Five Questions to Ask When Moving Postgres to Microsoft Azure

Aislinn Wright | Director of Product Management

Benjamin Anderson | CTO Cloud

July 2021

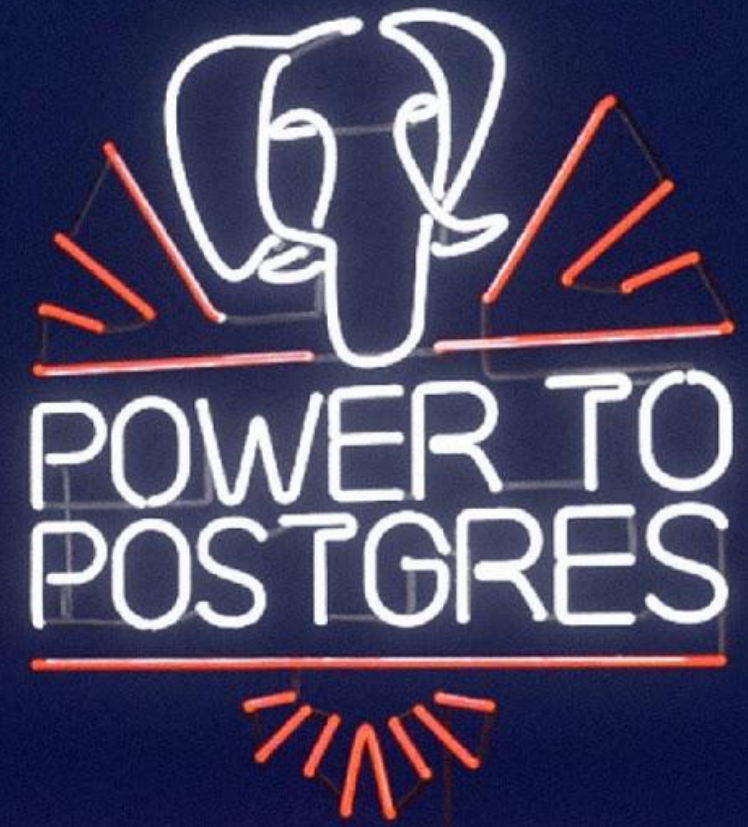

EDB

EDB

Clouds are complicated!

Databases in the cloud
add even more complexity.

We get it.

EDB

Today, we'll explore 5 key questions to consider when moving Postgres to Microsoft Azure.

EDB

1. Security

How are you going to keep your databases secure?

Identity & Access Management

Service Identity

- Azure AD is the default solution for Azure accounts
- Integration with your existing on-premises identity provider

Account Management

- Understand the account hierarchy with directories, subscriptions, and resource groups
- Standardize how subscriptions and resource groups should be used

More auth{n,z} power, more complexity

Who can delete your PostgreSQL? Who can turn off its firewall?

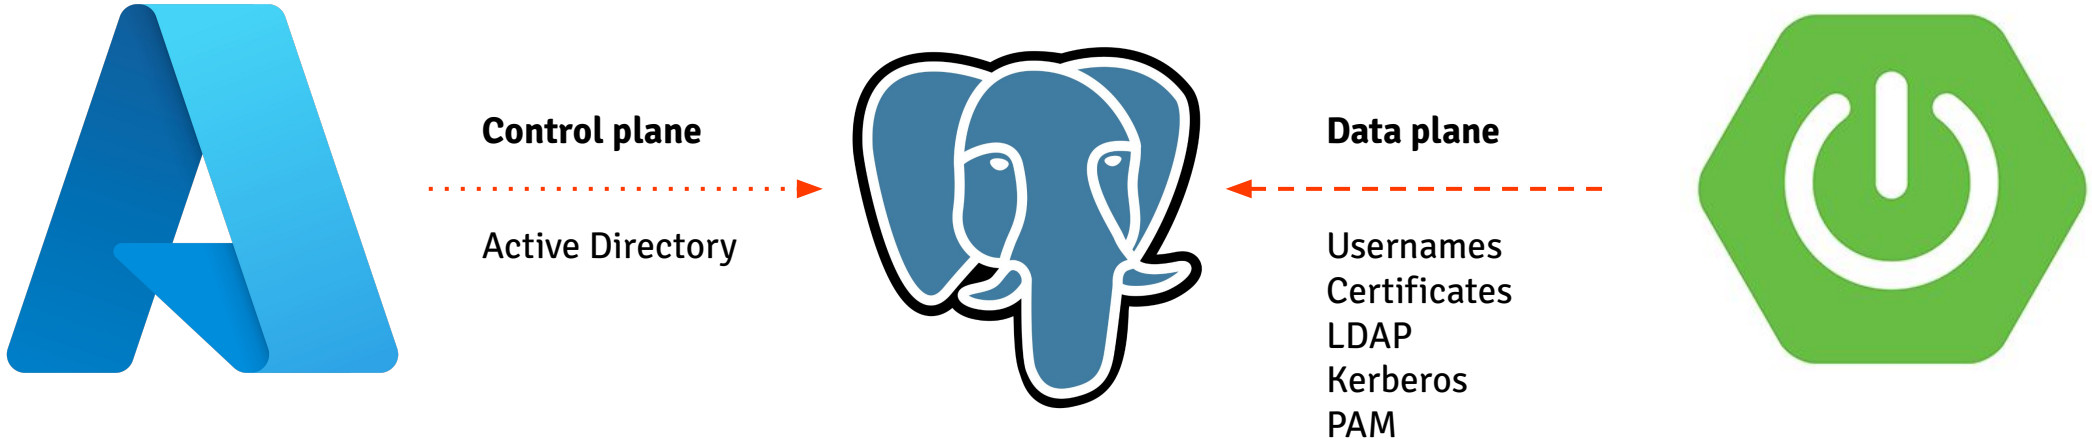

Network Security

Transitioning from a software & hardware based networking setup to a software defined networking architecture

These are all real Azure networking services - and they all have an impact on the security of your system.

Your data is the most important thing to think about ...AND the hardest to get right.

EDB

2. Performance

Are your databases going to slow down when you move them to the cloud?

Performance in the cloud is tied to how fast you want to spend.

Compute selection between on-premises and cloud: gain flexibility, but sacrifice some transparency

E Series Comparison

What are the differences?

Size	vCPU	Memory GiB	Temp Storage SSD (GiB)	Max data disks	Max temp storage throughput: IOPS / Read MBps / Write MBps	Max NICs / Network bandwidth
Standard_E20_v3	20	160	500	32	30000/469/234	8/10000

Size	vCPU	Memory GiB	Temp Storage SSD (GiB)	Max data disks	Burst cached and temp storage throughput: IOPS/MBps	Max uncached disk throughput: IOPS/MBps	Burst uncached disk throughput: IOPS/MBps	Max NICs/Expected network bandwidth (Mbps)
Standard_E20s_v3	20	160	320	32	30000/469/234	8/10000	40000/1000	8/10000

D Series Comparison

What are the differences?

Size	vCPU	Memory GiB	Temp Storage SSD (GiB)	Max data disks/throughput: IOPS	Max temp storage throughput: IOPS/Read MBps/Write MBps	Max NICs	Expected network bandwidth (Mbps)
Standard_D15_v2	20	140	1000	64/64x500	60000/937/468	8	25000

Size	vCPU	Memory GiB	Temp Storage SSD (GiB)	Max data disks	Max cached and temp storage throughput: IOPS/MBps (cache size in GiB)	Max uncached disk throughput: IOPS/MBps	Max NICs	Expected network bandwidth (Mbps)
Standard_DS15_v2	20	140	280	64	80000/640 (720)	64000/960	8	25000

Why should you care about the compute comparisons?

The slight differences could have **meaningful** impacts on your workloads.

Storage

Storage Types

Azure Disk Storage

Azure Blob Storage

Azure Data Lake Storage

Azure Files

Azure NetApp Files

Data Box

Microsoft Azure Confidential Ledger

Premium SSD
Ultra Disks

Comparison between Premium SSD & Ultra Disks

Type	Disk Size	IOPS	Throughput (MB/s)	Price per Month
Premium SSD	1024	5000	200	$135.17
Ultra Disks	1024	5000	200	$440.73

You can tune the knobs on an ultra disk to save money, but it's hard to determine what the right level is.

Are you using local SSDs today?

Network attached storage will likely be slower

- Local NVMes are coupled to certain instance types (LSv2)
- Local NVMe disks are ephemeral, you'll experience data loss if you stop a VM
- Tuning DB for NVMe compared to local SSD requires different database parameters, you cannot swap one out for the other

Databases are special, of course

Performance tuning anywhere - especially in cloud - is a lot harder for databases than it is for applications - stateless applications can be scaled up and down much faster, easier to best-fit, easier to capacity plan

Don’t underestimate the difficulty of dealing with state!

EDB

3. Availability

What sorts of failures do you need to tolerate?

Focus on your RPO and RTO requirements as a baseline for decisions related to performance.

RPO: How far back?

RTO: How long till recovery?

Let’s walk through an example...

Two physically isolated sites, database replicated between 2 sites with synchronous rep. Zero RPO and near zero RTO.

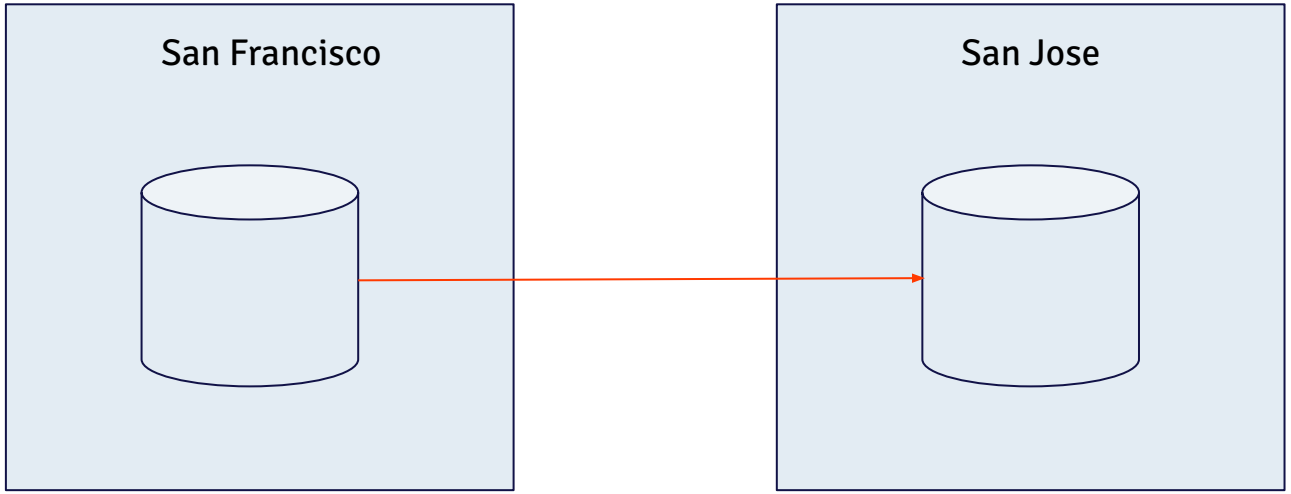

This can be replicated in the cloud. And if you’re not there, you can get there with the cloud.

Take Advantage of Azure Global Infrastructure

- **Use 2 Availability Zones within a region**
  - Get physical and power grid isolation
  - They will be geographically very close so you can tolerate synchronous replication
  - Where's the leader?

- **Want to get more sophisticated?**
  - Store backups in multiple regions
  - Use EDB's BDR (Bi-directional replication) product to run cross-regional active/active database workloads

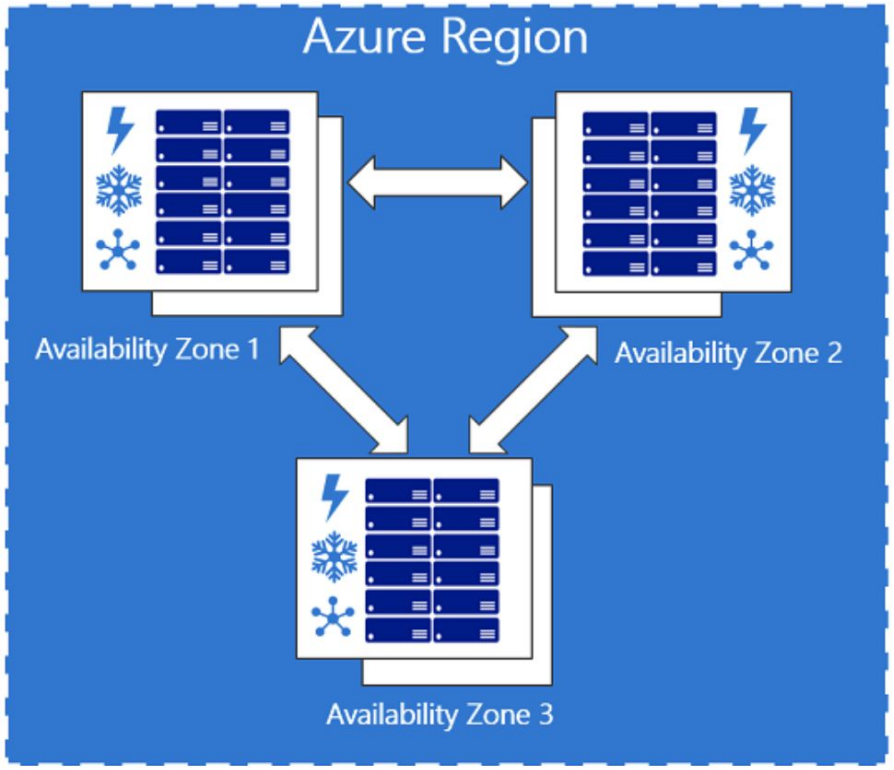

Credit: Microsoft Azure AZs

Cross-Region Considerations

- Your app devs will think this is great and go wild with it because cloud is easier to take advantage of for application developers.
- It can be expensive to fail a database across regions

EDB

4. Cost

What should you be aware of running Postgres on Azure?

Cloud is *expensive* and *transparent*

Spend is visible, thus you should expect increased cost pressure

- Storage costs
  - Ultra disks are charged based on the amount of storage, IOPS, and throughput, unlike premium disks where each disk comes with a guaranteed amount of IOPS and throughput
  - Be cognizant of the tradeoffs between increasing IOPS and throughput and the highers costs
- Networking
  - Intracontinental and intercontinental data transfer fees
  - In 2022, Azure will start charging for cross AZ traffic…a reminder Microsoft can change the pricing!
- Managed Services
  - For Azure Database for PostgreSQL, IOPS are constrained by the VM type meaning you may need to provision a larger instance to achieve a desirable amount of IOPS

EDB

5. What about database as a service?

You are probably not in the database business, database operations are undifferentiated heavy lifting.

Database as a service could be the right fit for you.

Azure's Managed PostgreSQL DBaaS

Azure Database for PostgreSQL

- 3 flavors: Single Server, Flexible Server, Hyperscale
- Different properties and capabilities between flavors
  - Version support drastically differs
  - Supported sizes for databases
  - Different SLAs

How can EDB help?

EDB is the PostgreSQL company.
We support our customers in any environment.

EDB Cloud

Database as a service from the Postgres experts

Postgres Expertise

EDB's expertise goes above the infrastructure; we steer the database roadmap and patch its bugs

Oracle Compatibility

Leave Oracle and further your cloud journey with a fully managed Postgres service

Continuous availability

High availability of your PostgreSQL clusters so you're always on, always available

Single pane of glass for Postgres, everywhere

EDB

EDB Cloud

Preview launching in August

Sign up for the preview!
edbcloud@enterprisedb.com