

ARTIGO: 11938

MySQL NDB Cluster

Configure three node as Manager + MySQL/API + Data Node

- Create tables on cluster:

Use `ENGINE=NDBCLUSTER`

- Pré configs

- Permite all connections between nodes
- Configure `/etc/hosts` to all nodes

- Configuration expected

All nodes must be same configurations. Only changes:

- Data node should connect on local manager: `connect-string=ndb2`
- MySQL/API node show connect on local manager: `ndb-connectstring=ndb2`

- Create folder

```
sudo mkdir /var/lib/mysql-cluster
```

- Install dependencies

```
sudo apt install libclass-methodmaker-perl
sudo apt install libaio1 libmecab2
```

- Install manager node

```
sudo dpkg -i mysql-cluster-community-management-server_8.0.21-1ubuntu18.04_amd64.deb
```

- Install data node

```
sudo dpkg -i mysql-cluster-community-data-node_8.0.21-1ubuntu18.04_amd64.deb
```

- Instal MySQL/API node

```
sudo dpkg -i mysql-common_8.0.21-1ubuntu18.04_amd64.deb
sudo dpkg -i mysql-cluster-community-client-core_8.0.21-1ubuntu18.04_amd64.deb
sudo dpkg -i mysql-cluster-community-client_8.0.21-1ubuntu18.04_amd64.deb
sudo dpkg -i mysql-client_8.0.21-1ubuntu18.04_amd64.deb
sudo dpkg -i mysql-cluster-community-server-core_8.0.21-1ubuntu18.04_amd64.deb
sudo dpkg -i mysql-cluster-community-server_8.0.21-1ubuntu18.04_amd64.deb # request mysql password
sudo dpkg -i mysql-server_8.0.21-1ubuntu18.04_amd64.deb
```

- Cluster configuration file

```
sudo vim /var/lib/mysql-cluster/config.ini
```

```
config.ini
```

```
# TCP PARAMETERS
```

```
[tcp default]
SendBufferMemory=2M
ReceiveBufferMemory=2M

# Increasing the sizes of these 2 buffers beyond the default values
# helps prevent bottlenecks due to slow disk I/O.

# MANAGEMENT NODE PARAMETERS

[ndb_mgmd default]
DataDir=/var/lib/mysql-cluster

# It is possible to use a different data directory for each management
# server, but for ease of administration it is preferable to be
# consistent.

[ndb_mgmd]
HostName=ndb1
# NodeId=management-server-A-nodeid

[ndb_mgmd]
HostName=ndb2
# NodeId=management-server-B-nodeid

[ndb_mgmd]
HostName=ndb3
# NodeId=management-server-B-nodeid

# Using 2 management servers helps guarantee that there is always an
# arbitrator in the event of network partitioning, and so is
# recommended for high availability. Each management server must be
# identified by a HostName. You may for the sake of convenience specify
# a NodeId for any management server, although one will be allocated
# for it automatically; if you do so, it must be in the range 1-255
# inclusive and must be unique among all IDs specified for cluster
# nodes.

# DATA NODE PARAMETERS

[ndbd default]
NoOfReplicas=2

# Using 2 replicas is recommended to guarantee availability of data;
# using only 1 replica does not provide any redundancy, which means
# that the failure of a single data node causes the entire cluster to
# shut down. As of NDB 8.0.19, it is also possible (but not required) to
# use more than 2 replicas, although 2 replicas are sufficient to provide
# high availability.

LockPagesInMainMemory=1

# On Linux and Solaris systems, setting this parameter locks data node
# processes into memory. Doing so prevents them from swapping to disk,
# which can severely degrade cluster performance.

DataMemory=2048M

# The value provided for DataMemory assumes 4 GB RAM
# per data node. However, for best results, you should first calculate
# the memory that would be used based on the data you actually plan to
# store (you may find the ndb_size.pl utility helpful in estimating
# this), then allow an extra 20% over the calculated values. Naturally,
# you should ensure that each data node host has at least as much
# physical memory as the sum of these two values.

# ODirect=1

# Enabling this parameter causes NDBCLUSTER to try using O_DIRECT
# writes for local checkpoints and redo logs; this can reduce load on
# CPUs. We recommend doing so when using NDB Cluster on systems running
# Linux kernel 2.6 or later.

NoOfFragmentLogFiles=300
DataDir=/usr/local/mysql/data
MaxNoOfConcurrentOperations=100000

SchedulerSpinTimer=400
SchedulerExecutionTimer=100
RealTimeScheduler=1
# Setting these parameters allows you to take advantage of real-time scheduling
# of NDB threads to achieve increased throughput when using ndbd. They
# are not needed when using ndbmtdd; in particular, you should not set
# RealTimeScheduler for ndbmtdd data nodes.
```

```

TimeBetweenGlobalCheckpoints=1000
TimeBetweenEpochs=200
RedoBuffer=32M

# CompressedLCP=1
# CompressedBackup=1
# Enabling CompressedLCP and CompressedBackup causes, respectively, local
#checkpoint files and backup files to be compressed, which can result in a space
#savings of up to 50% over noncompressed LCPs and backups.

# MaxNoOfLocalScans=64
MaxNoOfTables=10000
MaxNoOfAttributes=1000000
MaxNoOfOrderedIndexes=5000

[ndbd]
HostName=ndb2
# NodeId=data-node-A-nodeid

LockExecuteThreadToCPU=1
LockMaintThreadsToCPU=0
# On systems with multiple CPUs, these parameters can be used to lock NDBCLUSTER
# threads to specific CPUs

[ndbd]
HostName=ndb3
# NodeId=data-node-B-nodeid

LockExecuteThreadToCPU=1
LockMaintThreadsToCPU=0

# You must have an [ndbd] section for every data node in the cluster;
# each of these sections must include a HostName. Each section may
# optionally include a NodeId for convenience, but in most cases, it is
# sufficient to allow the cluster to allocate node IDs dynamically. If
# you do specify the node ID for a data node, it must be in the range 1
# to 144 inclusive and must be unique among all IDs specified for
# cluster nodes. (Previous to NDB 8.0.18, this range was 1 to 48 inclusive.)

# SQL NODE / API NODE PARAMETERS

[mysqlid]
# HostName=sql-node-A-hostname
# NodeId=sql-node-A-nodeid

[mysqlid]
[mysqlid]

# Each API or SQL node that connects to the cluster requires a [mysqlid]
# or [api] section of its own. Each such section defines a connection
# "slot"; you should have at least as many of these sections in the
# config.ini file as the total number of API nodes and SQL nodes that
# you wish to have connected to the cluster at any given time. There is
# no performance or other penalty for having extra slots available in
# case you find later that you want or need more API or SQL nodes to
# connect to the cluster at the same time.
# If no HostName is specified for a given [mysqlid] or [api] section,
# then any API or SQL node may use that slot to connect to the
# cluster. You may wish to use an explicit HostName for one connection slot
# to guarantee that an API or SQL node from that host can always
# connect to the cluster. If you wish to prevent API or SQL nodes from
# connecting from other than a desired host or hosts, then use a
# HostName for every [mysqlid] or [api] section in the config.ini file.
# You can if you wish define a node ID (NodeId parameter) for any API or
# SQL node, but this is not necessary; if you do so, it must be in the
# range 1 to 255 inclusive and must be unique among all IDs specified
# for cluster nodes.

```

- Cluster service (mgmd)

```
sudo vim /etc/systemd/system/ndb_mgmd.service
```

```
ndb_mgmd.service
```

```

[Unit]
Description=MySQL NDB Cluster Management Server
After=network.target auditd.service

[Service]

```

```
Type=forking
ExecStart=/usr/sbin/ndb_mgmd -f /var/lib/mysql-cluster/config.ini
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure

[Install]
WantedBy=multi-user.target
```

```
sudo systemctl daemon-reload
sudo systemctl enable ndb_mgmd
sudo systemctl stop/start/restart ndb_mgmd
```

- Data nodes configuration

```
sudo vim /etc/my.cnf
```

```
my.cnf
# provide connection string for management server host (default port: 1186)
[ndbd]
connect-string=ndb1,ndb2,ndb3
```

```
sudo systemctl daemon-reload
sudo systemctl enable ndbd
sudo systemctl stop/start/restart ndbd
```

- Data node service (ndbd)

```
sudo vim /etc/systemd/system/ndbd.service
```

```
ndbd.service
[Unit]
Description=MySQL NDB Data Node Daemon
After=network.target auditd.service

[Service]
Type=forking
ExecStart=/usr/sbin/ndbd
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure

[Install]
WantedBy=multi-user.target
```

- MySQL/API configuration

```
sudo vim /etc/mysql/mysql.conf.d/mysqld.cnf
```

```
mysqld.cnf
[mysqld]
pid-file           = /var/run/mysqld/mysqld.pid
socket             = /var/run/mysqld/mysqld.sock
datadir           = /var/lib/mysql
log-error          = /var/log/mysql/error.log
ndbcluster        # run NDB storage engine
ndb-connectstring=ndb1,ndb2,ndb3 ## connect to local mgr or other if need
```

```
sudo systemctl daemon-reload
sudo systemctl enable mysql
sudo systemctl stop/start/restart mysql
```

After all installed and configure, execute: (each configs update too)

Step 1 - Node 1, Node 2, Node 2

```
sudo ndb_mgmd -f /var/lib/mysql-cluster/config.ini --initial --configdir=/var/lib/mysql-cluster
```

Step 2 - Node 1, Node 2, Node 2

```
sudo ndbd
```

Step 2 - Node 1, Node 2, Node 2

```
sudo systemctl restart mysql
```

Verify cluster logs success:

```
-- Starting initial configuration change
2020-08-19 02:11:13 [MgmtSrvr] INFO -- Configuration 1 committed
2020-08-19 02:11:13 [MgmtSrvr] INFO -- Config change completed! New generation: 1
```

Important!!! Each update round should display this message

On this point, all data nodes, MySQLs and Managers should be connected:

Execute in Node 1, Node 2 and Node 3. Should be same result.

```
$ ndb_mgm

ndb_mgm> show
Cluster Configuration
-----
[ndbd(NDB)] 2 node(s)
id=4 @172.30.0.14 (mysql-8.0.21 ndb-8.0.21, Nodegroup: 0, *)
id=5 @172.30.0.183 (mysql-8.0.21 ndb-8.0.21, Nodegroup: 0)

[ndb_mgmd(MGM)] 3 node(s)
id=1 @172.30.0.38 (mysql-8.0.21 ndb-8.0.21)
id=2 @172.30.0.14 (mysql-8.0.21 ndb-8.0.21)
id=3 @172.30.0.183 (mysql-8.0.21 ndb-8.0.21)

[mysqld(API)] 5 node(s)
id=6 @172.30.0.14 (mysql-8.0.21 ndb-8.0.21)
id=7 @172.30.0.183 (mysql-8.0.21 ndb-8.0.21)
id=8 @172.30.0.38 (mysql-8.0.21 ndb-8.0.21)
id=9 (not connected, accepting connect from any host)
id=10 (not connected, accepting connect from any host)
```

Utilities:

Remove all mysql cluster

```
sudo apt-get remove --purge mysql-*
```

Kill services

```
sudo pkill -f ndbd
sudo pkill -f ndb_mgmd
```

Manager Script

```
#!/bin/bash
```

```

stop_mgm(){
    pkill -f ndb_mgmd
}

stop_ndb(){
    pkill -f ndbd
}

stop_mysql(){
    systemctl stop mysql
}

status(){
    ps aux | grep ndb
}

start_mgm(){
    ndb_mgmd -f /var/lib/mysql-cluster/config.ini --configdir=/var/lib/mysql-cluster
}

start_init_mgm(){
    ndb_mgmd -f /var/lib/mysql-cluster/config.ini --configdir=/var/lib/mysql-cluster --initial
}

start_ndb(){
    ndbd
}

start_mysql(){
    systemctl start mysql
}

restart_mysql(){
    systemctl restart mysql
}

case $1 in
    "restart")
        case $2 in
            "mgm")
                stop_mgm
                start_mgm
                ;;
            "data")
                stop_ndb
                start_ndb
                ;;
            "mysql")
                restart_mysql
                ;;
            *)
                echo "use [mgm |
data | mysql]"
                ;;
        esac
        ;;

    "stop")
        case $2 in
            "mgm")
                stop_mgm
                ;;
            "data")
                stop_ndb
                ;;
            "mysql")
                stop_mysql
                ;;
            *)
                echo "use [mgm |
data | mysql]"
                ;;
        esac
        ;;

    "start")
        case $2 in
            "mgm")

```

```

        start_mgm
        ;;
"mgm_init")
    start_init_mgm
    ;;
"data")
    start_ndb
    ;;
"mysql")
    start_mysql
    ;;
*)
    echo "use [mgm |
data | mysql]"
    ;;
    esac
    ;;
"logs")
    case $2 in
        "mgm")
            tail -n 100 -f /var/lib/mysql-cluster/*cluster.log
            ;;
        "mysql")
            tail -n 100 -f /var/log/mysql/error.log
            ;;
    *)
        echo "use [mgm |
mysql]"
        ;;
    esac
    ;;
"vim")
    case $2 in
        "mgm")
            vim /var/lib/mysql-cluster/config.ini
            ;;
        "data")
            vim /etc/my.cnf
            ;;
        "mysql")
            vim /etc/mysql/mysql.conf.d/mysqld.cnf
            ;;
    *)
        echo "use [mgm |
data | mysql]"
        ;;
    esac
    ;;
"ps")
    ps aux | grep ndb
    ;;
"bash")
    ndb_mgm
    ;;
*)
    echo "use [stop | start | restart | ps | bash]"
    ;;
esac

```

Links:

Configuration example:

<https://dev.mysql.com/doc/refman/5.7/en/mysql-cluster-config-starting.html>

Start MySQL Cluster

<https://www.digitalocean.com/community/tutorials/how-to-create-a-multi-node-mysql-cluster-on-ubuntu-18-04#step-3---configuring-and-starting-the-mysql-server-and-client>

Add new nodes on cluster

