

Astra DB

Big Data Systems

Dr. Rubi Boim

Motivation (for this course)

- Managing a cluster of Cassandra is not trivial
Not the focus of this course
- Astra DB is a managed Cassandra DB service
 - Free version
 - We will use it to learn basic Cassandra programming

DataStax

- A commercial company established at 2010
Cassandra was released at 2008

Provides

- Support for Cassandra
- DSE (DataStax Enterprise)
own version of Cassandra with extra features
- Managed Database-as-a-service
based on Apache Cassandra

DataStax Enterprise (DSE)

- Commercial product - NOT open source
an extension of Cassandra

Provides extra features / tools on top of Cassandra

- Cluster management
- GraphQL support
- Analytics
- Search capabilities
- ...

Astra DB

- A managed Cassandra database service
- Launched at 2020
- Support Cassandra +3.11

Astra DB can be configure to run on

- AWS
- GCP
- Azure

Create an account

<https://astra.datastax.com/>

Register



Sign Up With GitHub



Sign Up With Google

OR

First Name

Last Name

Work Email

Password



I agree to the [DataStax MSA](#), including the [Astra Supplement](#).

Create Account

Already have an account? [Sign In](#)

Astra DB

Multi-cloud DBaaS Built on Apache Cassandra™

- Start in minutes. No credit card required. Up to 80 GB free monthly.
- Build faster with REST, GraphQL, CQL and JSON/Document APIs.
- Deploy multi-tenant or dedicated databases on AWS, Azure, or GCP.

Introducing Astra Streaming Beta!

Astra Streaming is an open, multi-cloud event-streaming and data-streaming cloud service powered by Apache Pulsar™, now available within Astra when you sign up for an account.

Welcome to Astra! Let's finish signing up.

Tell us more about you

To provide you with a better experience, we need to know a bit more about you. Afterwards, we may be able to customize your experience, tools, and content.

What's your primary goal in using Astra? *

Learning Cassandra

What is your primary programming language?

Java


What role best describes you?

Student / Intern

Are you currently running a Cassandra workload?

Yes No

Next



Current Organization
boim@post.tau.ac.il

Dashboard

Databases

Create Database

You currently have no databases.
Get started by [creating your first one](#).

Streaming

BETA

Create Streaming



Sample App Gallery

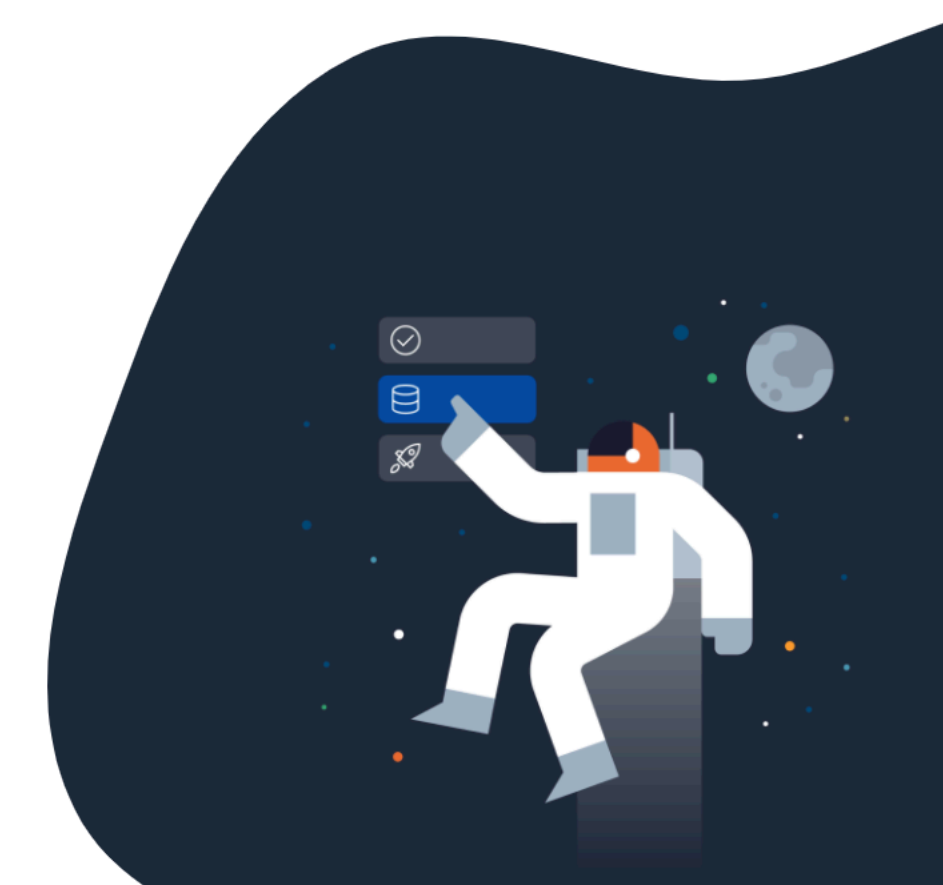
Looking to get up and running?
Get started with a sample app!

Get Started Now

Create your First Database

Get started in a matter of seconds. Deploy to Amazon Web Services, Google Cloud Platform, or Microsoft Azure. 99% uptime. Regional support. Liftoff!

Create Database



Dashboard

Current Plan

Free

Credits Remaining

\$0.00

\$0.00 spent

[Apply Promo Code](#)

Want to [upgrade your Astra serverless experience](#)? Add a credit card to get more.

Upgrade



New! Invite Friends, Get More Free Astra Credits

Both you and your friends will get \$25 in Astra credits when they sign up. That's worth 20 million reads/writes, and 80GB storage!

[Invite Friends](#)

Usage For Current Billing Period

Read Requests

0

Write Requests

0

Storage Consumed

0.00

Data Transfer

0.00

Databases (0)

Create Serverless Database

You haven't created a serverless database yet. Once you create a serverless database, it will appear here.

Other Resources

Documentation

Create a Database

1 Enter the Basic Details

Database Name <input type="text" value="GCP-Belgium-HW2"/> <small>Give it a memorable name - this can't be changed later.</small>	Keyspace Name ⓘ <input type="text" value="bigdatacourse"/> <small>Want to know more? Read our Docs to learn more about keyspaces.</small>
--	--

Current Plan

Free

You're currently on our free plan, which gives you free credits monthly. That recurring credit should be more than sufficient for your development needs, running sample code or apps, building proof-of-concepts, hackathon participation – even running small production workloads.

[Learn more](#) about our free accounts.

Create Database

2 Select a Provider and Region

You've got access to 3 free regions in GCP. Unlock all regions by [upgrading to the Pay as you go plan](#).

Google Cloud	Amazon Web Services	Microsoft Azure	Unlock All Regions
Select an Area	Select a Region		
North America 0 of 7 regions selected	<input checked="" type="radio"/> St. Ghislain, Belgium europ...		
Europe, Middle East, ... 1 of 2 regions selected >	<input type="radio"/> West Europe2 (London, E... europ... 🔒		
Asia Pacific 0 of 3 regions selected			

For HW choose GCP - Belgium



Current Organization
boim@post.tau.ac.il

[Dashboard](#) / TAU / AWS-Ireland [Load Data](#) [Connect](#)

[Dashboard](#)

[Databases](#) [Create Database](#)

● TAU / AWS-Ireland

[Streaming](#) **BETA** [Create Streaming](#)

Sample App Gallery
Looking to get up and running?
Get started with a sample app!

[Overview](#) [Health](#) [Connect](#) [CQL Console](#) [Settings](#)

Usage For Current Billing Period for TAU / AWS-Ireland

● Status Pending



Can take a few minutes...

Regions

Our Pay as you go plan allows you to add multiple regions. [Add Region](#)

Provider	Area	Region	Datacenter ID	Region Availability
Amazon Web Services	Europe, Middle East, and Africa	eu-west-1	8f9f3636-7e51-4a26-ad15-e7442194064e-1	Online

Keyspaces

Managing multiple applications? Consider keeping each application in a separate keyspace -- whatever a [keyspace](#) is. [Add Keyspace](#)

Keyspace
bigdatacourse

Other Resources

[Documentation](#)

Current Organization
boim@post.tau.ac.il

Dashboard / TAU / AWS-Ireland

Load Data

Connect

Dashboard

Databases

Create Database

TAU / AWS-Ireland

Streaming

BETA

Create Streaming



Sample App Gallery

Looking to get up and running?
Get started with a sample app!

Other Resources

Documentation

Live Chat

Overview Health Connect CQL Console Settings

```
Connected as boim@post.tau.ac.il.  
Connected to cndb at cassandra.ingress:9042.  
[cqqlsh 6.8.0 | Cassandra 4.0.0.6816 | CQL spec 3.4.5 | Native protocol v4]  
Use HELP for help.  
token@cqqlsh> describe keyspaces;  
  
system_virtual_schema  system_views  data_endpoint_auth  
system_schema          system       system_traces  
system_auth             datastax_sla bigdatacourse  
  
token@cqqlsh> █
```

Useful commands

- Autocomplete with the “tab” key
- `describe keyspaces;`
- `use <keyspace>`
- `describe tables`
- `describe <table>`
- `clear / CTRL+L` (clear screen)
- `create table / select / ...`

Demo - Create table

```
CREATE TABLE users_by_country (  
    country          TEXT,  
    user_id         TEXT,  
    name            TEXT,  
    age             INT,  
    PRIMARY KEY ((country), user_id)  
);
```

users_by_country	
country	K
user_id	▲C
name	
age	

Demo - Insert

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('Israel', 'user_123', 'Rubi Boim', 22);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('USA', 'user_123', 'Rubi Boim', 22);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('Israel', 'user_124', 'Tova Milo', 18);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('USA', 'user_125', 'Lebron James', 35);
```

```
INSERT INTO users_by_country(country, user_id, name)
VALUES ('USA', 'user_125', 'Michael Jordan');
```

users_by_country	
country	K
user_id	▲C
name	
age	

Demo - Insert

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('Israel', 'user_123', 'Rubi Boim', 22);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('USA', 'user_123', 'Rubi Boim', 22);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('Israel', 'user_124', 'Tova Milo', 18);
```

```
INSERT INTO users_by_country(country, user_id, name, age)
VALUES ('USA', 'user_125', 'Lebron James', 35);
```

```
INSERT INTO users_by_country(country, user_id, name)
VALUES ('USA', 'user_125', 'Michael Jordan');
```

users_by_country	
country	K
user_id	▲C
name	
age	

What is Michael's age?

Demo - Select

Which queries are valid?

```
SELECT * FROM users_by_country;
```

```
SELECT * FROM users_by_country WHERE country = 'Israel';
```

```
SELECT * FROM users_by_country WHERE name = 'Rubi Boim';
```

```
SELECT * FROM users_by_country WHERE user_id = 'user_123';
```

```
SELECT * FROM users_by_country WHERE country = 'Israel' AND user_id =  
'user_123';
```

```
SELECT * FROM users_by_country WHERE country = 'Israel' AND name =  
'Rubi Boim';
```

users_by_country	
country	K
user_id	▲C
name	
age	

Demo - Select

Which queries are valid?

users_by_country

country K

user_id ▲C

name

age

```
SELECT * FROM users_by_country;
```

```
SELECT * FROM users_by_country WHERE country = 'Israel';
```

```
SELECT * FROM users_by_country WHERE name = 'Rubi Boim';
```

```
SELECT * FROM users_by_country WHERE user_id = 'user_123';
```

```
SELECT * FROM users_by_country WHERE country = 'Israel' AND user_id =  
'user_123';
```

```
SELECT * FROM users_by_country WHERE country = 'Israel' AND name =  
'Rubi Boim';
```

Remember - "ALLOW FILTERING is ANTI PATTERN"

Demo - Update

`UPDATE is similar to INSERT
(*)`

`If you prefer to use the UPDATE syntax, "read the..."`

Demo - Delete

Which queries are valid?

users_by_country	
country	K
user_id	▲C
name	
age	

```
DELETE FROM users_by_country WHERE name = 'Rubi Boim';
```

```
DELETE FROM users_by_country WHERE user_id = 'user_123';
```

```
DELETE FROM users_by_country WHERE country = 'Israel';
```

```
DELETE FROM users_by_country WHERE country = 'Israel'  
AND name = 'Rubi Boim';
```

```
DELETE FROM users_by_country WHERE country = 'Israel'  
AND user_id = 'user_123';
```

Demo - Drop table

```
DROP TABLE <table_name>;
```