

[← Blog posts](#)

Azure Storage Performance Showdown (Post 1)

TIPS AND TRICKS .NET DEVELOPMENT, AZURE



Allan Thraen | 3 Years Ago | PDF |



Almost every project has some data you want to persist, then read, search through, update and eventually delete. With Azure there are loads of great possibilities - for example Blob Storage, Table Storage, CosmosDb, SQL Azure. I've decided to do some simple and fairly naive tests to compare these for some typical usage scenarios and see how they perform.


SQL Server used to be the hammer in the toolbox that made us think every problem we saw was a nail. But it's been years since I thought that way. Sure, it's still a great tool for some transactional problems or in situations where it makes sense to explore data in a relational way. More and more often I find myself turning to regular storage, table storage, document db's and NO-SQL when I need to deal with data at scale. Somehow writing SQL today just feel so 1990'ies, right? Just like working with XML doesn't feel nearly as cool as working with json. But that's another story.

However, this made me wonder: in a typical web scenario how do the different services really match up when considering performance, scale - and of course keeping their price in mind? And does it matter how you use them? How will they work if you have 20 workers constantly hitting them as opposed to one? These are questions that sometimes keep me up wondering late at night - so I decided to venture into a little fun experimentation. Luckily the friendly folks at Azure provide each new subscription with ~\$200 which I figure should be plenty to have some fun.

I'm going to do this as a series of Blog posts. In this, first blog post I'll explain the setup and look at some initial results for storing (create) the data. I took a quick Twitter Poll to see what the expected results would be.

Allan Thraen
@athraen

Developers! Which @Azure service do you think it's fastest to write a lot of data-rows to? (blogpost with measurements coming up)



I guess I might not be the only one that has moved on from good old SQL Server?!

I'm not surprised that people have great confidence in **Table Storage**. It's awesome, easy to work with and can take pretty much any amount of data, as the great Troy Hunt has shown again and again:
<https://www.troyhunt.com/working-with-154-million-records-on/>

Azure Blob storage is great for binary files, but my expectation is that it's not in the running, so I'm surprised to see those that think it is.

CosmosDB is the new kid on the block - and it seems to have taken the world by storm. I recently did a prototype with a complete CRUD Episerver Content Provider for Cosmos (that I might blog/share one of these days) and it certainly seemed pretty powerful. One of the things I really like about Cosmos is that you can use any number of API approaches for it - be it SQL (ewwr), Mongo, Gremlin, Cassandra or - as in my case Table. In fact, the Table API for it is so identical to the Table Storage API that I've been able to use the exact same code, just with different connection strings. That's pretty neat. Another cool thing is that you can easily adjust how many Request Units (R/U) you want to be able to use - that's how it's rate limiting works which means that you can scale it up like crazy.

The data

Post Comments 0

[Azure Storage Performance Showdown \(Post 2\)](#)

[Azure Storage Performance Showdown \(Post 3\)](#)

[Storage Performance Aftermath - ElasticSearch Joins the Fight](#)