

# Database On Demand Service (DBOD)

Daniel Gomez  
Ruben Gaspar  
Ignacio Coterillo\*  
Dawid Wojcik

\*CERN/CSIC funded by Spanish Ministry of Research EIC 2010 program. Researcher code  
SEII-2010-00237



- Objectives
- Main functionality
- Architecture
- Implementation
- Current status



- Allow the CERN user community to **request and manage database instances**
- Give users **full DBA privileges** on their instances
- Offer **different DBMS vendors**
  - MySQL and Oracle are our starters, but the service is designed to scale easily
- The DBOD service takes care of system administration operations and providing **upgrade functionality for DBMS**
- The **DBOD service does not provide DBA or application support**
- Users are **entirely responsible** for ensuring that their systems, and the use of their systems, are fully compliant with the Rules for the use of CERN Computing Facilities.









## DB On Demand

[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco ▾

Are you experiencing problems with DB On Demand? [Contact us](#)[Home](#) [Admin](#) [Help](#) [Admin Help](#)

## Actions

Star up



Shut down



File mngmt



Backup mngmt



Restore



Upgrade



Monitoring



## Information for instance testinstance

Refresh instance

**State:** Running **DB Type:** MySQL**Creation Date:** 24/11/2011**Category:** Test**NO Connections:** 50**Description:** Test instance until February edited.**Username:** icoteril**e-Group:** it-dep-db**Expiry Date:** 08/02/2012**Project:** DBOD**DB Size:** 500 GB

## Information for job

Create a backup 09/12/2011 9:09:36 ▾



**State:** OK **Creation Date:** 09/12/2011 9:09:36**Log:**

```

Fri Dec 9 09:09:39 CET 2011 : Main: Starting
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.GetGeneralVars: host <dbvrtg046.cern.ch>
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.GetGeneralVars: port <5500>
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.GetGeneralVars: socket </var/lib/mysql/mysql.sock>
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.GetGeneralVars: mysql_datafiles </ORA/dbs03/DRUPAL/mysql>
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.GetGeneralVars: backup_dest </ORA/dbs02/DRUPAL/mysql>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_user_nastorag>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_db_dod_mysql>
Fri Dec 9 09:09:40 CET 2011 : Main: my.cnf </ORA/dbs03/DRUPAL/mysql/my.cnf>
Fri Dec 9 09:09:40 CET 2011 : RunTime.GetControllerVol: searching for <(^dbnasw+);/vol/(.*?03)/b>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /bin/mount
Fri Dec 9 09:09:40 CET 2011 : Main: working with following controller:volume <dbnasb301:dodtest03>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running ps -elf | grep -i datadir=/ORA/dbs03/DRUPAL/mysql | grep -v grep
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock </ORA/dbs01/syscontrol/projects/dod/bin/actual_binarylog.sql
Fri Dec 9 09:09:40 CET 2011 : Main: <binlog>.<000299> actual binary log sequence
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -pXXXXXXXXX --socket=/var/lib/mysql/mysql.sock </ORA/dbs01/syscontrol/projects/dod/bin/presnap.sql
Fri Dec 9 09:09:40 CET 2011 : Main: presnap actions completed successfully.
Fri Dec 9 09:09:40 CET 2011 : Main: controller: <dbnasb301> volume: <dodtest03>
Fri Dec 9 09:09:40 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0x1faccd70)>
Fri Dec 9 09:09:40 CET 2011 : GeneralMysql.RunSSHStr running snap create dodtest03 snapscrip_09122011_090940_300
Fri Dec 9 09:09:43 CET 2011 : Main: stdout:
Fri Dec 9 09:09:43 CET 2011 : Main: creating snapshot...
Fri Dec 9 09:09:43 CET 2011 : Main: snapshots have been created!
Fri Dec 9 09:09:43 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock </ORA/dbs01/syscontrol/projects/dod/bin/postsnap.sql
Fri Dec 9 09:09:43 CET 2011 : Main: postsnap actions completed successfully.
Fri Dec 9 09:09:43 CET 2011 : Main: mysql_snapshot is over.
Fri Dec 9 09:09:43 CET 2011 : mysql_snapshot.Main: State: [0]

```





DB On Demand








Home Admin Help Admin Help

CERN Home | IT Home | IT-DB Home | Phonebook

Welcome Daniel Gomez Blanco ▾


Are you experiencing problems with DB On Demand? [Contact us](#)

Actions




Information for instance testinstance


Refresh Instance ↻


State: Running 


DB Type: MySQL


Creation Date: 24/11/2011


Category: Test 


NO Connections: 50 


Description: Test instance until February edited. 

Username: icoleril 


e-Group: it-dep-db 

Expiry Date: 08/02/2012 

Project: DBOD 

DB Size: 500 GB 


Information for job Create a backup 09/12/2011 9:09:36

State: OK 

Creation Date: 09/12/2011 9:09:36

Log:

Single backup


 Create a single backup now


Backup configuration

☐ Perform automatic backups every 24 hours

☐ Backup database to tape every week starting on 11/04/2012 12:18:14

Please contact us before enabling this option for the first time to prepare the system for tape backups

Apply changes 

Close 

Fri Dec 9 09:09:39 CET 2011 : Main: <binlog>: <000299> actual binary log sequence

Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod\_mysql -pXXXXXXXXX --mysql\_socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin /actual\_binarylog.sql

Fri Dec 9 09:09:40 CET 2011 : Main: <binlog>: <000299> actual binary log sequence

Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod\_mysql -pXXXXXXXXX --socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin /presnap.sql

Fri Dec 9 09:09:40 CET 2011 : Main: presnap actions completed successfully.

Fri Dec 9 09:09:40 CET 2011 : Main: controller: <dbnasb301> volume: <dodtest03>

Fri Dec 9 09:09:40 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0x1faccd70)>

Fri Dec 9 09:09:40 CET 2011 : GeneralMysql.RunSSHStr running snap create dodtest03 snapscrip\_09122011\_090940\_300

Fri Dec 9 09:09:43 CET 2011 : Main: stdout:

Fri Dec 9 09:09:43 CET 2011 : Main: creating snapshot...

Fri Dec 9 09:09:43 CET 2011 : Main: snapshots have been created!

Fri Dec 9 09:09:43 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod\_mysql -pXXXXXXXXX --mysql\_socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin /postsnap.sql

Fri Dec 9 09:09:43 CET 2011 : Main: postsnap actions completed successfully.

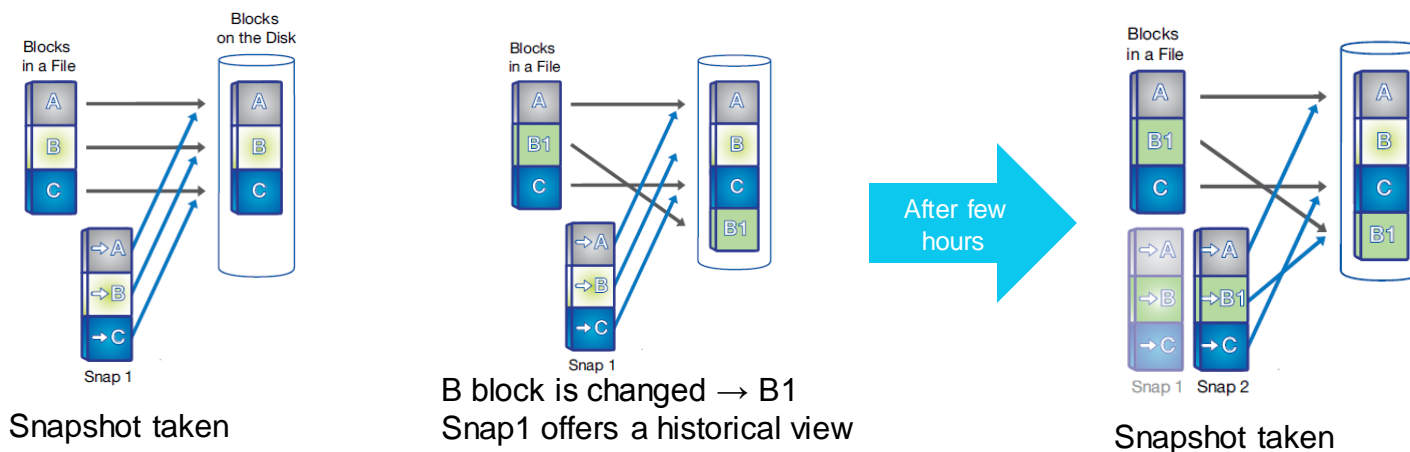
Fri Dec 9 09:09:43 CET 2011 : Main: mysql\_snapshot is over.

Fri Dec 9 09:09:43 CET 2011 : mysql\_snapshot.Main: State: [0]

6





- Based on snapshots.
- Full database backup in few seconds.



- How many snapshots are stored? For how long?
  - Extra 20% space provided for snapshots (e.g.: active file system 100GB, snapshot area 20 GB → total: 120GB)
  - Depends on DBOD instance activity (e.g.: changing 2% active file system will allow to have 10 days of snapshots)
- Possibility to send to a backup to tape once a week
- For MySQL: InnoDB storage engine recommended










[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco


## DB On Demand

[Home](#)
[Admin](#)
[Help](#)
[Admin Help](#)

Are you experiencing problems with DB On Demand? [Contact us](#)

**Actions**








**Information for instance testinstance**
[Refresh Instance](#)

**State:** Running 

**DB Type:** MySQL

**Creation Date:** 24/11/2011

**Category:** Test

**NO Connections:** 50

**Description:** Test instance until February edited.

**Username:** icoleril

**e-Group:** it-dep-db


**Expiry Date:** 08/02/2012

**Project:** DBOD

**DB Size:** 500 GB

**Information for job**

Create a backup 09/12/2011 9:09:36

**State:** OK 

**Creation Date:** 09/12/2011 9:09:36

**Log:**

```

Fri Dec 9 09:09:39 CET 2011 : Main: Start
Fri Dec 9 09:09:39 CET 2011 : GeneralM
Fri Dec 9 09:09:39 CET 2011 : GeneralM
Fri Dec 9 09:09:39 CET 2011 : GeneralM
Fri Dec 9 09:09:39 CET 2011 : GeneralM
Fri Dec 9 09:09:39 CET 2011 : GeneralM
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : Main: my
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : Main: wor
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : RunTime
Fri Dec 9 09:09:40 CET 2011 : Main: <binlog>.<000299> actual binary log sequence
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXXX --socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin/presnap.sql
Fri Dec 9 09:09:40 CET 2011 : Main: presnap actions completed successfully.
Fri Dec 9 09:09:40 CET 2011 : Main: controller: <dbnasb301> volume: <dodtest03>
Fri Dec 9 09:09:40 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0x1faccd70)>
Fri Dec 9 09:09:40 CET 2011 : GeneralMysql.RunSSHStr running snap create dodtest03 snapscrip_09122011_090940_300
Fri Dec 9 09:09:43 CET 2011 : Main: stdout:
Fri Dec 9 09:09:43 CET 2011 : Main: creating snapshot...
Fri Dec 9 09:09:43 CET 2011 : Main: snapshots have been created!.
Fri Dec 9 09:09:43 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXXX --socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin/postsnap.sql
Fri Dec 9 09:09:43 CET 2011 : Main: postsnap actions completed successfully.
Fri Dec 9 09:09:43 CET 2011 : Main: mysql_snapshot is over.
Fri Dec 9 09:09:43 CET 2011 : mysql_snapshot.Main: State: [0]
          
```

**Dispatch a restore**

Requester: dgomezbl

**Available snapshots:**

Apr 2012						
Mon	Tue	Wed	Thu	Fri	Sat	Sun
26	27	28	29	30	31	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	1	2	3	4	5	6

**Snapshots for selected day:**  
[13:49:32](#)

**Select a day and time to restore to:**

11/04/2012

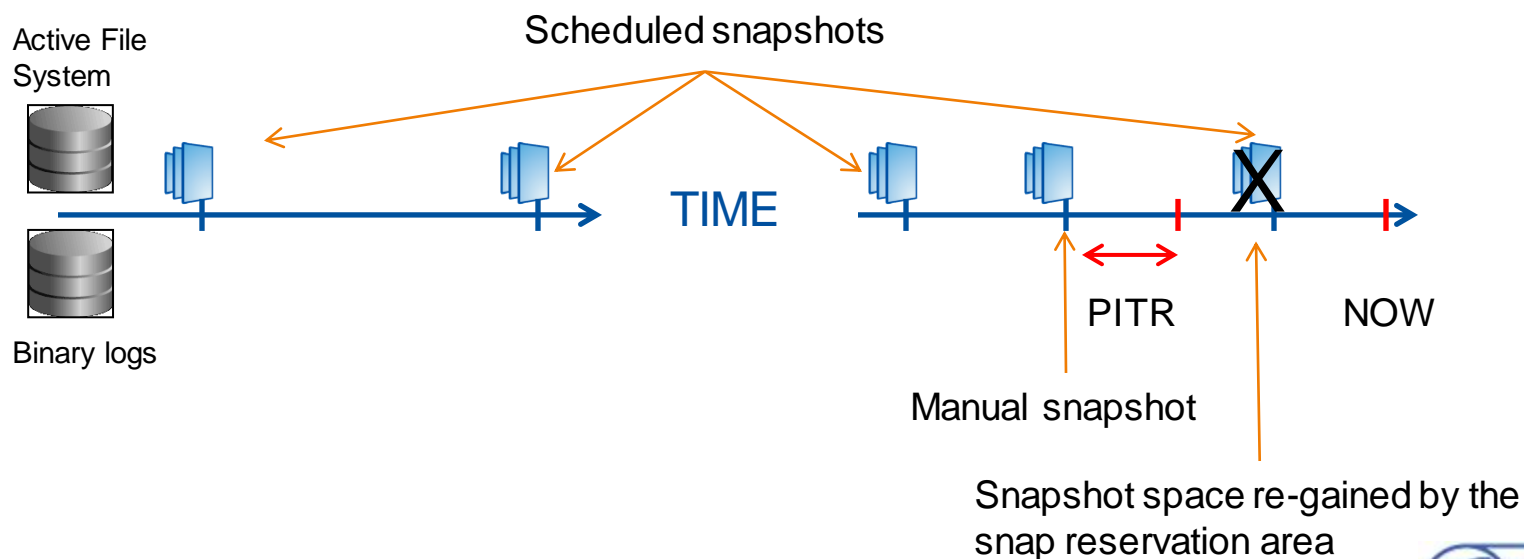
13:49:32

Cancel



Accept



- Based on snapshots
  - Restore time reduced to **a few seconds**
- Full restore of whole database
- DBOD instances configured with binary logs
  - Point-In-Time Recovery












[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco

## DB On Demand

[Home](#)
[Admin](#)
[Help](#)
[Admin Help](#)

Are you experiencing problems with DB On Demand? [Contact us](#)

**Actions**








**Information for instance testinstance**
[Refresh Instance](#)

**State:** Running
 **DB Type:** MySQL
 **Creation Date:** 24/11/2011
 **Category:** Test
 **NO Connections:** 50
 **Description:** Test instance un

**Username:** icoteril
 **e-Group:** it-dep-db
 **Expiry Date:** 08/02/2012
 **Project:** DBOD
 **DB Size:** 500 GB

**Information for job**

**State:** OK
 **Creation Date:** 09/12/2011
 **Log:**

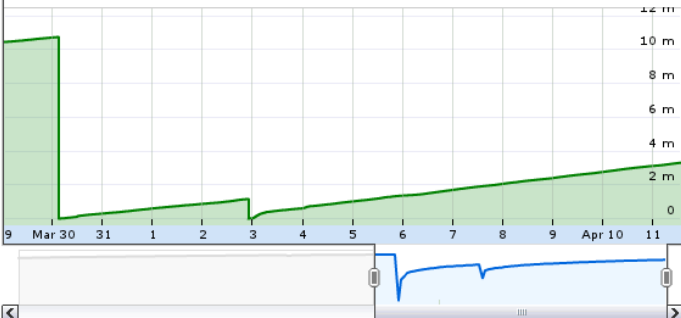
```

Fri Dec 9 09:09:39 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0x1faccd70)>
Fri Dec 9 09:09:40 CET 2011 : GeneralMysql.RunSSHStr running snap create dodtest03 snapscrip_09122011_090940_300
Fri Dec 9 09:09:43 CET 2011 : Main: stdout:
Fri Dec 9 09:09:43 CET 2011 : Main: creating snapshot...
Fri Dec 9 09:09:43 CET 2011 : Main: snapshots have been created!
Fri Dec 9 09:09:43 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin /postsnap.sql
Fri Dec 9 09:09:43 CET 2011 : Main: postsnap actions completed successfully.
Fri Dec 9 09:09:43 CET 2011 : Main: mysql_snapshot is over.
Fri Dec 9 09:09:43 CET 2011 : mysql_snapshot.Main: State: [0]
          
```

**Display metrics**

**Select metric to display:**
Number of active connections

**Zoom:** 1' 5' 1h 1d 5d 1m 3m 6m 1y Max
 13:53 April 11, 2012
 3,308,156



You can also monitor your instance on [Lemon Monitoring](#)

Close



- Service Level Status (SLS)

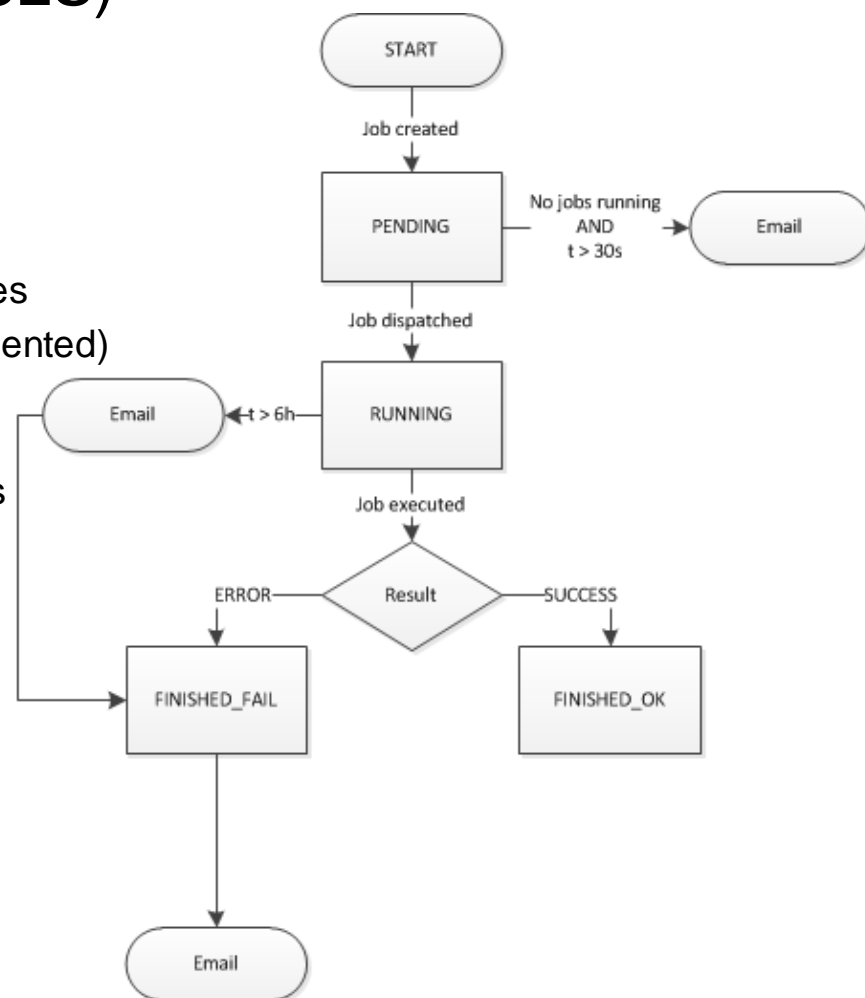
- LEMON

- RACMon

- In-house developed system
- Easy to adapt to different use cases
- Weekly notifications (to be implemented)

- Job monitoring

- Queries jobs table every 5 minutes
- Informs admins in case of:
  - Pending jobs not executed
  - Timed out jobs
  - Failed jobs











## DB On Demand















































































































[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco ▾

[Home](#) [Admin](#) [Help](#) [Admin Help](#)Are you experiencing problems with DB On Demand? [Contact us](#)

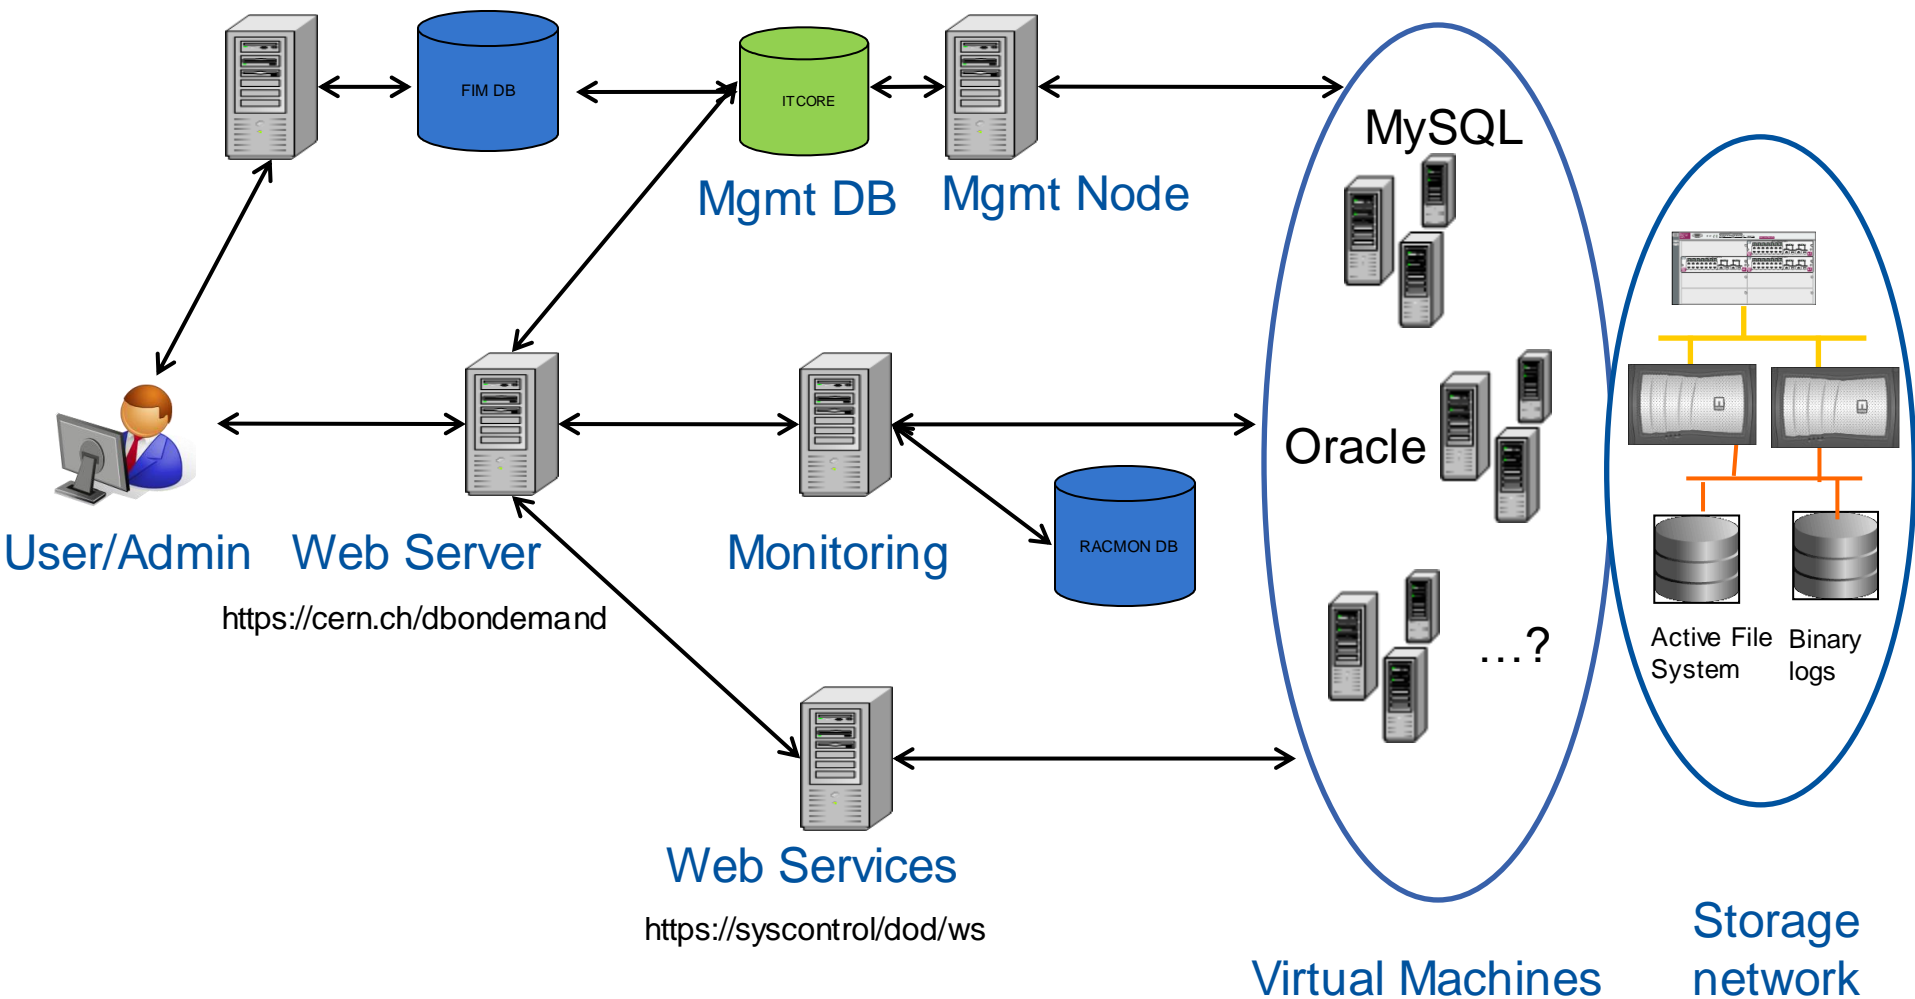
## Instances overview

Collective actions    Refresh Instances 

																					<input type="text" value="All"/>
<input type="checkbox"/>	DB Name	Username	e-Group	Category	Project	Creation Date	Expiry Date	DB Type	DB Size	NO Conn.	State	Actions									
<input type="checkbox"/>	<a href="#">boinc</a>	pete	project-lhcat-home-it	Official	BOINC	20/12/2011	-	MySQL	400 GB	300											
<input type="checkbox"/>	<a href="#">drupal</a>	rgaspar	drupal-admins	Test	DRUPAL	01/09/2011	-	MySQL	500 GB	100											
<input type="checkbox"/>	<a href="#">drupalprod</a>	jpolok	drupal-admins	Official	DRUPAL	19/01/2012	-	MySQL	200 GB	1000											
<input type="checkbox"/>	<a href="#">drupalprod02</a>	jpolok	drupal-admins	Official	DRUPAL	30/01/2012	-	MySQL	200 GB	1000											
<input type="checkbox"/>	<a href="#">drupalsiv</a>	dgomezbl	drupal-admins	Test	DRUPAL	03/02/2012	-	MySQL	500 GB	100											
<input type="checkbox"/>	<a href="#">hc_atlas</a>	rmedrano	hammercloud-dba-atlas...	Official	Hammer...	20/12/2011	-	MySQL	500 GB	500											
<input type="checkbox"/>	<a href="#">hc_cms</a>	rmedrano	hammercloud-dba-cms	Official	Hammer...	08/02/2012	-	MySQL	500 GB	1000											
<input type="checkbox"/>	<a href="#">hc_lhcb</a>	rmedrano	hammercloud-dba-lhcb	Official	Hammer...	01/03/2012	-	MySQL	500 GB	500											
<input type="checkbox"/>	<a href="#">mastro</a>	domarack	service-avc	Official	AVC	19/03/2012	-	MySQL	100 GB	100											
<input type="checkbox"/>	<a href="#">nova</a>	belmiro	ai-admins	Official	Agile	02/04/2012	-	MySQL	100 GB	1000											



## Account Management (FIM)

<https://cern.ch/account>



- Implementation

- J2EE Web Application running on CERN central web servers
  - ZK Framework 5.0.10 (Ajax based)
  - SSO (Single Sign On) + SSL for authentication/authorisation
  - JDBC + Apache BDCCP connection pooling via JNDI
  - Apache Axis 1.4 for SOAP Web Services
  - Google Visualization API for graphs

- User operations

- Request for a new instance (via FIM)
- Instance administration: start/stop, backup/restore, upgrade, configure...
- Documentation for users (integrated with TWiki)

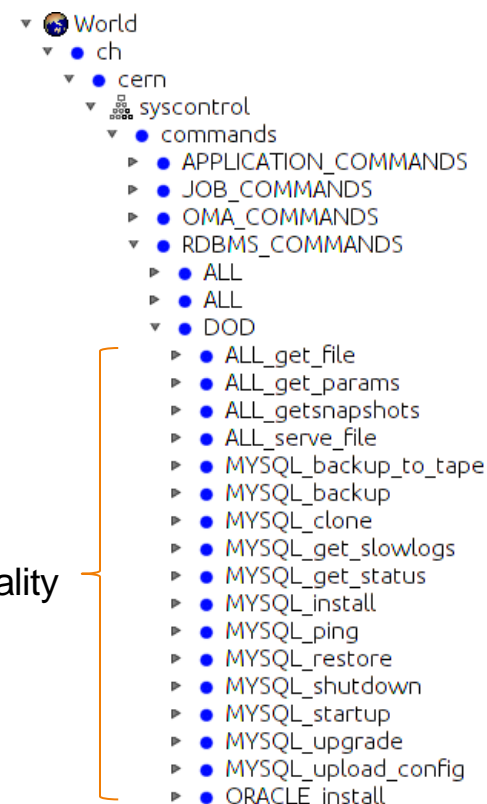
- Service Administration operations

- Approval of requests (via FIM)
- Special collective actions
- System administration maintenance, manage upgrades, edit info...
- Documentation for admins (integrated with TWiki)



- Perl + IT-DB framework
  - Main process (daemon):
    - Queries jobs table every 5 seconds
    - FIFO queue per instance
    - Launches script on instance
    - Captures output and updates instance
  - Worker processes:
    - Executes a job in a certain instance
  - Operations implemented:
    - **ALL**: common to any instance
    - **MYSQL**: specific to MySQL instances
    - **ORACLE**: under development

Basic DBOD core functionality





- Based on **Oracle VM**, a Xen-based hypervisor
  - Required for support of Oracle DBOD instances
  - Easy to integrate with our IT-DB infrastructure
  - Keeping track of CERN Agile Infrastructure project
- Each DBOD instance is delivered in a preconfigured virtual machine
  - 2 cores, 16 GB memory, 64-bit platform
    - Comparable to Amazon EC2 Extra Large Instance
    - Possibility to upgrade depending on user needs



- Production services already testing the system and giving feedback
  - PH-LBC
  - IT (PES, ES, CIS, OIS)
- More coming
  - Other DBMS (Oracle in the works)
  - On-the-fly VM creation
- Possible future developments
  - High availability with master/slave replication
  - Multi-instance entities



# Questions





CERN home &gt; IT Department

daniel.gomez.blanco@cern.ch (dgomezbl)

[\[Logout\]](#) [\[Details\]](#)

## Account Management

[Home](#) [My Accounts](#) [Account Status](#) [Change Password](#) [Approve Requests](#) [My Requests](#) [Applications and Resources](#) [CERN External Accounts](#) [Help](#) [Statistics](#)

## Applications and Resources (dgomezbl)

## Common Services

Mail Services	Automatic Subscription	<a href="#">[Manage]</a>
Windows Desktops	Automatic Subscription	<a href="#">[Manage]</a>
Terminal Services	Automatic Subscription	<a href="#">[Manage]</a>
DFS Workspaces	Automatic Subscription	<a href="#">[Manage]</a>
Web Services	Automatic Subscription	<a href="#">[Manage]</a>
Linux and AFS	Subscribed	<a href="#">[Manage]</a>

## Special Services

Audio conference service	<a href="#">[Subscribe]</a>	
Fax Service	<a href="#">[Subscribe]</a>	
Oracle	Subscribed	<a href="#">[Manage]</a>
DB On Demand	Subscribed	<a href="#">[Manage]</a>
Vidyo	Subscribed	<a href="#">[Manage]</a>

## DB On Demand

DB On Demand is a service which empowers users to create and manage database instances using a simple and intuitive web interface.

MySQL is the only database system currently supported. For more details, check the service home page.

## Service common properties

Associated E-Groups:dbondemand-users (authorized users)  
Home Page: [DB On Demand](#)

The DB On Demand 'drupal' was successfully created.

## DB On Demand resources owned by Daniel Gomez Blanco (dgomezbl)

drupal [\[details\]](#) [\[delete\]](#)

## Resource details: drupal

Name: drupal  
Description: Assigning drupal to dgomezbl for HEPIX

Change Owner

## Create a new DB On Demand resource

[Click to show name limitations.](#)

DB Name:

Description:



Please provide a description for your account. DB On Demand administrators will use this information to evaluate your request (max 200 characters).

Last update: April, 2012 - Contact: [service-desk@cern.ch](mailto:service-desk@cern.ch)



20










[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco

## DB On Demand

[Home](#) | [Admin](#) | [Help](#) | [Admin Help](#)

Are you experiencing problems with DB On Demand? [Contact us](#)

**Actions**








**Information for instance testinstance**
[Refresh Instance](#)

**State:** Running
 **DB Type:** MySQL
 **Creation Date:** 24/11/2011
 **Category:** Test
 **NO Connections:** 50
 **Description:** Test instance until February edited.

**Username:** icoteril
 **e-Group:** it-dep-db
 **Expiry Date:** 08/02/2012
 **Project:** DBOD
 **DB Size:** 500 GB

**Information for job**

Create a backup 09/12/2011 9:09:36

**State:** OK
 **Requester:** dgomezbl
 **Completion Date:** 09/12/2011 9:09:44


**Log:**

```

Fri Dec 9 09:09:39 CET 2011 : Main: Starting
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:39 CET 2011 : GeneralMysql.Get
Fri Dec 9 09:09:40 CET 2011 : RunTime.Retrieve
Fri Dec 9 09:09:40 CET 2011 : RunTime.Retrieve
Fri Dec 9 09:09:40 CET 2011 : Main: my.cnf </ORA/dbs03/DRUPAL/mysql/my.cnf>
Fri Dec 9 09:09:40 CET 2011 : RunTime.GetControllerVol: searching for <(^dbnasw+);vol/(.*?03)jb>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /bin/mount
Fri Dec 9 09:09:40 CET 2011 : Main: working with following controller:volume <dbnasb301:dodtest03>
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running ps -elf | grep -i datadir=/ORA/dbs03/DRUPAL/mysql | grep -v grep
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin
/actual_binarylog.sql
Fri Dec 9 09:09:40 CET 2011 : Main: <binlog>.<000299> actual binary log sequence
Fri Dec 9 09:09:40 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXX --socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin
/presnap.sql
Fri Dec 9 09:09:40 CET 2011 : Main: presnap actions completed successfully.
Fri Dec 9 09:09:40 CET 2011 : Main: controller: <dbnasb301> volume: <dodtest03>
Fri Dec 9 09:09:40 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0x1faccd70)>
Fri Dec 9 09:09:40 CET 2011 : GeneralMysql.RunSSHStr running snap create dodtest03 snapscrip_09122011_090940_300
Fri Dec 9 09:09:43 CET 2011 : Main: stdout:

Fri Dec 9 09:09:43 CET 2011 : Main: creating snapshot...

Fri Dec 9 09:09:43 CET 2011 : Main: snapshots have been created!.
Fri Dec 9 09:09:43 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock < /ORA/dbs01/syscontrol/projects/dod/bin
/postsnap.sql
Fri Dec 9 09:09:43 CET 2011 : Main: postsnap actions completed successfully.
Fri Dec 9 09:09:43 CET 2011 : Main: mysql_snapshot is over.
Fri Dec 9 09:09:43 CET 2011 : mysql_snapshot.Main: State: [0]
          
```

**Upgrade database version**

**Upgrade your database from version 5.5.20 to version 5.5.21?**

Cancel
Accept





## DB On Demand

[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco ▾

[Home](#) [Admin](#) [Help](#) [Admin Help](#)Are you experiencing problems with DB On Demand? [Contact us](#)

## Help Contents

- [Manifesto](#)
- [Disclaimer](#)
- [Request an instance](#)
- [Getting started](#)
  - [MySQL](#)
- [Management of your instance](#)
  - [Basics](#)
  - [Start and stop an instance](#)
  - [File management](#)
  - [Manage backups](#)
  - [Restore](#)
  - [Upgrade](#)
  - [Monitoring](#)
  - [Destroy](#)
- [Presentations](#)

**Manifesto**

DB On Demand is a service which empowers users to create and manage database instances using a simple and intuitive web interface. Although MySQL is the only database system supported at present, the architecture has been designed to be flexible and we should be able to support other relational database systems in the future.

DB On Demand has been developed to provide, in a scalable way, database platform support for people who cannot use, the Oracle relational database service. If you need to manage your own MySQL instance today to support a simple database application, the DB On Demand service offers automated backup and recovery services as well as guarantees of service continuity in case of hardware failure.

The need to offer support based on different database systems is on the rise within our user community. For example, one use case includes certain applications certified for running with non-Oracle databases (i.e MySQL). Additional use cases are based on the desire expressed by user groups in being able to easily deploy and maintain different DBMS systems for their internal applications development. DB On Demand will cover this need offering users an easy way to create and manage their databases through a system that is fully integrated with different CERN technologies.

The following [document](#) provides a brief introduction to DB On Demand and its architecture.

**DB On Demand offers**

- Full DBA privileges on a database running in a preconfigured virtual machine.
- Configurable backups. Possibility to send consistent full backup to tape.
- Configurable point in time restore with binary logs utilization.
- System administrative operations and database system updates for you to install at a schedule to suit you—or, if you prefer, to be installed automatically by us.
- Configurable client access to the database either using operating system capabilities or database vendor ones.
- In case of hardware failure based on the underlying virtualization platform database will be migrated and restarted on a different host.

**DB On Demand limitations**

- No application support is provided.
- No DBA support is provided.
- Only single instance databases are supported. Applications with specific high availability requirements should use the more functional Oracle database service.
- Similarly, applications requiring advanced features, such as replication or clustering should use the Oracle Database service.
- No support is provided for linking between DB On Demand databases or between DB On Demand databases and centrally managed Oracle databases.

**DB On Demand limitations for MySQL**

- The InnoDB engine must be used to ensure correct functioning of the backup service.

**Important Note:**

DB On Demand users have DBA privileges on their databases and may be responsible to the scheduling of operating system and database system patches. As such, DB On Demand users, are entirely responsible for ensuring that their systems, and the use of their systems comply fully with the Rules for the use of CERN Computing Facilities and also with any specific security rules for databases. Failure to comply with these guidelines may lead to a DB On Demand database being shutdown with little or no advance warning.





## DB On Demand

[CERN Home](#) | [IT Home](#) | [IT-DB Home](#) | [Phonebook](#)

Welcome Daniel Gomez Blanco ▾

Are you experiencing problems with DB On Demand? [Contact us](#)[Home](#) [Admin](#) [Help](#) [Admin Help](#)

## Help Contents

- [Documentation](#)
- [Development](#)
  - [Web Application](#)
    - [Java packages](#)
  - [Database](#)
    - [Database account](#)
    - [Triggers](#)
    - [Stored procedures](#)
    - [Job scheduler](#)
  - [Web services](#)
- [Administration](#)
  - [Specific configuration](#)
  - [Daemon](#)
  - [MySQL ping](#)
  - [Backup/Restore to Tape](#)
  - [Upgrade Drupal: physical move plus Mysql version upgrade](#)
  - [Manage upgrades](#)
  - [Master Slave Replication](#)
- [Instance creation workflow](#)
  - [Common procedures](#)
  - [MySQL installation](#)
  - [Oracle installation](#)
- [DBOD meetings](#)

## Documentation

DB On Demand's documentation is managed using CERN's Twiki. This approach takes advantage of all the features provided by Twiki, like versioning, topic linking, user permissions, etc., to then render the content of the created topics on an IFrame inside the DB On Demand help web pages (user and administrator). There are two main parent topics defined for user and administrator documentation:

- [User documentation](#): This topic is stored on the public DB web in order to allow any authenticated user to view the content. It contains all the information needed to use DB On Demand: manifesto, disclaimer, user manual, etc...
- [Admin documentation](#): This topic is stored on the private DB web in order to limit access to these documents to members of the IT-DB group. It contains information for development and administration purposes.

Documents must be created on the corresponding parent topics depending on their content. One of these two templates must be used:

- [User template](#)
- [Admin template](#)

Both templates are almost the same, there is only one difference: *User template* defines the Twiki parameter `DENYTOPICVIEW = TWikiGuest` in order to force the authentication of the user. This is necessary due to attachments used as links in topics (e.g. pictures). If a user is not authenticated, he will be using the TWikiGuest identity, and he will not be able to see the content correctly. As SSO is used for DB On Demand, this is not an issue. The user is automatically authenticated on Twiki because he is already authenticated on DB On Demand.

These templates also define a style sheet, it must be used to be able to render the content correctly on the DB On Demand website. They also define the format for topic titles and section titles. The tag `<noautolink>` must contain the whole content to avoid automatic creation of links to Twiki pages, which would result in a bad presentation and uncontrolled behaviour inside the IFrame.

When a new topic is created a link to this topic must also be placed in one of the two parent topics. **NOTE: Don't use the WYSIWYG editor on these tables of contents. It does not interpret the document properly.** These parent topics act as a table of contents for the entire documentation, so it is crucial to maintain them when a new document is created. To make content edition easier, there are two groups of links in each of these parent topics, conditionally rendered depending on the URL used to view these tables of contents:

- The first group is used to display the links using the parameter `target="dodhelpcontent"` to load the content on that specific IFrame, and the request parameter `template=viewplain` to present the content with no Twiki frames or extra content. This block of links is only rendered when the request parameter `template` is defined for this topic. That behaviour is achieved using the following expression on the container DIV: `style="display: %IF{ "defined 'template'" then="block" else="none"}%"`
- The second group is used to display the links to normal Twiki pages, so they can be edited. This block of links is rendered when the topic is viewed using a normal URL. That behaviour is achieved using the following expression on the container DIV: `style="display: %IF{ "defined 'template'" then="none" else="block"}%"`

After the creation of a new topic, a link to that topic must be placed in both blocks to allow correct navigation both on Twiki and on the DB On Demand website.

If you want to create a link to a website that is not part of this documentation, please use `target="_blank"` to force the web browser to open it in a different window.



Admin (ignacio.cern.ch) x SQL Editor (ignacio.cern.ch) x

File Edit View Query Database Plugins Scripting Community Help

ACTIONS

- Execute SQL File
- Add Schema
- Add Table
- Add View
- Add Routine

SCHEMAS

- dbtest
- dod\_dbmon
- phpmyadmin
- test
  - Tables
  - Views
  - Routines
- tpcc**

Query 1 x

```
1 use tpcc;  
2 select * from district limit 10;
```

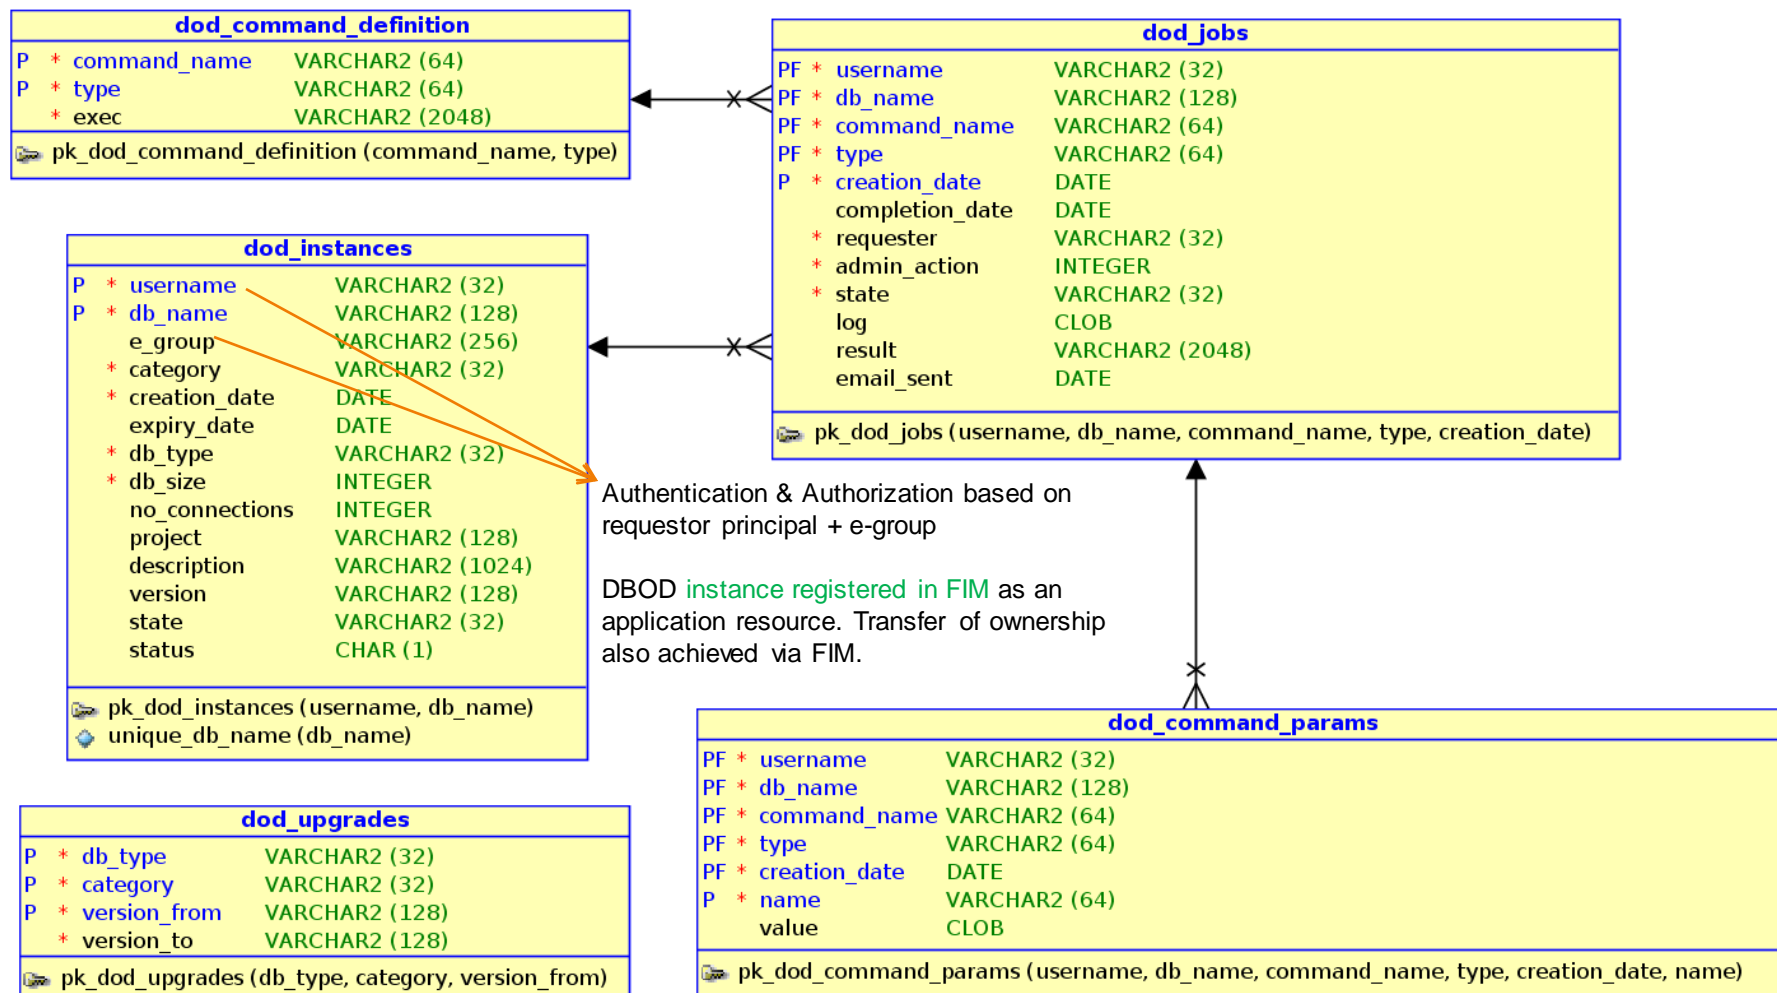
Overview | Output | Snippets | Query 1 Result x

Filter: Fetched 10 records, more available. Duration: 3.692 sec, fetched in: 0.004 sec

#	d_id	d_w_id	d_name	d_street_1	d_street_2	d_city	d_state	d_zip	d_tax	d_y
1	1	1	ff5uELo	lOfFpj2FQz	b9beweMwehBZK	eEfx0oEMTT9CT	dT	263013893	0.13	508
2	2	1	4B1gNhcNVq	wHr7kGiDtuo6	JyuiRmKWvKAJ1mgUaOO	GVcUDnN7ZUps	ka	499316880	0.20	508
3	3	1	pj2gzY	RI4prdHgUfO4Usv	I8G0KULyJ51y	ZCWIGJ7sY6MDCof3KTLs	jl	111685894	0.16	508
4	4	1	sZdi5P	GaTwdlNwB5FwHsekc	APWmd2cDdLG3NXbo	EZo8ltNf8GJtqc1D8	r8	850440398	0.12	505
5	5	1	S5fBw1q	F8w5ixjoi8w3	fkS0mGjeUH4ZuF	v6ZACVEuQVi6dc7nFpDd	TN	822382584	0.10	505
6	6	1	mzkeNfMT	yZdDOpp0ICCEmSiGw	eUABE7vsV9Cw	AVdMuQ9sij4v5UbjRE	vO	486831793	0.20	508
7	7	1	kflAoBq7	jAhKU8ZQ2vXAPZAOjLKs	jcK4m7DcCQMN7epEBFG7	QULYt2sBuXuUpwGw9	9Z	421327086	0.13	508
8	8	1	Yd3Pgta	QU3DOH8X8F4Hjj5	Z4Yi8qCgRD656	cU9fVuq1155lmmNmMP	2X	065825863	0.11	508
9	9	1	ohZiVah	d0KOL4klMk	QwSvMiUrddsgRXD82	paNpul9wTsGD8	7Z	838707119	0.16	508
10	10	1	FiW2lrqA	z63pMGv0NUZpkb8	HJBOzwbhQdzHT7rQDuFX	9XTd7HLiPyz7Lg	uO	874351081	0.18	509

Query Completed








[Home](#)[Documentation](#)[Alarms](#)[Metrics](#)[Misc](#)[Help](#)

## Information for Clusters / database / dbvrts1029

## Host information

<b>operating system(s)</b>	Red Hat Enterprise Linux Server release 5.7 (Tikanga)
<b>architecture (kernel)</b>	x86_64 (2.6.18-274.7.1.el5)
<b>up time (since)</b>	6 days, 5h:21m (Thu, 05 Apr 2012 09:40:46 +0200)
<b>CPU (count/logical)</b>	Intel(R) Xeon(R) CPU L5520 @ 2.27GHz (1/2)
<b>memory (swap)</b>	12022 MB (8189 MB)
<b>cluster (subcluster)</b>	<b>database (database/DRUPAL)</b>
<b>IP address(es)</b>	10.18.2.220 (eth0) 172.31.128.51 (eth2) 172.31.133.51 (eth1)
<b>state</b>	maintenance
<b>status</b>	Available



Since 06.02.2012 the new **Lemon web** frontend is in production.


This version will remain active during 2012.

Hide forever


Search entities:

[Virtual Clusters](#) | [Clusters](#) | [Racks](#) | [Hardware types](#) | [Virtual Batch](#) | [Services](#)

## Use of host partitions



## Number of users and 1-min load average





- Based on snapshots
  - Restore time reduced to **a few seconds**
- Full restore of whole database
- DBOD instances configured with binary logs
  - Point-In-Time Recovery

