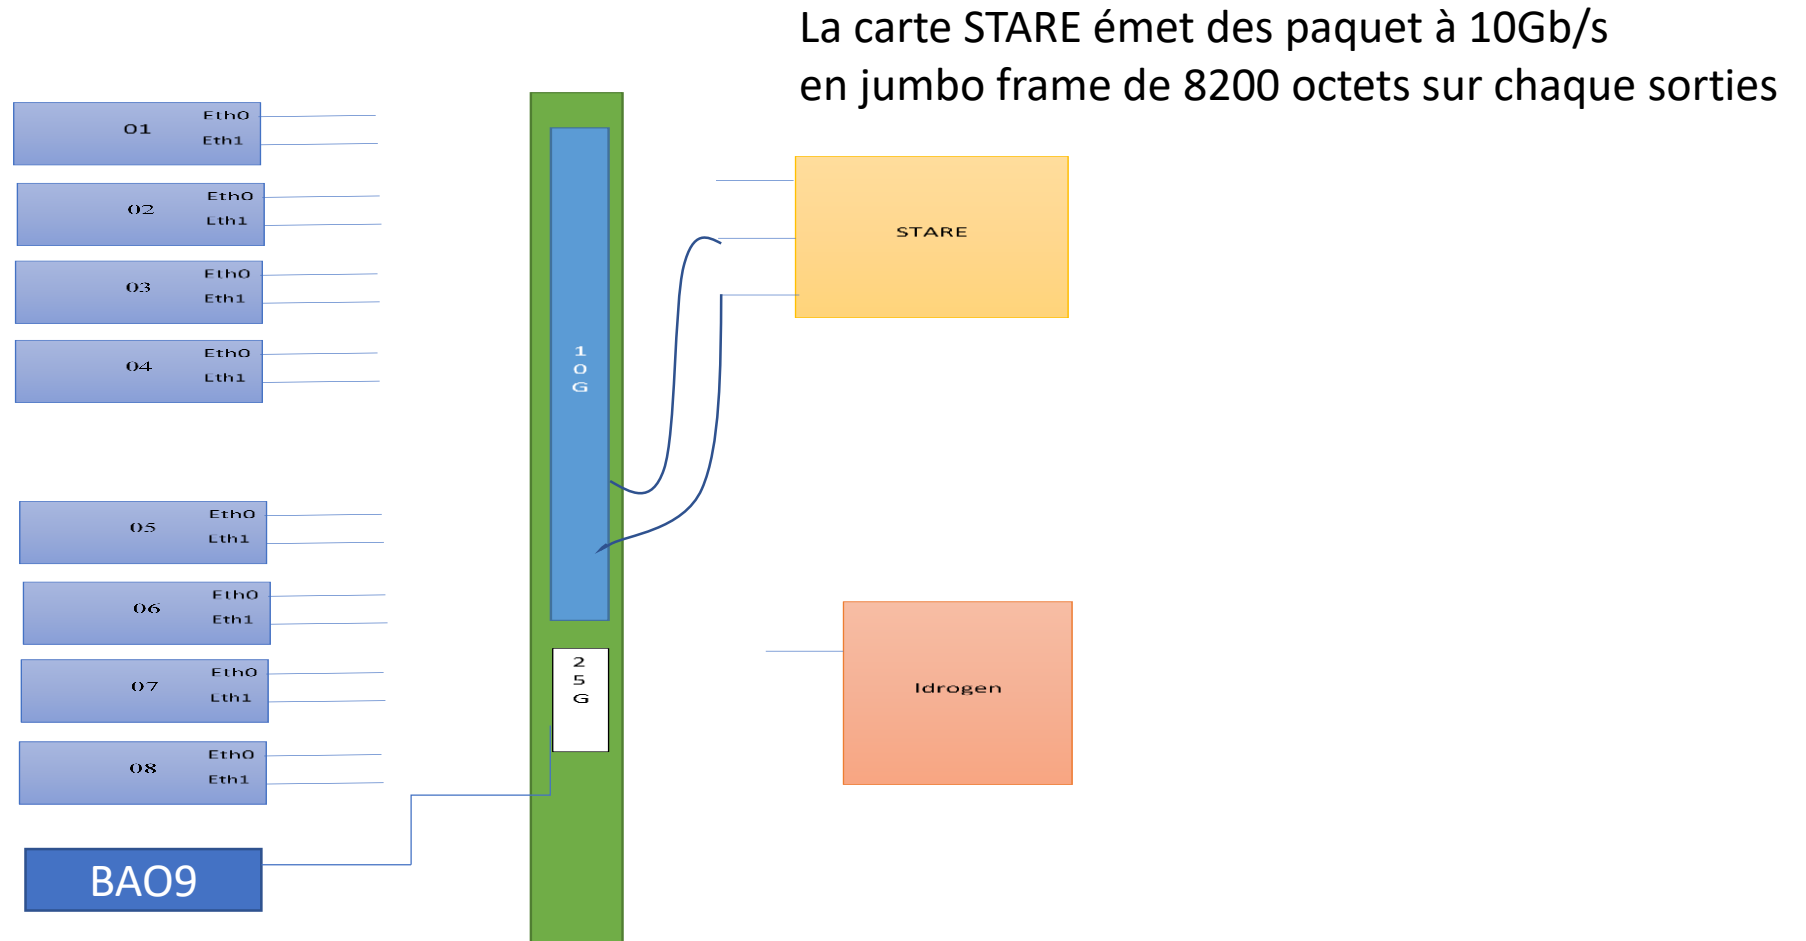


Bilan des tests pour la carte Idrogen

08/04/2022

Configuration de test



Lors de la dernière réunion j'avais constaté que je n'arrivais pas à un débit correct sur les serveurs AGATA

```
Terminal - taurigna@c6400-02: ~
Fichier  Édition  Affichage  Terminal  Onglets  Aide

ID] Interval      Transfer      Bitrate
5]  0.00-10.03   sec  10.9 GBytes  9.35 Gbits/sec
----- receiver
server listening on 5201
-----
Accepted connection from 10.10.1.8, port 38913
5] local 10.10.1.3 port 5201 connected to 10.10.1.8 port 38169
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-1.00     sec  226 MBytes  1.89 Gbits/sec  0.006 ms  0/163543 (0%)
5]  1.00-2.00     sec  236 MBytes  1.98 Gbits/sec  0.006 ms  0/170856 (0%)
5]  2.00-3.00     sec  236 MBytes  1.98 Gbits/sec  0.005 ms  0/170689 (0%)
5]  3.00-4.00     sec  236 MBytes  1.98 Gbits/sec  0.006 ms  0/170792 (0%)
5]  4.00-5.00     sec  236 MBytes  1.98 Gbits/sec  0.006 ms  0/170749 (0%)
5]  5.00-6.00     sec  240 MBytes  2.01 Gbits/sec  0.006 ms  0/173577 (0%)
5]  6.00-7.00     sec  240 MBytes  2.01 Gbits/sec  0.006 ms  0/173755 (0%)
5]  7.00-8.00     sec  240 MBytes  2.01 Gbits/sec  0.006 ms  0/173885 (0%)
5]  8.00-9.00     sec  240 MBytes  2.01 Gbits/sec  0.006 ms  0/173880 (0%)
5]  9.00-10.00    sec  240 MBytes  2.01 Gbits/sec  0.007 ms  0/173795 (0%)
5]  10.00-10.04   sec  9.01 MBytes  2.00 Gbits/sec  0.006 ms  0/6525 (0%)
----- receiver
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-10.04   sec  2.32 GBytes  1.99 Gbits/sec  0.006 ms  0/1722046 (0%)
----- receiver
server listening on 5201
-----
[5] Interval      Transfer      Bitrate
5]  0.00-10.00   sec  10.9 GBytes  9.38 Gbits/sec  113
5]  0.00-10.03   sec  10.9 GBytes  9.35 Gbits/sec
----- sender
----- receiver
perf Done.
aurigna@c6400-07:~$ iperf3 -c 10.10.1.3 -B 10.10.1.8 -b10G -u
connecting to host 10.10.1.3, port 5201
5] local 10.10.1.8 port 38169 connected to 10.10.1.3 port 5201
ID] Interval      Transfer      Bitrate      Total Datagrams
5]  0.00-1.00     sec  236 MBytes  1.98 Gbits/sec  170761
5]  1.00-2.00     sec  236 MBytes  1.98 Gbits/sec  170740
5]  2.00-3.00     sec  236 MBytes  1.98 Gbits/sec  170624
5]  3.00-4.00     sec  236 MBytes  1.98 Gbits/sec  170699
5]  4.00-5.00     sec  236 MBytes  1.98 Gbits/sec  170658
5]  5.00-6.00     sec  240 MBytes  2.01 Gbits/sec  173607
5]  6.00-7.00     sec  240 MBytes  2.01 Gbits/sec  173685
5]  7.00-8.00     sec  240 MBytes  2.01 Gbits/sec  173794
5]  8.00-9.00     sec  240 MBytes  2.01 Gbits/sec  173796
5]  9.00-10.00    sec  240 MBytes  2.01 Gbits/sec  173682
----- sender
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-10.00   sec  2.32 GBytes  1.99 Gbits/sec  0.000 ms  0/1722046 (0%) sender
5]  0.00-10.04   sec  2.32 GBytes  1.99 Gbits/sec  0.006 ms  0/1722046 (0%) receiver
perf Done.
aurigna@c6400-07:~$
```

```
perf Done.
aurigna@c6400-02:~$ iperf3 -c 10.10.1.8 -B 10.10.1.3 -b10G -u
connecting to host 10.10.1.8, port 5201
5] local 10.10.1.3 port 49981 connected to 10.10.1.8 port 5201
ID] Interval      Transfer      Bitrate      Total Datagrams
5]  0.00-1.00     sec  569 MBytes  4.78 Gbits/sec  412349
5]  1.00-2.00     sec  583 MBytes  4.89 Gbits/sec  422177
5]  2.00-3.00     sec  586 MBytes  4.92 Gbits/sec  424560
5]  3.00-4.00     sec  586 MBytes  4.91 Gbits/sec  424084
5]  4.00-5.00     sec  585 MBytes  4.91 Gbits/sec  423958
5]  5.00-6.00     sec  591 MBytes  4.96 Gbits/sec  428218
5]  6.00-7.00     sec  590 MBytes  4.95 Gbits/sec  427353
5]  7.00-8.00     sec  590 MBytes  4.95 Gbits/sec  427148
5]  8.00-9.00     sec  593 MBytes  4.98 Gbits/sec  429531
5]  9.00-10.00    sec  591 MBytes  4.96 Gbits/sec  428310
----- sender
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-10.00   sec  5.73 GBytes  4.92 Gbits/sec  0.000 ms  0/4247688 (0%) sender
5]  0.00-10.05   sec  2.43 GBytes  2.08 Gbits/sec  0.004 ms  2442522/4247533 (58%) receiver
perf Done.
aurigna@c6400-02:~$
[5] Interval      Transfer      Bitrate
5]  0.00-10.04   sec  9.36 GBytes  8.00 Gbits/sec
----- receiver
server listening on 5201
-----
Accepted connection from 10.10.1.3, port 52157
5] local 10.10.1.8 port 5201 connected to 10.10.1.3 port 49981
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-1.00     sec  241 MBytes  2.02 Gbits/sec  0.005 ms  219534/394315 (56%)
5]  1.00-2.00     sec  251 MBytes  2.10 Gbits/sec  0.003 ms  240610/422178 (57%)
5]  2.00-3.00     sec  248 MBytes  2.08 Gbits/sec  0.005 ms  244394/423786 (58%)
5]  3.00-4.00     sec  248 MBytes  2.08 Gbits/sec  0.005 ms  244576/424336 (58%)
5]  4.00-5.00     sec  250 MBytes  2.09 Gbits/sec  0.008 ms  243101/423831 (57%)
5]  5.00-6.00     sec  247 MBytes  2.07 Gbits/sec  0.005 ms  248982/427632 (58%)
5]  6.00-7.00     sec  249 MBytes  2.09 Gbits/sec  0.005 ms  246742/427030 (58%)
5]  7.00-8.00     sec  249 MBytes  2.09 Gbits/sec  0.006 ms  246847/427132 (58%)
5]  8.00-9.00     sec  253 MBytes  2.12 Gbits/sec  0.007 ms  246073/429324 (57%)
5]  9.00-10.00    sec  245 MBytes  2.06 Gbits/sec  0.006 ms  250256/428007 (58%)
5]  10.00-10.05   sec  11.8 MBytes  2.13 Gbits/sec  0.004 ms  11407/19962 (57%)
----- sender
ID] Interval      Transfer      Bitrate      Jitter      Lost/Total Datagrams
5]  0.00-10.05   sec  2.43 GBytes  2.08 Gbits/sec  0.004 ms  2442522/4247533 (58%) receiver
server listening on 5201
-----
```



Le problème a été trouvé : un alimentation manquait sur le bloc des 4 serveurs C6400-5 à 8
Et forcement mes tests sur le 25G étaient sur la 7 et 8.

De ce fait j'ai obtenu de bon résultats avec la carte STARE mais pas avec iperf (pas plus investigué)
J'ai ensuite fait le test de 2 fibres en provenance de STARE qui rentre sur la carte réseau (eth0) par
l'intermédiaire du switch:

```
taurigna@c6400-07:~/Stare_Queue_Manager$ ./Server 10.10.1.8 5002 8200 8192 1 3 30
===== Acquisition from Software Client =====
new socket 10.10.1.8 5002 Listening for packets at udp://10.10.1.8:5002
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-03-24 15:46:12
Search info event size in the first event frames
FirstFrame 16b8aea prevCounter 16b8ae9
started with frame ID 16b8aea with FirstFrame 16b8aea
Receive 9730000000 Paquets for total size 82000000000 rate 1186.79Mo/s Lost 9730 corrupted 0

started with frame ID f424d with FirstFrame f424d
Receive 10000000 Paquets for total size 82000000000 rate 1186.79Mo/s Lost 0 corrupted 0
Big Timeout ,last frame cpt ? 11d1991
q
Stop Sended
wait thread exits
End
taurigna@c6400-07:~/Stare_Queue_Manager$ ./Server 10.10.1.8 5001 8200 8192 1 3 24
===== Acquisition from Software Client =====
new socket 10.10.1.8 5001 Listening for packets at udp://10.10.1.8:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-03-24 15:45:43
Search info event size in the first event frames
FirstFrame 12c5bd4 prevCounter 12c5bd3
started with frame ID 12c5bd4 with FirstFrame 12c5bd4
Receive 9730000000 Paquets for total size 82000000000 rate 1186.79Mo/s Lost 3294 corrupted 0
```

9730 000 000 paquets reçus :
9730 + 3294 = 13024 perdus

1,3 e-4 pour 100 paquets reçus

Test sur la machine BAO9:

Cette machine n'a que 4 Cœurs , de plus la configuration des paramètres réseaux modifiés par rapport à La carte Broadcom , mais surtout les affectations des irq à un CPU n'était pas fait (sur c6400)

Test une fibre envoie à 10Gb/s pas de paquets perdus , sauf sur le CPU 3 utilisé par la gestion des RX buffer ne doit pas être utilisé.

```
bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5001 8200 8192 1 3 0
===== Acquisition from Software Client =====
new socket 10.10.1.29 5001 Listening for packets at udp://10.10.1.29:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 11:07:34
Search info event size in the first event frames
FirstFrame 26c416963a prevCounter 26c4169639
started with frame ID 26c416963a with FirstFrame 26c416963a
Receive 60000000 Paquets for total size 82000000000 rate 1124.33Mo/s Lost 0 corrupted 0

bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5001 8200 8192 1 3 1
===== Acquisition from Software Client =====
new socket 10.10.1.29 5001 Listening for packets at udp://10.10.1.29:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 09:52:14
Search info event size in the first event frames
FirstFrame 269f25ecb6 prevCounter 269f25ecb5
started with frame ID 269f25ecb6 with FirstFrame 269f25ecb6
Receive 230000000 Paquets for total size 82000000000 rate 1124.33Mo/s Lost 0 corrupted 0

bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5001 8200 8192 1 3 2
===== Acquisition from Software Client =====
new socket 10.10.1.29 5001 Listening for packets at udp://10.10.1.29:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 10:27:39
Search info event size in the first event frames
FirstFrame 26b08386fd prevCounter 26b08386fc
started with frame ID 26b08386fd with FirstFrame 26b08386fd
Receive 20000000 Paquets for total size 82000000000 rate 1124.33Mo/s Lost 0 corrupted 0

bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5001 8200 8192 1 3 3
===== Acquisition from Software Client =====
new socket 10.10.1.29 5001 Listening for packets at udp://10.10.1.29:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 10:21:50
Search info event size in the first event frames
FirstFrame 26ada9f7fd prevCounter 26ada9f7fc
started with frame ID 26ada9f7fd with FirstFrame 26ada9f7fd
Receive 30000000 Paquets for total size 82000000000 rate 813.885Mo/s Lost 10560775 corrupted 0
```

Et donc 2 x 10 Gb/s catastrophique

```
bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5002 8200 8192 1 3 2
===== Acquisition from Software Client =====
new socket 10.10.1.29 5002 Listening for packets at udp://10.10.1.29:5002
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 10:58:09
Search info event size in the first event frames
FirstFrame 2adf92f2bd prevCounter 2adf92f2bc
started with frame ID 2adf92f2bd with FirstFrame 2adf92f2bd
Receive 10000000 Paquets for total size 82000000000 rate 1161.22Mo/s Lost 214924 corrupted 0

bao@bao9:~/Stare_Queue_Manager$ ./Server 10.10.1.29 5001 8200 8192 1 3 1
===== Acquisition from Software Client =====
new socket 10.10.1.29 5001 Listening for packets at udp://10.10.1.29:5001
Type q to quit the program!
Acquisition in standalone , checking consistency of paquets
Ready, waiting datas
Found frame size of 8200 octets
wait stabilized rate
2022-04-08 10:58:08
Search info event size in the first event frames
FirstFrame 26bf761121 prevCounter 26bf761120
started with frame ID 26bf761121 with FirstFrame 26bf761121
Receive 10000000 Paquets for total size 82000000000 rate 791.253Mo/s Lost 4191116 corrupted 0
```

BAO9 n'est pas assez puissance pour traiter un tel débit

Format de paquet STARE : 8200 octets dont 8 octets de header:

Header :

80 00 xx xx xx xx xx xx + 8192 octets de data

Buffer est un unsigned char *

```
buffer[7]=0x80;
```

```
buffer[6]=0x0;
```

```
buffer[5] = (frame_id%256)&0xFF;
```

```
buffer[4] = (frame_id >> 8)&0xFF;
```

```
buffer[3] = (frame_id >> 16)&0xFF;
```

```
buffer[2] = (frame_id >> 24)&0xFF;
```

```
buffer[1] = (frame_id >> 32)&0xFF;
```

```
buffer[0] = (frame_id >> 40)&0xFF;
```

On met le compteur frame_id sur 40 bits.

Sur les machine c6400-07 et c6400-8 joignable à partir de la machine gw-online.ijclab.in2p3.fr sous le compte ansari ou perdereau mdp ??

Se connecter ssh paon@c6400-07 mdp paon4m.. ..= comme sur bao

Il y a aussi la machine bao9 : ssh bao@bao9 mdp comme d'habitude

2 repertoires :

Cd /home/paon/client_udp/client_udp

Ce client permet l'envoi de paquets UDP

`./client_udp_cpt 10.10.1.8 5001 8200 100000000 8`

usage `./client_udp_cpt hostname port paqsize nbPaquets tempo`

Si le client est sur bao9 tempo=0

Cd /home/paon/Stare_Queue_Manager

Source setup.sh

Usage: ./Server <ip> <port><frameSize (in octets) <Event Size in octets)> <Run Number> <FrameID

Format 0,1,2,3 (0 agata-v1 ,1 Cpt64,2 Idrogen,3 Agata-v2)> < CPU (> 19 even)

Sur une C6400 : ./Server 10.10.1.9 5001 8200 8192 1 3 20

Sur Bao9 : ./Server 10.10.1.9 5001 8200 8192 1 3 2

En argument paqsize=frameSize = (EventSize + header)

En cas de reboot des serveurs il faudra aller dans

Cd /home/bao/Stare_Queue_Manager/Scripts

Puis executer

Sudo ou en su sur bao9 ./nic_25.sh eth0 (ou enp8s0f0 sur bao9)

Puis ./irq_affinity.sh