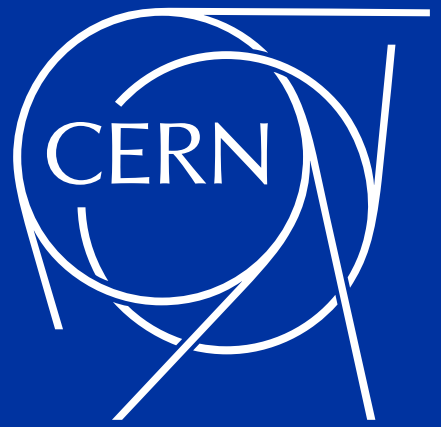

Database Services by IT-DA-DB

Introduction

Andrzej Nowicki

tCSC Ferney-Voltaire 2024

Andrzej Nowicki

12 years of Oracle DB experience
Database Engineer @ CERN since 2020

andrzej.nowicki@cern.ch

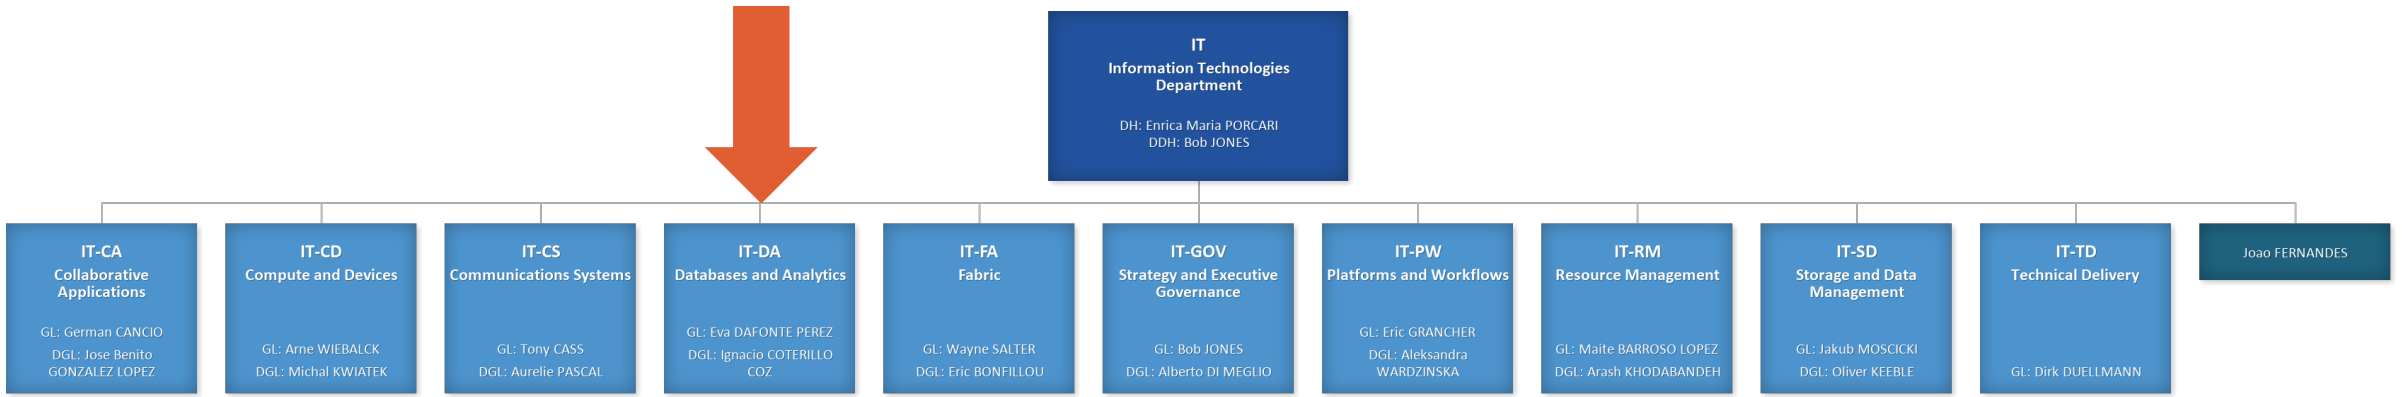
IT
Information Technologies Department
DH: Enrica Maria PORCARI
DDH: Bob JONES
IT-CA
Collaborative Applications
GL: German CANCIO
DGL: Jose Benito GONZALEZ LOPEZ
IT-CD
Compute and Devices
GL: Arne WIEBALCK
DGL: Michal KWIATEK
IT-CS
Communications Systems
GL: Tony CASS
DGL: Aurelie PASCAL
IT-DA
Databases and Analytics
GL: Eva DAFONTE PEREZ
DGL: Ignacio COTERILLO COZ
IT-FA
Fabric
GL: Wayne SALTER
DGL: Eric BONFILLOU
IT-GOV
Strategy and Executive Governance
GL: Bob JONES
DGL: Alberto DI MEGLIO
IT-PW
Platforms and Workflows
GL: Eric GRANCHER
DGL: Aleksandra WARDZINSKA
IT-RM
Resource Management
GL: Maite BARROSO LOPEZ
DGL: Arash KHODABANDEH
IT-SD
Storage and Data Management
GL: Jakub MOSCICKI
DGL: Oliver KEEBLE
IT-TD
Technical Delivery
GL: Dirk DUELLMANN
Joao FERNANDES

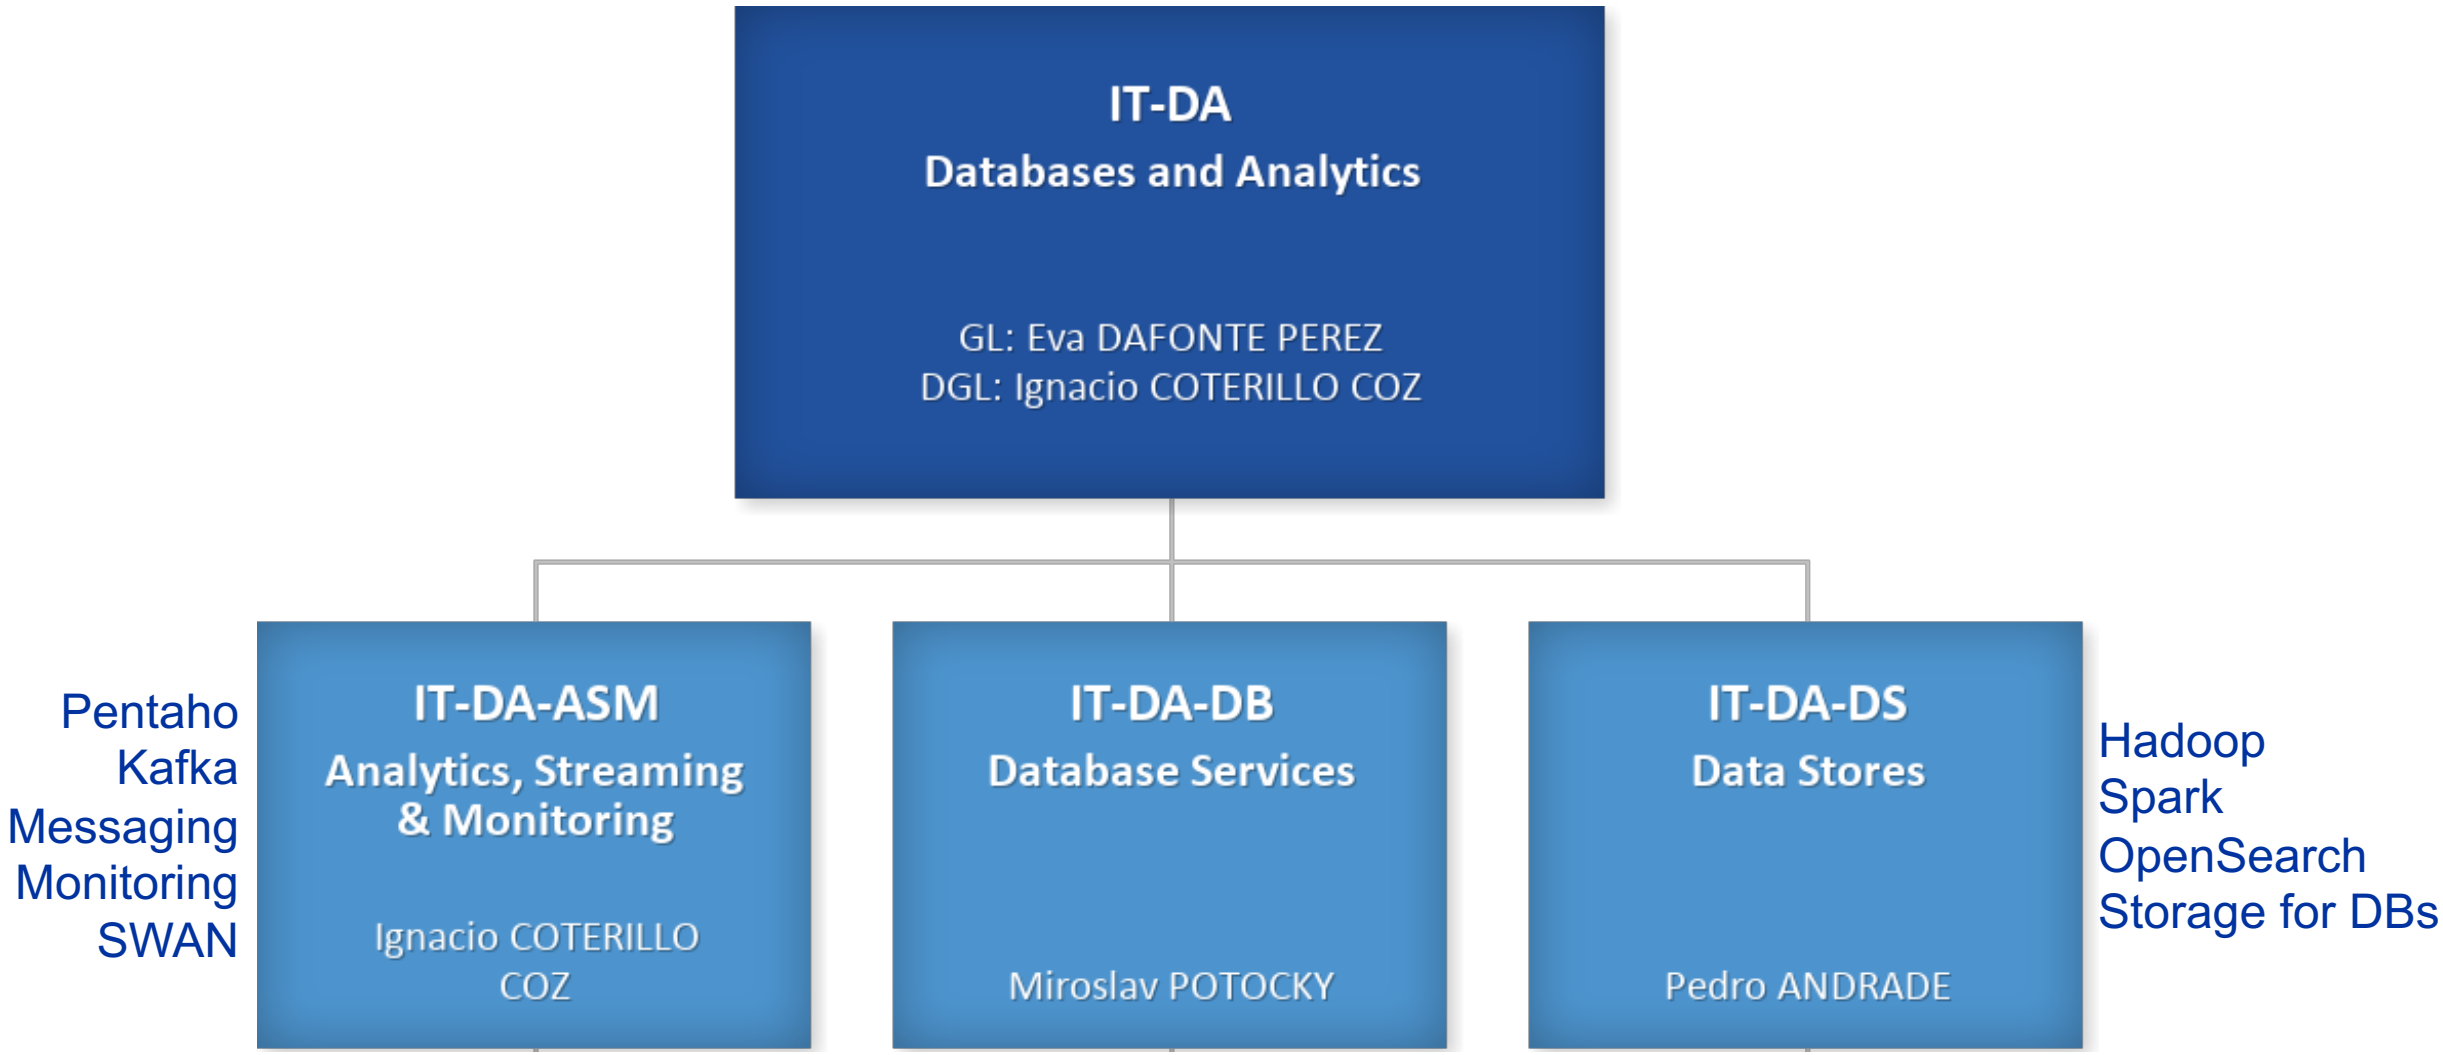
IT-DA
Databases and Analytics
GL: Eva DAFONTE PEREZ
DGL: Ignacio COTERILLO COZ
IT-DA-ASM
Analytics, Streaming & Monitoring
Ignacio COTERILLO COZ
IT-DA-DB
Database Services
Miroslav POTOCKY
IT-DA-DS
Data Stores
Pedro ANDRADE
Pentaho
Kafka
Messaging
Monitoring
SWAN
Hadoop
Spark
OpenSearch
Storage for DBs

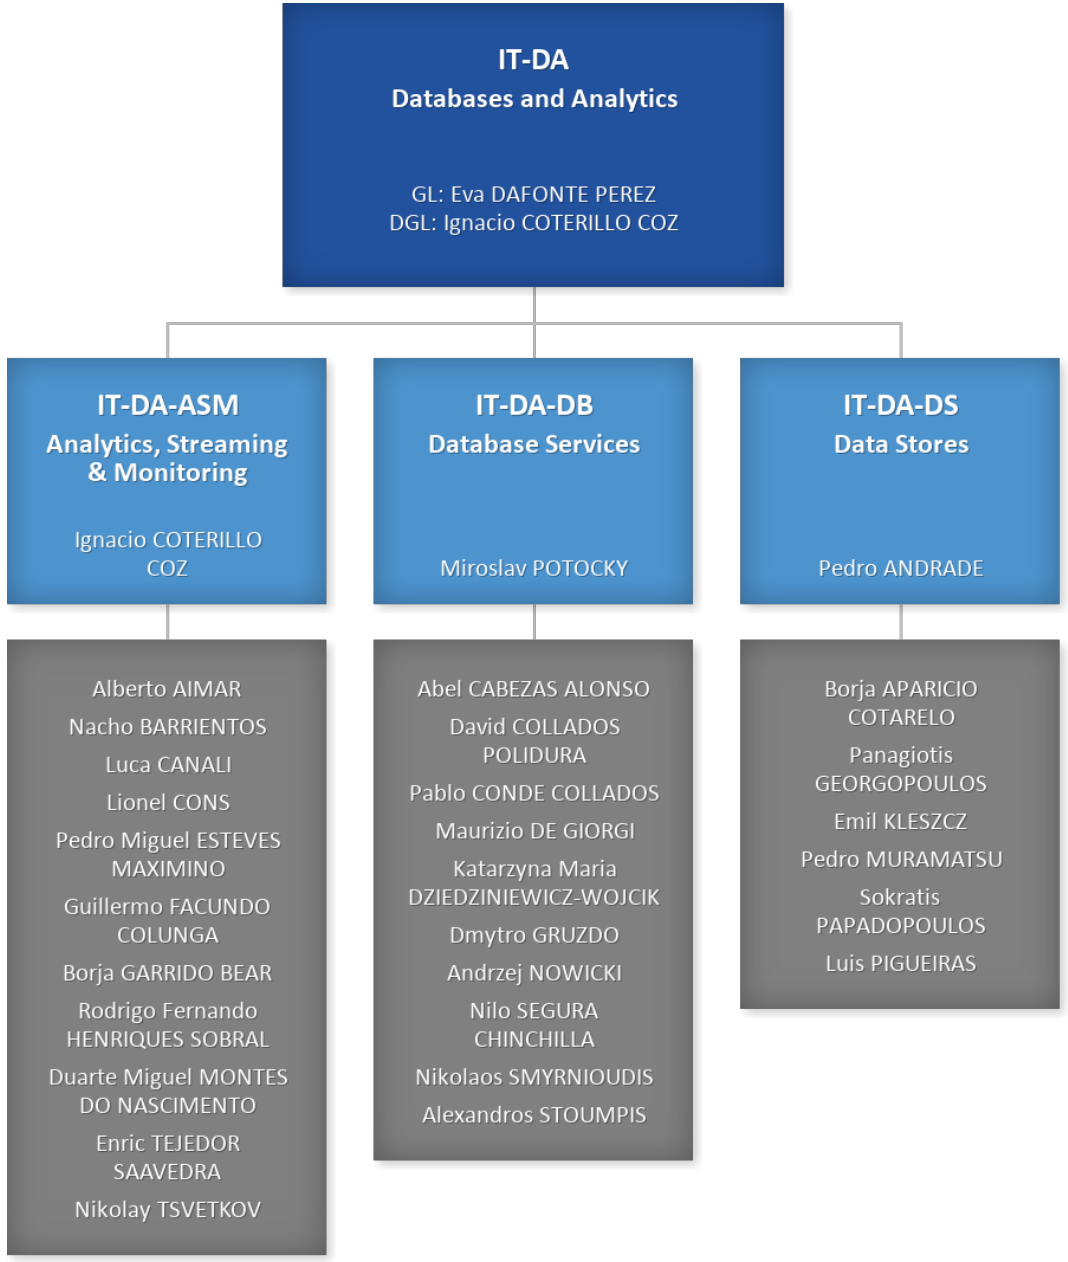

https://organigram.web.cern.ch/organigram/IT-DA

Monday 4 November 2024	Tuesday 5 November 2024	Wednesday 6 November 2024	Thursday 7 November 2024	Friday 8 November 2024
17:00 Database Services (part 1 of 4) - Introduction to DBoD		09:00 Database Services (part 2 of 4) - DBoD maintenance exercises 11:30 Database Services (part 3 of 4) - DBoD maintenance exercises		12:30 Database Services (part 4 of 4) - Oracle Database

Databases

A **database** is:

- an organized collection of data or a type of data store based on the use of a **database management system** (**DBMS**),

- the software that interacts with end users, applications, and the database itself to capture and analyze the data.

- The DBMS additionally encompasses the core facilities provided to administer the database. The sum total of the database, the DBMS and the associated applications can be referred to as a **database system**.

Relational databases

A relational model:

- **organizes data into one or more tables (or "relations") of columns and rows**
- a unique key identifying each row

Generally, each table/relation represents one "entity type" (such as customer or product). The rows represent instances of that type of entity (such as "Lee" or "chair") and the columns represent values attributed to that instance (such as address or price).

Customer ID	Name
123	Alice
456	Bob
789	Carol

SQL

Is a language used to manage data, especially in a relational database management system (RDBMS)

```
create table customers (CustomerId number, Name varchar2(20));

insert into customers values (123, 'Alice');

select CustomerId from customers where Name = 'Alice';

update customers set Name = 'Unnamed';

delete from customers where CustomerId = 123;
```

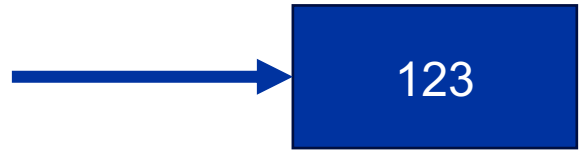

Customer ID	Name
~~123~~	~~Unnamed~~
456	Unnamed
789	Unnamed

A tiny bit of theory…

CAP theorem AKA Brewer's theorem

You can only have 2 out of 3

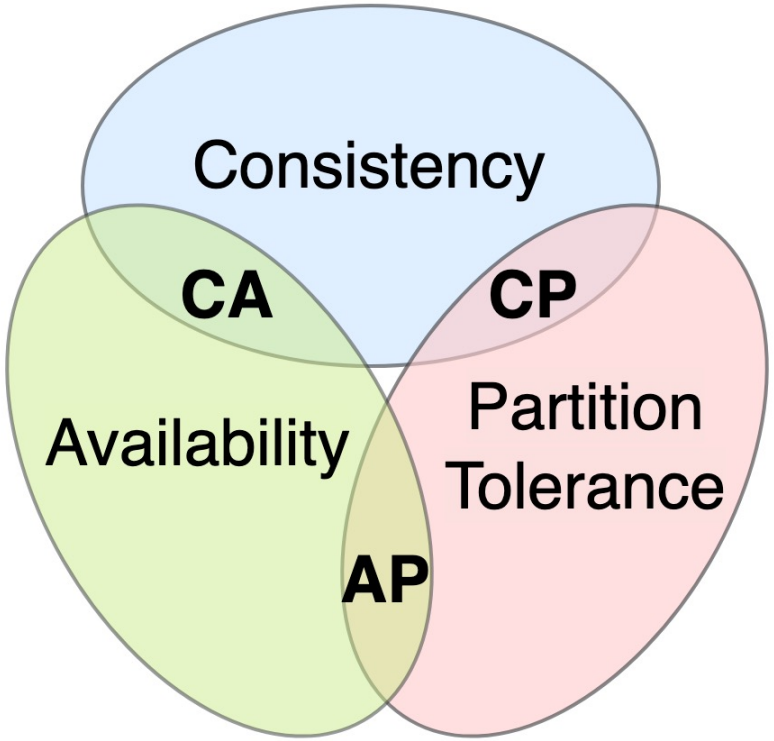

https://en.wikipedia.org/wiki/CAP_theorem

CAP theorem

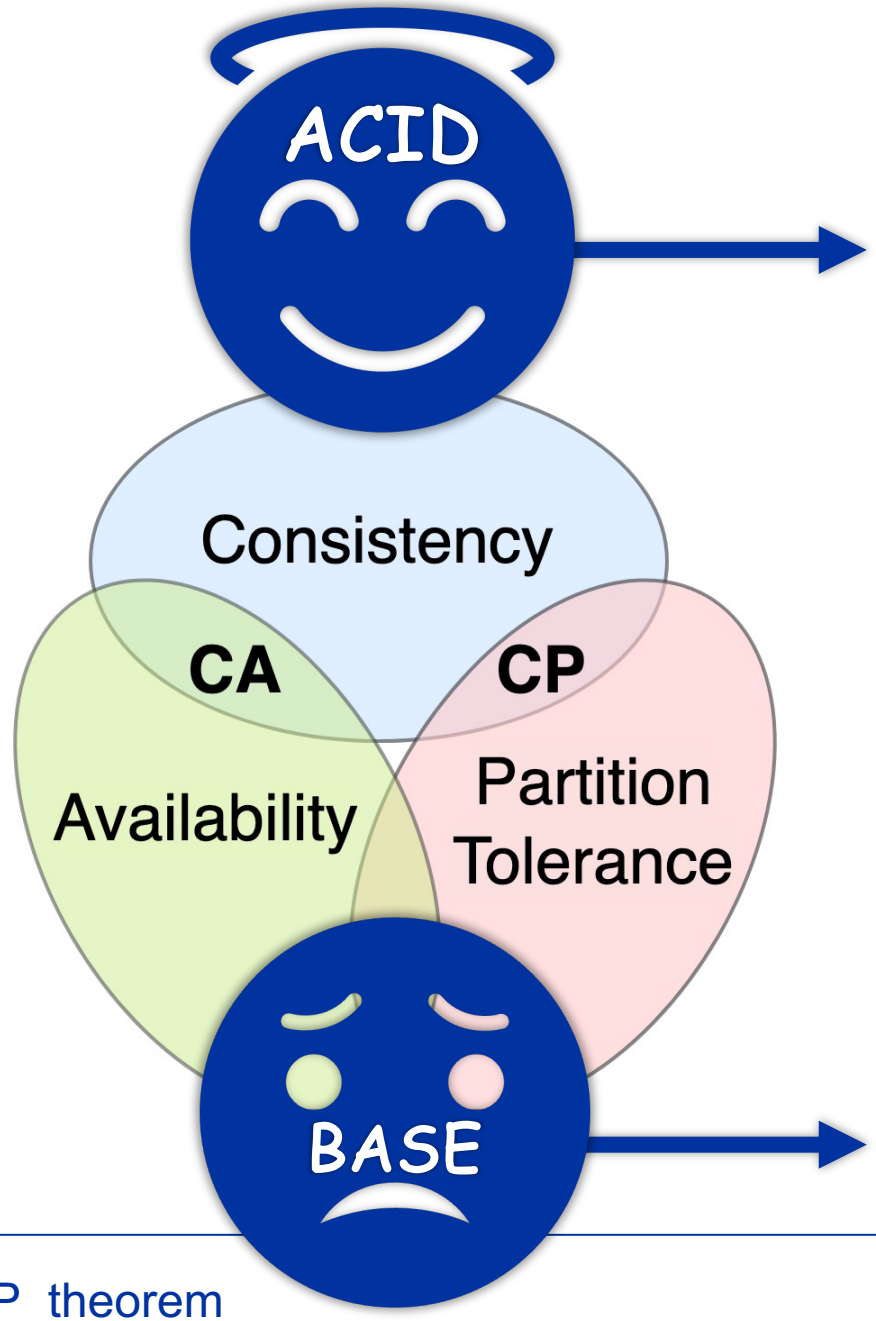

Atomicity
Consistency
Isolation
Durability

Eventual Consistency

https://en.wikipedia.org/wiki/CAP_theorem

How to make your database resilient?

Basic database system

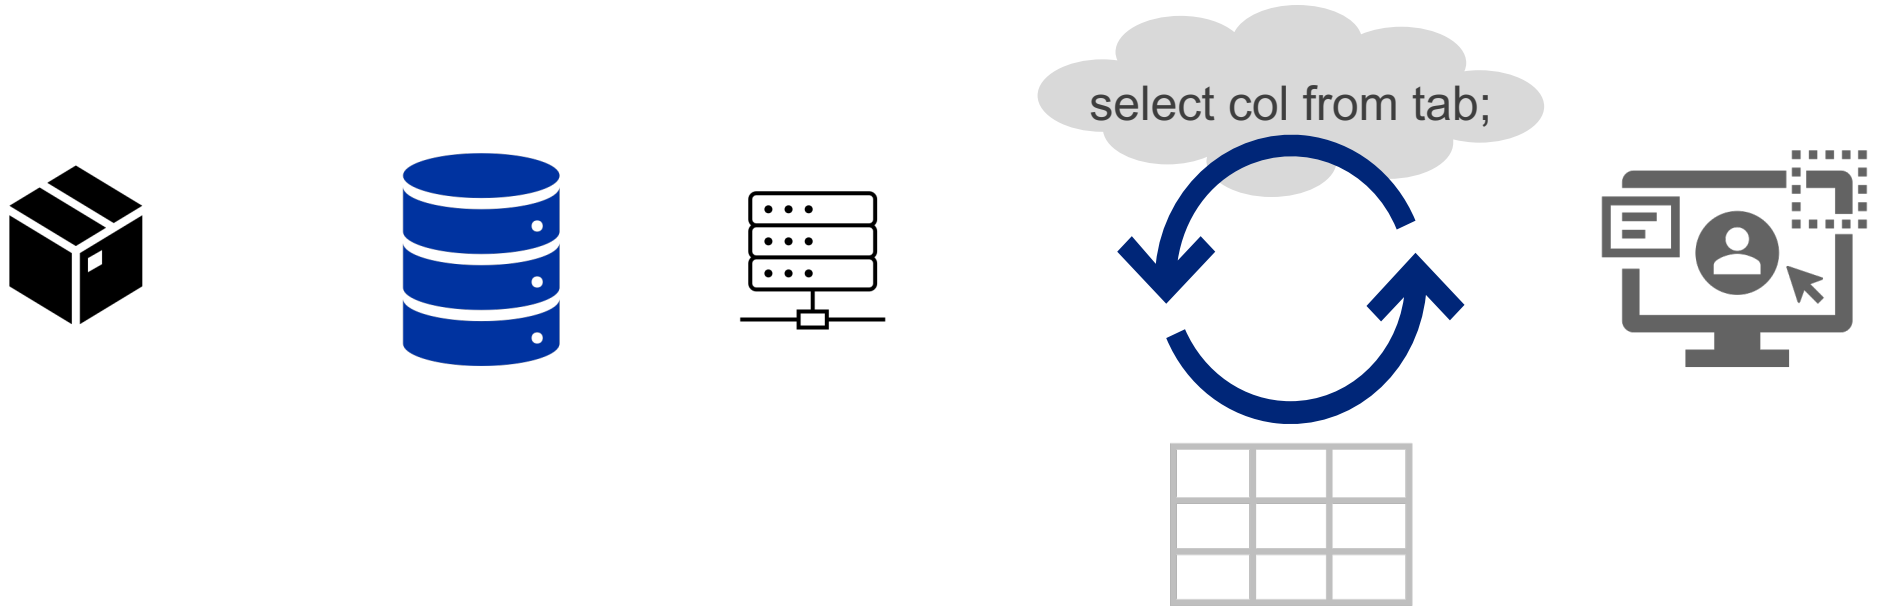

What if something breaks?

Basic database system

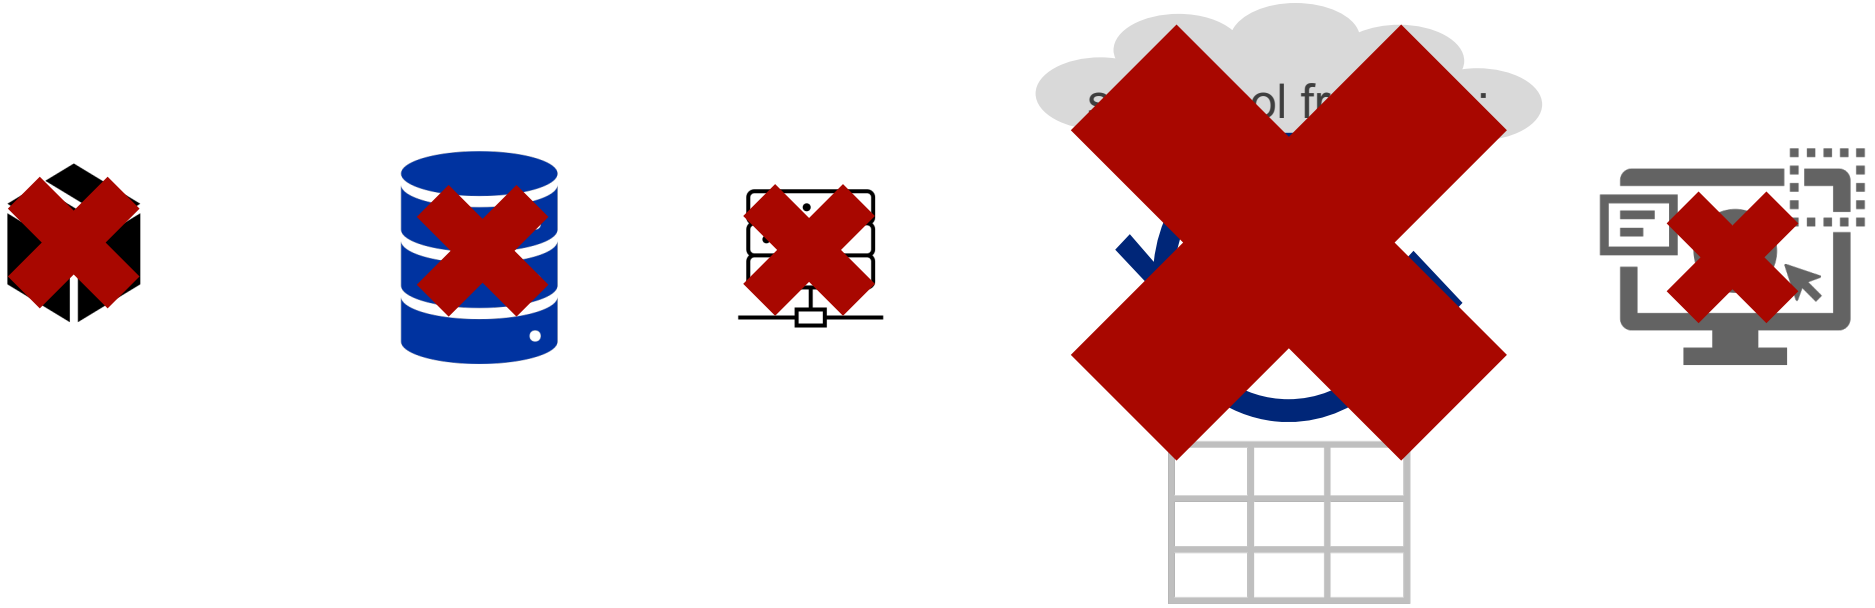

STORAGE DATAFILES SERVER APPLICATION

Highly available database system

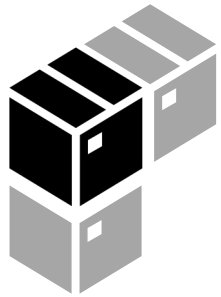

STORAGE

DATAFILES

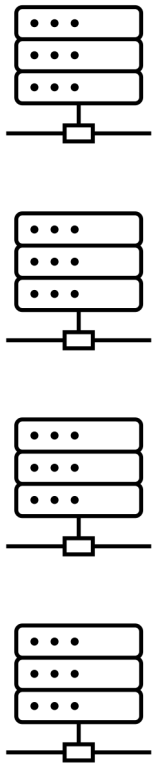

SERVERS

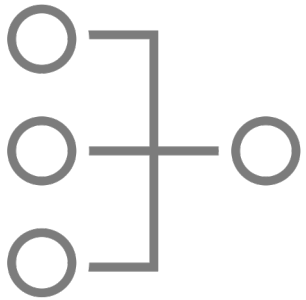

LISTENER
(LOAD BALANCER)

APPLICATION

Is that enough?

What if your datacenter is on fire?

Highly available database system with DR

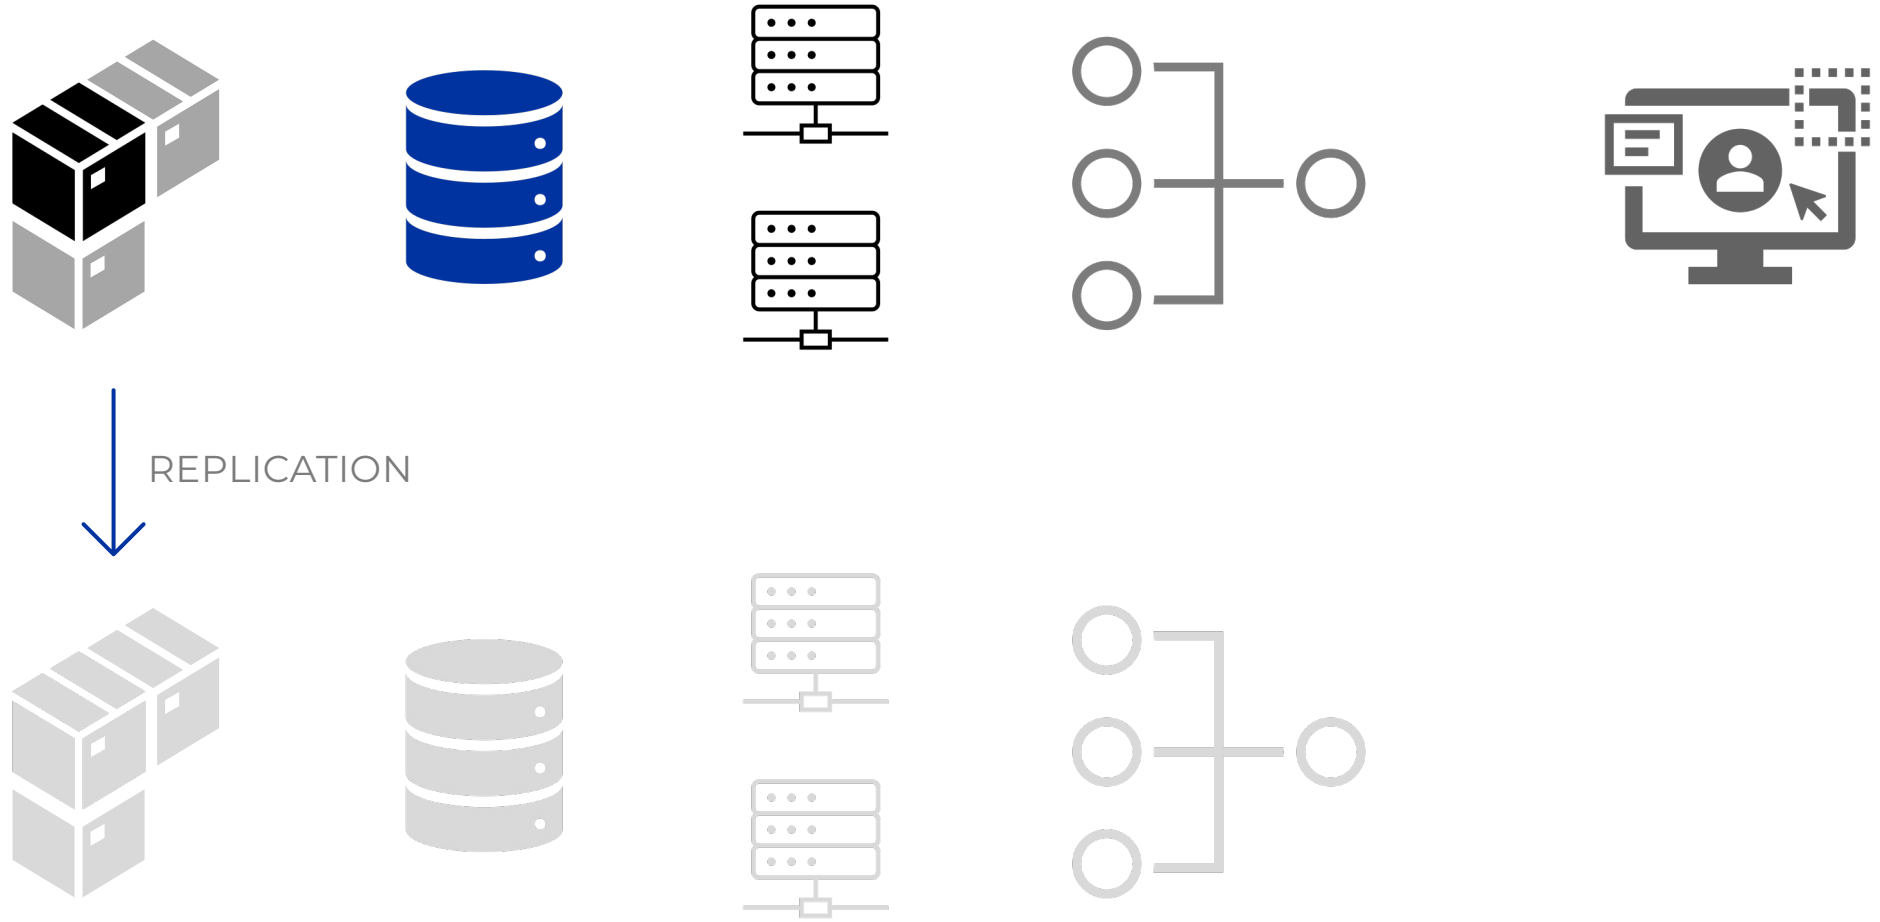

Highly available database system with DR

Highly available database system with DR

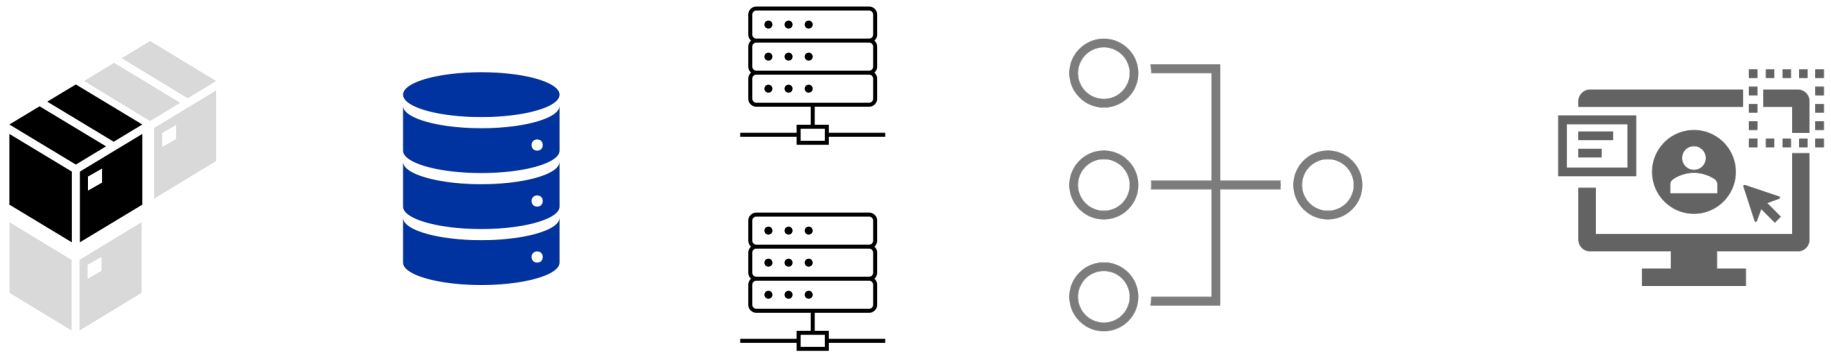

Is that enough?


WHEN YOU DELETE A TABLE
FROM A PRODUCTION DATABASE
imgflip.com

Backups

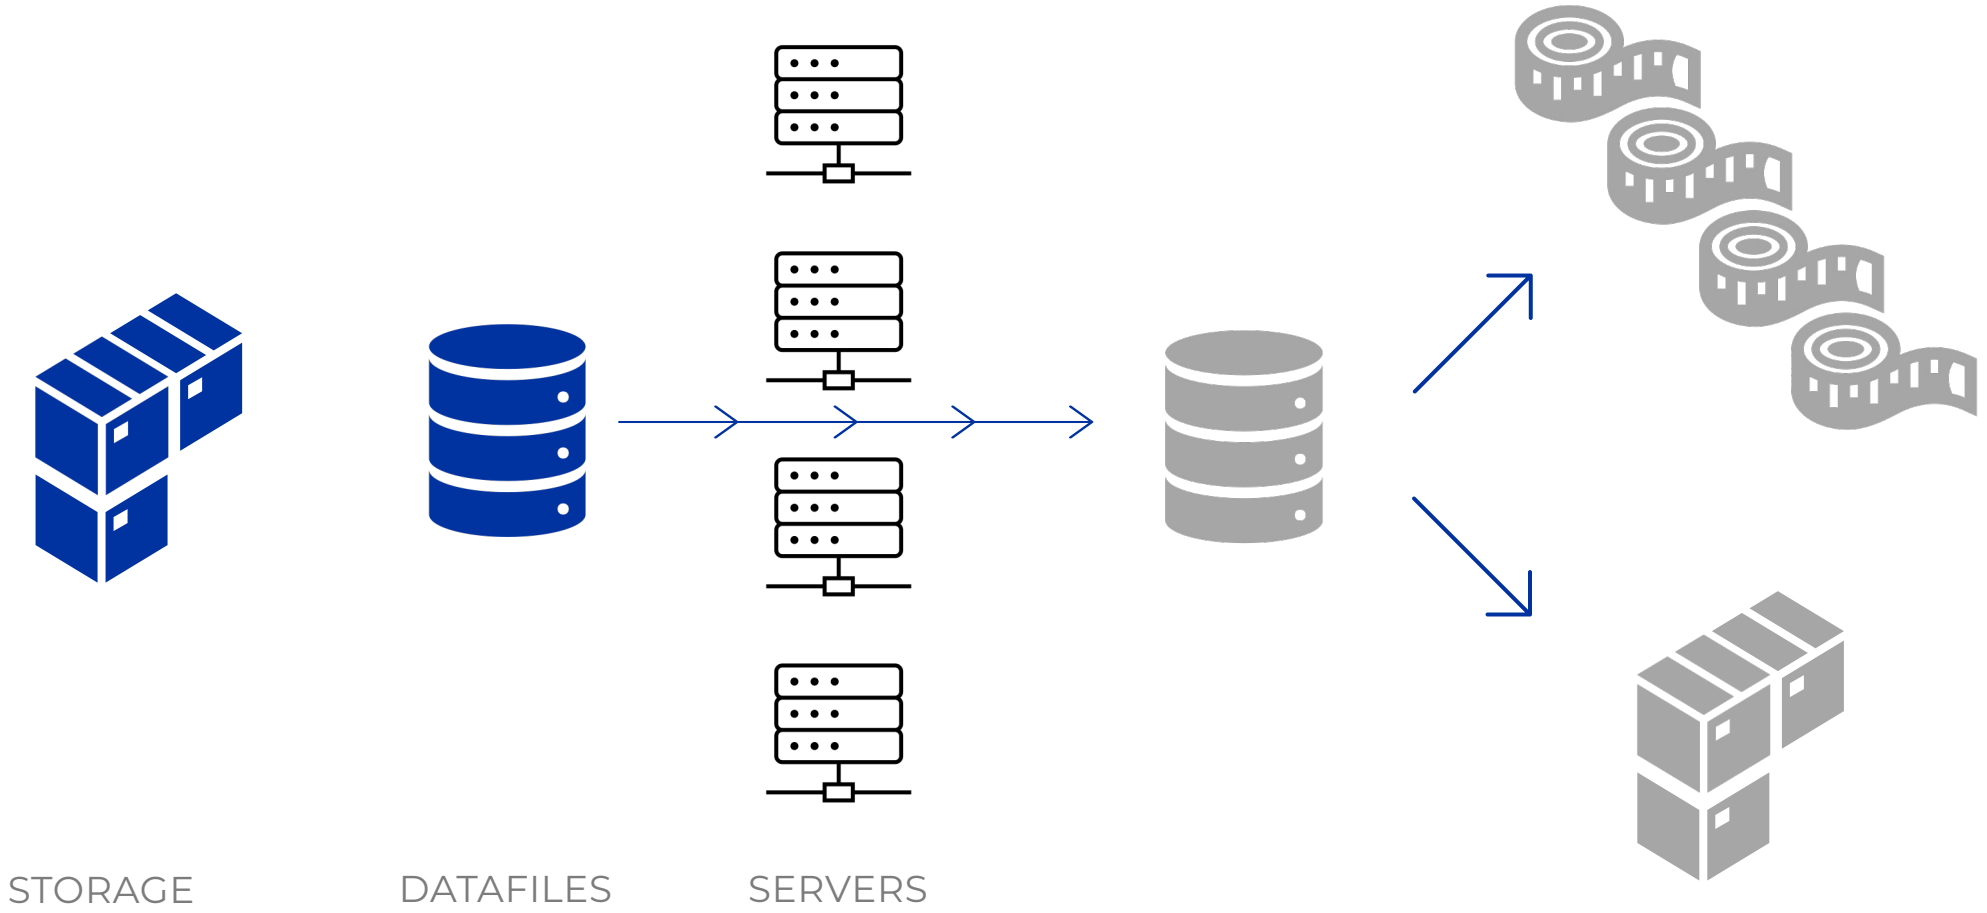

Backups

Modern database systems are allowing us to restore to any point in time

You need to backup the data files but also the log of all the changes

Each system has different approach to solving this issue:

Oracle:	**archivelogs**
Postgres:	**write ahead log – WAL**
MySQL:	**binlog**

Backups

Modern data

You need to

Each system

- Oracl
- Postg
- MySQ

How does our infra look on the application side?

Application Side

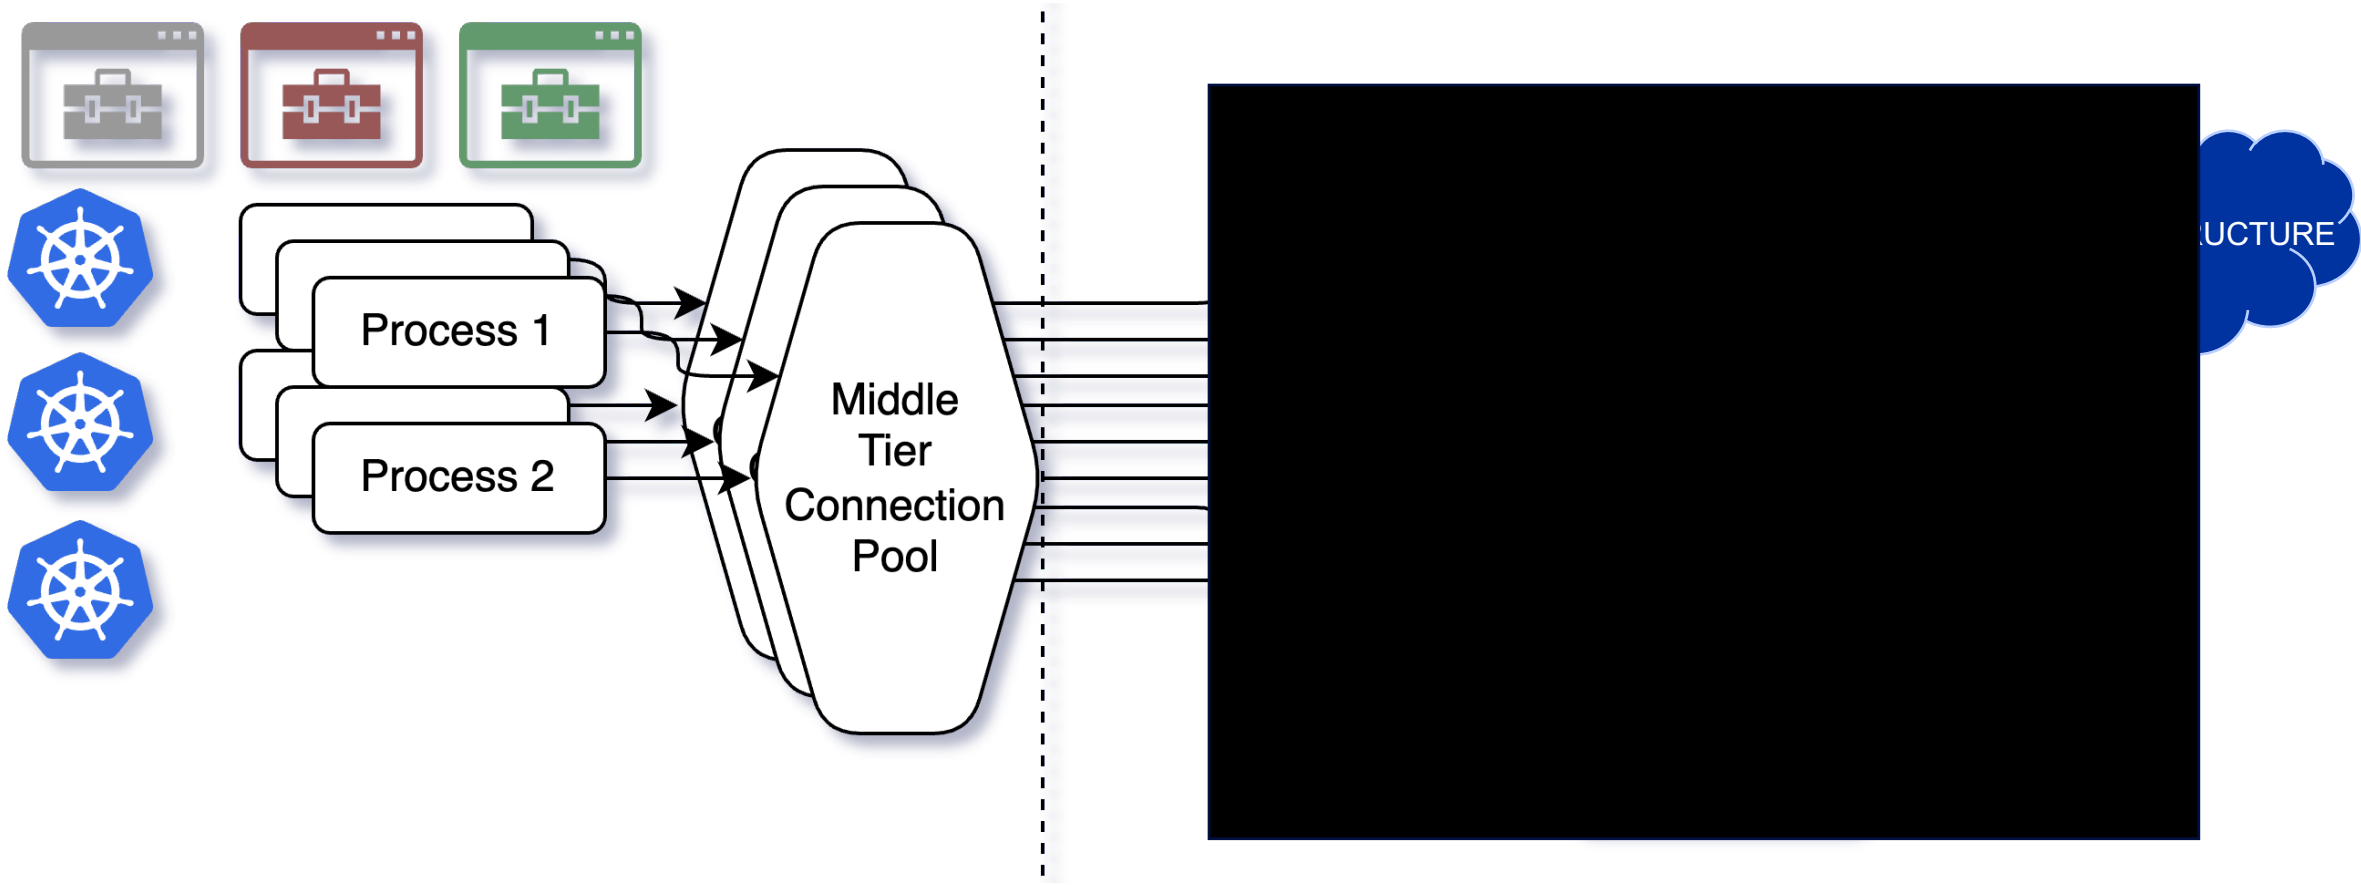

Databases at CERN

Oracle since 1982

- 105 Oracle databases, more than 11.800 Oracle accounts
- *RAC, Active Data Guard, GoldenGate, OEM, RMAN, APEX, Cloud…*
- Complex environment, ≈100 physical servers

Database on Demand (DBoD) since 2011

- ≈600 MySQL, ≈400 PostgreSQL, ≈200 InfluxDB
- *Automated backup and recovery services, monitoring, clones, replicas*
- *HA MySQL clusters (Proxy + primary replica)*
- *≈200 servers (both VMs and physical)*

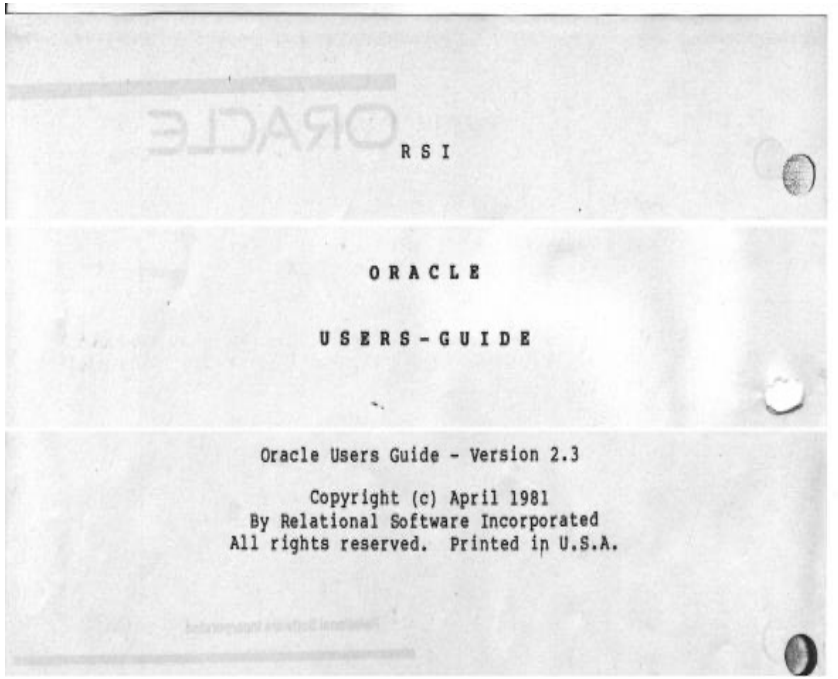
RSI

ORACLE

USERS-GUIDE

Oracle Users Guide - Version 2.3

Size of the database environment

	Total size
Oracle	≈5 PB
DBoD (MySQL, PostrgeSQL, InfluxDB)	≈150 TB
Backups	≈3 PB

Resource management

Each resource should have an owner, whom you can contact in case of problems.

We give two options:

- **individual owner**
- **service owner (which might belong to a team, group, etc.)**

What do you do when people leave the organisation?

Oracle

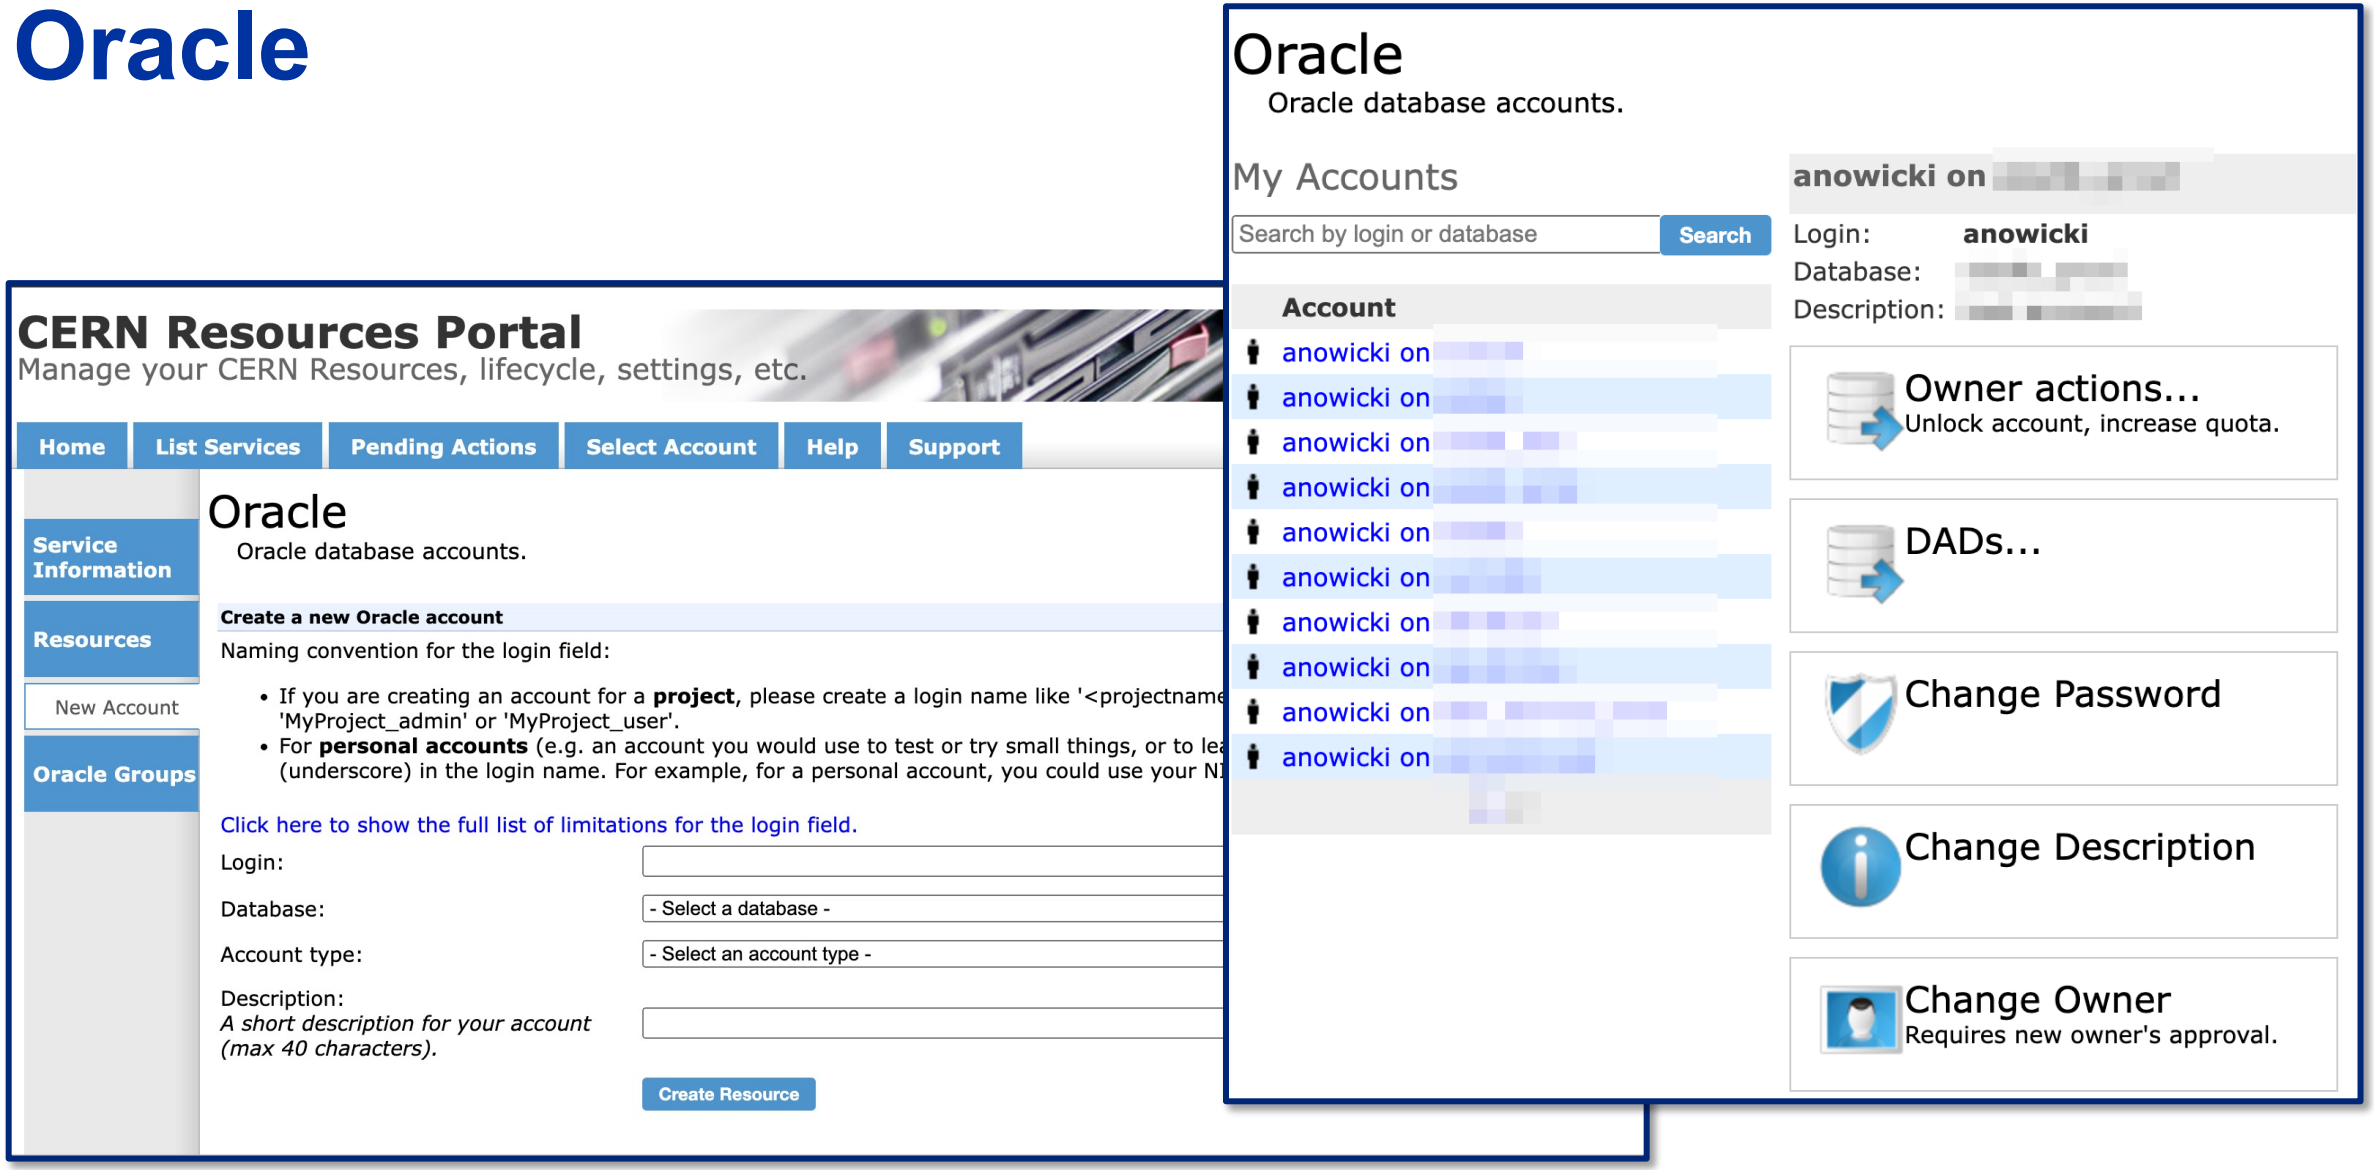
CERN Resources Portal
Manage your CERN Resources, lifecycle, settings, etc.
Home
List Services
Pending Actions
Select Account
Help
Support
Service Information
Resources
New Account
Oracle Groups
Oracle
Oracle database accounts.
Create a new Oracle account
Naming convention for the login field:
If you are creating an account for a project, please create a login name like '<projectname
'MyProject_admin' or 'MyProject_user'.
For personal accounts (e.g. an account you would use to test or try small things, or to lea
(underscore) in the login name. For example, for a personal account, you could use your NI
Click here to show the full list of limitations for the login field.
Login:
Database:
- Select a database -
Account type:
- Select an account type -
Description:
A short description for your account
(max 40 characters).
Create Resource
Oracle
Oracle database accounts.
My Accounts
Search by login or database
Search
Account
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
anowicki on
Login: anowicki
Database:
Description:
Owner actions...
Unlock account, increase quota.
DADs...
Change Password
Change Description
Change Owner
Requires new owner's approval.

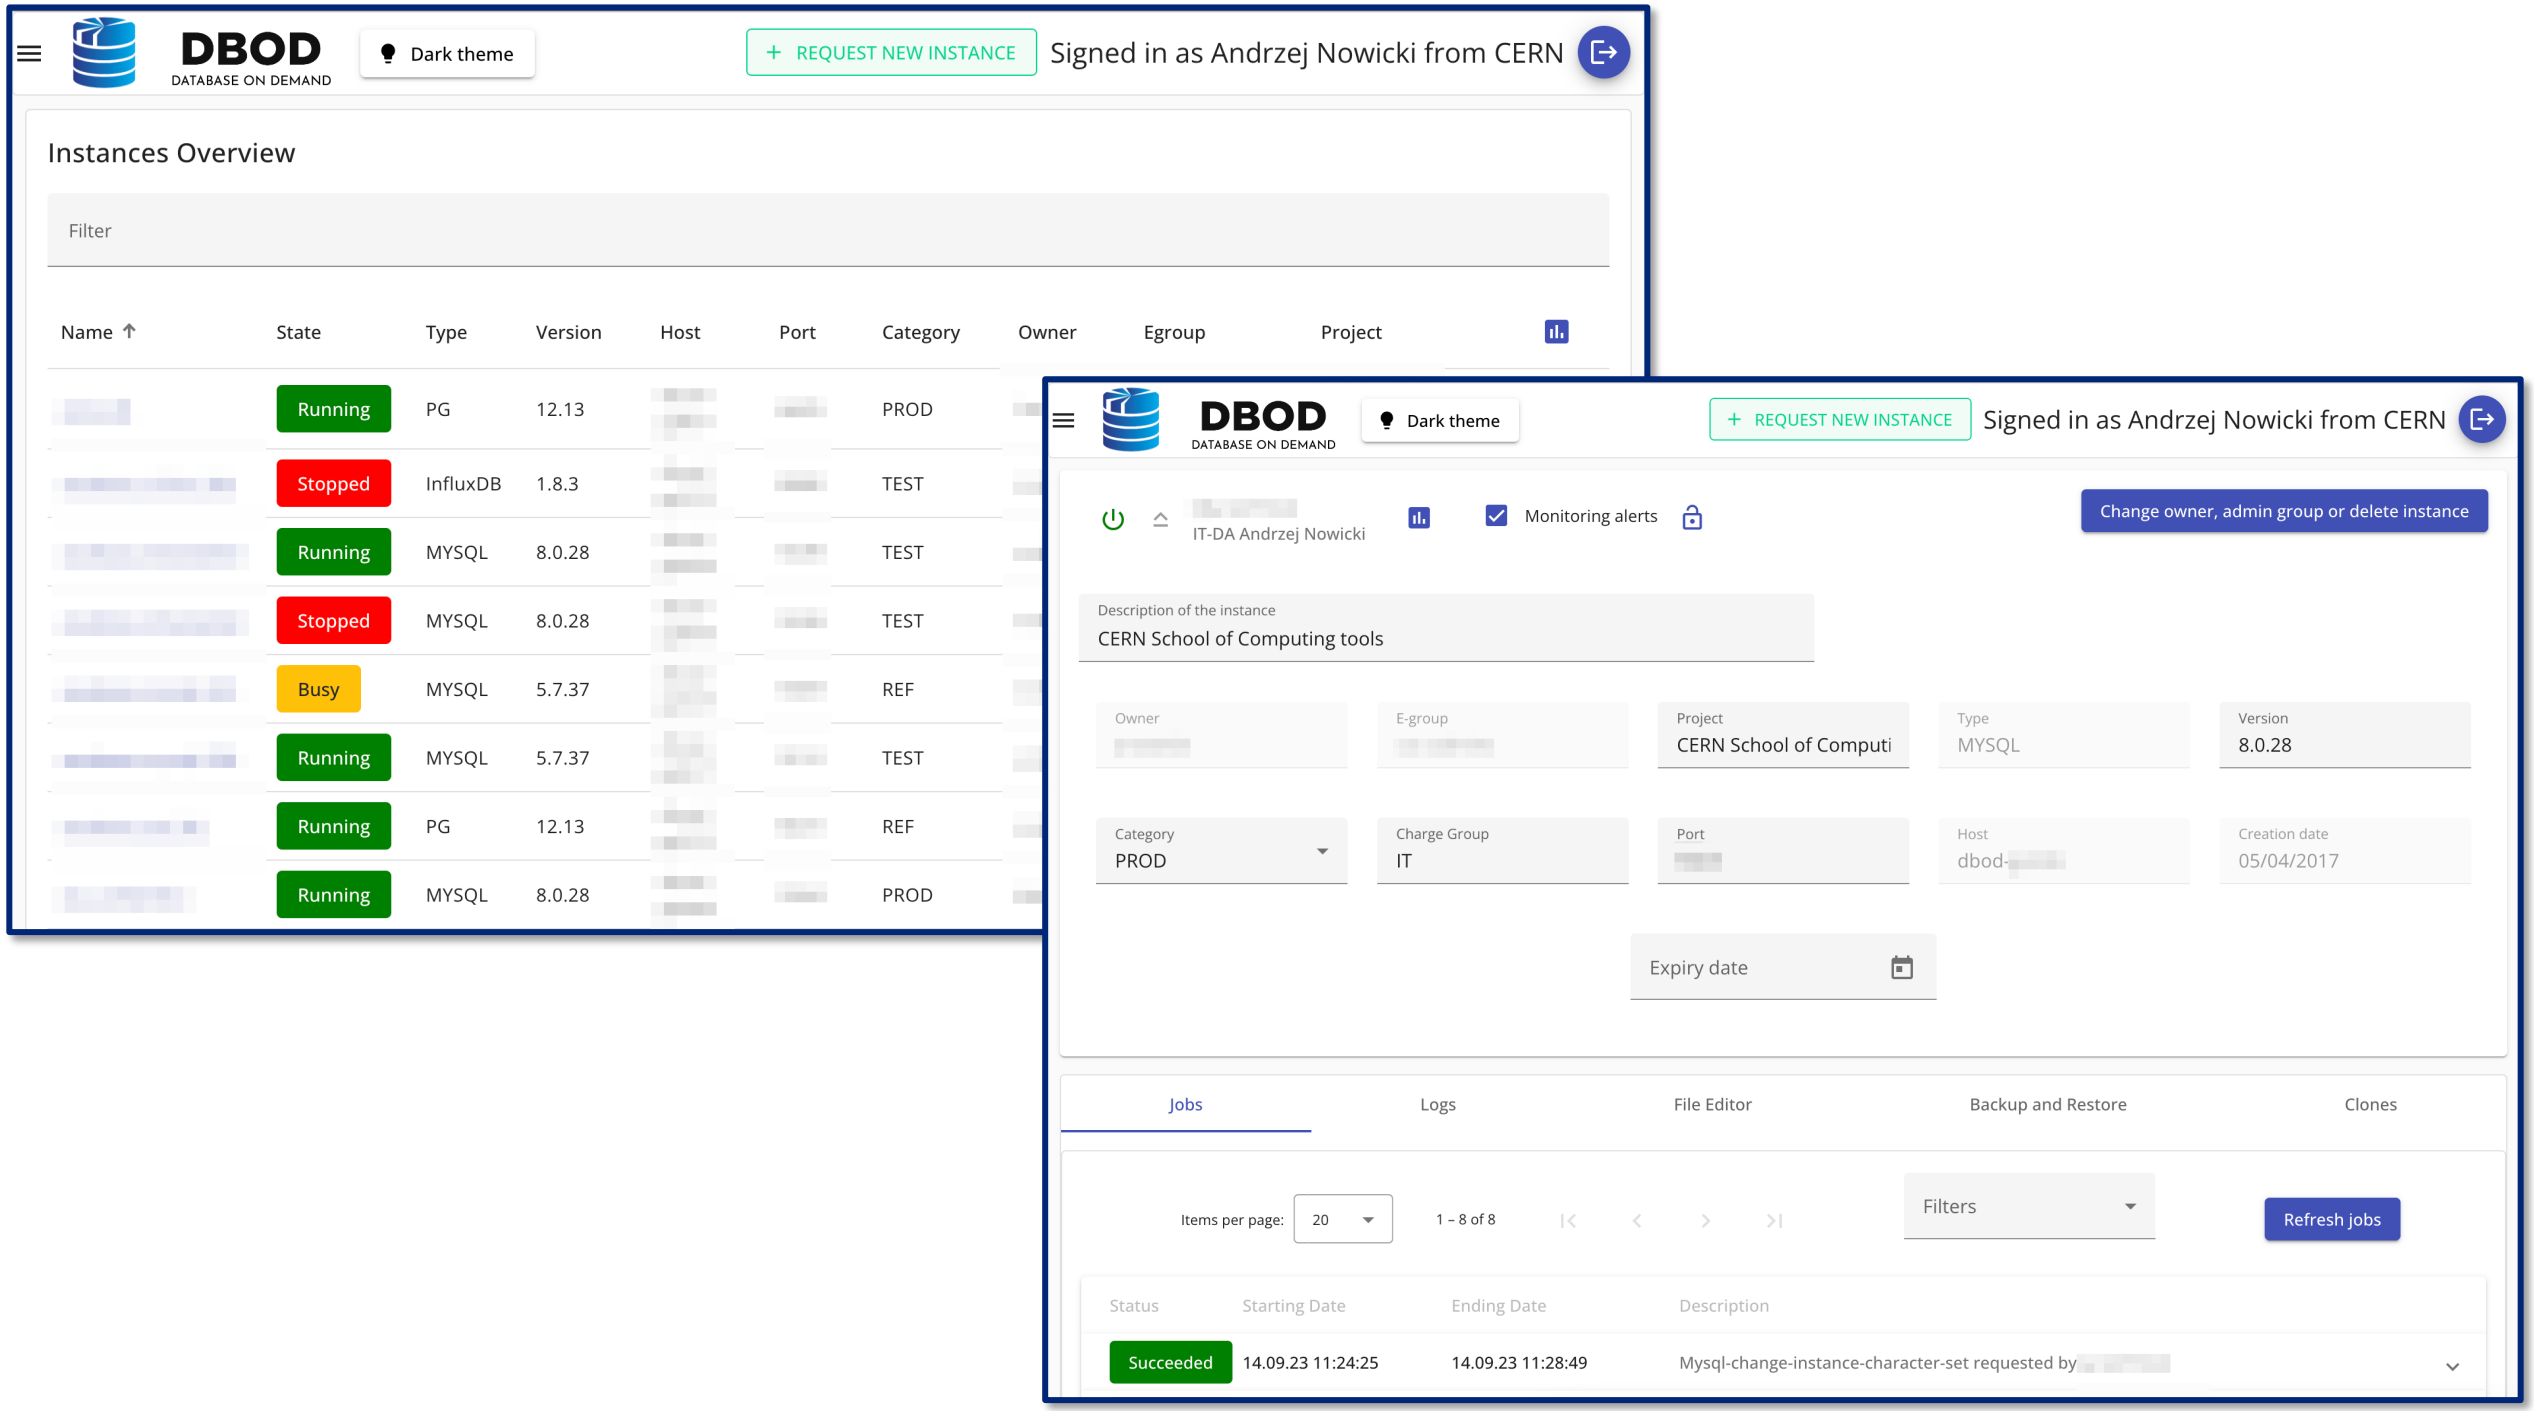
DBOD
DATABASE ON DEMAND
Dark theme
+ REQUEST NEW INSTANCE
Signed in as Andrzej Nowicki from CERN
Instances Overview
Filter
Name
State
Type
Version
Host
Port
Category
Owner
Egroup
Project
Running PG 12.13 PROD
Stopped InfluxDB 1.8.3 TEST
Running MYSQL 8.0.28 TEST
Stopped MYSQL 8.0.28 TEST
Busy MYSQL 5.7.37 REF
Running MYSQL 5.7.37 TEST
Running PG 12.13 REF
Running MYSQL 8.0.28 PROD
DBOD
DATABASE ON DEMAND
Dark theme
+ REQUEST NEW INSTANCE
Signed in as Andrzej Nowicki from CERN
IT-DA Andrzej Nowicki
Monitoring alerts
Change owner, admin group or delete instance
Description of the instance
CERN School of Computing tools
Owner
E-group
Project
CERN School of Computi
Type
MYSQL
Version
8.0.28
Category
PROD
Charge Group
IT
Port
Host
dbod-
Creation date
05/04/2017
Expiry date
Jobs
Logs
File Editor
Backup and Restore
Clones
Items per page:
20
1 – 8 of 8
Filters
Refresh jobs
Status
Starting Date
Ending Date
Description
Succeeded
14.09.23 11:24:25
14.09.23 11:28:49
Mysql-change-instance-character-set requested by

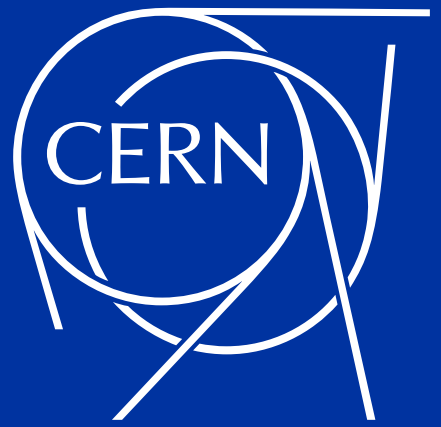

Database Services by IT-DA-DB

Database on Demand

Andrzej Nowicki

tCSC Ferney-Voltaire 2024

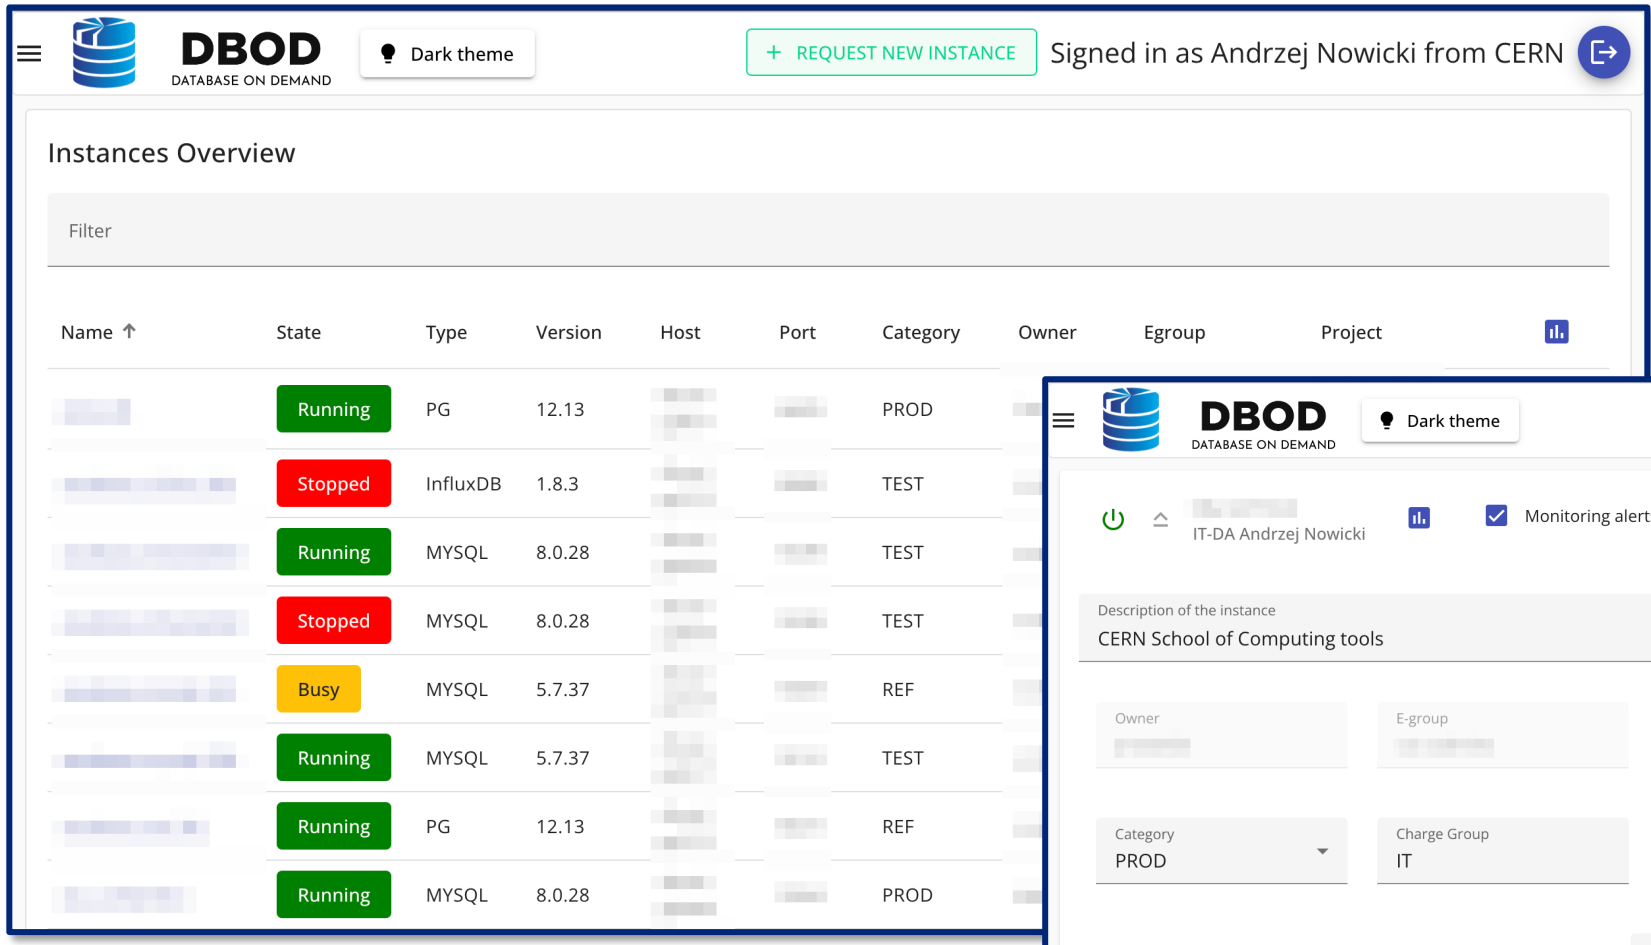

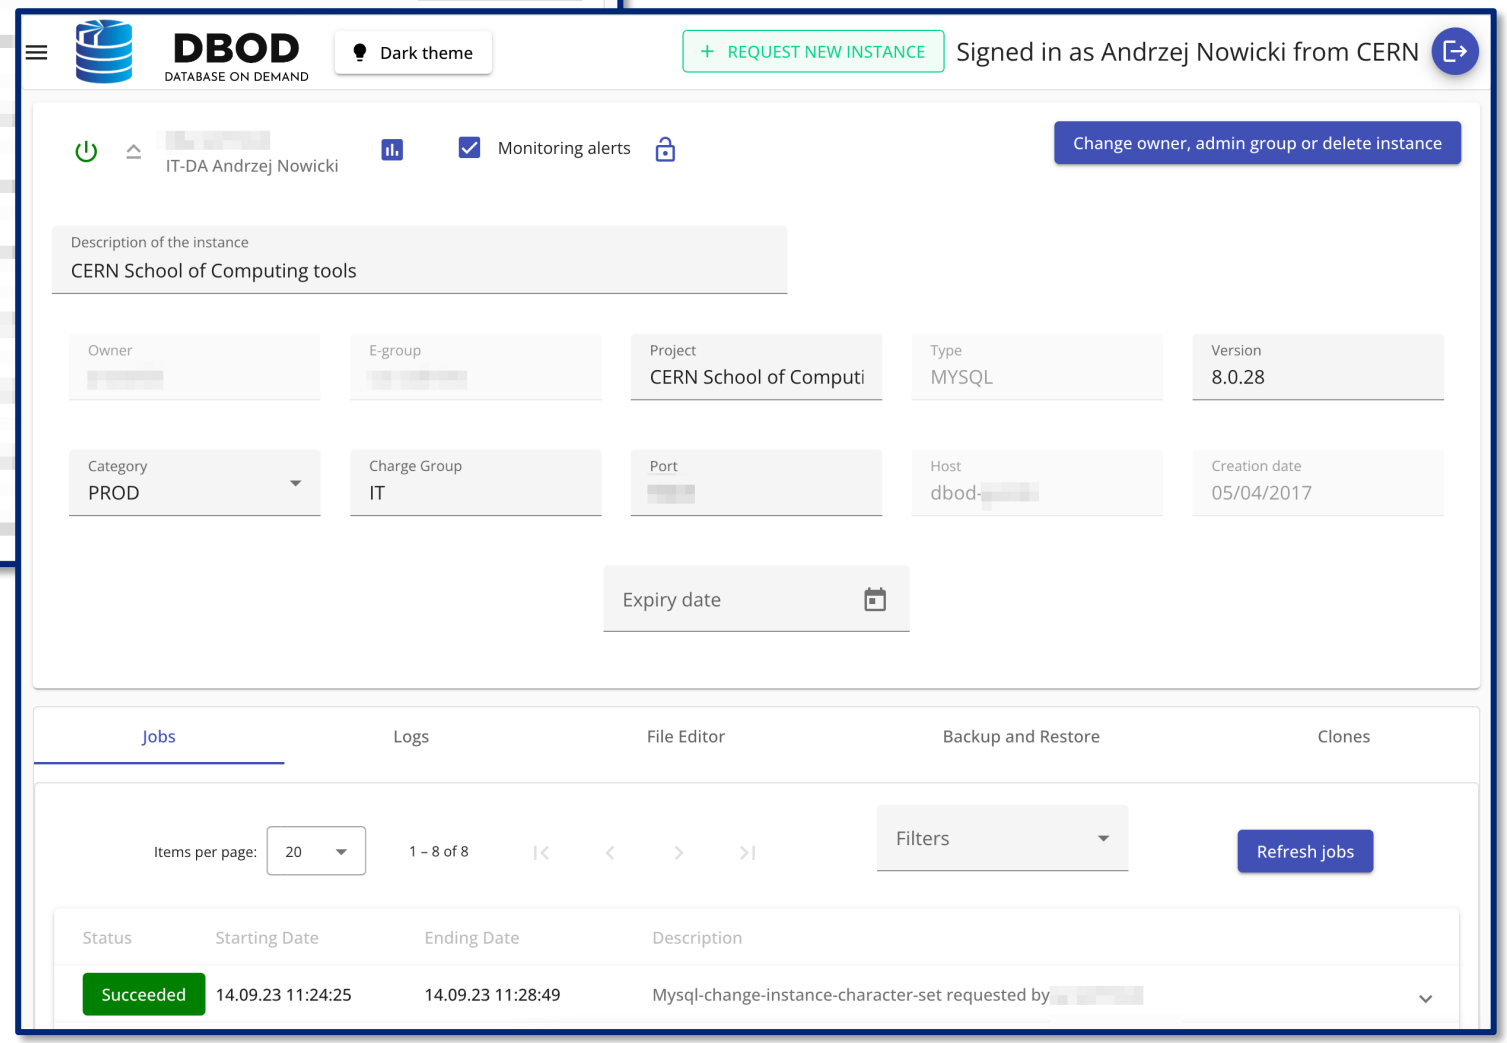

https://dbod.web.cern.ch/

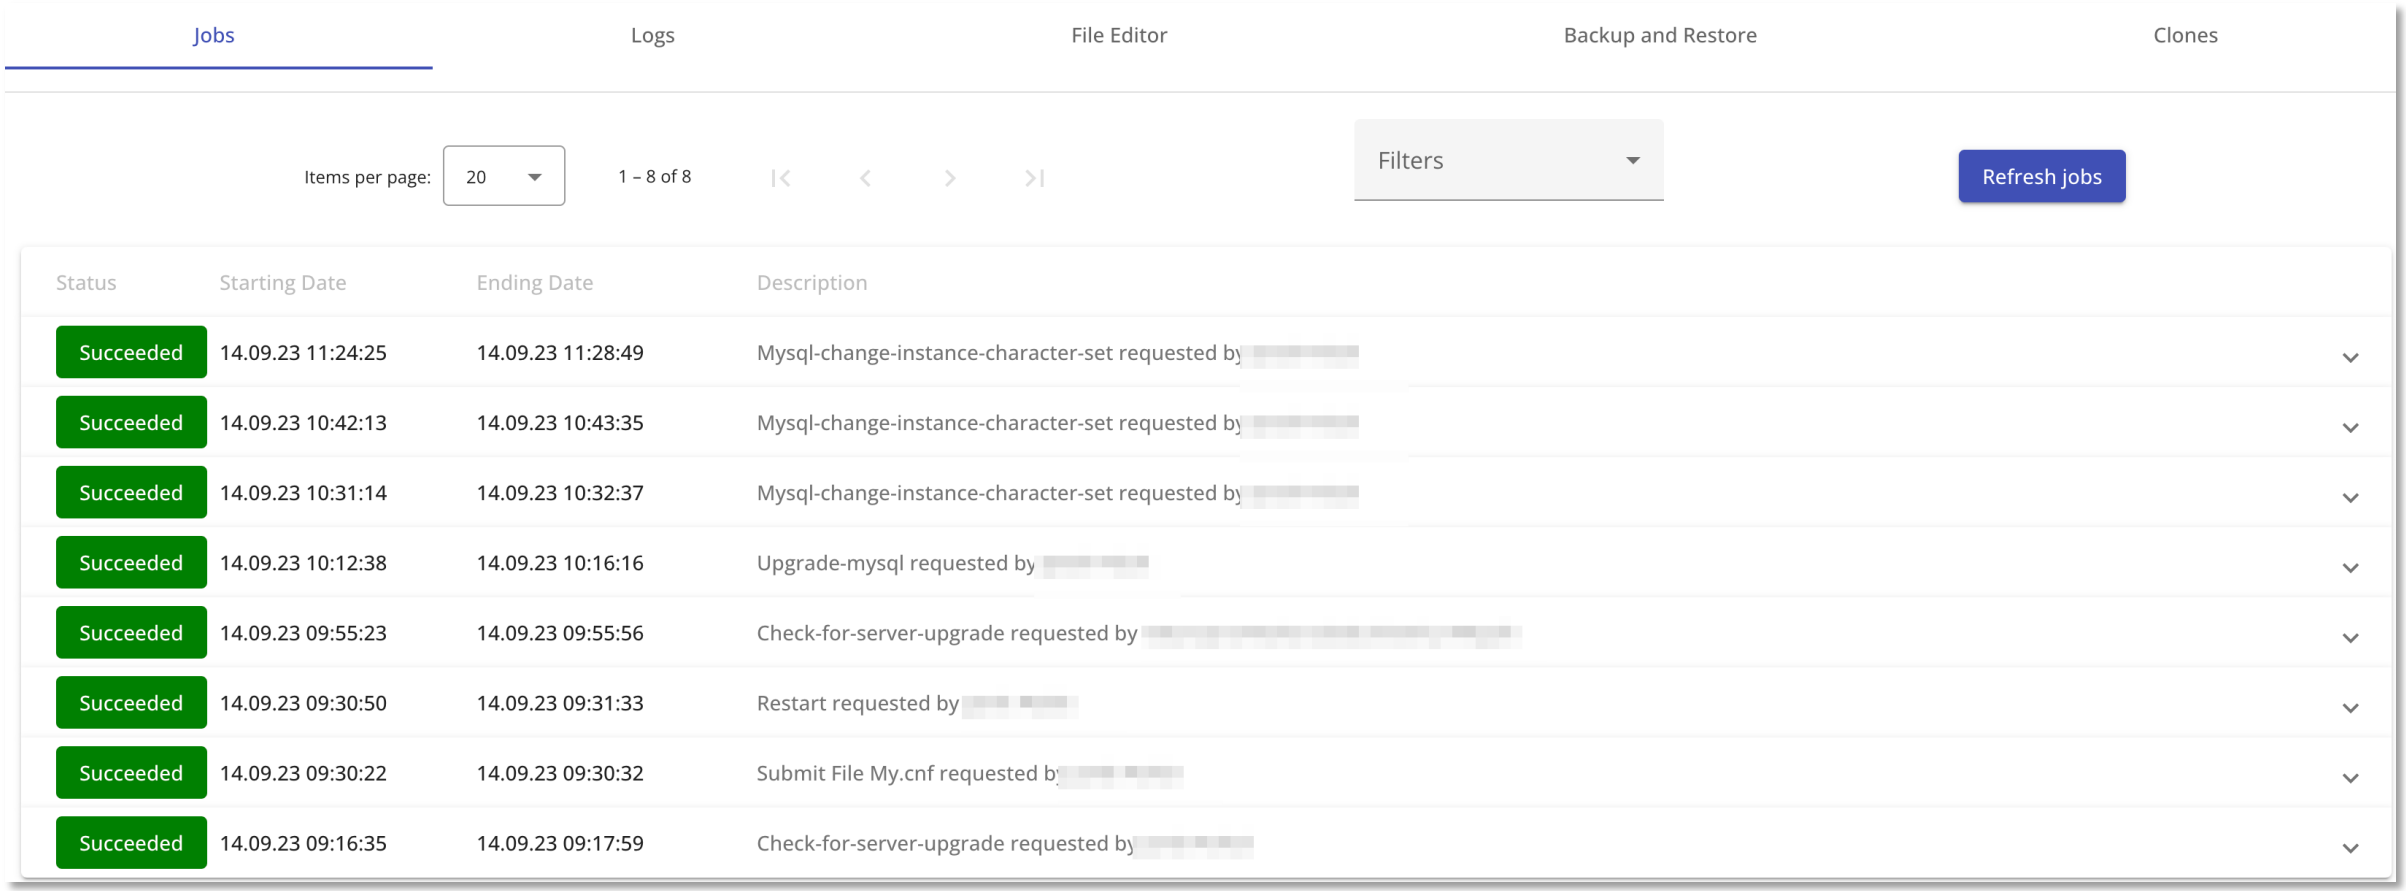

https://dbod.web.cern.ch/

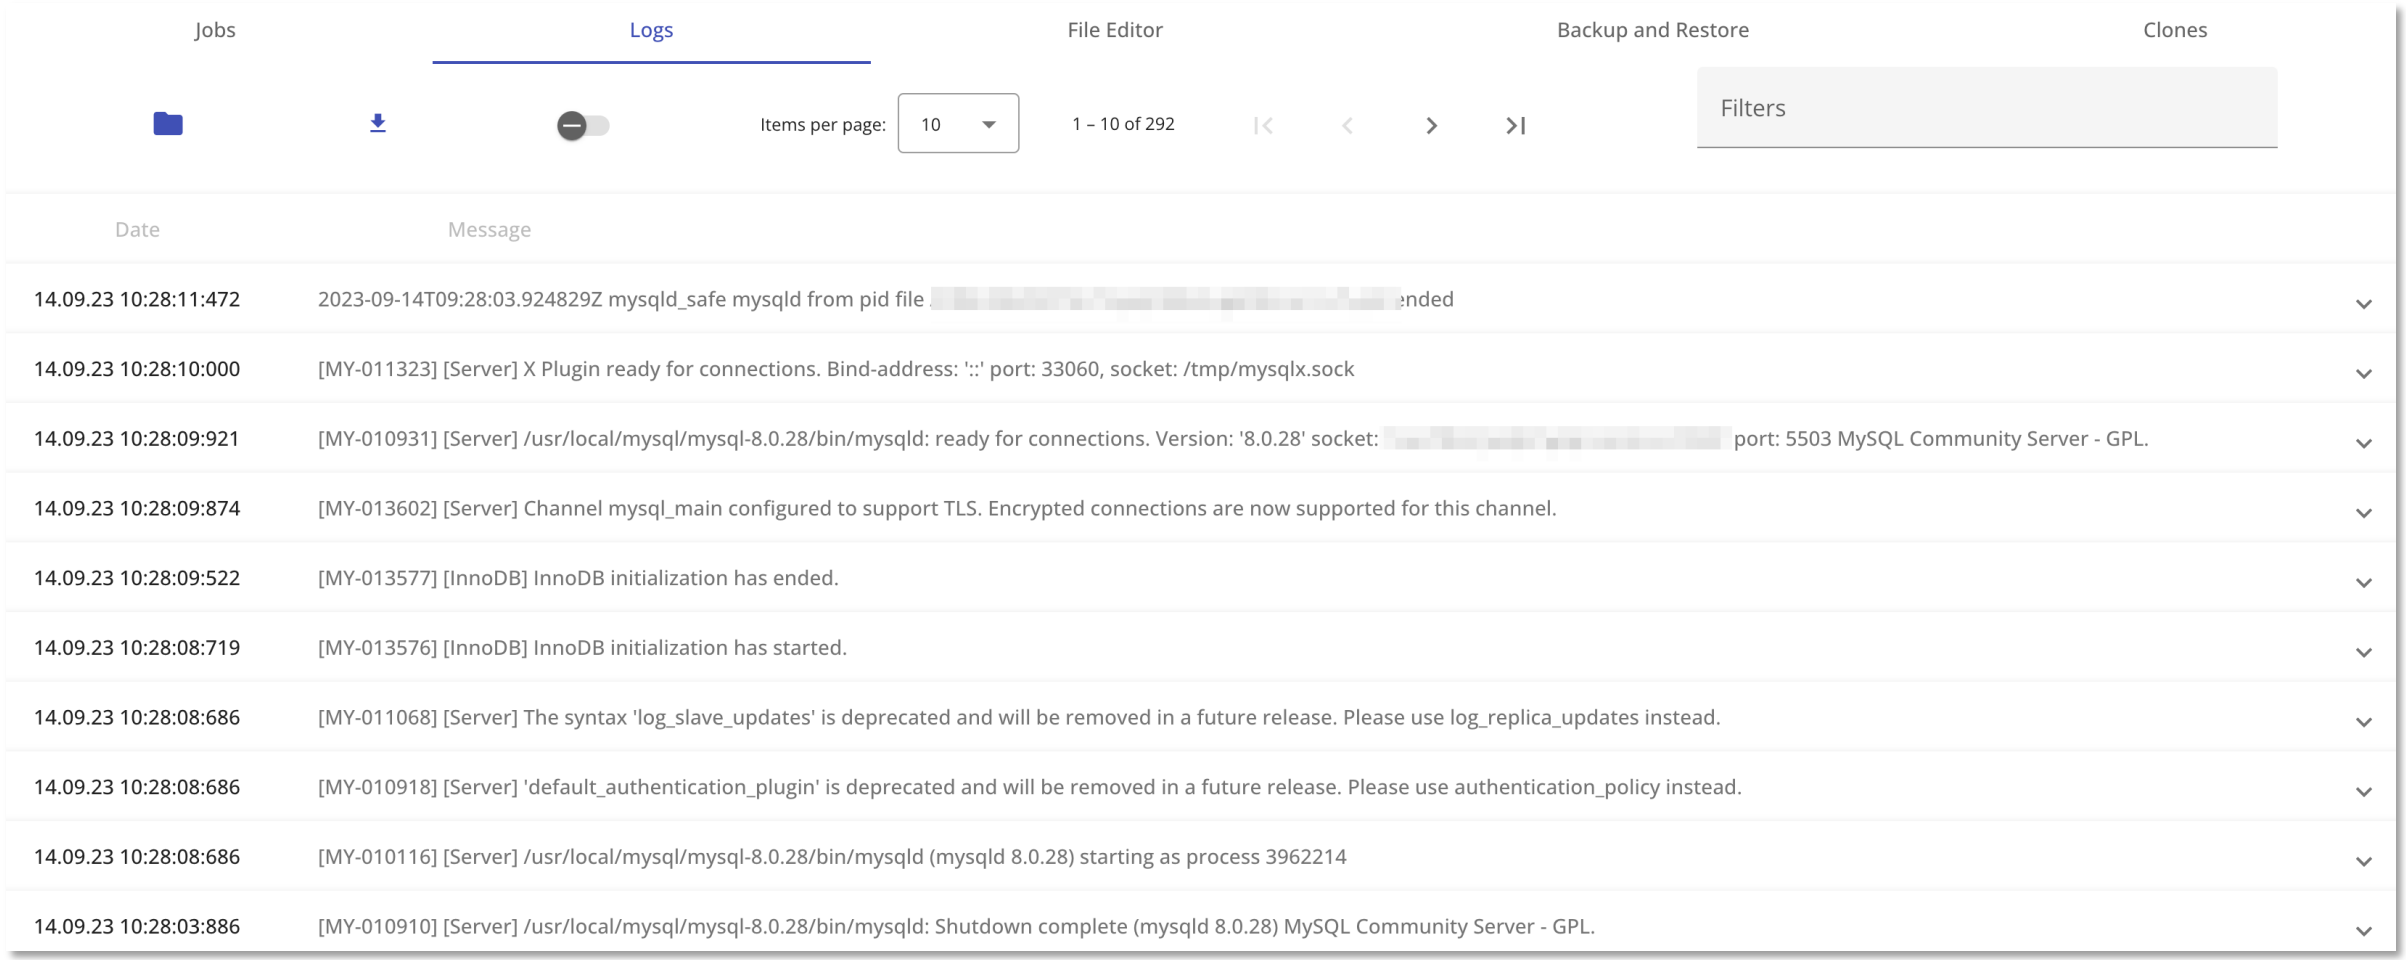

https://dbod.web.cern.ch/

DBoD Maintenance operations

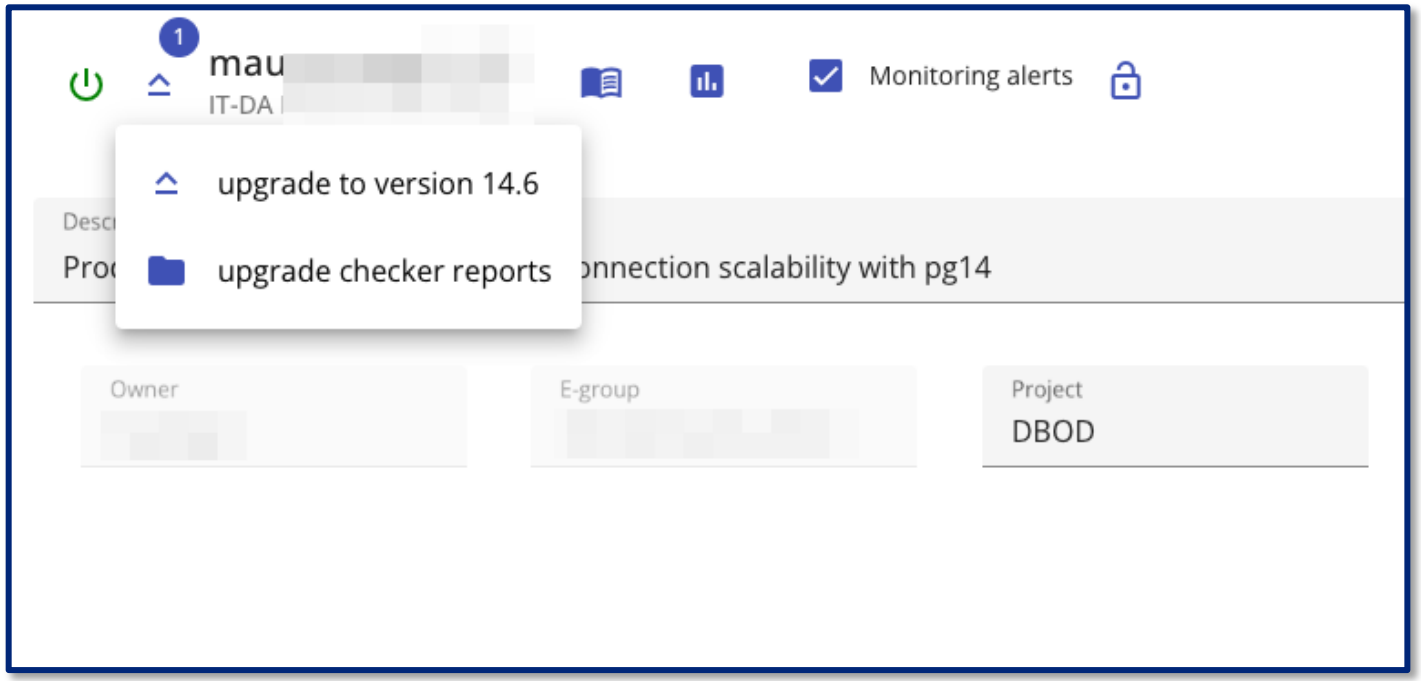

https://dbod.web.cern.ch/

DBoD Configuration management

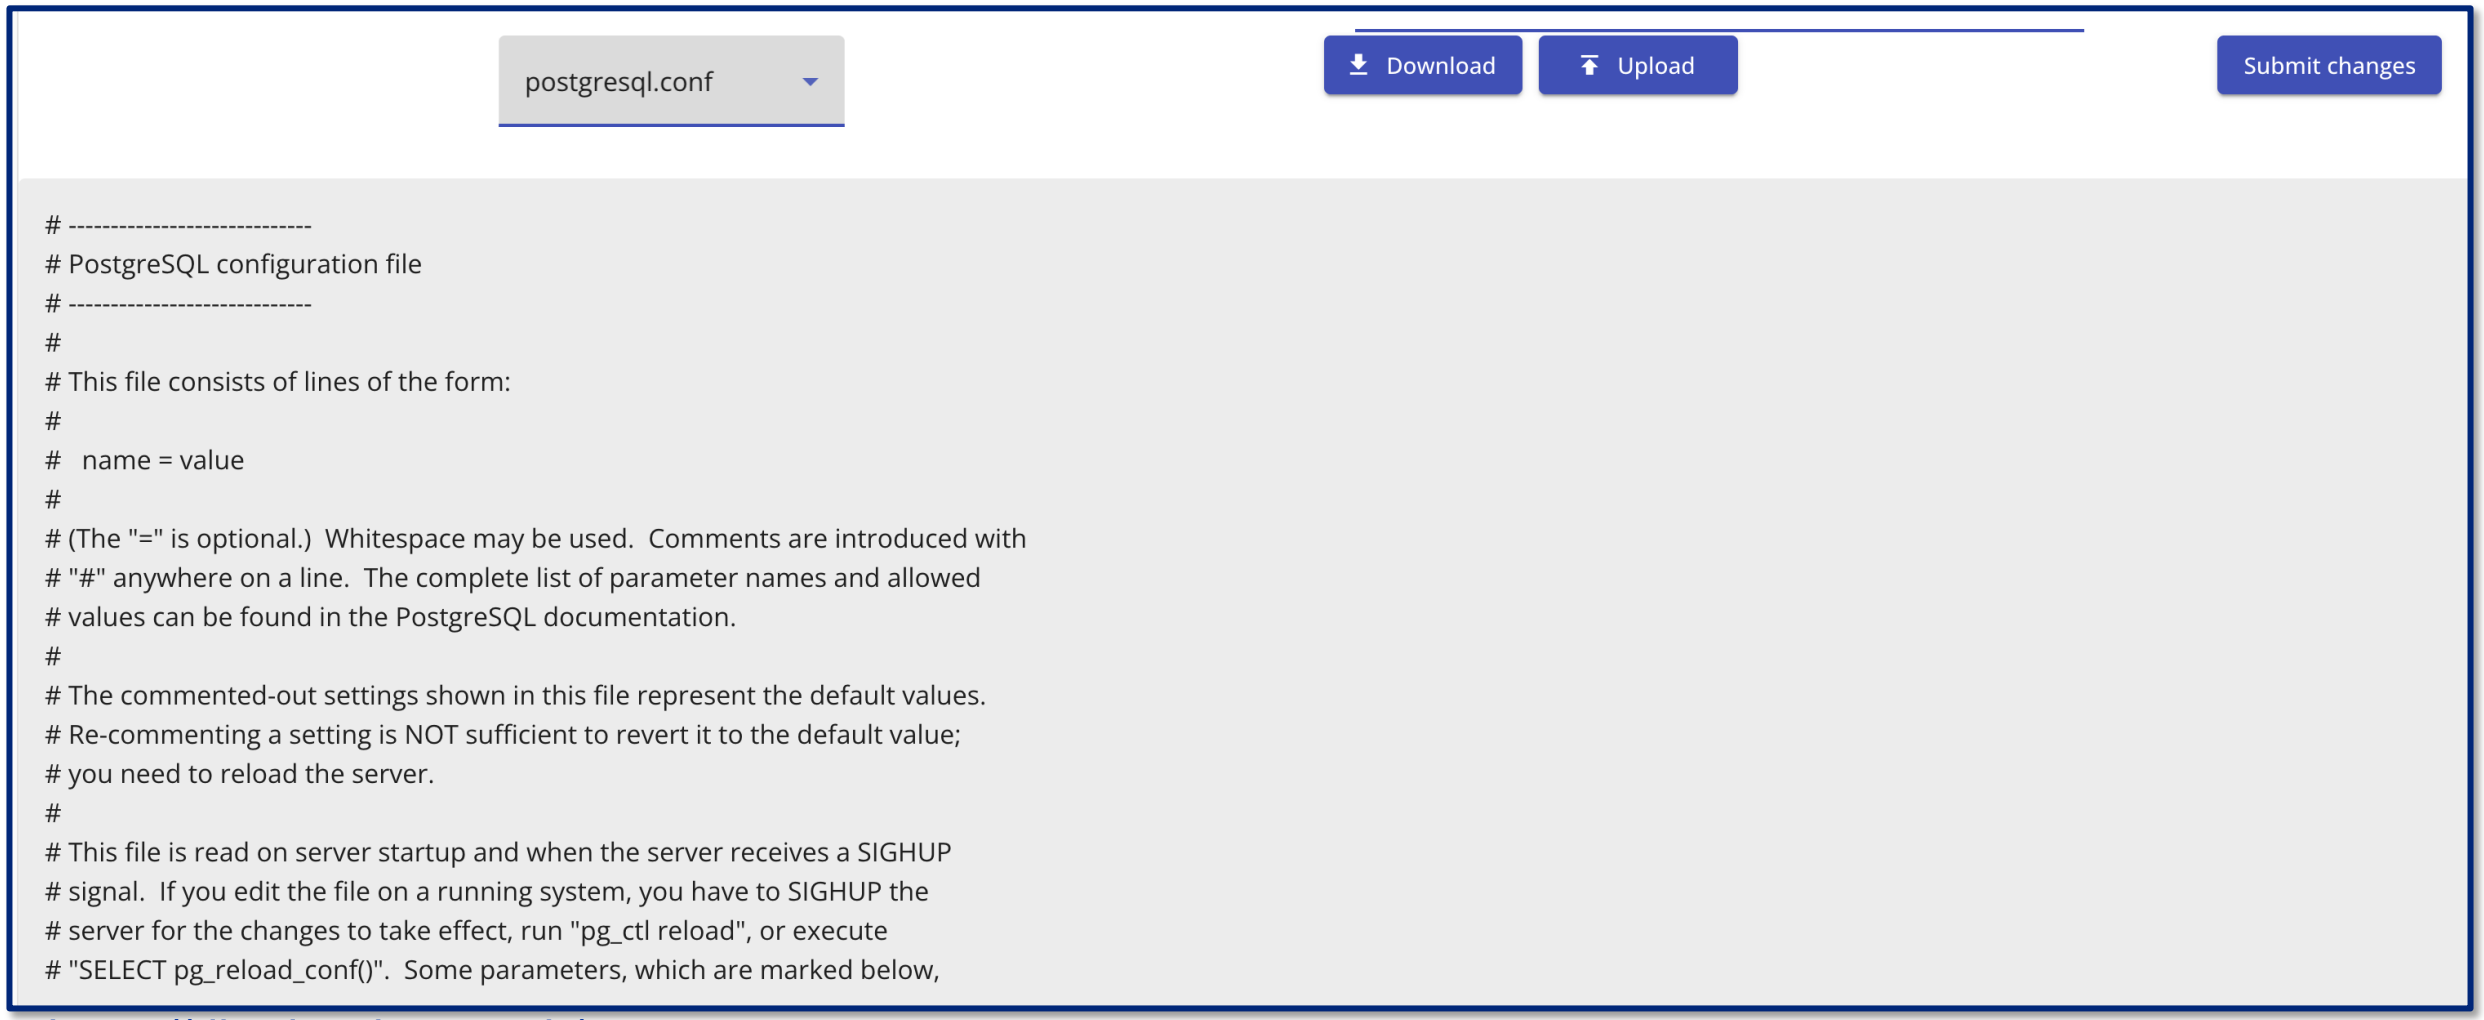

https://dbod.web.cern.ch/

System Memory

System Load

System CPU (%)

Kernel

System Metrics (Net, Disk IO)

Network usage

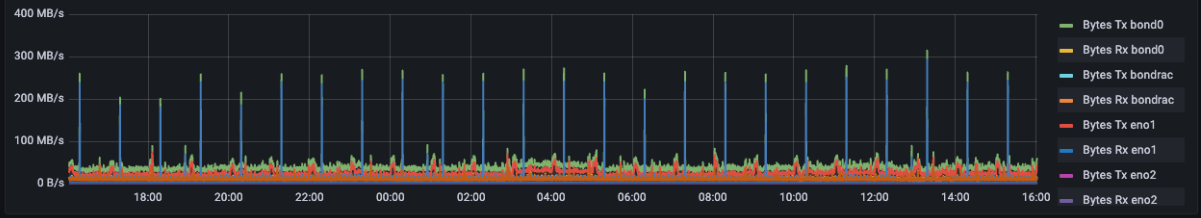

Disk IO

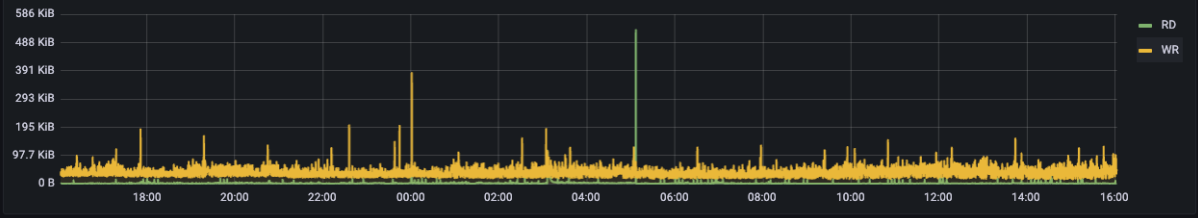

Database Process Metrics

Process CPU Time

Process MEM

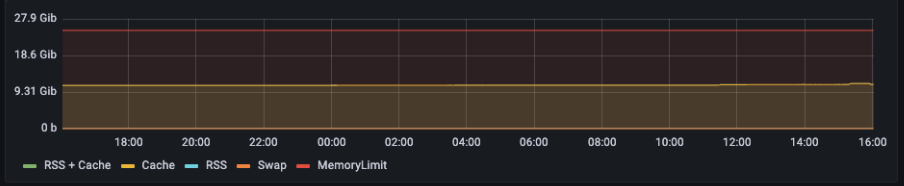

Replication lag

No data

Database Activity

Tuples

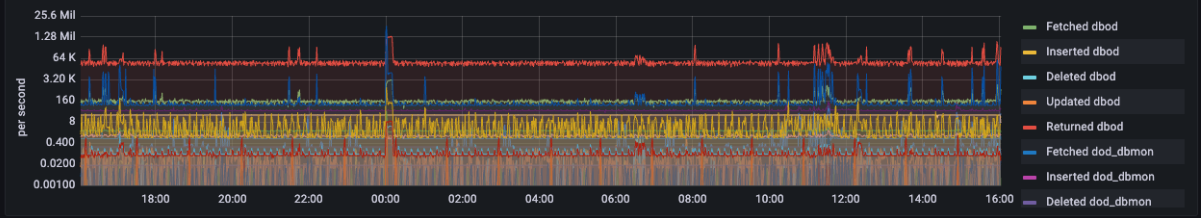

Transactions per Second

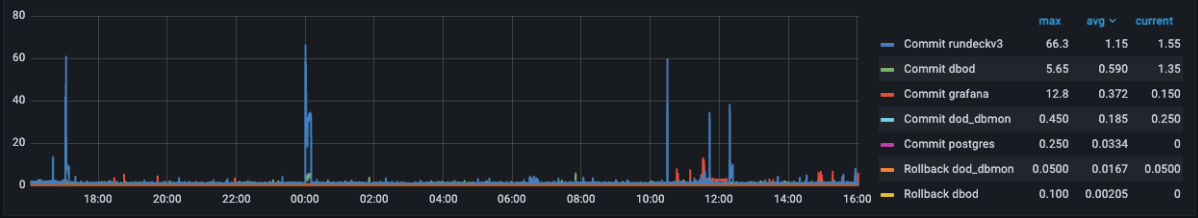

	max	avg	current
Commit rundeckv3	66.3	1.15	1.55
Commit dbod	5.65	0.590	1.35
Commit grafana	12.8	0.372	0.150
Commit dod_dbmon	0.450	0.185	0.250
Commit postgres	0.250	0.0334	0
Rollback dod_dbmon	0.0500	0.0167	0.0500
Rollback dbod	0.100	0.00205	0

Buffers

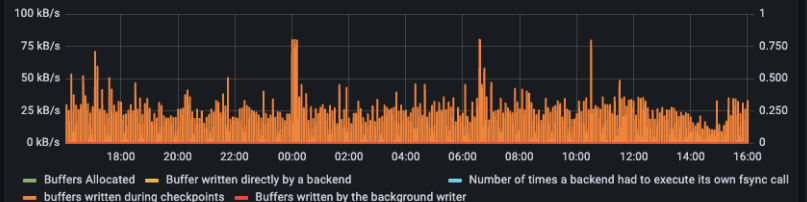

Checkpoint Time

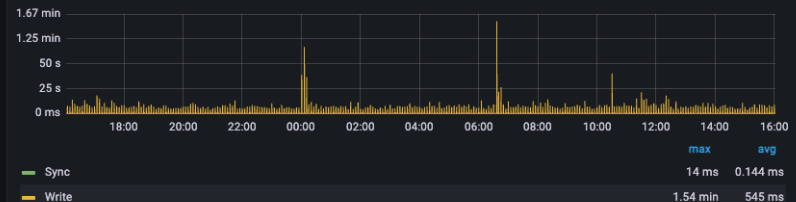

	max	avg
Sync	14 ms	0.144 ms
Write	1.54 min	545 ms

Blocks Hit (in Memory) /Read (on disk) per second

Services relying on DBoD

- Indico (meeting planning, room booking system)
- Autorization Service (SSO, 2FA)
- Configuration Management (puppet, foreman), Secrets Vault
- Jira, Gitlab, Openstack
- Websites (https://home.cern/)
- CERN Document Server (https://cds.cern.ch)
- File Transfer System
- ATLAS Panda (Production and Distributed Analysis System)
- and more…

DBoD architecture

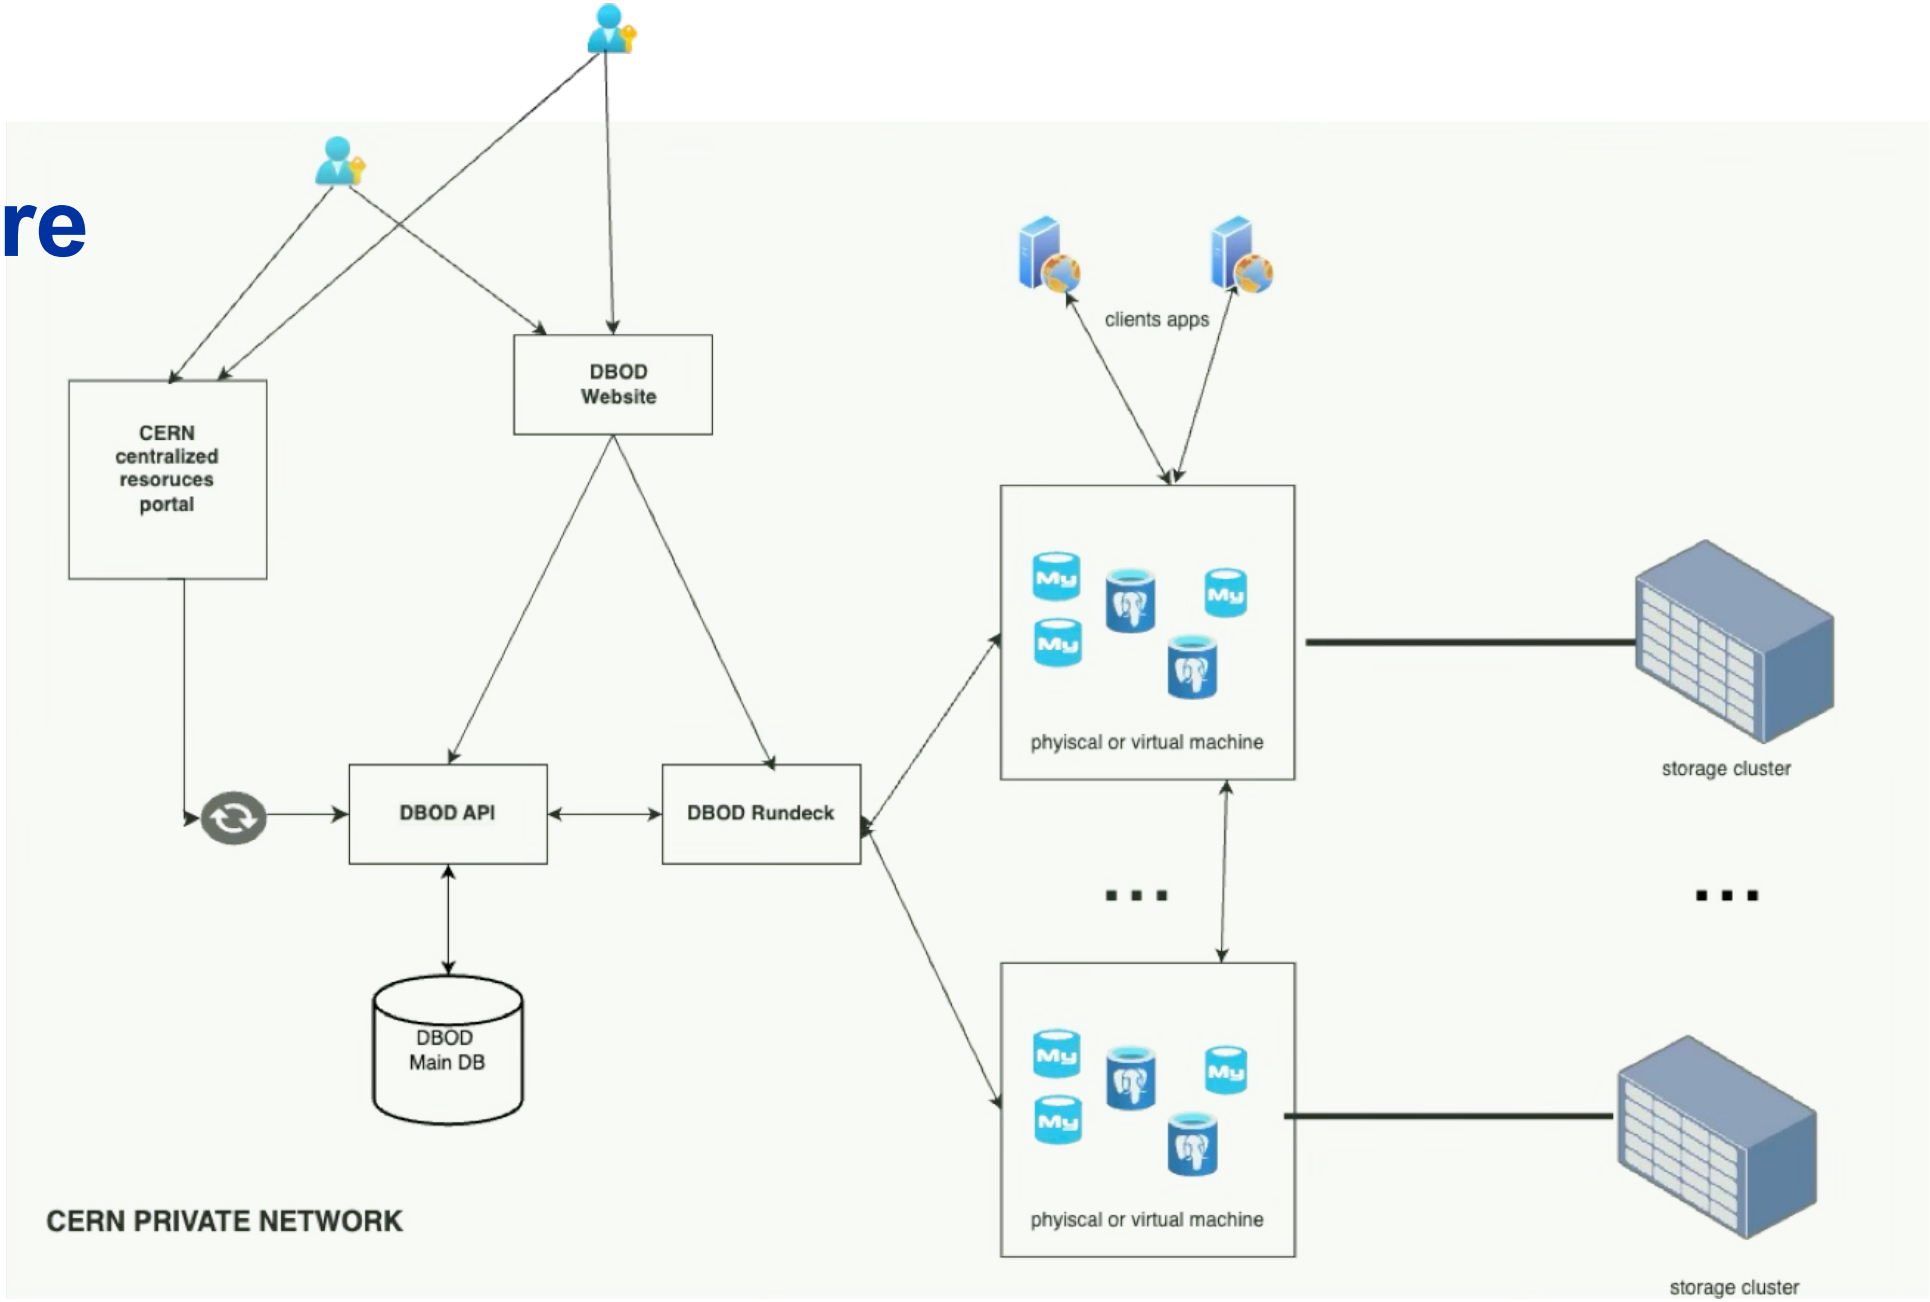

How to use DBoD?

DBoD manifesto

- https://dbod-user-guide.web.cern.ch/
- In the DBoD platform users are responsible for managing their databases, ensuring security and compliance with CERN policies, and handling application-level issues.
- The DBoD team, in contrast, provides and maintains the platform, handles infrastructure-level issues, ensures hardware reliability, performs automated backups, and offers tools for self-service.
- DBoD team doesn't provide support for application-level problems but ensure the database environment remains functional and up-to-date.

How to get access to DBoD?

- https://auth-resources.web.cern.ch/

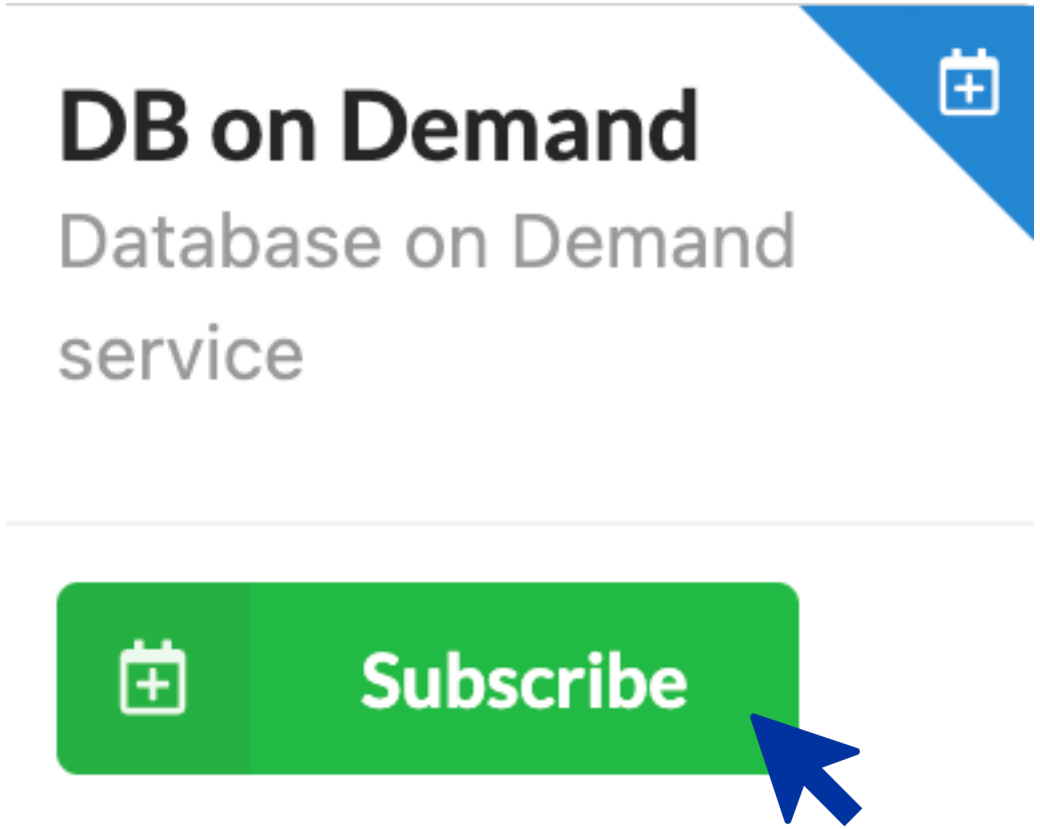

DBoD offer

- **PostgreSQL 13.13, 14.10, 15.7**
- **MySQL 8.0.35, 8.4.2**
- **InfluxDB 1.8.3**

DATA COLLECTED 10/2024

DB support schedule vs LHC schedule

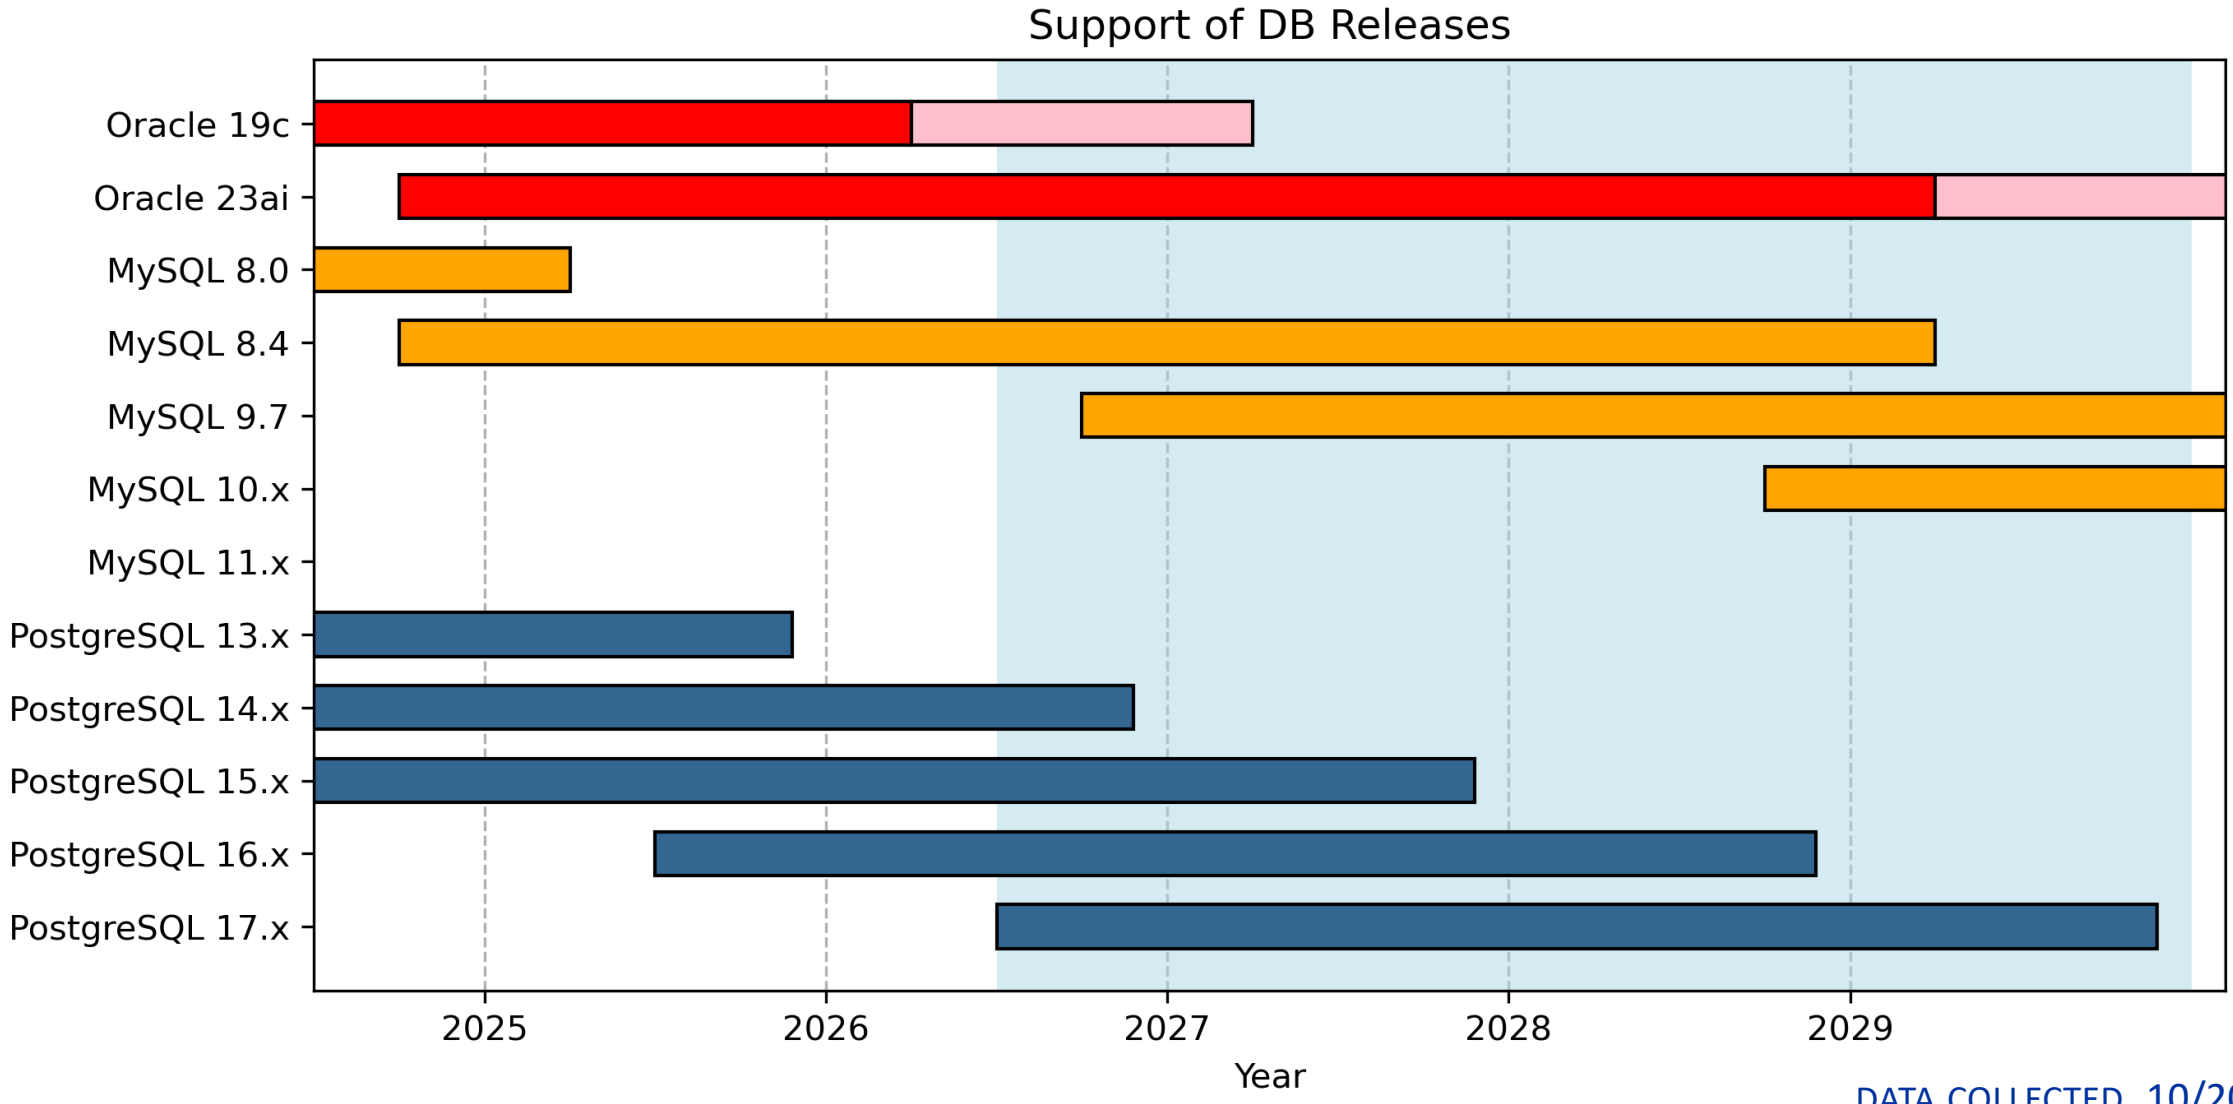

DATA COLLECTED 10/2024

DBoD creation

Create a new DB on Demand resource

* DB Name: anowicki_test01

* Category: PROD | TEST

* Admin group: csc-cc-admins

* Project: CSC

* Database type: PG / 13.13

* Description: Test for the preparation of CSC on IT services exercises

* Manifesto: By requesting a database, we consider you have fully read and accepted our manifesto.

Note: The database will be available after verification and approval from the DBOD team.

Save | Cancel

RQF2891802 "Request to create DB On Demand instance *anowicki_test01*"

CS

CERN Service Desk <service-desk@cern.ch>

Today at 15:08

To: Andrzej Nowicki

****
*** Please do not reply to this email unless you are unsatisfied with the proposed solution. ***
*** For feedback (e.g. a "thank you" or a complaint) please use the smiley at the bottom. ***
****

Dear Andrzej,
There is a proposed solution for you.

Ticket No: RQF2891802, **Opened:** 28-10-2024 15:01:24
Short description: Request to create DB On Demand instance *anowicki_test01*
Message from Pablo Conde Collados:

Dear Andrzej,

Your instance has been created.

You can connect to it using the following configuration parameters:

host: dbod-anowicki-test01.cern.ch, port: 6616
user: admin, password: changeme

Please follow the specific instructions for your database type described in the User Guide [1].

If you are currently connected to the DBOD web interface [3] and you cannot see your new instance, you will have to logout and login to the again, as your instances are loaded to your profile at login time.

Please do not hesitate to report to DB On Demand admins any feedback. You can report an incident or submit a request using the CERN Service Portal [2] at , or sending an email to service-desk@cern.ch stating "Database on demand Demand" as the service.

Both links are available in the DB On Demand Web Interface [3].

Please make sure a set of people are included in the admin e-group, rather than only the owner, to ensure important notifications and announcement are received and taken care by someone all the time and the instance is managed properly.

Best regards,

The DB On Demand team

[1] https://cern.ch/dbod-user-guide
[2] https://cern.service-now.com/service-portal?id=service_element&name=database-on-demand
[3] https://cern.ch/dbod

DBoD maintenace tasks

- Upgrade
- Backup
- Restore
- Clone
- Extension of the expiration date

DBoD upgrade

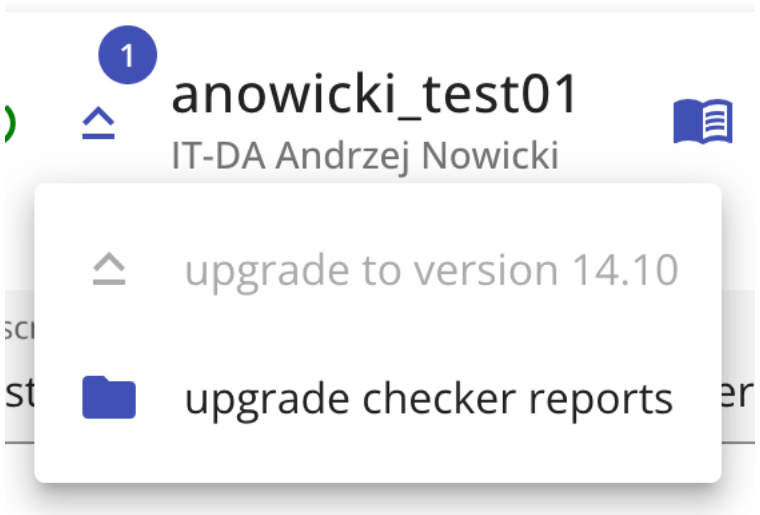

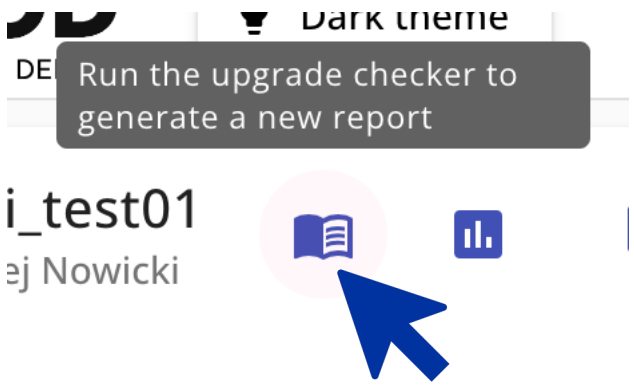

DBoD recap

- **The DBoD team provides infrastructure-level support**
- **DBoD team doesn’t provide traditional DBA support. You get full DBA rights.**
  - DBoD implements best configuration practices, etc.
  - It’s a shared service, so there are some limitations. Heavy users are normally identified and separated.
- **SNOW Service Element: Database on Demand Service**

Questions?

You'll work in pairs

- **Go to Google Sheets and write down your CERN Logins:**
- **https://cern.ch/dbod-allocation**

Before we finish

- **Go to the DBoD portal:**
- **https://dbod.web.cern.ch/**

Exercise 1. Requesting an instance
Go to the DBoD portal and request a new PgSQL instance.

DB Name:	csc_pg***<NN>***
Category:	TEST
Admin group*:	csc-cc-users
Database Type:	Postgres 13.13

The instance will be created and you will get an e-mail when your instance is ready.

DBoD exercises

- Requesting an instance
- Using the DBoD portal
- Logging into the MySQL/PostgreSQL instance
- Managing your instances
- Backups, Clones
- Upgrades

Thank you !

andrzejnowicki

andrzej.nowicki@cern.ch