Cassandra Backup/Store

Friday, October 24, 2014 10:04

Backups

- **The snapshot full backups in parallel**
  a. To perform the steps below on each node within a cluster (one time task):
     1) To create a shell script as /home/cassandra/daily_backup.sh
        # to clear incremental backup
        sudo rm -f /var/lib/cassandra/data/*/*/backups/*.*
        # to take a daily full backup @ each node staring on the same time
        /usr/bin/nodetool -h localhost -p 7199 snapshot -t `/bin/date +'%Y%m%d%H%M'` > /home/cassandra/backup_log.`/bin/date +'%Y%m%d%H%M'`
     2) To make this script executable
        sudo chown cassandra:cassandra /home/cassandra/daily_backup.sh
        sudo chmod 755 /home/cassandra/daily_backup.sh
     3) To create a cron job @ each node
        0 5 * * * /home/cassandra/./daily_backup.sh
  b. To check the result when needed
     1) To check the backup process log: /home/cassandra/backup_log.<timestamp>
     2) To check the backup hard links under ./snapshot: cd /var/lib/cassandra/data/<keyspace_name>/<table_name>/snapshot
     3) To check the Cassandra system log: /var/log/cassandra/system.log
  c. Note:
     1) The setting of start time (mm hh * * *) in the cron job at each node MUST be identical within a cluster. For the regional cluster the UTC is used; for the cluster at each CM the local time (5:00 am.) is used.
     2) The snapshot full backup will be done within 30 sec. per node.
     3) The DBA will receive a daily basis email that indicates whether the snapshot full backup on each node has been done successfully or not.

- **The snapshot incremental backups**
  a. To enable the incremental backup on each node (one time task):
     1) To modify the configuration file:
        sudo vi /etc/dse/cassandra/cassandra.yaml
        incremental_backups: true
     2) To restart cassandra instance
        sudo service dse restart
  b. To check the result when needed:
     1) To check the backup hard links under ./backups: cd /var/lib/cassandra/data/<keyspace_name>/<table_name>/backups
     2) To check the Cassandra system log: /var/log/cassandra/system.log
  c. Note:
     1) The incremental backup process is running all times automatically by Cassandra.
     2) The incremental backup data will be removed prior to a new full backup process is kicked off.
     3) The DBA will receive an alter email in case the incremental backup fails.

- **The rsync full backups in parallel**
  a. To create an initial copy of data files on each node (one time task):
     1) To create a destination directory
        sudo mkdir /data/backup
        sudo chown cassandra:Cassandra /data/backup
        sudo chmod 600 /data/backup
     2) To initialize the rsync process: rsync -avz /var/lib/cassandra/data/*/*/*.* /data/backup
  b. To perform the steps below on each node (one time task):
     1) To create a shell script as /home/cassandra/daily_rsync_backup.sh
        # sync up any DB changes @ daily basis
        rsync -avz /var/lib/cassandra/data/*/*/*.* /data/backup
     2) To make this script executable
        sudo chown cassandra:cassandra /home/cassandra/daily_rsync_backup.sh
        sudo chmod 755 /home/cassandra/daily_rsync_backup.sh
     3) To create a cron job @ each node
        0 22 * * * /home/cassandra/./daily_rsync_backup.sh
  c. To check the result when needed:
     1) To check the rsync backup process log
        /home/cassandra/rsync_backup_log.<timestamp>
  d. Note:
     1) The setting of start time (mm hh * * *) in the cron job at each node MUST be identical within a cluster. For the regional cluster the UTC is used; for the cluster at each CM the local time (22:00) is used.
     2) The initial rsync full backup (step a)) could take a few days per node.
     3) The ongoing daily rsync full backup (step b)) could take a few hours per node.
     4) The DBA will receive a daily basis email that indicates the result of the daily rsync full backups on each node.

Restore

- **Restore the entire cluster**
  a. To shut down c* instance at all nodes within the cluster: sudo service dse stop
  b. To execute the commands below at each nodes:
     1) To remove files under /var/lib/cassandra/commitlog
        sudo rm -f /var/lib/cassandra/commitlog/*.*
     2) To remove all *.db files
        sudo rm /var/lib/cassandra/data/*/*/*.*
     3) To copy rsync full backup files over
        sudo rsync -avz /data/backup/ /var/lib/cassandra/data/
     4) To start c* instance
        sudo service dse start
     5) To check the cluster status
        nodetool status
- **Point-in-time recovery one node**
  - To recovery a damaged node prior to putting it back to a cluster
     1) To shut down c* instance: sudo service dse stop
     2) To clear up:

- To remove files under /var/lib/cassandra/commitlog: sudo rm -f /var/lib/cassandra/commitlog/*.*
- To remove all *.db files: sudo rm /var/lib/cassandra/data/*/*/*.*
- To copy the local snapshot full dump files over:
  Loop on each keyspace/table:
  cd /var/lib/cassandra/data/<keyspace_name>/<table_name>/
  cp snapshots/<yyyy-mm-dd>/*.* .
- To copy the local incremental dump file over
  Loop on each keyspace/table:
  cp /var/lib/cassandra/data/<keyspace_name>/<table_name>/backups/*.* /var/lib/cassandra/data/<keyspace_name>/<table_name>/
- To start c* instance
  sudo service dse start
- To check the cluster status
  nodetool status