Integrate DB2 into HA Env using Shared Storage on Linux

June 2019

Author: Haibo Wang
@DB Doctor

TABLE OF CONTENTS

Introduction

This document demonstrates how to prepare the shared storage in Linux VMs environment and then setup HA env using db2haicu. DB2 integrates Tivoli System Automation MP to achieve high availability.

The following 3 virtual machines are used in this setup.

VM Name	Role	IP Address	DB2 Version
yaya2	SAN Server	192.168.245.183	N/A
yaya	HA active server	192.168.245.129	V10.1 FP3
yaya1	HA passive server	192.168.245.184	V10.1 FP3

Setup

1. Setup SAN Server

The following settings are on yaya2 – SAN server.

(1) In VM Settings, add a Hard Disk device :

- Settings > Add… > Hard disk

Following other instructions to create a 20GB disk device.

(2) On yaya2, you should see your new disk with Yast > Partitioner.

- **"Format partition" to ext3 and "Do not mount partition"**

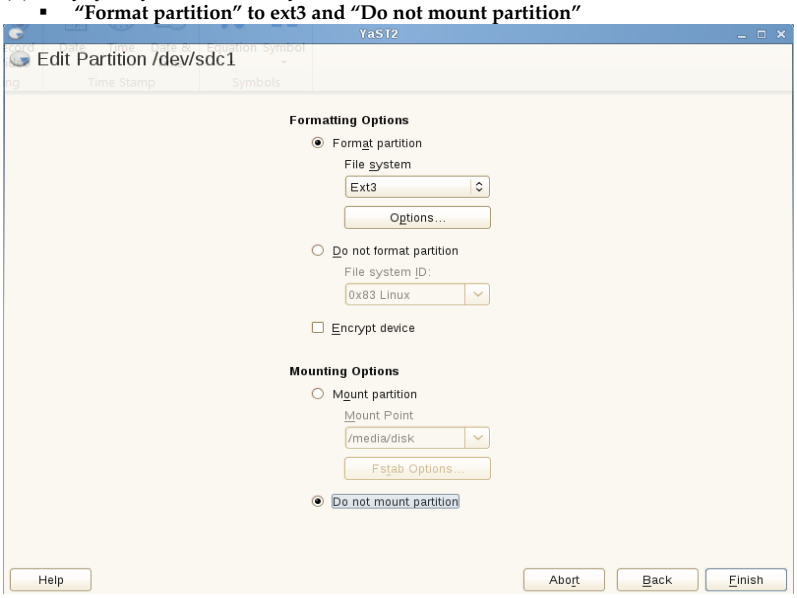

(3) Install iscsi if you can't find iSCSI-Target in YaST.

- YaST > Software > Software Management > search iscsi
- Install iSCSI Target and Open-iscsi

(4) Open iSCSI Target
- Choose “Service start when booting”
- Choose “No authentication” in Global tab.
- Add a target with the path you saw in Partitioner (ex. /dev/sdc1)

(5) iSCSI Target should restart after all settings finish. If it doesn’t restart automatically you can issue:
/etc/init.d/iscsitarget restart

2. Setup SAN Clients

(1) Install iscsi on yaya and yaya1 if iscsi hasn’t been installed yet.
(2) On both machines, Start iSCSI Initiator
- Choose “Service start when booting”
- Choose “No authentication” in Global tab.
- In Discovered Targets, click Discovery, put the IP address of SAN server into “IP Address” box and then follow other instructions to add the target

(3) iSCSI should restart after all settings finish. If it doesn’t restart automatically you can issue:
/etc/init.d/open-iscsi restart
(4) Verify on both machines (yaya and yaya1) with lsscsi

```
YAYA:/opt/ibm/db2/V10.1FP3/instance # lsscsi
[0:0:0:0]    disk    VMware,  VMware Virtual S 1.0   /dev/sda
[2:0:0:0]    cd/dvd  NECVMWar VMware IDE CDR10 1.00  /dev/sr0
[3:0:0:0]    disk    IET      VIRTUAL-DISK     0     /dev/sdc
[3:0:0:1]    disk    IET      VIRTUAL-DISK     0     /dev/sdd
[4:0:0:0]    disk    IET      VIRTUAL-DISK     0     /dev/sdb
```

In the above example, the device added is /dev/sdd
(5) On both machines, add the following line to /etc/fstab (#device mountpoint fstype options dump fsck) :
/dev/sdd /db2home ext3 noauto 0 0
(6) Mount the disk on both machines:
YAYA:~/Desktop # mkdir /db2home
YAYA:~/Desktop # mount /db2home

3. Create instance @ passive server

(1) on yaya1, install DB2 10.1FP3
(2) create user
/usr/sbin/useradd -m -d /db2home/db2sd -g 1000 -u 1006 -s /bin/ksh db2sd
YAYA1:~/Desktop # chown db2sd.db2igrp /db2home/db2sd
(3) create the instance
./db2icrt -s ese -u db2v10 db2sd
mv sqllib sqllib.yaya1

4. Create instance @ active server

(1) on yaya, install DB2 10.1FP3
(2) create user
/usr/sbin/useradd -m -d /db2home/db2sd -g 1000 -u 1006 -s /bin/ksh db2sd
YAYA:~/Desktop # chown db2sd.db2igrp /db2home/db2sd
(3) create the instance

./db2icrt -s ese -u db2v10 db2sd

5. Prepare TSA Cluster

(1) If any of the VMs are cloned, run
/usr/sbin/rsct/install/bin/recfgct
(2) Prepare the hosts for cluster

```
$ preprpnode yaya yaya1
```

Refer to Troubleshooting section if you have errors here.

(3) @BOTH add the following to /etc/profile
export CT_MANAGEMENT_SCOPE=2

6. Create HA env using db2haicu

(1) Start db2haicu

```
db2haicu is collecting information on your current setup. This step may take some time as db2haicu will need to
 activate all databases for the instance to discover all paths ...
When you use db2haicu to configure your clustered environment, you create cluster domains. For more information
, see the topic 'Creating a cluster domain with db2haicu' in the DB2 Information Center. db2haicu is searching
the current machine for an existing active cluster domain ...
db2haicu did not find a cluster domain on this machine. db2haicu will now query the system for information abou
t cluster nodes to create a new cluster domain ...

db2haicu did not find a cluster domain on this machine. To continue configuring your clustered environment for
high availability, you must create a cluster domain; otherwise, db2haicu will exit.
```

(2)Create domain

```
Create a domain and continue? [1]
1. Yes
2. No
1
Create a unique name for the new domain:
HA_YAYA
Nodes must now be added to the new domain.
How many cluster nodes will the domain 'HA_YAYA' contain?
2
Enter the host name of a machine to add to the domain:
yaya
Enter the host name of a machine to add to the domain:
yaya1
db2haicu can now create a new domain containing the 2 machines that you specified. If you choose not to create
a domain now, db2haicu will exit.

Create the domain now? [1]
1. Yes
2. No
1
Creating domain 'HA_YAYA' in the cluster ...
Creating domain 'HA_YAYA' in the cluster was successful.
You can now configure a quorum device for the domain. For more information, see the topic "Quorum devices" in t
he DB2 Information Center. If you do not configure a quorum device for the domain, then a human operator will h
ave to manually intervene if subsets of machines in the cluster lose connectivity.
```

(3)Create quorum

```
You can now configure a quorum device for the domain. For more information, see the topic "Quorum devices" in t
he DB2 Information Center. If you do not configure a quorum device for the domain, then a human operator will h
ave to manually intervene if subsets of machines in the cluster lose connectivity.

Configure a quorum device for the domain called 'HA_YAYA'? [1]
1. Yes
2. No
1
The following is a list of supported quorum device types:
  1. Network Quorum
Enter the number corresponding to the quorum device type to be used: [1]
1
Specify the network address of the quorum device:
192.168.245.183
Configuring quorum device for domain 'HA_YAYA' ...
Configuring quorum device for domain 'HA_YAYA' was successful.
```

(4)Create network

- Network card for node 1

```
The cluster manager found the following total number of network interface cards on the machines in the cluster
domain: '4'.  You can add a network to your cluster domain using the db2haicu utility.

Create networks for these network interface cards? [1]
1. Yes
2. No
1
Enter the name of the network for the network interface card: 'eth1' on cluster node: 'yaya'
1. Create a new public network for this network interface card.
2. Create a new private network for this network interface card.
Enter selection:
1
Are you sure you want to add the network interface card 'eth1' on cluster node 'yaya' to the network 'db2_publi
c_network_0'? [1]
1. Yes
2. No
1
Adding network interface card 'eth1' on cluster node 'yaya' to the network 'db2_public_network_0' ...
Adding network interface card 'eth1' on cluster node 'yaya' to the network 'db2_public_network_0' was successfu
l.
```

-Network card for node 2

```
Enter the name of the network for the network interface card: 'eth0' on cluster node: 'yaya1'
1. db2_public_network_0
2. Create a new public network for this network interface card.
3. Create a new private network for this network interface card.
Enter selection:
1
Are you sure you want to add the network interface card 'eth0' on cluster node 'yaya1' to the network 'db2_publ
ic_network_0'? [1]
1. Yes
2. No
1
Adding network interface card 'eth0' on cluster node 'yaya1' to the network 'db2_public_network_0' ...
Adding network interface card 'eth0' on cluster node 'yaya1' to the network 'db2_public_network_0' was successf
ul.
```

-a little problem here

```
Adding network interface card 'eth0' on cluster node 'yaya1' to the network 'db2_public_network_0' was successf
ul.
Enter the name of the network for the network interface card: 'eth0' on cluster node: 'yaya1'
1. db2_public_network_0
2. Create a new public network for this network interface card.
3. Create a new private network for this network interface card.
Enter selection:
1
Network adapter 'eth0' on node 'yaya1' is already defined in network 'db2_public_network_0' and cannot be added
 to another network until it is removed from its current network.
Enter the name of the network for the network interface card: 'eth1' on cluster node: 'yaya'
1. db2_public_network_0
2. Create a new public network for this network interface card.
3. Create a new private network for this network interface card.
Enter selection:
2
Are you sure you want to add the network interface card 'eth1' on cluster node 'yaya' to the network 'db2_publi
c_network_1'? [1]
1. Yes
2. No
2
```

(5)Cluster manager

```
Retrieving high availability configuration parameter for instance 'db2ha' ...
The cluster manager name configuration parameter (high availability configuration parameter) is not set. For mo
re information, see the topic "cluster_mgr - Cluster manager name configuration parameter" in the DB2 Informati
on Center. Do you want to set the high availability configuration parameter?
The following are valid settings for the high availability configuration parameter:
  1.TSA
  2.Vendor
Enter a value for the high availability configuration parameter: [1]
1
Setting a high availability configuration parameter for instance 'db2ha' to 'TSA'.
```

(6)Policy

```
Now you need to configure the failover policy for the instance db2ha. The failover policy determines the machin
es on which the cluster manager will restart the database manager if the database manager is brought offline un
expectedly.

The following are the available failover policies:
  1. Local Restart -- during failover, the database manager will restart in place on the local machine
  2. Round Robin -- during failover, the database manager will restart on any machine in the cluster domain
  3. Active/Passive -- during failover, the database manager will restart on a specific machine
  4. M+N -- during failover, the database partitions on one machine will failover to any other machine in the c
luster domain (used with DPF instances)
  5. Custom -- during failover, the database manager will restart on a machine from a user-specified list
Enter your selection:
3
You can identify mount points that are noncritical for failover. For more information, see the topic 'Identifyi
ng mount points that are noncritical for failover' in the DB2 Information Center.  Are there any mount points t
hat you want to designate as noncritical? [2]
1. Yes
2. No
1
Enter the full path of the mount to be made non-critical:
/tmp
Adding path '/tmp' to the non-critical path list ...
Adding path '/tmp' to the non-critical path list was successful.
Do you want to add more paths to the non-critical path list? [1]
1. Yes
2. No
2
Active/Passive failover policy was chosen. You need to specify the host names of an active/passive pair.
Enter the host name for the active cluster node:
yaya
Enter the host name for the passive cluster node:
yaya1
Adding DB2 database partition '0' to the cluster ...
Adding DB2 database partition '0' to the cluster was successful.
```

(7)Virtual IP

```
Do you want to configure a virtual IP address for the DB2 partition: '0'? [2]
1. Yes
2. No
1
Enter the virtual IP address:
192.168.245.200
Enter the subnet mask for the virtual IP address '192.168.245.200': [255.255.255.0]
255.255.255.0
Select the network for the virtual IP '192.168.245.200':
1. db2_public_network_0
Enter selection:
1
Adding virtual IP address '192.168.245.200' to the domain ...
Adding virtual IP address '192.168.245.200' to the domain was successful.
All cluster configurations have been completed successfully. db2haicu exiting ...
```

6. How the resource model looks like

```
YAYA:/opt/ibm/db2/V10.1FP3/instance # lssam
Online IBM.ResourceGroup:db2_db2sd_0-rg Nominal=Online
        |- Online IBM.Application:db2_db2sd_0-rs
                |- Offline IBM.Application:db2_db2sd_0-rs:yaya1
                '- Online IBM.Application:db2_db2sd_0-rs:yaya
        '- Online IBM.Application:db2mnt-db2home-rs
                |- Offline IBM.Application:db2mnt-db2home-rs:yaya1
                '- Online IBM.Application:db2mnt-db2home-rs:yaya
Online IBM.Equivalency:db2_db2sd_0-rg_group-equ
        |- Online IBM.PeerNode:yaya:yaya
        '- Online IBM.PeerNode:yaya1:yaya1
Online IBM.Equivalency:db2_public_network_0
        |- Online IBM.NetworkInterface:eth1:yaya
        '- Online IBM.NetworkInterface:eth0:yaya1
```

* The above resource model doesn’t include a Virtual IP.

7. Setup HA database

(1) You may create a shared storage path /db2data similarly to /db2home
(2) Create a db on shared storage path
db2 create db ha_db on /db2data
(3)Verify the resource model

```
$ lssam
Online IBM.ResourceGroup:db2_db2sd_0-rg Nominal=Online
        |- Online IBM.Application:db2_db2sd_0-rs
                |- Offline IBM.Application:db2_db2sd_0-rs:yaya1
                '- Online IBM.Application:db2_db2sd_0-rs:yaya
        |- Online IBM.Application:db2mnt-db2data-rs
                |- Offline IBM.Application:db2mnt-db2data-rs:yaya1
                '- Online IBM.Application:db2mnt-db2data-rs:yaya
        '- Online IBM.Application:db2mnt-db2home-rs
                |- Offline IBM.Application:db2mnt-db2home-rs:yaya1
                '- Online IBM.Application:db2mnt-db2home-rs:yaya
```

Commented [I1]: Instance resource group

Commented [I2]: Instance resource

Commented [I3]: db2data mount resource

Commented [I4]: db2home mount resource

8. Shared Storage HA database and HADR database Co-existence

(1) Setup HADR pair
@Both machines
/usr/sbin/useradd -d /home/db2hadr -g 1000 -u 1007 -s /bin/ksh db2hadr
mkdir /home/db2hadr
chown db2hadr.db2igrp /home/db2hadr
/opt/ibm/db2/V10.1FP3/instance # ./db2icrt -s ese -u db2v10 db2hadr
@Yaya
> db2 update db cfg for hadr_db using LOGARCHMETH1
"DISK:/home/db2hadr/archlogs"
> db2 backup db hadr_db to /dev/null
> db2 update db cfg for hadr_db using LOGINDEXBUILD ON
> db2 update db cfg for hadr_db using HADR_LOCAL_HOST yaya
> db2 update db cfg for hadr_db using HADR_LOCAL_SVC hadr_db2hadr
> db2 update db cfg for hadr_db using HADR_REMOTE_HOST yaya1
> db2 update db cfg for hadr_db using HADR_REMOTE_SVC hadr_db2hadr
> db2 update db cfg for hadr_db using HADR_REMOTE_INST db2hadr
> db2 backup db hadr_db
scp HADR_DB.0.db2hadr.DBPART000.20140701102501.001
root@yaya1:/home/db2hadr
@Yaya1
db2hadr@YAYA1:/home/db2hadr> db2 restore db hadr_db
> db2 update db cfg for hadr_db using HADR_LOCAL_HOST yaya1
> db2 update db cfg for hadr_db using HADR_REMOTE_HOST yaya
> db2 start hadr on db hadr_db as standby
@Yaya
> db2 start hadr on db hadr_db as primary
(2) Enable HADR_PEER_WINDOW on both Primary and Standby
db2 update db cfg for hadr_db using HADR_PEER_WINDOW 120

(3) Run db2haicu on Standby instance first
db2hadr@yaya1> db2haicu

```
Retrieving high availability configuration parameter for instance 'db2hadr' ...
The cluster manager name configuration parameter (high availability configuration parameter) is not set. For mo
re information, see the topic "cluster_mgr - Cluster manager name configuration parameter" in the DB2 Informati
on Center. Do you want to set the high availability configuration parameter?
The following are valid settings for the high availability configuration parameter:
  1.TSA
  2.Vendor
Enter a value for the high availability configuration parameter: [1]
1
Setting a high availability configuration parameter for instance 'db2hadr' to 'TSA'.
Adding DB2 database partition '0' to the cluster ...
Adding DB2 database partition '0' to the cluster was successful.
Do you want to validate and automate HADR failover for the HADR database 'HADR_DB'? [1]
1. Yes
2. No
1
Adding HADR database 'HADR_DB' to the domain ...
HADR database 'HADR_DB' has been determined to be valid for high availability. However, the database cannot be
added to the cluster from this node because db2haicu detected this node is the standby for HADR database 'HADR_
DB'. Run db2haicu on the primary for HADR database 'HADR_DB' to configure the database for automated failover.
All cluster configurations have been completed successfully. db2haicu exiting ...
```

(4) At this stage, standby instance is added to the resource model:

```
YAYA1:/opt/ibm/db2/V10.1FP3/instance # lssam
Online IBM.ResourceGroup:db2_db2hadr_yaya1_0-rg Nominal=Online
        '- Online IBM.Application:db2_db2hadr_yaya1_0-rs
                '- Online IBM.Application:db2_db2hadr_yaya1_0-rs:yaya1
Online IBM.ResourceGroup:db2_db2sd_0-rg Nominal=Online
        |- Online IBM.Application:db2_db2sd_0-rs
                |- Offline IBM.Application:db2_db2sd_0-rs:yaya1
                '- Online IBM.Application:db2_db2sd_0-rs:yaya
        |- Online IBM.Application:db2mnt-db2data-rs
                |- Offline IBM.Application:db2mnt-db2data-rs:yaya1
                '- Online IBM.Application:db2mnt-db2data-rs:yaya
        '- Online IBM.Application:db2mnt-db2home-rs
                |- Offline IBM.Application:db2mnt-db2home-rs:yaya1
                '- Online IBM.Application:db2mnt-db2home-rs:yaya
```

(5) Run db2haicu from Primary instance
db2hadr@yaya> db2haicu

```
Select an administrative task by number from the list below:
  1. Add or remove cluster nodes.
  2. Add or remove a network interface.
  3. Add or remove HADR databases.
  4. Add or remove an IP address.
  5. Move DB2 database partitions and HADR databases for scheduled maintenance.
  6. Create a new quorum device for the domain.
  7. Destroy the domain.
  8. Exit.
Enter your selection:
3
Do you want to add or remove another HADR database to or from the domain? [1]
1. Add
2. Remove
1
Do you want to validate and automate HADR failover for the HADR database 'HADR_DB'? [1]
1. Yes
2. No
1
Adding HADR database 'HADR_DB' to the domain ...
Adding HADR database 'HADR_DB' to the domain was successful.
Do you want to make any other changes to the cluster configuration? [1]
1. Yes
2. No
2
All cluster configurations have been completed successfully. db2haicu exiting ...
```

(6)Verify the resource model

```
$ lssam
Online IBM.ResourceGroup:db2_db2hadr_db2hadr_HADR_DB-rg Nominal=Online
      '- Online IBM.Application:db2_db2hadr_db2hadr_HADR_DB-rs
            |- Offline IBM.Application:db2_db2hadr_db2hadr_HADR_DB-rs:yaya1
            '- Online IBM.Application:db2_db2hadr_db2hadr_HADR_DB-rs:yaya
…
```

Commented [I5]: HADR_DB rsc group

Commented [I6]: HADR_DB rsc

HA Tests

1. Failover test

(1) run “rgreq” command below

```
YAYA:/opt/ibm/db2/V10.1FP3/instance # rgreq -o move db2_db2sd_0-rg
Action on resource group "db2_db2sd_0-rg" returned Token "0x79677cdc3a1446670f28b1532e3f0f00" .
```

(2) This may take a few minutes. Verify in lssam until the resource group is online again:

```
YAYA:/opt/ibm/db2/V10.1FP3/instance # lssam
Online IBM.ResourceGroup:db2_db2sd_0-rg Nominal=Online
        |- Online IBM.Application:db2_db2sd_0-rs
                |- Online IBM.Application:db2_db2sd_0-rs:yaya1
                '- Offline IBM.Application:db2_db2sd_0-rs:yaya
        |- Online IBM.Application:db2mnt-db2data-rs
                |- Online IBM.Application:db2mnt-db2data-rs:yaya1
                '- Offline IBM.Application:db2mnt-db2data-rs:yaya
        '- Online IBM.Application:db2mnt-db2home-rs
                |- Online IBM.Application:db2mnt-db2home-rs:yaya1
                '- Offline IBM.Application:db2mnt-db2home-rs:yaya
```

All resource moved from yaya(active server) to yaya1(passive server).

(3)Power off yaya1and after a while see the resource group is online on yaya

```
YAYA:/opt/ibm/db2/V10.1FP3/instance # lssam
Online IBM.ResourceGroup:db2_db2sd_0-rg Control=MemberInProblemState Nominal=Online
        |- Online IBM.Application:db2_db2sd_0-rs Control=MemberInProblemState
                |- Failed offline IBM.Application:db2_db2sd_0-rs:yaya1 Node=Offline
                '- Online IBM.Application:db2_db2sd_0-rs:yaya
        |- Online IBM.Application:db2mnt-db2data-rs Control=MemberInProblemState
                |- Failed offline IBM.Application:db2mnt-db2data-rs:yaya1 Node=Offline
                '- Online IBM.Application:db2mnt-db2data-rs:yaya
        '- Online IBM.Application:db2mnt-db2home-rs Control=MemberInProblemState
                |- Failed offline IBM.Application:db2mnt-db2home-rs:yaya1 Node=Offline
                '- Online IBM.Application:db2mnt-db2home-rs:yaya
```

2. Instance software failure test

(1) Kill the db2sysc process

```
db2sd@YAYA1:/db2home/db2sd> db2_kill
Killed
```

(2) db2 instance is restarted on the same server

Troubleshooting

1. preprpnode returns error

```
$ preprpnode host1 host2
```

/usr/sbin/rsct/bin/lsrsrc-api: 2612-022 A session could not be established with the RMC daemon on host2.
preprpnode: 2602-344 Unable to obtain the public key from host2.

Reason:
There is a firewall between the servers.
Action:
Turn off Firewall

2. Create instance failed

YAYA:/opt/ibm/db2/V10.1FP3/instance # ./db2icrt -u db2v10 db2ha
DBI1446I The db2icrt command is running.

DB2 installation is being initialized.

The user "db2ha" already has a sqllib directory in the home directory. This user cannot be used as the instance-owning user.

A major error occurred during the execution that caused this program to terminate prematurely. If the problem persists, contact your technical service representative.

For more information see the DB2 installation log at "/tmp/db2icrt.log.24615".
DBI1264E Errors were encountered in running db2icrt. Please
refer to the installation log file /tmp/db2icrt.log.24615 for more
information.

Explanation:

All processed and failed operations have been saved into this log file.
Reason:
Sqllib already exists.
Action:
Rename sqllib

3. Diagnostic information

(1) OS log
/var/adm/log/messages
(2) db2diag.log