
Subject: dm-band: The I/O bandwidth controller: Performance Report

Posted by [Ryo Tsuruta](#) on Fri, 25 Jan 2008 07:07:20 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

Now I report the result of dm-band bandwidth control test I did yesterday.
I've got really good results that dm-band works as I expected. I made
several band-groups on several disk partitions and gave them heavy I/O loads.

Hardware Spec.

=====

DELL Dimention E521:

Linux kappa.local.valinux.co.jp 2.6.23.14 #1 SMP

Thu Jan 24 17:24:59 JST 2008 i686 athlon i386 GNU/Linux

Detected 2004.217 MHz processor.

CPU0: AMD Athlon(tm) 64 X2 Dual Core Processor 3800+ stepping 02

Memory: 966240k/981888k available (2102k kernel code, 14932k reserved,
890k data, 216k init, 64384k highmem)

scsi 2:0:0:0: Direct-Access ATA ST3250620AS 3.AA PQ: 0 ANSI: 5

sd 2:0:0:0: [sdb] 488397168 512-byte hardware sectors (250059 MB)

sd 2:0:0:0: [sdb] Write Protect is off

sd 2:0:0:0: [sdb] Mode Sense: 00 3a 00 00

sd 2:0:0:0: [sdb] Write cache: enabled, read cache: enabled,
doesn't support DPO or FUA

sdb: sdb1 sdb2 < sdb5 sdb6 sdb7 sdb8 sdb9 sdb10 sdb11 sdb12 sdb13 sdb14
sdb15 >

The results of bandwidth control test on partitions

=====

The configurations of the test #1:

- o Prepare three partitions sdb5, sdb6 and sdb7.
- o Give weights of 40, 20 and 10 to sdb5, sdb6 and sdb7 respectively.
- o Run 128 processes issuing random read/write direct I/O with 4KB data on each device at the same time.
- o Count up the number of I/Os and sectors which have done in 60 seconds.

The result of the test #1

```
-----  
| device   | sdb5   | sdb6   | sdb7   |  
| weight   | 40 (57.0%) | 20 (29.0%) | 10 (14.0%) |  
|-----+-----+-----+-----|  
| I/Os (r/w) | 6640( 3272/ 3368)| 3434( 1719/ 1715)| 1689( 857/ 832)|  
| sectors (r/w) | 53120(26176/26944)| 27472(13752/13720)| 13512(6856/6656)|  
| ratio to total | 56.4%   | 29.2%   | 14.4%   |  
-----
```

The configurations of the test #2:

- o The configurations are the same as the test #1 except this test doesn't run any processes issuing I/Os on sdb6.

The result of the test #2

```

-----
| device | sdb5 | sdb6 | sdb7 |
| weight | 40 (57.0%) | 20 (29.0%) | 10 (14.0%) |
|-----+-----+-----+-----|
| I/Os (r/w) | 9566(4815/ 4751)| 0( 0/ 0)| 2370(1198/1172)|
| sectors (r/w) | 76528(38520/38008)| 0( 0/ 0)| 18960(9584/9376)|
| ratio to total | 76.8% | 0.0% | 23.2% |
-----

```

The results of bandwidth control test on band-groups.

The configurations of the test #3:

- o Prepare three partitions sdb5 and sdb6.
- o Create two extra band-groups on sdb5, the first is of user1 and the second is of user2.
- o Give weights of 40, 20, 10 and 10 to the user1 band-group, the user2 band-group, the default group of sdb5 and sdb6 respectively.
- o Run 128 processes issuing random read/write direct I/O with 4KB data on each device at the same time.
- o Count up the number of I/Os and sectors which have done in 60 seconds.

The result of the test #3

```

-----
|dev|          sdb5          | sdb6 |
|---+-----+-----+-----|
|usr|  user1  |  user2  | other users | all users |
|wgt| 40 (50.0%) | 20 (25.0%) | 10 (12.5%) | 10 (12.5%) |
|---+-----+-----+-----|
|I/O| 5951( 2940/ 3011)| 3068( 1574/ 1494)| 1663( 828/ 835)| 1663( 810/ 853)|
|sec|47608(23520/24088)|24544(12592/11952)|13304(6624/6680)|13304(6480/6824)|
|% | 48.2% | 24.9% | 13.5% | 13.5% |
-----

```

The configurations of the test #4:

- o The configurations are the same as the test #3 except this test doesn't run any processes issuing I/Os on the user2 band-group.

The result of the test #4

```

-----
|dev|          sdb5          | sdb6 |
|---+-----+-----+-----|

```

```

|usr|  user1    |  user2    | other users | all users |
|wgt| 40 (50.0%) | 20 (25.0%) | 10 (12.5%) | 10 (12.5%) |
|----+-----+-----+-----+-----|
|I/O| 8002( 3963/ 4039)| 0( 0/ 0)| 2056(1021/1035)| 2008( 998/1010)|
|sec|64016(31704/32312)| 0( 0/ 0)|16448(8168/8280)|16064(7984/8080)|
| % | 66.3%    | 0.0%    | 17.0%    | 16.6%    |
-----

```

Conclusions and future works

=====

Dm-band works well with random I/Os. I have a plan on running some tests using various real applications such as databases or file servers. If you have any other good idea to test dm-band, please let me know.

Thank you,
Ryo Tsuruta.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
