



DB On Demand

What do you want from it?

`ignacio.coterillo@cern.ch`

29th May 2017

Table of Contents

Introduction

What is it?

Who uses it?

Main Offerings

Database types

Backup and Recovery functionality

Second level backup to EOS

High Availability

Monitoring

What do you need?

E.g. Time Series Databases

E.g. Technical Network instances

E.g. Migrating from Oracle?

E.g: Web Interface

Let us know!

Table of Contents

Introduction

What is it?

Who uses it?

Main Offerings

Database types

Backup and Recovery functionality

Second level backup to EOS

High Availability

Monitoring

What do you need?

E.g. Time Series Databases

E.g. Technical Network instances

E.g. Migrating from Oracle?

E.g: Web Interface

Let us know!

DB On Demand: What is it?

An operations automation platform

- ▶ Administrators: Configuration, creation, destruction, migration, ...
- ▶ Some operations exposed to users: backup, recovery, upgrades, ...
- ▶ (DB) System agnostic

What are the main differences from the User PoV (w/ Oracle service)?

- ▶ Instance owners get administrative privileges on their databases
- ▶ Some DBA classical operations are delegated to the Instance Owner.

Who uses it?

Number of instances per Group (FIM data, 523 total instances)




Table of Contents

Introduction

What is it?

Who uses it?

Main Offerings

Database types

Backup and Recovery functionality

Second level backup to EOS

High Availability

Monitoring

What do you need?

E.g. Time Series Databases

E.g. Technical Network instances

E.g. Migrating from Oracle?

E.g: Web Interface

Let us know!

Database types

For new instances

Relational

- ▶ MySQL 5.7
- ▶ PostgreSQL 9.6

Time Series

- ▶ InfluxDB 1.2

Backup and Recovery functionality

Based on storage snapshots

- ▶ Common work-flow for all supported database types
- ▶ Fast (seconds/minutes)
- ▶ Point in time recovery supported for MySQL and PostgreSQL

About the storage

- ▶ NetApp based. Cluster configuration highly available and capable of surviving multiple disk failures on the disk aggregates
- ▶ Using two separate volumes. One for **datafiles** and one for **Write Ahead Logs (WAL)**

NEW: Second level backup to EOS

For Production instances

WAL logs Archiving

Allow for longer recovery windows (Going back further in time) while:

- ▶ Maintaining (or reducing) NetApp footprint
- ▶ Reducing configuration complexity

Periodical Full database export

Full database (copied from a snapshot) + WAL logs

High Availability

	MySQL	PostgreSQL	InfluxDB
Master-Slave Replication	Used	Used	TBD ¹
Proxy	Implemented	Implemented	TBD
Multi-master (Cluster)	First user tests	Delayed ²	Not yet ³
Sharding		Citus Extension (v.9.6.2)	

¹influxdb-relay

²Depends on BDR 2.0 by 2ndQuadrant

³Won't look at it until use cases come



New Monitoring system

- ▶ Based on centralized services (OpenShift, DBOD, Monitoring)
- ▶ Single point of access for:
 - ▶ System metrics
 - ▶ Database metrics
 - ▶ NFS Volume metrics
 - ▶ **WIP**: Database logs
- ▶ Already available for InfluxDB and **TEST** MySQL and PostgreSQL instances

> System Metrics (Mem, Load, CPU)

> System Metrics (Net, Disk IO)



> Instance Process



> Database



> Queries



> HTTP

> Collectd

> NFS: Data Volume






Table of Contents

Introduction

What is it?

Who uses it?

Main Offerings

Database types

Backup and Recovery functionality

Second level backup to EOS

High Availability

Monitoring

What do you need?

E.g. Time Series Databases

E.g. Technical Network instances

E.g. Migrating from Oracle?

E.g: Web Interface

Let us know!

What do you need?

Please let us know

- ▶ Do you have additional requirements for the current service offerings?
- ▶ Do you use or plan to use a different database for your system?
- ▶ Do you think more people are/would be interested?

E.g. Time Series Databases

How it happened?

- ▶ Several users asked us if we were going to support TSDB (different backgrounds)
- ▶ We went around looking for additional uses cases looking for critical mass
- ▶ Internal evaluation of different options available at the time (Prometheus, OpenTSDB, Graphite/Carbon, InfluxDB)
- ▶ ~ Two months later, the firsts instances were offered to users

E.g. Technical Network instances

Collecting Requests

- ▶ Two individual requests during the last two years → No action
- ▶ Several requests in the last couple of months → Where we are

Where we are

- ▶ Analyzing requirements
- ▶ Studying integration of TN based servers in the platform
- ▶ Working in resource allocation

E.g. Migrating from Oracle?

Misc

- ▶ There are tools to help migrating database objects. Results may vary
- ▶ Database side logic (PL/SQL) → For PostgreSQL, there are companies⁴ offering compatibility extensions which emulate Oracle PL/SQL modules

PostgreSQL Foreign Data Wrappers

- ▶ Like Oracle DB links
- ▶ We use them (DBOD) as part of migrating the DB On Demand backend database from Oracle to PostgreSQL

⁴Enterprise DB

E.g: Web Interface

Will be reimplemented during Summer

Feature Requests: <https://github.com/cerndb/dbod-webapp>

Feature requests:

- ▶ Support for named snapshots
- ▶ Admin password reset from the interface

Let us know!

Administrators

- ▶ dbondemand-admin@cern.ch
- ▶ SNOW Ticket to the **Database on Demand Instances FE**

General Discussion

- ▶ dbondemand-users@cern.ch
- ▶ Mattermost DBOD Channel (IT-Dep Team)



www.cern.ch