

Performance Testing PostgreSQL® on Azure NetApp Files

Executive Summary

Instaclustr PostgreSQL® on Azure NetApp Files (ANF) is up to 325% faster (TPS) and 70% cheaper (\$/TPS) than equivalent Instaclustr Managed PostgreSQL instances using Azure Premium Disk. Using the higher IOPS available on ANF storage we can increase read/write transactions per second (TPS) by 135% to 167% and read-only by 193% to 325% vs an Azure Premium Disk based Instaclustr PostgreSQL. Using ANF provides a 50%+ reduction in the cost per TPS performance of heavy read/write workloads on a PostgreSQL database.

Introduction

Since NetApp's acquisition of InstaClustr, we have been busy working on ways to bring the best of both our technologies to our customers. One of the first examples of this has recently been released: InstaClustr's Managed PostgreSQL running on Azure NetApp Files. This white paper demonstrates the high-end performance and price-performance advantages that are available today with InstaClustr Managed PostgreSQL on ANF. The following sections provide details on the testing methodology that we applied, and results achieved to demonstrate this performance advantage.

However, performance is only one advantage of running PostgreSQL on ANF. Future releases are planned to allow us to make use of the snapshots and mirroring features of ANF to offer super-fast backups, forking, and multi-region disaster recovery as we progress our roadmap.

Testing Methodology

We used pgbench to run our tests. We created a separate VM running pgbench in the same network as our PostgreSQL instance to simulate a customer application talking to the database across network.

We tested 2 instances to gather comparisons. The DS13 instance runs on a Microsoft Premium SSD while the E8s_v4 runs on ANF premium storage.

- InstaClustr PGS-PRD-Standard_DS13_v2-2000 (8 vCPU 56GB RAM 2TB Storage)
- InstaClustr PGS-PRD-Standard_E8s_v4-2048-ANF (8 vCPU 64GB RAM 2TB Storage)

Note: As part of the release of Postgres-ANF we upgraded the VM hardware to the more modern Esv4 from the older Dsv2. This offers a slightly lower cost and our testing demonstrated that it does not impact performance for the premium disk comparison. Results from "other managed postgres" running on E8s_v4 hardware show extremely similar TPS to the DS13_v2 hardware.

Before testing, we loaded the databases with 1.5TB of data to ensure this was a good simulation of a production workload. We have found in the past that if we load only small amounts of data prior to benchmarking tests that this data can be cached in RAM and does not accurately represent real world performance.

Testing was driven by a bash script to loop over different client counts to research scalability. The test was performed 3 times per configuration and the results averaged.


Minimal variance was seen between repeated runs. In our pgbench configuration, scale was set to 10,000, clients were looped through 4, 8, 16, 32, 64, and 96. Full scripts we used for testing can be found on our [GitHub](#).

Performance Tuning

Initial testing showed that the Postgres-ANF server performed very well when clients=cores, but performance would drop off unexpectedly beyond that. Investigation and metric logging showed that bottlenecks were being created in the WAL compression and `bgwriter_delay`. Adjusting these parameters allowed us to push well past the 8-client count and get the results you see below. These tuning optimizations are now implemented in Instacluster Managed PostgreSQL by default and the results below are based on the current default configuration.


Results

Read/Write Workload



The increased IOPS available on the ANF storage showed strong improvement on the performance of the PostgreSQL application. The highest increase was 167% faster TPS than the equivalent Azure Premium Disk based Instacluster server at 8 clients, and the worst we saw was a 127% increase at 32 clients. ANF storage consistently delivered higher TPS, was durably much faster at all client counts tested, and this result was consistent across many repeated runs over the course of days.

Read-Only Workload



The increased IOPS available on the ANF storage had an even more dramatic impact on the read-only performance of the PostgreSQL application. The highest increase was 325% faster TPS than a Azure Premium Disk based Instacluster server at 64 clients, and the worst we saw was a 193% increase at 16 clients. Once again TPS was consistently higher at all client counts tested and this result was demonstrated across many repeated runs over the course of days.

Analysis

Price vs. Performance

Region	Azure-AZ-Central-US
Cores	8
RAM	64GB
Disk	2048GB
Clients	64

Instaclustr RIYOA Price

Instaclustr offers a "Run In Your Own Account" package, where you pay Instaclustr for the provisioning and management of your cluster, and all other costs are included in your cloud bill to Microsoft (in green). You may have negotiated better rates for infrastructure from Microsoft which will reduce the TCO.

	PGS-PRD-Standard_DS13_v2-2000-an	PGS-PRD-Standard_E8s_v4-ANF-2048
TPS (Read only)	4030	17140
TPS (Read/Write)	1963	4479
Management Units	3	3
Management Units Price	1500	1500
Instance Cost (On-Demand)	541.3	368.17
Storage Cost (On-Demand)	259.05	602.50
Network Cost (estimated)	340	348.16
Backup Cost	52.8	54.07
Price per month	\$2,693.15	\$2,872.90
\$/TPS (Read only)	\$0.67	\$0.17
\$/TPS (Read/Write)	\$1.37	\$0.64

Postgres-ANF does attract a premium price compared to an Azure Premium Disk based Instaclustr deployment due to the high-speed storage. However, at full utilization, the price per TPS is at least 50% cheaper for read/write workloads and 70% cheaper for read-only. For customers with demanding read or read/write heavy workload, we recommend using ANF storage to back your PostgreSQL cluster. Instaclustr will continue to offer Azure Premium Disk based PostgreSQL nodes for customers with smaller, less demanding workloads.

Summary

Using faster storage such as Azure NetApp Files (ANF) to back your PostgreSQL database can substantially increase the TPS workload the server can handle. We have shown for a relatively small cluster (8 cores) the uplift is up to 270%, and are preparing to test on much large VMs for the most demanding workloads. Even if your workload requirement does not exceed the capabilities of standard storage, there may be price/performance advantages from ANF. This and the additional backup and forking management capabilities that will be available soon with Instacluster's PostgreSQL on ANF make it a compelling offering for many PostgreSQL use cases on Azure.

About Instacluster

Instacluster helps organizations deliver applications at scale through its managed platform for open source technologies such as [Apache Cassandra[®]](#), [Apache Kafka[®]](#), [Apache Spark[™]](#), [Redis[™]](#), [OpenSearch[®]](#), [PostgreSQL[®]](#), and [Cadence[®]](#).

Instacluster combines a complete data infrastructure environment with hands-on technology expertise to ensure ongoing performance and optimization. By removing the infrastructure complexity, we enable companies to focus internal development and operational resources on building cutting edge customer-facing applications at lower cost. Instacluster customers include some of the largest and most innovative Fortune 500 companies.

© 2023 Instacluster Copyright | Apache[®], Apache Cassandra[®], Apache Kafka[®], Apache Spark[™], and Apache ZooKeeper[™] are trademarks of The Apache Software Foundation. Elasticsearch[™] and Kibana[™] are trademarks for Elasticsearch BV. Kubernetes[®] is a registered trademark of the Linux Foundation. OpenSearch is a registered trademark of Amazon Web Services. Postgres[®], PostgreSQL[®] and the Slonik Logo are trademarks or registered trademarks of the PostgreSQL Community Association of Canada, and used with their permission. Redis[™] is a trademark of Redis Labs Ltd. Any rights therein are reserved to Redis Labs Ltd. Cadence is a trademark of Uber Technologies, Inc. Any use by Instacluster Pty Limited is for referential purposes only and does not indicate any sponsorship, endorsement or affiliation between Redis and Instacluster Pty Limited. All product and service names used in this website are for identification purposes only and do not imply endorsement.