

Today we are going to talk about the **Azure Cache for Redis** service.

Given the name, I guess it is based on Redis ?!

Exactly.

Azure Cache for Redis is a managed service from Microsoft that relies on Redis, which is an **in-memory key/value** data store.

One of its best-known roles is to **speed up search results** on a website.

Azure Cache for Redis offers either **Open Source Redis (OSS Redis)** or a commercial product, **Redis Enterprise**.




Why is data stored in memory?

Storing data in memory allows them to be transmitted more quickly than if they were stored in a database (DB) or on a disk.

So it's a **low latency data storage system**?

Absolutely.

This makes it a service widely used by applications that require **low latency** and **high storage throughput**.



And can we couple it with a database system (DB)?

No problem.

It can be used as a standalone service or alongside other database services like SQL Database for example.

next

Azure Cache for Redis **improves application performance** through the various use cases it offers.

Obviously, only the most used data should be stored, mainly for cost reasons.

Azure Cache for redis can also be used for caching **static content** that rarely changes.

But also, **user sessions** such as for online shopping carts for example

It is therefore necessary to clearly identify the type of information to be stored in Azure Cache for Redis.

You're absolutely right.

This service is also used to store **jobs, messages in queues**, or for **distributed transactions**.

I assume that Azure offers different service tiers?

Obviously.

You have the **Basic** tier in which OSS Redis runs on a VM, without any SLA. It's perfect for development or test phases.

The **Standard** tier in which OSS Redis runs on 2 replicated VMs.

The **Premium** tier in which OSS Redis runs on more powerful VMs to deliver better performances.

The **Enterprise** tier with Redis Enterprise edition which offers additional Redis modules and provides even greater availability.

And finally the **Enterprise Flash** tier, with Redis Enterprise, with larger storage capacities on SSD disks to reduce the cost of the service.



You must therefore know what you want to do before choosing your service tier.

As always, you have to think about it a little bit before to start your project.

You can find a comparison between the different service tiers on the Microsoft* site to guide you according to your use case.

Yes, I think it's safer to take a look at it.

Knowing that you can upgrade your service tier at any time, however you cannot go to a lower level, for example from Premium to Standard.

Another constraint, it is **impossible to switch** from an **OSS Redis** version to a version with **Redis Enterprise**.

That seems rather logic to me, but you do well to mention it.



Concerning versions, **Azure Cache for Redis** supports OSS Redis version **4.0.x** and version **6.0.x**.

I specify it because they don't offer the same features.

I can update from version 4 to version 6, if I want?

Yes absolutely, it is supported.

But then you can't go back.

Understood.



Regarding the Enterprise and Enterprise Flash versions, they offer additional modules like **RedisBloom**, **RedisTimeSeries** and **RedisSearch**.

Thank you for the info, I will document myself to know what it brings as features.

Good answer!

next



I assume Azure offers **high availability** and **disaster recovery** mechanisms?

Absolutely.

These mechanisms are related with the service tier you have chosen.

The **Standard Replication** mode replicates the data on the 2 Redis nodes within a DC.

Zone Redundancy mode replicates data across multiple Redis nodes spread across different Availability Zones (ZA).

Geo-replication mode replicates data to a 2nd Azure region.

The **Import/Export** mode, as its name suggests, allows you to import/export data from/to a storage account.

And finally the **Persistence** mode which allows you to save your data in a storage account.

Persistence mode is manual or automatic?

Completely automatic.

There are also 2 persistence options.


The 1st, the **RDB** persistence (**Redis Database**) which saves a snapshot of the cache in binary format. The user defines the frequency of snapshots, which ranges from 15 minutes to 24 hours.

And the 2nd, the **AOP** persistence (**Append Only File**) which records each write operation in a log. It necessarily implies an impact on the performance of your Redis, especially if it is very busy.

Really interesting if you want to restore data in the event of a hardware failure.

Exactly and know that if necessary, we can have up to **3 Redis replicas**, on the Premium tier.

It's very important if we want afford high availability!



In addition to the service tiers, we saw earlier, you can also define the size of your Redis cache.

That is to say the **size of the memory allocated** to your instance.

And on what elements should we base to define its size and potentially scale to another one?

You can base yourself on different options, and if necessary, scale:

First the Redis **server load**. If it is often high, it is necessary to consider either a cluster mode, that is to say, several instances to distribute the load, or else a scaling.

Then, the quantity of **memory space** that you plan or consider to use. We are talking in fact about the size of the cache.

There are also **client connections**, which are limited by cache size. If necessary, you can therefore either use cluster mode or change the cache size.

The **network bandwidth** which may be insufficient to handle the number of requests. It is then necessary to consider a change in the size of the cache.

And today the scaling mechanisms are done manually?

Yes, mainly for technical reasons.



You can do it in 2 different ways:


Either with a **Scale-up**, ie you go to a higher cache size.

Either with a **Scale-down**, by adding one or more replica to your Redis cluster.

Great, I got it!

next





There is also a memory management mode before considering a cache size change.

Redis natively integrates a mechanism to delete certain data, these are called **eviction policies**, there are several:

noeviction: New values aren't saved when memory limit is reached. When a database uses replication, this applies to the primary database

allkeys-lru: Keeps most recently used keys; removes least recently used (LRU) keys

allkeys-lfu: Keeps frequently used keys; removes least frequently used (LFU) keys

volatile-lru: Removes least recently used keys with the expire field set to true.

volatile-lfu: Removes least frequently used keys with the expire field set to true.

allkeys-random: Randomly removes keys to make space for the new data added.

volatile-random: Randomly removes keys with expire field set to true.

And finally, **volatile-ttl:** Removes least frequently used keys with expire field set to true and the shortest remaining time-to-live (TTL) value.

Very ingenious as a mechanism for cleaning up.

In addition to eviction policies, there are 2 other important settings in Redis:


Maxmemory-reserved, which determines the amount of memory allocated to non-cache operations like data replication.

And **Maxfragmentationmemory-reserved** which determines the amount of memory allocated for memory fragmentation.

Thank you Thomas, today I discovered a new service that seems really awesome!



J'❤️ Redis



Thank you!



If you want to continue **learning** in a fun way about the **Azure ecosystem**, and not miss any of our illustrations ...

... Feel free to subscribe at:



<https://aka.ms/grow-una>



<https://tinyurl.com/youtube-growuna>

If you like our work, please share it ;o)

See you soon!

