



OpenStack Trove

Project overview and update

Bartosz Żurkowski

Cloud and Big Data Engineer,
Samsung R&D Institute Poland

IRC:
bzurkowski

EMAIL:
b.zurkowski@samsung.com

Outline

- What is Trove?
- Core features overview
- Deployment stats
- Rocky highlights
- Stein goals
- Community update

What is Trove?

- Database as a Service for OpenStack
- Provides full database lifecycle management
 - Provisioning, configuration, backups, scaling
- Multi-datastore support
 - 11 database engines
 - Relational, non-relational
 - Single-instance and clustered deployments
- Unified management interface
- Built entirely on OpenStack
 - Synergy of Nova, Cinder, Swift, Glance and Neutron

Application optimization
DB performance tuning
Replication and clustering
Scaling
Periodic backups
DB software upgrades
Hardening
DB software setup and config
Virtual resource provisioning

■ DBA responsibilities

■ DBaaS (Trove) responsibilities

Core features overview

- Instance provisioning
- Instance resizing (volume, flavor)
- Database and user management
- Configuration groups
- Backups (full, incremental, scheduled)
- Datastore upgrades
- Logs (guest, database)
- Security groups management
- Flavors management
- Cluster provisioning
- Cluster sizing (grow, shrink)
- Replication setup
- Replication failover (promote read replica, eject source)

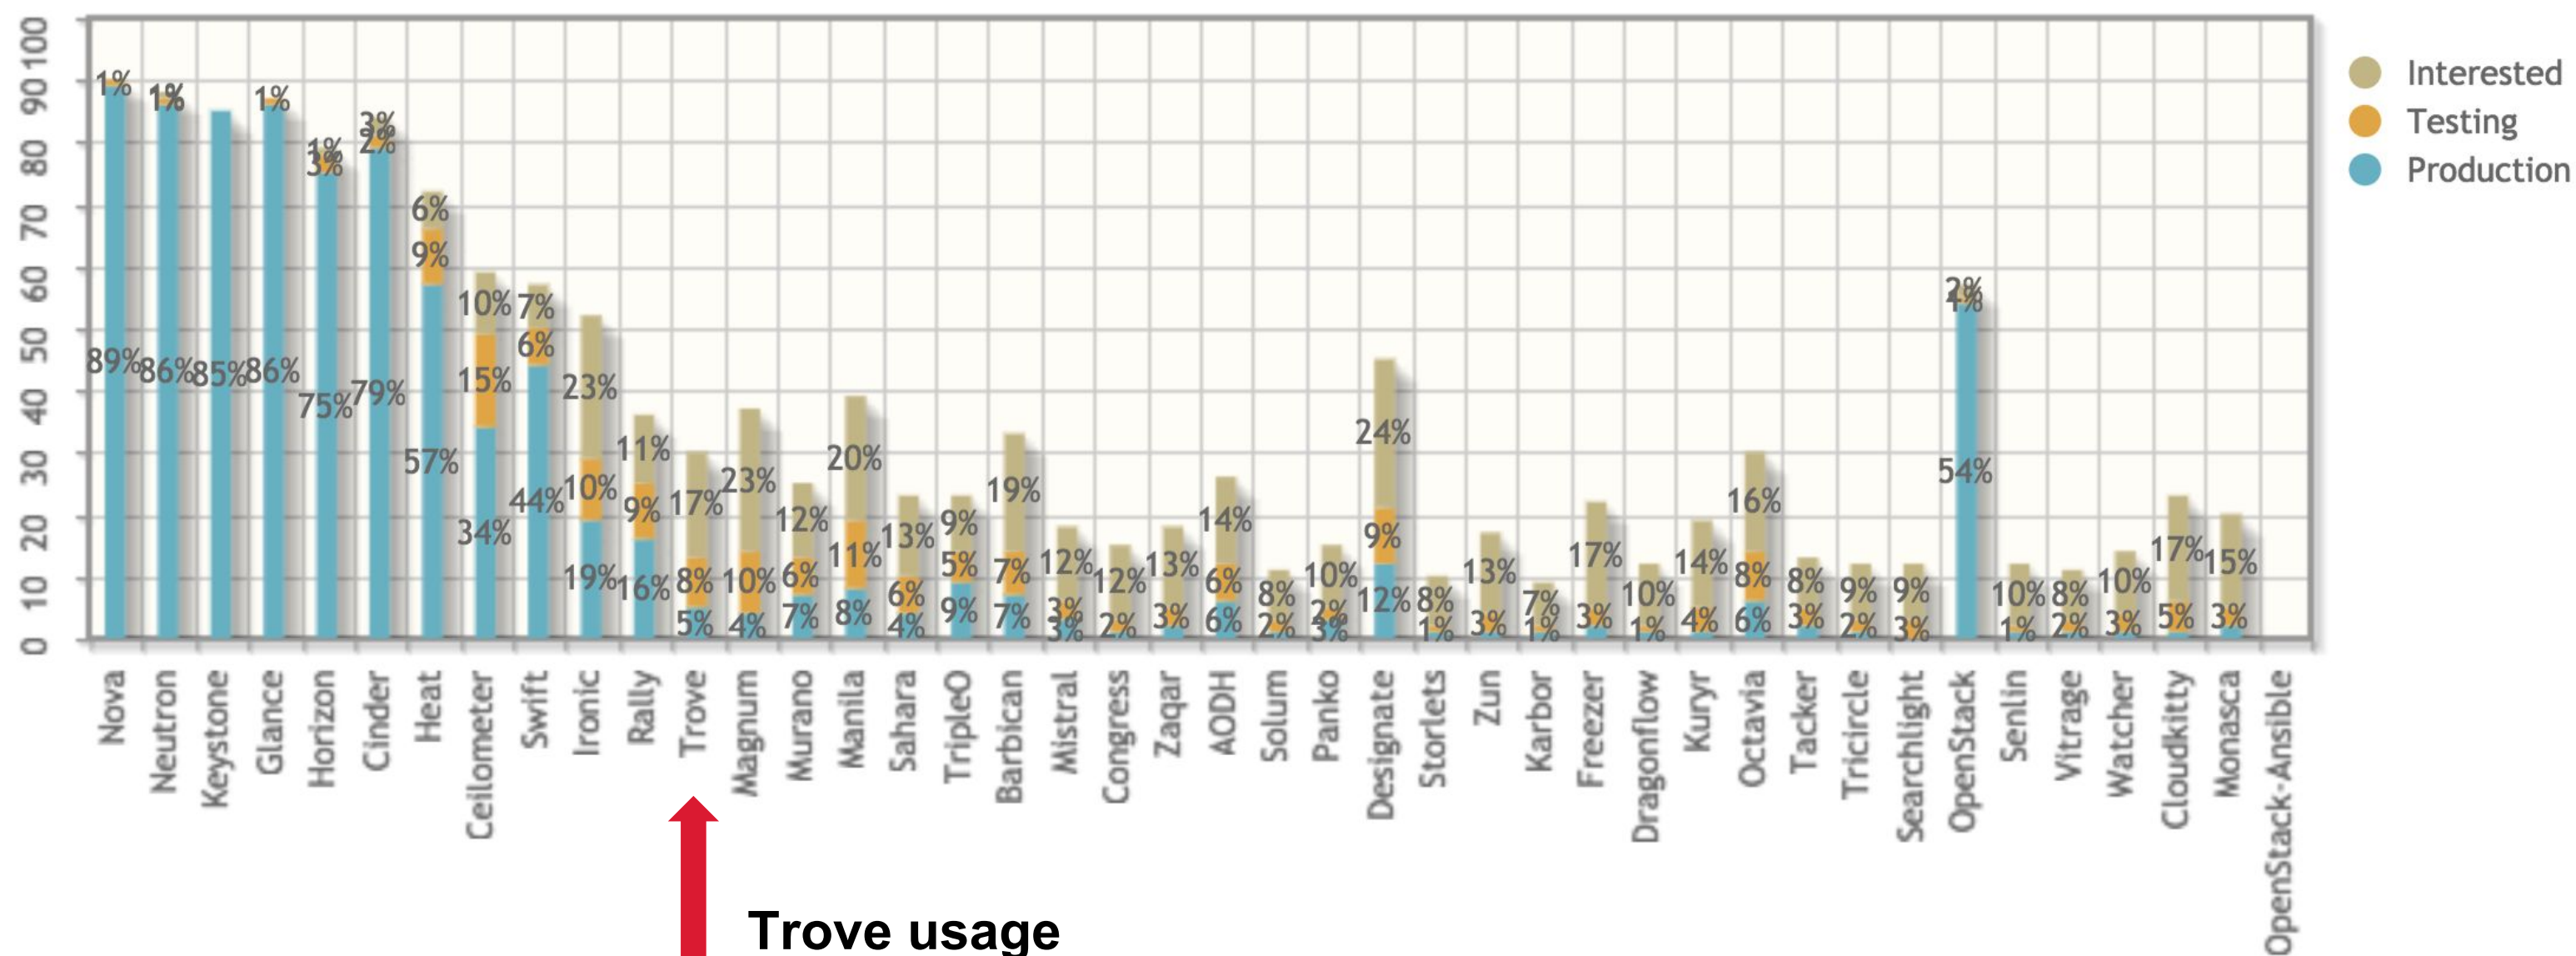
Did you find your favourite database on the slide?

Datastore capability matrix

Datastore	Type	Provisioning	Schema management	Backup and restore	Resizing	Replication	Clustering
MySQL	Relational	✓	✓	✓	✓	✓	✓
MariaDB		✓	✓	✓	✓	✓	✓
PostgreSQL		✓	✓	✓	✓	✓	✗
Percona		✓	✓	✓	✓	✓	✓
DB2		✓	✓	✓	✓	✗	✗
Redis	Key-value	✓	—	✓	✓	✓	✓
Cassandra	Column	✓	✓	✓	✓	—	✓
Vertica		✓	✓	✗	✓	—	✓
MongoDB	Document	✓	✓	✓	✓	—	✓
CouchDB		✓	✓	✓	✓	—	✗
Couchbase	Multi-model	✓	✗	✓	✓	—	✗

Deployment stats

Projects used in Production Deployments



Trove usage
 Production: 5%
 Testing: 8%
 Interested: 17%

n=357

Rocky highlights

➔ [OPENSTACK WIDE GOAL] ENABLE MUTABLE CONFIGURATION

- Some configuration options can now be changed without restarting Trove services

➔ [OPENSTACK WIDE GOAL] REMOVE USE OF MOX/MOX3 FOR TESTING

- Mox package is no longer actively maintained and never updated to add support for Python 3
- To provide a clean path towards Python 3 support we migrated all projects to mock

➔ MIGRATION TO OPENSTACK CLIENT

- Progressing with alignment of OpenStack CLI (`openstack database . . .`) with Trove CLI
- Added support for new commands in scope of: clustering, replication, configuration and executions
- <https://etherpad.openstack.org/p/trove-support-python-openstackclient>

Rocky highlights

➔ NOVA FILE INJECTION DEPRECATION

- Nova deprecated personality files used by Trove to inject guest config into VM
- Prepared spec proposing adaptation of user data instead
- <https://specs.openstack.org/openstack/trove-specs/specs/rocky/adapt-to-file-injection-deprecation-in-nova.html>


➔ INTEGRATION TESTS IN TEMPEST

- Long-term goal: replace all tests from Trovestack framework (based on Probasic) with Tempest
- Added Tempest job to Zuul CI pipeline - now running basic API test suite

➔ GATE IMPROVEMENTS

- Migrated legacy jobs to Zuul V3

Rocky stats



- Commits: 108
- LOCs: 11k
- Contributors: 31 members
- Reviewers: 47 members
- Core team: 14 members

Stein goals

➔ [OPENSTACK WIDE GOAL] RUN UNDER PYTHON 3 BY DEFAULT

- Run all CI jobs using Python 3 by default
- <https://etherpad.openstack.org/p/trove-python-3>

➔ [OPENSTACK WIDE GOAL] SUPPORT PRE UPGRADE CHECKS

- Provide upgrade checks that can be run prior to upgrade that will help identify any known issues that would result in upgrade failing (e.g. deprecated configuration options)

➔ MIGRATION TO OPENSTACK CLIENT

- Support more Trove commands in OpenStack CLI
- <https://etherpad.openstack.org/p/trove-support-python-openstackclient>

Stein goals

➔ NOVA FILE INJECTION DEPRECATION

- Implement user data adaptation spec
- <https://specs.openstack.org/openstack/trove-specs/specs/rocky/adapt-to-file-injection-deprecation-in-nova.html>

➔ INTEGRATION TESTS IN TEMPEST

- Setup testing environment (build guest image, generate Tempest config)
- Start covering Trove with Tempest scenarios

➔ DOCUMENTATION IMPROVEMENTS

- Update pages about architecture, data flow, instance states etc.
- Update datastore capability matrix

Stein goals

➔ GUEST IMAGE BUILDING IN DEVSTACK

- Enhance Devstack plugin by building guest image
- Parametrize datastore type and version

➔ NEW COMMUNICATION CHANNEL BETWEEN CONTROL AND DATA PLANE

- Disable guest access to message broker
- Inspired by Octavia (LBaaS)
- Spec design in progress: <https://review.openstack.org/#/c/553679/>

Community update

- Samsung R&D joined the project in the mid of Rocky release
- 3 new active core contributors
- New Project Technical Lead: Dariusz Król (dkrol) dkrol3@gmail.com
- Plan for revitalization (<https://etherpad.openstack.org/p/trove-revitalization>)
 - Evangelize
 - Understand user requirements
 - Improve Trove documentation
 - Simplify guest image building

We need Your help!



- Code reviews
- Feature requests and blueprints
- Bug fixes
- Documentation
- Testing

Weekly meetings

Wed 14:00 UTC on IRC

openstack-meeting-alt channel

IRC channel

openstack-trove

Questions & Answers



openstack



@OpenStack



openstack



OpenStackFoundation

Thank You!



openstack



@OpenStack



openstack



OpenStackFoundation