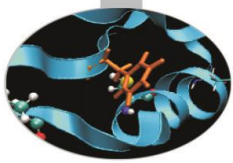

Outline

- MARCONI: Sanity checks launched during maintenance and production on June
- MARCONI100: Sanity checks on GPU (dgemm, stream, nvlink, p2p) during maintenance on May and production on June

MARCONI – June 2022

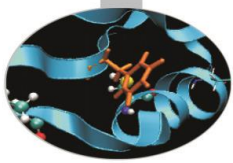

- Tests were performed during maintenance on June 14th
- Tests: linpack and stream on each single node of SKL
- Switch-level Linpack tests were performed during maintenance
- Slow nodes have been removed from production for investigation and fixes

Linpack SKL

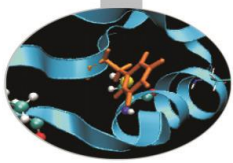

- Sanity threshold: 1800 Gflops (no slow nodes found)

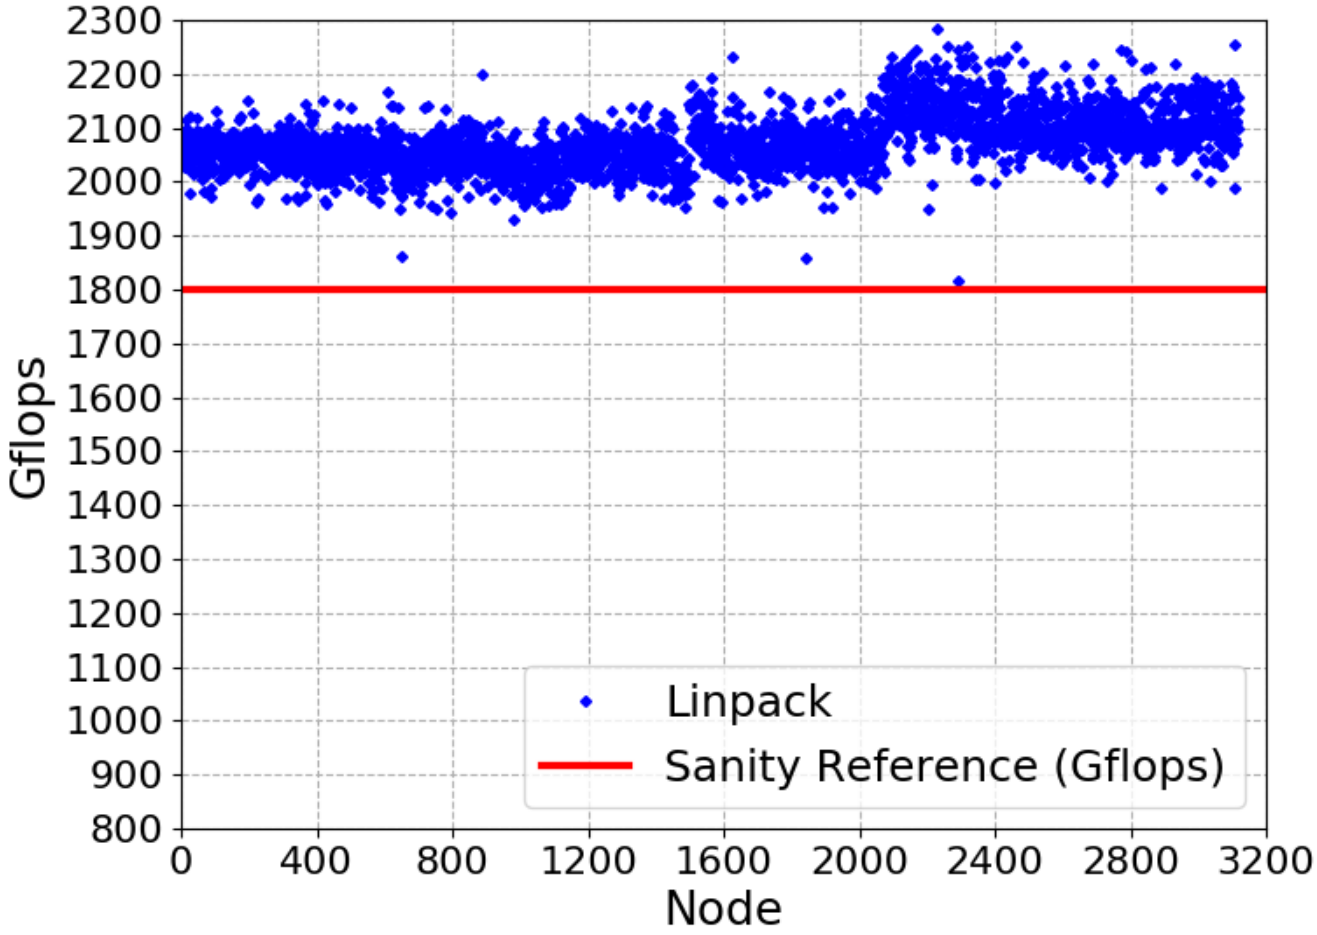

Stream SKL

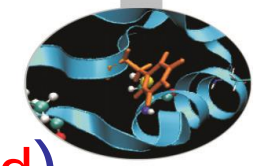

- Sanity threshold: 130000 MB/s (no slow nodes found)

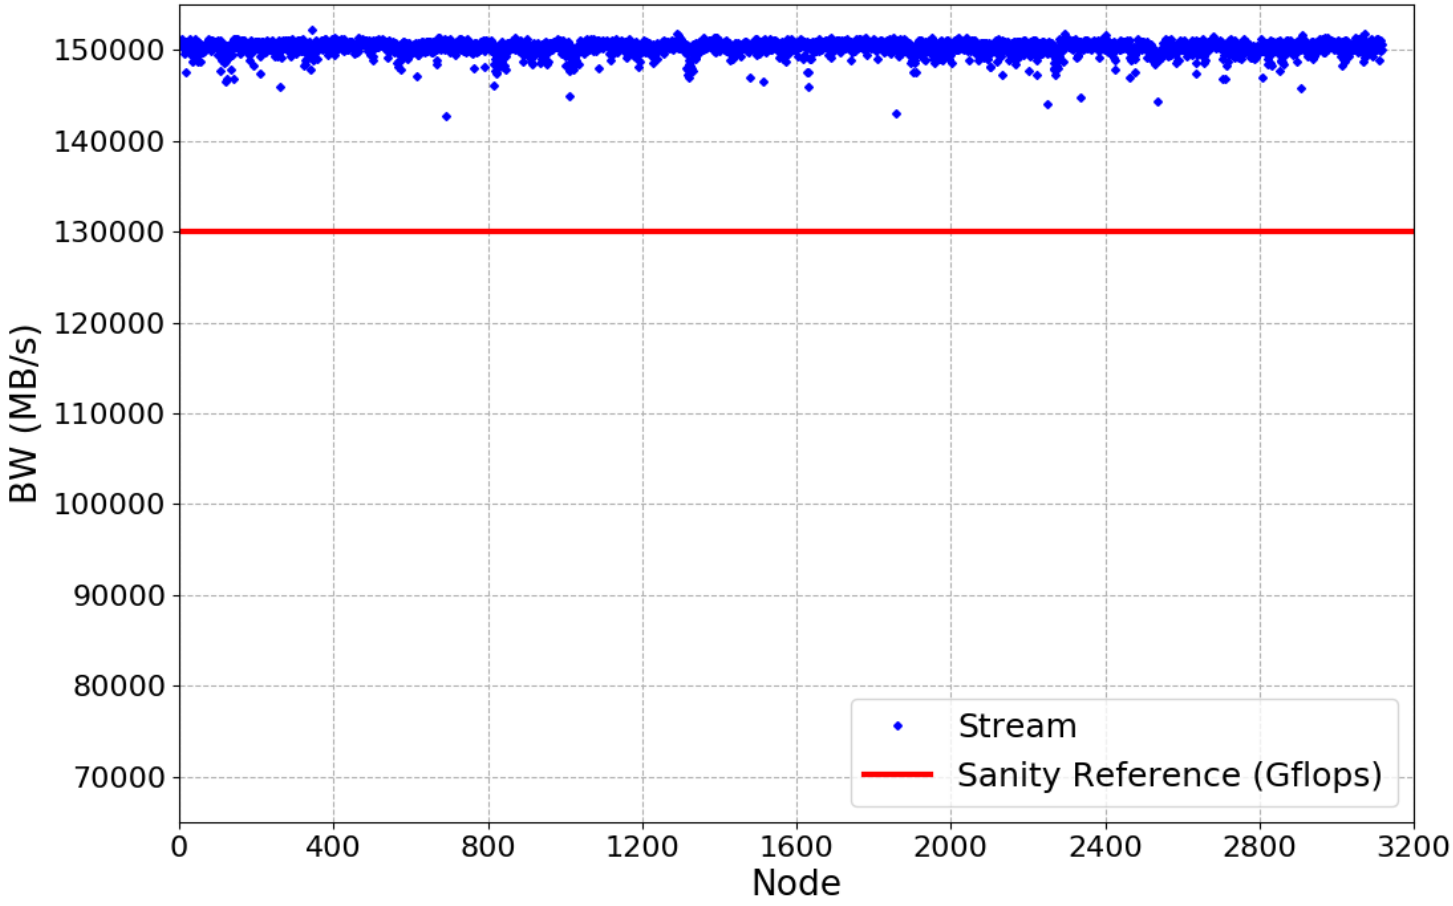

Switch-level Linpack SKL

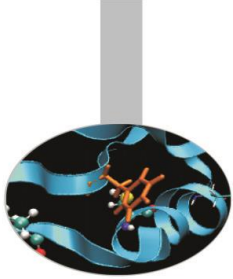

- 93 switches tested
- Overall stable performance

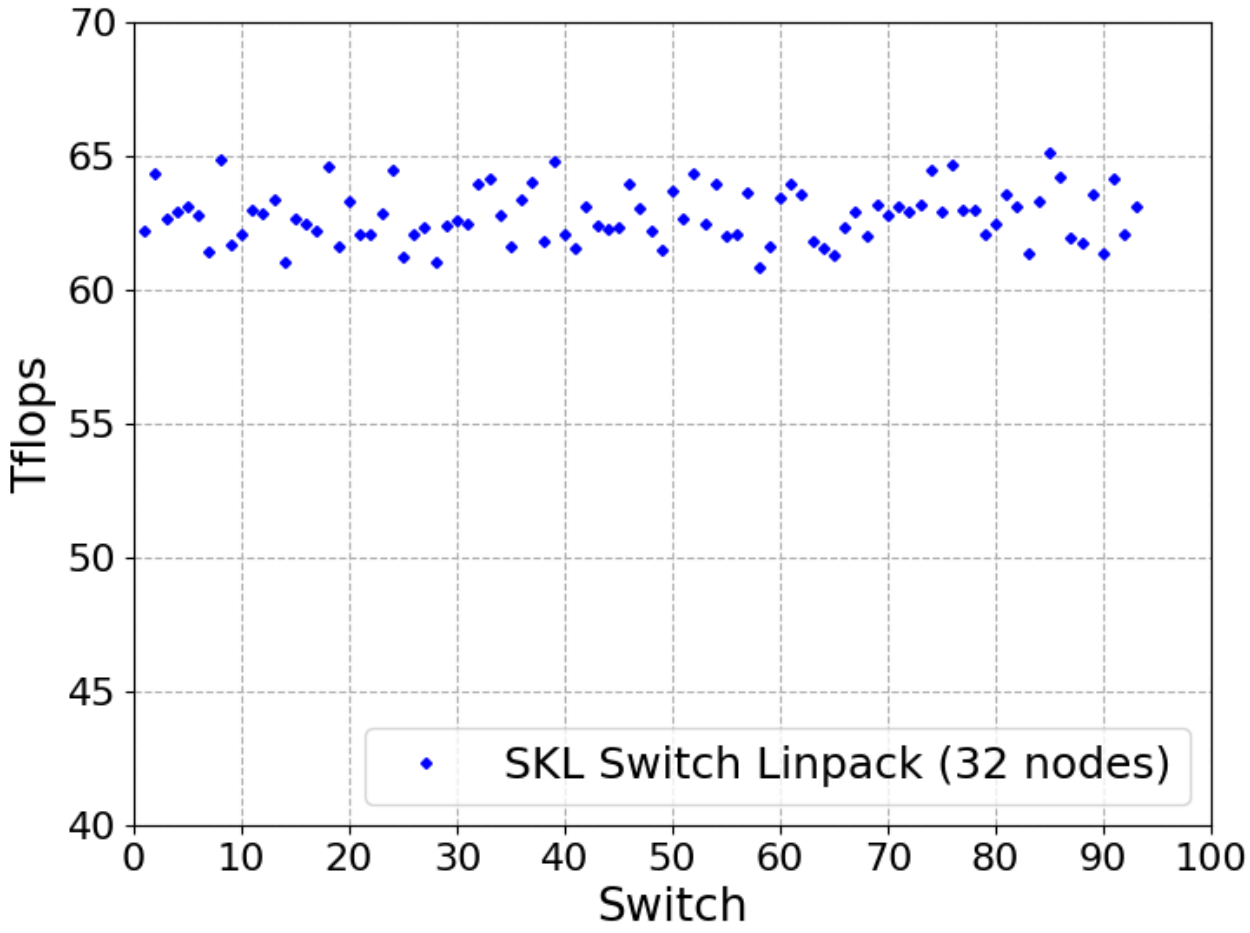

MARCONI – June 2022

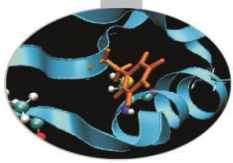

- Tests were performed during production and launched on June 24th
- Tests: linpack and stream on each single node of SKL
- 1597 were tested in time for this Ticket Meeting (only EF nodes tested)
- No Switch-level Linpack tests with production runs
- Slow nodes have been removed from production for investigation and fixes

Linpack SKL

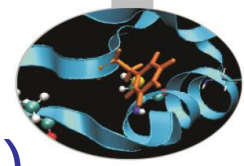

- Sanity threshold: 1800 Gflops (three slow nodes found)

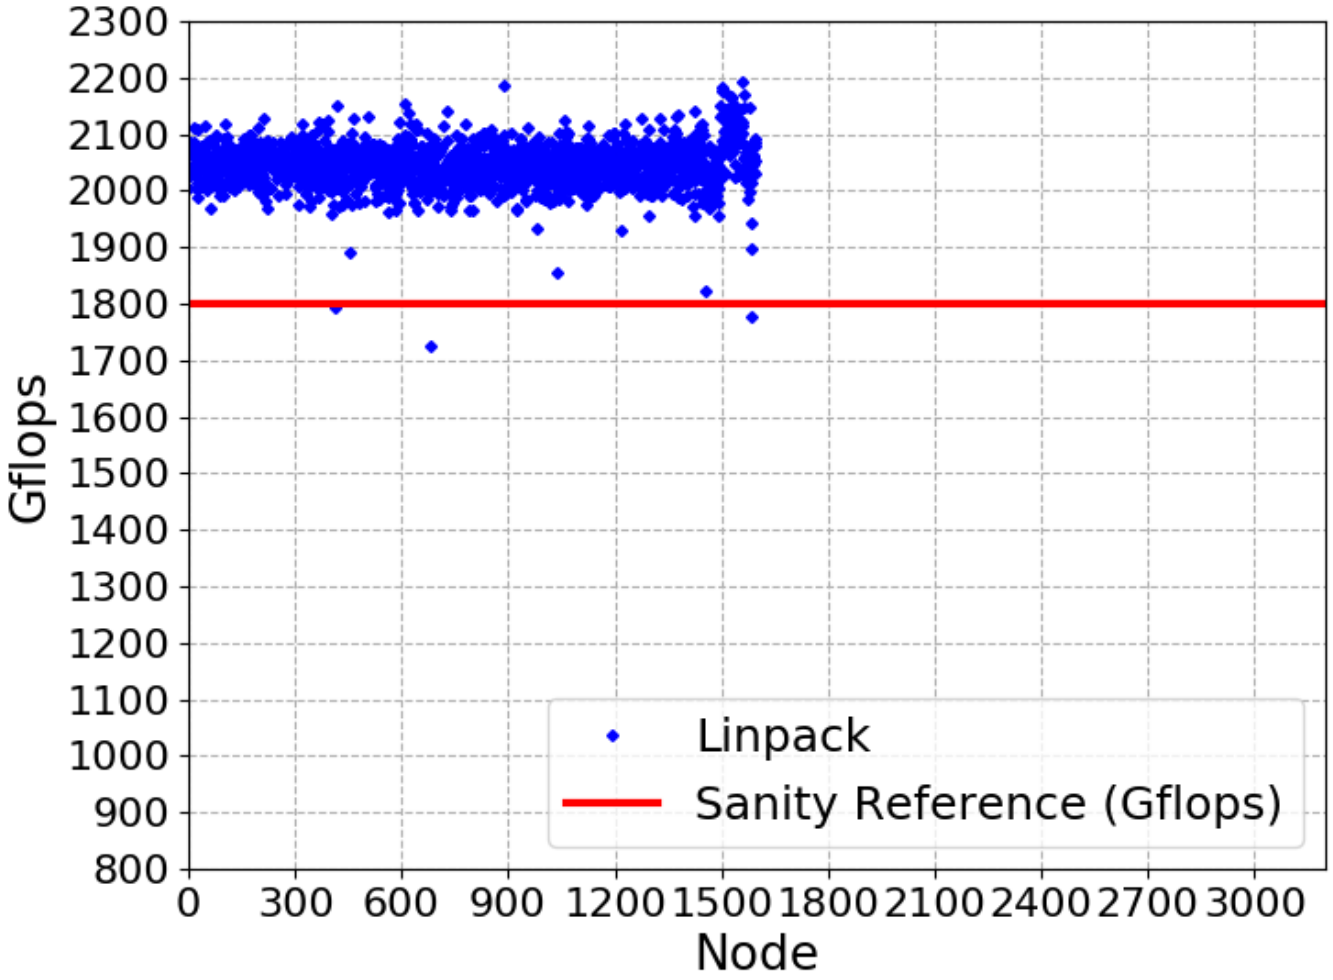

Stream SKL

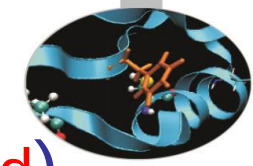

- Sanity threshold: 130000 MB/s (no slow nodes found)

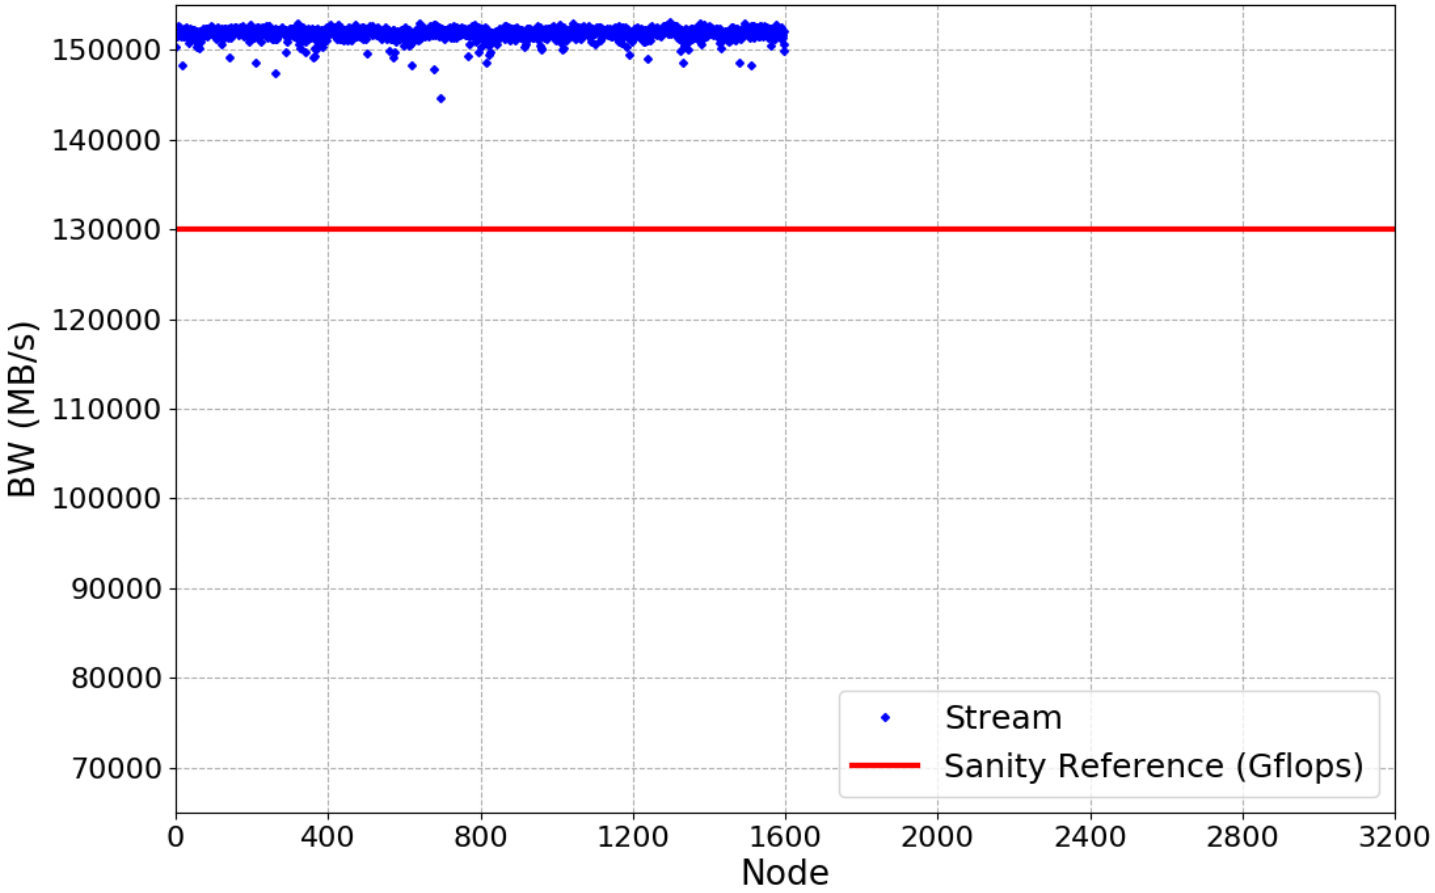

MARCONI100 – June 2022

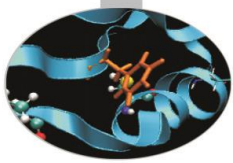

- Sanity checks were performed during production and launched on June 30th
- Tests: dgemm, stream-GPU, nvlink and p2p on each single node of M100
- Results from the nodes of EUROFusion partition are shown

Dgemm

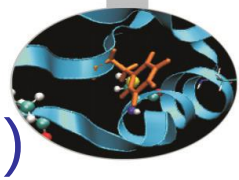

- Sanity threshold: 6.5 Tflops (two slow GPUs found)

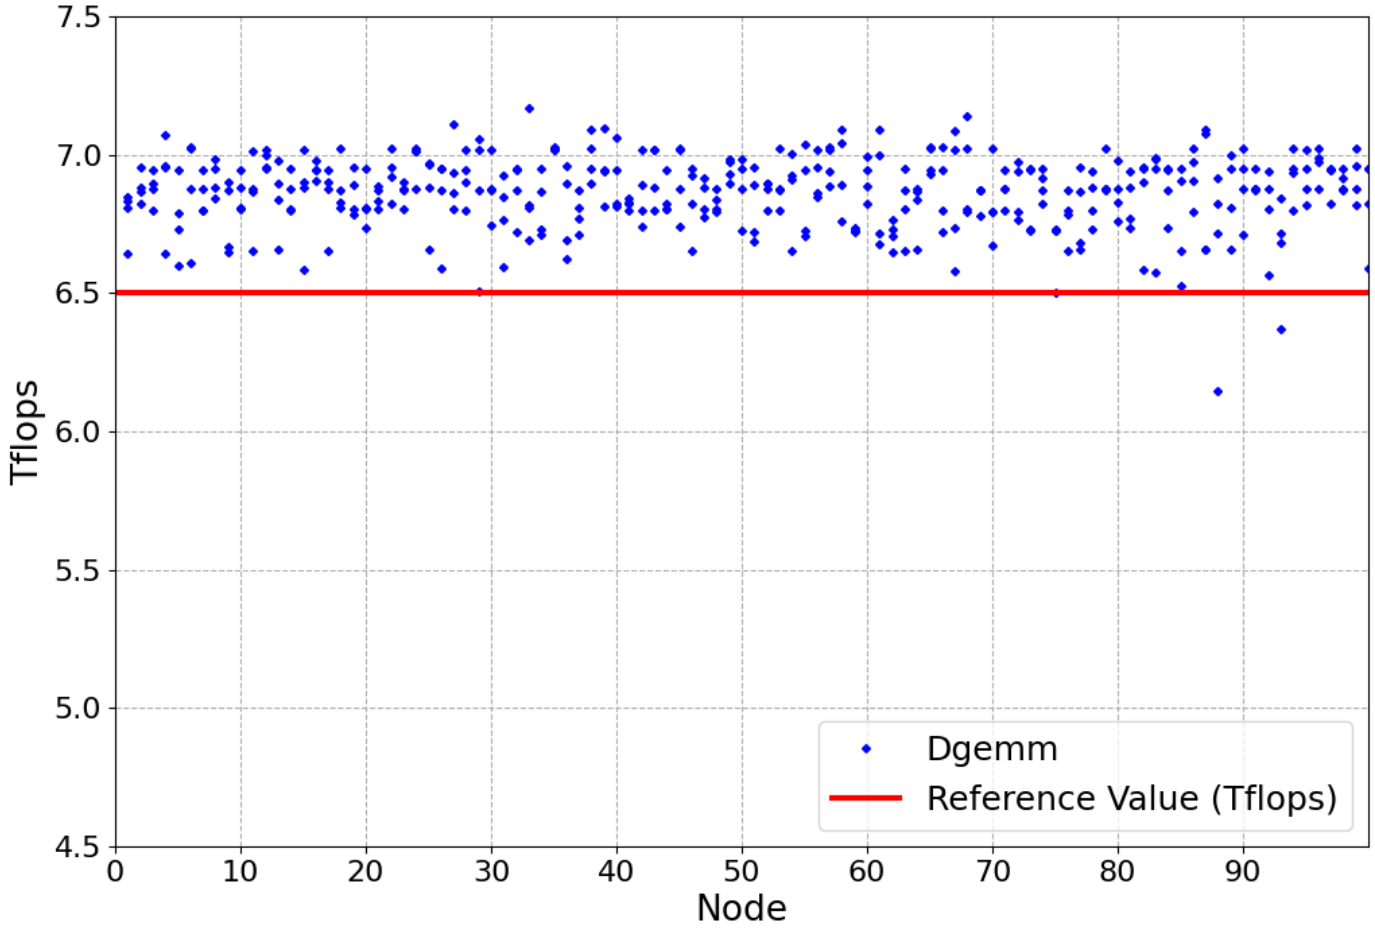

Stream-GPU

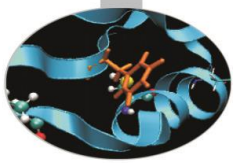

- Sanity threshold 750 Gb/s (no slow GPUs found)

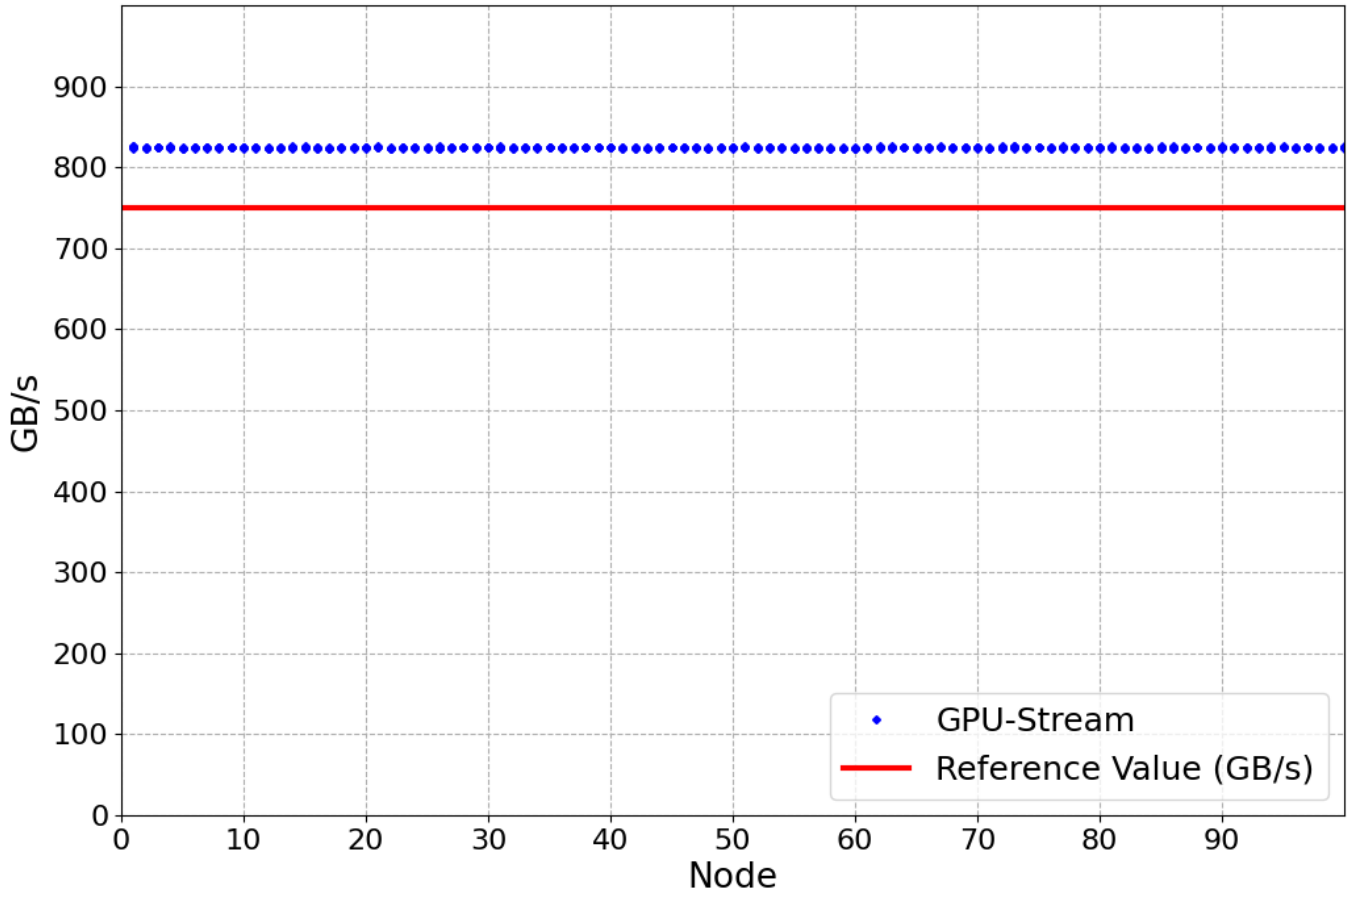

Nvlink-GPU

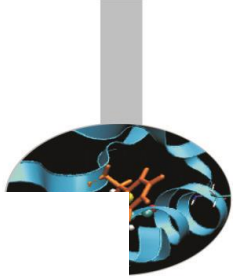

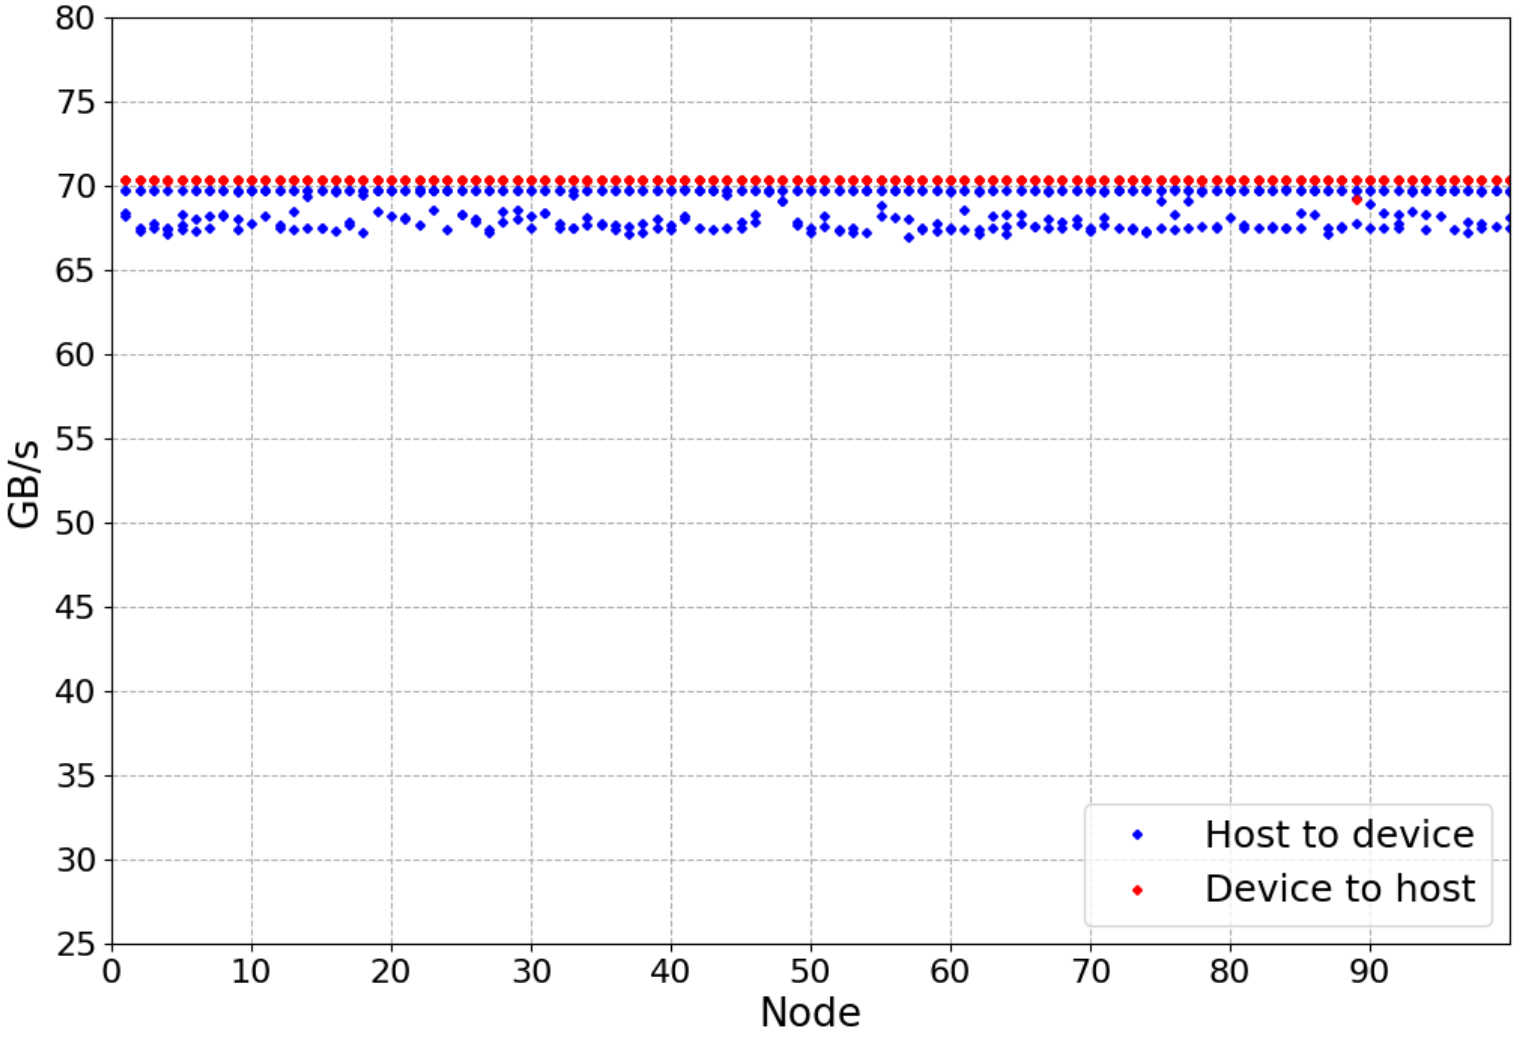
GB/s
Node
Host to device
Device to host
80
75
70
65
60
55
50
45
40
35
30
25
0
10
20
30
40
50
60
70
80
90

P2P-GPU

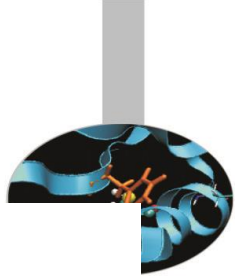

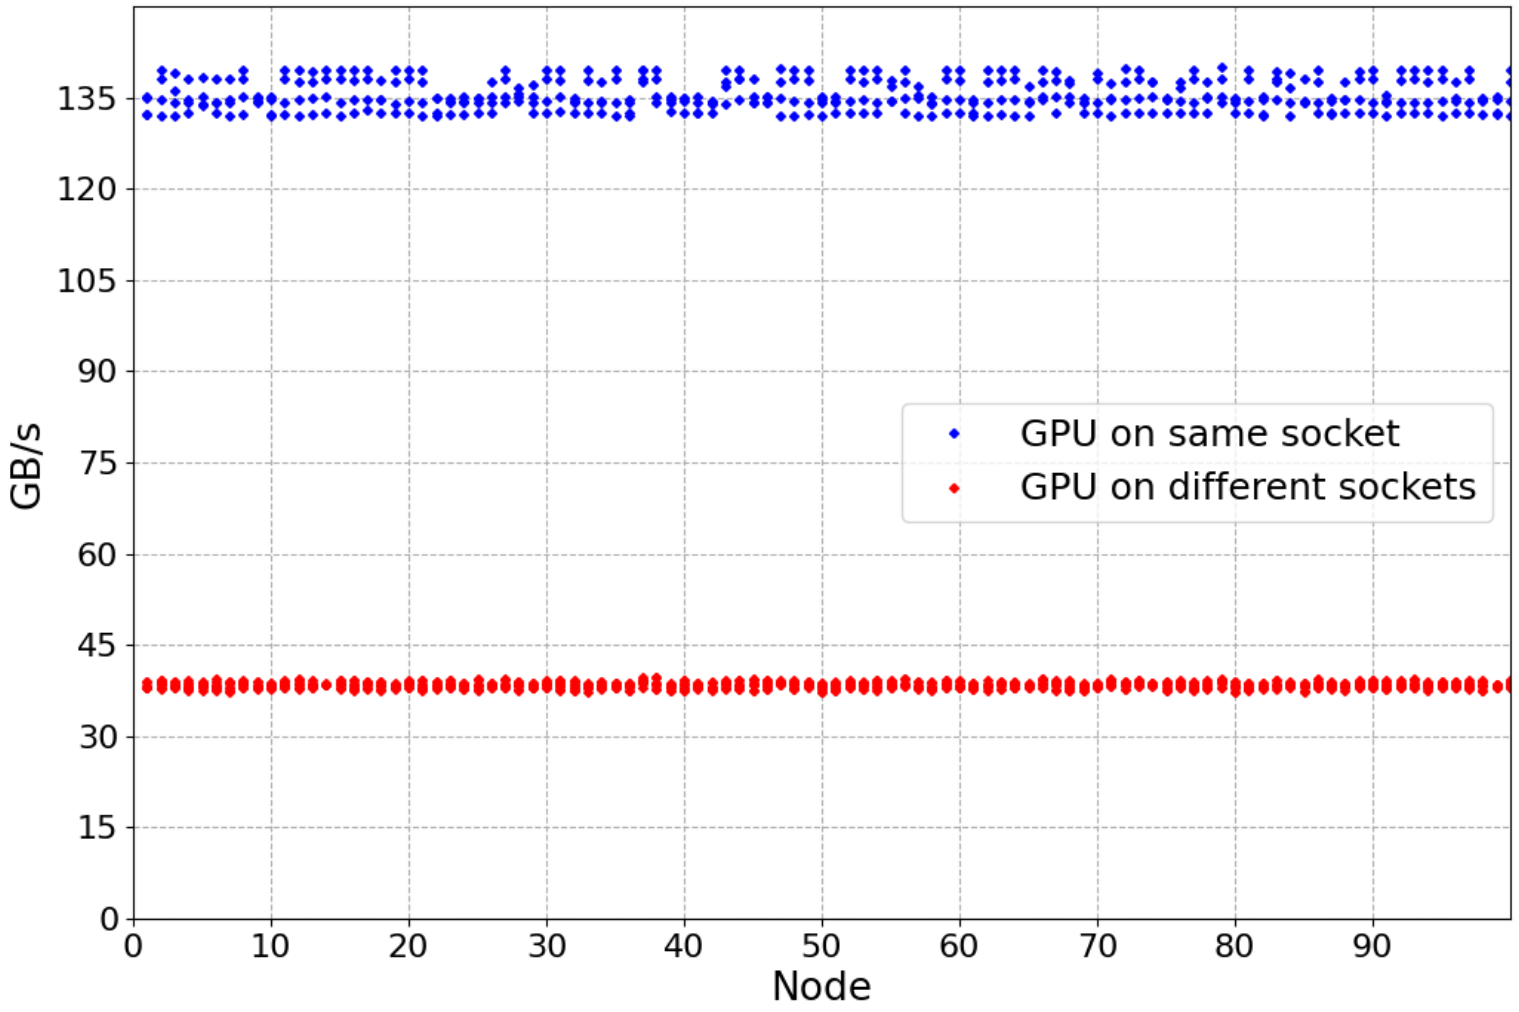
GB/s
135
120
105
90
75
60
45
30
15
0
0
10
20
30
40
50
60
70
80
90
Node
GPU on same socket
GPU on different sockets

Updates on other topics

- Power capping on Marconi: no news, sysadmins are working on improving the power capping monitoring system. Waiting for a node to study (it may take some time)