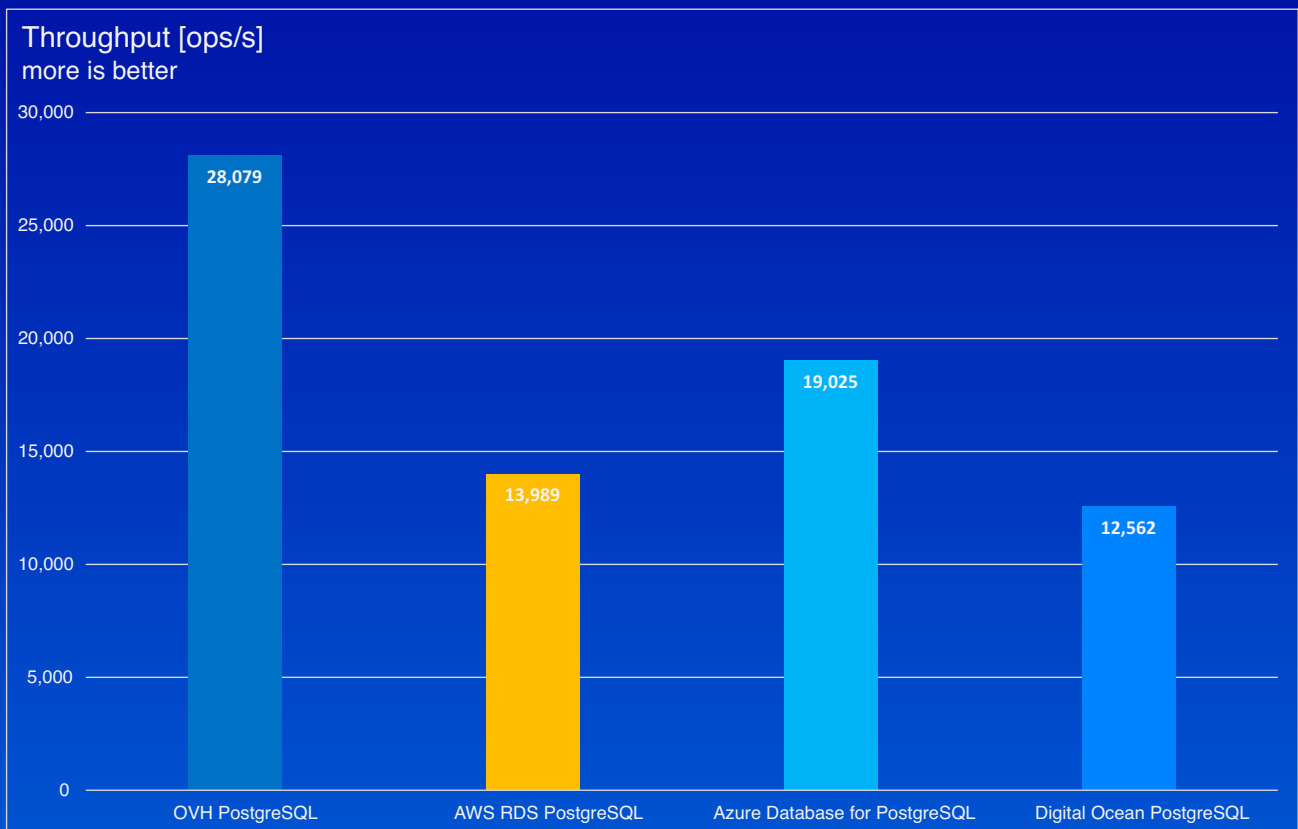

Figure 2. Throughput for Read-heavy YCSB Workload

The throughput shown in Figure 2 does not differentiate between the two operations. The ranking amongst the four competitors is clear: OVH PostgreSQL achieves more than 28,000 operations / s making it the clear winner for this workload. OVH PostgreSQL outperforms the second – Azure Database for PostgreSQL – by more than 9,000 operations per second, more than 30%. AWS RDS PostgreSQL (13,989 operations per second) and DigitalOcean PostgreSQL (12,562 operations per second) achieve less than 50% of OVH PostgreSQL's performance.
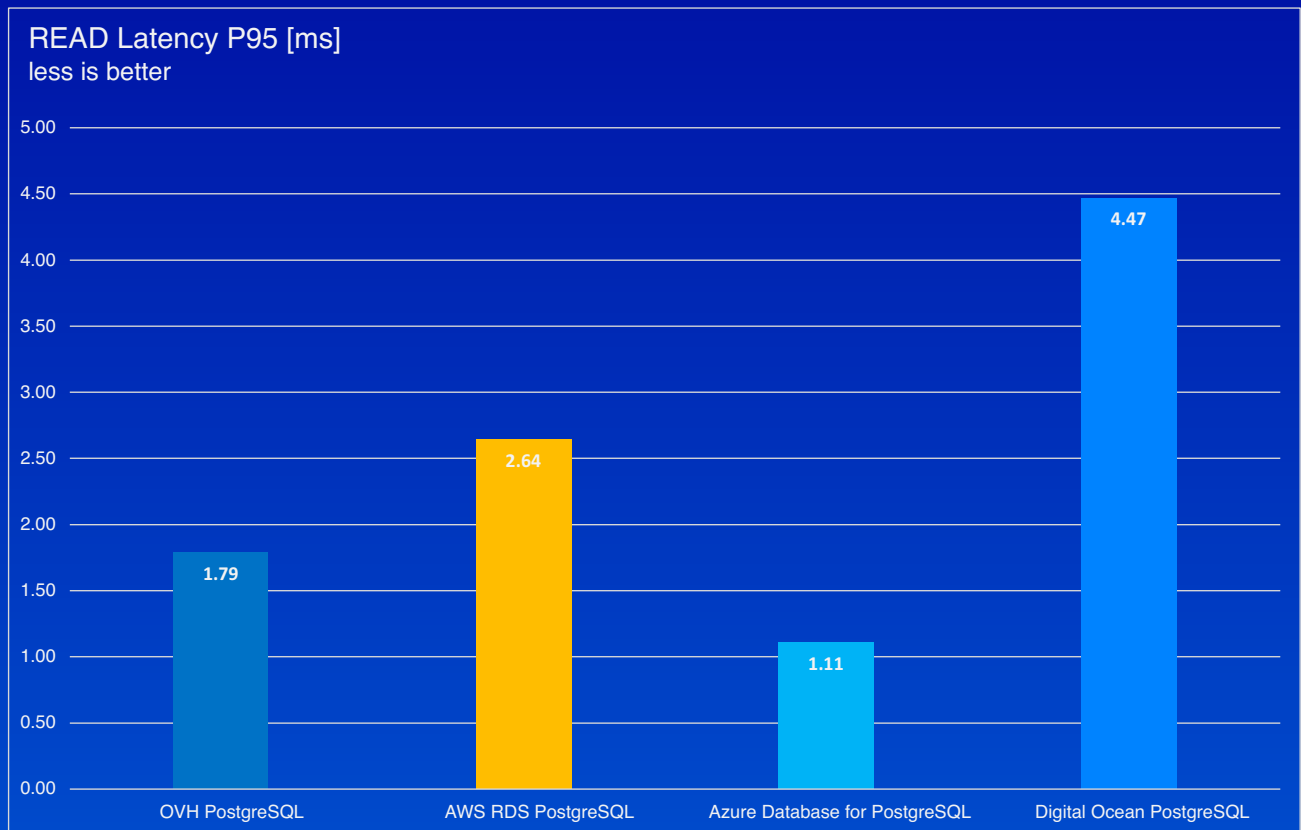
Case-A:
Read-heavy YCSB Workload

Figure 3. P95 of READ Latency for Read-heavy YCSB Workload

In contrast to throughput, we discuss the latencies of READ and UPATE operations separately. Their respective 95th percentiles are shown in Figure 3 and Figure 4 respectively. For the READ latencies depicted in Figure 3, we can state that Azure Database for PostgreSQL has the lowest 95th latency percentile, 1.11 ms, right before OVH PostgreSQL with 1.78 ms. AWS RDS PostgreSQL ranges third with 2.64 ms more than doubling the winner's latency. DigitalOcean PostgreSQL ranges way behind with almost 4.5 ms. The latter two were to be expected considering the throughput results discussed above. Yet, the very low latency shown by Azure Database for PostgreSQL comes as a surprise.

Case-A:
Read-heavy YCSB Workload

UPDATE Latency P95[ms]
less is better

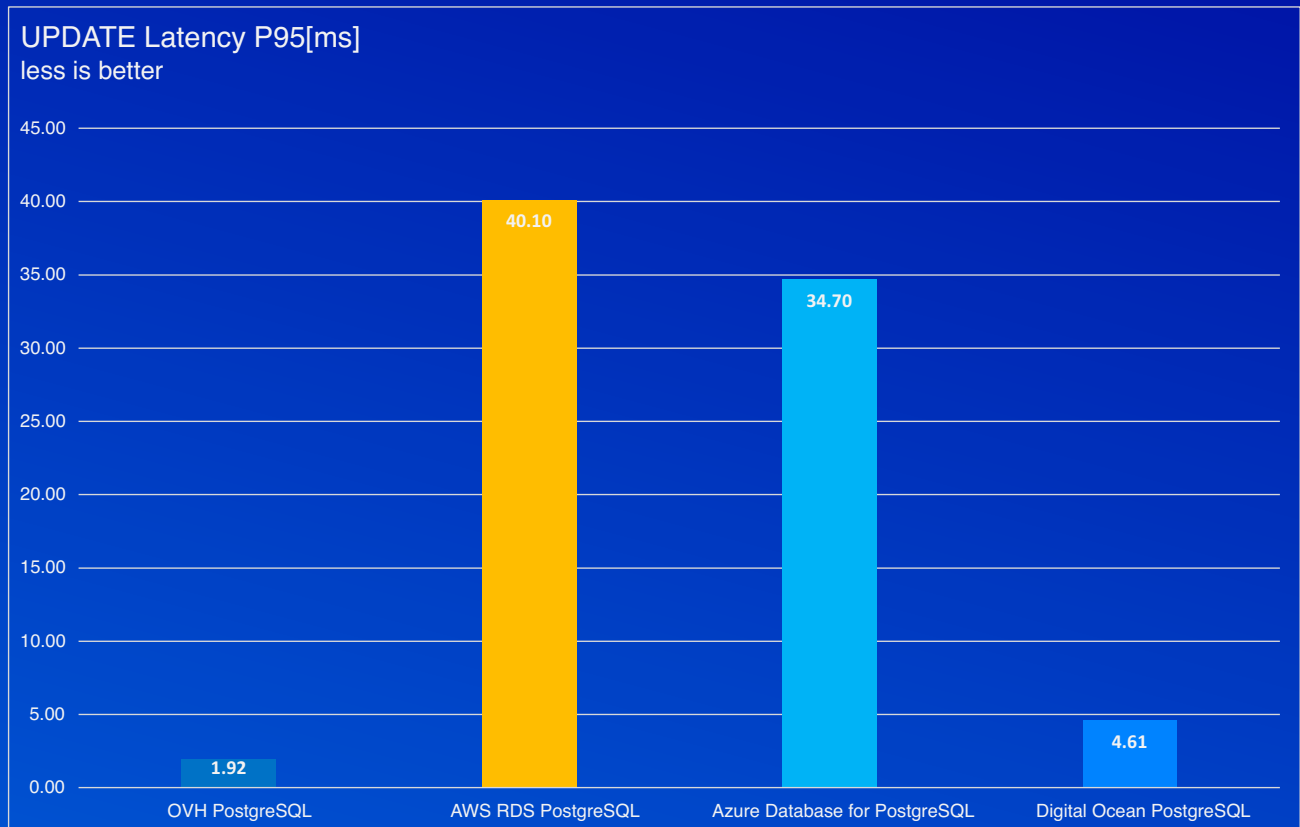|	
40.10	AWS RDS PostgreSQL
34.70	Azure Database for PostgreSQL
1.92	OVH PostgreSQL
4.61	Digital Ocean PostgreSQL

Figure 4. P95 of UPDATE Latency for Read-heavy YCSB Workload

The 95th percentiles of the UPDATE operations (cf. Figure 4) show a completely different picture than for READ. Here, OVH PostgreSQL and Digital Ocean PostgreSQL are the only products that provide comparable latencies for both operations. OVH PostgreSQL is clearly better (1.92 ms in contrast to 4.61 ms) making it the winner in this category. The 95th latency percentile for AWS RDS PostgreSQL for UPDATEs is more than 15x higher than for READs (40.10ms) and almost 21x higher than for OVH PostgreSQL. For Azure Database for PostgreSQL the UPDATE latency is even 31x higher than its READ latency, 18x higher than OVH PostgreSQL. This high UPDATE latency explains why Azure Database for PostgreSQL only ranges second in throughput despite its superior READ performance with very low READ latencies.

Case-A:
Read-heavy YCSB Workload

Price-Performance

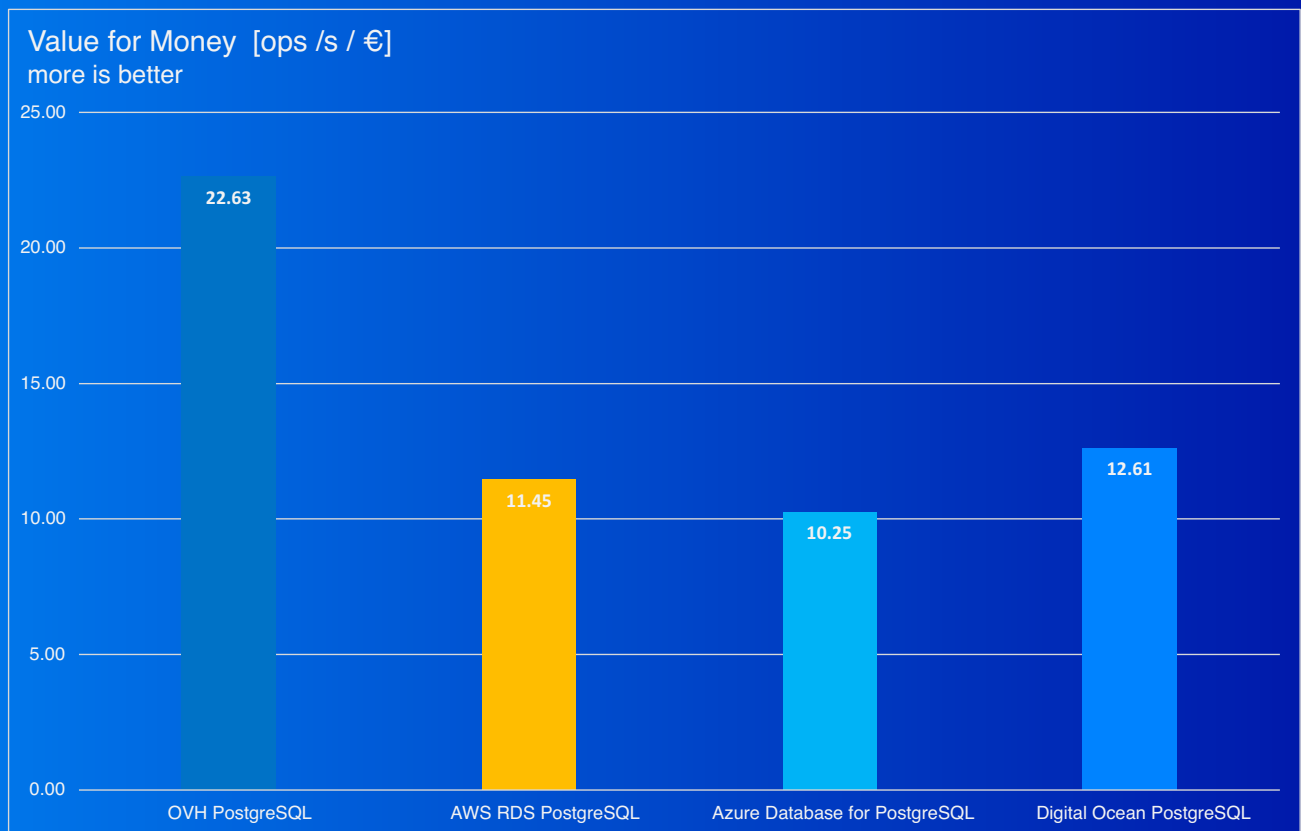Value for Money [ops /s / €]
more is better

Figure 5. Value for Money for Read-heavy YCSB Workload

Figure 5 plots the performance per € for all four competitors under the given workload. With OVH PostgreSQL being amongst the cheapest offers and leading the throughput ranking for this workload, it does not come as a surprise that it also provides the best value for money: 22.63 operations per second per €. What is surprising, though, is the large gap between OVH PostgreSQL and all three competitors who have a price-performance between 10.25 operations per second per € (Azure Database for PostgreSQL) and 12.61 operations per second per € (DigitalOcean PostgreSQL). AWS RDS PostgreSQL is right between those two and 11.45 operations per second per €, barely half of OVH PostgreSQL.

Case-B:
TPC-C Workload by BenchBase

Figure 6 plots the throughput results for the BenchBase TPC-C workload. Due to the nature of the workload and the benchmark, the absolute numbers are a lot lower compared to Case A. OVH PostgreSQL, AWS RDS PostgreSQL and Azure Database for PostgreSQL are all within 140 operations per second: OVH PostgreSQL achieves 2,546 operations per second, while AWS RDS PostgreSQL achieves 2,638 ops / s (+3.6%), and Azure Database for PostgreSQL achieves 2,680 ops / s (+5.3% compared to OVH PostgreSQL). DigitalOcean is trailing behind with only 2,070 ops / s (-22.8% ops/s than Azure Database for PostgreSQL and -18.7 ops/s compared to OVH PostgreSQL).
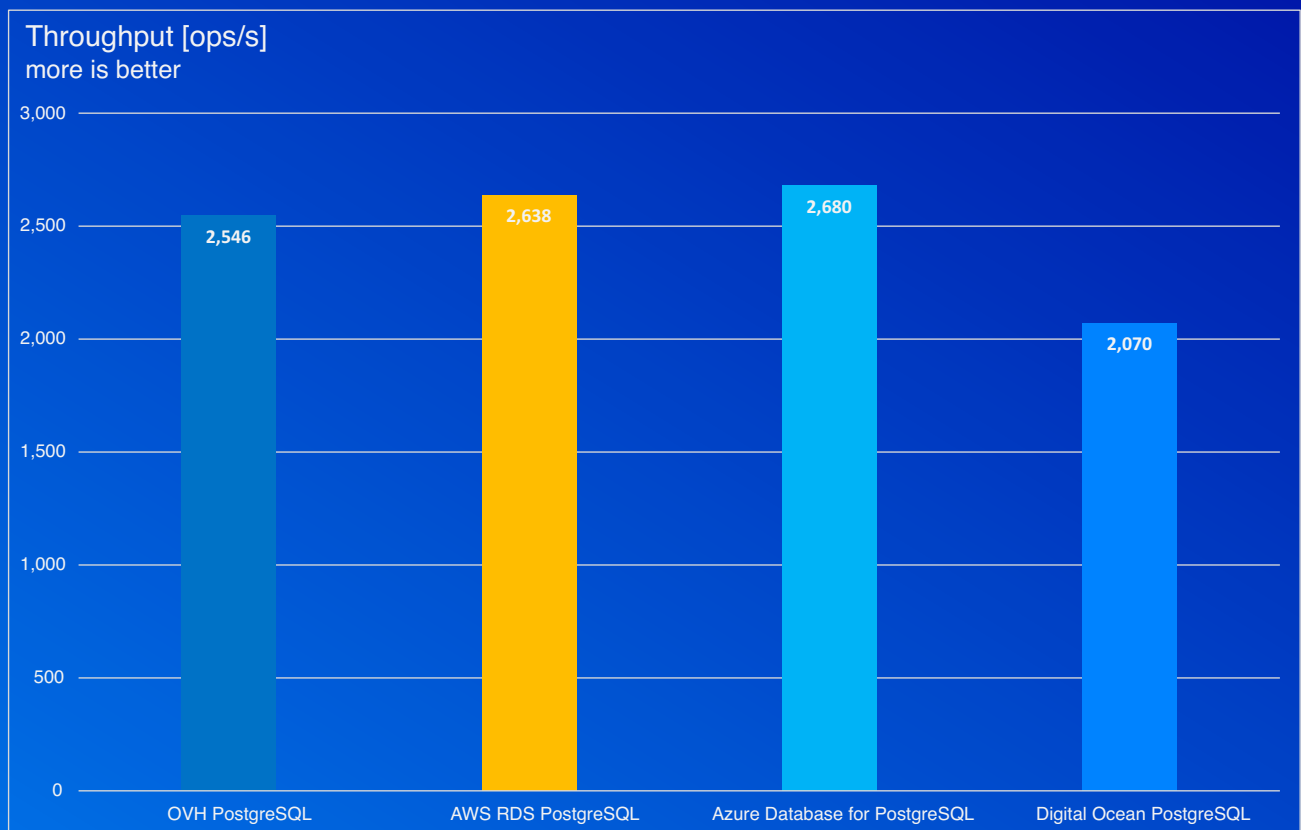

Figure 6. Throughput for BenchBase TPC-C Workload

Case-B:
TPC-C Workload by BenchBase

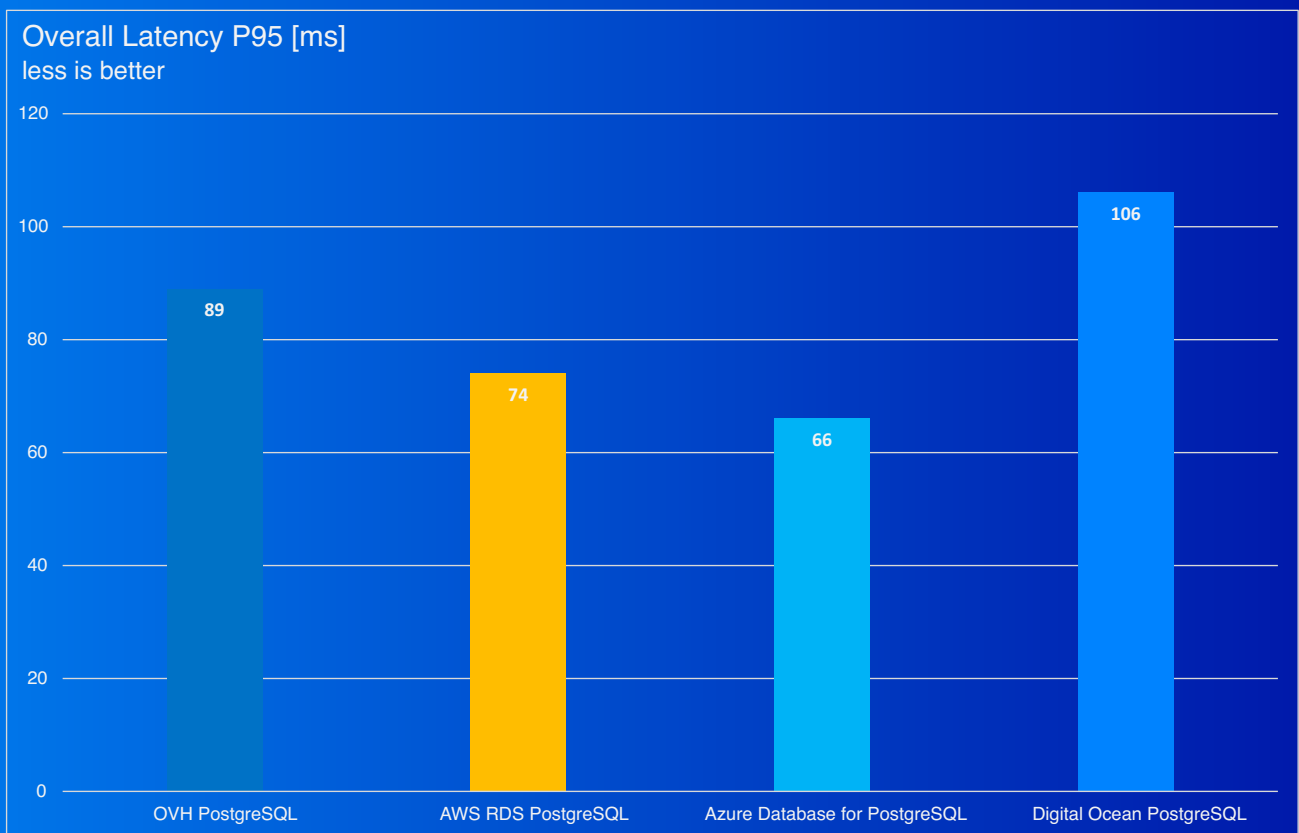**Overall Latency P95 [ms]**
less is better

Figure 7. P95 of Latency for BenchBase TPC-C Workload

These results are mirrored by the latency results shown in Figure 7. Again DigitalOcean shows the worst performance, i.e. highest latency. With 106 ms, it is the only competitor whose P95 latency is beyond 100 ms. In contrast, Azure Database for PostgreSQL is best in class with only 66 ms. AWS RDS PostgreSQL comes in second with 74ms (+12%), while OVH PostgreSQL ranks third with 89ms (+35% compared to Azure Database for PostgreSQL).

Case-B:
TPC-C Workload by BenchBase

Price-Performance

Value for Money [ops /s / €]
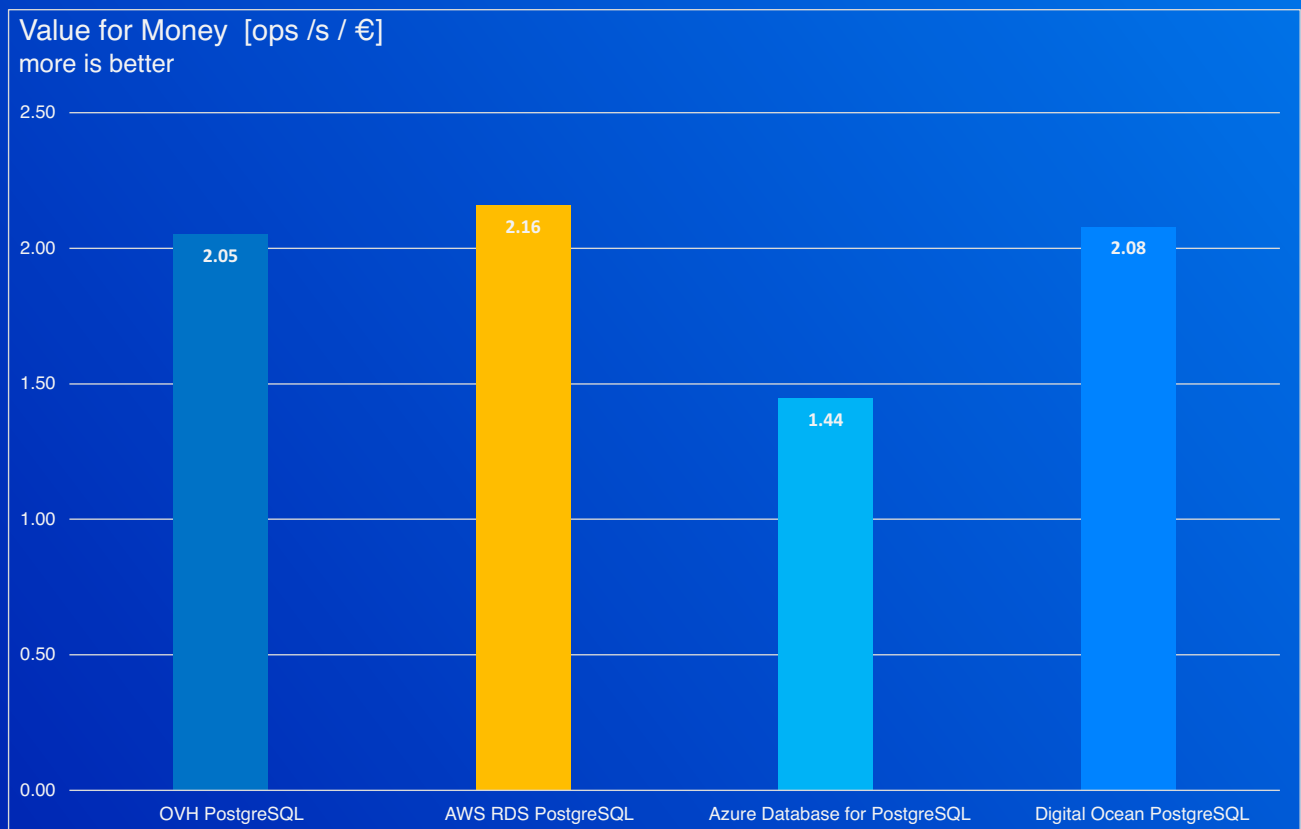more is better

Figure 8. Value for Money for BenchBase TPC-C Workload

Figure 8 breaks down the performance per € for all four competitors. In this normalized metric, the contender with the highest performance, Azure Database for PostgreSQL, is actually the worst due to their high-priced offering. All three other competitors are in the same range: The best amongst them is AWS RDS PostgreSQL with 2.16 operations per second per €, followed by DigitalOcean with 2.08 (-3.7%) and OVH PostgreSQL with 2.05 (-5.1%).

Conclusions

This paper compares the performance and cost performance of OVHcloud's PostgreSQL service against comparable services from Azure, AWS, and DigitalOcean. It shows that OVHcloud's PostgreSQL is truly competitive in the field from a price and performance point of view. In some scenarios it even surpasses established players. Also, price-performance ratio is strong compared to hyperscaler services. This makes OVHcloud PostgreSQL a true alternative to products offered by American companies making it a building block for European customers aiming to increase their digital sovereignty.

Why OVHcloud Might Be Your PostgreSQL Power Player

Disclaimer

This benchmarking project carried out by benchANT was sponsored by OVHcloud with the goal to provide a fair, transparent, and reproducible comparison of the selected database technologies.

Discover OVHcloud databases

See the benchmark online

About benchANT

benchANT is a consulting and analytics firm specializing in comparative performance analysis of database management systems with a focus on cloud hosted databases and Database-as-a-Service technologies. benchANT provides services to database vendors, cloud providers, and end users taking the role of an unbiased analyst, researcher, and evaluator. The experiments described in this paper have been designed, executed, and written-up by two of benchANT's key employees.

Dr. Jörg Domaschka is one of benchANT's co-founders. He has been trying to understand distributed systems for more than two decades. Performance engineering and benchmarking help him with that task and allow him to educate others on his findings.

Dr. Daniel Seybold is a co-founder and the CTO of benchANT. Daniel started his career as a researcher with a focus on distributed systems and databases. He has extensive experience in the field of database performance testing and has been working with NoSQL databases such as MongoDB, Cassandra and ScyllaDB for more than a decade.

About OVHcloud

OVHcloud is a global cloud player and the leading European cloud provider operating over 450,000 servers within 43 data centers across 4 continents to reach 1,6 million customers in over 140 countries. Spearheading a trusted cloud and pioneering a sustainable cloud with the best performance-price ratio, the Group has been leveraging for over 20 years an integrated model that guarantees total control of its value chain: from the design of its servers to the construction and management of its data centers, including the orchestration of its fiber-optic network. This unique approach enables OVHcloud to independently cover all the uses of its customers so they can seize the benefits of an environmentally conscious model with a frugal use of resources and a carbon footprint reaching the best ratios in the industry. OVHcloud now offers customers the latest-generation solutions combining performance, predictable pricing, and complete data sovereignty to support their unfettered growth.