Building MySQL Cluster in a Box Using Solaris 11 Zones

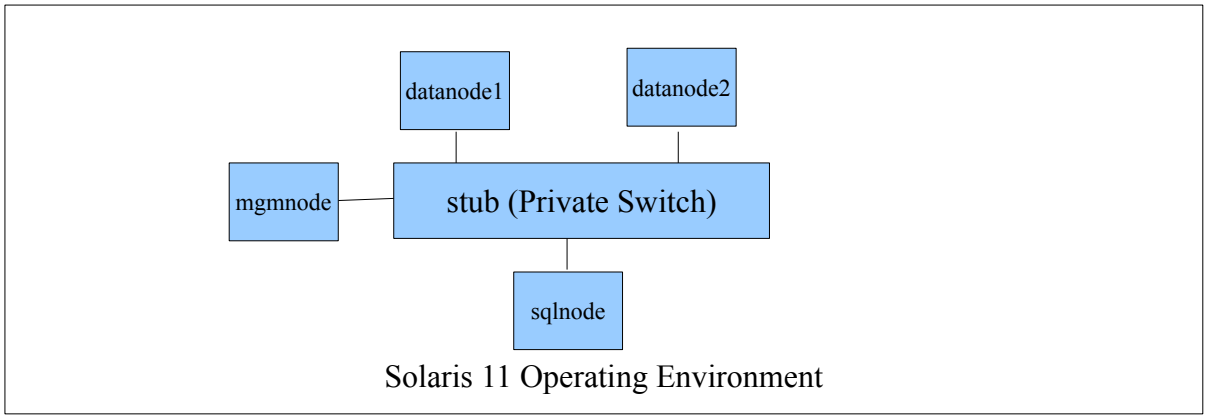

In the Solaris 11 Operating Environment, four zones are created namely:

- mgmnode
- datanode1
- datanode2
- sqlnode

All of the four zones are connected via a private switch (stub) in Solaris 11. stub has four virtual network interface card (VNIC) each of them assigned to individual zones

- mgmnode => vnic0
- datanode1 => vnic1
- datanode2 => vnic2
- sqlnode => vnic3

Each of the zone is assigned an IP address on the vnic so that all of them can communicate with each other:

- mgmnode => vnic0 => 192.168.0.100
- datanode1 => vnic1 => 192.168.0.101
- datanode2 => vnic2 => 192.168.0.102
- sqlnode => vnic3 => 192.168.0.103

Details on download location of Solaris 11 and the zone configuration are mentioned as an addendum to this note.

MySQL Cluster in a box – A Walk Through:

The following is to setup on individual zones to bring up the Management Node, the data nodes and the SQL nodes:

[On Management Node 'mgmnode']

1. Prepare a directory for MySQL installation
mkdir -p /usr/local
2. Copy the MySQL Cluster binary to the installation directory
cp /softwares/mysql-cluster-gpl-7.2.10-solaris11-x86_64.tar.gz /usr/local/
3. Navigate to the installation directory of MySQL, unzip the file and extract the tar
cd /usr/local/
gunzip mysql-cluster-gpl-7.2.10-solaris11-x86_64.tar.gz
tar xvf /sqlnode2/root/usr/local/mysql-cluster-gpl-7.2.10-solaris11-x86_64.tar.gz
4. Rename the installation directory to a convenient name preferably 'mysql'
mv mysql-cluster-gpl-7.2.10-solaris11-x86_64 mysql
5. Remove the binary to save space
rm /usr/local/mysql-cluster-gpl-7.2.10-solaris11-x86_64.tar.gz
6. Navigate to the mysql installation directory (if required) and initialize the 'mysql' database
cd usr/local/mysql
scripts/mysql_install_db
7. Change the ownership of the data directory of mysql installation
chown -R mysql:mysql /sqlnode2/root/usr/local/mysql

8. Copy the startup script from the default location to the standard location (/etc/init.d)
cp support-files/mysql.server /etc/init.d/mysql
9. Create the Data Directory for the nodes in the cluster
mkdir -p /var/lib/mysql-cluster
10. Create the config.ini file under the /etc directory
vi /etc/config.ini

```
[NDBD DEFAULT]
NoOfReplicas=2
DataDir=/var/lib/mysql-cluster

[NDBD_MGMD]
Hostname=192.168.0.100
DataDir=/var/lib/mysql-cluster

[NDBD]
Hostname=192.168.0.101
DataDir=/var/lib/mysql-cluster
[NDBD]
Hostname=192.168.0.102
DataDir=/var/lib/mysql-cluster

[MYSQLD]
Hostname=192.168.0.103
```

11. Start the management daemon
ndb_mgmd -f /etc/config.ini
12. Connect to the management console

```
ndb_mgm
ndb_mgm> show
ndb_mgm> exit
```

[On Data Node(s)]

1. Repeat step 1 to step 9 as mentioned above in the mgmnode configuration
2. Create the my.cnf file under /etc/ directory

```
vi my.cnf
[mysqld]
ndbcluster
ndb-connectstring=192.168.0.100

[mysql_cluster]
ndb-connectstring=192.168.0.100

[ndbd]
ndb-connectstring=192.168.0.100
```

3. Start the mysqld daemon (in case if the datanode also acts as the SQL node)

```
/etc/init.d/mysql start
```

4. Start the ndbd daemon

```
ndbd -n --initial
```

[On SQL Node]

Repeat the steps 1 - 3 from the Data Node configuration
Finally, connect to the mgmnode and start the cluster

```
ndb_mgm
ndb_mgm> ALL START
ndb_mgm> SHOW
```

Addendum:

Take your first steps with Solaris 11 here (download location, installation on VBOX etc.):
http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-112-s11-first-steps-524819.html

Also read, Getting Started with Solaris Zones:
http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-092-s11-zones-intro-524494.html

Configuring Private Switch (stub) in Solaris 11 and creating & configuring zones:

Create a private switch (stub0) / Create vnics on the stub:

1. Get back to the server by exiting out of the zone1 and confgure an etherstub:

```
root@mys11:~#  dladm create-etherstub stub0
root@mys11:~# dladm show-etherstub
LINK
stub0
root@mys11:~# dladm show-link
LINK     CLASS    MTU   STATE OVER
net0     phys     1500  up    --
---output omitted for brevity--
stub0    etherstub      9000  unknown  --
```

2. Create vnics on stub

```
root@mys11:~# dladm create-vnic -l stub0 vnic0
root@mys11:~# dladm create-vnic -l stub0 vnic1
root@mys11:~# dladm create-vnic -l stub0 vnic2
root@mys11:~# dladm create-vnic -l stub0 vnic3

root@mys11:~# dladm show-link
LINK     CLASS    MTU   STATE OVER
net0     phys     1500  up    --
---output omitted for brevity--
---output omitted for brevity--
stub0    etherstub      9000  unknown  --
```

```
vnic0    vnic           9000  up stub0
vnic1    vnic           9000  up stub0
vnic2    vnic           9000  up stub0
vnic3    vnic           9000  up stub0
```

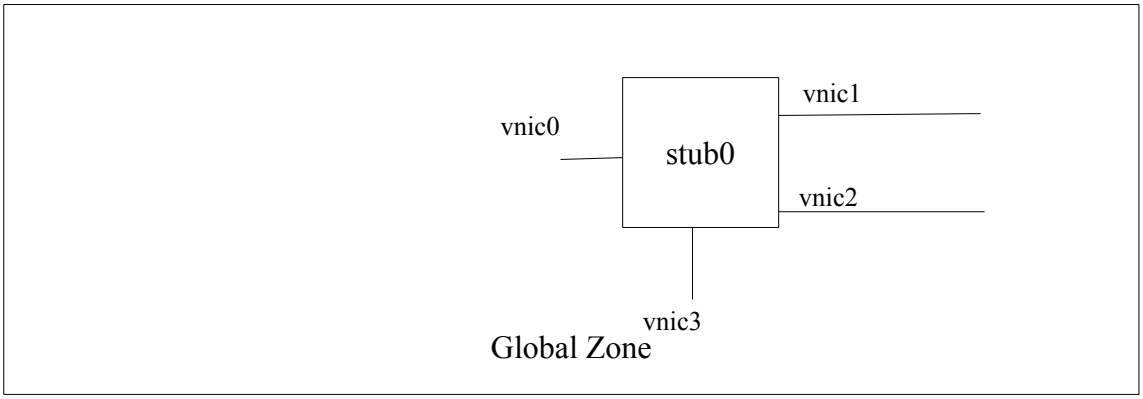

Configure mgmzone

1. Configure mgmzone and assigning the vnic0 to it:

```
root@mys11:~# zonecfg -z mgmgzone
Use 'create' to begin configuring a new zone.
zonecfg: zone1> create
create: Using system default template 'SYSdefault'
zonecfg:zone2> set zonepath=/mgmzone
zonecfg:zone2> add net
zonecfg:zone2:net> set physical=vnic0
zonecfg:zone2:net> end
zonecfg:zone2> exit
root@mys11:~# zoneadm list -cv
ID NAME  STATUS         PATH           BRAND       IP
0  global   running     /              solaris     shared
-  mgmzone  configured  /mgmzone       solaris     excl
```

2. Install the mgmzone

```
root@mys11:~# zoneadm -z mgmzone install
root@mys11:~# zoneadm list -cv
ID NAME     STATUS         PATH              BRAND       IP
```

```
0  global   running        /                solaris     shared
-  mgmzone  installed      /mgmzone         solaris     excl
root@mys11:~# zoneadm -z mgmzone boot
root@mys11:~# zoneadm list -cv
ID NAME  STATUS            PATH             BRAND       IP
0  global   running        /                solaris     shared
1  mgmzone  running        /mgmzone         solaris     excl
```

3. Login to mgmzone and complete the system identification:

```
root@mys11:~# zlogin -C mgmzone
```

Hit 'Enter' if it takes too much of time to show the configuration screens
Hit 'Esc' + 2 to navigate the screens and provide the following values:
Computer Name: **mgmzone**
Network Configuration: Manually
Network Card: **vnic0**
IP address: **192.168.0.100**
DNS: Do Not Configure DNS
Name Service: None
Region: Choose appropriately
Time Zone: Choose appropriately
Root password: oracle1
Let the system (mgmzone) reboot

Logout of zone console by using the keys ~. (Note: this may take you out of the Solaris 11 Server itself. Do SSH again to the Solaris 11 server if that happens)
In similar fashion install and configure the datanode1, datanode2 and sqlnode1 using the appropriate vnics and ip addresses.

Note: For the other nodes using the following zonepath:

datanode1 => /datanode1
datanode2 => /datanode2
datanode3 => /datanode3

Please note that the zonepaths that you use like 'mgmnode', 'datanode1', 'datanode2', 'sqlnode1' MUST be a ZFS file system in Solaris 11. Create ZFS filesystems using the following command(s). Please bear in mind the configuration below is a non-redundant one, for a serious setup, please consider using Mirror Pools / RAID Z pools.

```
# zpool create mcluster c1t1d1
```

The above command creates a pool by name 'mcluster' using the disk c1t1d1
Then run the following commands to create individual file systems on the pool

```
# zfs create mcluster/mgmpool
# zfs create set mountpoint=/mgmpool mcluster/mgmpool

# zfs create mcluster/datanode1
# zfs create set mountpoint=/datanode1 mcluster/datanode1

# zfs create mcluster/datanode2
# zfs create set mountpoint=/datanode2 mcluster/datanode2

# zfs create mcluster/sqlnode1
# zfs create set mountpoint=/sqlnode1 mcluster/sqlnode1
```

ZFS File System: http://docs.oracle.com/cd/E19253-01/819-5461/zfsover-2/index.html