

SAN & iSCSI

SAN - Storage Area Network, iSCSI - Internet Small Computer System Interface

- [iSCSI and multipath on CentOS 7](#)
- [Learning](#)
- [Multipath on RedHat \(NetApp FAS2240\)](#)
- [FAQ](#)

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-
```



```

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)/
host          lun          device
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

```


Learning

iSCSI

- [How to configure iSCSI target & initiator on RHEL/CentOS 7.6](#)

Multipath on RedHat (NetApp FAS2240)

Multipath ? CentOS/RedHat ?????????? port ? HBA ??? Storage????????????????

????

- RedHat 4.7
- NetApp FAS2240
- SAN Switch x 2
- HBA Card: QLogic, port x 2

```
#>lspci
```

```
Fibre Channel: QLogic Corp. ISP2532-based 8Gb Fibre Channel to PCI Express HBA (rev 02)
```

RedHat ? HBA Card ??

???????????? RedHat ???? HBA ?????? RedHat ?????? HBA ??

????????? RedHat ???? /etc/modprobe.conf????????

```
alias scsi_hostadapter2 qla2500
```

```
“ scsi_hostadapterX ??????????
```

?? HBA ?? WWN

```
#> cat /proc/scsi/qla2xxx/2
```

```
SCSI Device Information:
```

```
scsi-qla0-adapter-node=5001438004253c81;
```

```
scsi-qla0-adapter-port=5001438004253c80; <---- WWN
```

```
scsi-qla0-target-0=50001fe1500e28fd;
```

```
scsi-qla0-target-1=50001fe1500e28f9;
```

```
scsi-qla0-target-2=500a09829140a66c;
```

“ ??? HBA port ??? /proc/scsi/qla2xxx/3?????

?? port ??? WWN

NetApp storage ?? LUN ?? WWN ?? Host Mapping ?????

????????

```
# lsmod | grep qla
qla2500                209985  0
qla2xxx                203681  13 qla2500
scsi_transport_fc     19521   1 qla2xxx
scsi_mod              153489  7
ib_srp,qla2xxx,scsi_transport_fc,sr_mod,libata,cciss,sd_mod
```

SAN Storage ?????

1. ? NetApp storage?SAN switch?RedHat ????????
2. ?? RedHat ? HBA ? ? port ??? WWN
3. ?? NetApp?? LUNs ????? WWN Hosts

?? Multipath

“ ???RedHat ?????????? multipath.conf ??????????????????????
<https://access.redhat.com/labs/multipathhelper/>

????????

```
rpm -q device-mapper
rpm -q device-mapper-multipath
yum install device-mapper
yum install device-mapper-multipath
```

CentOS 5.x)

?? /etc/multipath.conf

```
# Blacklist all devices by default. Remove this to enable multipathing
# on the default devices.
devnode_blacklist {
#     devnode "*" << 
}

## Use user friendly names, instead of using WWIDs as names.
defaults {
    user_friendly_names yes
}
```

“ ???????? multipathd ?????????????? LUNs ??????????????”

?? multipathd ??

```
#> service multipathd start
#> chkconfig multipathd on
```

CentOS 6/7)

???????? /etc/multipath.conf???????????????????????? multipathd ??????????

```
#> mpathconf --enable --with_multipathd y
```

“ ???? multipath.conf

?? multipah ???

CentOS 5.x

```
[root@cdb3-b ~]# multipath -ll
mpath1 (360a98000383035537824474d3936556a) dm-4 NETAPP,LUN
[size=700G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=4][active]
\_ 1:0:1:0 sdb      8:16 [active][ready]
\_ 2:0:1:0 sdd      8:48 [active][ready]
```

```
\_ round-robin 0 [prio=1][enabled]
\_ 1:0:0:0 sda      8:0  [active][ready]
\_ 2:0:0:0 sdc      8:32 [active][ready]
```

CentOS 6.x

```
[root@cdb3 ~]# multipath -ll
mpathb (360a98000383035537824474d39365568) dm-2 NETAPP,LUN
size=700G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwhandler='0' wp=rw
|-+- policy='round-robin 0' prio=4 status=active
|  |- 1:0:1:0 sdc 8:32 active ready running
|  `-- 2:0:1:0 sde 8:64 active ready running
`--+- policy='round-robin 0' prio=1 status=enabled
    |- 1:0:0:0 sdb 8:16 active ready running
    `-- 2:0:0:0 sdd 8:48 active ready running
```

?? NetApp ?????

?????? LUN ????? - NetApp Linux Host Utilities

```
yum install libnl
rpm -ivh QConvergeConsoleCLI-1.0.01-32.x86_64.rpm
rpm -ivh netapp_linux_host_utilities-6-2.x86_64.rpm
```

“ - ?????????????????? sanlun ??

- ?? CentOS 5.x/6.x

?? LUN ????????

???????????? LUN disks (? dual path SAN ??????? LUN disks)???????????????? mpathXX
???????

```
#> multipath -l
mpath1 (360a98000424732615124444150534a35)
[size=100 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [active]
\_ 2:0:2:0 sde 8:64 [active]
```

```
\_ 3:0:2:0 sdf 8:80 [active]
```

```
mpath0 (3600508b400068e5c00012000001a0000)
```

```
[size=100 GB][features="1 queue_if_no_path"][hwhandler="0"]
```

```
\_ round-robin 0 [active]
```

```
\_ 2:0:0:1 sda 8:0 [active]
```

```
\_ 3:0:0:1 sdc 8:32 [active]
```

```
\_ round-robin 0 [enabled]
```

```
\_ 2:0:1:1 sdb 8:16 [active]
```

```
\_ 3:0:1:1 sdd 8:48 [active]
```

“ mpathX ?????? LVM ? PV ??????? filesystem?????? /dev/mapper/mpathX

2:0:0:1 ????? host : channel : ID : LUN

Host 2 ? Host 3 ?? HBA ????? port ??? port ?????

??/? LUN ????????

?????

- <https://library.netapp.com/ecmdocs/ECMP1654943/html/GUID-F88A86D2-F3A9-4440-8C32-539BFEA6BED1.html>
- <https://access.redhat.com/solutions/3941>

CentOS/RedHat 5.x ??

? CentOS/RedHat 5 ????? SLES 10 ????????????????

```
yum install sg3_utils
```

????????? LUN Disk

```
□□□□□□ LUN Disks
```

```
#> sanlun lun show all
```

```
□□□□ LUN
```

```
#> scsi-rescan
```

```
NOTE: □□□□□□ LUN□□□□□ -a
```

```
#> scsi-rescan -a
```

```
□□□□□□ LUN Disks
```

```
#> sanlun lun show all
```

?? LUN ???

?????? LUN disk ????????????

1. ?? `umount` ??
2. ?? `lvremove` ?? LV
3. ?????? `vgremove` ?? VG
4. ?? `pvremove` ?? PV
5. ?? `multipath -f mpathXX` ?? multipathXX
NOTE: ?????????????? LUN ?????? Hang?
6. ? Storage ?????????? LUN disk
NOTE: ? LUN ???????? `sanlun lun show all` ?????? LUN ??? unknown?
7. ?? `scsi-rescan -r` ?????????? LUN disk
NOTE: ??????????????????
4 device(s) removed.
[5:0:1:0]
[5:0:3:0]
[6:0:1:0]
[6:0:3:0]
????? `sanlun lun show all` ??? unknown ??????

CentOS/RedHat 4.x?

- [Vmware Linux Guest Add a New Hard Disk Without Rebooting Guest](#)

????????? script

1. remove-LUNs.sh
2. rescan-LUNs.sh

?? LUN disk ??

1. ? NetApp storage ??? LUN Mapping ???
2. ? RedHat ??? `rescan-LUNs.sh`
3. ? RedHat ??? `multipath -l`

???? LUN disk

1. ? RedHat ?????? mpathX

2. ? RedHat ??? `sanlun lun show all` ?? LUN ??? NetApp ???
 ????? sanlun ?????? NetApp ??? `netapp_linux_host_utilities-6-2.x86_64.rpm`
3. ? NetApp ?????? LUN
4. ? RedHat ??? `multipath -l`
5. ????? LUN disks ? `host:channel:id:LUN`
6. ?? `remove-LUNs.sh <host> <channel> <id> <LUN>`

?? LUN disk ????

1. ?? LUN ID ? `sd[X]` ???

```
smartctl -a /dev/sd[X]
```

2. ?? LUN disk???????

```
echo 1 > /sys/block/sd[X]/device/rescan
```

```
lsblk | grep sd[X]
```

3. Resize PV

```
pvresize /dev/sd[X]
```

4. ?????? SOP ?? LV ? Filesystem ????

LVM ??

- ?? LUN disk ??? PV

```
#> pvcreate /dev/mapper/mpath6
```

“ PV ?????? `pvdisplay` ?????? `dm-X?`

- ???? VG (??? PV)

```
#> vgcreate dbVG /dev/mapper/mpath6
```

??????

multipath

| | |
|--|--|
| <code>/etc/init.d/multipathd start</code> | Manually starts the multipathing service on the Linux host. |
| <code>multipath -ll</code> | Displays the multipathing configuration. |
| <code>multipath -v2</code> | Creates the initial multipathing configuration. |
| <code>multipath -f <i>device_map</i></code> | Removes an existing LUN from the multipathing configuration. |
| <code>multipath -ll -v2 -d</code> | Displays the multipathing configuration in verbose mode. |
| <code>multipath -t ? multipathd show config</code> | ?????????? |

Accessing LUNs on host

| | |
|--|--|
| <code>sanlun fcp show adapter -v</code> | Displays HBA information. |
| <code>sanlun fcp show adapter</code> | Display WWPN of HBA cards |
| <code>sanlun lun show all</code> | Lists available LUNs as seen by the sanlun script. |
| <code>sanlun lun show -p</code> | Provides multipath information for both DM-Multipath and Veritas Storage Foundation TM dynamic multipathing solutions. |
| <code>ls -lh /dev/mapper/</code> | Lists DM-Multipath devices generated for LUNs. |
| <code>e2label <i>device label</i></code> | Creates a persistent label for a file system. <i>device</i> is the partition; <i>label</i> is a unique string up to 16 characters. Not needed with DM-Multipath. |

Q & A

Q: RHEL 7 ????? alias path

“ ?? NetApp Storage??? multipath -ll

```

360a98000443437576f2444726e556d76 dm-2 NETAPP ,LUN
size=800G 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
|`- 6:0:0:0 sdb 8:16 active ready running
`+- policy='service-time 0' prio=10 status=enabled
  `- 6:0:3:0 sdc 8:32 active ready running

```

???? alias name: mpathXX

???????https://access.redhat.com/solutions/2061463

?? multipath ????????

```
multipath -t
```

“???? NetApp ????”

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

“NOTE: ??? NetApp Storage?”

```
defaults {
    user_friendly_names yes
    find_multipaths yes
}

devices {
    device {
        vendor "NETAPP"
        product "LUN.*"
        path_grouping_policy "group_by_prio"
        path_checker "tur"
        features "3 queue_if_no_path pg_init_retries 50"
        hardware_handler "0"
        prio "ontap"
        failback immediate
        rr_weight "uniform"
        rr_min_io 128
        flush_on_last_del "yes"
        dev_loss_tmo "infinity"
        user_friendly_names yes
        retain_attached_hw_handler yes
        detect_prio yes
    }
}
```

?????

```
#> systemctl reload multipathd

#> multipath -ll
mpatha (360a98000443437576f2444726e556d76) dm-2 NETAPP ,LUN
size=800G 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
|  |- 4:0:0:0 sdb 8:16 active ready running
|   `-- 6:0:0:0 sdd 8:48 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
    |- 4:0:3:0 sdc 8:32 active ready running
    `-- 6:0:3:0 sde 8:64 active ready running
```

Q: ?? netapp_linux_host_utilities ??????????

“ Warning: libnl.so library not found, some sanlun commands may not work. Refer Linux Host Utilities Installation and Setup Guide for more details
Warning: libHBAAPI.so library not found, some sanlun commands may not work. Refer Linux Host Utilities Installation and Setup Guide for more details


Solution: ??? libnl

```
yum install libnl
```

???????????

- http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/ResourceByOS.aspx?productid=1043&oemid=372&oemcatid=111969
- http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByOs.aspx?ProductCategory=39&OsCategory=2&Os=65&OsCategoryName=Linux&ProductCategoryName=Fibre+Channel+HBAs&OSName=Linux+Red+Hat+%2864-bit%29
- QConvergeConsole CLI (x64)
- SANsurfer CLI (x86/x64)

Q: ?? pvscan?vgscan ??????????????



3. ?????????????????????????????

???????????????????????????????? grub.conf ? default=1?

FAQ

???? /dev/sda ? Local Disk ?? SAN Disk

Solution:

```
#> lsscsi
[2:0:0:0]    cd/dvd  hp          DVD-ROM DTA0N    WBD0    /dev/sr0
            state=running queue_depth=1 scsi_level=6 type=5 device_blocked=0 timeout=0
[3:0:0:0]    tape    IBM        ULT3580-HH5    D2AD    /dev/st0
            state=running queue_depth=32 scsi_level=7 type=1 device_blocked=0 timeout=900
[3:0:1:0]    tape    IBM        ULT3580-HH5    D2AD    /dev/st1
            state=running queue_depth=32 scsi_level=7 type=1 device_blocked=0 timeout=900
[3:0:1:1]    mediumx IBM        3573-TL        C.20    -
            state=running queue_depth=32 scsi_level=6 type=8 device_blocked=0 timeout=0
[3:0:2:0]    disk    IBM        1726-4xx  FAStT  0617    /dev/sda          <=====
            state=running queue_depth=32 scsi_level=6 type=0 device_blocked=0 timeout=60
[3:0:2:31]   disk    IBM        Universal Xport 0617    -
            state=running queue_depth=32 scsi_level=6 type=0 device_blocked=0 timeout=60
```

“ IBM FAStT ? IBM Storage ?????

???? lsscsi ?????? yum install lsscsi

???? HBA Card ? port WWN

???: ? systool

```
# Install
yum install sysfsutils
#
systool -c fc_host -v
systool -c fc_transport -v
```

???? /sys/class/fc_host/{host#}/port_name

```
#> cat /sys/class/fc_host/host5/port_name
0x2100001b32837bcd
```

???? /proc/scsi/{adapter-type}/{host#}

```
#> cat /proc/scsi/qla2xxx/2

SCSI Device Information:
scsi-qla0-adapter-node=5001438004253c81;
scsi-qla0-adapter-port=5001438004253c80; <---- WWN
scsi-qla0-target-0=50001fe1500e28fd;
scsi-qla0-target-1=50001fe1500e28f9;
scsi-qla0-target-2=500a09829140a66c;
```

mpathX v.s. dm-X

????? multipath??? LUN ?????? /dev/mapper/mpathX??????? pvcreate ? LUN ??? PV ?????
pvdisplay -C ???????????? /dev/dm-X??????????????????? mpathX ????? dm-X?

???? dmsetup

```
[root@bdb3 ~]# dmsetup ls
mpath1 (253, 4)
rootVG-varLV (253, 1)
rootVG-swapLV (253, 3)
rootVG-rootLV (253, 0)
rootVG-worktmpLV (253, 2)
```

“ mpath1 (253,4)?4 ?????? /dev/dm-4

???? multipath

```
[root@cdb3-b ~]# multipath -ll
mpath1 (360a98000383035537824474d3936556a) dm-4 NETAPP,LUN <===== []
[size=700G][features=3 queue_if_no_path pg_init_retries 50][hwhandler=0][rw]
\_ round-robin 0 [prio=4][active]
\_ 1:0:1:0 sdb 8:16 [active][ready]
\_ 2:0:1:0 sdd 8:48 [active][ready]
```

```
\_ round-robin 0 [prio=1][enabled]
\_ 1:0:0:0 sda      8:0  [active][ready]
\_ 2:0:0:0 sdc      8:32 [active][ready]
```

? partition ????

? mpathb ?? partition ??????????????????????????

“ WARNING: Re-reading the partition table failed with error 22: Invalid argument.
The kernel still uses the old table. The new table will be used at the next reboot or after you run `partprobe(8)` or `kpartx(8)`

???????????????????? partition ????

```
partprobe /dev/mapper/mpathb
```