



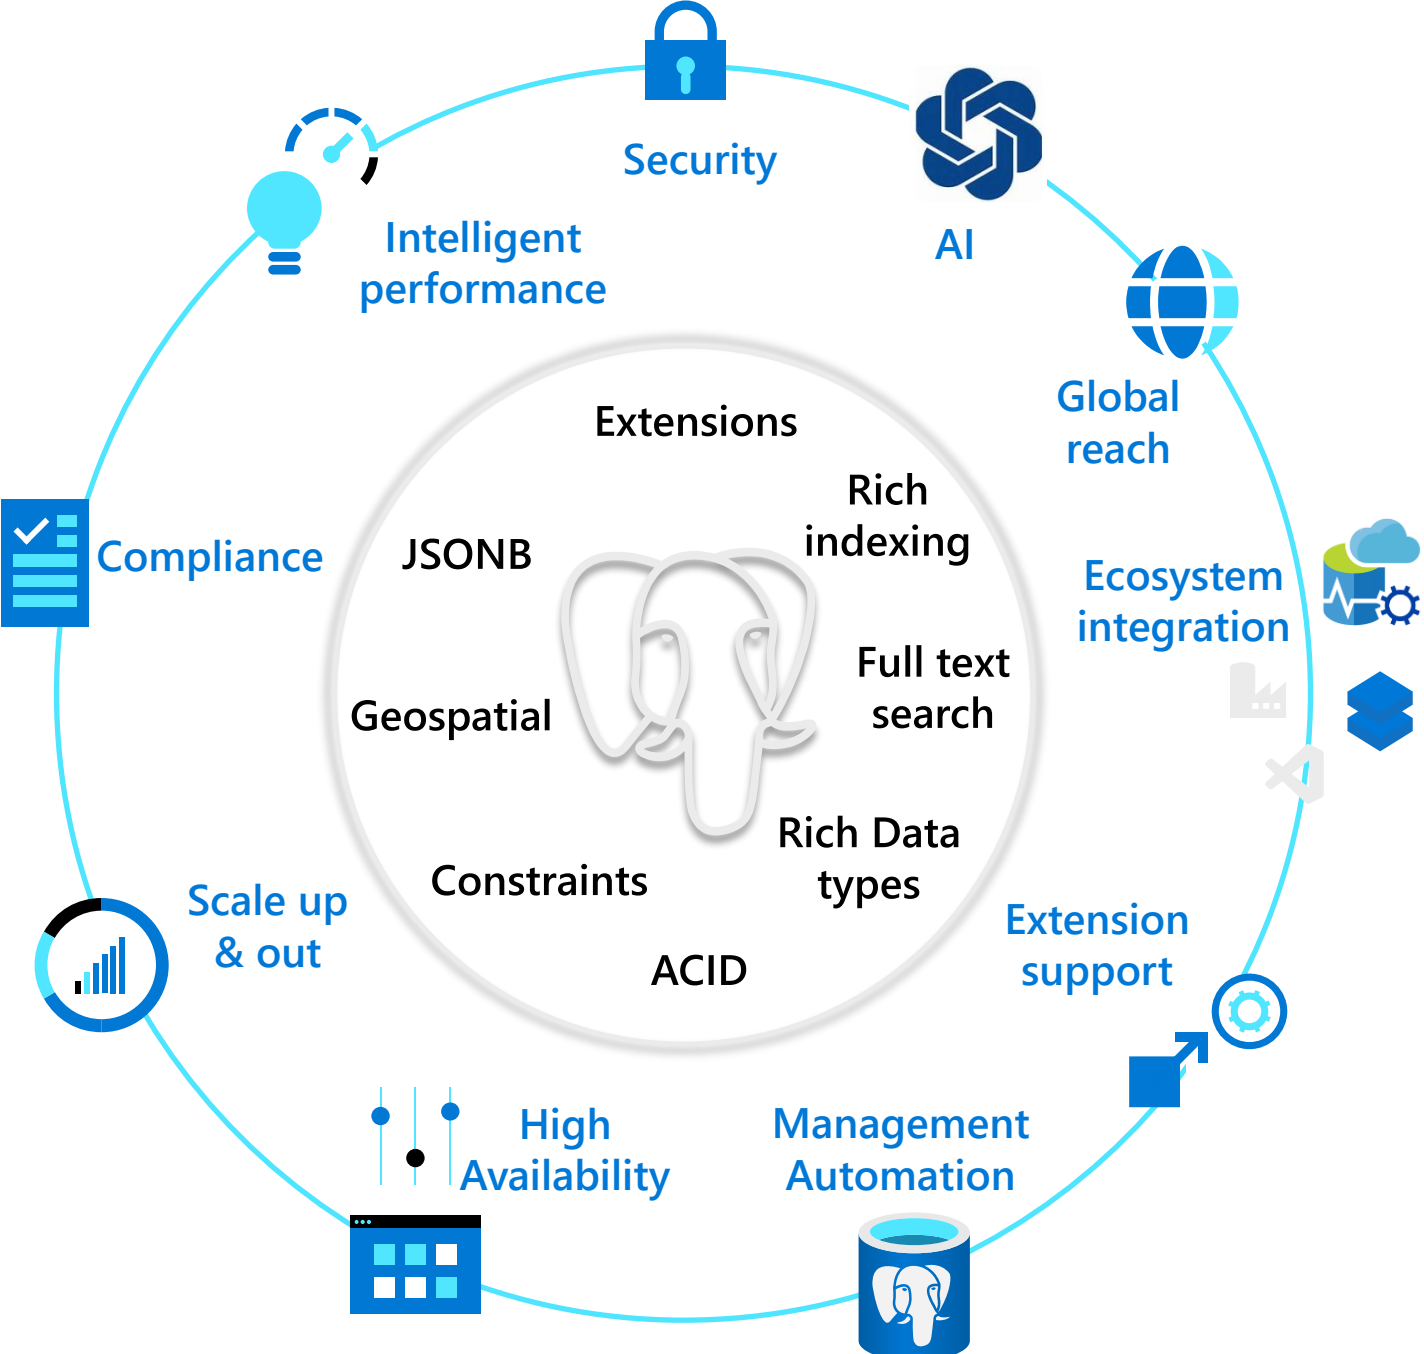
# Microsoft Azure



# Azure Database for PostgreSQL

Operational Databases - GBB team  
February 2025


# Azure builds upon PostgreSQL



# Postgres is the most popular database for professional developers

## PostgreSQL extended lead 2024

Which **database environments** have you done extensive development work in over the past year, and which do you want to work in over the next year? (If you both worked with the database and want to continue to do so, please check both boxes in that row.)



Source:  
Stack Overflow Developer Survey 2024

# Azure Database for PostgreSQL

Designed to meet both enterprise and developer needs

## Enterprise

---

High Availability  
Disaster Recovery  
Scalability  
Security



## Developers

---

Ecosystem Integration  
Idiomatic SDKs  
Community led Open Source  
Generative AI Apps

# Azure Database for PostgreSQL is Enterprise Ready Today

Entra  
ID



Private  
Endpoints



Virtual  
Networks



Data  
Encryption



Azure  
Defender



Availability  
Zones



Geo DR



Performance



Azure  
Policy



Azure  
Advisor

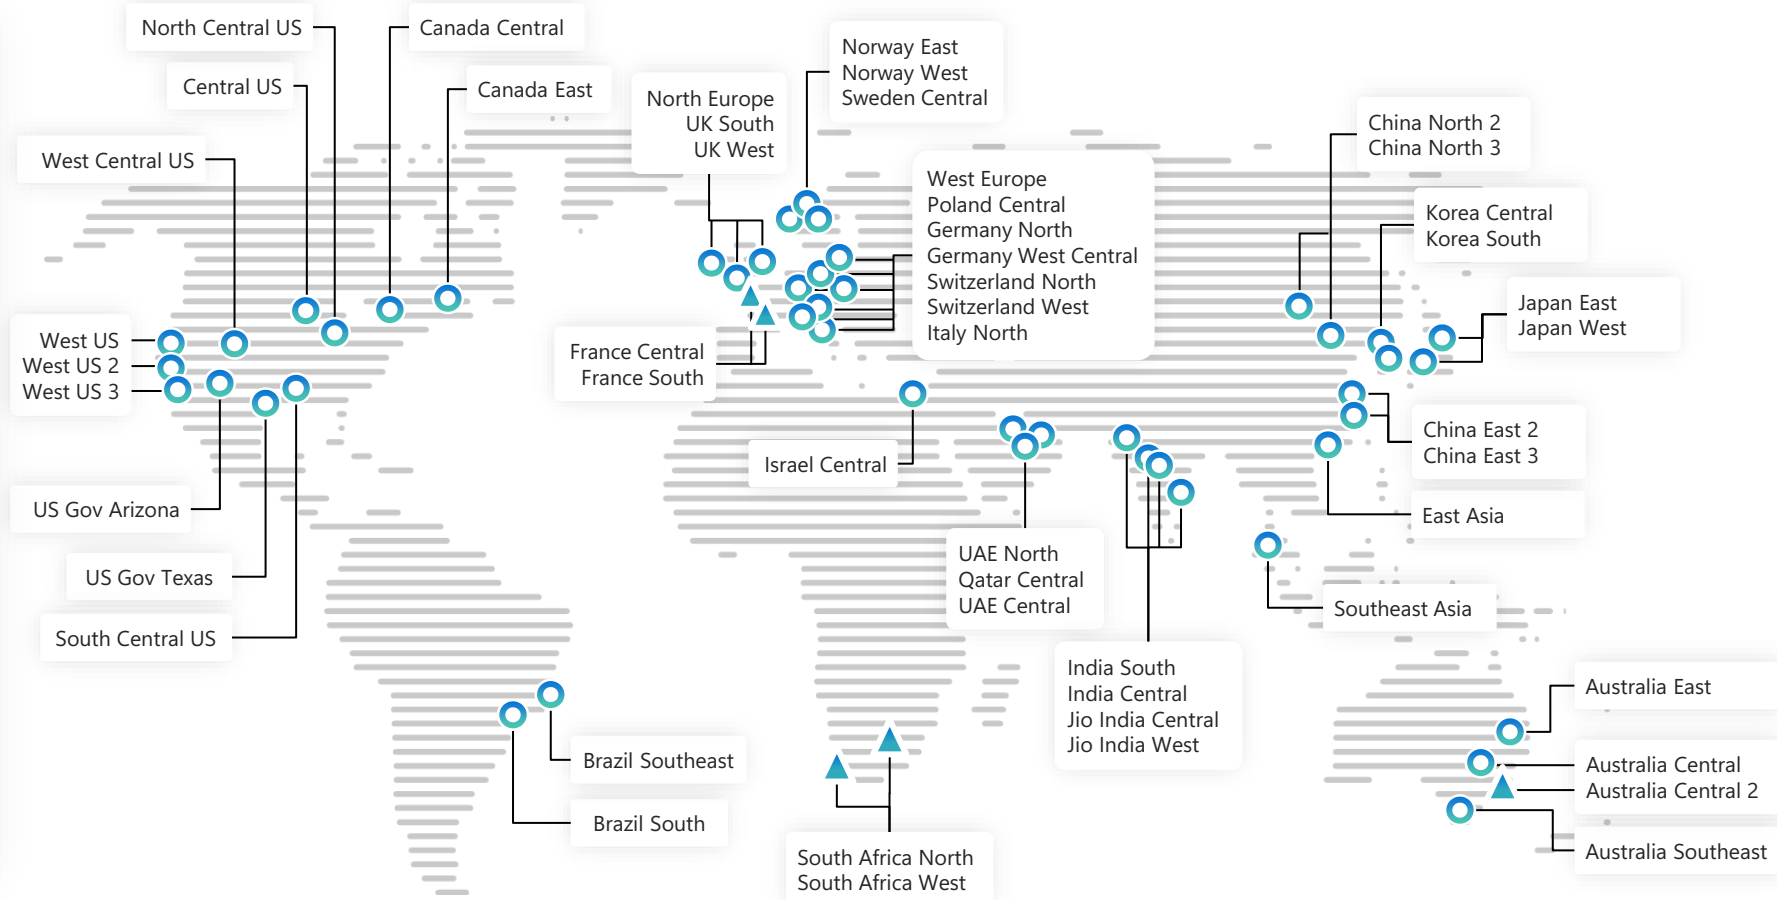


# Industry Leading Global Reach

# 55+

Azure regions available

- Available
- ▲ Coming



# Postgres 17 Contributions

Top level metrics (as of April 2024)

353

Commits by Microsoft  
committers

27.2%

Microsoft commits  
done on behalf of  
other authors

331


Commits credited to  
Microsoft authors





## Azure Database for PostgreSQL





# Industry leading AI for building intelligent applications



[1.5, -0.4, ..., 20.2]

Native  
Vectors



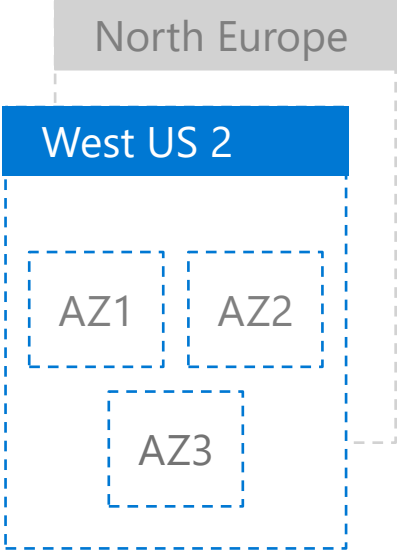
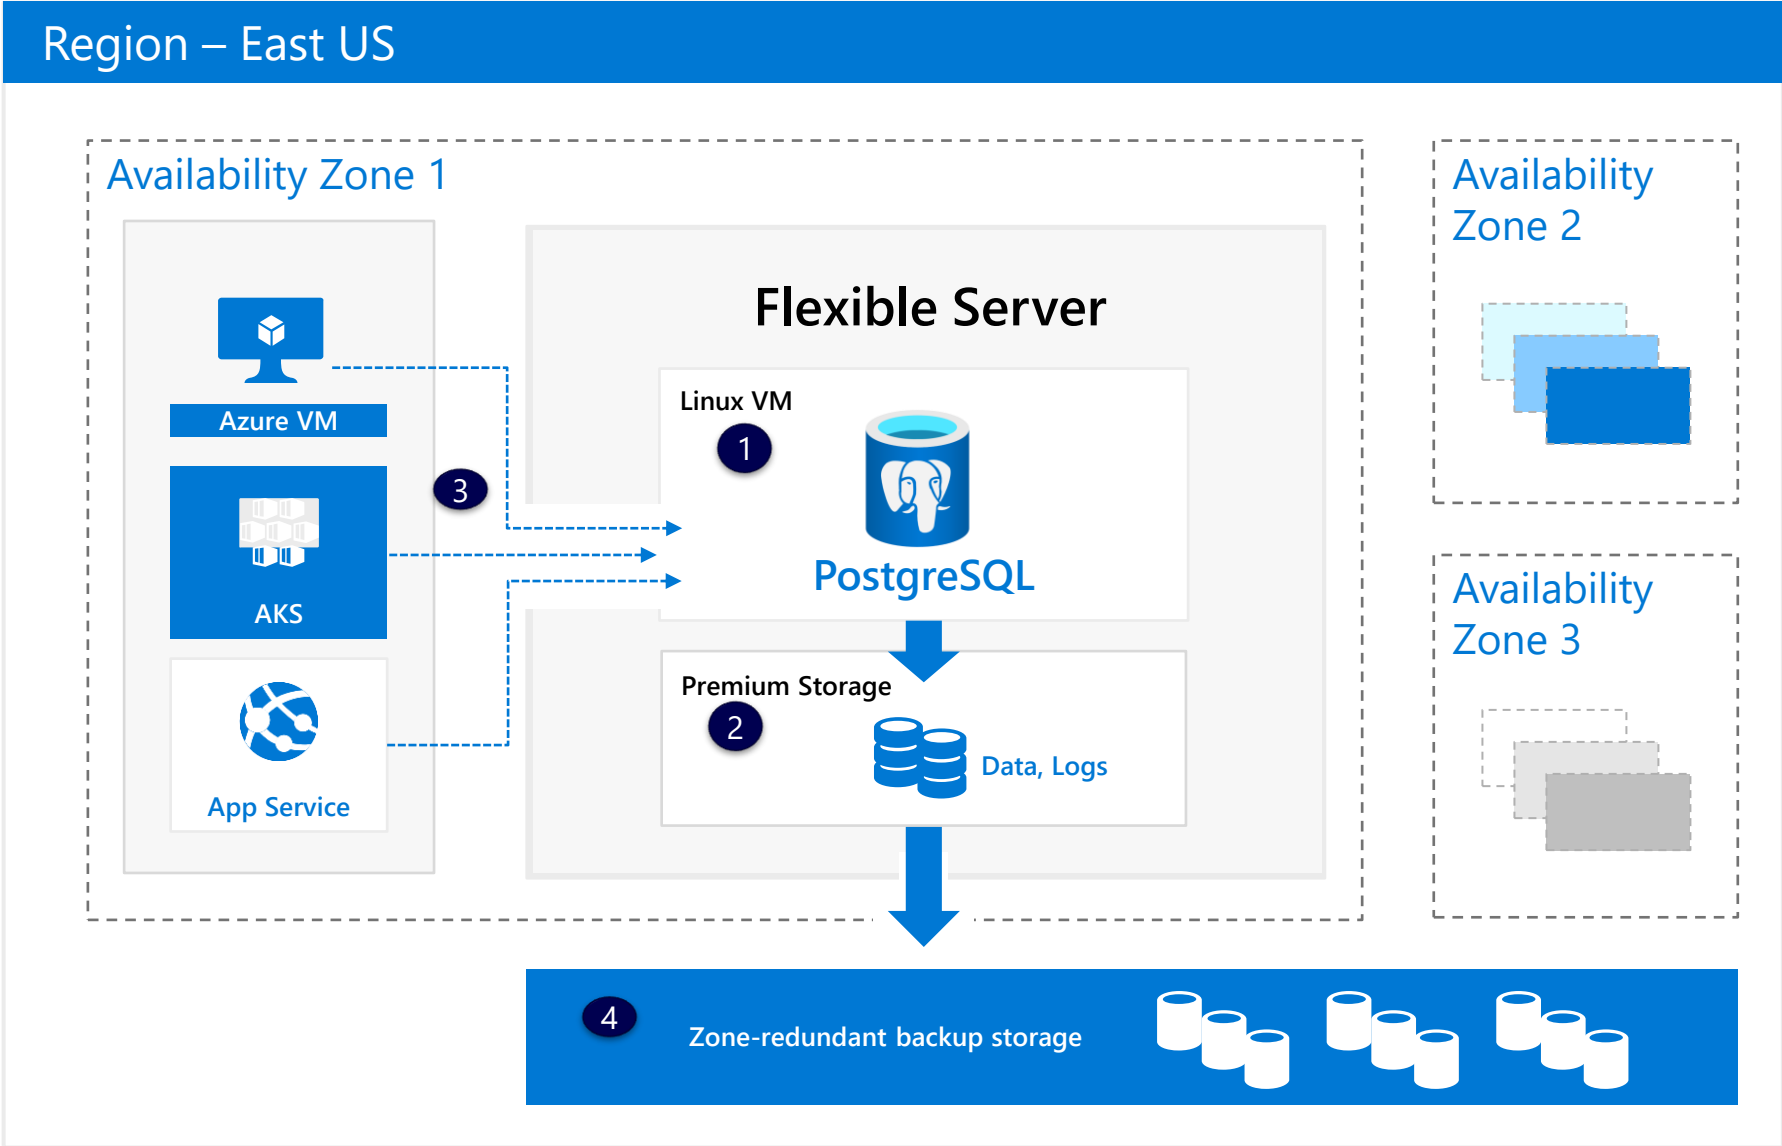
Azure AI  
Integrated



Copilot  
Powered

# Service Overview

# Flexible Server Architecture



- 1 Linux VM
- 2 Premium managed disks (3 copies)
- 3 AZ co-location with applications
- 4 Zone-redundant backup storage

# Workload Optimized Compute SKU's

Cost optimized for different workloads

---

Each switch between any SKU in minutes

---

Stop/Start during inactive periods

---

Reserved Capacity



## Memory Optimized

Up to 96 vCores with 1:8 CPU to Memory ratio optimized for best performance of IO intensive workloads



## General Purpose

Up to 96 vCores with 1:4 CPU to Memory ratio suitable for most database workloads



## Burstable

Highly cost effective, ideal for Development and Testing

# Elastic Compute and Storage

Scale compute in under 20s

---

Scale storage online


---

Scale storage size and IOPS separately

---

99.99% SLA with Availability Zones

## Compute



## Storage



# All Community Supported Postgres Versions

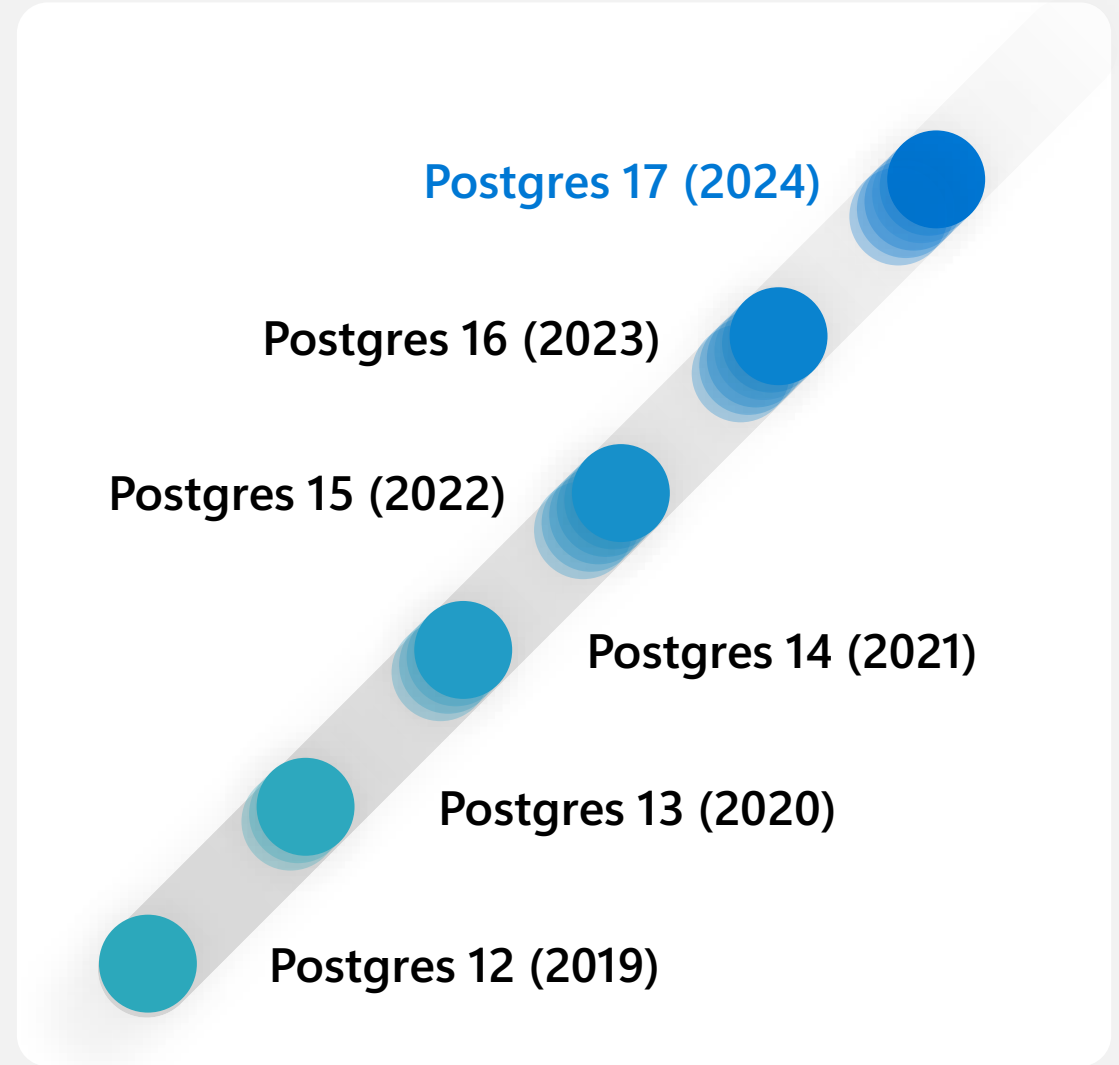
Major versions available **within weeks** of community GA

---

Minor versions **maintained automatically**

---

**Upgrade in-place** in minutes



# Broad support for common Postgres extensions

60+ Postgres extensions supported

Enables developers to **extend the functionality of Postgres** beyond core capabilities

Microsoft **automatically maintains** extensions versions

address_standardizer	pg_freespacemap
address_standardizer_data_us	pg_hint_plan
amcheck	pglogical
azure_ai	<b>pg_partman</b>
azure_local_ai (Preview)	pg_prewarm
<b>azure_storage</b>	pg_repack
bloom	pgrouting
btree_gin	pgrowlocks
btree_gist	pg_squeeze
citext	pg_stat_statements
cube	pgstattuple
dblink	pg_trgm
dict_int	pg_visibility
dict_xsyn	plpgsql
earthdistance	plv8
fuzzystrmatch	<b>postgis</b>
hstore	postgis_raster
hypopg	postgis_sfcgal
intagg	postgis_tiger_geocoder
intarray	postgis_topology
isn	postgres_fdw
lo	semver
login_hook	session_variable
ltree	sslinfo
orafce	tablefunc
pageinspect	tds_fdw
<b>pgaudit</b>	timescaledb
pg_buffercache	tsm_system_rows
<b>pg_cron</b>	tsm_system_time
pgcrypto	unaccent
pg_failover_slots (Preview)	uuid-oss
	vector



# Extensive Monitoring

Rich metrics and logs provide observability into the entire database workload

Access to detailed metrics and logs

Quickly diagnose performance issues

Make informed scaling decisions

Set up alerts and auto-scaling for quick responses

Visualize data using the Portal, Power BI, Grafana, or Log Analytics

## Azure PG-Flexible-Svr - Monitoring


Private dashboard

+ Create Upload Refresh Full screen Edit Share Export Clone Assign tags Delete Feedback


Auto refresh: Off UTC Time: Past 30 minutes Add filter

Get access to an improved dashboard experience and new mobile presence. [Try it now](#)


### Avg CPU percent and Avg Memory percent for varund-m33-april2024




### Avg Active Connections for varund-m33-april2024



### Avg Read IOPS and Avg Write IOPS for varund-m33-april2024



### Avg Database Size (Preview) for varund-m33-april2024 by DatabaseName



# Integrated with Azure Advisor

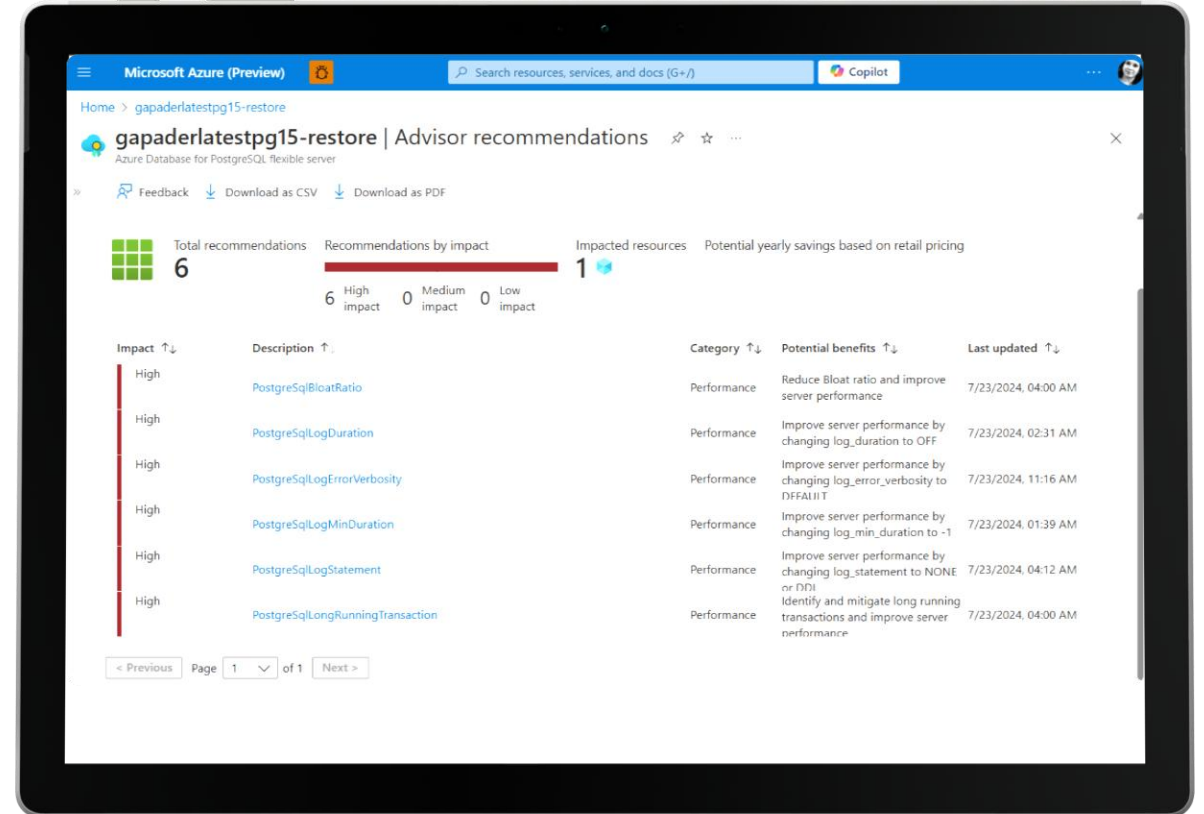
## Built-in recommendations for workload optimization



Azure Advisor supports over 16 recommendations for optimizing workloads on Azure Database for PostgreSQL

Performance optimizations are based on actual usage history to prioritize highest impact

Recommendations include optimizations for **logging, disk space, and memory usage** amongst others



**Enterprise-Ready**

# Enterprise Ready



## Security

Identity

---

Encryption

---

Networking



## High Availability

Availability Zones

---

RPO = 0

---

Automatic Failovers



## Disaster Recovery

Backups

---

Geo-Redundancy

---

Long Term Retention

# Enterprise Security

Azure Database for PostgreSQL is **the only Postgres offering** with support for Entra Id Authentication

---

Service or Customer Managed encryption keys to **protect data at rest**

---

**Secure network connectivity** via Private Endpoints



## Entra Id

Enterprise Identity built into Postgres



## Customer Managed Keys

Key Vault integration with optional HSM support



## Private Endpoints

Network interface that uses a private IP address from your virtual network



# Enterprise Security

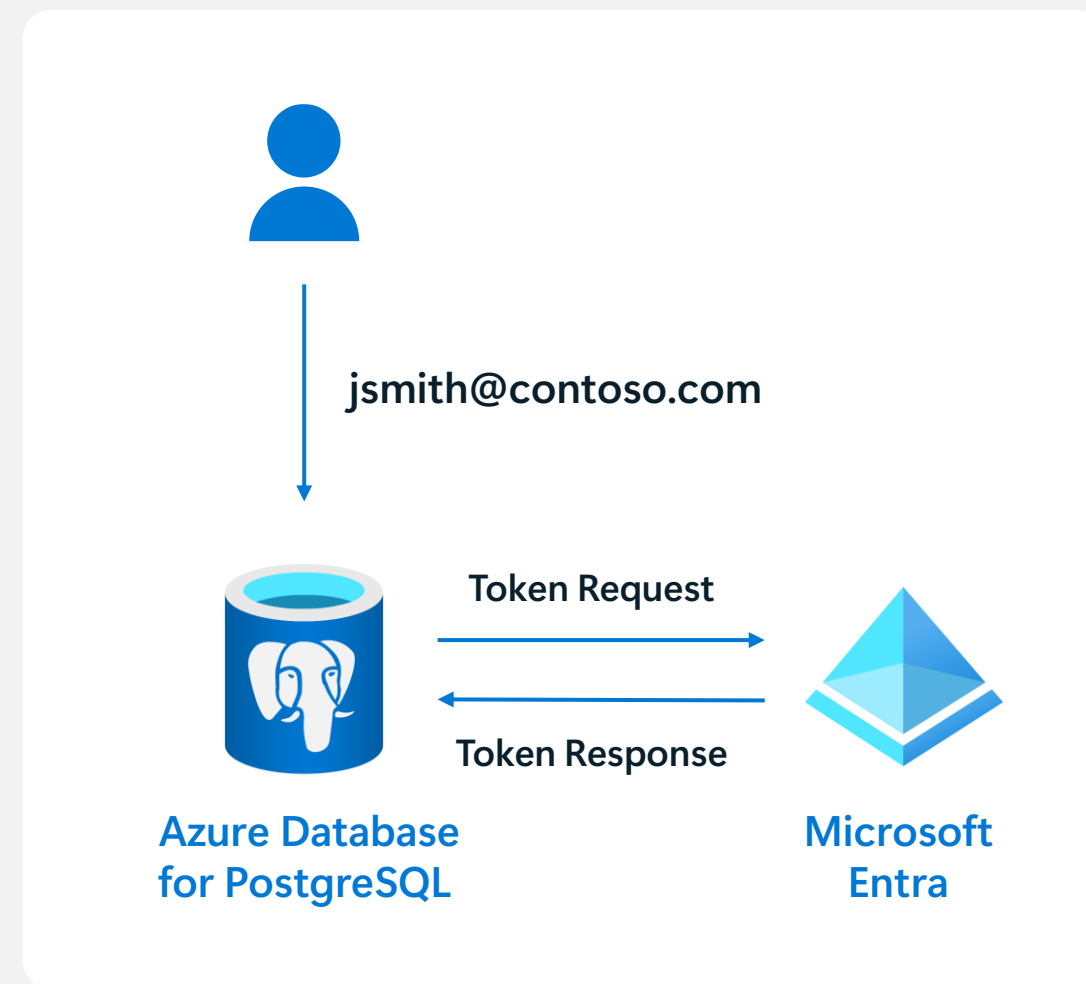
## Enterprise Identity

Azure Database for PostgreSQL is **the only Postgres offering** with support for **Entra Id Authentication**

**No passwords stored** in Postgres service

Users authenticate with their **Entra (formerly AAD) enterprise credentials**

Manage access within Entra to **grant/revoke permissions efficiently**



# Enterprise Security

## Enterprise Identity

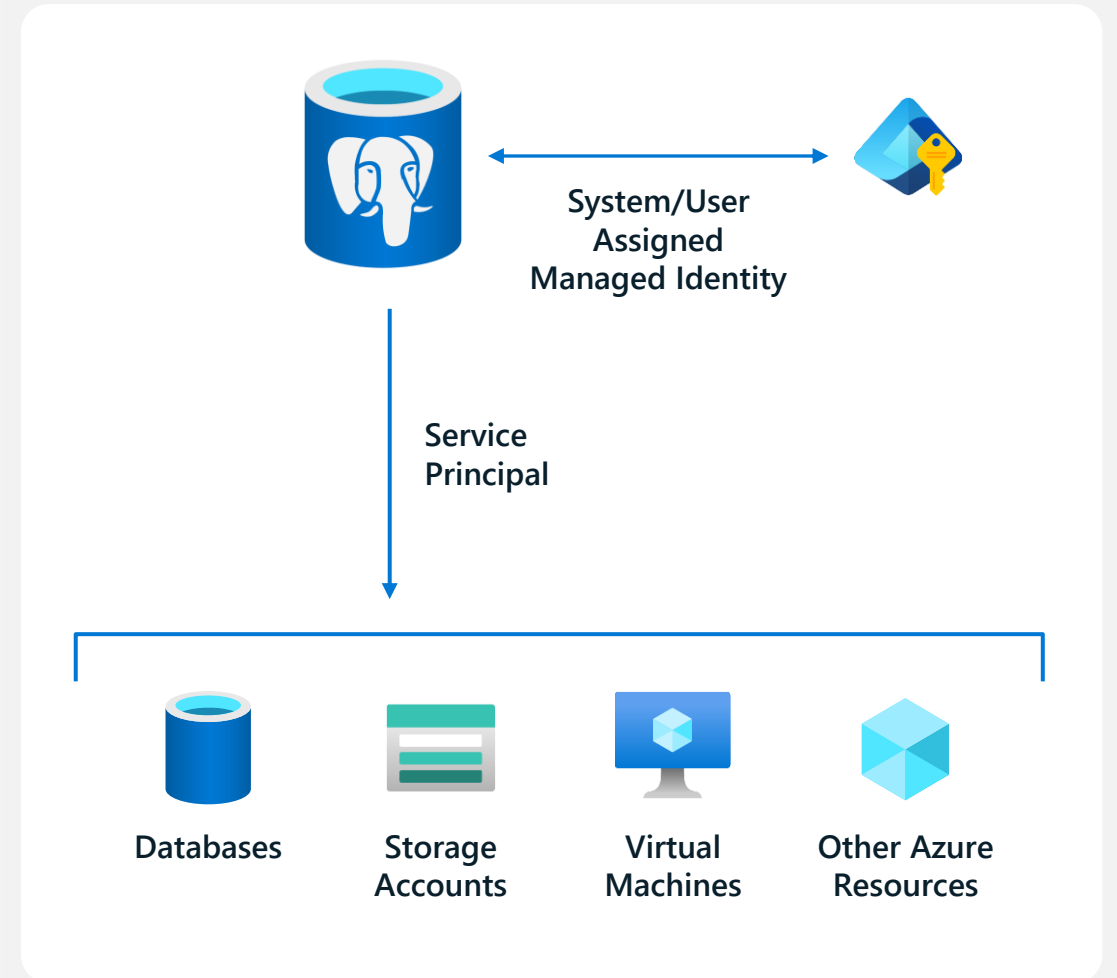
Use Managed Identities to **connect to any Azure resource** that supports Entra ID

---

**Eliminates the need** to store and manage secrets within Postgres service

---

Support for both **User and System Assigned** Managed Identities



# Enterprise Security

## Customer Managed Encryption Keys

Flexible Server supports **Customer Managed Encryption Keys** for securing data

---

**Rotate and revoke keys** for Postgres as needed

---

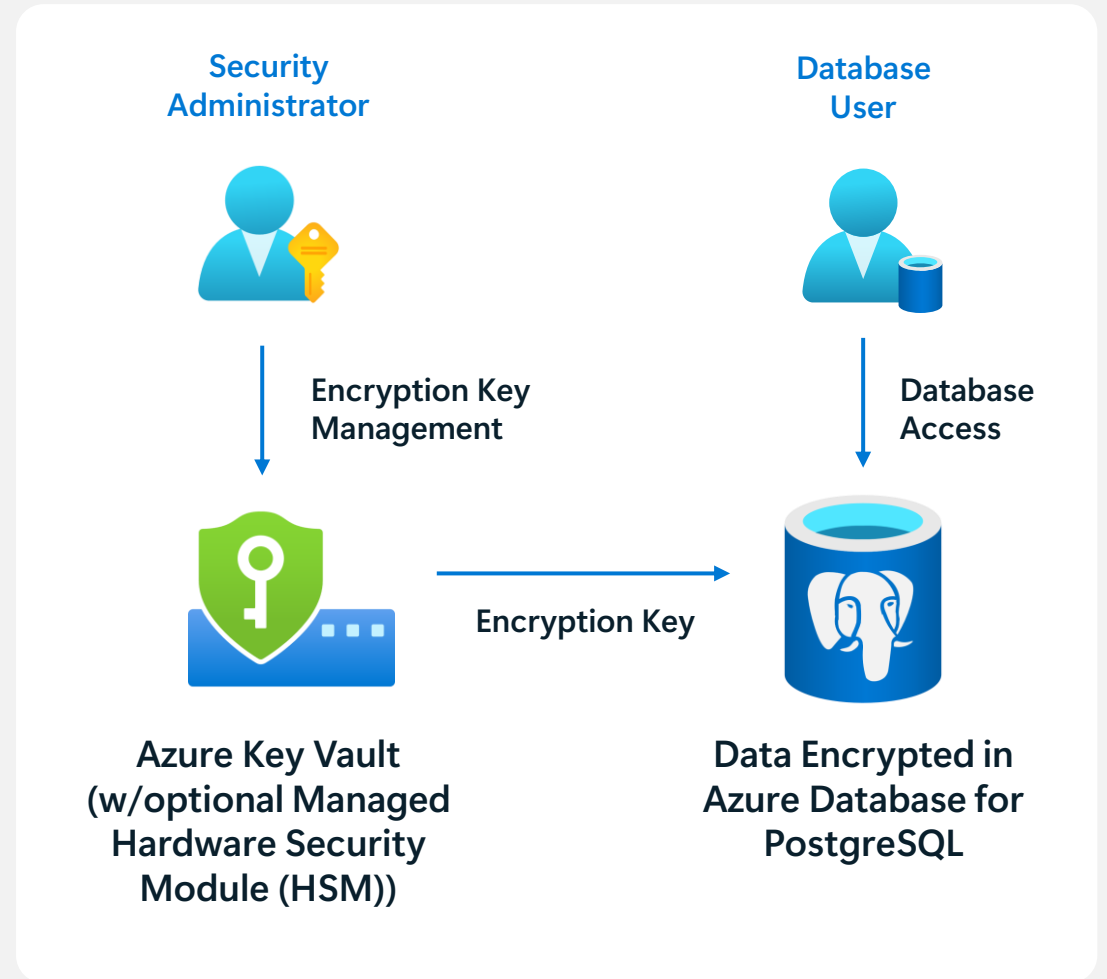
Meet regulatory and compliance requirements by managing **key lifecycle and access policies**

---

Provides **separation of duties** to key management from data management

---

**Centralized management** of keys in Azure Key Vault or Managed HSM





# Enterprise Security

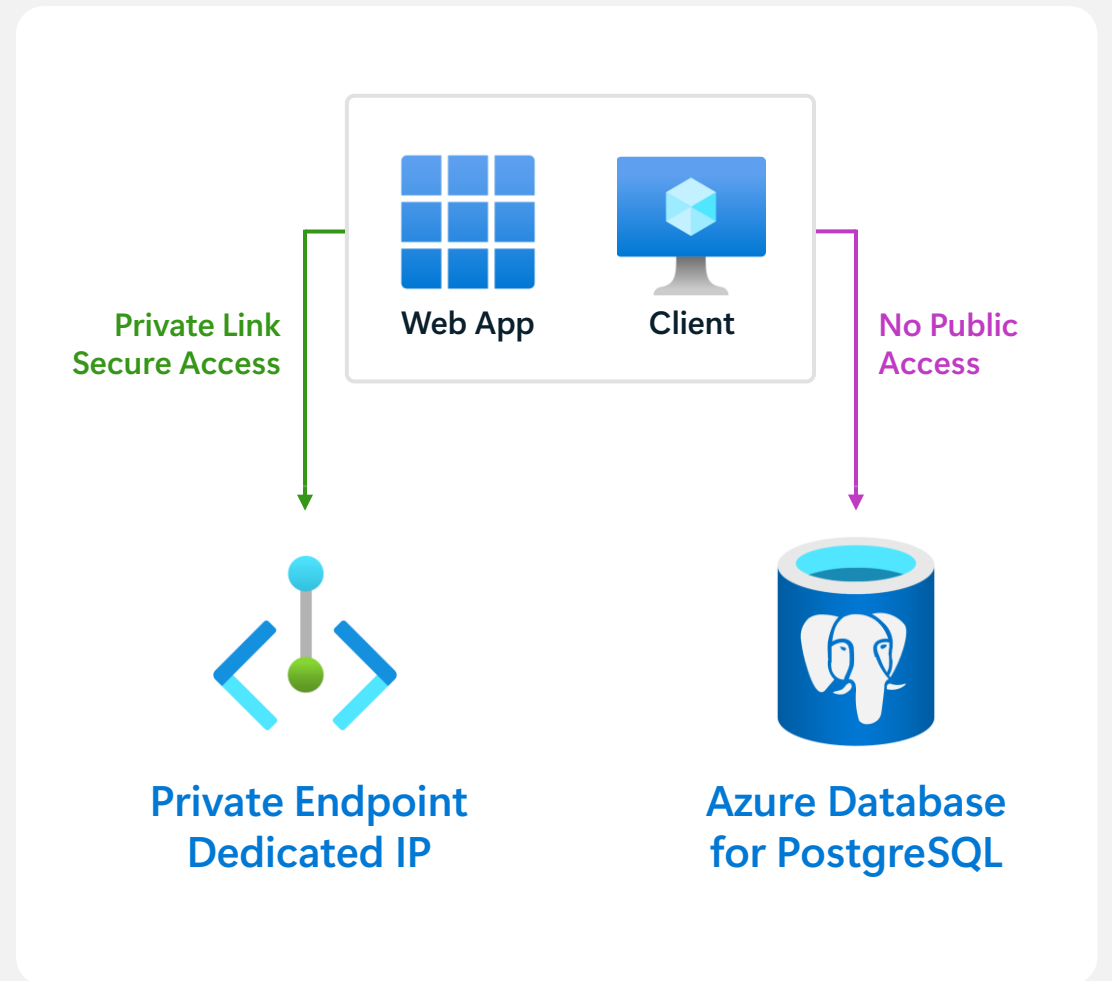
## Private Endpoints

Provides a **private IP address** from your virtual network for accessing Flexible Server.

**Prevents data exfiltration** from your virtual network

**Improved performance** using a direct route to Azure resources vs. connecting via Internet

Fine-grained access control using **network security groups (NSG's)** to restrict access to the private endpoint



# High Availability

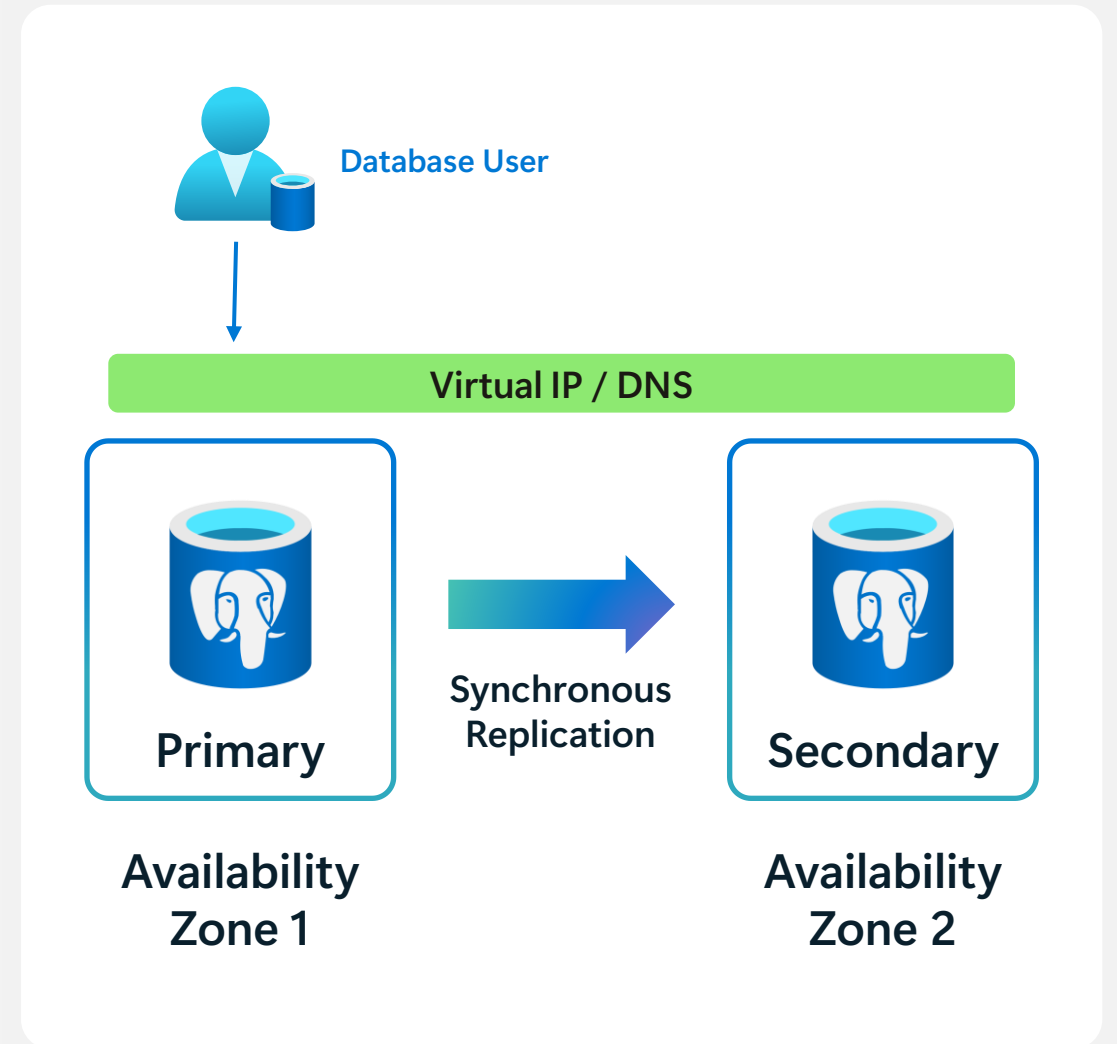
99.99% Uptime SLA

Provides a replica of the production database across Availability Zones with a (no data loss)

**Recovery Point Objective = 0**

**Automatic Failover** in scenarios where servers or zones fail

**Automatic rebuild** of the Secondary HA instance after failover occurs



# Disaster Recovery

Backup protection for major disruptions

**On-Demand** or **Automated** backups protect data from loss with 35 Day retention

---

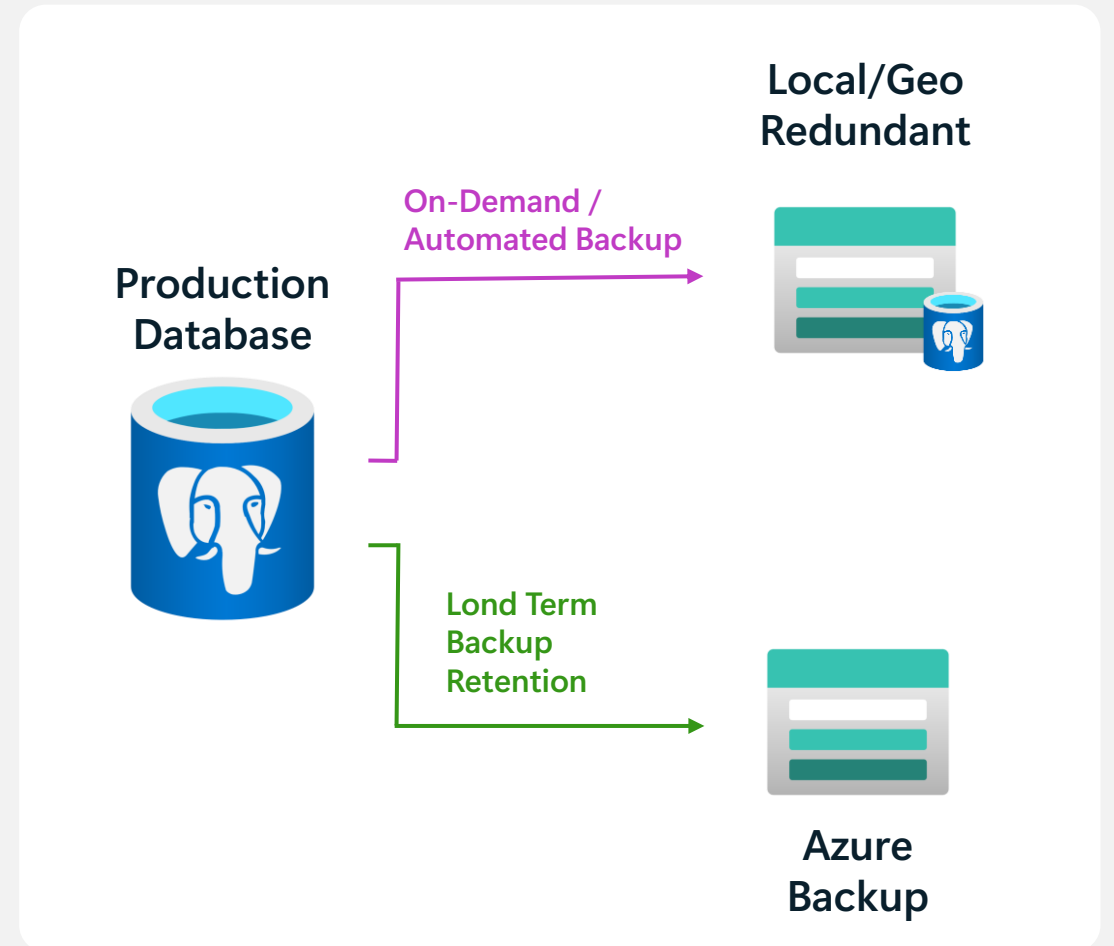
Configure backups for **geo-redundant storage**, and restore to different regions in event of failure

---

Long-term backup retention for **up to 10 years** supports compliance requirements for data retention

---

RPO of 5 minutes



# Disaster Recovery

## Cross-region Failover

**Planned** or **Unplanned** failover to replicas running in any other Azure region

---

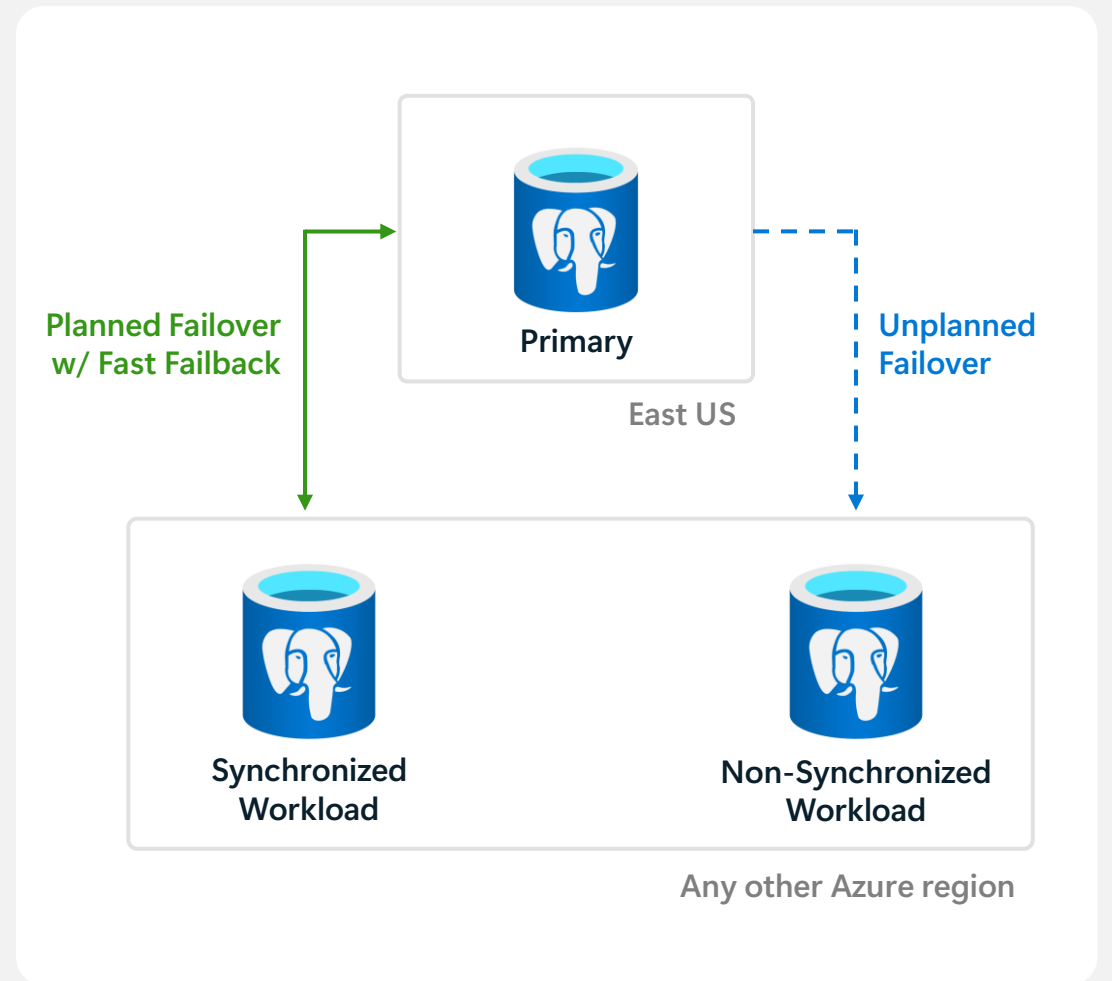
Execute a **Planned Failover** to **synchronize data between regions** prior to failover

---

Execute **Unplanned Failover** for **high impact emergency events** to move workload immediately without data synchronization

---

Both approaches support Virtual IP with no application changes required



# Disaster Recovery

## Cross-region Failover – Portal Experience

**Planned or Unplanned** failover to replicas running in any other Azure region

---

Execute a Planned Failover to **synchronize data between regions** prior to failover

---

Execute Unplanned Failover for **high impact emergency events** to move workload immediately without data synchronization

---

Both approaches support Virtual IP with no application changes required

## Promote

✕

Promote this read replica. Choose to make it the primary server or an independent server.

[Learn more](#)

### Replica server to promote

Server

rmendizabal-  
testreplica01.postgres.database.azure.com

Action

- Promote to primary server.
- Promote to independent server and remove from replication. This won't impact the primary server.

Data sync

- Planned - sync data before promoting.
- Forced - don't sync data, promote as soon as possible.

I understand that this read replica will become an independent server and that its data won't be synced first. This action can't be undone.

Promote

Cancel

# Comprehensive IaC

Infrastructure-as-Code

Use your **preferred IaC provider** to programmatically deploy and manage Flexible Server resources

---

Integrate with existing **infrastructure workflows**

---

**Version control** IaC for robust resource change management



## Terraform

Open-source IaC tool for configuring and deploying cloud infrastructure.



## Bicep

A domain-specific language (DSL) that uses declarative syntax to deploy Azure resources



## Ansible

Declaratively manage your Azure using a simple configuration language



## ARM Templates

Declaratively manage your Azure using a simple configuration language

+ Rest API's, CLI, and PowerShell also

# Elastic Clusters

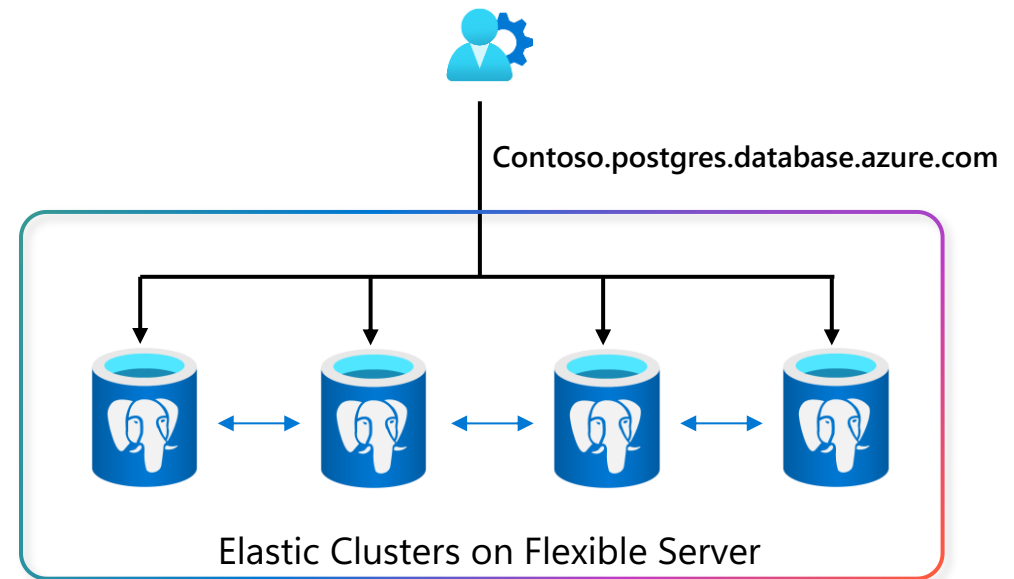
Multi-node Cluster – sharded databased

**Horizontal scaling:** Now scale out to multiple nodes to power multi-tenant and AI applications

**Simplicity:** Offload sharding complexity to a managed service

**Efficiency:** Simplified management of server fleet at no additional cost

**AI ready:** Sharding models scale out vast amounts of vectorized data needed for modern AI applications



- Single cluster endpoint
- Scale a single database horizontally
- Shared nothing architecture
- Powered by Citus extension

# Elastic Clusters – Row/Schema sharding

Public preview



## Flexible Server Cluster

**NODE 1**  
*Coordinator Role*

Metadata Tables & Functions

Shard 1

Shard 2

Shard 3

**NODE 2**

Metadata Tables & Functions

Shard 8

Shard 5

Shard 4


**NODE 2**

Metadata Tables & Functions

Shard 31

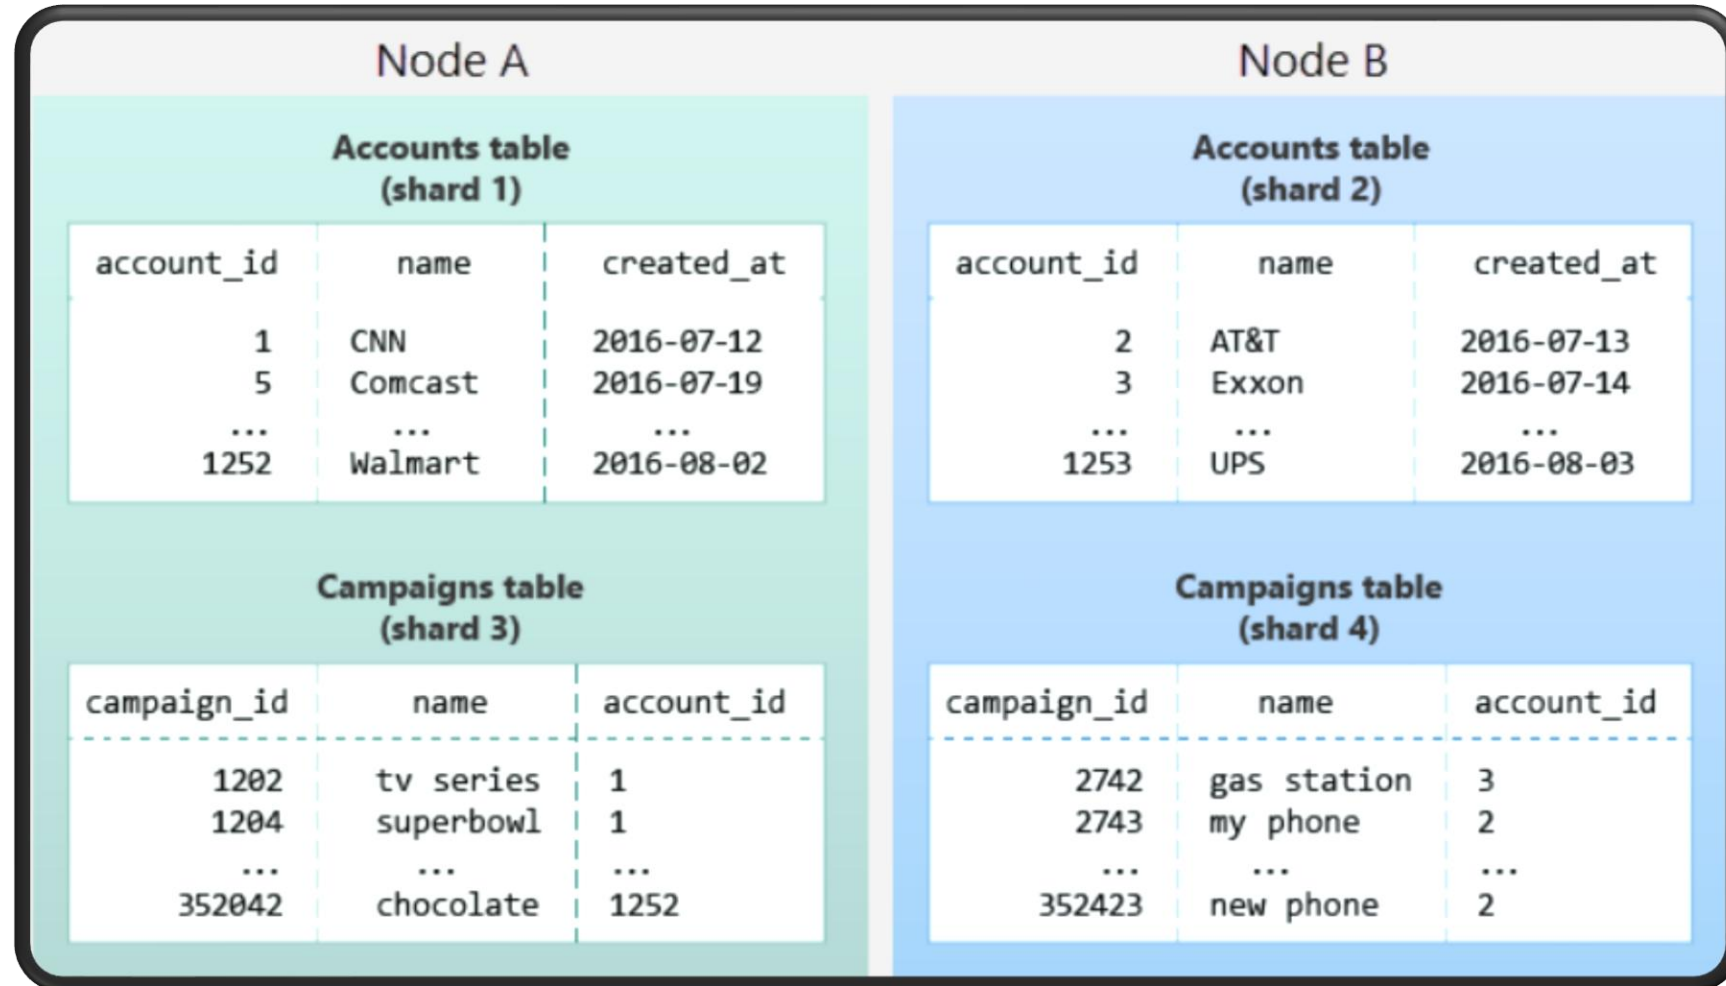
Shard 41

Shard 51





# Row based sharding



# Built in Optimization - Automatic Tuning



## Index Recommendations

Automatically determine optimal indexes based on user activity



## Server Parameter Tuning

Refines server parameters to maximize workload performance



**Minimize Server Resources**



**Maximize Workload Performance**

# Autonomous Optimization

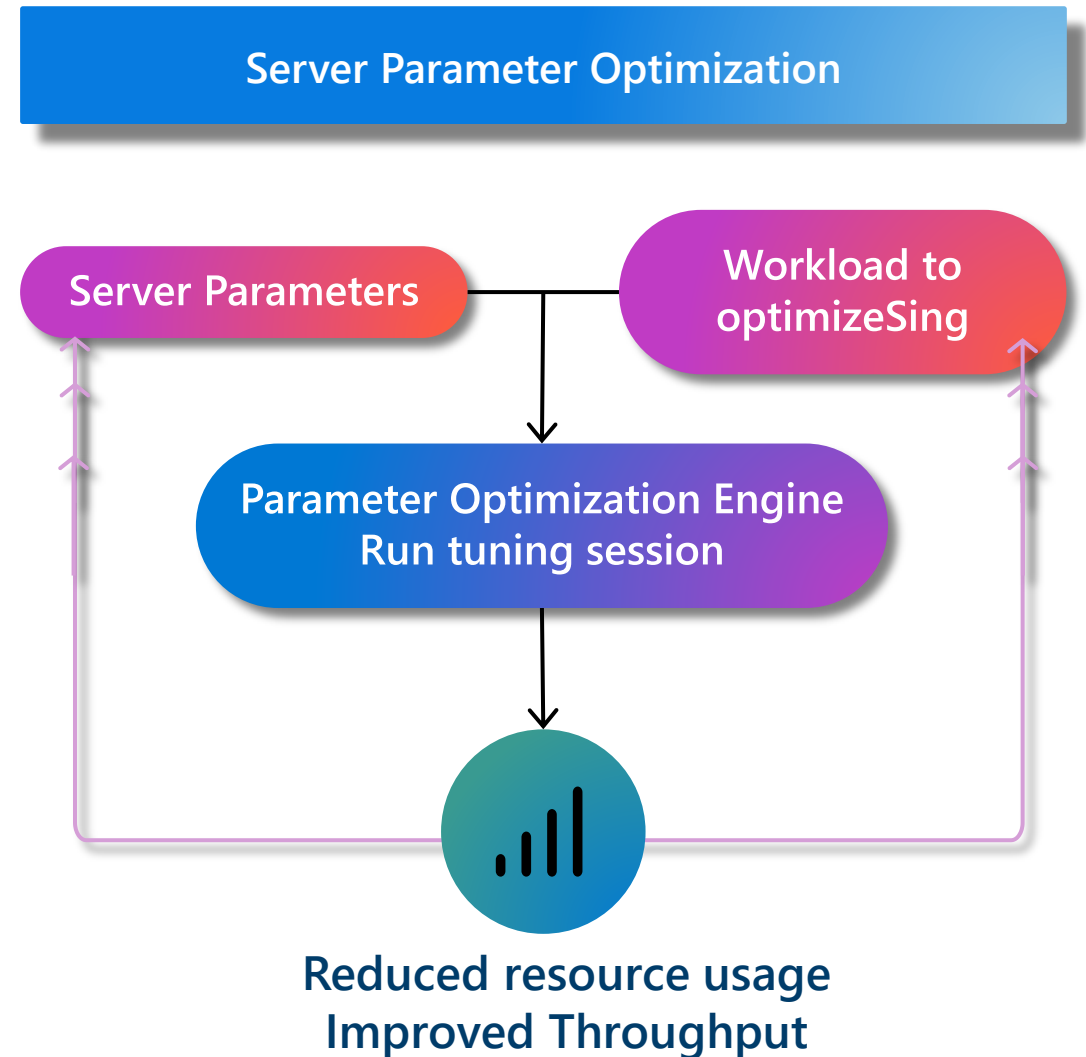
## Index recommendations & server parameter tuning

**Optimize index and server configuration** with autonomous tuning

**Improves workload performance** by analyzing queries tracked by Query Store and providing index recommendations

**Suggests indexes that could be added or removed** during an index tuning session that can improve or reduce performance impact

Identifies indexes not used in a configurable period that could be removed to **reduce unnecessary consumption**



**AI-Ready**

# Generative AI uses-cases on Postgres



Chat history



Retrieval  
Augmented  
Generation (RAG)



Multi-tenant  
AI apps



Real-time  
Recommendations



Real-time  
Anomaly Detection



Multi-Agent  
AI

# Built-in AI capabilities to empower your AI apps

Modern, AI apps on an integrated data platform

## Vector Search



Enable new AI-powered user experiences and chat with your data with pgvector extension and DiskANN for vector indexing and vector search

## Azure AI extension




Build rich PostgreSQL generative AI applications with Azure AI extension integrated with Azure OpenAI, Azure ML, and other AI services.

## In-database embedding




Run faster AI apps by generating embeddings within the database for highly confidential or private workloads.

# How Vectors Work: Important in the AI Era



Vector similarity search empowers Generative AI apps



# Azure Database for PostgreSQL: Native Vector Search

**Open-source Pgvector extension** provides support to store, index, and query vectors for similarity search scenarios

---

Supports multiple vector **distance functions**

---

Enables AI solutions to seamlessly integrate into **existing OLTP Postgres apps** without exporting data to specialize systems

---

Access control, encryption, high availability, disaster recovery all **“just work”**

## Generative AI apps

### RAG (Retrieval Augmented Generation) apps

Retrieve private data to ground LLM model responses

### Recommendation/Semantic Search

Retrieve similar documents by distance between vectors

### Hybrid Search

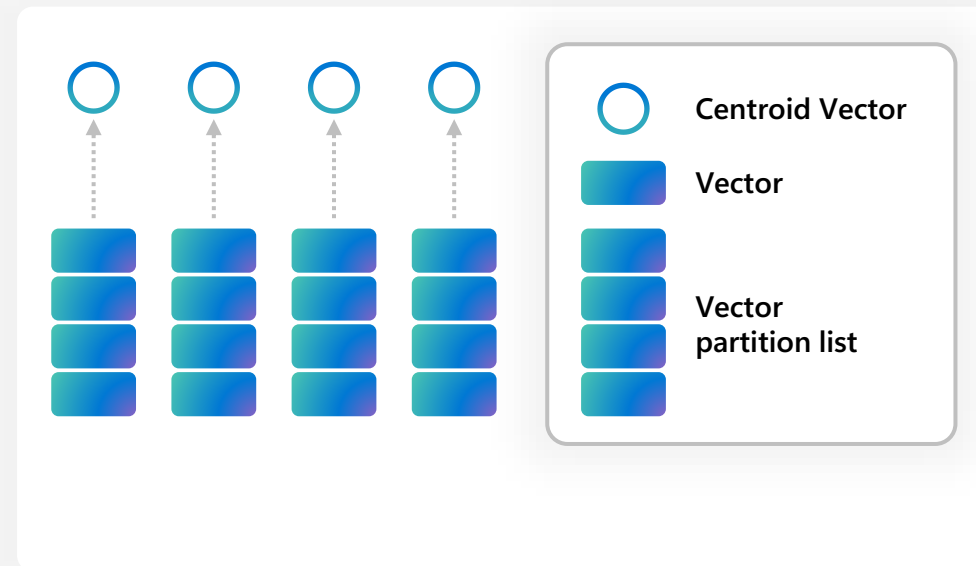
Combine vector search, row filtering, and full-text search



# Vector indexes supported today

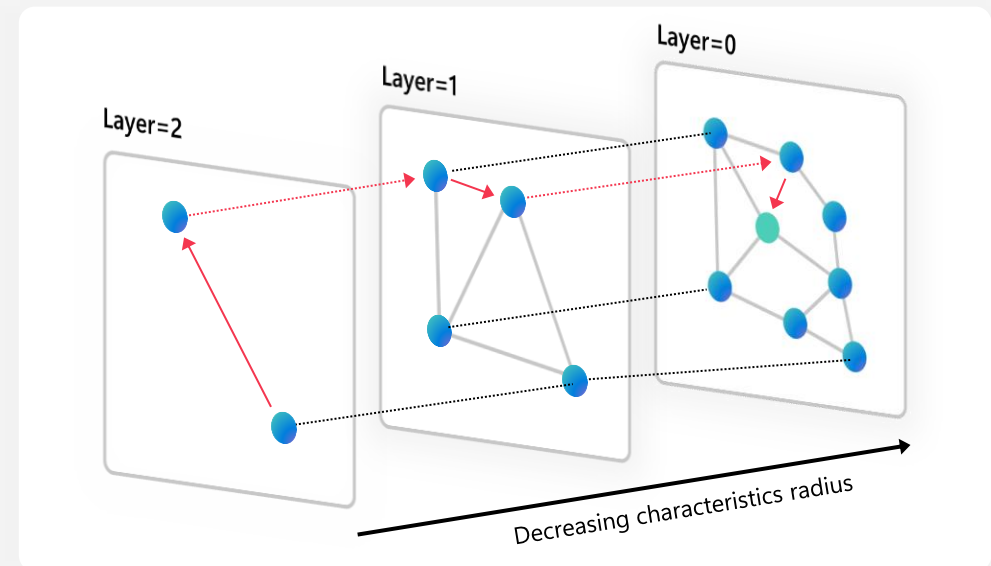
## IVFFlat

- Clusters vectors by applying k-means clustering.
- Memory efficient but requires index rebuilds.



## HNSW

- Builds a multi-layer graph with long and short connections between the vectors.
- The graph can be incrementally updated.



# DiskANN Vector Index

Preview



Highly performant, scalable, and accurate index for vectors

Superior to native PGVector IVFLAT & HNSW index types


Reduced memory footprint by storing vectors on SSD

Compression and quantization improve speed and accuracy of vector search

Accuracy retained as data changed

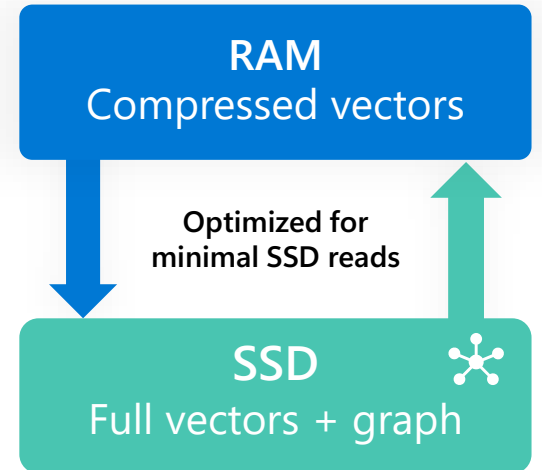
## Vector compression

Large Vectors  
{ D1, D2, D3, D4, D5, ..., D99, D100 }



Compressed Vectors  
{ D1, D2 .., D10 }

## Optimized storage



# Semantic ranker improves recall accuracy

Azure Database for PostgreSQL [Semantic Ranker Solution Accelerator](#) is now available

---

Semantic ranker goes to the **text-level**, providing deep analysis of **semantic relevance** between two text strings

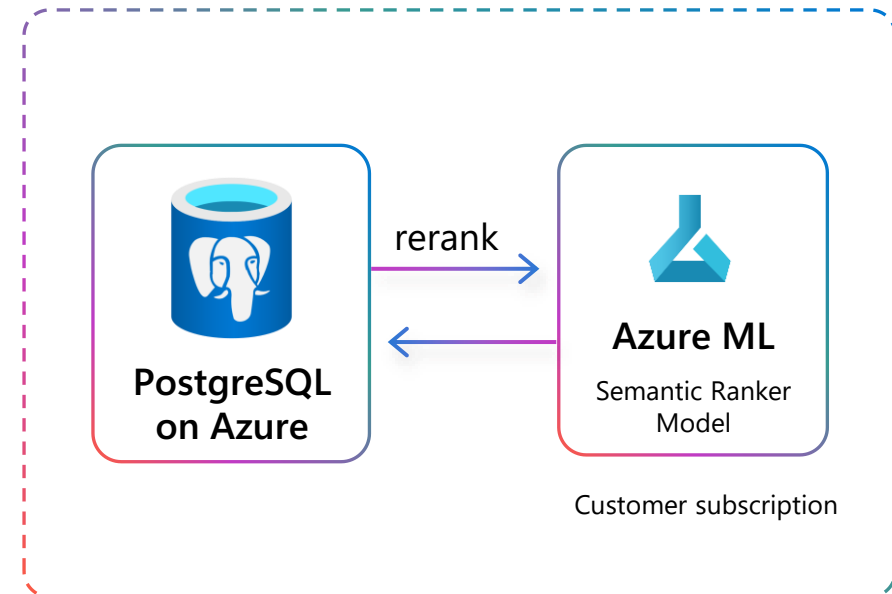
---

**SQL Integration** makes use of azure ai extension to make remote calls to the **Azure Machine Learning** model

---

Perform semantic ranking **directly in the SQL query language**

## Semantic Ranker Solution Accelerator for PostgreSQL



# GraphRAG Solution Accelerator for Postgres

Preview

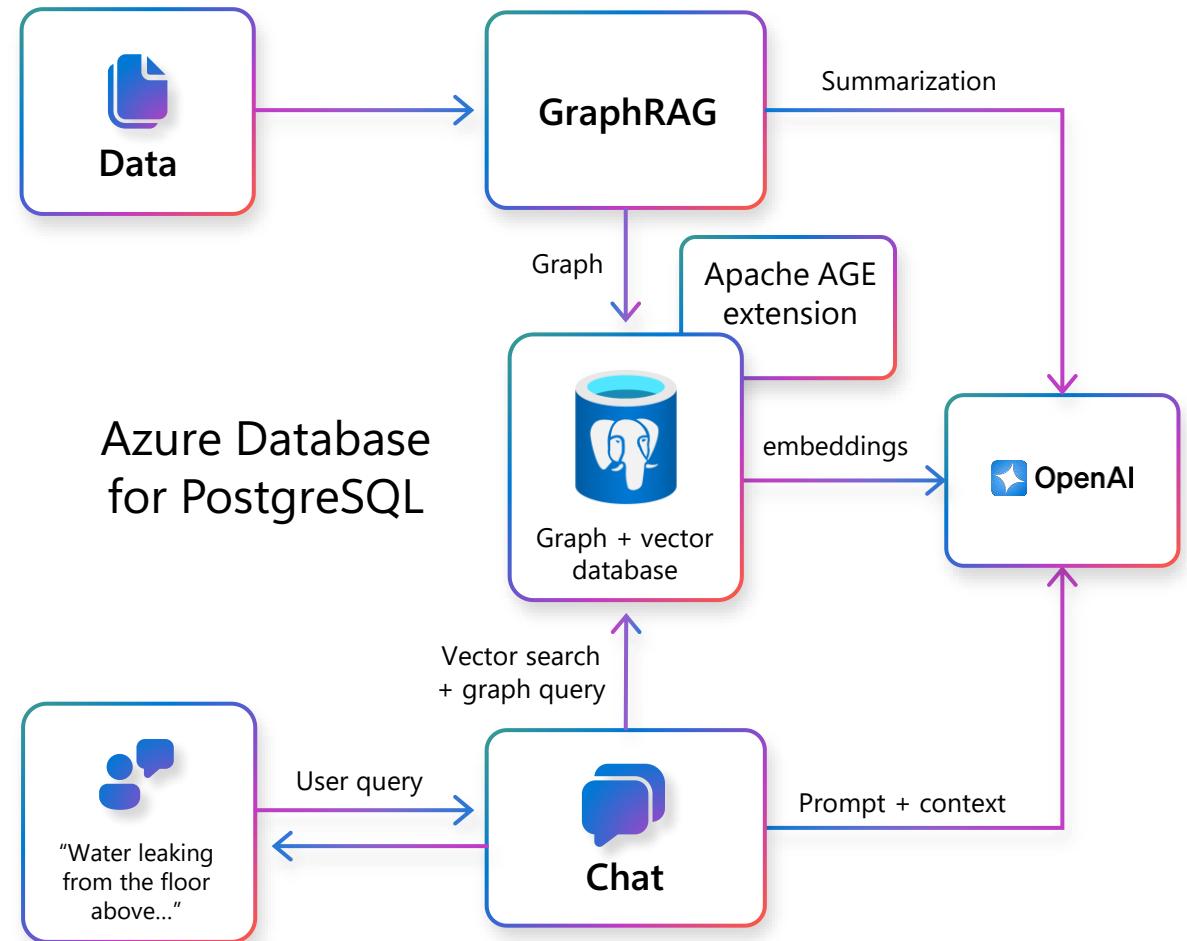
**GraphRAG** is an advanced RAG technique from Microsoft Research to **improve the quality of RAG system processes**

Stores graph natively in PostgreSQL using the **Apache AGE extension**

Consists of three basic steps:

- 1 Graph extraction
- 2 Entity summarization
- 3 Graph query generation at query time

The Solution Accelerator for GraphRAG is [now available](#)



# AI Services integrated with Azure Postgres

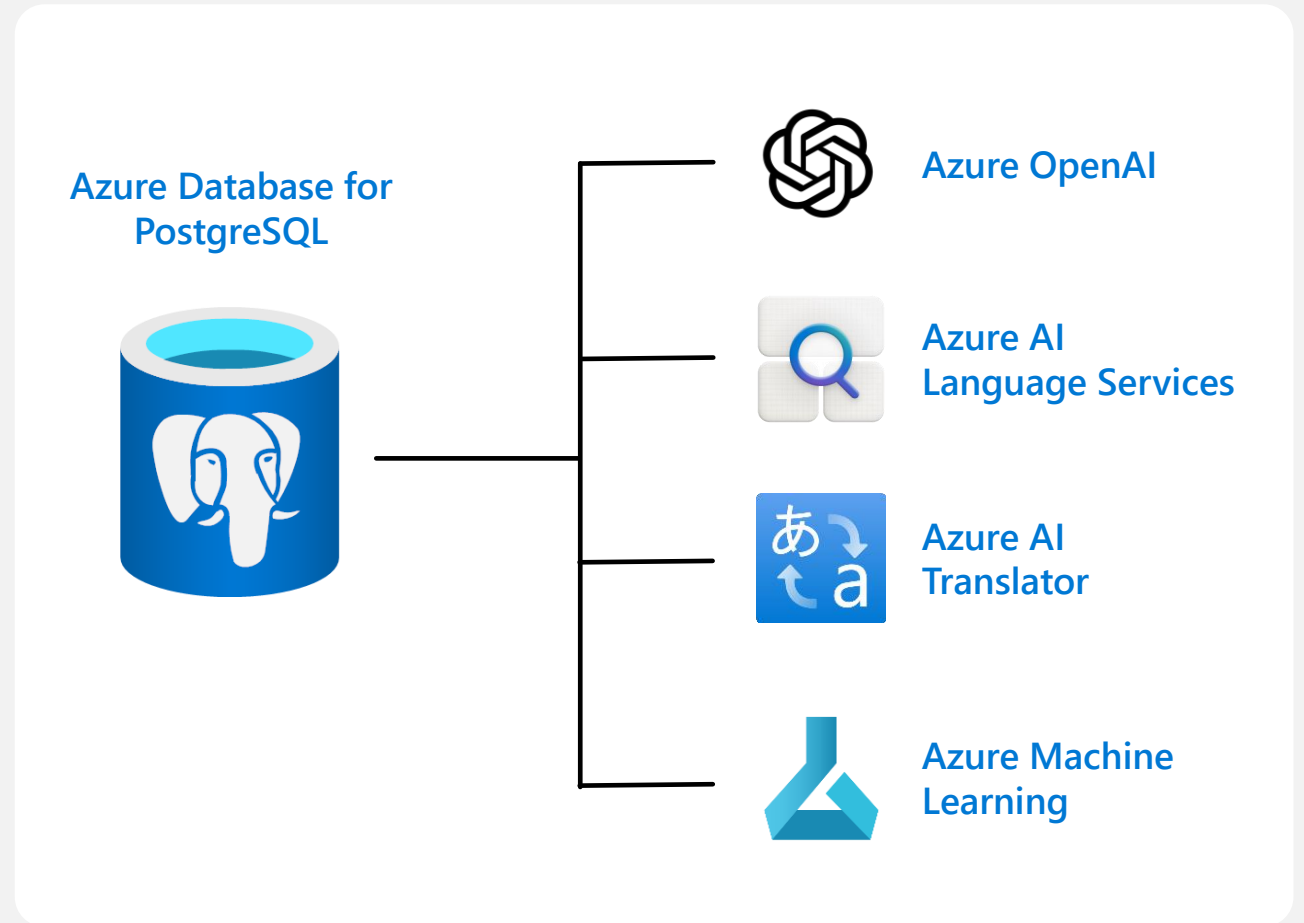
Make remote calls directly from PostgreSQL

The `azure_ai` extension provides a SQL-based interface to integrate with AI services

Supports:

- Azure OpenAI
- Azure AI Language Services
- Azure AI Translator
- Azure Machine Learning

Enables developers to **rapidly integrate AI capabilities** into their app without complex re-architecture or refactoring



# In-Database Embedding Models

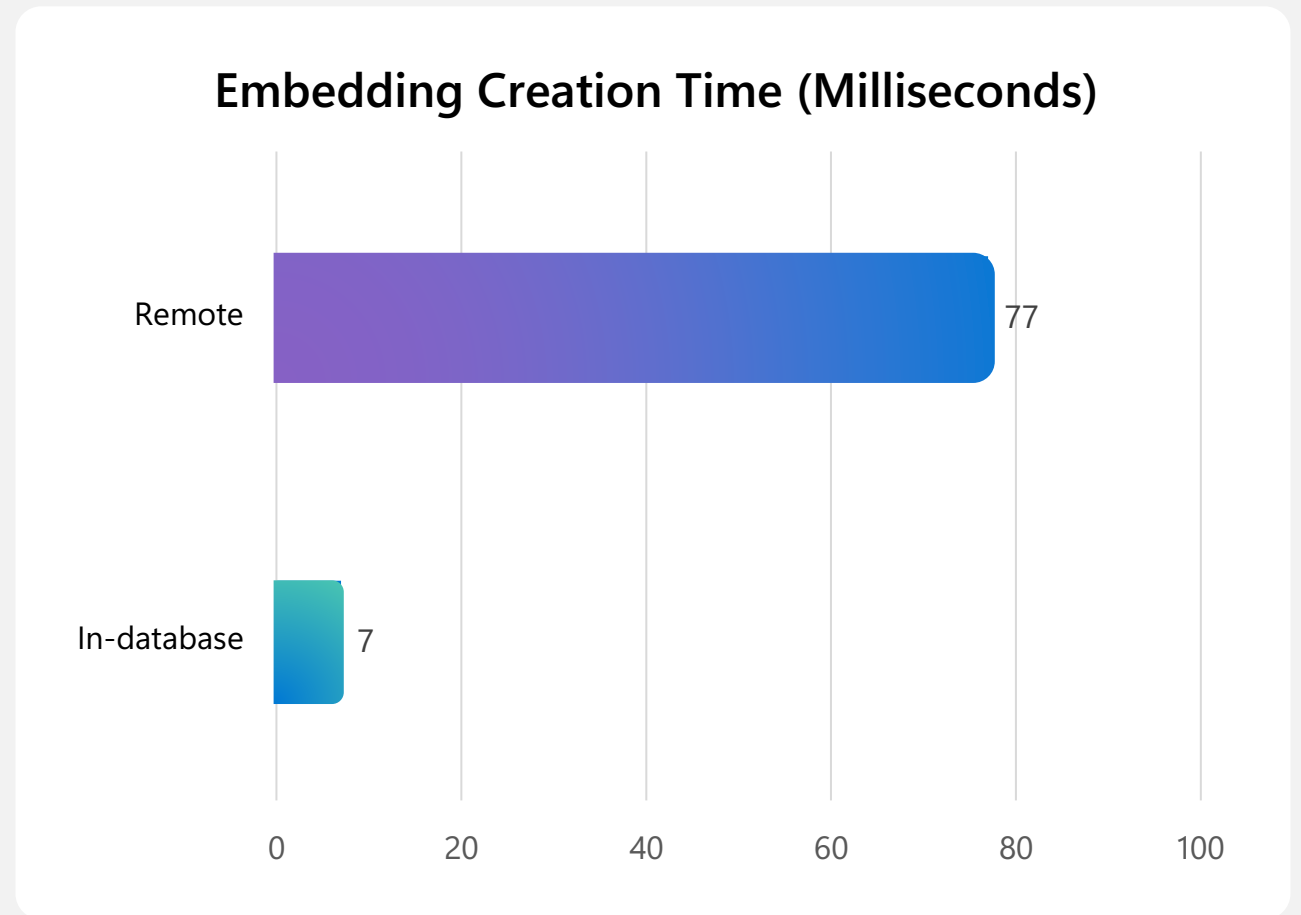
Low-latency embedding creation for OLTP workloads

The `azure_local_ai` extension enables vector embeddings to be generated locally within the Postgres server

Based on the Microsoft **open-source E5** embedding model

Benefits:

- **~10x faster** creation time
- No external service setup, maintenance, or transaction costs
- All data remains **within Postgres**
- Perfect for workloads where the **underlying data changes frequently**

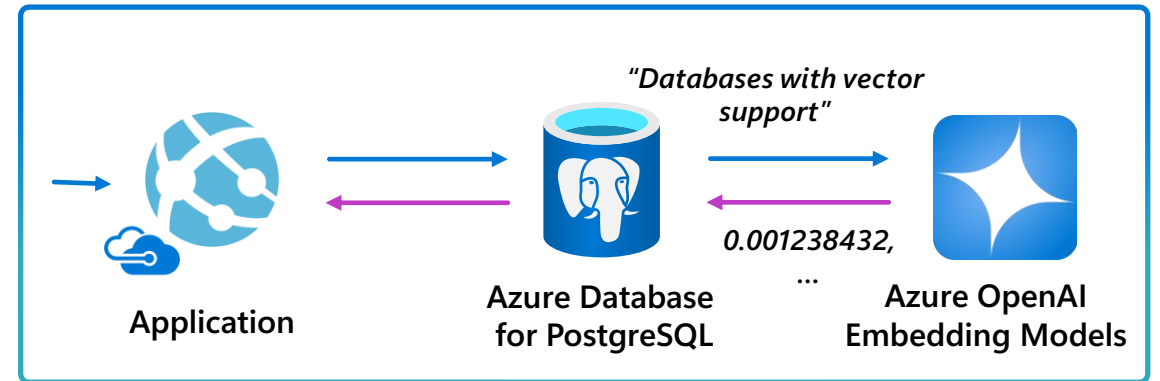


# Vector Generation

## Unique Remote + In-Database Embedding Models

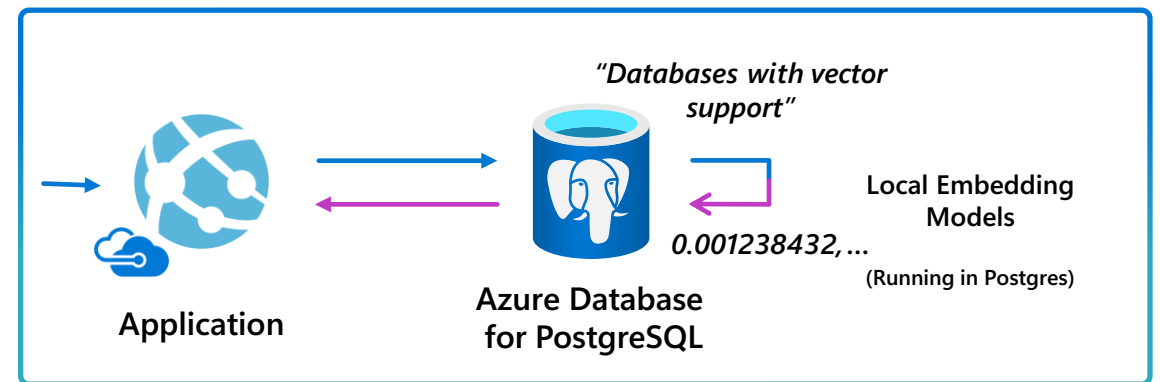
### Remote Embedding Models

```
SELECT * FROM <table>
ORDER BY
database_description <->
azure_openai.create_embeddings(
'text-embedding-ada-002',
'Databases with vector support')
```



### In-Database Embedding Models (Preview)

```
SELECT * FROM <table>
ORDER BY
recipe_embedding <#>
azure_local_ai.create_embeddings(
'multilingual-e5-small:v1',
'Databases with vector support')
```



# Azure Copilot for PostgreSQL Flexible Server

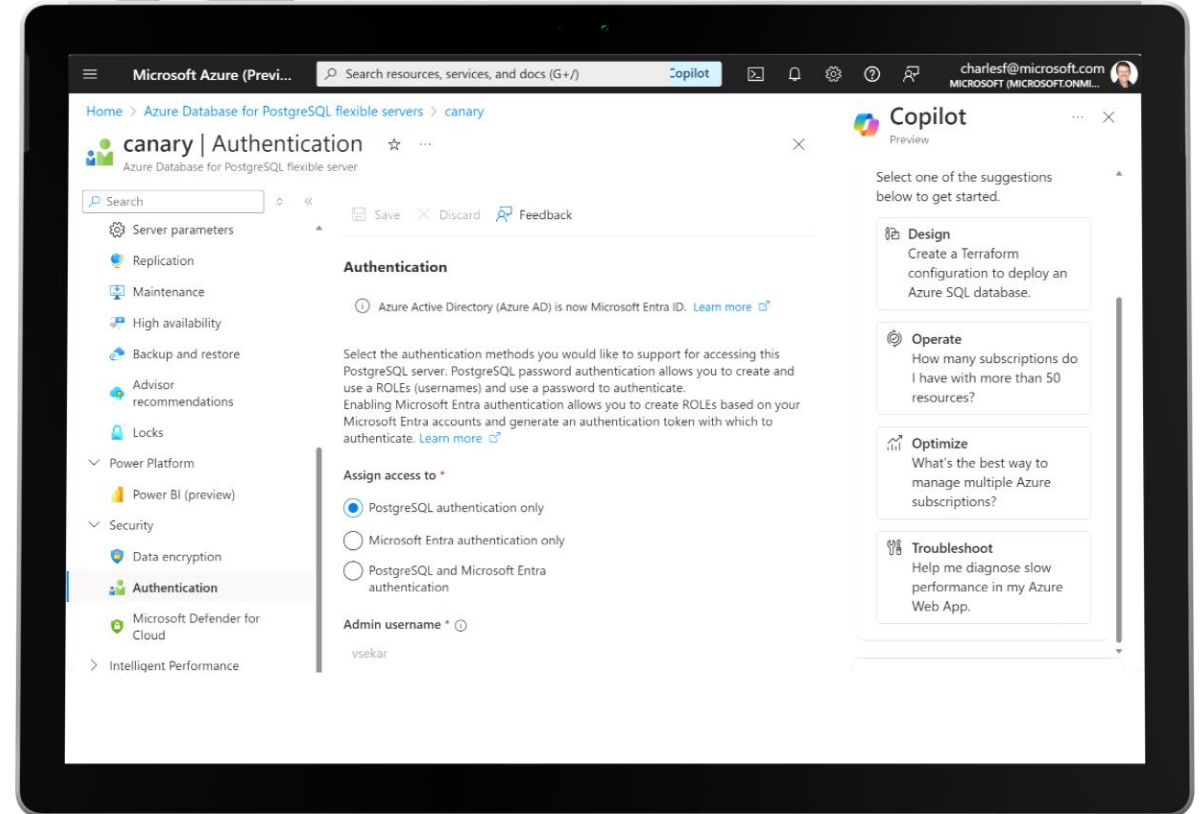
Expert advice when you need it

Coming Soon

Chat-based interface for **analyzing or debugging** the Flexible Server with the Azure portal

Contextual awareness of the individual server


Knowledge based on both **Flexible Server and Postgres community documentation** to provide service and database engine specialization





# Migration resources

# PostgreSQL Migration has never been easier



100% compatible, built on open source Postgres

---

Major versions of PostgreSQL available within weeks



No application code changes needed to migrate

---

List of **supported extensions**



Free migration service built into Azure Database for PostgreSQL

---

Migrate from directly within the database with one-click

***"With Azure, we've doubled if not tripled the performance of our applications. We were able to handle 200,000 transactions per second during our performance tests. We decreased latency by 50% across the networks, all the while decreasing infrastructure costs by 40%, amounting to millions in Turkish Lira. Conversions also increased by 20%, and Boyner saw a 190% rise in revenue from online channels."***

- **[Boyner Group \(Microsoft Customer Story\)](#)**

# Comprehensive Pre-migration Validation

Rules-based assessment of source systems to determine compatibility for migration for Azure Database for PostgreSQL

---

No impact to existing source system running

---

Provides compatibility recommendation for Online vs. Offline migration

---

Recommended to run prior to migration to identify and remediate issues

AuthenticationAndConnectivityValidation

SourceVersionValidation –

Server Parameters validation –

SKU validation –

Extensions validation

Schema validation

Collation validation –

Microsoft Entra Id Validation -

AuthPermissionsValidation

**Offline  
Migration  
Checks**

Online Settings Validation

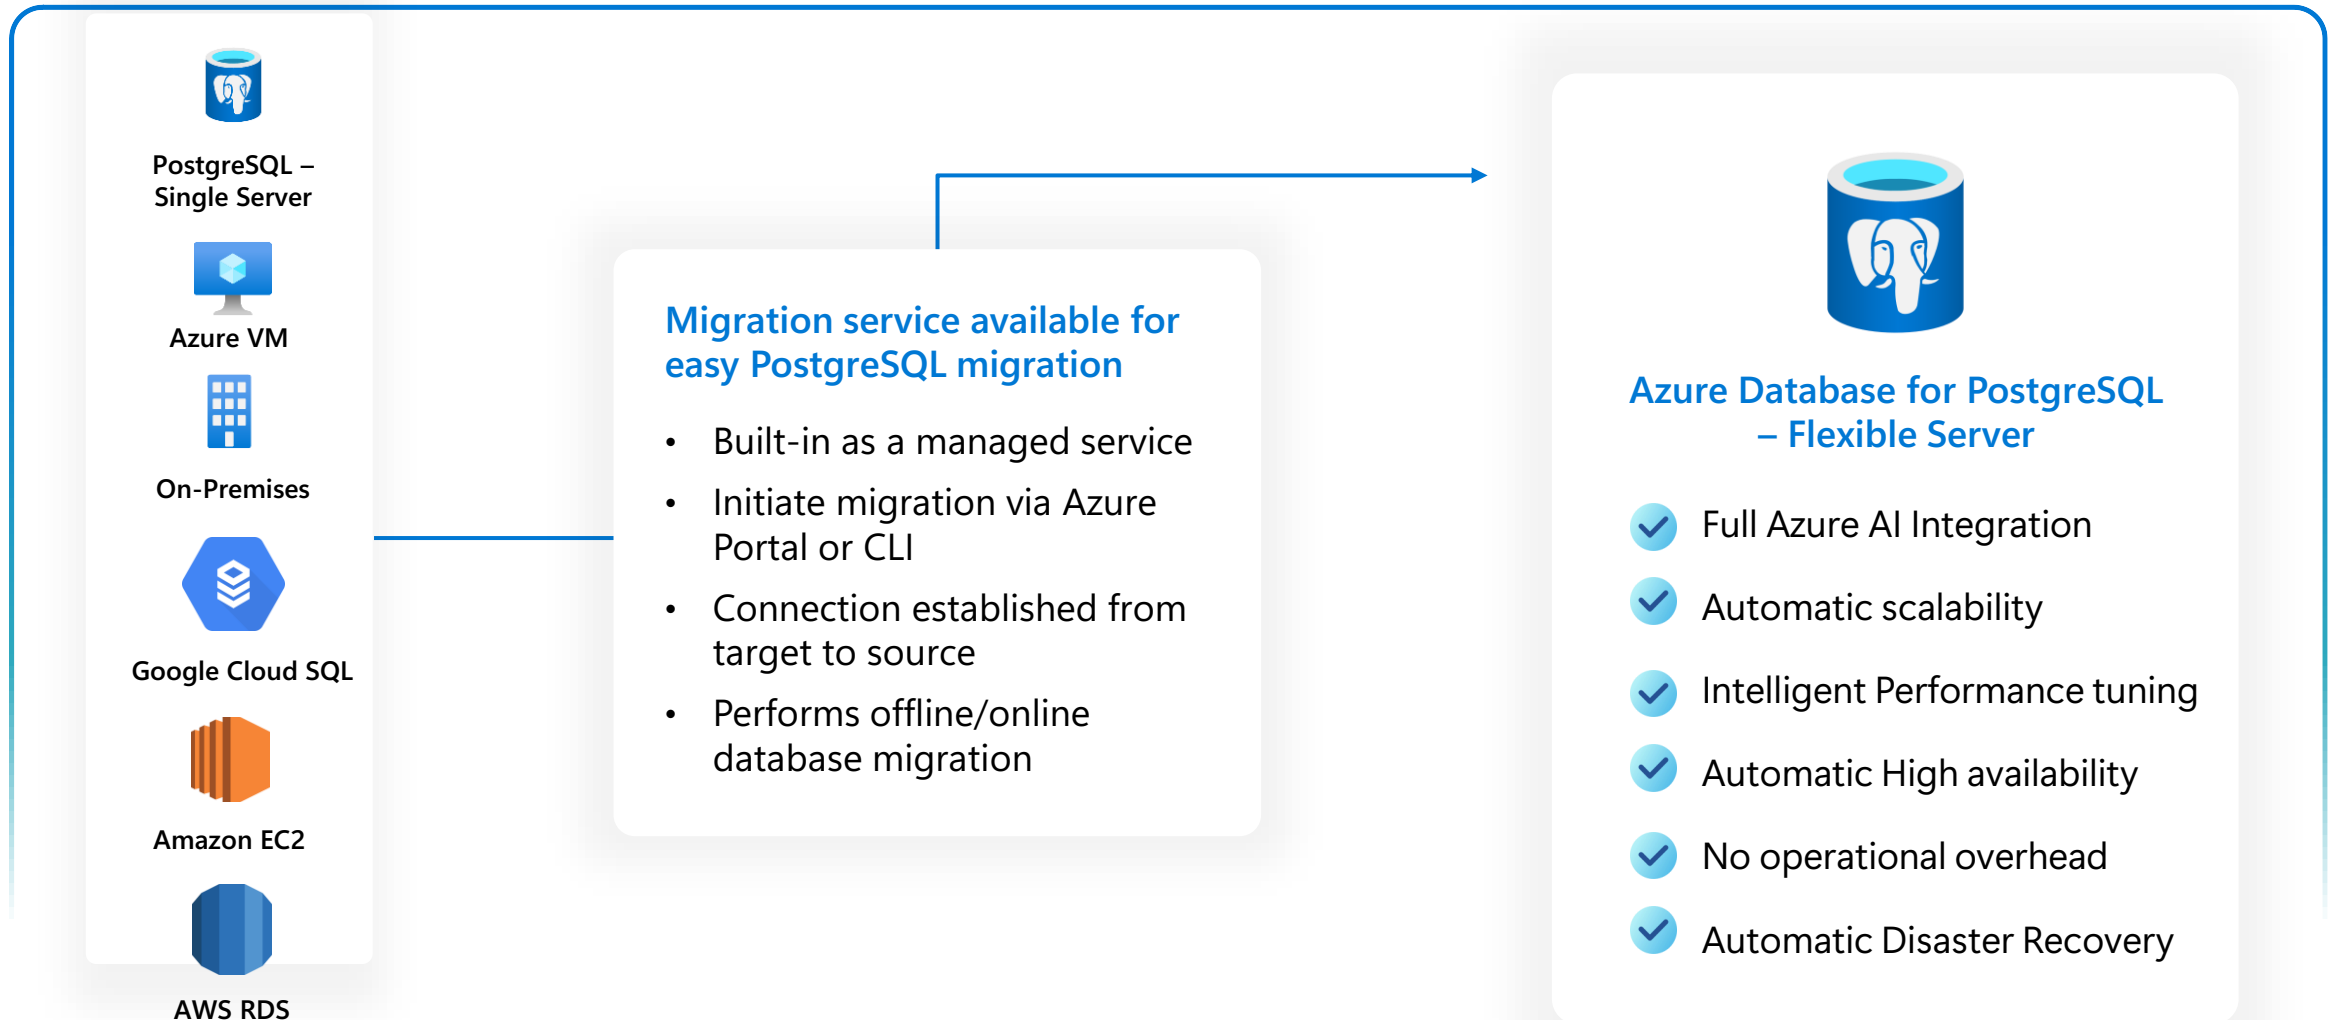
Replication Role Permission Validation:

Missing Primary Keys Validation

**Additional  
Online  
Migration  
Checks**

# Modernize using fully managed, AI-ready PostgreSQL

Migration made easy with built-in tools



# Modes of migration supported using migration service in Azure Database for PostgreSQL

Mode	PROs	CONs	Recommended For
Offline	<ul style="list-style-type: none"><li>• Simple, easy, and less complex to execute.</li><li>• Very fewer chances of failure.</li><li>• No restrictions regarding database objects it can handle</li></ul>	<ul style="list-style-type: none"><li>• Downtime to applications.</li></ul>	<ul style="list-style-type: none"><li>• Best for scenarios where simplicity and a high success rate are essential.</li><li>• Ideal for scenarios where the database can be taken offline without significant impact on business operations.</li><li>• Suitable for databases when the migration process can be completed within a planned maintenance window.</li></ul>
Online	<ul style="list-style-type: none"><li>• Very minimal downtime to application.</li><li>• Ideal for large databases and customers having limited downtime requirements.</li></ul>	<ul style="list-style-type: none"><li>• Replication used in online migration has multiple <a href="#">restrictions</a> (e.g. Primary Keys needed in all tables).</li><li>• Complex to execute than offline migration.</li><li>• Greater chances of failure due to the complexity of migration.</li><li>• There's an impact on the source instance's storage and computing if the migration runs for a long time.</li></ul>	<ul style="list-style-type: none"><li>• Best suited for businesses where continuity is critical, and downtime must be kept to an absolute minimum.</li></ul>

# Migration service in Azure Database for PostgreSQL

## Benefits:

- ✓ Fast, Reliable Migration.
- ✓ Simple UI Wizard in Azure Portal with minimum parameters.
- ✓ Automation via CLI/ARM Templates
- ✓ Comprehensive Migration: Schema, Data, Users/Roles, Privileges, etc.
- ✓ Offline and Online Migration Options Available.
- ✓ Validation Feature to Assess Business Rules and Identify Limitations.

## Cons:

Couple of scenarios are restricted:

- ✗ 4 extensions which broke backward compatibility are restricted and only migrated on case by case basis unless they are being migrated to the same version.
- ✗ Only 8 databases can be migrated at a time through portal (any number of databases can be scheduled through CLI).

# Azure Migrate and Modernize & Azure Innovate



## Comprehensive resources in one place

- Extensive guidance optimized approach from start to finish with assessments, proof of concepts, pilots, tooling, deployment
- Free automated tooling provides you with discovery, assessment, business case analysis, planning, migration, and modernization capabilities
- Proven technical frameworks to help design optimized workloads with security and cost recommendations built throughout



## Direct access to experts and funding

- Access to validated, specialized partners with advanced capabilities to help with all stages from planning to deployment
- Benefit from funding for specialized partners to help offset your project costs



## Extensive coverage - from migration to innovation

- End-to-end coverage including migrating or modernizing Windows Server & SQL Server, PostgreSQL, MySQL, Linux, Oracle, SAP, HPC, analytics, AI and more
- Microsoft-led delivery for rapid rehost migrations, and specialized partners for more complex workloads.

Learn more!

[aka.ms/AzureHeroOfferings](https://aka.ms/AzureHeroOfferings)



# Azure PostgreSQL Resources



**Azure Database for Postgres homepage**  
<http://aka.ms/postgres>



**Get Started for Free with an Azure Free Account**  
<https://aka.ms/try-postgres-free>



**Azure Database for Postgres Docs**  
<http://aka.ms/postgresdocs>



**Azure Database for Postgres Blog**  
<https://aka.ms/azurepostgresblog>



**Azure Postgres on X**  
[@AzureDBPostgres](https://twitter.com/AzureDBPostgres)



**Azure Postgres on LinkedIn**  
<https://www.linkedin.com/company/azure-database-for-postgresql/>



**Azure Postgres Feedback Forum**  
<https://aka.ms/pgfeedback>





**Thank you**

