

Author: Zhang Yang Alex_doesAThotmail.com

Last Updated: 2004-7-14

General Information

The command and code execution in the following environments compiled by:

- The IBM xSeries 345 (Inter 32-bit)
- RedHat Enterprise Linux 3.0
- DB2 V8.1 for Linux

Host connected to an IBM EMP300 disk cabinet for the storage of database data. A total of six disks in disk cabinet, single block 146GB, 5 made RAID5, the remaining one as spares.

RAID5, the parity information is not the sole occupation of a disk, but according to the algorithm evenly distributed on each disk, thus avoiding the bottleneck of read and write parity information.

Assumptions on the host has installed a DB2 services, create an instance, named db2 database has not been established, and the disk cabinet has been installed in the system a good driver, we get the equipment named sdb, not yet partition .

Planning disk partition

The following step by step instructions to create a raw device disk array and data planning and placement steps:

Login as root partition on the disk array, the root prompt, type:

```
# Fdisk / dev / sdb
```

3 new primary partition sdb1, sdb2 sdb3, each partition 20GB, respectively allocated to the system catalog table space, the default system temporary table space and 32kb new temporary table space, The remaining space of about 527GB allocated to all four extended partition sdb4, increase five logical partitions sdb5. In the extended partition the sdb6, sdb7 sdb8 sdb9, the first two logical partitions each 160GB, used as a table stored in the user table space , and the third logical partition is 20GB, the user table space used as a special place index, the remaining two logical partitions for 80GB, 93GB, two logical partitions using ext3 format for storing logs and database backup.

Note: sdb1-sdb7 were not to format, ready to be used as a raw device, self-managed by DB2.

The following are points to the good disk status:

```
Disk / dev / sdb: 587.2 GB, 587257085952 bytes
254 heads, 63 sectors / track, 71677 cylinders
Units = cylinders of 16002 * 512 = 8193024 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/ Dev/sdb1	1	2501	20010469	+	83	Linux
/ Dev/sdb2	2502	5002	20010501		83	Linux
/ Dev/sdb3	5003	7503	20010501		83	Linux

```

/ Dev/sdb4 7504 71677 513456174 5 Extended
/ Dev/sdb5 7504 27502 160011967 + 83 Linux
/ Dev/sdb6 27503 47501 160011967 + 83 Linux
/ Dev/sdb7 47502 50002 20010469 + 83 Linux
/ Dev/sdb8 50003 60002 80009968 + 83 Linux
/ Dev/sdb9 60003 71677 93411643 + 83 Linux

```

Create a RAW device

Going to each logical partition DB2 used as a raw device can not be used directly, using the raw command to create a raw device name:

At the root prompt, type:

```
# raw -a
```

Raw device name query has been established, now empty, we can use from the first device name from the

At the root prompt, type:

```

# Raw / dev/raw/raw1 / dev/sdb1 # for system catalog tablespace (4K)
# Raw / dev/raw/raw2 / dev/sdb2 # used for system temp tablespace (4K)
# Raw / dev/raw/raw3 / dev/sdb3 # for system temp tablespace (32K)
# Raw / dev/raw/raw5 / dev/sdb5 # the for user TABLESPACE1 (32K), the storage table
# raw / dev/raw/raw6 / dev/sdb6 # the for user TABLESPACE1 (32K) Storage Table
# raw / dev / raw/raw7 / dev/sdb7 # for user tablespace2 (32K) store the index

```

At this point, we have created six raw device name directly to DB2 management, a look at the raw device name has been created:

```
# Raw-a
```

```

/ Dev/raw/raw1: bound to major 8, minor 17
/ Dev/raw/raw2: bound to major 8, minor 18
/ Dev/raw/raw3: bound to major 8, minor 19
/ Dev/raw/raw5: bound to major 8, minor 21
/ Dev/raw/raw6: bound to major 8, minor 22
/ Dev/raw/raw7: bound to major 8, minor 23

```

If you want to delete a raw device name (for example, / dev/raw/raw1), the root prompt, type:

```
# Raw / dev/raw/raw1 0 0
```

Then modify the system initialization file, add the following line to / etc / rc.d / sysinit file:

```

raw / dev/raw/raw1 / dev/sdb1
raw / dev/raw/raw2 / dev/sdb2

```

```
raw / dev/raw/raw3 / dev/sdb3
raw / dev/raw/raw5 / dev/sdb5
raw / dev/raw/raw6 / dev/sdb6
raw / dev/raw/raw7 / dev/sdb7
```

So that the system the next time you start, you will need to re-establish the raw device name.

Created by default raw device owners disk, database instance owner must be allowed to have access to this device, the root prompt, type:

```
# Chown db2: disk / dev / raw / raw [1,2,3,4,5,6,7]
```

DB2 user to change the seven bare owner of the device this point, these raw device has references to DB2.

Create a database

The following ready-to-use more raw device to create a newdb database contains the page size for 4kb system directory table space, the system temporary table space and user table space;

And then create a new page size is 8KB, 16KB, 32KB buffer pool;

Delete 4kb user table space, the new the two 32KB page size of the user table space for the storage of the user table;

New users of a 32KB page size table space used to store the index;

New a 32KB page size of the system temporary table space;

We will accomplish the above script saved as newdb.sql This script is performed as follows:

The root prompt, type the following command to change the DB2 user:

```
# Su - db2
```

Run the above script:

```
$ Db2-td ";"-f newdb.sql
```

At this point the database is created.

The planning logs and backup files are stored

Next, you need to plan for the storage of log and database backup, the first will be used to log and database backup format the partition (this process takes a long time):

At the root prompt, type:

```
# Mkfs.ext3 / dev/sdb8
# Mkfs.ext3 / dev/sdb9
```

To create a DB2 database directory:

```
# Mkdir / db2
```

```
# Chown db2: db2 / db2
```

```
# Su - db2
```

```
$ Mkdir / db2/db2log
```

```
$ Mkdir / db2/db2backup
```

```
$ Exit
```

mount logical partition to the directory:

```
# Mount / dev/sdb8 / db2/db2log
```

```
# Mount / dev/sdb9 / db2/db2backup
```

Update the database log file path:

```
# Su - db2
```

```
$ Db2 connect to newdb
```

```
$ Db2 update db cfg using newlogpath '/ db2/db2log'
```

Pause database changes to take effect, then to use defer parameters continue to implement the existing transaction, reject the new connection, pause after the success of released active database command:

```
$ Db2 quiesce db defer
```

```
The $ db2 unquiesce db
```

```
$ Exit
```

So far we have completed the task of the raw device data planning and placement.

Additional information

Has been created in the raw device container, if the table space where the database is not removed through the normal steps, there is no the direct delete container or the container table space where, again use these raw device, DB2 will complain:

The container is already in use

You should use release of db2untag command container mark, and with re-partition the disk or format operation is useless.

Release has in raw device raw1 establish the Container flag:

```
# Su - db2
```

```
$ Db2untag / dev/raw/raw1
```

```
????
```

- [???????](#)
- [redhat] [Basics of Raw Devices](#)