

OpenStack: How do I install OpenStack's Trove into Bright?

How do I install OpenStack's Trove into Bright?

Trove is OpenStack's reference database service. At the time of writing it is undergoing heavy development. This means that administrators may need to roll up their sleeves and fix issues with it.

How to install Trove Database service on an existing Bright OpenStack installation.

The following guide got Trove from around mid October 2014 up and running with Bright. It is based on the OpenStack Installation Guide (http://docs.openstack.org/icehouse/install-guide/install/yum/content/ch_trove.html) of around that time, with some appropriate text updates.

1. Install the packages using yum:

```
# yum install openstack-trove python-troveclient
```

2. Create the Trove user and give the user an admin role on the service tenant:

```
keystone user-create --name trove --pass PASSWORD
keystone user-role-add --user trove --tenant service --role admin
```

3. Protect & create the Trove service endpoints

The service endpoints in Trove are stored and extracted from a database. CMDaemon needs to be prevented from overwriting these settings. This can be done by modifying `openstackcommands.py` (path `/cm/local/apps/cmd/scripts/cloudproviders/openstack/openstackcommands.py`). In this file, the function `handleActionEnsureEndpoints()` should be commented out:

Thus:

```
#def handleActionEnsureEndpoints(request):
#    start = time.time()
#    [...]
#    logd("Duration of %s: %f seconds" % (inspect.stack()[0][3],
time.time() - start))
#    return output
```

OpenStack: How do I install OpenStack's Trove into Bright?

The endpoints can now be created manually with:

```
keystone service-create \  
--name=trove --type=database --description="OpenStack Database  
Service"  
keystone endpoint-create \  
--service-id=$(keystone service-list | awk '/ trove / {print  
$2}') \  
--publicurl=http://master:8779/v1.0/%\(tenant_id\)s \  
--internalurl=http://master:8779/v1.0/%\(tenant_id\)s \  
--adminurl=http://master:8779/v1.0/%\(tenant_id\)s
```

4. Create a database for Trove and give the system user trove access to it:

```
$ mysql -u root -p  
mysql> CREATE DATABASE trove;  
mysql> GRANT ALL PRIVILEGES ON trove.* TO trove@'localhost'  
IDENTIFIED BY 'TROVE_DBPASS';  
mysql> GRANT ALL PRIVILEGES ON trove.* TO trove@'%' IDENTIFIED BY  
'TROVE_DBPASS';
```

4. Edit the files `trove.conf`, `trove-taskmanager`, `trove-conductor`. Add the following configuration to [Default]:

```
log_dir = /var/log/trove  
trove_auth_url = http://master:5000/v2.0  
nova_compute_url = http://master:8774/v2  
cinder_url = http://master:8776/v1  
swift_url = http://master:8080/v1/AUTH_  
sql_connection = mysql://trove:TROVE_DBPASS@master/trove  
notifier_queue_hostname = master  
rpc_backend = trove.common.rpc.impl_fake      ### some  
don'ts:  
  
### don't add rabbit as a rpc_backend, Trove doesn't get  
that  
  
### also trove.openstack.common.rpc.impl_kombu was broken  
at the time  
  
### of writing, but this will probably have changed in a  
few days or weeks.  
rabbit_host = master  
rabbit_password = g3dtrA45t2liS5rMv1sXEWoV
```

OpenStack: How do I install OpenStack's Trove into Bright?

5. **Edit** `/usr/share/trove/trove-dist-paste.ini`, by adding under `[filter:authtoken]`:

```
auth_host = master
auth_port = 35357
auth_protocol = http
admin_user = admin
admin_password = system
admin_tenant_name = service
```

6. **Edit** `trove.conf`, by adding:

```
[DEFAULT]

default_datastore = mysql

add_addresses = True

network_label_regex = ^NETWORK_LABEL$
```

7. **Edit** `trove-taskmanager`, by adding the following lines under `[Default]`:

```
nova_proxy_admin_user = admin

nova_proxy_admin_pass = 3de4922d8b6ac5a1aad9
nova_proxy_admin_tenant_name = admin
```

8. **Create the database schema for Trove:**

```
#trove-manage db_sync
```

9. **Now create the Datastore.** Each datastore type like mysql, cassandra, mongodb has to be run separately:

```
$ su -s /bin/sh -c "trove-manage datastore_update mysql ''"
trove
```

```
$ su -s /bin/sh -c "trove-manage datastore_update cassandra ''"
```

OpenStack: How do I install OpenStack's Trove into Bright?

trove

10. **Create the database image**, The image should have the trove-guestagent installed and configured to access the trove service on the controller. This is done by installing trove-guestagent, and configuring /etc/trove-guestagent.conf as follows:

```
rpc_backend = trove.common.rpc.impl_fake
qpidd_host = controller
nova_proxy_admin_user = admin
nova_proxy_admin_pass = 3de4922d8b6ac5a1aad9
nova_proxy_admin_tenant_name = service
trove_auth_url = http://master:35357/v2.0
```

11. **Notify and update the datastore that you have your image ready:**

```
trove-manage --config-file /etc/trove/trove.conf datastore_version_update mysql mysql-5.5
mysql 8cac48d8-2c29-4b2e-[Image-ID]9e86-c8efddc2db87 mysql-server-5.5 1
```

12. **Restart the Trove services:**

```
# cd /etc/init.d; for i in $(ls': cd /etc/init.d; for i in $(ls
openstack-trove*); do /etc/init.d/$i restart; done;
```

Note: At the time this was tried out (Oct 23 2014) there was a bug in Trove, in that Flavors with ids of the form 346346-24gre2-kh434 caused Trove to crash. Ids should be integers like 1, 2, 3. Like was said earlier, Trove is very much in development, so this may be fixed by now.

Your feedback is always welcome.

Unique solution ID: #1249

Author: Frank Furter

Last update: 2014-12-31 14:12