



Introducing DocumentDB – A NoSQL Database for Microsoft Azure

Ken Casada
kcasada@microsoft.com
Technical Evangelist
Microsoft Switzerland

Definition

What is Azure DocumentDB?

It's a fully managed NoSQL document database service running on Azure

NoSQL

- "Non-Relational" would probably be a better term
- Different types:
 - Key-value store (Azure Table Storage, Redis)
 - column store (Cassandra)
 - graph databases (Neo4)
 - document databases (MongoDB, Azure DocumentDB)
- Built for scale and performance (designed to scale out)
- Great at managing big amount of data
- No enforced schema

Document Databases

Great for these documents...

```
{
  id: "mspress",
  name: "Microsoft Press",
  yearOfFirstPublication: 1984,
  books: [
    { id:"1",
      name:"DocumentDB 101",
      numberOfPages: 545
    },
    { id:"2",
      name:"DocumentDB for RDBMS users",
      numberOfPages: 700
    }
  ]
}
```

Document Databases

Definitely not for these
type of documents...



TechDays

2017 | Baden

Getting started with DocumentDB

Demo

DocumentDB

- Store schema-free JSON documents
 - Searchable using SQL grammar (New: with support for aggregate functions!)
 - Try out the Query Playground: <https://www.documentdb.com/sql/demo>
 - Automatically indexes every property in a document
 - Optional: reduce storage and processing overhead by defining own indexing policy!
(<https://docs.microsoft.com/en-us/azure/documentdb/documentdb-indexing-policies>)
 - includes or exclude properties to be indexed
 - consider using asynchronous (lazy) index updates (default: synchronous)
- Scalability & predictable performance (SLA guaranteed) in terms of:
 - Throughput (RU) @99.99
 - Estimate Request Unit and Data Storage: <https://www.documentdb.com/capacityplanner>
 - Low latency
 - Reads < 10ms @99
 - Writes < 15ms @99
- SLA also for availability and consistency @99.99
 - DocumentDB SLA: [SLA: https://azure.microsoft.com/en-us/blog/azure-documentdb-service-level-agreements/](https://azure.microsoft.com/en-us/blog/azure-documentdb-service-level-agreements/)

DocumentDB

- How to add documents to a collection?
 - a. Use the Azure Portal
 - b. Use the DocumentDB Data Migration Tool in order to import from external data sources:
 - ✓ JSON files
 - ✓ MongoDB
 - ✓ SQL Server
 - ✓ CSV files
 - ✓ Azure Table storage
 - ✓ Amazon DynamoDB
 - ✓ HBase

Reference: <https://docs.microsoft.com/en-us/azure/documentdb/documentdb-import-data>

DocumentDB

- How to add documents to a collection?
 - c. Programmatically
 - REST API
(<https://docs.microsoft.com/en-us/rest/api/documentdb/?redirectedfrom=MSDN>)
 - Client SDK (.NET / .NET core / Java / Python / Node.js / JavaScript)
(<https://msdn.microsoft.com/library/azure/dn781482.aspx>)
 - Server Side Programming Model
(<https://docs.microsoft.com/en-us/azure/documentdb/documentdb-programming>)
 - Stored Procedure, Trigger (do not fires automatically) and UDF
 - Fully Transactional (ACID guarantee)

Using the .NET Client SDK / Server-Side Programming Model

Demo

DocumentDB tunable consistency

- 4 different consistency levels supported (CONSISTENCY vs PERFORMANCE)

- STRONG**

- you are sure all queries return the latest (committed) data → pay in terms of performance (slowest)
 - write operation is only visible once it is committed durably by the majority of replicas
 - Only available when you are deploying DocumentDB within a single region!

- EVENTUAL**

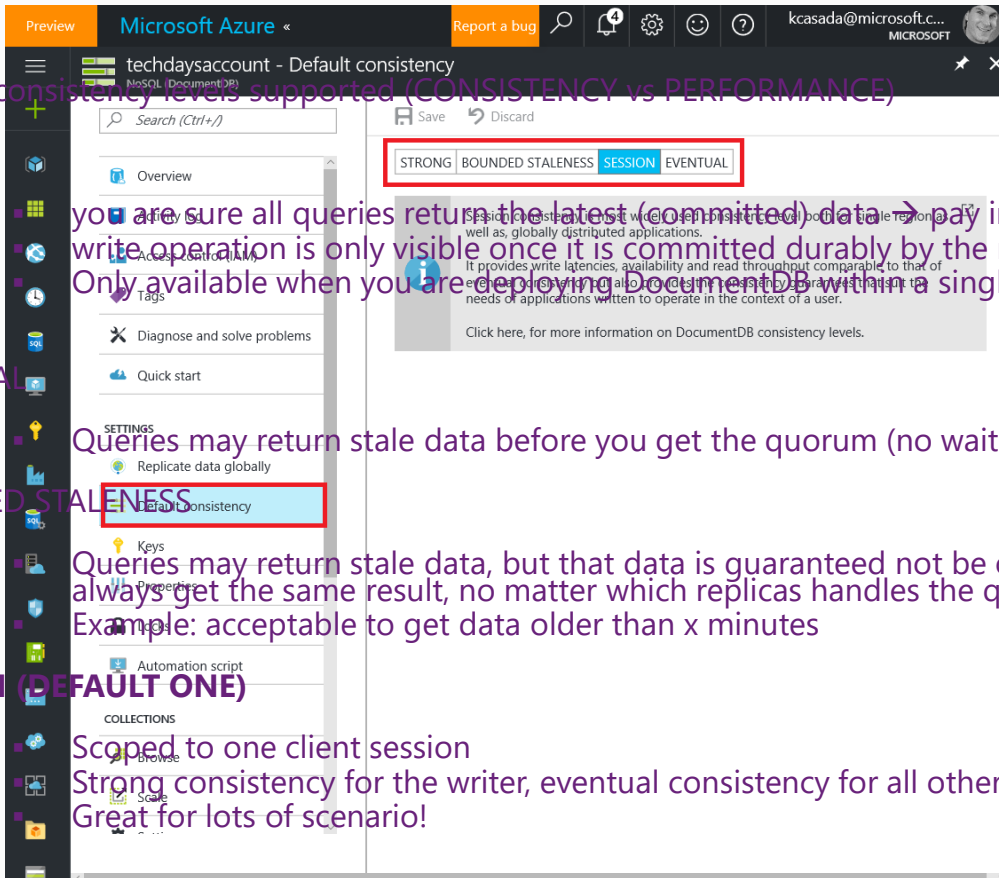
- Queries may return stale data before you get the quorum (no waiting → fastest performance)

- BOUNDED STALENESS**

- Queries may return stale data, but that data is guaranteed not be older than a specified time period (you always get the same result, no matter which replicas handles the query)
 - Example: acceptable to get data older than x minutes

- SESSION (DEFAULT ONE)**

- Scoped to one client session
 - Strong consistency for the writer, eventual consistency for all other users
 - Great for lots of scenario!

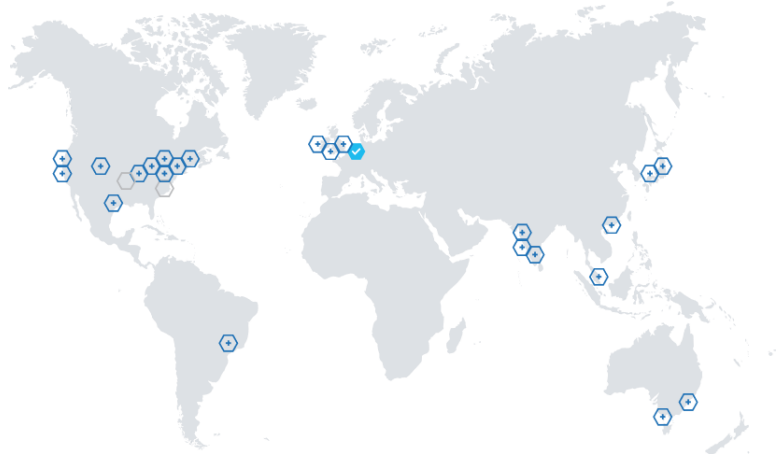


DocumentDB Global Distribution

Search (Ctrl+/)

Save Discard Manual Failover Failover Priorities

Click on a location to add or remove regions from your DocumentDB account.
* Each region is billable based on the throughput and storage for the account. [Learn more](#)



WRITE REGION

West Europe

READ REGIONS

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Quick start

SETTINGS

Replicate data globally

Default consistency

Keys

Properties

Locks

Automation script

COLLECTIONS

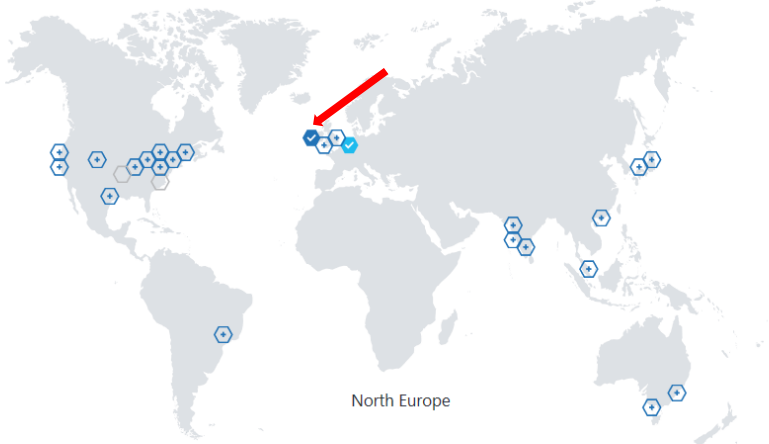
- DocumentDB deployed within 1 single region (all read and writes operations are redirected to the only available region)

DocumentDB Global Distribution

Search (Ctrl+/)

Save Discard Manual Failover Failover Priorities

Click on a location to add or remove regions from your DocumentDB account.
* Each region is billable based on the throughput and storage for the account. [Learn more](#)



North Europe

| WRITE REGION |
|--------------|
| West Europe |

| READ REGIONS |
|--------------|
| North Europe |

The screenshot shows the AWS IAM console interface for DocumentDB. On the left is a navigation sidebar with sections for Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Quick start, SETTINGS (Replicate data globally, Default consistency, Keys, Properties, Locks, Automation script), and COLLECTIONS (Browse). The main content area displays a world map with various regions marked by hexagonal icons. A red arrow points to the 'North Europe' region. Below the map, a configuration table is shown with a red box highlighting the 'WRITE REGION' and 'READ REGIONS' sections. The 'WRITE REGION' is currently set to 'West Europe', and 'North Europe' is listed under 'READ REGIONS'. At the top of the main content area, there are buttons for 'Save', 'Discard', 'Manual Failover', and 'Failover Priorities'. A search bar is located at the top left of the main content area.

- Click on + to add or remove regions
- Click Save → replication takes place (up to 100TBs within 30 min)

DocumentDB Global Distribution

Search (Ctrl+/)

Save Discard **Manual Failover** Failover Priorities

Click on a location to add or remove regions from your DocumentDB account.
* Each region is billable based on the throughput and storage for the account. [Learn more](#)

WRITE REGION

West Europe

READ REGIONS

North Europe

- More regions allows Manual Failover (Read region becomes the new Write region)

DocumentDB Global Distribution

Search (Ctrl+/)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems
- Quick start


SETTINGS

- Replicate data globally
- Default consistency
- Keys
- Properties
- Locks
- Automation script

COLLECTIONS

- Browse

Failover Priorities

Drag-and-drop read regions items to reorder the failover priorities.
Tip: Drag  on the left of the hovered row to reorder the list.

WRITE REGION

North Europe

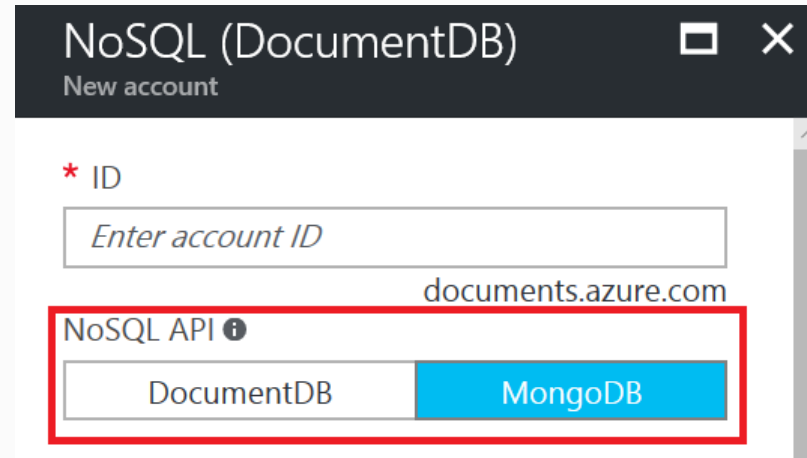
| READ REGIONS | PRIORITIES |
|--------------|------------|
| West Europe | 1 |
| Canada East | 2 |

OK

- Control sequence of regional failover (need at least 3 regions)

DocumentDB support for MongoDB App

- Use existing MongoDB SDKs and tools with DocumentDB
- DocumentDB used as data store for apps written for MongoDB
- No code changes!!!
- <https://docs.microsoft.com/en-us/azure/documentdb/documentdb-protocol-mongodb>
- <https://azure.microsoft.com/en-us/blog/documentdb-api-for-mongodb-now-generally-available/>



DocumentDB: useful Tools

- DocumentDB local emulator (public preview)
 - Local environment that emulates the Azure DocumentDB service
 - No need for an Azure Subscription
 - For dev/test purposes
 - Works on Windows 10, Windows Server 2012 R2, Window Server 2016
 - Download and more information: <https://docs.microsoft.com/en-us/azure/documentdb/documentdb-nosql-local-emulator>
- Azure DocumentDB Studio
 - A client management viewer/explorer for Microsoft Azure DocumentDB
 - Easily browse DocumentDB resources and learn DocumentDB resource model
 - Support Create, Read, Update, Delete (CRUD) and Query operations
 - and much more...
 - Download and more information: <https://github.com/mingaliu/DocumentDBStudio>

Thank you!

© 2017 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.

The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.