

DRBD 3

DRBD 3

 DRBD 3

 3

DRBD 3

 Single-primary mode : : 4

 Dual-primary mode : : 4

 4

DRBD 5

 5

 CentOS 5

 5

DRBD 6

 global 7

 common 7

 resource 7

 8

 initial device synchronization : 8

 Using truck based replication : 8

 4 9

 12

 12

DRBD 13

DRBD

: <https://docs.linbit.com/docs/users-guide-9.0/>

DRBD


DRBD

DRBD (, ,) . DRBD .

- 가 가
- 가 가
- () 가

DRBD Linux , DRBD 가 DRBD DRBD I/O 가 가

DRBD Linux 가 가 가 DRBD DRBD가 DRBD ext3 XFS 가




DRBD

Single-primary mode :

(ext3, ext4, XFS)
DRBD 가 (가)

Dual-primary mode :

가 가
GFS OCFS2가
DRBD 가
DRBD (: 가

 DRBD-9.0 (가). DRBD-9.1
2 가

DRBD 3 가 3 가

A : Async

TCP 가 (failover)가
가 DRBD A 가
DRBD

B : Semi-Async

() 가 (failover)

가 가

가

C : Sync

가

(가)가 가

DRBD

가

C

가

DRBD

SLES 가

(HAE)

DRBD가

. SLES

DRBD

YaST2

가

```
# yast -i drbd
```

```
# zypper install drbd
```

CentOS

CentOS

5

DRBD 8 가

. DRBD 9

EPEL, ELRepo

. DRBD

yum(

가

).

```
# yum install drbd kmod-drbd
```

Ubuntu LTS

LINBIT

<https://launchpad.net/~linbit/+archive/ubuntu/linbit-drbd9-stack>

PPA

PPA

가

```
# apt-get install drbd-utils python-drbdmanage drbd-dkms
```

DRBD

- DRBD

```
global {
    usage-count yes;
}
common {
    net {
        protocol C;
    }
}
```

- DRBD

```
resource r0 {
    on alice {
        device    /dev/drbd1;
        disk      /dev/sda7;
        address   10.1.1.31:7789;
        meta-disk internal;
    }
    on bob {
        device    /dev/drbd1;
        disk      /dev/sda7;
        address   10.1.1.32:7789;
        meta-disk internal;
    }
}
```

- DRBD

```
resource r0 {
    volume 0 {
        device    /dev/drbd1;
        disk      /dev/sda7;
        meta-disk internal;
    }
    volume 1 {
        device    /dev/drbd2;
        disk      /dev/sda8;
        meta-disk internal;
    }
}
```

```

on alice {
  address 10.1.1.31:7789;
}
on bob {
  address 10.1.1.32:7789;
}
}

```

global

/etc/drbd.d/global_common.conf

가

usage-count

DRBD

DRBD

DRBD

HTTP

usage-count no;

usage-count ask; DRBD

DRBD

(<http://usage.drbd.org>).

common

/etc/drbd.d/global_common.conf

common

common net { protocol C; }

resource

가

resource

/etc/drbd.d/<resource>.res

가 *.res

DRBD

/etc/drbd.conf

on host

(

common

(

) DRBD

)

resource

```

resource r0 {

```

```
device /dev/drbd1;  
disk /dev/sda7;  
meta-disk internal;  
on alice {  
    address 10.1.1.31:7789;  
}  
on bob {  
    address 10.1.1.32:7789;  
}  
}
```

initial device synchronization :

가

```
# drbdadm primary --force <resource>
```

Using truck based replication :

가

```
# drbdadm new-current-uuid --clear-bitmap <resource>/<volume>
```

```
# drbdsetup new-current-uuid --clear-bitmap <minor>
```

dd RAID1 1
(file-copy .)

가

```
# drbdadm new-current-uuid <resource>
```


-clear-bitmap

r0 0 2 1

V=r0/0
NODE_FROM=2
NODE_TO=1

```
drbdadm -- --force dump-md $V > /tmp/md_orig.txt
sed -e "s/node-id $NODE_FROM/node-id $NODE_TO/" \
    -e "s/^peer.$NODE_FROM. /peer-NEW /" \
    -e "s/^peer.$NODE_TO. /peer[$NODE_FROM] /" \
    -e "s/^peer-NEW /peer[$NODE_TO] /" \
    < /tmp/md_orig.txt > /tmp/md.txt
```

```
drbdmeta --force $(drbdadm sh-minor $V) v09 $(drbdadm sh-ll-dev $V) internal
restore-md /tmp/md.txt
```

drbdadm up <resource>

가

4

4

```
resource r0 {
  device      /dev/drbd0;
  disk        /dev/vg/r0;
  meta-disk   internal;

  on store1 {
    address    10.1.10.1:7100;
    node-id    1;
  }
  on store2 {
    address    10.1.10.2:7100;
    node-id    2;
  }
}
```

```
on store3 {
  address 10.1.10.3:7100;
  node-id 3;
}
on store4 {
  address 10.1.10.4:7100;
  node-id 4;
}

# All connections involving store1
connection {
  host store1 port 7012;
  host store2 port 7021;
}
connection {
  host store1 port 7013;
  host store3 port 7031;
}
connection {
  host store1 port 7014;
  host store4 port 7041;
}

# All remaining connections involving store2
connection {
  host store2 port 7023;
  host store3 port 7032;
}
connection {
  host store2 port 7024;
  host store4 port 7042;
}

# All remaining connections involving store3
connection {
  host store3 port 7034;
  host store4 port 7043;
}

# store4 already done.
}
```

4 가 Full-mesh

1-2, 1-3, 1-4, 2-3, 2-4, 3-4

```
resource r0 {
```

```

device      /dev/drbd0;
disk        /dev/vg/r0;
meta-disk   internal;

on store1 {
  address    10.1.10.1:7100;
  node-id    1;
}
on store2 {
  address    10.1.10.2:7100;
  node-id    2;
}
on store3 {
  address    10.1.10.3:7100;
  node-id    3;
}
on store4 {
  address    10.1.10.4:7100;
  node-id    4;
}

connection-mesh {
  hosts      store1 store2 store3 store4;
}
}

```

NIC가

IP

```

resource r0 {
  ...

  # store1 has crossover links like 10.99.1x.y
  connection {
    host store1 address 10.99.12.1 port 7012;
    host store2 address 10.99.12.2 port 7021;
  }
  connection {
    host store1 address 10.99.13.1 port 7013;
    host store3 address 10.99.13.3 port 7031;
  }
  connection {
    host store1 address 10.99.14.1 port 7014;
    host store4 address 10.99.14.4 port 7041;
  }

  # store2 has crossover links like 10.99.2x.y
  connection {
    host store2 address 10.99.23.2 port 7023;
    host store3 address 10.99.23.3 port 7032;
  }
}

```

```
}  
connection {  
    host store2 address 10.99.24.2 port 7024;  
    host store4 address 10.99.24.4 port 7042;  
}  
  
# store3 has crossover links like 10.99.3x.y  
connection {  
    host store3 address 10.99.34.3 port 7034;  
    host store4 address 10.99.34.4 port 7043;  
}  
}
```

(/etc/drbd.d/*.res)

```
# drbdadm create-md <resource>
```

```
[root@node1 drbd.d]# drbdadm create-md lv_voll  
md_offset 32212250624  
al_offset 32212217856  
bm_offset 32211234816  
  
Found ext3 filesystem  
    31457280 kB data area apparently used  
    31456284 kB left usable by current configuration  
  
initializing activity log  
initializing bitmap (960 KB) to all zero  
Writing meta data...  
New drbd meta data block successfully created.  
[root@node1 drbd.d]#
```

가

```
Device size would be truncated, which  
would corrupt data and result in  
'access beyond end of device' errors.
```

```
You need to either
* use external meta data (recommended)
* shrink that filesystem first
* zero out the device (destroy the filesystem)
Operation refused.
```

```
Command 'drbdmeta 0 v09 /dev/mapper/vg_data-lv_vol1 internal create-md 1'
terminated with exit code 40
```

DRBD

가

DRBD

DRBD

drbd-overview

```
nina# drbd-overview
0:r0/0 Connected(*) Seco(*)/Prim(nina) UpTo(*)/Disk(nono)
    /mnt ext3 1008M 18M 940M 2%
1:r1/0 Connected(*) Secondary(*) UpTo(*)/Disk(nono)
5:r2/0 Connected(*) Seco(*)/Prim(nini) UpTo(*)/Disk(nono)
6:r2/1 Connected(*) Seco(*)/Prim(nini) UpTo(*)/Disk(nono)
```

- r0 0 nina() Primary ext3 /mnt
- nono .(DRBD)
- UpToDate .
- r1 .
- r2 nini .

From:
<https://atl.kr/dokuwiki/> - AllThatLinux!

Permanent link:
https://atl.kr/dokuwiki/doku.php/drbd_%EC%82%AC%EC%9A%A9%EC%9E%90_%EC%95%88%EB%82%B4%EC%84%9C?rev=1527731362

Last update: 2018/05/31 01:49

