DAQ Overview

+ selected Topics

Beat Jost
Cern EP

Overall Architecture

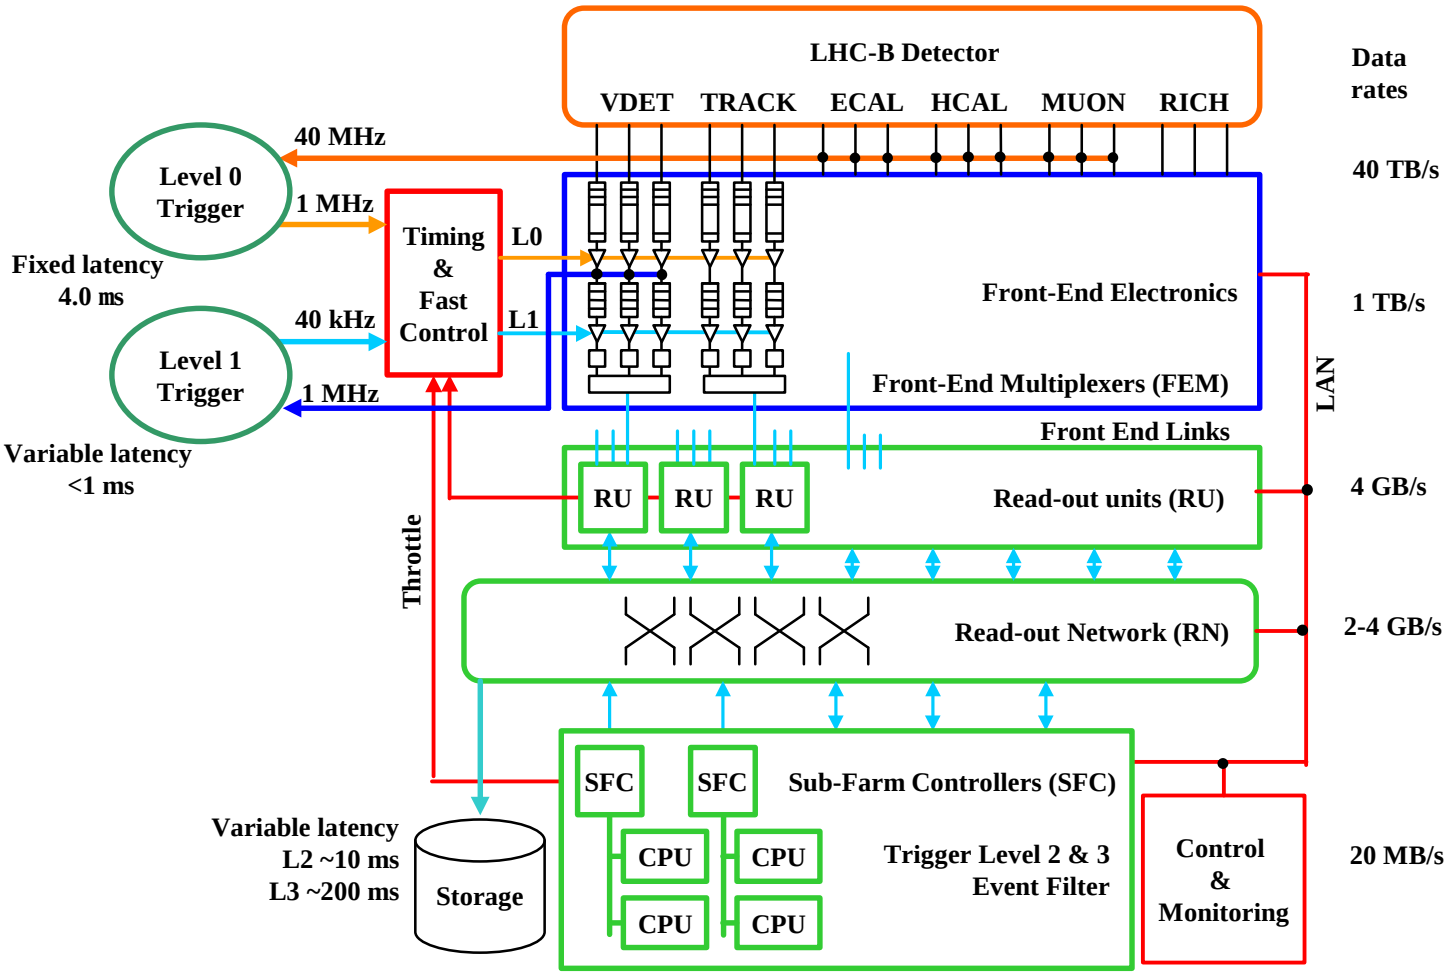

Event Builder Architecture

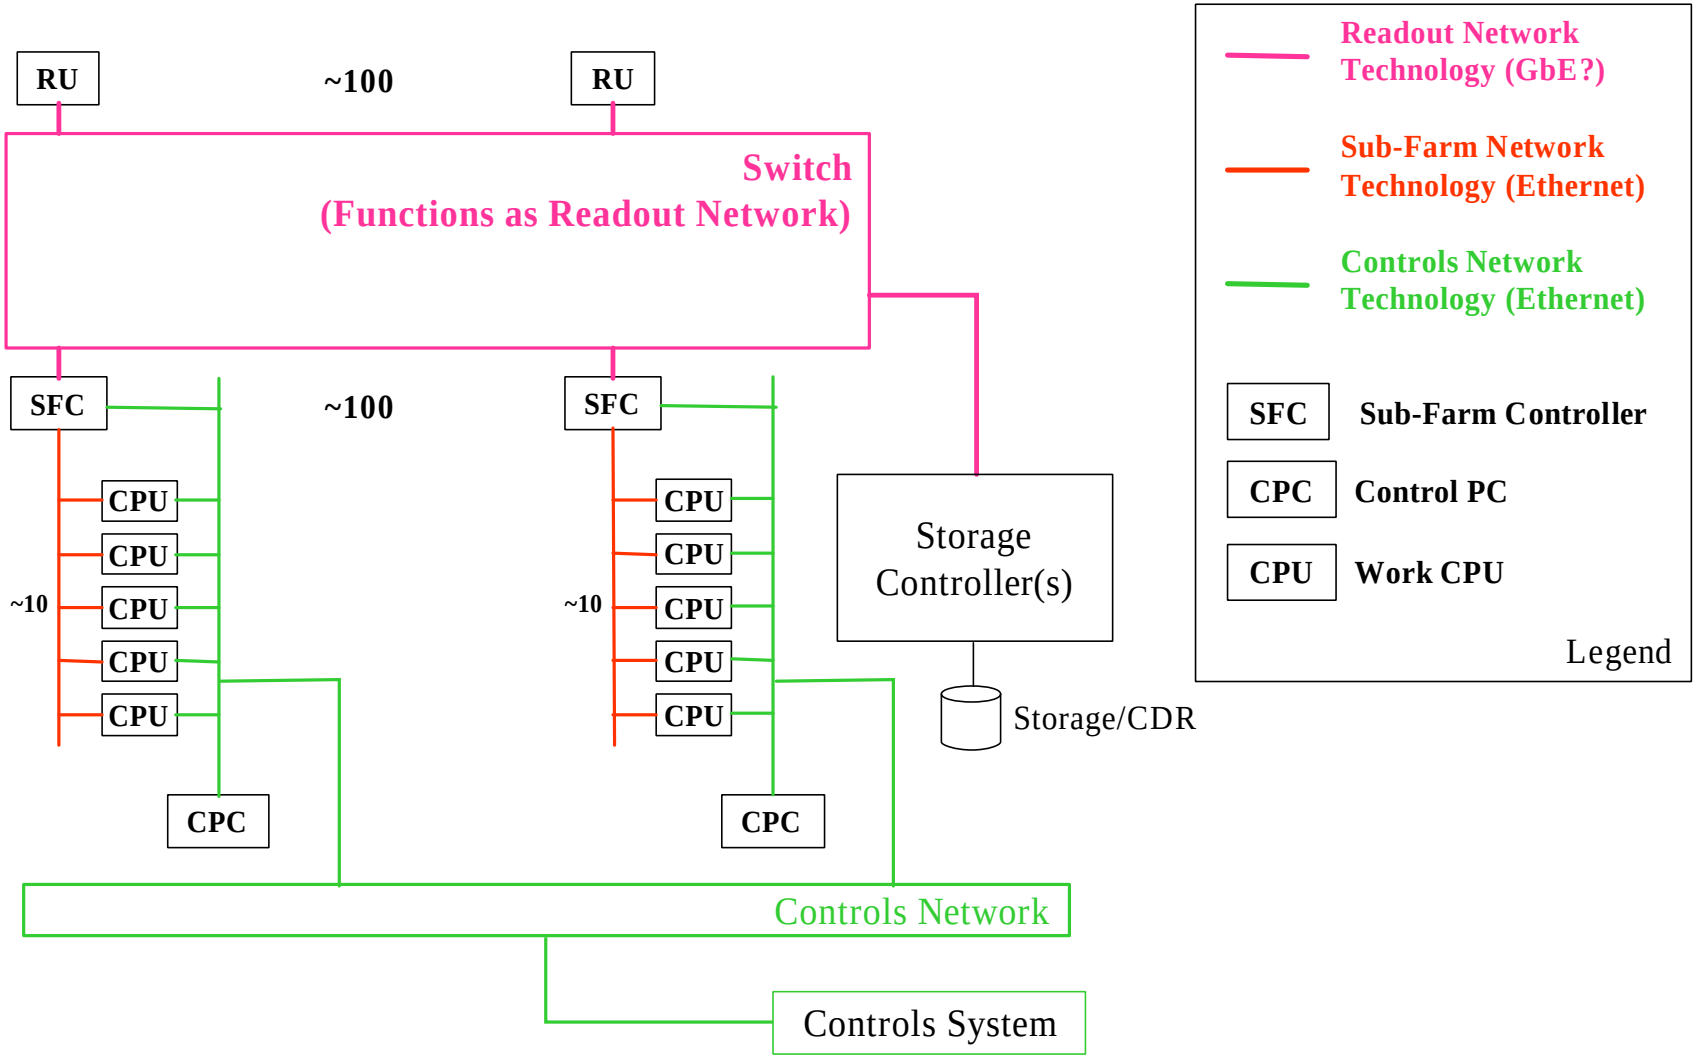

Event Building Protocol

- Pure “push” protocol
  - Each source (RU) of the event-builder sends data autonomously to a destination as soon as they are available
  - The destination is assigned through a static algorithm based on the event number. The identical algorithm runs in all sources (belonging to the same partition) -> The RUs belong to only one partition
  - The destinations (SFCs) receive the the fragments of the RUs and assemble them to form complete events. The SFCs will be configured to treat only fragments from the sources of the same partition -> The SFCs belong to only one partition
  - No communication (other than event fragments) between RUs and SFCs

Event Fragment Assembly

- The current idea is that the SFCs will be equipped with intelligent NICs (Network Interface Cards)
  - The NICs will handle all the fragment assembly to form complete events.
  - The host computer will only receive fully assembled events (or events flagged as erroneous) -> load on host reduced by ~100 (number of SFCs).

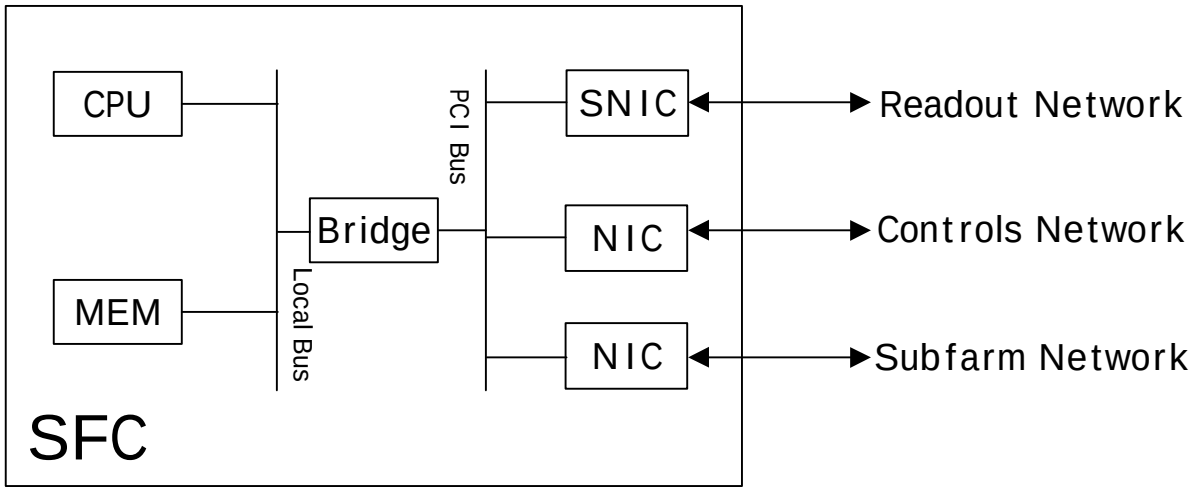

SFC Functionality

- Assemble complete events
- Monitor the state of the CPUs connected to it, implement dynamic load balancing among them and send raw data to 'free' CPUs
- Receive the accepted and reconstructed events from the CPUs and send them to the storage controller
- Monitor the buffer occupancy in the sub-farm and issue L1-Throttle in case buffers start getting full
- Provide monitoring information to ECS

Data Flow

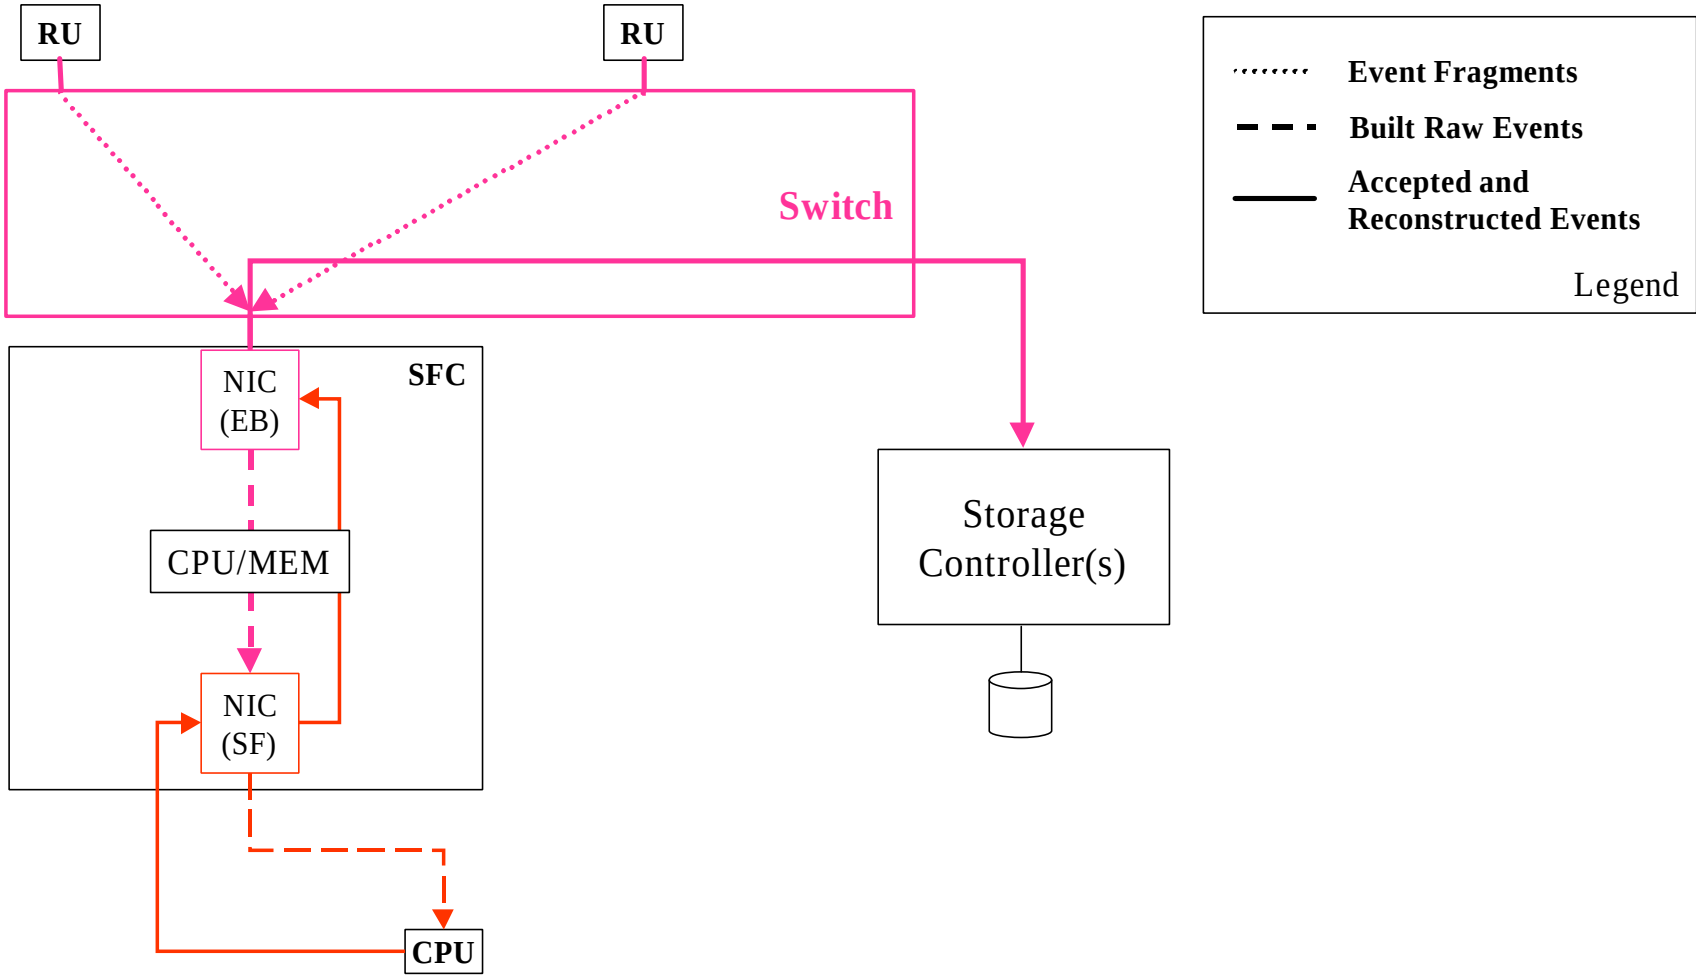
RU
RU
Switch
SFC
NIC
(EB)
CPU/MEM
NIC
(SF)
CPU
Storage
Controller(s)
Event Fragments
Built Raw Events
Accepted and
Reconstructed Events
Legend

Readout Network Technologies

Limit speed of end-node connections to ~1 Gb/s

Current Candidate Technologies

- ATM
  - 'no' smart NICs
  - enough performance
  - **too expensive**
- Myrinet
  - smart NICs exist
  - works in principle
  - cheap switches, ~expensive interfaces
  - No buffering in switches
  - **have to cascade switches -> have to build FIFOs to put between levels of switches (scalability)**
- Gb Ethernet
  - under study
  - has smart NICs
  - prob. works
  - cheap
  - Prob. have to cascade switches, should improve with 10GbE

Other Networks

- Data Network
  - ➢ Cheapest possible compatible with performance requirements
    - ↳ Input rate into CPU: ~5 MB/s
    - ↳ Output rate from CPU: ~25 kB/s
  - ➜ Fast Ethernet should be sufficient
- Controls Network
  - ➢ Cheapest possible (prob. Fast Ethernet)

Open Issues

- Non Zero-suppressed readout
  - Memory on smart NICs is limited (currently 1 MB, including code)
  - Possible solution could be to do fragment assembly in the SFC itself? (Not too nice…)