

A Deep Dive into Trove

SCALE 13x, Los Angeles, CA




Kenneth Rugg

February 22, 2015

Who Am I?

- Ken Rugg, Founder & CEO of Tesora
- Tesora: the Trove Company
 - #1 contributor to Trove project
 - Tesora DBaaS Platform
 - packaged, supported version of Trove
 - Nearly all work is upstream first
 - Database & distribution certifications
- Database Virtualization Engine (DVE)
 - Open source, transparent sharding for MySQL

Trove Contributors (Juno)



Diverse and growing community

- 128 contributors from 28 companies
- 1723 commits, and 231829 lines of code

Transformation of Cloud Data Management

Traditional IT

- Provisioning by DBA's
- Database management by specialists
- Waterfall development
- Few large machines / bare metal
- Oracle enterprise licenses
- Captive audience

Cloud

- Self-service provisioning
- Developers manage their own databases
- Agile development
- Many small machines / virtualization
- Many data management technologies
- Competition with AWS

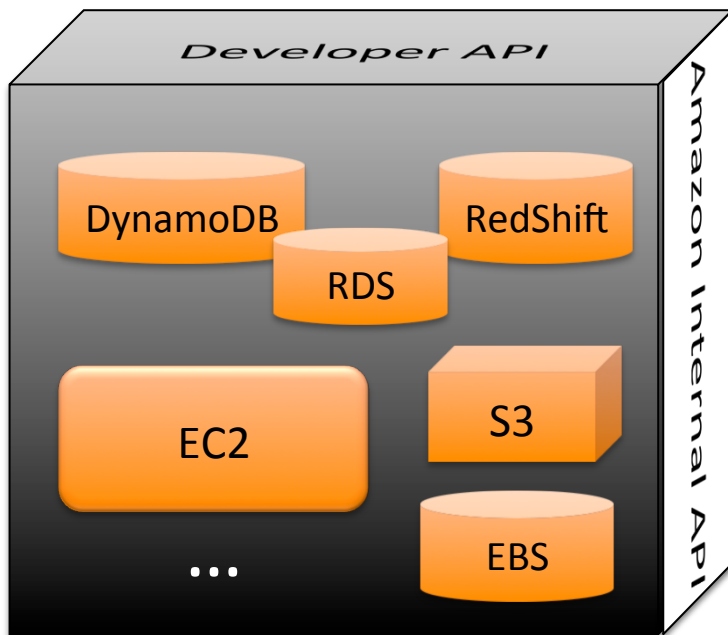
What is Trove?

- Database as a Service for OpenStack
- API's for both development and operations
- Self service database provisioning
- Full database lifecycle management
- Multi-database support
- Both Relational and NoSQL

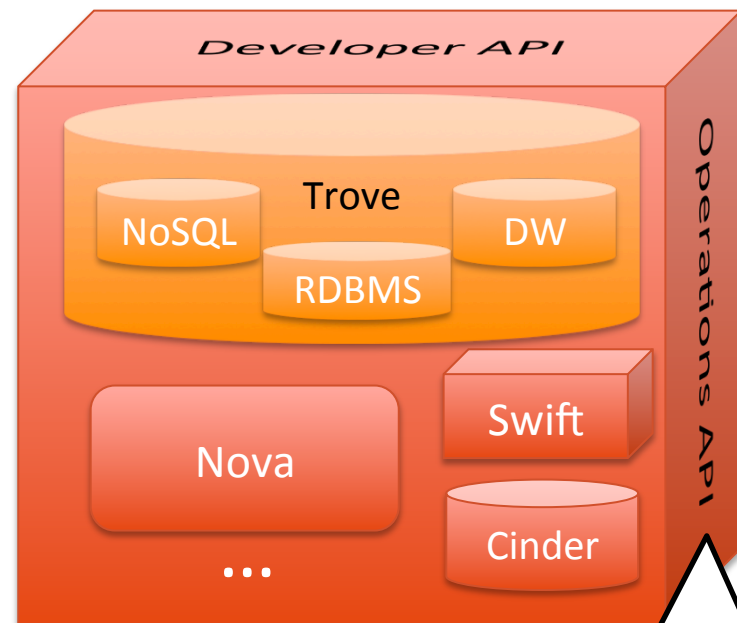


What is Trove? RDS for OpenStack ...and DynamoDB ...and RedShift

Amazon AWS



OpenStack



...and API's for operations

What is Trove? - The Trove Mission

OpenStack Trove mission statement:

*“To provide **scalable** and **reliable** Cloud Database as a Service provisioning functionality for both **relational** and **non-relational** database engines, and to continue to improve its fully-featured and extensible **open source** framework.”*



What is Trove? - More than just provisioning...

Tuning

- Automatically tune my.cnf
 - Buffer Pool Size
 - Log file size
 - max_connections
- Sane defaults
 - InnoDB only
 - Disable load data infile
 - Disable select into outfile
- API to programmatically set configuration groups

Management

- Create database / schema
- Create users
- Grant permissions to a Schema
- Enable root user
- Resize flavor
- Resize volume
- Full and incremental backups
- Create replica

Security

- Security groups
- Turn off SSH
- Remove anonymous user
- Remove non-localhost users
- Remove local file access
- Mangle root user password
- Apply security patches automatically

Trove In Production Today

- Private Cloud: eBay

- Began mid 2013



- Multiple Databases

- MySQL, MongoDB, Redis
Cassandra, Couchbase

- Multi-region + HA

- Working on Clustering



- Public Cloud

- HP Cloud Relational Database

- Launched May 2012



RACKSPACE®
CLOUD DATABASES

- Rackspace Cloud Databases

- Launched August 2012

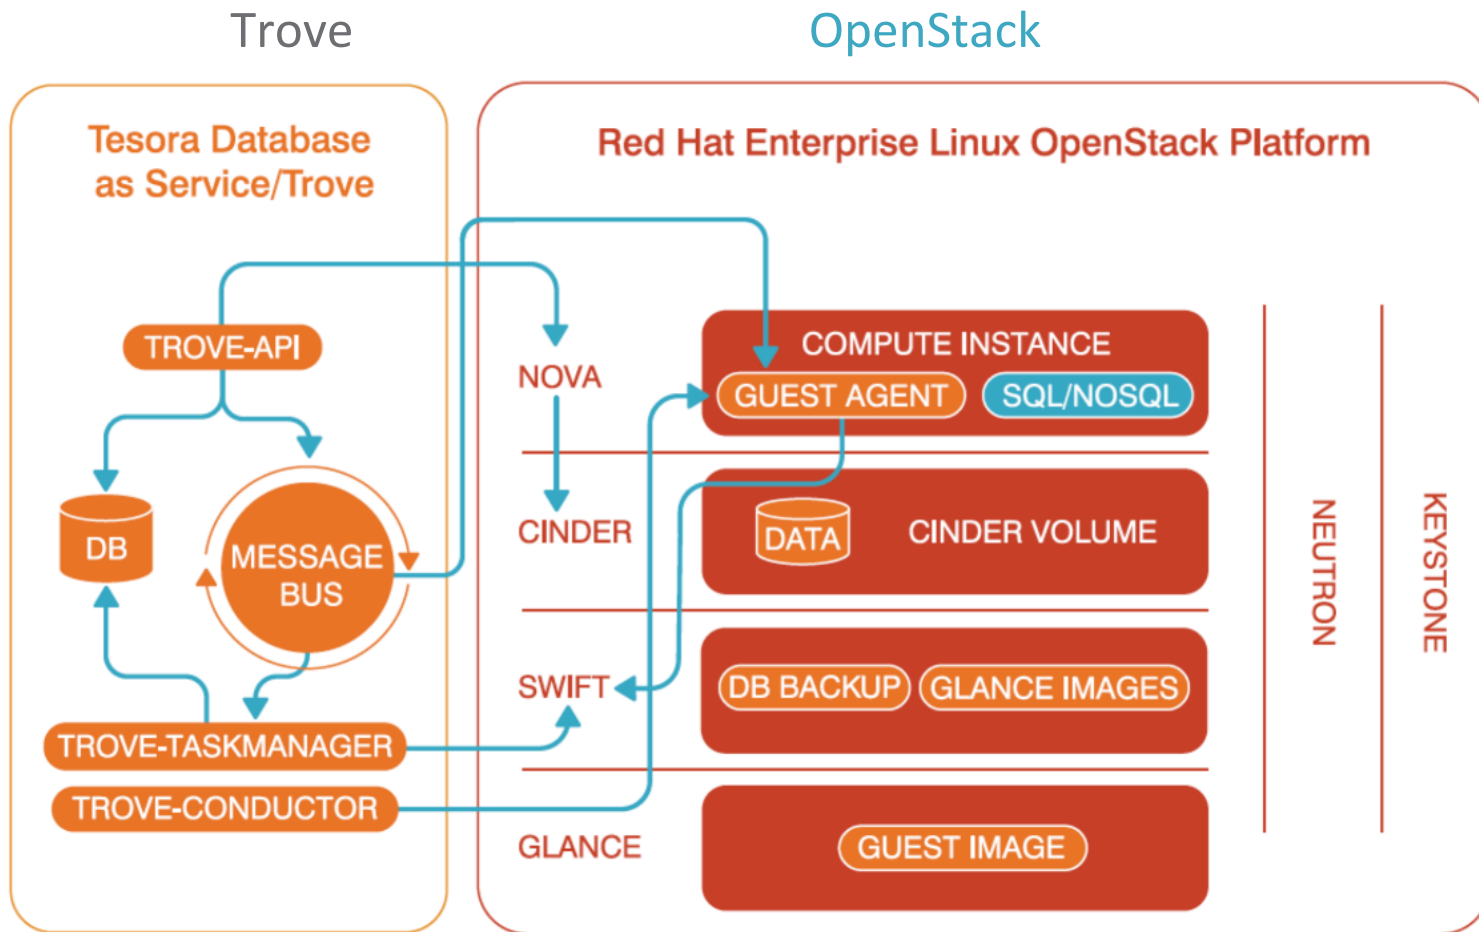
- Key Use Cases

- Development & test
 - Web application hosting
 - On-demand analytics

- Critical Capabilities

- Self-service provisioning & management
 - Fleet wide configuration
 - Multi-datastore architecture

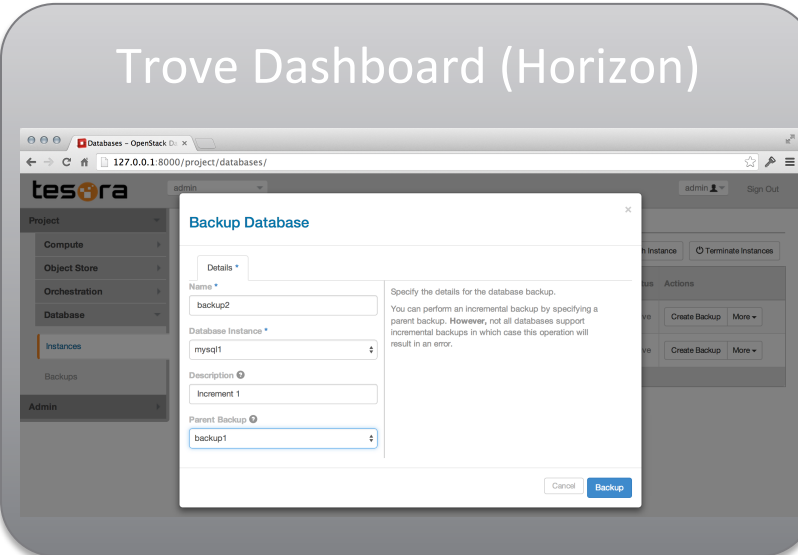
What is Trove? - OpenStack Architecture



What is Trove? Multi-Datastore Architecture

Datastore agnostic code in Trove Controller & Dashboard

Trove Dashboard (Horizon)



Trove Controller

Datastore specific code isolated to Guest Agents



- Guest Agent: MySQL
- Guest Agent: cassandra
- Guest Agent: PERCONA
- Guest Agent: mongoDB
- Guest Agent: MariaDB
- Guest Agent: Couchbase
- Guest Agent: PostgreSQL
- Guest Agent: redis

Trove Architecture: Code Modularity

trove/guestagent/

[...]
backup/
common/
datastore/
[...]
strategies/
[...]

trove/guestagent/strategies/backup

base.py
base.pyc
couchbase_impl.py
couchbase_impl.pyc
__init__.py
__init__.pyc
mysql_impl.py
mysql_impl.pyc
postgresql_impl.py

trove/guestagent/strategies

backup/
[...]
replication/
restore/
storage/

grep class guestagent/strategies/backup/mysql_impl.py

```
class MySQLDump(base.BackupRunner):  
class InnoBackupEx(base.BackupRunner):  
class InnoBackupExIncremental(InnoBackupEx):
```

The background is a vibrant orange color. It features a complex network diagram with numerous nodes of varying sizes connected by thin lines. In the background, there is a faint, repeating pattern of binary code (0s and 1s). A teal-colored shape, resembling a stylized arrow or a corner piece, is located in the bottom-left corner.

Using OpenStack Trove

Getting started with OpenStack Trove

- As a Trove user
 - OpenStack distribution that includes Trove (such as HP Helion Dev Platform)
 - <http://www8.hp.com/us/en/cloud/hphelion-openstack.html>
 - Tesora DBaaS platform, a Trove packaging tailored for the enterprise
 - <http://tesora.com/download-tesora-dbaas-platform-community-edition>
- As a Trove developer

```
$ git clone http://github.com/openstack/trove-integration
$ cd trove-integration/scripts
$ ./redstack install
$ ./redstack kick-start mysql
```
- On top of DevStack
 - Add to localrc:
 - `ENABLED_SERVICES+=,trove,tr-api,tr-tmgr,tr-cond`
 - Swift should also be enabled for Backup and Restore.
 - `ENABLED_SERVICES+=,s-proxy,s-object,s-container,s-account`
 - `SWIFT_HASH=<swift-hash-here>`

Provisioning a database instance with Trove

```
$ trove create <instance-name> <flavor-id>  
--size <volume-size>
```

- Support for flavors
- Support for volumes using Cinder
- Optional parameters to create
 - Image per ‘datastore type’ and ‘version’
 - Support AZs using --availability_zone
 - Support for Neutron using --nic

Managing a database with Trove

- Resize flavor
- Resize volume
- Datastore specific extensions:
 - Create Database / Schema
 - Create Users
 - Grant Users Permissions on Databases
- Enable a Root User

Backup and Restore

```
$ trove backup-create  
  <backup-name> <instance-id>
```

- Optional params:
 - Description through --description
 - Incremental backups using --parent
- To Restore backup use create:

```
$ trove create  
  <instance-name> <flavor-id>  
  --size <volume-size>  
  --backup <backup-id>
```

- Fully managed
- Triggered and tracked via API
- Streamed to Swift (OpenStack Object Storage)
- Incremental & full backups
- Multiple formats per datastore supported via strategies:
 - XtraBackup (Percona)
 - mysqldump

OpenStack Trove: Completed in Juno

- Async MySQL replication (master-slave)
- Clusters for MongoDB
- Neutron Support
- Support for PostgreSQL
- Config-groups enhancements
 - Configuration groups per datastore / version
 - Config-groups for MongoDB
- Backups for Cassandra and Couchbase
- Additional Tempest Tests

New in Juno: Replication

```
$ trove create <instance-name> <flavor-id>  
  --size <volume-size>  
  --replica_of <instance-id>
```

- Support for async MySQL replication (MySQL slave instances)
- Manual detach using

```
$ trove update <instance-id> --detach-replica-  
source
```

New in Juno: Clusters

```
$ trove cluster-create <datastore> <ds-version>
```

- Optional parameters to cluster-create
 - `--instance <flavor_id=flavor_id,volume=volume>`
 - Specify multiple times to create multiple instances for your cluster
- Initial support in Juno added for MongoDB Clusters
 - Sets up mongo config server, and mongo query routers
 - Transparent and driven by configuration options
 - Support adding shards to existing cluster for horizontal scale out.

Planned for Kilo

- Building out clusters
 - Semi Synchronous MySQL clusters (Galera)
- Async Replication v2
 - GTID based replication
 - Manual failover support
- Associate flavors with datastores
- Access datastore logs via API
- Removing deprecated oslo-incubator code
- Upgrade testing through grenade

Where can I get OpenStack Trove?

- Get the source
 - OpenStack Trove Project
 - <https://git.openstack.org/cgit/openstack/trove>
 - <http://github.com/openstack/trove.git>
 - Trove python client binding and command line client
 - <https://git.openstack.org/cgit/openstack/python-troveclient>
 - <http://github.com/openstack/python-troveclient.git>
 - Trove design specifications for blueprints
 - <https://git.openstack.org/cgit/openstack/trove-specs>
 - <http://github.com/openstack/trove-specs>
 - Trove scripts for installation and testing, and elements for building guest images
 - <https://git.openstack.org/cgit/openstack/trove-integration>
 - <http://github.com/openstack/trove-integration.git>
- Installable packages and guest images from Tesora
 - Community Edition:
<http://tesora.com/download-tesora-dbaas-platform-community-edition>
 - Enterprise Edition trial:
<http://tesora.com/download-tesora-dbaas-platform-enterprise-edition>

More about Trove

- Trove Wiki
 - <https://wiki.openstack.org/wiki/Trove>
- Trove Source
 - <https://git.openstack.org/cgit/openstack/trove>
- On IRC
 - #openstack-trove
- Trove Day
 - Tesora.com/troveday
 - Slideshare.net/Tesoracorp
- LinkedIn
 - OpenStack Trove Group

Contact information

Contact Tesora

info@tesora.com

www.tesora.com

@tesoracorp

Contact Ken

krugg@tesora.com

www.tesora.com

@kenrugg

The background is a vibrant orange color. It features a complex network of thin, light-colored lines connecting various sized circular nodes, creating a web-like pattern. In the bottom-left corner, there is a teal-colored triangular shape pointing towards the center. The text "Thank You!" is centered in the middle of the orange area.

Thank You!