Netmist over Lustre FS

Configuration

Lustre Config

2 MDS active/passive

2 FC ports per MDS, 1 to each SP

2 OSS Servers

4 FC ports per OSS, 2 to each SP

3 OST's per OSS

VNX7500

1 SAS busses per SP
1 Voyagers, 60 drives total
RAID 6 8+2, 6 total Luns
1 25 Drive Tray of 10K 2.5" 600GB Drives
Two RAID 1/0 RAID Groups, 1 for MDT, 1 for the 24 Journal LUNs
3 4 port 8Gb FC Slics Per SP

Clients

Dell R810
4 x 6 Core Processors
128GB Memory
1 QDR Infiniband port
CentOS 5.6

SFS2008

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

  ************************************************************************
  |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
  |                        by                              |
  |    Author:      Don Capps. Email: capps@iozone.org              |
  |    Maintainer:   Don Capps. Email: capps@iozone.org              |
  |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
  |               George Dowding, Gary Little, Terry Capps         |
  |               Robin Miller                                |
  ************************************************************************

   Netmist       $Revision: 17 $

   Test run time = 300 seconds, Warmup = 300 seconds.
   Running 64 copies of the test on 8 clients
   Results directory: /mnt/jpedone/netmist2012/tests
   Clients have a total of 64 GiBytes of memory
   Clients have 1024 MiBytes of memory size per process
   Clients each have 8 processes
   Adjustable aggregate data set value set to 1 GiBytes
   Each process file size = 1056 kbytes
   Client data set size      = 74250 MiBytes
   Total starting data set size = 594000 MiBytes
   Total initial file space   = 594000 MiBytes
   Total max file space      = 660000 MiBytes

   Starting tests: Fri May 11 09:06:49 2012

     Launching 64 processes.
     Tests finished: Fri May 11 09:46:59 2012
     Shutting down clients, and communications layer...
     ----------------------------------------------------
     Overall average latency          7.36 Milli-seconds
     Overall Netmist_2012 Ops/sec    8841.85 Ops/sec
     Overall Read_throughput     ~  42099.27 Kbytes/sec
     Overall Write_throughput    ~  34750.35 Kbytes/sec
     Overall throughput          ~  76849.62 Kbytes/sec
     Public Finger Print       2071940488
     ----------------------------------------------------
```

BACKUP

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

  ************************************************************************
  |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
  |                        by                              |
  |    Author:      Don Capps. Email: capps@iozone.org              |
  |    Maintainer:   Don Capps. Email: capps@iozone.org             |
  |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
  |               George Dowding, Gary Little, Terry Capps         |
  |               Robin Miller                                     |
  ************************************************************************

   Netmist       $Revision: 17 $

   Test run time = 300 seconds, Warmup = 300 seconds.
   Running 64 copies of the test on 8 clients
   Results directory: /mnt/jpedone/netmist2012/tests
   Clients have a total of 64 GiBytes of memory
   Clients have 1024 MiBytes of memory size per process
   Clients each have 8 processes
   Adjustable aggregate data set value set to 1 GiBytes
   Each process file size = 1056 kbytes
   Client data set size      = 74250 MiBytes
   Total starting data set size = 594000 MiBytes
   Total initial file space   = 594000 MiBytes
   Total max file space      = 660000 MiBytes

   Starting tests: Fri May 11 09:50:07 2012

     Launching 64 processes.
     Tests finished: Fri May 11 10:31:44 2012
     Shutting down clients, and communications layer...
     ----------------------------------------------------
     Overall average latency          7.59 Milli-seconds
     Overall Netmist_2012 Ops/sec    8480.45 Ops/sec
     Overall Read_throughput     ~  42804.28 Kbytes/sec
     Overall Write_throughput    ~  34694.94 Kbytes/sec
     Overall throughput          ~  77499.22 Kbytes/sec
     Public Finger Print       2071940488
     ----------------------------------------------------
```

STREAM

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

  ************************************************************************
  |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
  |                        by                              |
  |    Author:      Don Capps. Email: capps@iozone.org              |
  |    Maintainer:   Don Capps. Email: capps@iozone.org              |
  |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
  |              George Dowding, Gary Little, Terry Capps         |
  |              Robin Miller                                |
  ************************************************************************

   Netmist       $Revision: 17 $

   Test run time = 300 seconds, Warmup = 300 seconds.
   Running 64 copies of the test on 8 clients
   Results directory: /mnt/jpedone/netmist2012/tests
   Clients have a total of 64 GiBytes of memory
   Clients have 1024 MiBytes of memory size per process
   Clients each have 8 processes
   Adjustable aggregate data set value set to 1 GiBytes
   Each process file size = 1056 kbytes
   Client data set size      = 74250 MiBytes
   Total starting data set size = 594000 MiBytes
   Total initial file space   = 594000 MiBytes
   Total max file space      = 660000 MiBytes

   Starting tests: Fri May 11 10:34:14 2012

     Launching 64 processes.
     Tests finished: Fri May 11 11:14:55 2012
     Shutting down clients, and communications layer...
     ----------------------------------------------------
     Overall average latency          7.27 Milli-seconds
     Overall Netmist_2012 Ops/sec    9033.49 Ops/sec
     Overall Read_throughput     ~ 527346.63 Kbytes/sec
     Overall Write_throughput    ~  94063.58 Kbytes/sec
     Overall throughput         ~ 621410.22 Kbytes/sec
     Public Finger Print        2071940488
     ----------------------------------------------------
```

HOMEDIR

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

 ************************************************************************
 |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
 |                        by                              |
 |    Author:      Don Capps. Email: capps@iozone.org              |
 |    Maintainer:   Don Capps. Email: capps@iozone.org             |
 |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
 |               George Dowding, Gary Little, Terry Capps         |
 |               Robin Miller                                     |
 ************************************************************************

  Netmist       $Revision: 17 $

  Test run time = 300 seconds, Warmup = 300 seconds.
  Running 64 copies of the test on 8 clients
  Results directory: /mnt/jpedone/netmist2012/tests
  Clients have a total of 64 GiBytes of memory
  Clients have 1024 MiBytes of memory size per process
  Clients each have 8 processes
  Adjustable aggregate data set value set to 1 GiBytes
  Each process file size = 1056 kbytes
  Client data set size      = 74250 MiBytes
  Total starting data set size = 594000 MiBytes
  Total initial file space   = 594000 MiBytes
  Total max file space      = 660000 MiBytes

  Starting tests: Fri May 11 11:17:00 2012

    Launching 64 processes.
    Tests finished: Fri May 11 11:57:49 2012
    Shutting down clients, and communications layer...
    ----------------------------------------------------
    Overall average latency          8.95 Milli-seconds
    Overall Netmist_2012 Ops/sec    7200.14 Ops/sec
    Overall Read_throughput     ~  48161.59 Kbytes/sec
    Overall Write_throughput    ~  43002.28 Kbytes/sec
    Overall throughput          ~  91163.87 Kbytes/sec
    Public Finger Print        2071940488
    ----------------------------------------------------
```

DATABASE

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

  ************************************************************************
  |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
  |                        by                              |
  |    Author:      Don Capps. Email: capps@iozone.org              |
  |    Maintainer:   Don Capps. Email: capps@iozone.org             |
  |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
  |               George Dowding, Gary Little, Terry Capps         |
  |               Robin Miller                                     |
  ************************************************************************

   Netmist       $Revision: 17 $

   Test run time = 300 seconds, Warmup = 300 seconds.
   Running 64 copies of the test on 8 clients
   Results directory: /mnt/jpedone/netmist2012/tests
   Clients have a total of 64 GiBytes of memory
   Clients have 1024 MiBytes of memory size per process
   Clients each have 8 processes
   Adjustable aggregate data set value set to 1 GiBytes
   Each process file size = 1056 kbytes
   Client data set size      = 74250 MiBytes
   Total starting data set size = 594000 MiBytes
   Total initial file space   = 594000 MiBytes
   Total max file space      = 660000 MiBytes

   Starting tests: Fri May 11 12:12:41 2012

     Launching 64 processes.
     Tests finished: Fri May 11 12:51:18 2012
     Shutting down clients, and communications layer...
     ----------------------------------------------------
     Overall average latency          6.73 Milli-seconds
     Overall Netmist_2012 Ops/sec    9547.12 Ops/sec
     Overall Read_throughput     ~  47854.62 Kbytes/sec
     Overall Write_throughput    ~   4747.13 Kbytes/sec
     Overall throughput          ~  52601.75 Kbytes/sec
     Public Finger Print        2071940488
     ----------------------------------------------------
```

VIRTUAL

```
[root@fdg72144 tests]# netmist -b 8 -d 8 -B 1 -t 300 -w 300 -f config_file64_60drives

  ************************************************************************
  |   Networked Maximum I/O Sustained Throughput: Netmist_2012        |
  |                        by                              |
  |    Author:      Don Capps. Email: capps@iozone.org              |
  |    Maintainer:   Don Capps. Email: capps@iozone.org             |
  |    Contributors: Don Capps, Carol Capps, Darren Sawyer, Jerry Lohr |
  |               George Dowding, Gary Little, Terry Capps         |
  |               Robin Miller                                     |
  ************************************************************************

   Netmist       $Revision: 17 $

   Test run time = 300 seconds, Warmup = 300 seconds.
   Running 64 copies of the test on 8 clients
   Results directory: /mnt/jpedone/netmist2012/tests
   Clients have a total of 64 GiBytes of memory
   Clients have 1024 MiBytes of memory size per process
   Clients each have 8 processes
   Adjustable aggregate data set value set to 1 GiBytes
   Each process file size = 1056 kbytes
   Client data set size      = 74250 MiBytes
   Total starting data set size = 594000 MiBytes
   Total initial file space   = 594000 MiBytes
   Total max file space      = 660000 MiBytes

   Starting tests: Fri May 11 12:52:45 2012

     Launching 64 processes.
     Tests finished: Fri May 11 13:32:48 2012
     Shutting down clients, and communications layer...
     ----------------------------------------------------
     Overall average latency          17.68 Milli-seconds
     Overall Netmist_2012 Ops/sec    3650.71 Ops/sec
     Overall Read_throughput     ~  18902.85 Kbytes/sec
     Overall Write_throughput    ~  24457.36 Kbytes/sec
     Overall throughput          ~  43360.21 Kbytes/sec
     Public Finger Print        2071940488
     ----------------------------------------------------
```

Custom workloads

This section describes how to create a custom workload. This includes customized op type distributions as well as custom transfer size distributions. This feature is not available in the freeware version. It is included in the commercial version of Netmist from SPEC.

To use the custom distribution capability, first run

```
netmist -E  > workloads
```

The file "workloads" will now contain something like: (example below shortened to only one type of workload. There are more in the actual –E output. (The full output contains workloads for **SFS2008**, **STREAM, BACKUP, HOMEDIR, DATABASE, VIRTUAL** and **USERDEF** workloads)

For any given workload one can specify the percentage of each op type. Also there are 15 I/O slots. Each slot contains a range for transfer size, and a percentage of the I/O transfers that will use each I/O slot. There are 15 I/O slots for read operations, and another 15 I/O slots for write operations.
As I/O operations are generated, the transfer sizes used will be in accordance with the percentages and the ranges specified.

```
Workload name SFS2008
Percent read 10
Percent mmap read 0
Percent rand read 8
Percent write 7
Percent mmap write 0
Percent rand write 2
Percent rmw 1
Percent mkdir 2
Percent unlink 4
Percent append 3
Percent locking 6
Percent access 24
Percent stat 26
Percent chmod 2
Percent readdir 4
Percent copyfile 0
Percent rename 0
Percent statfs 1
Percent pathconf 0
Read elem 0 xfer min size 1
Read elem 0 xfer max size 511
Read elem 0 xfer percent 3
Read elem 1 xfer min size 512
Read elem 1 xfer max size 1023
Read elem 1 xfer percent 1
Read elem 2 xfer min size 1024
Read elem 2 xfer max size 2047
Read elem 2 xfer percent 2
Read elem 3 xfer min size 2048
Read elem 3 xfer max size 4095
Read elem 3 xfer percent 1
```

Read elem 4 xfer min size 4096
Read elem 4 xfer max size 4096
Read elem 4 xfer percent 16
Read elem 5 xfer min size 4097
Read elem 5 xfer max size 8191
Read elem 5 xfer percent 6
Read elem 6 xfer min size 8192
Read elem 6 xfer max size 8192
Read elem 6 xfer percent 36
Read elem 7 xfer min size 8193
Read elem 7 xfer max size 16383
Read elem 7 xfer percent 7
Read elem 8 xfer min size 16384
Read elem 8 xfer max size 16384
Read elem 8 xfer percent 7
Read elem 9 xfer min size 16385
Read elem 9 xfer max size 32767
Read elem 9 xfer percent 2
Read elem 10 xfer min size 32768
Read elem 10 xfer max size 32768
Read elem 10 xfer percent 9
Read elem 11 xfer min size 65536
Read elem 11 xfer max size 65536
Read elem 11 xfer percent 4
Read elem 12 xfer min size 98304
Read elem 12 xfer max size 98304
Read elem 12 xfer percent 3
Read elem 13 xfer min size 131072
Read elem 13 xfer max size 131072
Read elem 13 xfer percent 2
Read elem 14 xfer min size 262144
Read elem 14 xfer max size 262144
Read elem 14 xfer percent 1
Write elem 0 xfer min size 1
Write elem 0 xfer max size 511
Write elem 0 xfer percent 13
Write elem 1 xfer min size 512
Write elem 1 xfer max size 1023
Write elem 1 xfer percent 3
Write elem 2 xfer min size 1024
Write elem 2 xfer max size 2047
Write elem 2 xfer percent 7
Write elem 3 xfer min size 2048
Write elem 3 xfer max size 4095
Write elem 3 xfer percent 5
Write elem 4 xfer min size 4096
Write elem 4 xfer max size 4096
Write elem 4 xfer percent 11
Write elem 5 xfer min size 4097
Write elem 5 xfer max size 8191
Write elem 5 xfer percent 3

Write elem 6 xfer min size 8192
Write elem 6 xfer max size 8192
Write elem 6 xfer percent 30
Write elem 7 xfer min size 8193
Write elem 7 xfer max size 16383
Write elem 7 xfer percent 7
Write elem 8 xfer min size 16384
Write elem 8 xfer max size 16384
Write elem 8 xfer percent 5
Write elem 9 xfer min size 16385
Write elem 9 xfer max size 32767
Write elem 9 xfer percent 1
Write elem 10 xfer min size 32768
Write elem 10 xfer max size 32768
Write elem 10 xfer percent 6
Write elem 11 xfer min size 65536
Write elem 11 xfer max size 65536
Write elem 11 xfer percent 4
Write elem 12 xfer min size 98304
Write elem 12 xfer max size 98304
Write elem 12 xfer percent 2
Write elem 13 xfer min size 131072
Write elem 13 xfer max size 131072
Write elem 13 xfer percent 2
Write elem 14 xfer min size 262144
Write elem 14 xfer max size 262144
Write elem 14 xfer percent 1
Percent write commit 33
Percent direct 0
Percent osync 0
Percent poisson 10
Percent compress 0
Background 0
Sharemode 0

Remember..

- Total percent for the op_mix for each workload type must be 100.
- Total percent for the read-transfer size distribution for each workload type must be 100.
- Total percent for the write-transfer size distribution for each workload type must be 100.
- Percent for the write commit can be 0 to 100 percent, for each of the workload types..
- Percent for direct can be from 0 to 100 percent, for each of the workload types.
- Percent for osync can be from 0 to 100 percent, for each of the workload types.
- Percent for compress can be from 0 to 100 percent, for each of the workload types.
- When one specifies “Backgound 1” then workload will not be included in the overall results. The individual workload logs will contain the results for the workload, but the overall summary will not include Background workloads.
- Sharemode will enable sharing for all processes that are using this workload object.

Now simply edit the "workloads" file to the desired distribution of workload types, ops and transfer sizes. Then later, simply use the -I option and specify the filename. e.g. -I workload