


CPU on the farms

Yen-Chu Chen, Roman Lysak,
Miro Siket, Stephen Wolbers


June 4, 2003

Plot found on fnpcc (farms home page)
What does it mean?
Why isn't it higher?
Can it get higher?



Each CPU can only deliver 100% total
(user + system)


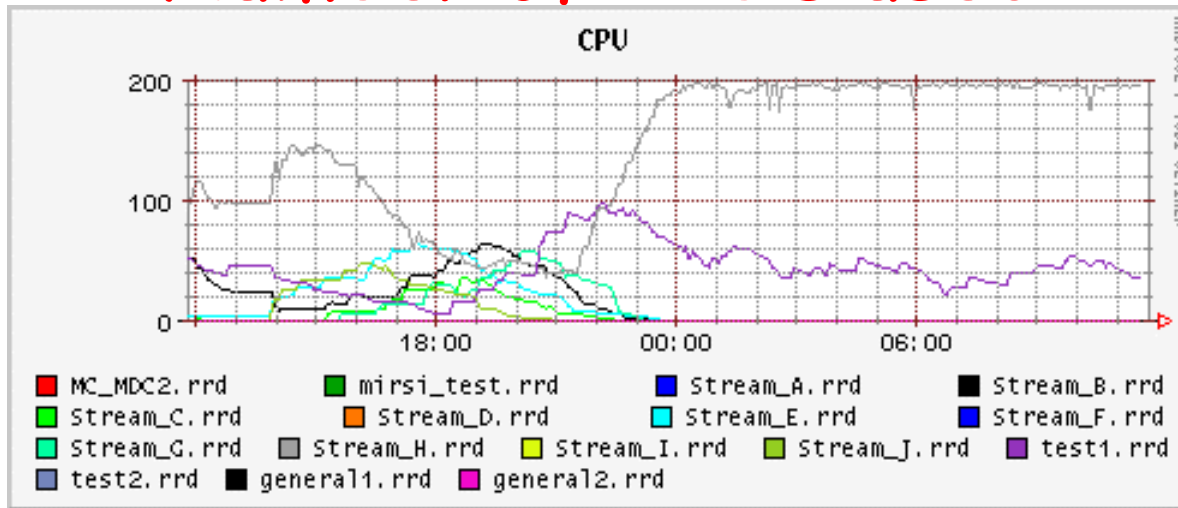
So one has to add the two lines:



I/O nodes - Not available for processing

- There are 16 "I/O" nodes on the farms.
 - 8 input
 - 8 output
- There is no CPU available for reconstruction on those nodes.
- The amount of CPU used on those I/O nodes is quite low.
- Subtract 32 processors and the total available is 270, not 302 CPUs. This adds 10% to the utilization graph.
- So the 80% (70+10) becomes **88%**.

Now look at the plots behind: Number of CPU's used

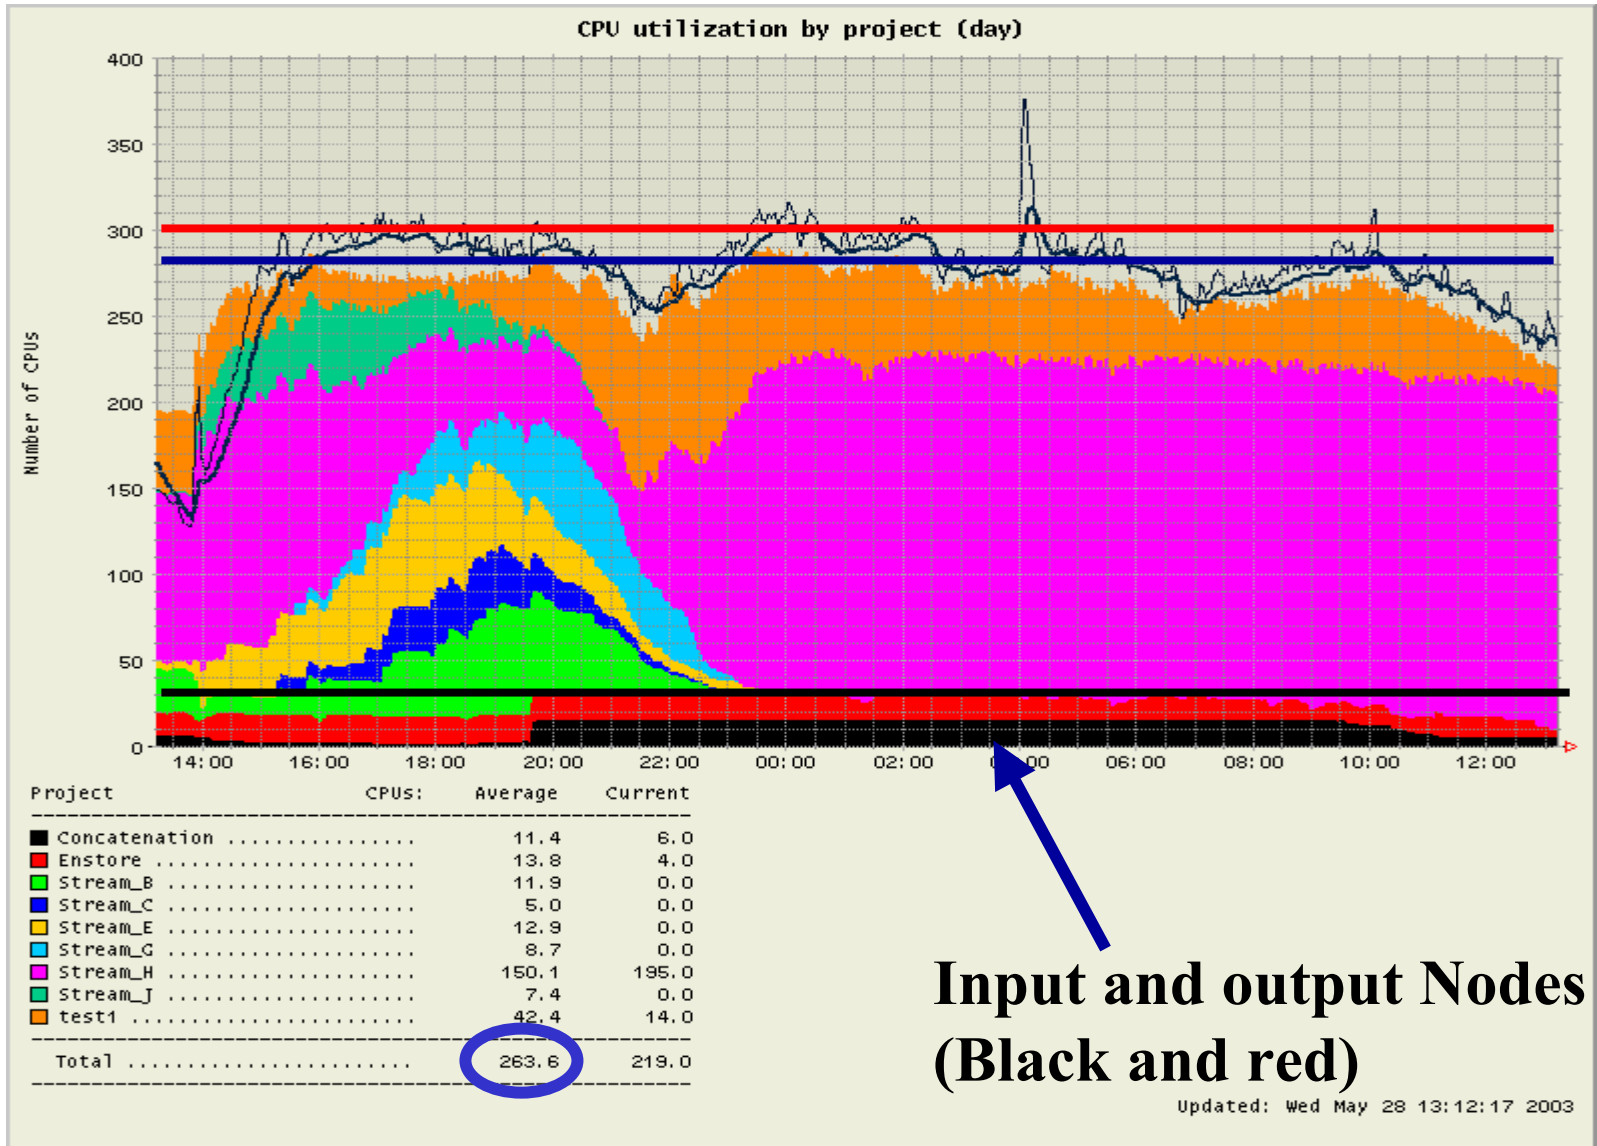


$270 * 0.88$
 $= 238$
(consistent)

270

238

FBS Information



302
280
(Max Used)

32
(I/O)

Input and output Nodes (Black and red)

FBS utilization

- 248 nodes used for processing at maximum (280 (total) - 32 (I/O)).
- 270 available for processing.
- Ratio is $248/270 = 92\%$.
- Consistent with 88%.



What does it mean?

- While the farm is being used fully (Friday May 23 through Wed May 28) the utilization of the farm was about **88-92%**, whether it is measured by fraction of CPU used on the 270 CPU's available for processing or by the number of CPU's used as assigned by the batch system.

Consistency of the result with the number of events processed

- One can calculate the expected number of events/day that can be reconstructed on this farm using the CPU time/event and the total CPU.
- Total capacity is $353-16-20 = 317$ GHz
 - (23 800 MHz, 64 1 GHz, 32 1.26 GHz, 32 1.67 GHz).
- CPU time/event = 3.3 GHz-s.
 - Still working on this number. Yen-Chu may talk about this next week.
- Max. No. of events possible = $(317/3.3) * 24 * 3600$
 - Max. Total No. of events = 8.3M events/day.
- Average processed during this time was 7.5M/day.
- This gives $7.5/8.3 = 90\%$

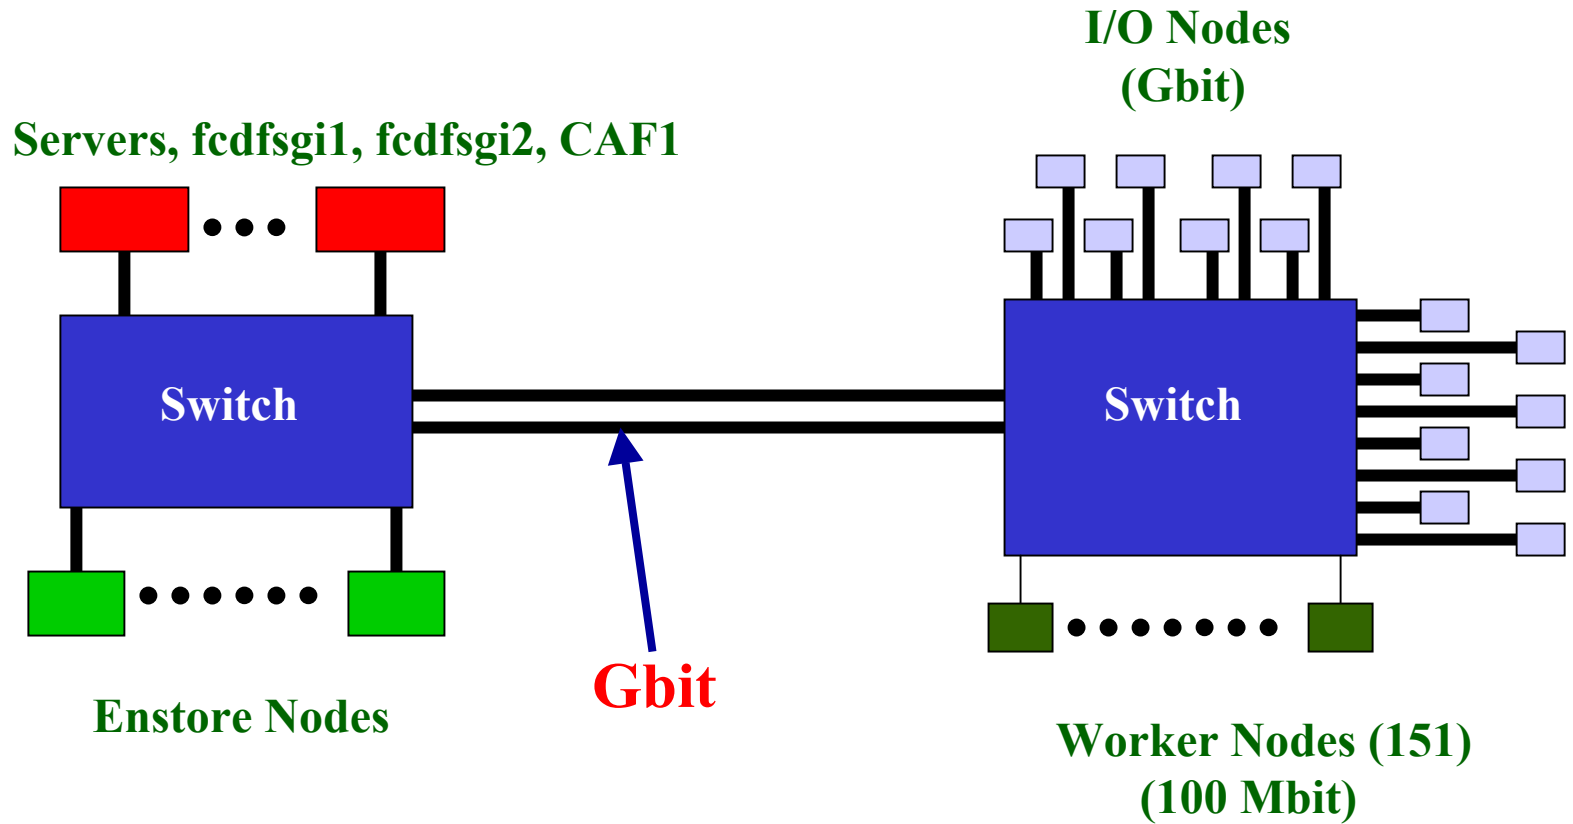
CPU Time measurements



Can the farm expand?

- The issue is whether we are "bottlenecked" in any part of the system:
 - Input
 - Batch System
 - Output
 - Internal transfers
 - Database accesses

Farm Topology



Farms Expansion

- Given the current hardware, no bottlenecks are expected in:
 - Input (8 simultaneous Gbit connections)
 - Output (8 simultaneous Gbit connections)
 - But there is the switch-to-switch bandwidth limitation to take into account.
 - Batch system (FBSNG should be OK)
 - Internal Transfers (32 Gbit backplane speed on switch)
- Not sure about DB access
 - Internal Farms DB
 - CDF DB (calibrations and DFC access)

Conclusion

- The Production Farm is running at about 90% efficiency, we think.
- We may be able to increase the efficiency by a little but there is no huge increase available.
- The graph on the fnpcc home page is misleading. We should update it.
- The farms can expand by adding CPUs or by replacing old slow CPUs with newer ones.
- The FY03 plan for the farms adds more CPUs to be able to handle increased data from online and reprocessing.