

iSCSI and multipath on CentOS 7

????

- iSCSI Storage: NetApp FAS8200
- Linux OS: CentOS 7
- Server H/W: HP DL380 Gen10
- iSCSI ?????: 10G SFP ????? x 2

?????

???? nmtui ?????????? IP ? Mask

“ ??: ????????? Bonding?

?? iSCSI

?? iscsi ??

```
yum install iscsi-initiator-utils
```

?? iscsid.conf: `/etc/iscsi/iscsid.conf`

```
# ?? timeout
#node.session.timeo.replacement_timeout = 120
node.session.timeo.replacement_timeout = 5

# ?? node ?? Login
node.startup = automatic
```

?? initiator name: `/etc/iscsi/initiatorname.iscsi`

?????????

- iqn.<year-to-date>.<your-domain>.<your-hostname>:redhat-disk
?? InitiatorName=iqn.2019-02.mycompany.bdb7-a:redhat-disk
- initiator name ??? Storage ??????
- ??? Host ??????????

?? multipath ? netapp_linux_unified_host_utilities

```
yum install device-mapper-multipath
yum install libhbaapi
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64.rpm
```

“ ?????? libhbaapi??? sanlun ??????

?? iSCSI storage

??????

```
iscsiadm -m iface -I iscsi_ens2f0 --op=new
iscsiadm -m iface -I iscsi_ens3f0 --op=new
```

Discovery for the storage

```
iscsiadm -m discovery -t sendtargets -p <ip-to-iscsi-storage>:3260
```

“ ?? NetApp storage ????? iSCSI IPs???????????? IP ????????????????? IP?

```
[root@bdb7-b RPMs]# iscsiadm -m discovery -t sendtargets -p 10.18.104.32:3260
10.18.104.32:3260,1038 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.34:3260,1040 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.33:3260,1039 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.31:3260,1037 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.32:3260,1038 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.34:3260,1040 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.33:3260,1039 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
10.18.104.31:3260,1037 iqn.1992-08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16
```

Login nodes

```
iscsiadm -m node -l all
```

“ ???????????????

```
[root@bdb7-b RPMs]# iscsiadm -m node -l all
  Logging in to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.32,3260] (multiple)
  Logging in to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.32,3260] (multiple)
  Logging in to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.34,3260] (multiple)
  Logging in to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.34,3260] (multiple)
  Logging in to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.33,3260] (multiple)
  Logging in to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.33,3260] (multiple)
  Logging in to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.31,3260] (multiple)
  Logging in to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.31,3260] (multiple)
  Login to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.32,3260]
successful.
  Login to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.32,3260]
successful.
  Login to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.34,3260]
successful.
  Login to [iface: iscsi_ens3f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.34,3260]
successful.
  Login to [iface: iscsi_ens2f0, target: iqn.1992-
08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.33,3260]
successful.
  Login to [iface: iscsi_ens3f0, target: iqn.1992-
```

08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.33,3260]

successful.

Login to [iface: iscsi_ens2f0, target: iqn.1992-

08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.31,3260]

successful.

Login to [iface: iscsi_ens3f0, target: iqn.1992-

08.com.netapp:sn.58042162340f11e9892700a098e3ab45:vs.16, portal: 10.18.104.31,3260]

successful.

?? sanlun ?? LUN Disks

```
[root@bdb7-b RPMs]# sanlun lun show
      controller(7mode/E-Series)/
host          lun
      vserver(cDOT/FlashRay)      lun-pathname      filename      adapter
protocol  size  product
-----
-----
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdag      host19
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdaf      host20
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdae      host17
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdad      host18
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdac      host16
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdab      host15
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_3_vol/BDB7_B_3      /dev/sdz       host20
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_3_vol/BDB7_B_3      /dev/sdaa      host19
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_4_vol/BDB7_B_4      /dev/sdy       host14
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_3_vol/BDB7_B_3      /dev/sdx       host18
iSCSI      500.1g  cDOT
      TPEDBISCSISVM01      /vol/BDB7_B_3_vol/BDB7_B_3      /dev/sdw       host17
```

```

iSCSI      500.1g  cDOT
    TPEDBISCSISVM01          /vol/BDB7_B_3_vol/BDB7_B_3    /dev/sdt          host16
iSCSI      500.1g  cDOT
    TPEDBISCSISVM01          /vol/BDB7_B_3_vol/BDB7_B_3    /dev/sdu          host15
iSCSI      500.1g  cDOT
...

```

?? multipath

????????????? multipath.conf

```
mpathconf --enable --with_multipathd y
```

?? device alias name ? LUN WWID ???

```

“ ???? RedHat 7.6 ? multipathd ?????? user_friendly_names yes
   ??? NetApp ? LUN disks ?????? mpath X ?????????? WWID
   ??????????
   ?????????????????????? WWID ???

```

?? /etc/multipath.conf , ??????????????

```

...
multipaths {
    multipath {
        wwid      3600a0980383136312d3f4e387455564b
        alias     mpath1
    }
    multipath {
        wwid      3600a0980383136312d3f4e387455564c
        alias     mpath2
    }
    multipath {
        wwid      3600a0980383136312d3f4e387455564d
        alias     mpath3
    }
    multipath {
        wwid      3600a0980383136312d3f4e387455564e
        alias     mpath4
    }
}

```



```
## CentOS 7
systemctl enable remote-fs.target
```

```
## CentOS 4/5/6
chkconfig netfs on
```

?? MTU ??

????? MTU ? 1500????? iSCSI ?????????? 9000??? Linux ?????????? Switch ? iSCSI Storage (iSCSI Target) ??????????

RedHat 7.x) ?????????????? frame ????????

UNSUCCESSFUL

```
# ping -M do -c 4 -s 8196 10.1.1.21
PING 10.1.1.21 (10.1.1.21) 8185(8213) bytes of data.
PING 10.1.1.21(10.1.1.21) 8196(8224) bytes of data.
ping: local error: Message too long, mtu=1500
ping: local error: Message too long, mtu=1500
```

SUCCESSFUL

```
# ping -M do -c 4 -s 1408 10.1.1.21
PING 10.1.1.21 (10.1.1.21) 1408(1436) bytes of data.
1416 bytes from 10.1.1.21: icmp_seq=1 ttl=58 time=224 ms
```

????? MTU (NOTE: ???????)

```
ip link set dev ens2f0 mtu 9000
```

????????? iSCSI ?????????? ifcfg-ethXX?????????

```
MTU="9000"
```

?????????

????? LUN disk

1. iSCSI Storage ?? LUN disk ?????????? Linux host?

2. ?????? Linux host ??????

3. ??(discover)????? LUN disk

```
[root@bdb7-b ~]# /usr/bin/rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 3 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 3 0 0 0 ...
OLD: Host: scsi3 Channel: 00 Id: 00 Lun: 00
      Vendor: HPE      Model: DVDROM DUD0N      Rev: UMD1
      Type:   CD-ROM          ANSI SCSI revision: 05
Scanning host 4 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 5 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 6 for all SCSI target IDs, all LUNs
Scanning host 7 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 7 0 0 0 ...
```

4. ?? sanlun ?????? LUN disk

```
[root@bdb7-b ~]# sanlun lun show all
controller(7mode/E-Series)/                device
host                lun
vserver(cDOT/FlashRay)    lun-pathname    filename    adapter
protocol  size  product
-----
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdaw    host16
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdav    host15
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdau    host14
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdat    host13
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdas    host12
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdar    host11
```

```

iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdaq      host10
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_6_vol/BDB7_B_6    /dev/sdap      host9
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_5_vol/BDB7_B_5    /dev/sdao      host15
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_5_vol/BDB7_B_5    /dev/sdan      host16
iSCSI      500.1g  cDOT
TPEDBISCSISVM01          /vol/BDB7_B_4_vol/BDB7_B_4    /dev/sdam      host16
iSCSI      500.1g  cDOT
...

```

5. ?? multipath ????? LUN disk

```

[root@bdb7-b alang]# multipath -ll
mpath2 (3600a098038313631305d4e2d6f774744) dm-6 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 9:0:0:1 sdf 8:80 active ready running
| |- 12:0:0:1 sdt 65:48 active ready running
| |- 16:0:0:1 sdah 66:16 active ready running
| `-- 15:0:0:1 sdai 66:32 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
    |- 11:0:0:1 sdg 8:96 active ready running
    |- 10:0:0:1 sde 8:64 active ready running
    |- 14:0:0:1 sdu 65:64 active ready running
    `-- 13:0:0:1 sdv 65:80 active ready running
mpath1 (3600a098038313631305d4e2d6f774743) dm-5 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 9:0:0:0 sdc 8:32 active ready running
| |- 12:0:0:0 sdr 65:16 active ready running
| |- 16:0:0:0 sdaf 65:240 active ready running
| `-- 15:0:0:0 sdag 66:0 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
    |- 11:0:0:0 sdd 8:48 active ready running
    |- 10:0:0:0 sdb 8:16 active ready running

```

```

|- 13:0:0:0 sdq 65:0 active ready running
`- 14:0:0:0 sds 65:32 active ready running
3600a098038313631305d4e2d6f77474a dm-26 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
| |- 9:0:0:5 sdap 66:144 active ready running
| |- 12:0:0:5 sdas 66:192 active ready running
| |- 15:0:0:5 sdav 66:240 active ready running
| ` - 16:0:0:5 sdaw 67:0 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 10:0:0:5 sdaq 66:160 active ready running
|- 11:0:0:5 sdar 66:176 active ready running
|- 13:0:0:5 sdat 66:208 active ready running
`- 14:0:0:5 sdau 66:224 active ready running
...

```

6. ?? LUN WWID ? mpathX ???

?? `/etc/multipath.conf`: ??? wwid ? alias ???

```

# added by alang
multipaths {
    multipath {
        wwid    3600a098038313631305d4e2d6f774743
        alias   mpath1
    }
    multipath {
        wwid    3600a098038313631305d4e2d6f774744
        alias   mpath2
    }
    multipath {
        wwid    3600a098038313631305d4e2d6f774745
        alias   mpath3
    }
    multipath {
        wwid    3600a098038313631305d4e2d6f774746
        alias   mpath4
    }
    multipath {

```

```

        wwid    3600a098038313631305d4e2d6f774747
        alias   mpath5
    }
    multipath {
        wwid    3600a098038313631305d4e2d6f77474a
        alias   mpath6
    }
}

```

7. ?? multipath ??

```

[root@bdb7-b alang]# systemctl reload multipathd

[root@bdb7-b alang]# multipath -ll
mpath2 (3600a098038313631305d4e2d6f774744) dm-6 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 12:0:0:1 sdt  65:48  active ready running
|  |- 15:0:0:1 sdai 66:32  active ready running
|  |- 16:0:0:1 sdah 66:16  active ready running
|  `-- 9:0:0:1 sdf  8:80   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
    |- 10:0:0:1 sde  8:64   active ready running
    |- 11:0:0:1 sdg  8:96   active ready running
    |- 13:0:0:1 sdv  65:80  active ready running
    `-- 14:0:0:1 sdu  65:64  active ready running
mpath1 (3600a098038313631305d4e2d6f774743) dm-5 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 12:0:0:0 sdr  65:16  active ready running
|  |- 15:0:0:0 sdag 66:0   active ready running
|  |- 16:0:0:0 sdaf 65:240 active ready running
|  `-- 9:0:0:0 sdc  8:32   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
    |- 10:0:0:0 sdb  8:16   active ready running
    |- 11:0:0:0 sdd  8:48   active ready running
    |- 13:0:0:0 sdq  65:0   active ready running
    `-- 14:0:0:0 sds  65:32  active ready running

```

```

mpath6 (3600a098038313631305d4e2d6f77474a) dm-26 NETAPP ,LUN C-Mode
size=500G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 12:0:0:5 sdas 66:192 active ready running
| |- 15:0:0:5 sdav 66:240 active ready running
| |- 16:0:0:5 sdaw 67:0 active ready running
| ` - 9:0:0:5 sdap 66:144 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:0:5 sdaq 66:160 active ready running
  |- 11:0:0:5 sdar 66:176 active ready running
  |- 13:0:0:5 sdat 66:208 active ready running
  ` - 14:0:0:5 sdau 66:224 active ready running

```

?????? disk mpath6 ??? LVM ?????

????

?? disk path ????

- ??????? Storage LUN ????
- LUN ??? Size
- Storage ? path ????

```
[root@bdb7-a ~]# sanlun lun show -p
```

```
ONTAP Path: TPEDBISCSISVM01:/vol/BDB7_A_2_vol/BDB7_A_2
```

```
LUN: 1
```

```
LUN Size: 500.lg
```

```
Product: cDOT
```

```
Host Device: mpath2(3600a0980383136312d3f4e387455564c)
```

```
Multipath Policy: service-time 0
```

```
Multipath Provider: Native
```

```

-----
host      vserver
path      path      /dev/     host      vserver
state     type      node     adapter   LIF
-----
up        primary   sdo      host12    TPEDBISCSISVM01_iSCSI2
up        primary   sds      host13    TPEDBISCSISVM01_iSCSI2
up        primary   sdaa     host15    TPEDBISCSISVM01_iSCSI1

```

Revision #7

Created 2021-08-29 05:22:44 CST by A-Lang (Admin)

Updated 2022-01-04 06:09:27 CST by A-Lang (Admin)