

DB

Database Services

CERN IT
Department

Proposal for a Database On Demand Service (DBOD)

Ignacio Coterillo
Daniel Gomez
Dawid Wojcik
Ruben Gaspar

- Objectives
- Demo
- Behind the scenes
- Next

- Allow the CERN user community to **create and manage database instances**. DBOD Users will have **full DBA privileges** on their instance.
- Offer **different DBMS vendors**. **MySQL** is our **starter**. Service designed to scale.
- DBOD users will have a **set of operations** available via a web interface. More on that later 😊
- DBOD comes on a **preconfigured virtual machine**.
- The DBOD service will take care of system administration operations and provide **upgrade paths for DBMS**
- **DBOD 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.
- The **DBOD service will not provide DBA or application support**.
 - Although we could consider a **community oriented** support site (as for iOS/Android).
- Users needing **advanced features** such as performance analysis, clustering, replication, or guaranteed links to Oracle databases should use the **Oracle service**.

DB

DBOnDemand site

CERN IT
Department

<https://cern.ch/DBOnDemand/>
DEMO

Actions

-
-
-
-
-
-
-

Information for instance testinstance Refresh instance

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

Information for job Create a backup 09/12/2011 10:24:30

State: OK ● **Requester:** dgomezbl
Creation Date: 09/12/2011 10:24:30 **Completion Date:** 09/12/2011 10:24:39

Log:

```

Fri Dec 9 10:24:36 CET 2011 : Main: Starting
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.GetGeneralVars: host <dbvrtg046.cern.ch>
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.GetGeneralVars: port <5500>
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.GetGeneralVars: socket </var/lib/mysql/mysql.sock>
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.GetGeneralVars: mysql_datafiles </ORA/dbs03/DRUPAL/mysql>
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.GetGeneralVars: backup_dest </ORA/dbs02/DRUPAL/mysql>
Fri Dec 9 10:24:36 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_user_nastorag>
Fri Dec 9 10:24:36 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_db_dod_mysql>
Fri Dec 9 10:24:36 CET 2011 : Main: my.cnf </ORA/dbs03/DRUPAL/mysql/my.cnf>
Fri Dec 9 10:24:36 CET 2011 : RunTime.GetControllerVol: searching for <(^dbnas/w+):vol/(.*?03)lb>
Fri Dec 9 10:24:36 CET 2011 : RunTime.RunStr running /bin/mount
Fri Dec 9 10:24:36 CET 2011 : Main: working with following controller:volume <dbnasb301:dodtest03>
Fri Dec 9 10:24:36 CET 2011 : RunTime.RunStr running ps -elf | grep -i datadir=/ORA/dbs03/DRUPAL/mysql | grep -v grep
Fri Dec 9 10:24:36 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXX --mysql_socket=/var/lib/mysql/mysql.sock
< /ORA/dbs01/syscontrol/projects/dod/bin/actual_binarylog.sql
Fri Dec 9 10:24:36 CET 2011 : Main: <binlog>.<000304> actual binary log sequence
Fri Dec 9 10:24:36 CET 2011 : RunTime.RunStr running /usr/bin/mysql -u dod_mysql -pXXXXXXXXXXXX --socket=/var/lib/mysql/mysql.sock
< /ORA/dbs01/syscontrol/projects/dod/bin/presnap.sql
Fri Dec 9 10:24:36 CET 2011 : Main: presnap actions completed successfully.
Fri Dec 9 10:24:36 CET 2011 : Main: controller: <dbnasb301> volume: <dodtest03>
Fri Dec 9 10:24:36 CET 2011 : Main: working controller: <dbnasb301> volume: <ARRAY(0xa7abd70)>
Fri Dec 9 10:24:36 CET 2011 : GeneralMySQL.RunSSHStr running snap create dodtest03 snapscrip_09122011_102436_305
Fri Dec 9 10:24:38 CET 2011 : Main: stdout

Fri Dec 9 10:24:38 CET 2011 : Main: creating snapshot..

Fri Dec 9 10:24:38 CET 2011 : Main: snapshots have been created!
Fri Dec 9 10:24:38 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 10:24:38 CET 2011 : Main: postsnap actions completed successfully.
Fri Dec 9 10:24:38 CET 2011 : Main: mysql_snapshot is over.
Fri Dec 9 10:24:38 CET 2011 : mysql_snapshot.Main: State: [0]
        
```

backup

Takes 2 secs



DB On Demand

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

Home Admin Help Admin Help

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 edi

Information for job

Username: dgomezbl

e-Group: it-dep-db

Dispatch a restore

Available snapshots:

Dec 2011						
Mon	Tue	Wed	Thu	Fri	Sat	Sun
28	29	30	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	31	1

Snapshots for selected day:

- [9:09:19](#)
- [9:09:40](#)
- [10:19:21](#)
- [15:30:40](#)








Select a day and time to restore to:

09/12/2011





Restore


Actions
Information for instance testinstance Refresh instance ↻


State: Running ● **Username:** dgomezbl

DB Type: MySQL **e-Group:** it-dep-db 

Creation Date: 24/11/2011 **Expiry Date:** 08/02/2012 

Category: Test **Project:** - 

NO Connections: 50 **DB Size:** 500 GB

Description: Test instance until February edited. 

Information for job Dispatch a restore 09/12/2011 10:26:41 ▼

State: OK ● **Requester:** dgomezbl

Creation Date: 09/12/2011 10:26:41 **Completion Date:** 09/12/2011 10:27:12

Log:

```

Fri Dec 9 10:26:48 CET 2011 : Main: starting
Fri Dec 9 10:26:48 CET 2011 : GeneralMysql.GetGeneralVars: host <dbvrtg046.cern.ch>
Fri Dec 9 10:26:48 CET 2011 : GeneralMysql.GetGeneralVars: port <5500>
Fri Dec 9 10:26:48 CET 2011 : GeneralMysql.GetGeneralVars: socket </var/lib/mysql/mysql.sock>
Fri Dec 9 10:26:48 CET 2011 : GeneralMysql.GetGeneralVars: mysql_datafiles </ORA/dbs03/DRUPAL/mysql>
Fri Dec 9 10:26:48 CET 2011 : GeneralMysql.GetGeneralVars: backup_dest </ORA/dbs02/DRUPAL/mysql>
Fri Dec 9 10:26:48 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_user_nastorag>
Fri Dec 9 10:26:48 CET 2011 : RunTime.RetrievePasswordForUser: password found for <password_db_dod_mysql>
Fri Dec 9 10:26:48 CET 2011 : Main: my.cnf </ORA/dbs03/DRUPAL/mysql/my.cnf>
Fri Dec 9 10:26:48 CET 2011 : RunTime.GetControllerVol: searching for <(^dbnas/w+):vol/(.*?03)b>
Fri Dec 9 10:26:48 CET 2011 : RunTime.RunStr running /bin/mount
Fri Dec 9 10:26:48 CET 2011 : Main: working with following controller:volume <dbnasb301:dodtest03>
Fri Dec 9 10:26:48 CET 2011 : RunTime.RunStr running ps -elf | grep -i datadir=/ORA/dbs03/DRUPAL/mysql | grep -v grep
Fri Dec 9 10:26:48 CET 2011 : RunTime.RunStr running /usr/bin/mysqladmin -u dod_mysql -pXXXXXXXXXXXX --socket=/var/lib/mysql/mysql.sock shutdown
Fri Dec 9 10:26:52 CET 2011 : Main: Instance is down
Fri Dec 9 10:26:52 CET 2011 : Main: <303> binlog number at the time of the snapshot.
Fri Dec 9 10:26:52 CET 2011 : Main: binary logs to be used if PITR: <binlog.000303 binlog.000304 binlog.000305>
Fri Dec 9 10:26:52 CET 2011 : GeneralMysql.RunSSHStr running snap restore -f -t vol -s snapscrip_09122011_101921_303 dodtest03
Fri Dec 9 10:26:54 CET 2011 : Main: dodtest03 set to snapscrip_09122011_101921_303: <stdout:
Volume dodtest03: revert successful.
>
Fri Dec 9 10:26:54 CET 2011 : RunTime.RunStr running hostname
Fri Dec 9 10:26:54 CET 2011 : Main: running <system /usr/bin/mysqld_safe --defaults-file=/ORA/dbs03/DRUPAL/mysql/my.cnf --pid-file=/ORA/dbs03/DRUPAL/mysql/dbvrtg046.cern.ch.pid --socket=/var/lib/mysql/mysql.sock --skip-networking >/dev/null 2>&1 &&!!!
Fri Dec 9 10:27:04 CET 2011 : Main: file </ORA/dbs03/DRUPAL/mysql/dbvrtg046.cern.ch.pid> to check for.
Fri Dec 9 10:27:04 CET 2011 : File.CheckFile: <e> true on file </ORA/dbs03/DRUPAL/mysql/dbvrtg046.cern.ch.pid>
Fri Dec 9 10:27:04 CET 2011 : RunTime.RunStr running cd /ORA/dbs02/DRUPAL/mysql; /usr/bin/mysqlbinlog --stop-datetime="2011-12-09_10:21:21" binlog.000303 binlog.000304 binlog.000305 | /usr/bin/mysql -u dod_mysql -pXXXXXX --socket=/var/lib/mysql/mysql.sock
Fri Dec 9 10:27:05 CET 2011 : RunTime.RunStr running /usr/bin/mysqladmin -u dod_mysql -pXXXXXXXXXXXX --socket=/var/lib/mysql/mysql.sock shutdown
Fri Dec 9 10:27:07 CET 2011 : Main: Instance is down
Fri Dec 9 10:27:07 CET 2011 : RunTime.RunStr running /etc/init.d/mysql start
Fri Dec 9 10:27:11 CET 2011 : Main: Instance is up
Fri Dec 9 10:27:11 CET 2011 : Main: mysql_snaprestore is over.
Fri Dec 9 10:27:11 CET 2011 : mysql_restore.Main: State: [0]

```

Restore time: 2secs

Recovery time

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	wHr7kGiDtu06	JyuiRmKWvKAJ1mgUaOO	GVcUDnN7ZUps	ka	499316880	0.20	508
3	3	1	pj2gzY	RI4prdHgUf04Usv	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	Z4Y8qCgRD656	cU9fVuq1155lmmNmMP	2X	065825863	0.11	508
9	9	1	ohZiVah	d0KOL4kIMk	QwSvMiUrddsgRXD82	paNpui9wTsGD8	7Z	838707119	0.16	508
10	10	1	FIW2lrqA	z63pMGv0NUZpkb8	HJBOzwbhQdzHT7rQDuFX	9XTd7HLiPyz7Lg	uO	874351081	0.18	509

Connection Information

Name: ignacio.cern.ch
 Host: ignacio.cern.ch:5500
 Server: MySQL
 Version: 5.5.17-log
 User: dod_mysql

Query Completed

Admin (ignacio.cern.ch) x

File Edit View Database Plugins Scripting Community Help

MANAGEMENT

- Server Status
- Startup / Shutdown
- Status and System Variables
- Server Logs

CONFIGURATION

- Options File

SECURITY

- Users and Privileges

DATA EXPORT / RESTORE

- Data Export and Restore

Users and Privileges

Server Access Management Schema Privileges

Select a user and pick the privileges it has for a given Schema and Host combination.

Users	Host	Schema	Privileges
<anonymous>			
phpmyadmin	localhost	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
dod_mysql	pcitdb43.cern.ch	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
root	pcitdb43.cern.ch	tpcc2	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
mysql	localhost	tpcc2	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
dod_dbmon	localhost	tpcc1000	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
tpcc	localhost	tpcc1	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
tpcc_test	localhost	tpcc10	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
dbtest	dbpvt.cern.ch	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
newuser	dbpvt2.cern.ch	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
icoteril	dbpvt3.cern.ch	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO
	mydbptest.cern.ch	tpcc	ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE VIEW, DELETE, DRO

Schema and Host fields may use % and _ wildcards. The server will match specific entries before wildcarded ones.

Delete Entry Add Entry...

Object Rights

- SELECT
- INSERT
- UPDATE
- DELETE
- EXECUTE
- SHOW VIEW

DDL Rights

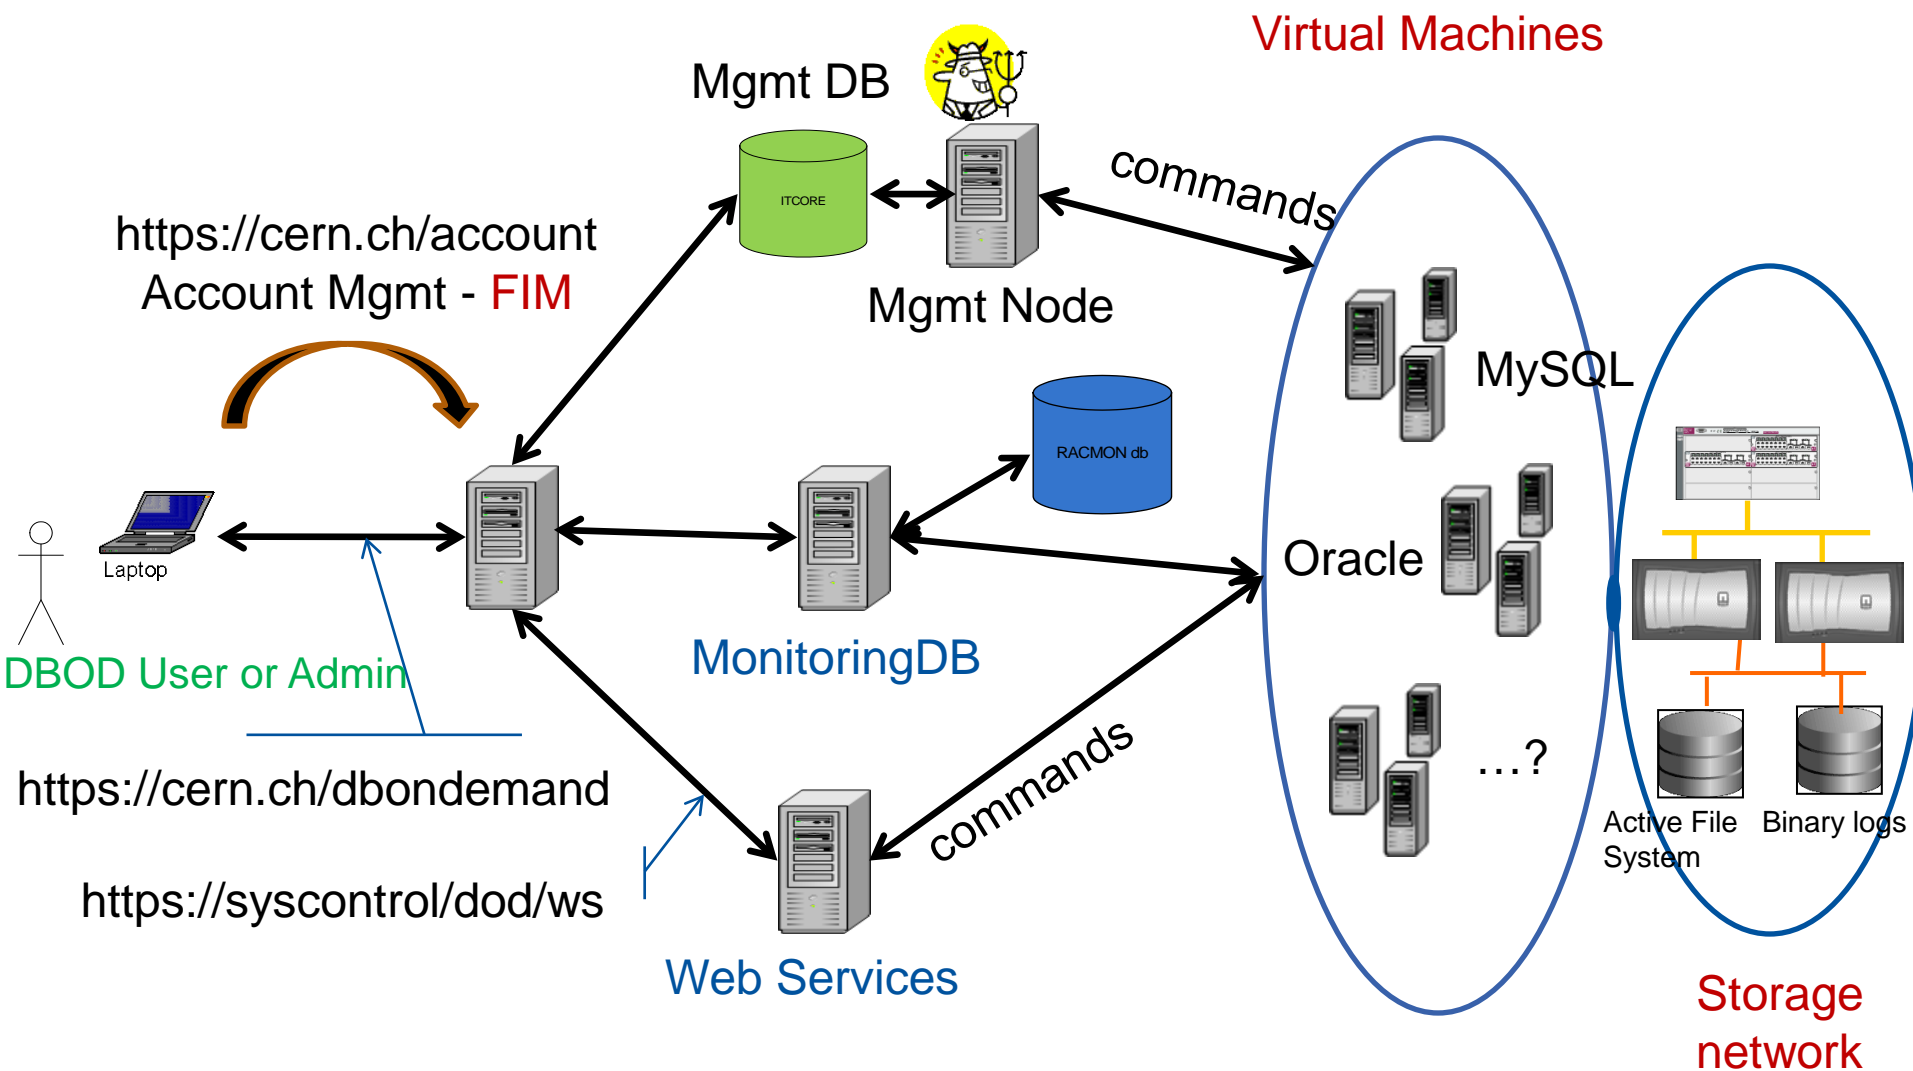
- CREATE
- ALTER
- REFERENCES
- INDEX
- CREATE VIEW
- CREATE ROUTINE
- ALTER ROUTINE
- DROP
- TRIGGER

Other Rights

- GRANT OPTION
- CREATE TEMPORARY TABLES
- LOCK TABLES

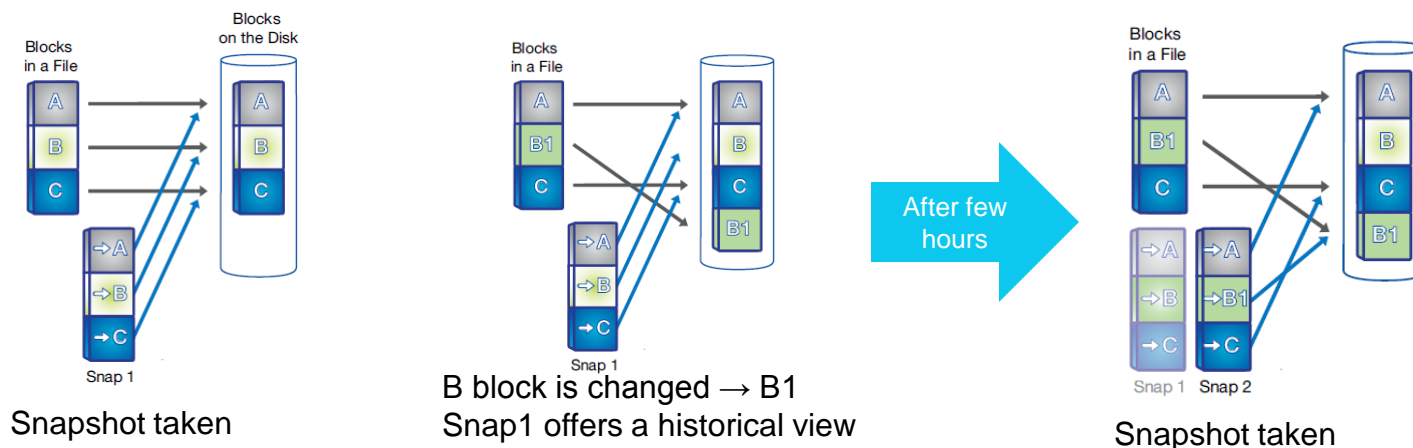
Select "ALL" Unselect All

Revert Save Changes



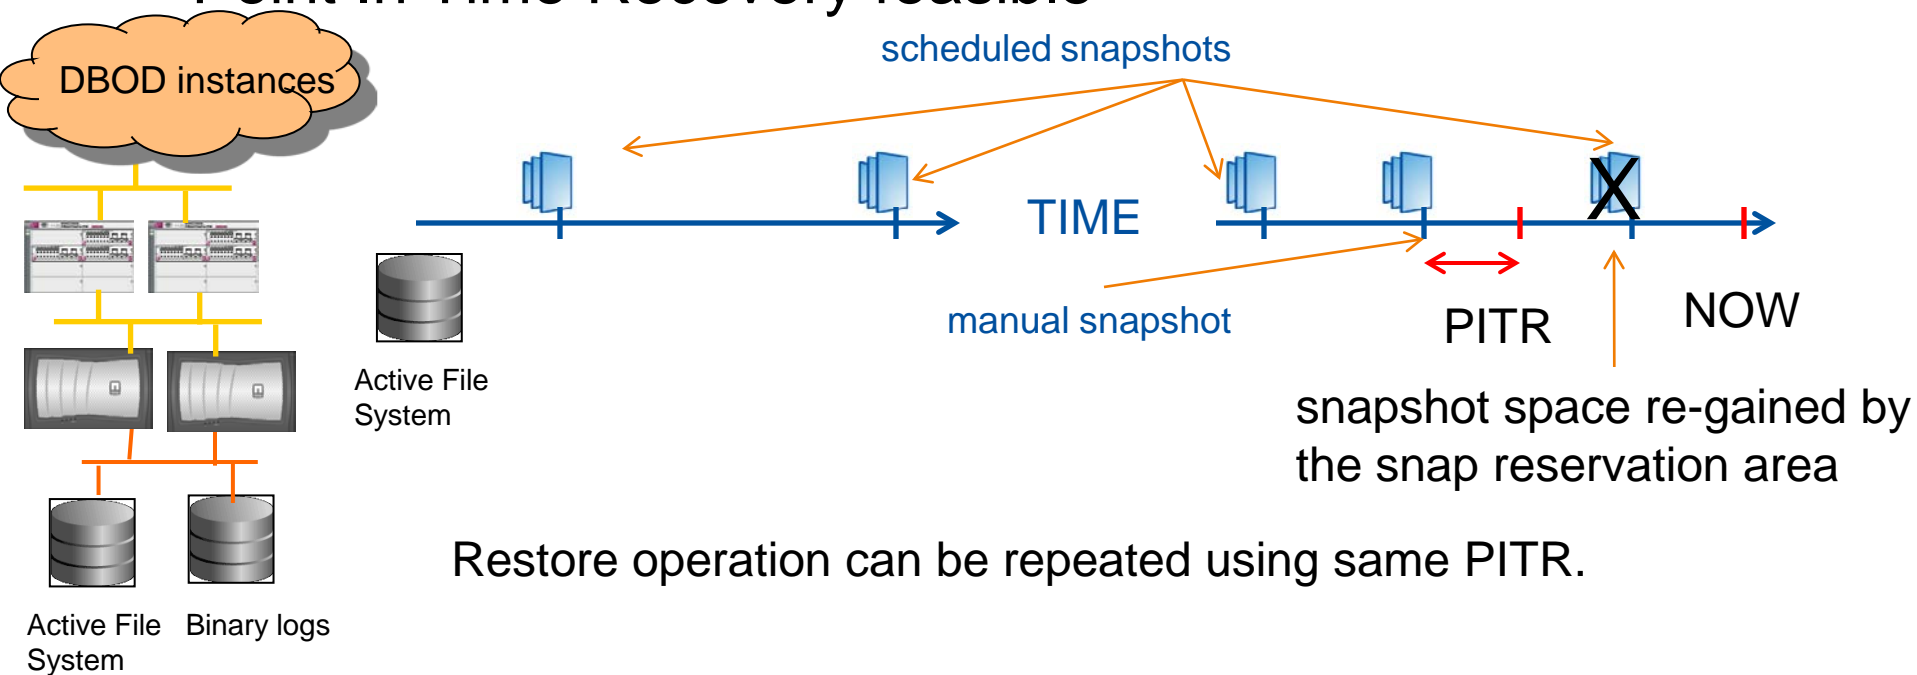
- **Implementation**
 - J2EE Web Application running on central web servers
 - ZK Framework 5.0.7.1 (Ajax based)
 - SSO (Single Sign On) + SSL for authentication/authorisation
 - JDBC + Apache BDCP connection pooling via JNDI
 - Apache Axis 1.4 for SOAP Web Services
- **User operations**
 - Request for a new instance
 - Instance administration: backup/restore, stop/start,...
 - Documentation (Twiki)
- **Service Administration operations**
 - Special collective actions
 - System administration maintenance, MySQL Upgrades,...
 - Approval of requests
 - Documentation (Twiki)

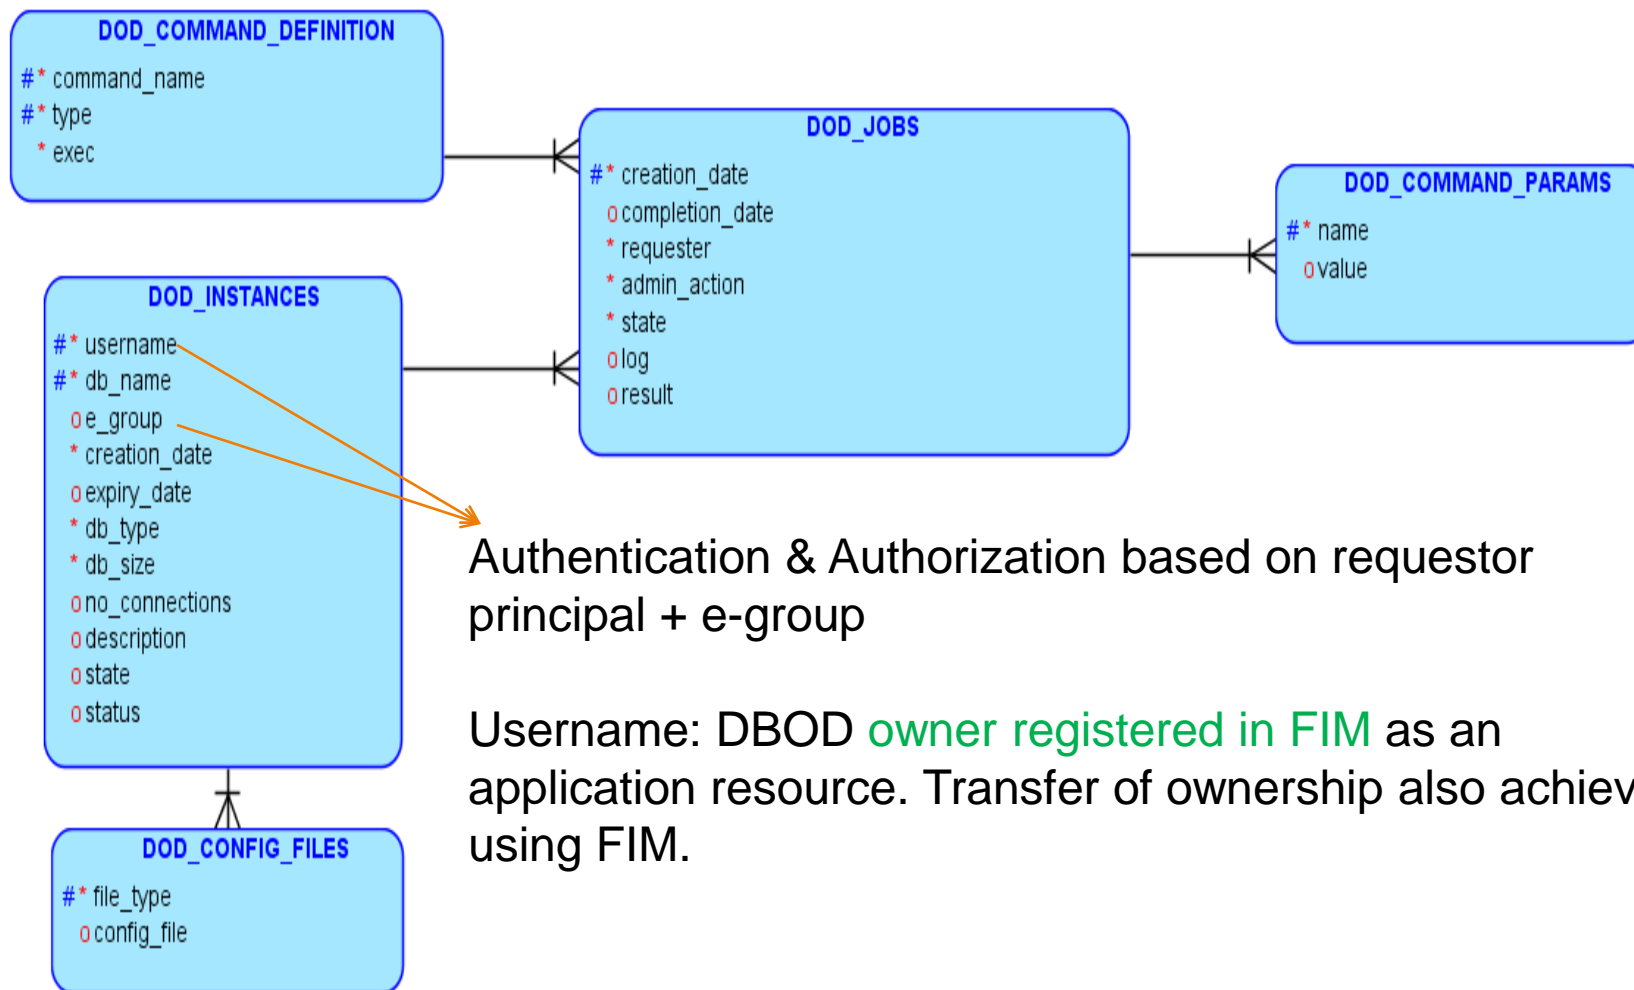
- It's a tool, good to know how to use it (see doc)
- Based on snapshots. **Full database backup in few secs.**



- How many snapshots can I have? How long backwards in time can I go?
 - **Extra 20%** space provided for **snapshots**: i.e active file system 100GB, snapshot area 20 GB → total: 120GB
 - Depends on DBOD instance activity, i.e. changing 2% active file system will allow to have 10 days of snapshots
- For MySQL: **storage engine 'innodb'** recommended. More in next slide.
- Soon possibility to send to **tape**: sorting out policy & software

- 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 feasible






Authentication & Authorization based on requestor principal + e-group

Username: DBOD **owner registered in FIM** as an application resource. Transfer of ownership also achieved using FIM.

- Lemon
- SLS
- In house monitoring
 - Easy to adapt to different use cases
 - Weekly notifications to DBOD owners (coming)
 - Accessible via <https://cern.ch/dbondemand>



- Technical: Perl + IT-DB framework

Main process:

Timed loop querying <dod_jobs> every 5secs.
FIFO type queue.

Worker processes:

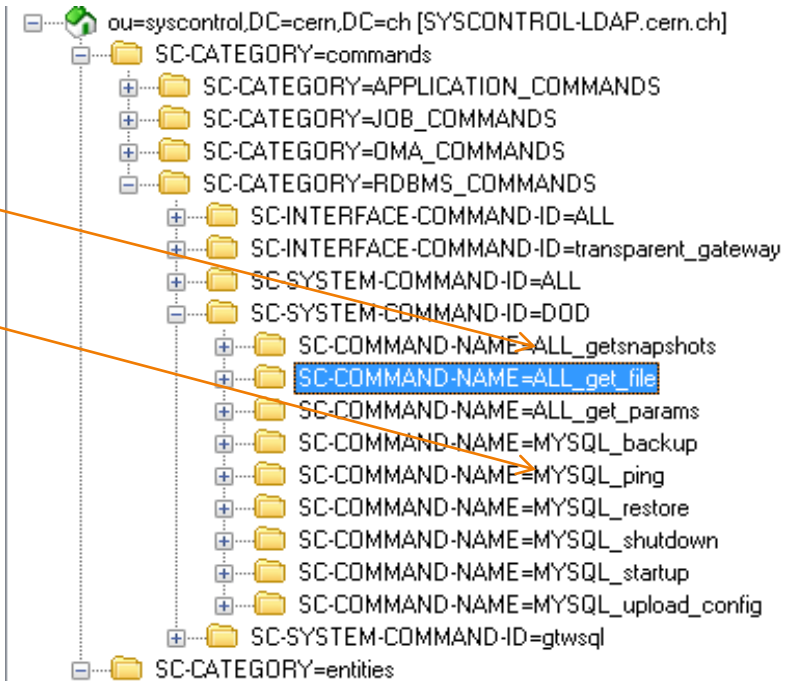
Executes and logs a job

Operations implemented:

ALL: all flavours

MYSQL: specific to

Basic core functionality of DBOD



- Based on **Oracle VM**, a Xen-based hypervisor
 - Required for support on Oracle DBOD instances.
- Provisioning: **Enterprise Manager 12cloud**
 - Easy to integrate on our IT-DB actual infrastructure. Running experience.
 - So far, creation of vm is a manual process
- **OpenStack...**
 - ... rackspace have recently announced a MySQL Cloud database API (REST interface)
 - We should aim to be compatible with this.
 - And also track **storage** (NetApp) integration with OpenStack

- Business **clients already testing** the system and giving feedback – Thank You!
 - PH-LBC
 - IT-PES
- Possibly move to **MySQL enterprise** next year
 - Admin monitoring, sql analyzer, security plug-in, backup.
- More coming:
 - Other RDBMS databases: oracle
 - Tape backup
 - Upgrade functionality
 - ...

DB

Thank You!