

Performance comparison between iSCSI and other hardware and software solutions

Mathias Gug
CERN, Geneva, Switzerland

We report on our investigations on some technologies that can be used to build disk servers and networks of disk servers using commodity hardware and software solutions. It focuses on the performance that can be achieved by these systems and gives measured figures for different configurations.

It is divided into two parts : iSCSI and other technologies and hardware and software RAID solutions.

The first part studies different technologies that can be used by clients to access disk servers using a gigabit ethernet network. It covers block access technologies (iSCSI, hyperSCSI, ENBD). Experimental figures are given for different numbers of clients and servers.

The second part compares a system based on 3ware hardware RAID controllers, a system using linux software RAID and IDE cards and a system mixing both hardware RAID and software RAID. Performance measurements for reading and writing are given for different RAID levels.

1. iSCSI, HyperSCSI and ENBD technologies

1.1. What is iSCSI ?

iSCSI is a protocol designed to transport SCSI commands over a TCP/IP network.

iSCSI can be used as a building block for network storage using existing IP infrastructure in a LAN/WAN environment. It can connect different types of block-oriented storage devices to servers.

iSCSI was initially standardized by ANSI T10 and further developed by the IP Storage working group of the IETF [1], which will publish soon an RFC. Many vendors in the storage industry as well as research projects are currently working on the implementation of the iSCSI protocol.

”The Small Computer Systems Interface (SCSI) is a popular family of protocols for communicating with I/O devices, especially storage devices. SCSI is a client-server architecture. Clients of a SCSI interface are called ”initiators”. Initiators issue SCSI ”commands” to request services from components, logical units, of a server known as a ”target”. A ”SCSI transport” maps the client-server SCSI protocol to a specific interconnect. Initiators are one endpoint of a SCSI transport and targets are the other endpoint.

The iSCSI protocol describes a means of transporting of the SCSI packets over TCP/IP, providing for an interoperable solution which can take advantage of existing Internet infrastructure, Internet management facilities and address distance limitations.”

draft-ietf-ips-iscsi-20

1.2. What is HyperSCSI ?

HyperSCSI is a protocol that sends SCSI commands using raw Ethernet packets instead of the TCP/IP packets used for *iSCSI*. Thus, it bypasses the TCP/IP stack of the OS and does not suffer from the shortcomings of TCP/IP.

HyperSCSI focuses on turning ethernet into a usable storage infrastructure by adding missing components such as flow control, segmentation, reassembly, encryption, access control lists and security. It can be used to connect different type of storage, such as SCSI, IDE and USB devices.

HyperSCSI is developed by the *Modular Connected Storage Architecture* group in the Network Storage Technology Division of the Data Storage Institute from the Agency for Science, Technology and Research of Singapore [2].

1.3. What is Enhanced Network Block Device (ENBD)?

ENBD is a linux kernel module coupled with a user space daemon that sends block requests from a linux client to a linux server using a TCP/IP connection. It uses multichannel communications and implements internal failover and automatic balancing between the channels. It supports encryption and authentication. This block access technology is only useful with a linux kernel because of the linux specific block request format.

It is developed by the linux community [3] under a GPL license.

2. Test configuration

2.1. Hardware and software configuration

The following hardware was used to perform the tests :

- *Test2* :
Dual Pentium 3 - 1 Ghz
3Com Gigabit Ethernet card based on BROAD-COM BCM 5700 chipset
1 Western Digital WD1800JB 180 Gbytes
3ware RAID controller 7000-series
- *Test11* :
Dual Pentium 4 - 2.4 Ghz (HyperThreading enabled)
6 Western Digital WD1800JB 180 Gbytes
3ware RAID controllers 7000-series or Promise Ultra133 IDE controllers
3Com Gigabit Ethernet card based on BROAD-COM BCM 5700 chipset
- *Test13* :
Dual AMD MP 2200+
6 Western Digital WD1800JB 180 Gbytes
3ware RAID controllers 7000-series or Promise Ultra133 IDE controllers
3Com Gigabit Ethernet card based on BROAD-COM BCM 5700 chipset
- iSCSI server : Eurologic eLANtra iCS2100 IP-SAN storage appliance - v1.0 [4]
3 SCSI drives

All the machines have a Redhat 7.3 based distribution, with kernel 2.4.19 or 2.4.20.

The following optimizations were made to improve the performance :

```
sysctl -w vm.min-readahead=127
sysctl -w vm.max-readahead=256
sysctl -w vm.bdflush =
'2 500 0 0 500 1000 60 20 0'
elvtune -r 512 -w 1024 /dev/hd{a,c,e,g,i,k}
```

2.2. Benchmarks and monitor

Two benchmarks were used to measure the IO bandwidth and the CPU load on the machines :

- *bonnie++* : v 1.03 [5]
This benchmark measures the performance of harddrives and filesystems. It aims at simulating a database like access pattern.
We are interested in two results : '*sequential output block*' and '*sequential input block*'.
Bonnie++ uses a filesize of 9GBytes with a chunksize of 8KBytes. Bonnie++ reports the

CPU load for each test. However, we found that the reported CPU load was incorrect. So we used a standard monitoring tool (*vmstat*) to measure the CPU Load during *bonnie++* runs instead.

- *seqent_random_io64* :
In this benchmark, we were interested in three results :
 - '*write*' performance : bandwidth measured for writing a file of 5 Gbytes, with a block-size of 1.5 Mbytes.
 - '*sequential reading*' performance : bandwidth measured for sequential reading of a file of 5 Gbytes with a blocksize of 1.5 Mbytes.
 - '*random reading*' performance : bandwidth measured for random reads within a file of 5 Gbytes with a blocksize of 1.5 Mbytes.

This benchmark is a custom program used at CERN to evaluate the performance of disk servers. It simulates an access pattern used by CERN applications.

vmstat has been used to monitor the CPU load on each machine.

3. Performance of iSCSI, HyperSCSI and ENBD

3.1. iSCSI performance for different software initiators

The server was the Eurologic iCS2100 IP-SAN storage appliance [4]. The client was *test13*, with kernel 2.4.19smp.

Two software initiators were used to connect to the iSCSI server : *ibmiscsi* [6] and *linux-iscsi* [7]. We used two versions of *linux-iscsi* : 2.1.2.9, implementing version 0.8 of the *iscsi* draft, and 3.1.0.6, implementing version 0.16 of the *iscsi* draft.

The results are given in the table below :

	seq output block	seq input block	write	seq read	random read
ibmiscsi 1.2.2	42 MB/s 37 % CPU	60 82 %	38 36 %	58 84 %	39 66 %
linux-iscsi 2.1.2.9	62 MB/s 43 %	62 78 %	60 44 %	60 82 %	42 42 %
linux-iscsi 3.1.0.6	64 MB/s 61 %	57 99 %	59 62 %	58 99 %	38 78 %

Comments :

- The maximum measured bandwidth of 60 MBytes/s corresponds to the maximum throughput that the disks in the server can deliver : there were only three SCSI drives in the server, each delivering a maximum throughput of around 20 Mbytes/s.
- In a storage infrastructure, the CPU load on the client should be taken into account. In order to increase the overall performance of the client, off-loading engine cards were also considered. Unfortunately, no driver for our Linux platform was available at that time.
- During the performance measurements which represent several days of continuous disk access, no crashes of either the client or the server were observed.

3.2. HyperSCSI and ENBD

The measurements were made using block client (*test11*) connected to either one or three block servers. The kernel used was 2.4.19.

The results are given in the table below :

	seq output block	seq input block	write	seq read	random read
HyperSCSI one server	46 MB/s 60 %CPU	43 81 %	46 61 %	44 79 %	21 50 %
ENBD one server	27 22 %	28 42 %	19 21 %	28 46 %	19 30 %
ENBD three server	74 46 %	44 18 %	43 45 %	51 19 %	30 22 %

Comments :

- Compared with local disk access, HyperSCSI is able to access the remote drive at its maximum throughput.
- HyperSCSI performs poorly if there are two or more connections to different servers. More generally, HyperSCSI lacks stability and works only on uniprocessor kernels 2.4.19.
- ENBD is able to handle several connections to different servers and integrates well within a software RAID array.

4. Local performance tests using hardware and software solutions.

4.1. Software RAID performance

This series of benchmarks was done on *test11* with 6 harddrives, using kernel 2.4.20. The hardware used

was only composed of IDE drives and IDE controllers. No hardware accelerator cards were used.

The results are given in the table below :

	seq output block	seq input block	write	seq read	random read
RAID 0	170 MB/s 23 %CPU	166 20 %	166 25 %	170 25 %	61 8 %
RAID 5	85 30 %	131 22 %	79 27 %	140 13 %	62 7 %

4.2. Software RAID and hardware RAID

This series of benchmarks were done on *test11* with 4 harddrives, using kernel 2.4.20.

The results are given in the table below :

	seq output block	seq input block	write	seq read	random read
Software RAID5	60 MB/s 30 %CPU	73 22 %	59 24 %	75 12 %	45 7 %
Hardware RAID5	47 6 %	59 10 %	49 2 %	61 6 %	38 4 %
Software RAID0 + Hardware RAID1	84 23 %	54 10 %	80 9 %	56 6 %	40 5 %

Comments :

SoftwareRAID delivers more bandwidth than HardwareRAID, but at a higher CPU cost.

5. Acknowledgments

I would like to thank all the people from the IT/ADC group at CERN for helping me in this study and particularly Markus Schulz, Arie Van Praag, Remi Tordeux, Oscar Ponce Cruz, Jan Iven, Peter Kelemen and Emanuele Leonardi for their support, comments and ideas.

I would also like to thank Eurologic for their iSCSI storage appliance that was used during the evaluation of the iSCSI initiators.

I am currently part of the "Volontaire Internatinal" program, which is funded by the French Foreign Ministry.

References

- [1] IETF IP Storage Working Group
<http://www.ietf.org/html.charters/ips-charter.html>
- [2] MCSA HyperSCSI
<http://nst.dsi.a-star.edu.sg/mcsa/hyperscsi/index.html>

- [3] Enhanced Network Block Device
<http://www.it.uc3m.es/~ptb/nbd/>
- [4] Eurologic eLANtra ICS2100
http://www.eurologic.com/products_elantra.htm
- [5] Bonnie++
<http://www.coker.com.au/bonnie++/>
- [6] ibmiscsi
<http://www-124.ibm.com/developerworks/projects/naslib/>
- [7] linux-iscsi
<http://linux-iscsi.sourceforge.net/>