Stand Alone Database MGMTPOSTGRES_SLAVE Failover to MGMTPOSTGRES_MASTER

Contents

Introduction

This document describes how to recover the MGMTPOSTGRES_SLAVE when it does not form a cluster with the MGMTPOSTGRESS_MASTER.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Linux Interface
- Virtual Machine Environment
- postgresql
- Pacemaker/Corosync Configuration System (PCS)

Components Used

The information in this document is based on these software versions:

- CloudCenter version 4.8.1.1
- MGMTPOSTGRES_SLAVE Component
- MGMTPOSTGRES_MASTER Component

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

If there is a failure on both MGMTPOSTGRES components, the MGMTPOSTGRES_SLAVE no longer forms a cluster with the MGMTPOSTGRES_MASTER.

Problem

The MGMTPOSTGRES_SLAVE does not form a cluster with the MGMTPOSTGRES_MASTER. In order to get both MGMTPOSTGRES to form a cluster, MGMTPOSTGRES_SLAVE database needs to be deleted. Then, the database will be recovered from the MGMTPOSTGRES_MASTER.

Error Logs

```
[root@mgmtpostgres_master etc]# pcs status
Cluster name: cliqrdbcluster
Stack: corosync
Current DC: dbmaster (version 1.1.15-11.e174ec8) – partition with quorum

Last updated: Mon Nov 13 19:15:30 2017                Last changed: Mon Nov 13 16:59:51 2017 by
root via crm_attribute on db master

2 nodes and 3 resources configured
Online: [ dbmaster dbslave ]

Full list of resrouces:
  Resrouce Group: VIPGroup
            PGMasterVIP             (ocf::heartbeat:IPaddr2):                    Started
dbmaster

Master/Slave Set: mspostgresql [pgsql]
            Masters: [ dbmaster ]
            Stopped: [ dbslave ]

Failed Actions:
* pgsql_start_0 on dbslave ‘unknown error’ (1): call=11, status=Timed Out, exitreason=’none’,
            last-rc-change=’Mon Nov 13 18:15:25 2017’, queued-0ms, exec=60003ms

Daemon Status:
            corosyn: active/disabled
            pacemaker: active/enabled
            pcsd: inactive/disabled
```

Solution

Proceed to recover the MGMTPOSTGRES_SLAVE database in order for the MGMTPOSTGRES to form a cluster.

Step 1. In the MGMTPOSTGRES_MASTER, ensure that the cluster is stopped.

```
pcs cluster stop
pcs status
```

Step 2. In MGMTPOSTGRES_SLAVE, delete the existing database.

```
rm -rf /var/lib/pgsql/9.5/data/*
```

Step 3. In the MGMTPOSTGRES_MASTER, start the cluster again.

```
pcs cluster start
pcs status
```

Step 4. In MGMTPOSTGRES_SLAVE, recover the database from the MGMTPOSTGRES_MASTER.

```
/usr/pgsql-9.5/bin/pg_basebackup -h <MGMTPOSTGRES_MASTER-IP> -D /var/lib/pgsql/9.5/data/ -U
replication -v -P --xlog-method=stream
```

Step 5. In MGMTPOSTGRES_SLAVE, change the ownership of the recovered database.

```
chown postgres:postgres -R /var/lib/pgsql/9.5/data/*
```

Step 6. In MGMTPOSTGRES_SLAVE, start the cluster.

```
pcs cluster start
pcs cluster status
```

Step 7. In the MGMTPOSTGRES_MASTER, clean up the resources and check the cluster status.

```
pcs resource cleanup
pcs cluster status
```

Step 8. In the MGMTPOSTGRES_MASTER, verify that there is replication (look for the IP in the MGMTPOSTGRES_SLAVE IP).

```
ps –ef | grep postgr
```