



Benjamin Lipp
IT-DB-DBB

Supervisor:
Ignacio Coterillo Coz

19/08/2014

Background image: Shutterstock

OpenStack Trove: Evaluation and Interfacing with the CERN Database on Demand Service

CERN Database on Demand

CERN Resources Portal

Manage your CERN Resources, lifecycle, settings, etc.

Home List Services Pending Actions Select Account Help Support

DB On Demand

DB On Demand

Service Information

Resources

DB On Demand resources owned by Benjamin Lipp (blipp)

Create a new DB On Demand resource

[Click to show name limitations.](#)

DB Name:

atlas_db

Description:

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

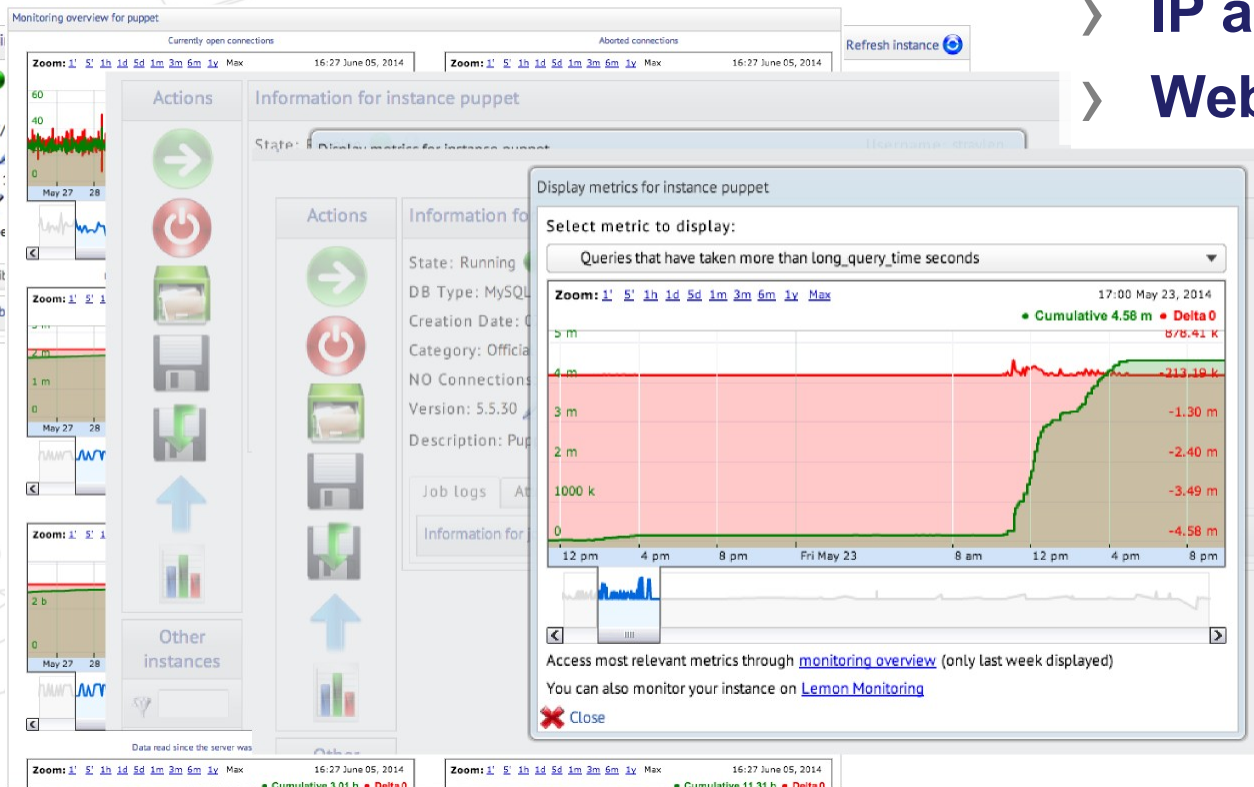
We would like to use a ~~MySQL~~ database for internal ATLAS data.

Create Resource

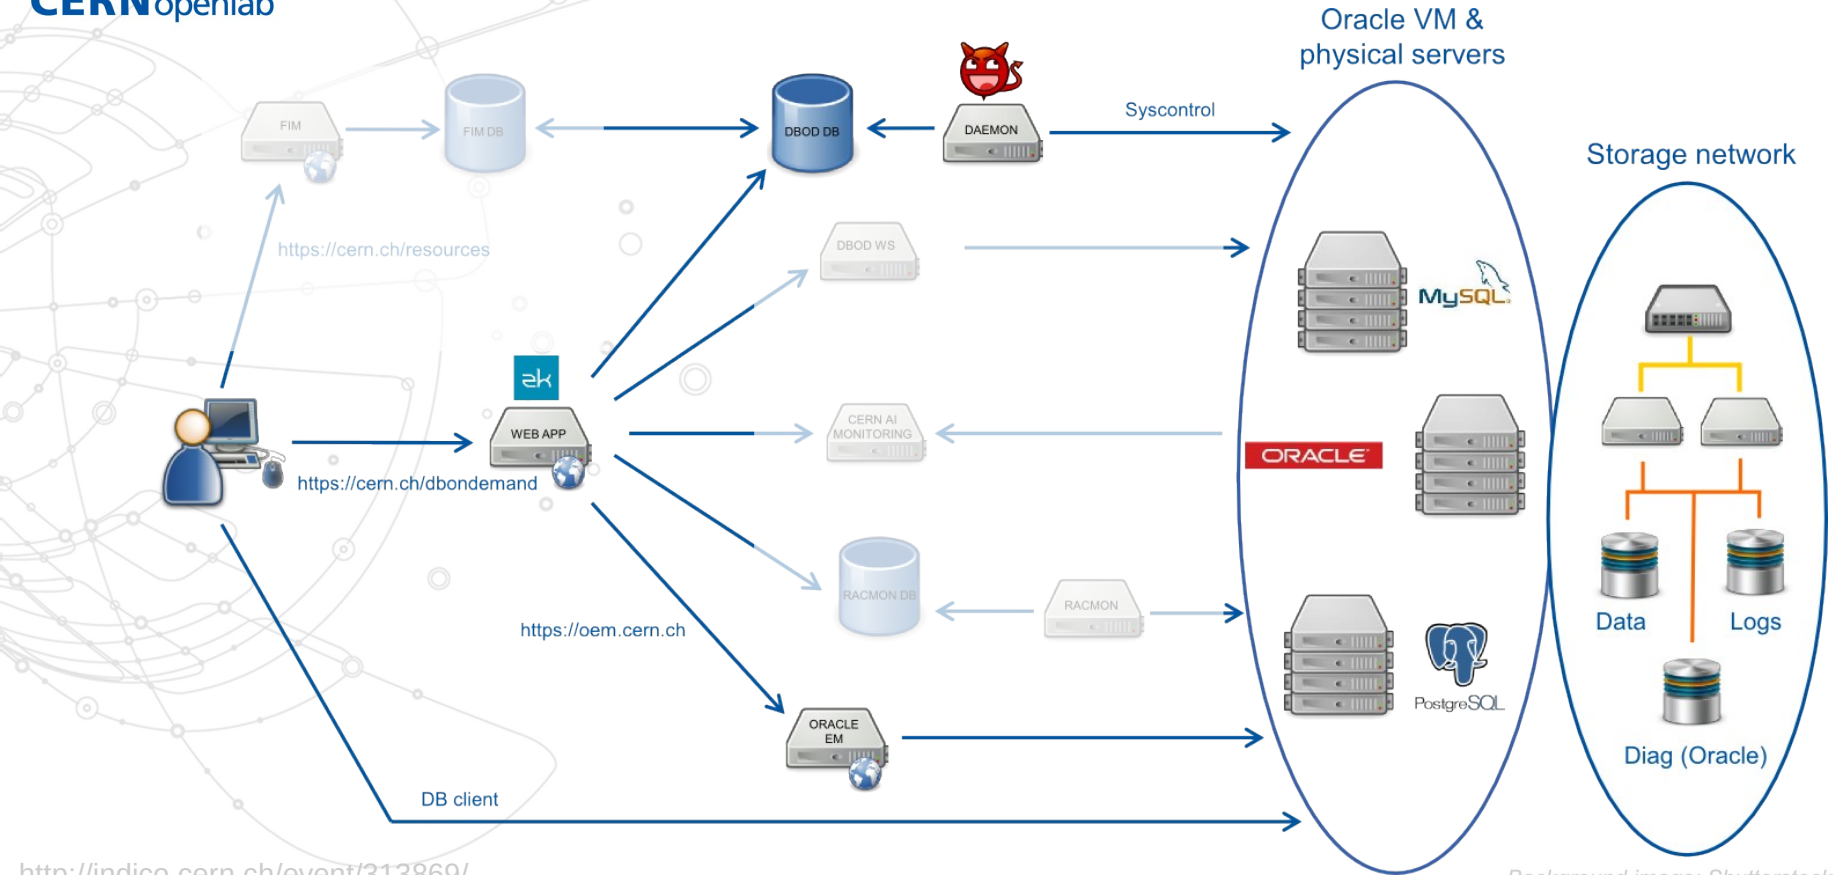
MySQL,
PostgreSQL
and Oracle are
available

CERN Database on Demand

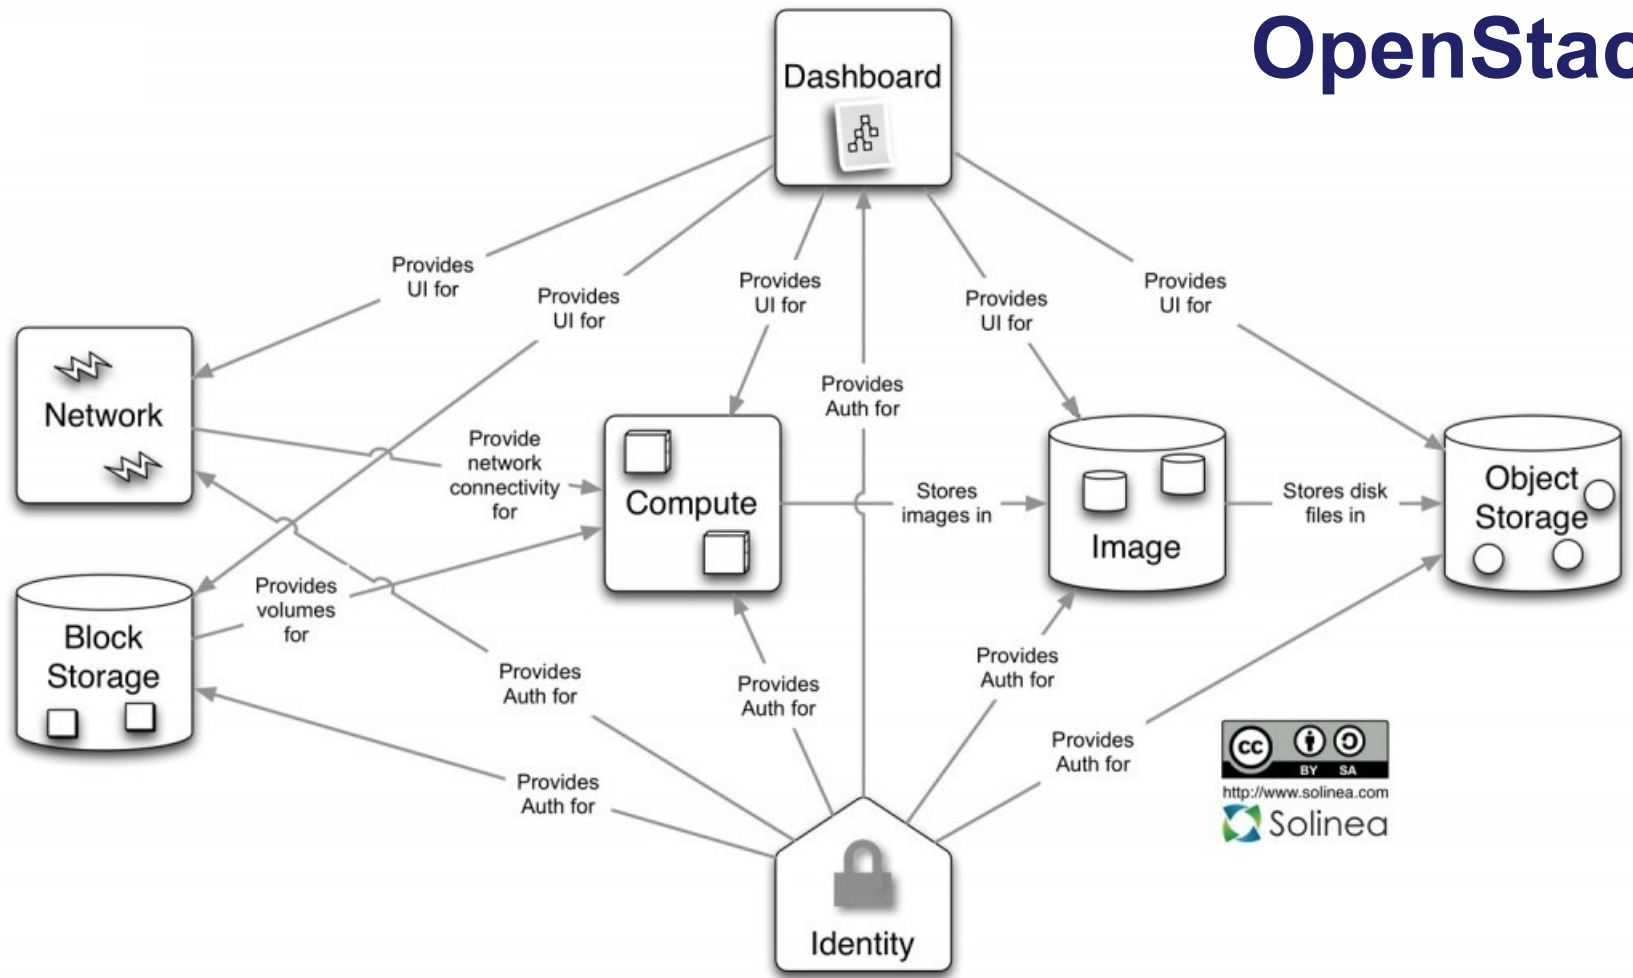
- › IP and credentials
- › Webinterface



CERN Database on Demand



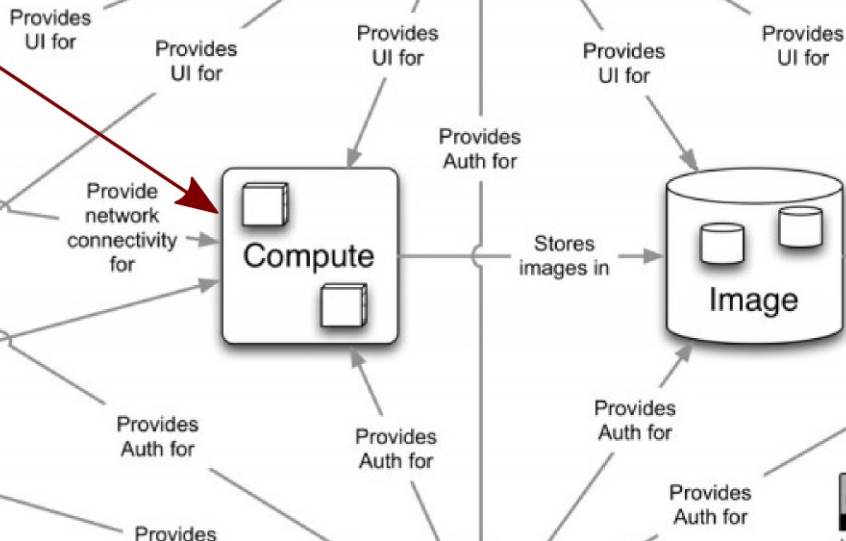




OpenStack



OpenStack



Creates virtual machines with



OpenStack Trove


Compute Node



OpenStack Trove

Compute Node

Virtual Machine 1



Virtual Machine 2



Virtual Machine 3



OpenStack Trove

Compute Node

Virtual Machine 1

MySQL Server

Virtual Machine 2


Redis Server

Virtual Machine 3

MongoDB Server

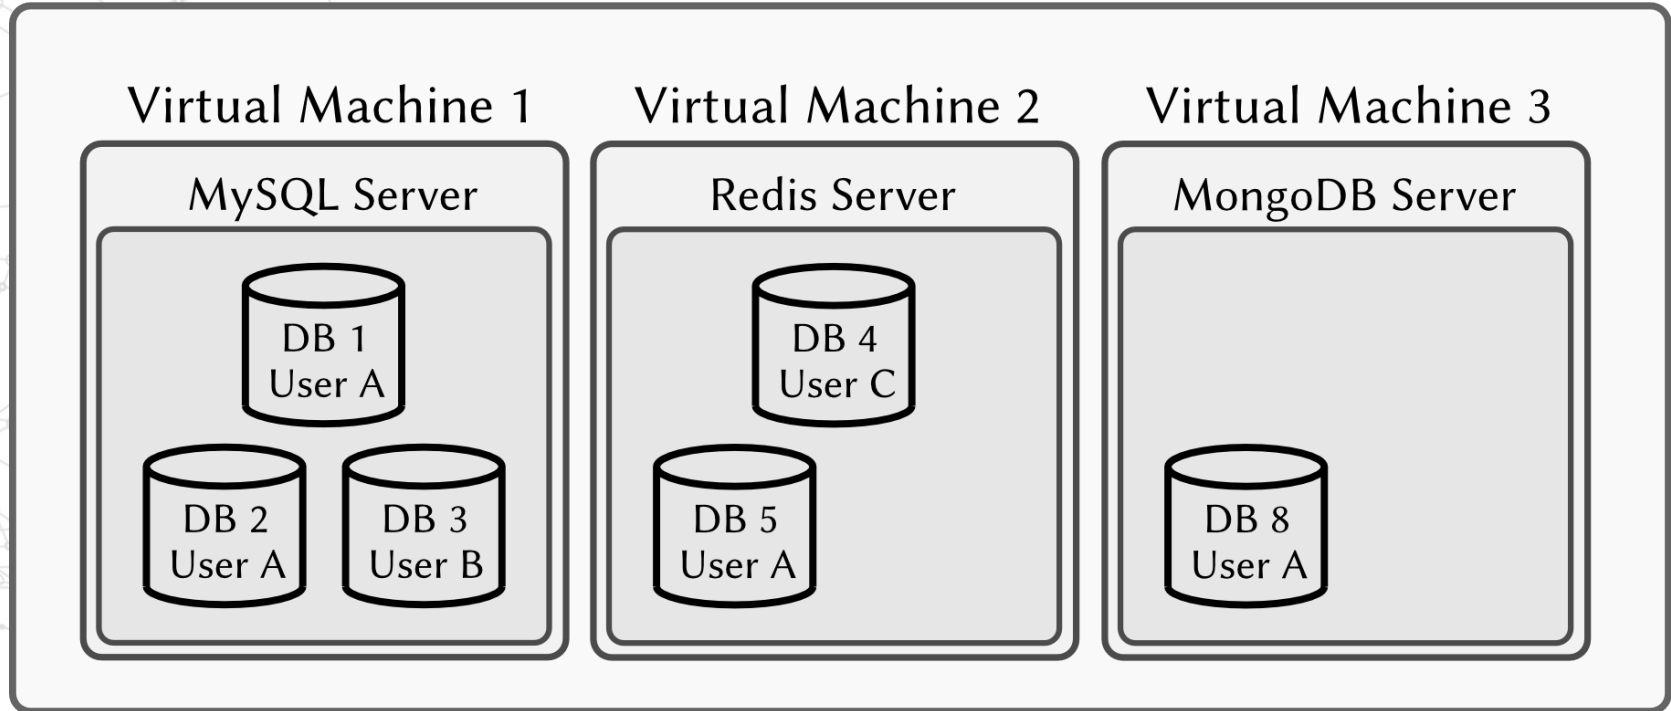
OpenStack Trove

Compute Node



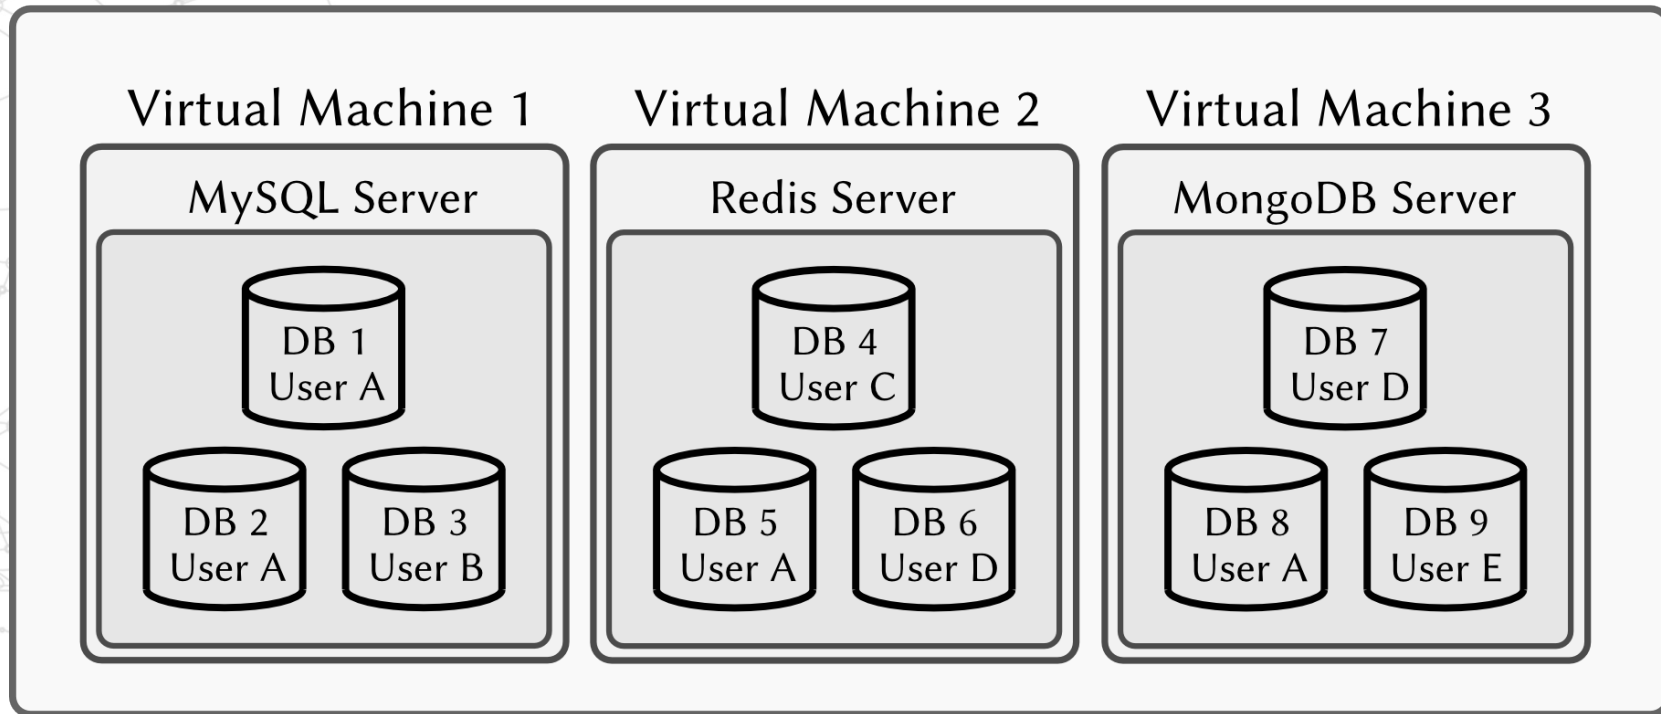
OpenStack Trove

Compute Node

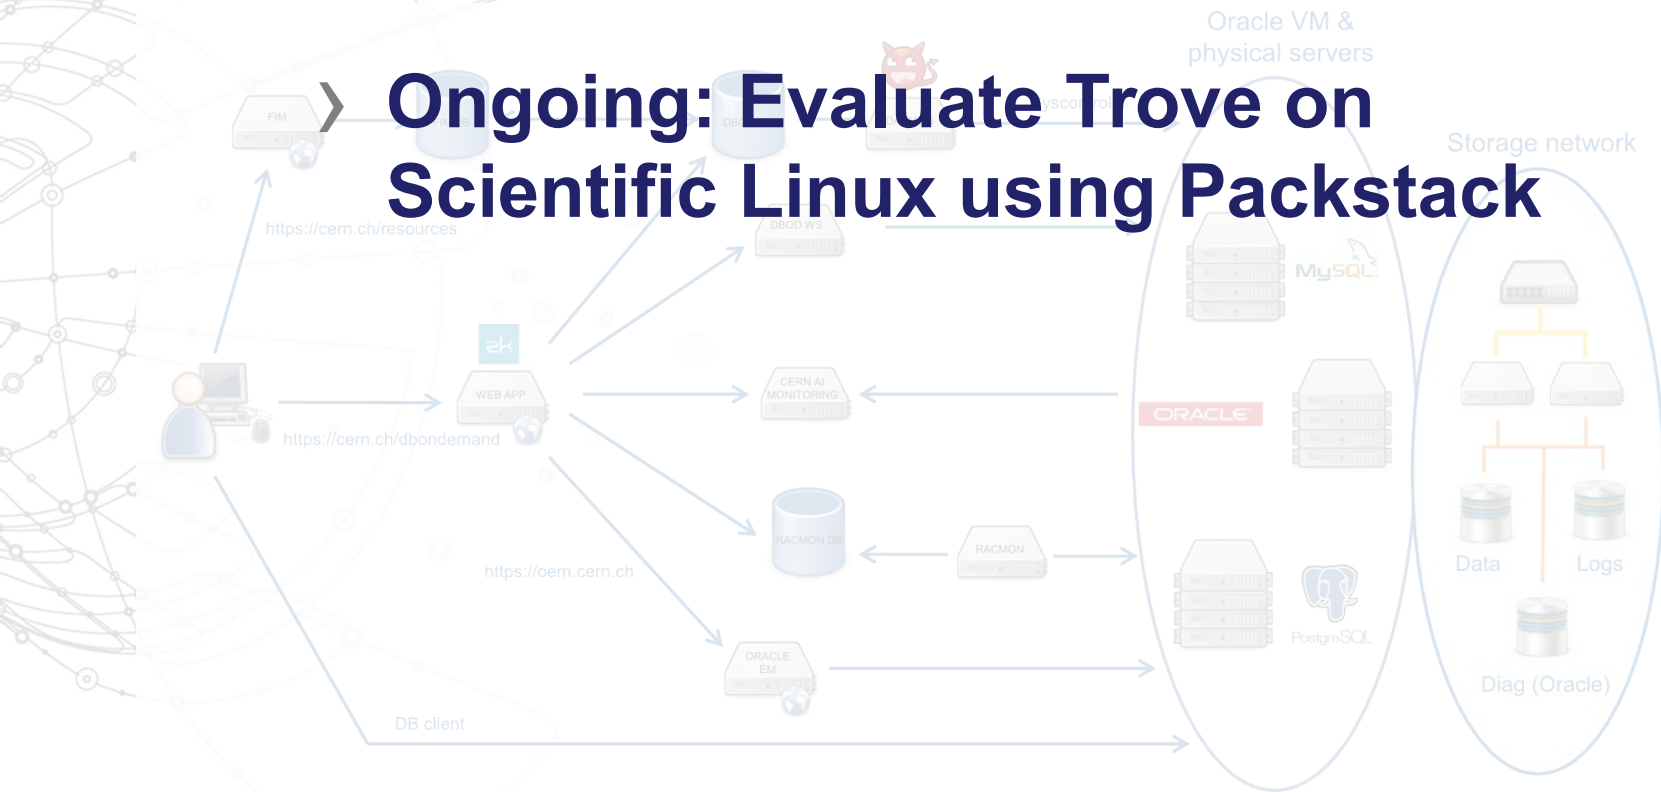


OpenStack Trove

Compute Node




Ongoing: Evaluate Trove on Scientific Linux using Packstack





RDO

 SIGN IN

HOME

QUICK START

DOCS

BLOG

ASK

GET INVOLVED

Q SEARCH

Networking in too much detail

This document describes the architecture that results from a particular OpenStack configuration, specifically:

- Quantum (or Neutron) networking using GRE tunnels;
- A dedicated network controller;
- A single instance running on a compute host

Much of the document will be relevant to other configurations, but details will vary based on your choice of layer 2 connectivity, number of running instances, and so forth.

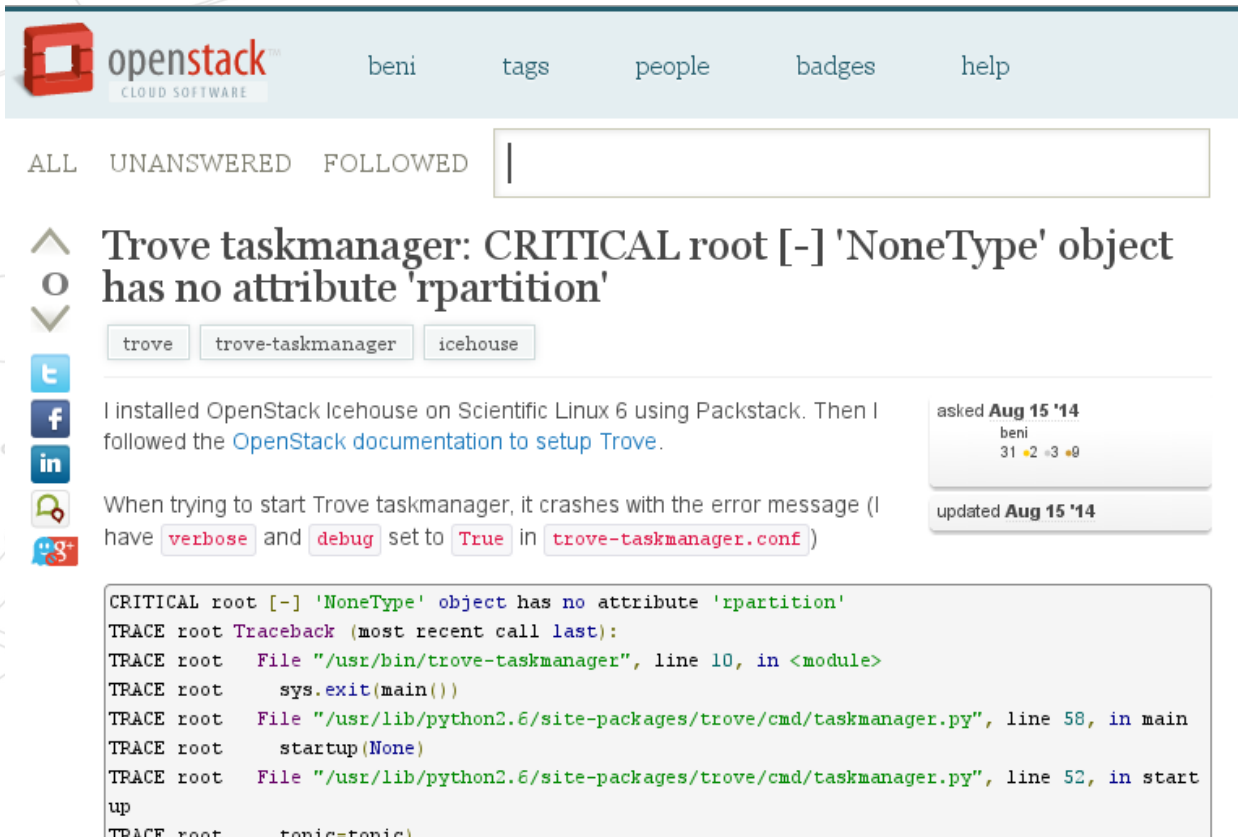
The examples in this document were generated on a system with Quantum networking but will generally match what you see under Neutron as well, if you replace `quantum` by `neutron` in names.

The lay of the land

http://openstack.redhat.com/Networking_in_too_much_detail

This is a simplified architecture diagram of network connectivity in a quantum/neutron managed world:

hutterstock



The screenshot shows the OpenStack Ask interface. At the top is the OpenStack logo and navigation links for 'beni', 'tags', 'people', 'badges', and 'help'. Below the navigation is a search bar and filter tabs for 'ALL', 'UNANSWERED', and 'FOLLOWED'. The main content area displays a question titled 'Trove taskmanager: CRITICAL root [-] 'NoneType' object has no attribute 'rpartition''. The question is tagged with 'trove', 'trove-taskmanager', and 'icehouse'. It includes social media sharing icons for Twitter, Facebook, LinkedIn, and Google+. The question text describes the installation of OpenStack Icehouse on Scientific Linux 6 and the error encountered when starting Trove taskmanager. The error message is shown in a code block. On the right side, there are statistics for the question: 'asked Aug 15 '14' by 'beni' with 31 votes, 2 answers, and 9 comments. Below that, it says 'updated Aug 15 '14'.

openstack™
CLOUD SOFTWARE

beni tags people badges help

ALL UNANSWERED FOLLOWED

Trove taskmanager: CRITICAL root [-] 'NoneType' object has no attribute 'rpartition'

trove trove-taskmanager icehouse

t f in

I installed OpenStack Icehouse on Scientific Linux 6 using Packstack. Then I followed the [OpenStack documentation to setup Trove](#).

asked **Aug 15 '14**
beni
31 2 9

updated **Aug 15 '14**

```
CRITICAL root [-] 'NoneType' object has no attribute 'rpartition'  
TRACE root Traceback (most recent call last):  
TRACE root   File "/usr/bin/trove-taskmanager", line 10, in <module>  
TRACE root     sys.exit(main())  
TRACE root   File "/usr/lib/python2.6/site-packages/trove/cmd/taskmanager.py", line 58, in main  
TRACE root     startup(None)  
TRACE root   File "/usr/lib/python2.6/site-packages/trove/cmd/taskmanager.py", line 52, in start  
TRACE root     up  
TRACE root     tonic=tonic)
```

```
function remove_trove {
    set -x

    service openstack-trove-api stop
    service openstack-trove-taskmanager stop
    service openstack-trove-conductor stop
    chkconfig openstack-trove-api off
    chkconfig openstack-trove-taskmanager off
    chkconfig openstack-trove-conductor off

    #yum remove openstack-trove python-troveclient

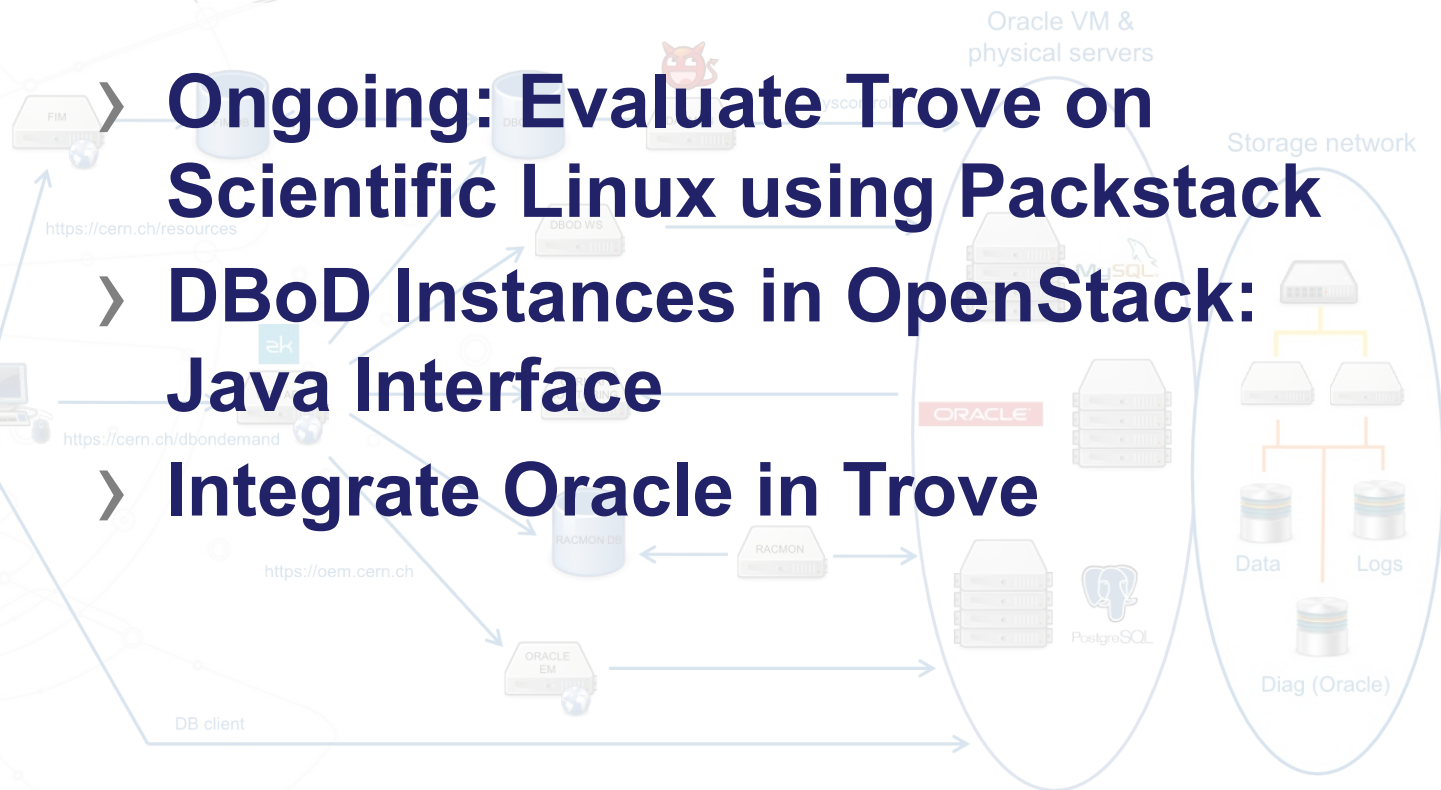
    keystone user-delete trove
    ENDPOINT_ID=$(get_field "endpoint-create-trove" " id " 4)
    keystone endpoint-delete $ENDPOINT_ID
    keystone service-delete trove
    keystone tenant-delete trove

    glance image-delete trove_mysql_ubuntu

    rm -rf /etc/truve/*

    echo "DROP DATABASE trove" | mysql -u root -p$MYSQL_ROOT_PASSWD
    echo "DROP USER trove@'localhost'" | mysql -u root -p$MYSQL_ROOT_PASSWD
    echo "DROP USER trove@'%" | mysql -u root -p$MYSQL_ROOT_PASSWD
}
```


- › Ongoing: Evaluate Trove on Scientific Linux using Packstack
- › DBoD Instances in OpenStack: Java Interface
- › Integrate Oracle in Trove



Interim Conclusion

› Trove Pros

- active development
- OpenStack

› Trove Cons

- not yet mature
- needed improvements:
documentation,
more database types

Benjamin Lipp
mail@benjaminlipp.de

 github.com/blipp

19/08/2014

Background image: Shutterstock

OpenStack Trove: Evaluation and Interfacing with the CERN Database on Demand Service