

Performance Test Report

GE Linux with iperf

Author: George Melikian

SysKonnnect GmbH
Dept. Development

Test and Verification

Siemensstraße 23
76275 Ettlingen

Version 0.01

12-Nov-2001

Document History

Version (Owner)	Comments
0.01 (George Melikian)	First version

Table of Contents

1	Introduction	4
2	Configuration and Settings	4
2.1	The Test Server	5
2.2	The Test Clients 1-4	5
2.3	HP Switch 4108gl	6
2.4	IPerf	6

1 Introduction

In order to find out the maximum throughput of the SysKonnnect SK-98xx Gigabit Ethernet Server Adapter on Linux platforms using a HP 4108GL switch, we run a test with Iperf (v1.20, <http://dast.nlanr.net/Projects/Iperf/>).

The following section gives a survey of the procedures and tests used.

2 Configuration and Settings

The test layout was as follows:

- SuSE Linux 7.2 with the kernel version 2.4.12 on all server and clients
- The Linux test server was connected through the HP 4108GL switch to 4 Linux test clients.

The test equipment was:

- One Linux test server with two SysKonnnect SK-9844 fiber adapters in DualNet mode (both ports active)
- Four Linux test clients, each equipped with 1 SysKonnnect SK-9822 copper adapter in default mode

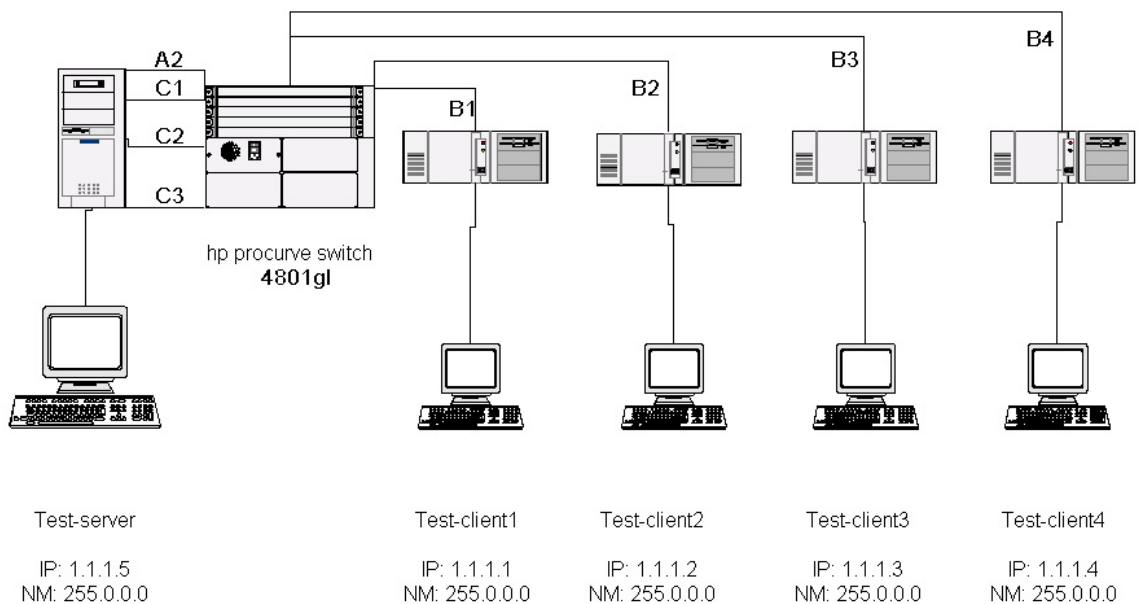


Figure 1: Test layout

2.1 The Test Server

The test server has been equipped with 2 Gigabit Ethernet fiber adapters SK-9844 (SK-NET GE-SX dual link 1000Base-SX 2 SC fiber connectors) configured in DualNet mode (both ports are active).

In addition, the patch veth-0.4.0.tar.gz (veth -- Ethernet Link Aggregation for Linux) has been installed as a link from: <http://www.st.rim.or.jp/~yumo/>.

All 4 ports are joined together in veth0 and have the same virtual ip-address=1.1.1.5.

The driver has been activated with the following commands:

```
# modprobe sk98lin RlmtMode=DualNet,DualNet
# ifconfig eth1 0 up (connected to C1 port of HP switch)
# ifconfig eth2 0 up (connected to C2 port of HP switch)
# ifconfig eth3 0 up (connected to C3 port of HP switch)
# ifconfig eth4 0 up (connected to A4 port of HP switch)
# insmod veth
# vethcfg add eth1 veth0
# vethcfg add eth2 veth0
# vethcfg add eth3 veth0
# vethcfg add eth4 veth0
# ifconfig veth0 inet 1.1.1.5 netmask 255.0.0.0 up
```

The driver status can be checked with the following command:

```
# cat /proc/net/veth
```

2.2 The Test Clients 1-4

Each test client is equipped with one SK-9822 Gigabit Ethernet copper adapter (SK-NET GE- T dual link 1000Base-T 2 RJ-45 copper connectors) configured in default mode (without LACP).

The driver has been loaded with the following commands:

Testclient 1:

```
# modprobe sk98lin
# ifconfig eth1 1.1.1.1 (connected to the port B1 of the HP switch)
```

Testclient 2

```
# modprobe sk98lin
# ifconfig eth1 1.1.1.2 (connected to the port B2 of the HP switch)
```

Testclient 3

```
# modprobe sk98lin
# ifconfig eth1 1.1.1.3 (connected to the port B3 of the HP switch)
```

Testclient 4

```
# modprobe sk98lin
# ifconfig eth1 1.1.1.4 (connected to the port B4 of the HP switch)
```

2.3 HP Switch 4108gl

The ports A2,C1,C2,and C3 are configured with LACP=yes and are connected to the test server. The ports B1,2,3, and 4 are set to LACP=no and are connected to the test clients (B1 to the test client 1, B2 to the test client 2 etc.).

LACP is working properly from the server to the clients. In the reverse direction, it does not work as intended. We can send data across 4 server ports to the HP switch and further to the clients. But if the clients put stress on the server only one port of the switch (A2) forwards the traffic to the server.

The performance was tested with iperf:

- with a PCI clock frequency of 66 MHz in 64-bit mode;
- with the default Maximum Frame Size (mtu = 1500);
- with the default driver parameters settings;

2.4 IPerf

Testing the performance throughput achievable in 10 seconds, only the UDP protocol could be used.

The following commands were used on the test server:

- On console 1 (tty1):
./iperf -c 1.1.1.1 -u -b 1000m -p5001
- On console 2 (tty2):
./iperf -c 1.1.1.2 -u -b 1000m -p5002
- On console 3 (tty3):
./iperf -c 1.1.1.3 -u -b 1000m -p5003
- On console 4 (tty4):
./iperf -c 1.1.1.4 -u -b 1000m -p5004

The following commands were applied on the test clients:

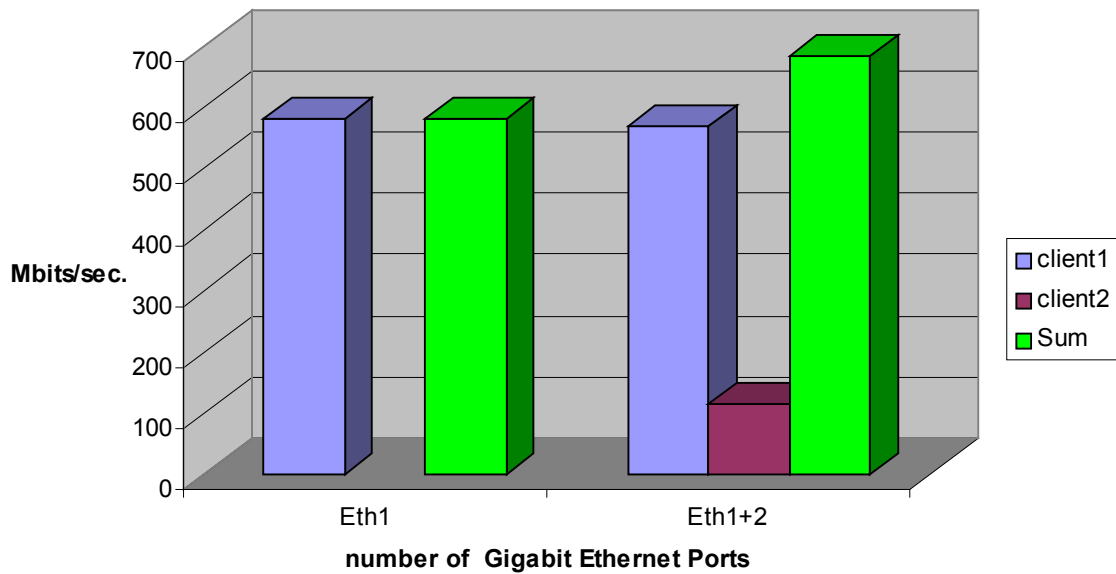
- Test client 1:
./iperf -s -u -i 1 -p 5001
- Test client 2:
./iperf -s -u -i 1 -p 5002
- Test client 3:
./iperf -s -u -i 1 -p 5003
- Test client 4:
./iperf -s -u -i 1 -p 5004

The following performance throughput was achieved:

When server and clients, both without LACP, were connected through the HP switch 4108gl:

Server	Client 1 Mbits/sec	Client 2 Mbits/sec	Client 3 Mbits/sec	Client 4 Mbits/sec	Sum Mbits/sec
Eth1	580				580
Eth1+2	568	116			684

**Performance throughput: SK-98xx without yumo patch;
through HP switch 4801gl**



When the server, equipped with LACP, was connected to the clients lacking LACP through the HP switch 4108gl:

Server	Client 1 Mbits/sec	Client 2 Mbits/sec	Client 3 Mbits/sec	Client 4 Mbits/sec	Sum Mbits/sec
Veth0=Eth1	1000				1000
Veth0=Eth1+2	460	784			1244
Veth0=Eth1+2+3	516	389	415		1320
Veth0=Eth1+2+3+4	280	462	220	446	1408

Performance throughput: SK-98xx with yumo patch; through HP switch 4801gl

