

Hi All,

I report new results of dm-ioband bandwidth control test. The previous test results were posted on Jan 25.

I've got really good results as well as the last report. dm-ioband works well with Xen virtual disk.

I also announce that dm-ioband website has launched. The patches, the manual, the benchmark results and other related information are available through this site. Please check it out:

<http://people.valinux.co.jp/~ryov/dm-ioband/>

Bandwidth control on a per partition basis

=====

Test procedure

- o Prepare three partitions sda11, sda12 and sda13.
 - o Create three ioband devices ioband1, ioband2 and ioband3 on each partition respectively.
 - o Give weights of 40, 20 and 10 to each ioband device respectively.
 - o Run 50, 100 and 200 processes issuing random read/write direct I/O with 4KB data on each ioband device at the same time respectively.
 - o Count up the number of I/Os which have done in 60 seconds.
-
- o For comparison, do this test under different conditions. The conditions are:
 - Direct access to the physical devices without dm-ioband.
 - Give weights in equal proportion to each ioband devices.

| Read/Write process | | Read/Write process | | Read/Write process | |
|---|---------------|--------------------|----------------|--------------------|--|
| x 50 | | x 100 | | x 200 | |
| | | | | | |
| +-----V-----+ +-----V-----+ +-----V-----+ | | | | | |
| ioband1 | ioband2 | ioband3 | ioband devices | | |
| +-----+ +-----+ +-----+ | | | | | |
| default group | default group | default group | ioband groups | | |
| (40) | (20) | (10) | (weight) | | |
| +----- -----+ +----- -----+ +----- -----+ | | | | | |
| +-----V-----+ +-----V-----+ +-----V-----+ | | | | | |
| /dev/sda11 | /dev/sda12 | /dev/sda13 | physical devs. | | |
| +-----+ +-----+ +-----+ | | | | | |

Results

Direct access without dm-ioband

| device | sda11 | sda12 | sda13 |
|----------------|------------|-------------|-------------|
| I/O processes | 50 (14.3%) | 100 (28.6%) | 200 (57.1%) |
| I/Os | 1469 | 2486 | 5032 |
| ratio to total | 16.3% | 27.7% | 56.0% |

Weights in inverse proportion to the number of processes

| device | sda11 | sda12 | sda13 |
|----------------|------------|------------|------------|
| weight | 40 (57.1%) | 20 (28.6%) | 10 (14.3%) |
| I/Os | 5023 | 2654 | 1369 |
| ratio to total | 55.5% | 29.3% | 15.1% |

Weights in equal proportion

| device | sda11 | sda12 | sda13 |
|----------------|------------|------------|------------|
| weight | 10 (33.3%) | 10 (33.3%) | 10 (33.3%) |
| I/Os | 2954 | 3004 | 2986 |
| ratio to total | 33.0% | 33.6% | 33.4% |

Bandwidth control on a per logical volume basis

Test procedure

- o Prepare two partitions sda11 and sdb11.
- o Create a volume group with the two partitions.
- o Create two striped logical volumes on the volume group.
- o Give weights of 20 and 10 to lv0 and lv1 respectively.
- o Run 128 processes issuing random read/write direct I/O with 4KB data on each ioband device at the same time respectively.
- o Count up the number of I/Os which have done in 60 seconds.

Block diagram

Read/Write process x 128 Read/Write process x 128

```

+-----V-----+ +-----V-----+
| /dev/mapper/ioband1 | | /dev/mapper/ioband2 | ioband devices
+-----+ +-----+
| default group | | default group | ioband groups
| (20) | | (10) | (weight)
+-----+ +-----+
+-----V-----+ +-----V-----+
| /dev/mapper/lv0 | | /dev/mapper/lv1 | striped logical
| | | | volumes
+-----+
| vg0 | volume group
+-----+
+-----V-----+ +-----V-----+
| /dev/sda11 | | /dev/sdb11 | physical devices
+-----+ +-----+

```

Result

```

-----
| device | lv0 | lv1 |
| weight | 20 (66.6%) | 10 (33.3%) |
+-----+
| I/Os | 13508 | 6779 |
| ratio to total | 66.6% | 33.3% |
-----

```

Bandwidth control on a per Xen virtual block device basis

Test procedure

- o Prepare two partitions sda11 and sda12.
- o Create two ioband devices ioband1 and ioband2 on each partition respectively.
- o Give weight of 20 and 10 to each ioband device respectively.
- o Create two virtual machines that using the ioband device as Xen virtual machine's disk.
- o Run 128 processes issuing random read/write direct I/O with 4KB data on each virtual machine at the same time respectively.
- o Count up the number of I/Os which have done in 60 seconds.

```

+-----+ +-----+
| Virtual Machine 1 | | Virtual Machine 2 | virtual machines
| | | |
| Read/Write process x 128 | | Read/Write process x 128 |
| | | |

```

```

| +-----V-----+ | | +-----V-----+ +
| | /dev/xvda1 | | | | /dev/xvda1 | | virtual block
| +-----|-----+ | | +-----|-----+ | devices
+-----|-----+ +-----|-----+
+-----V-----+ +-----V-----+
| /dev/mapper/ioband1 | | /dev/mapper/ioband2 | ioband devices
+-----+ +-----+
| default group | | default group | ioband groups
| (20) | | (10) | (weight)
+-----|-----+ +-----|-----+
+-----V-----+ +-----V-----+
| /dev/sda11 | | /dev/sda12 | physical device
+-----+ +-----+

```

Result

```

-----
| virtual machine | VM1 | VM2 |
| weight | 20 (66.6%) | 10 (33.3%) |
|-----+-----+-----|
| I/Os | 7140 | 3819 |
| ratio to total | 65.2% | 34.8% |
-----

```

Bandwidth control on a per Xen virtual block device basis

=====

Test procedure

- o Prepare one partition sda11.
- o Create two files on sda11 to use as virtual machine's disks.
- o Create an ioband devices on sda11.
- o Create two virtual machines that using the prepared files as Xen virtual machine's disk.
- o Create two extra ioband-groups on ioband1, the first is of process tapdisk1 and the second is of process tapdisk2.
- o Run 128 processes issuing random read/write direct I/O with 4KB data on each virtual machine at the same time respectively.
- o Count up the number of I/Os which have done in 60 seconds.

```

+-----+ +-----+
| Virtual Machine 1 | | Virtual Machine 2 | virtual machines
| | | |
| Read/Write process x 128 | | Read/Write process x 128 |
| | | |
| +-----V-----+ | | +-----V-----+ +

```

```

| | /dev/xvda1 | | | /dev/xvda1 | | virtual block
| +-----+-----+ | | +-----+-----+ | devices
+-----+-----+ +-----+-----+
+-----V-----+ +-----V-----+
| tapdisk1 | | tapdisk2 | | tapdisk
+-----+-----+ +-----+-----+ processes
+-----V-----+ +-----V-----+
| pid of tapdisk1 | | pid of tapdisk2 | | ioband groups
| (20) | | (10) | | (weight)
| /dev/mapper/ioband1 | | ioband device
+-----+-----+
| +-----V-----+ +-----V-----+ |
| | vm1.img | | vm2.img | | disk image files
| +-----+-----+ +-----+-----+ |
| | /dev/sda11 | | physical device
+-----+-----+

```

Result

```

-----
| virtual machine | VM1 | VM2 |
| weight | 20 (66.6%) | 10 (33.3%) |
|-----+-----+
| I/Os | 7486 | 3895 |
| ratio to total | 65.8% | 34.2% |
|-----+-----+

```

Thanks,
Ryo Tsuruta

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