

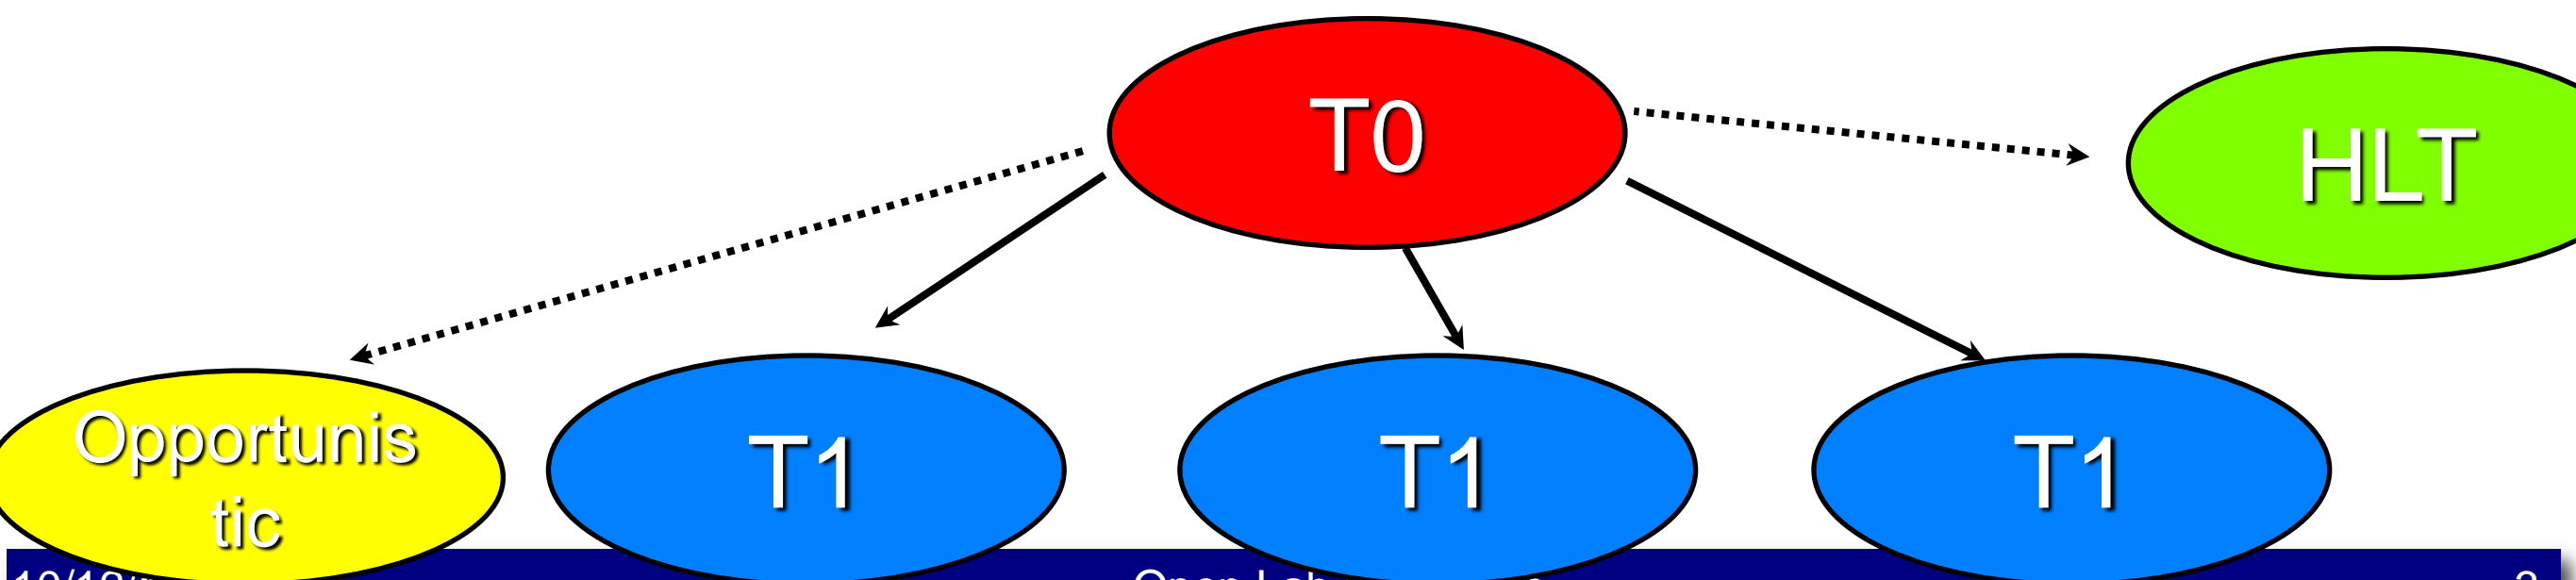
CMS Input for Networking

Daniele Bonacorsi/Ian Fisk/Maria Girone

- ▶ The network has improved and allowed PBs of LHC data per month to be delivered to analysis facilities from single, often distant, repositories. Distributed facilities can be effectively used by the program
- ▶ A good event data model and optimized read-ahead of IO has allowed CMS to have high CPU efficiency even if the data is served over the wide area to a running application.
- ▶ This opens a lot of possibility to utilize resources coherently even when they are separated over long distances

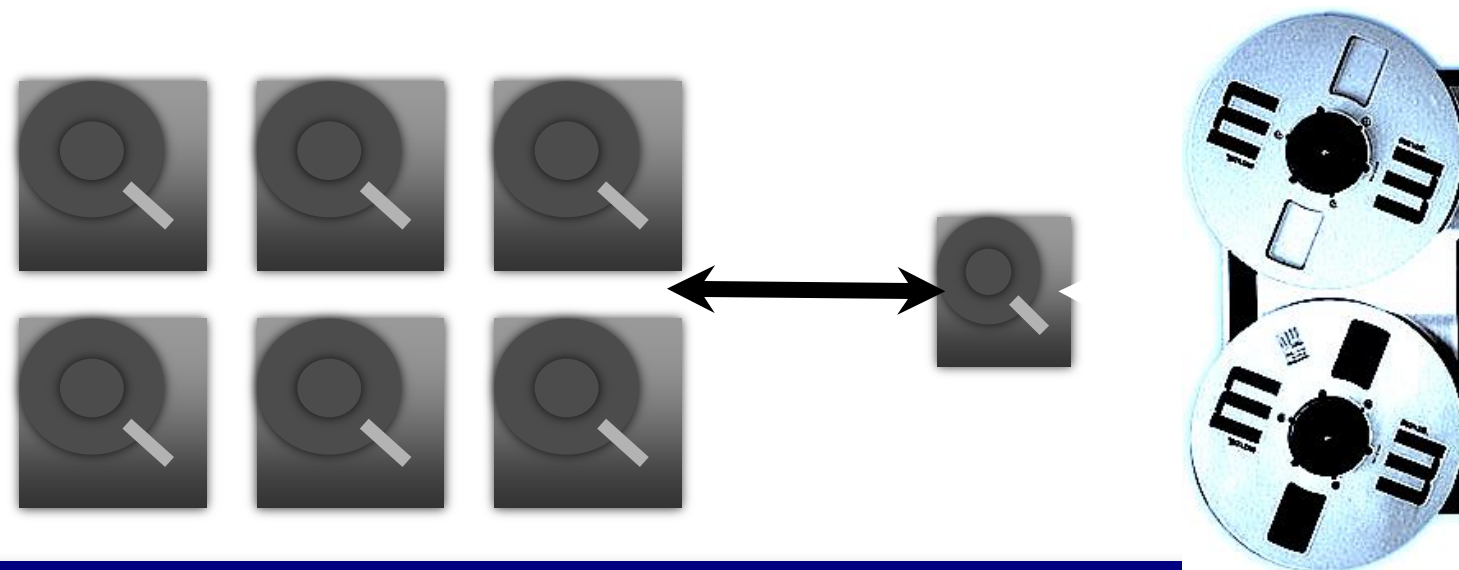
The OPN

- ▶ The Tier-1 centers have about 45% of the total disk storage in CMS, and about 30% of the CPU
- ▶ In 2015 there will be a close coupling with the Tier-0 as Tier-1s are expected to process with prompt reconstruction about half of the data
- ▶ We would like to see the OPN evolve to allow the Tier-0 and Tier-1s to work as a **single** processing center
- ▶ CMS has requested a better connection between worker nodes and the OPN to shared storage. Potentially a use case for IPV6 and an OPN upgrade



Separating Services

- ▶ CMS is in the process of separating the archival and disk storage solutions at Tier-1 sites.
- ▶ Tier-1s will have a large disk pool and a tape archive
- ▶ The tape archive will be able to accept data from multiple places, and stage data to local and remote disk storage
- ▶ Eventually we may consolidate to a smaller number of archival services
- ▶ This very organized data transfers will need reliable and predicable networking. Potentially a case for network reservation



- ▶ Before run 1 a brief network requirements document was prepared, it said a Tier-2 should be connected
 - ▶ Minimally at 1GB/s
 - ▶ Nominally at 5Gb/s
 - ▶ And leadership facilities 10Gb/s
- ▶ Looking forward to **Run2**
 - ▶ Minimally sites should be at 2Gbs/
 - ▶ Nominal connectivity is 10Gb/s
 - ▶ Leadership is 20Gb/s-100Gb/s
- ▶ This represents an interesting challenge because the best and worst connected Tier-2 sites will vary by a factor of 50
 - ▶ Will lead to different capabilities for each



Shared Centers and Connectivity

- ▶ Once we have sites connected at 100Gb/s, then there is little difference between the bandwidth on the WAN and LAN
 - ▶ This is being tested with the CERN/Wigner Centers
 - ▶ We expect to be able to share processing resources between Tier-1 and Tier-2 and use the disk from either
 - ▶ We expect to consolidate Tier-2 sites into large virtual sites where the resources will be a separated centers, but we see it as one site.
- ▶ Will save on local and central operations and reduce the number of service instances