



Best Practices for Deploying OpenStack Trove:

An Inside look at Database as a Service Architecture

OpenStack Summit at Barcelona, October 2016

Who are we?

- Sriram Kalyanasundaram, Director Implementations
 - Tesora Inc.

What are we covering in this session?

- Introduction to Trove
 - Trove Concepts
 - Tutorial demonstrating Trove functionality
 - Trove Architecture
- Access to Trove environment
 - DevStack VM
 - Trove VM
 - MySQL and MongoDB Guest Images
 - Step-by-step Lab Tutorial
- Questions

What is Trove?

Introduction to OpenStack Trove Project

- Motivation: Provide DBaaS within OpenStack framework
 - Incubated in Havana (October 2013)
 - Integrated in Icehouse (April 2014)
- Original project sponsors: HP and Rackspace
- Major contributors include Tesora, Rackspace, HP, IBM, Redhat, eBay, Mirantis

What does Trove provide?

- 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

Supported Databases

- Available today

- MySQL 5.6,5.7
- Percona 5.6, 5.7
- Percona XtraDB Cluster 5.6, 5.7
- MariaDB 10.1
- MongoDB 3.2
- Cassandra 2.2, 3.0
- Couchbase 3.0,4.0
- Couchdb 1.6.1
- Redis 3.0
- PostgreSQL 9.4
- DB2 Express 10.5
- Vertica

Available today

- Oracle MySQL Enterprise 5.6*
- EDB PostgreSQL Enterprise 9.4*
- Couchbase Enterprise 4.1*
- MongoDB Enterprise 3.2*
- DataStax Enterprise 4.8*
- Oracle 12c*, 11g*
- Oracle RAC*

Coming soon

- MS SQL Server

*Tesora DBaaS Enterprise Edition only


Complete Database Lifecycle Management

Provision

- Wide array of databases
- One click away
- Single instances to clusters

Secure

- Automated patching
- Granular permissions
- Restricted root access



Manage

- Databases
- Replica sets
- Users and backups

Tune

- Database images
- Optimized and tuned
- APIs for custom configs

Trove Terminology

- Guest Image
- Guest Agent
- Trove Instance
- Cluster
- Datastore
- Datastore Version
- Configuration Group
- Flavor

Provisioning a Database Instance with Trove

- Database instances can be deployed using dashboard or CLI
- Trove API allows options to customize the database instance
 - Flavor support
 - Cinder volume support
 - Create database(s)
 - Create user(s)
 - Manage configurations
 - Replicate from another instance
 - Restore from a backup

Trove Demo

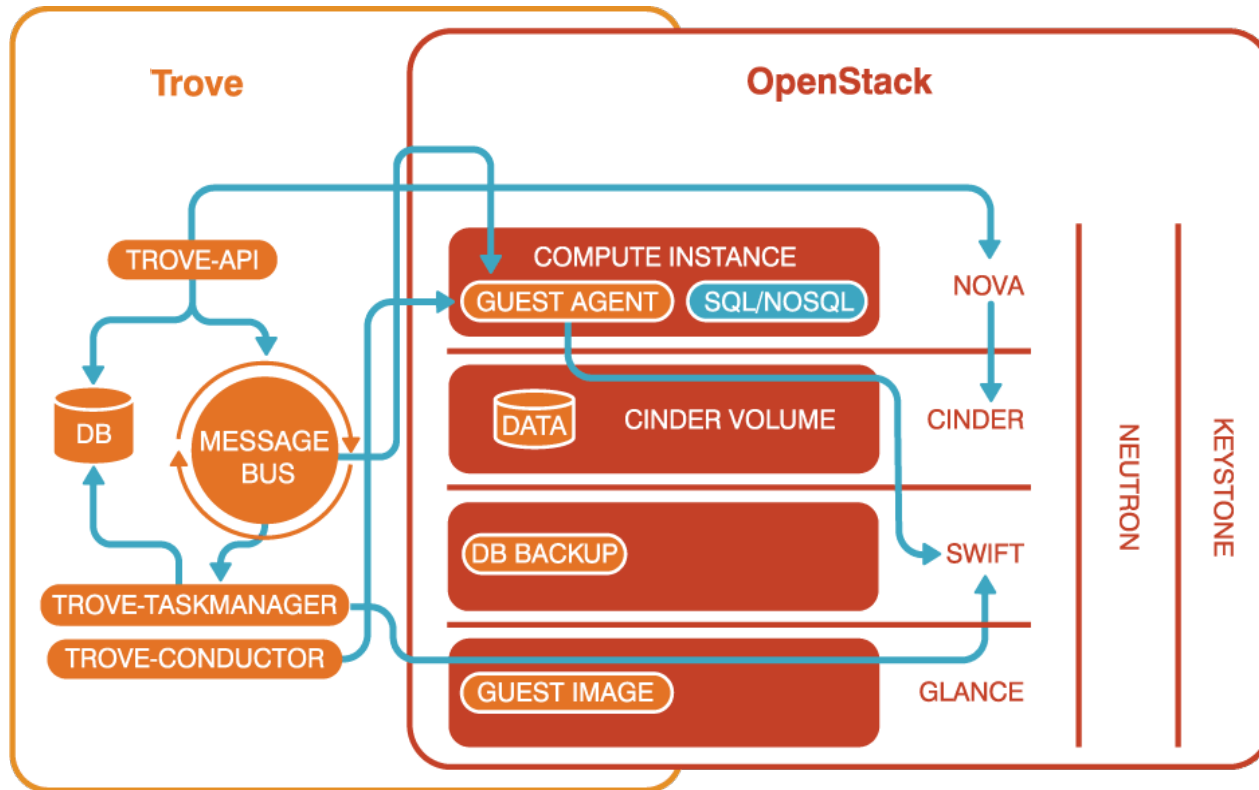
Trove in Action

In this demo we will...

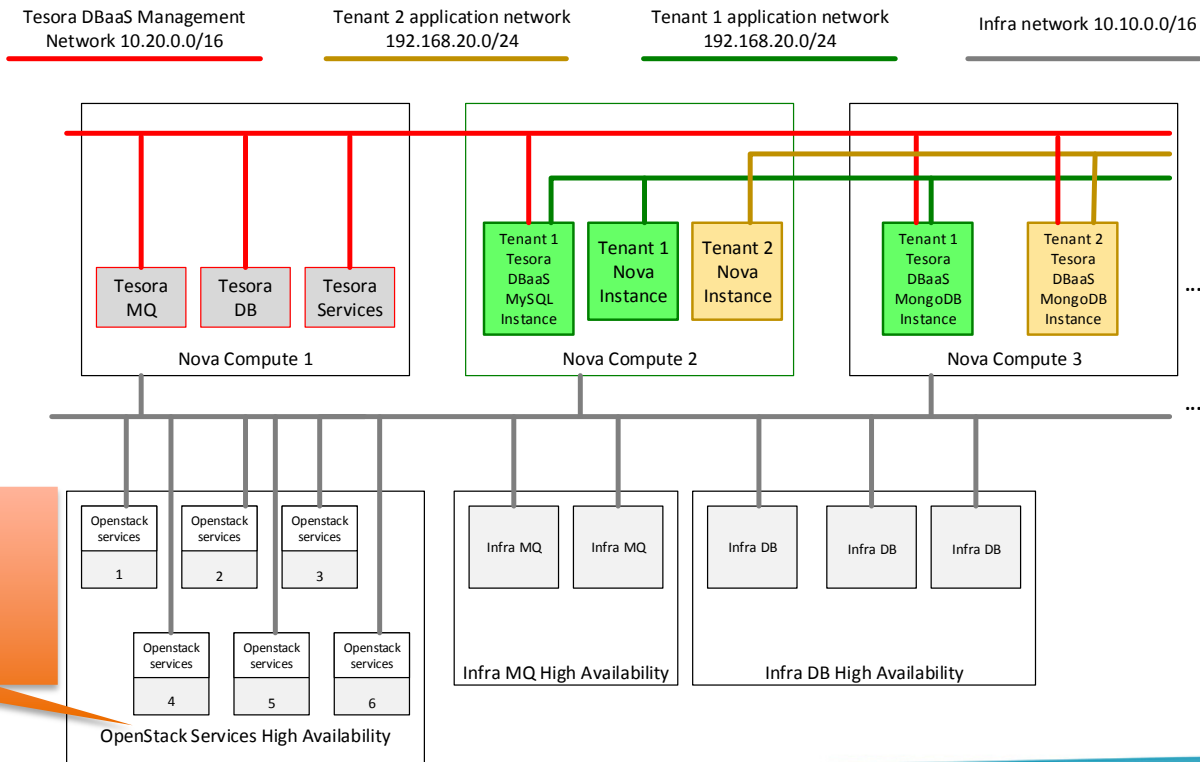
- Create a *MySQL* instance through Horizon
- Create a backup of the instance
- Create a *MongoDB* cluster

Trove Architecture and HA Deployment

Tesora DBaaS Platform Architecture

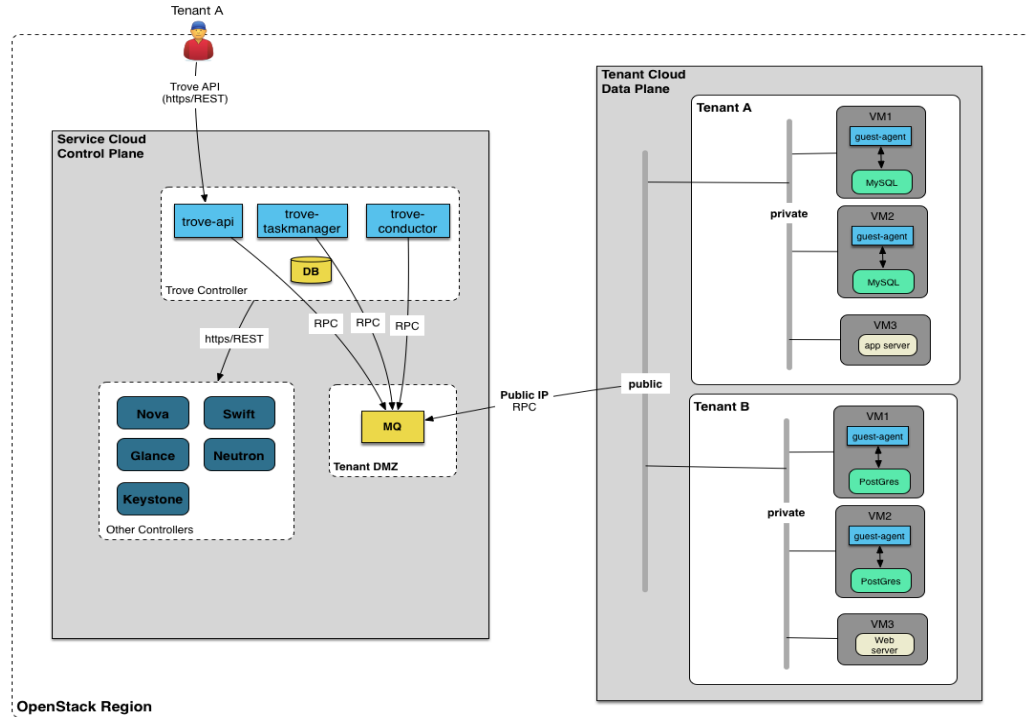


Sample POC architecture

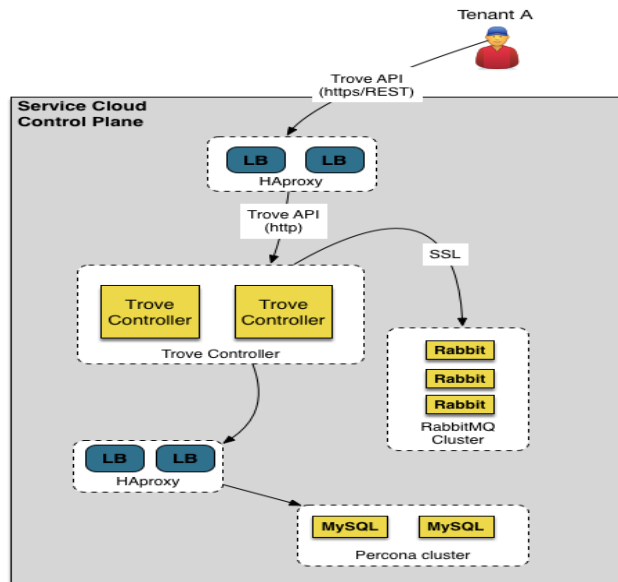


An existing
OpenStack
Deployment

Trove Deployment at Cisco



High Availability Configuration



What's new in Newton?

- Upgrade support
- Usability improvements
- Clustering improvements
- Locality support in instance creation and replication
- Improved DB2 Express-C Support
- Improved PostgreSQL Support
- New quota management [admin] API for reviewing and changing quota's for specific tenants

If you are interested in learning further:

Please attend the session on Thursday at

- 1:50 PM in P1 - Room 117
- *What's New with OpenStack Trove in Newton, What's On Deck for Ocata*



Questions?

If you are interested in deploying Trove

Contact Tesora for additional information

- <http://www.tesora.com/contact-us/>

If you want additional information

- Sriram Kalyanasundaram
 - sriram@tesora.com
 - 978-273-0607



Thank You