LHC Computir

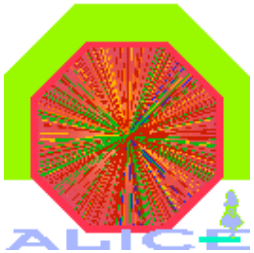

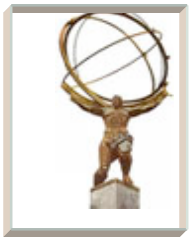


Ian Fisk
CHEP Conference
Victoria, Canada
September 3, 2007

CE, ATLAS, CMS and LHCB are
by LHC Computing

ta expected in late July of 2008

tive preparations for computing for 5-6
rs

increase the proposed scale of

stribution

ata Transfer

ata Access and Analysis

C experiments have enjoyed an
precedented level of support from grid

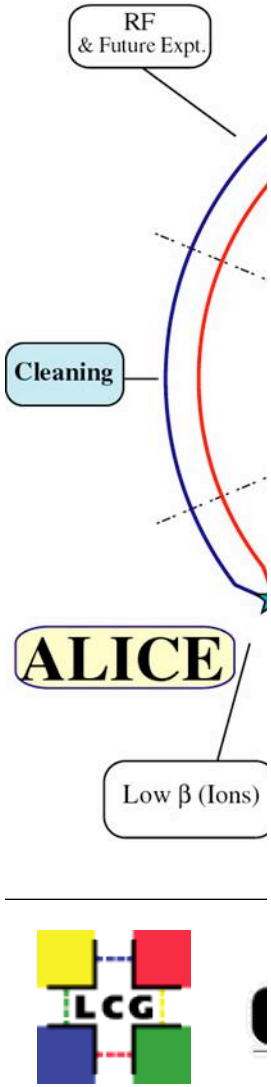

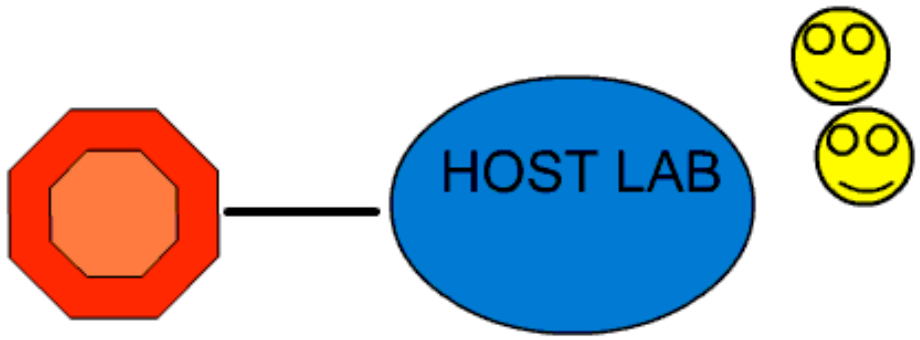

beginning the computing was centralized

ments began to develop distributed computing m

0 center at CERN
for prompt
nstruction, data
ve, low latency

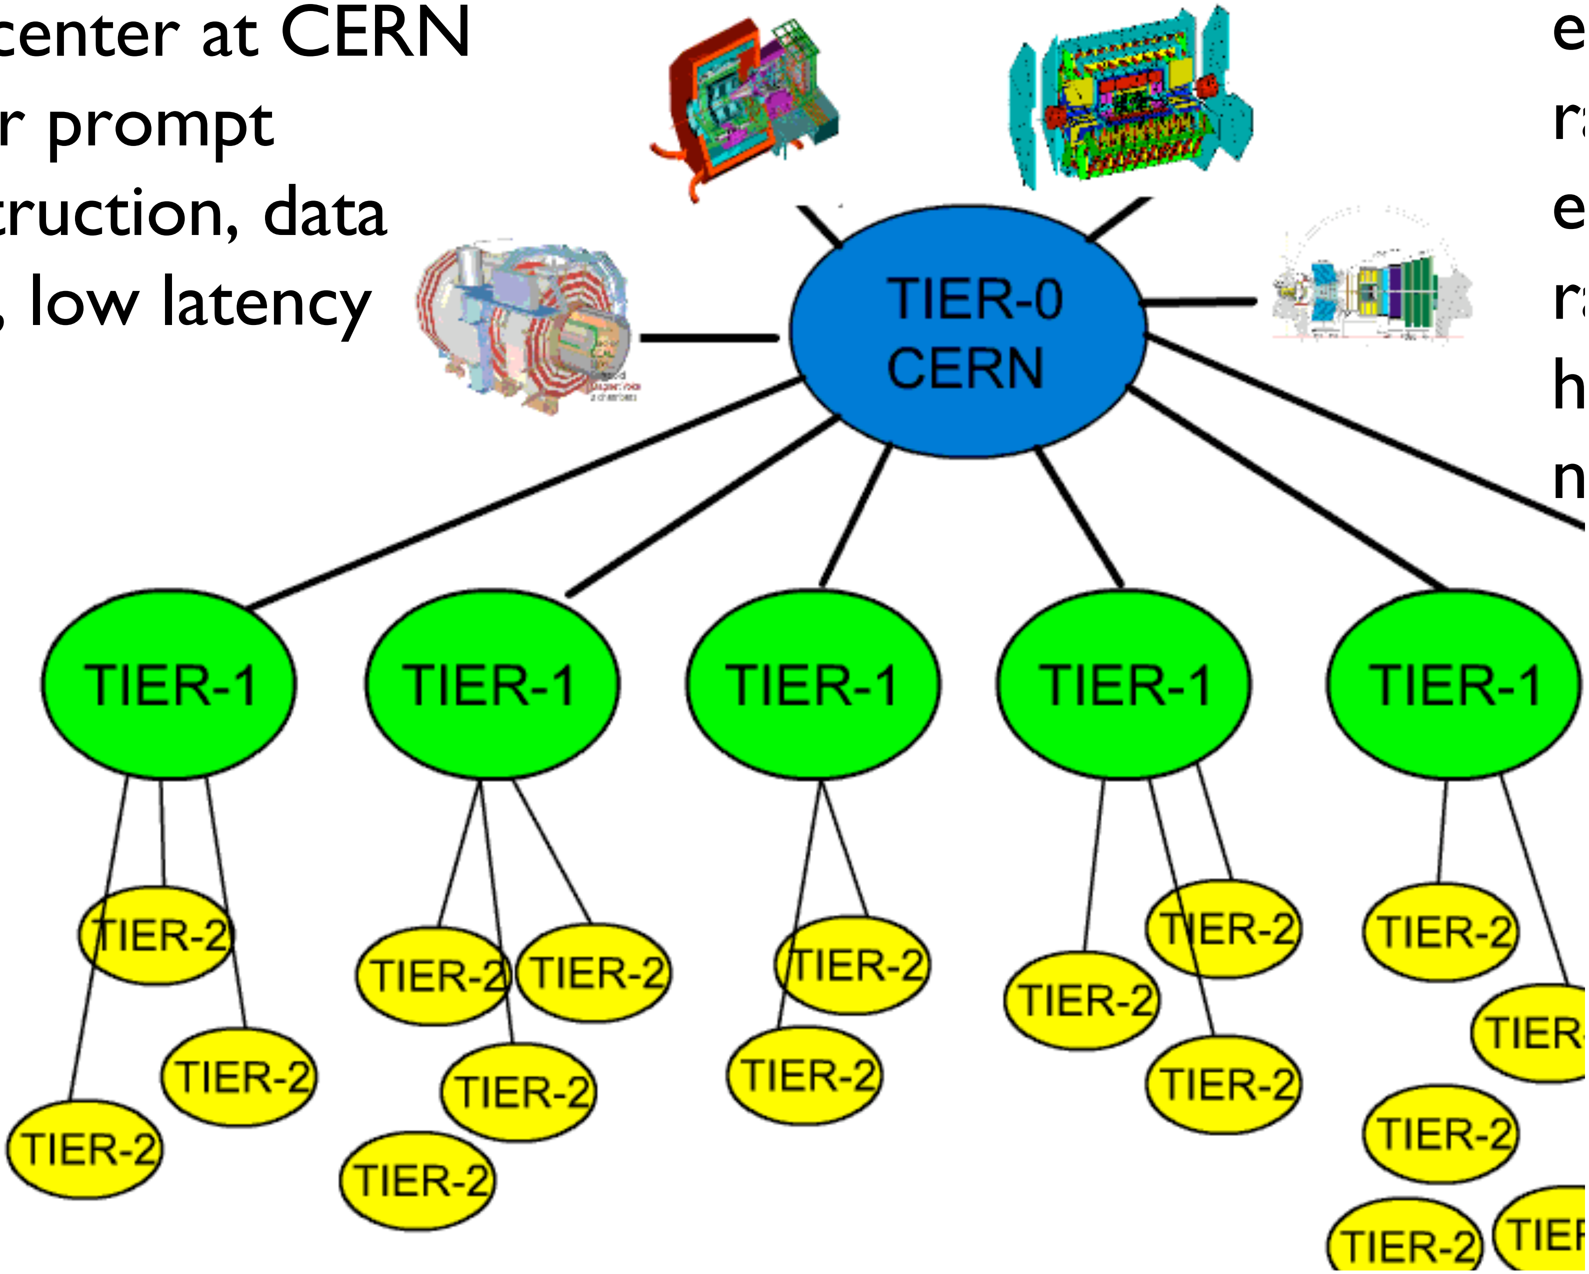

N
e
r
e
r
h
n

nters are primarily at
abs or large universites
econstruction
ping/skimming
serving
iving of simulation

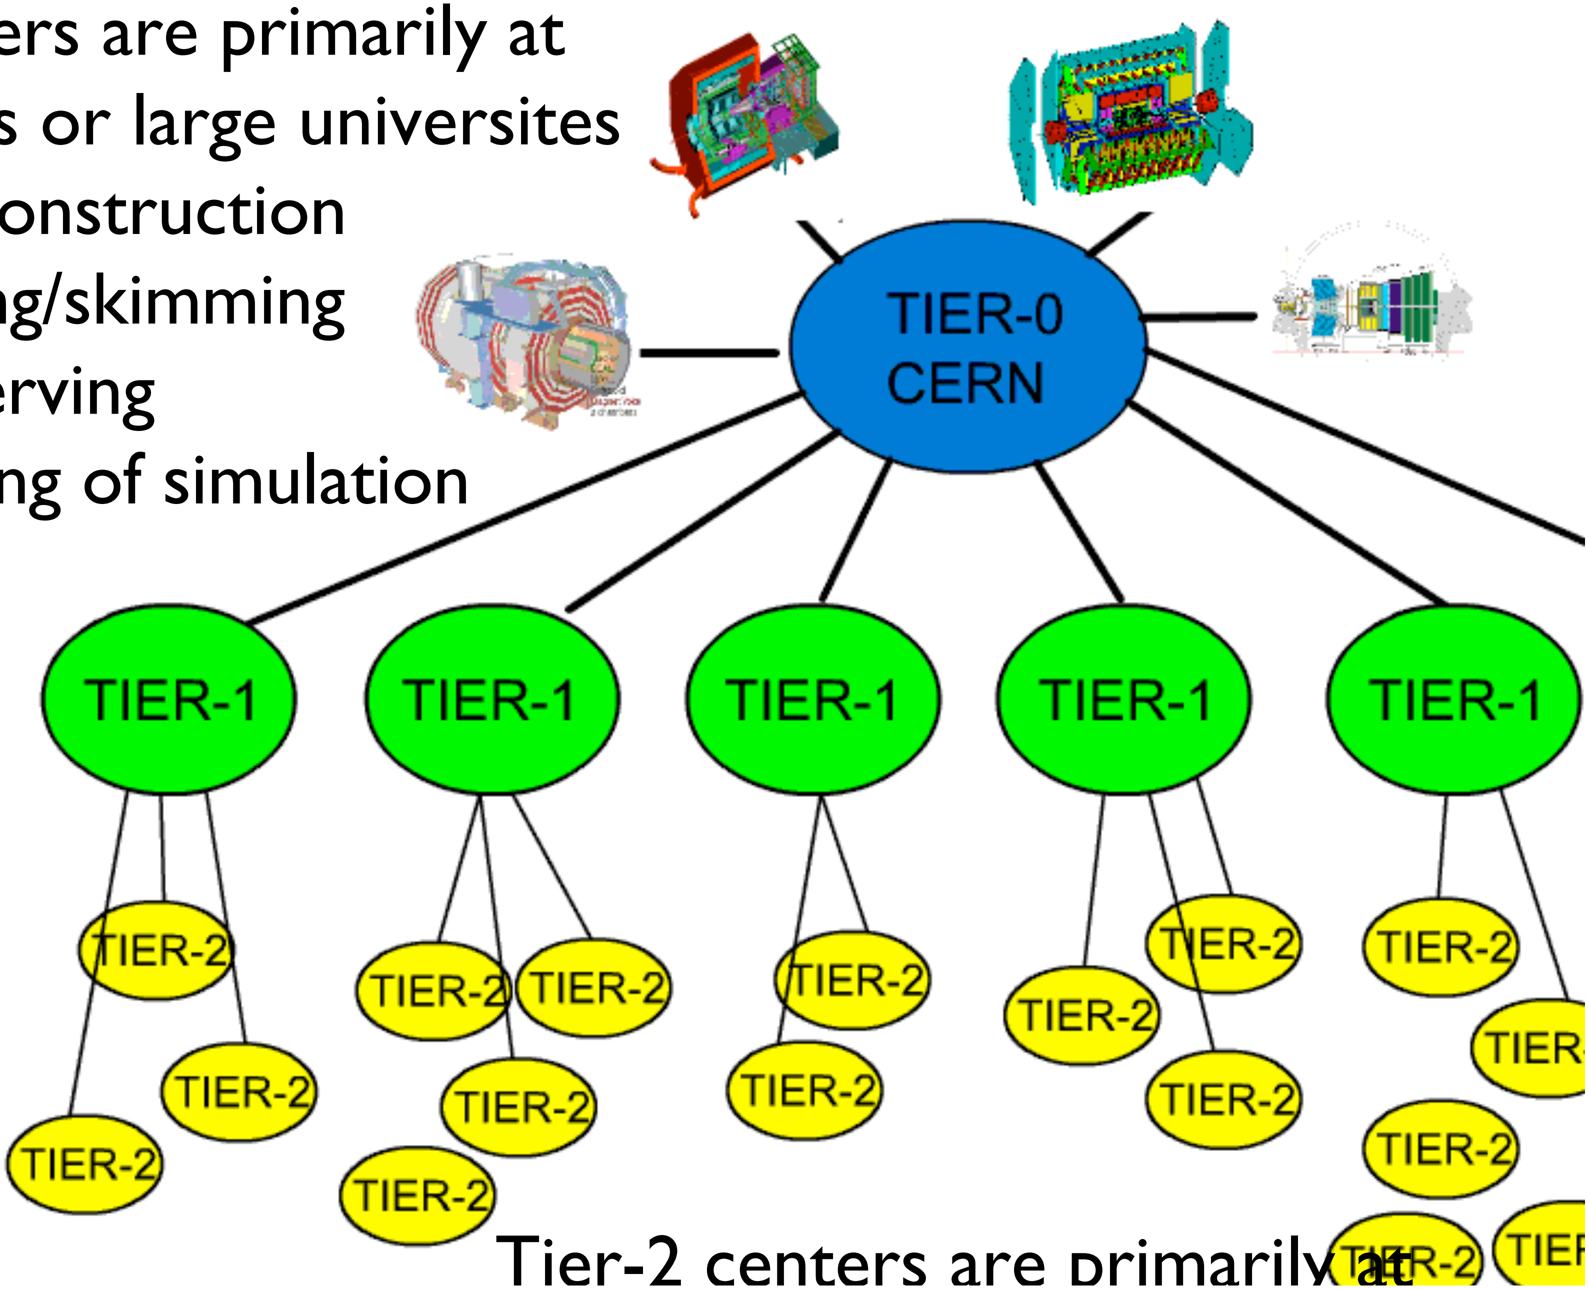

Tier-2 centers are primarily at

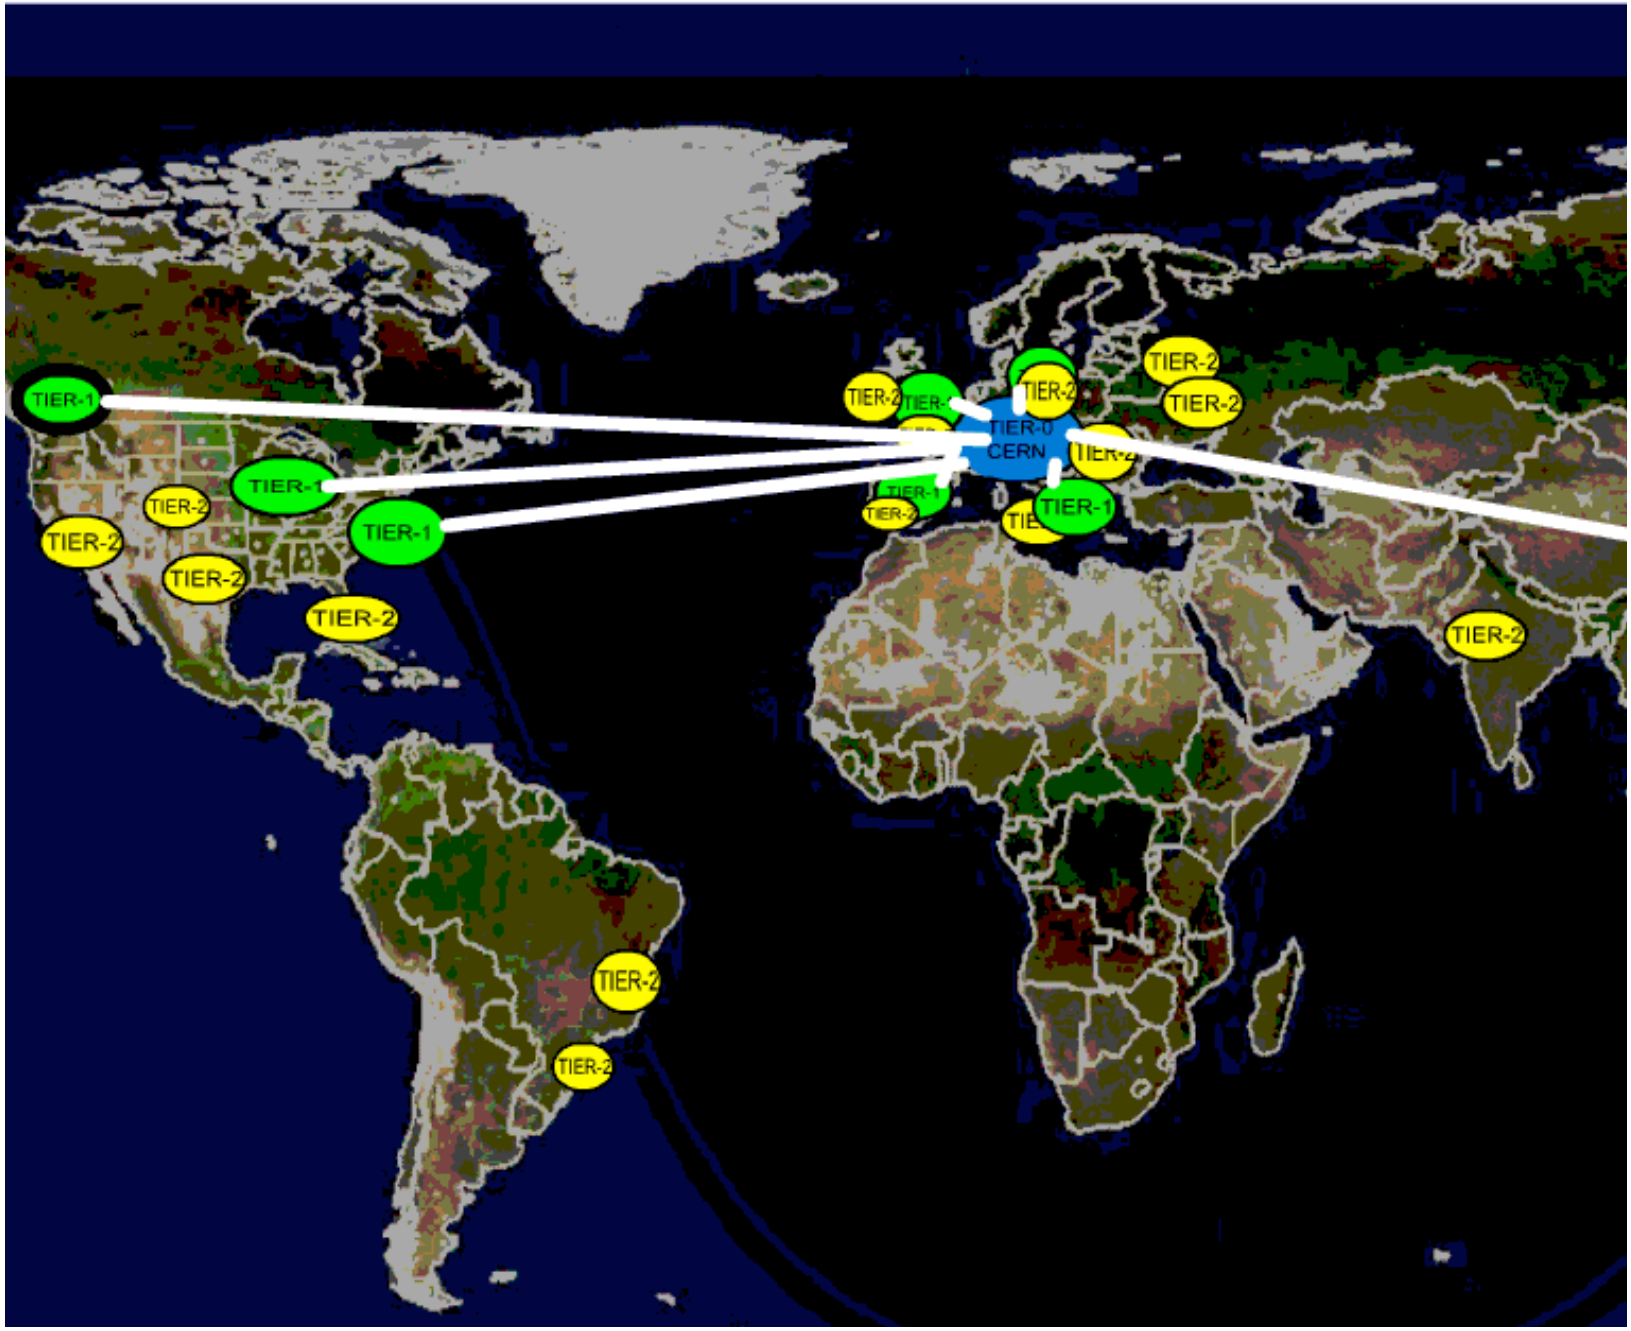
TIER-1
TIER-1
TIER-1
TIER-1
TIER-1
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-2
TIER-0
CERN

riginal CDF distributed analysis facility the servic
solution, authentication method, etc. were dev
:

e services were unique and the early farms were

LHC with the advent of grid services and interfa
ols was chosen

tire computing facilities could be shared betwee
cess would be possible, facility support would be
uld be improved…

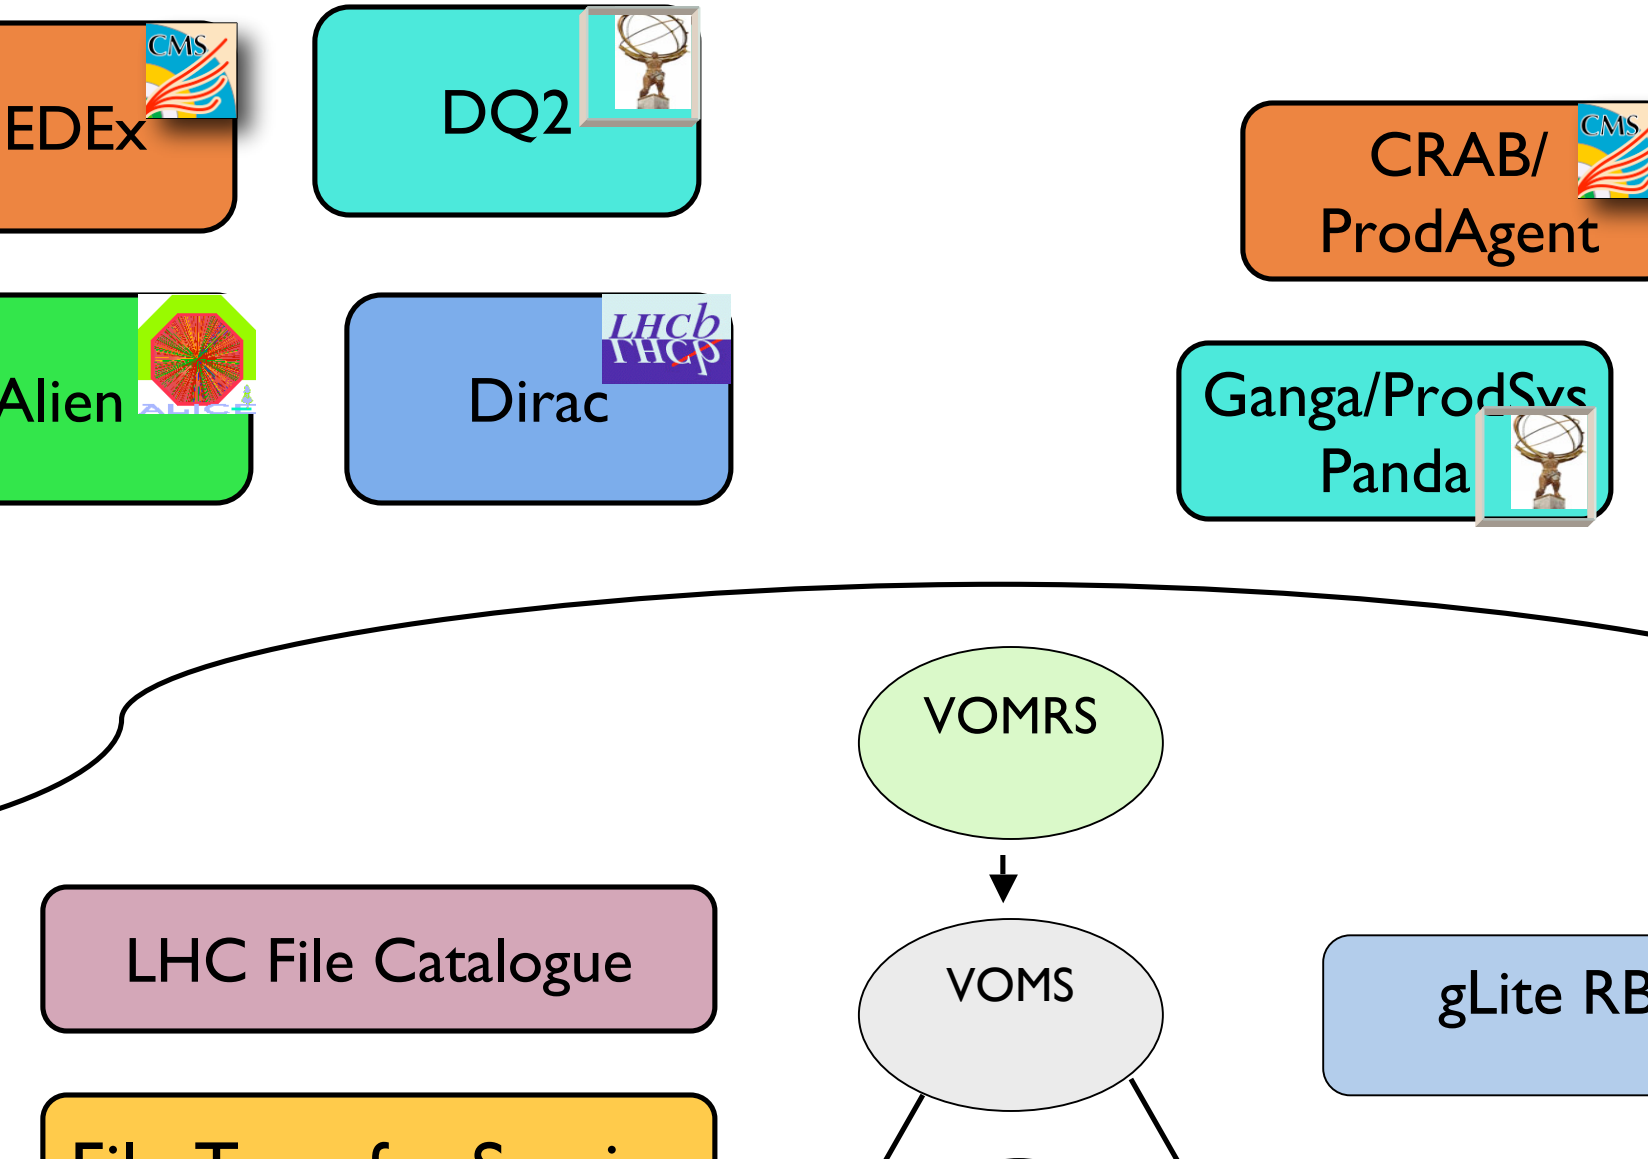
EDEx
DQ2
CRAB/
ProdAgent
Alien
Dirac
Ganga/ProdSys
Panda
VOMRS
VOMS
LHC File Catalogue
gLite RB

el of distribution and the number of services req
to check the health of the globally distributed sy
LCG has developed a series of Site Availability M
ries of automatically submitted and tracked tests
lidate the processing services all the way down
lidate storage services
formation systems
sts run every few hours and results are tracked
r VOs have begun to introduce their own tests
rify the experiment workflows within the SAM f

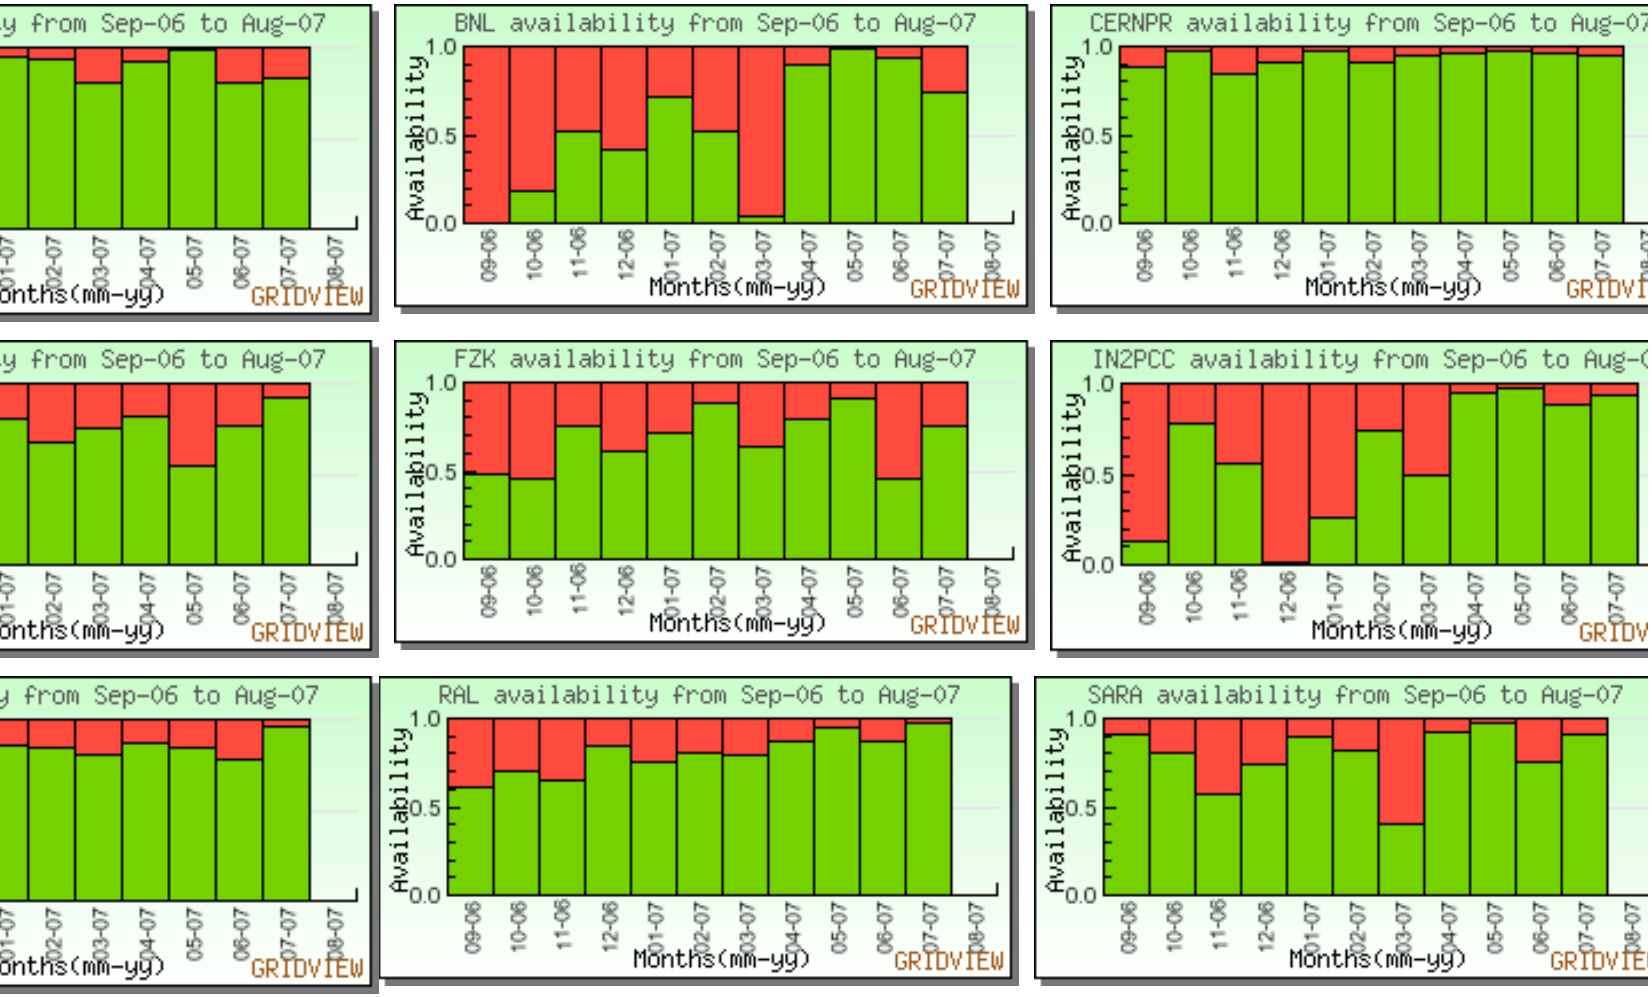

rly areas for improvement

derlying services need to end up in the much hi

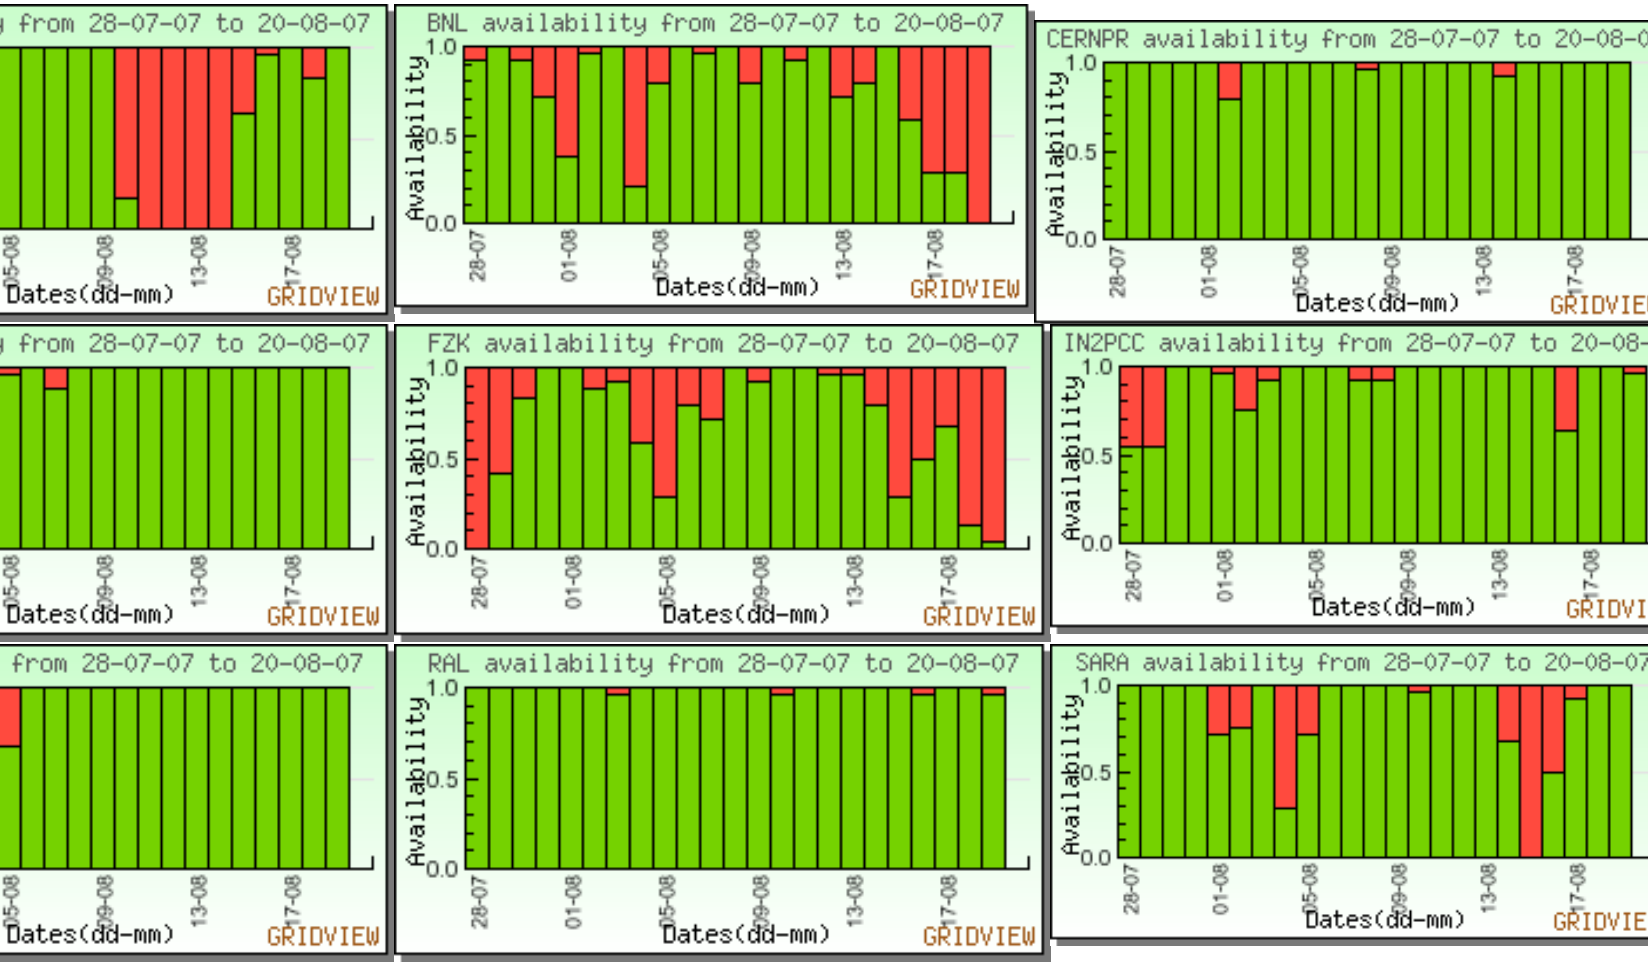

rly areas for improvement

derlying services need to end up in the much hig

ies in the data management functionality for the

experiments have services that sit on top of the

fine the mappings between events, files, and even

dataset is typically defined as a collection of logic

e files are immutable and can be replicated betv

TLAS and LHCb both use the LHC File Catalog

MS uses a TFC (Trivial File Catalog) technique si

Babar, where the storage element namespace is

gical file names to physical files names without a

periment data management systems drive the re

ls to define datasets tend to be experiment spe
nality is driven physics requirements and choices
ted

an be very flexible like ALICE's Event TAG servi
place cuts and receive a new list of files for that

- Datasets are more dynamic

LHCb the specialized data sample lends itself to
ripped datasets that are centrally produced

- Simplifies the definitions and access

MS is in-between with datasets being defined and
ookkeeping service, but operations and users car
needed

has the largest nominal CERN to Tier-1 transfer

ests this spring reached ~75% of the eventual tar

ccessful use of 11 Tier-1centers, successful demo

d FTS

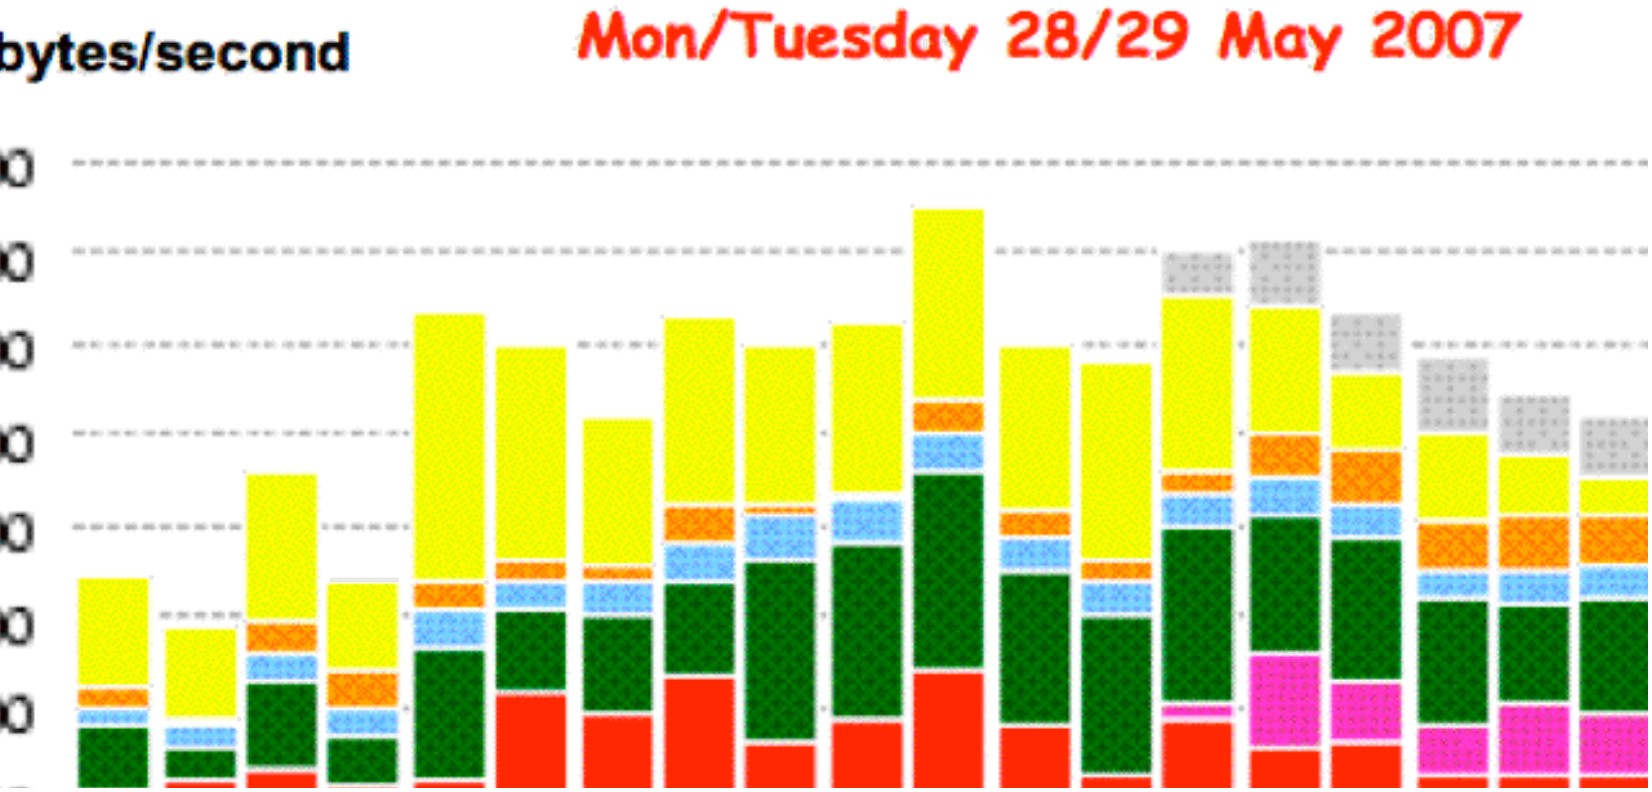

has the largest nominal CERN to Tier-1 transfer

ests this spring reached ~75% of the eventual tar

ccessful use of 11 Tier-1centers, successful demo

d FTS

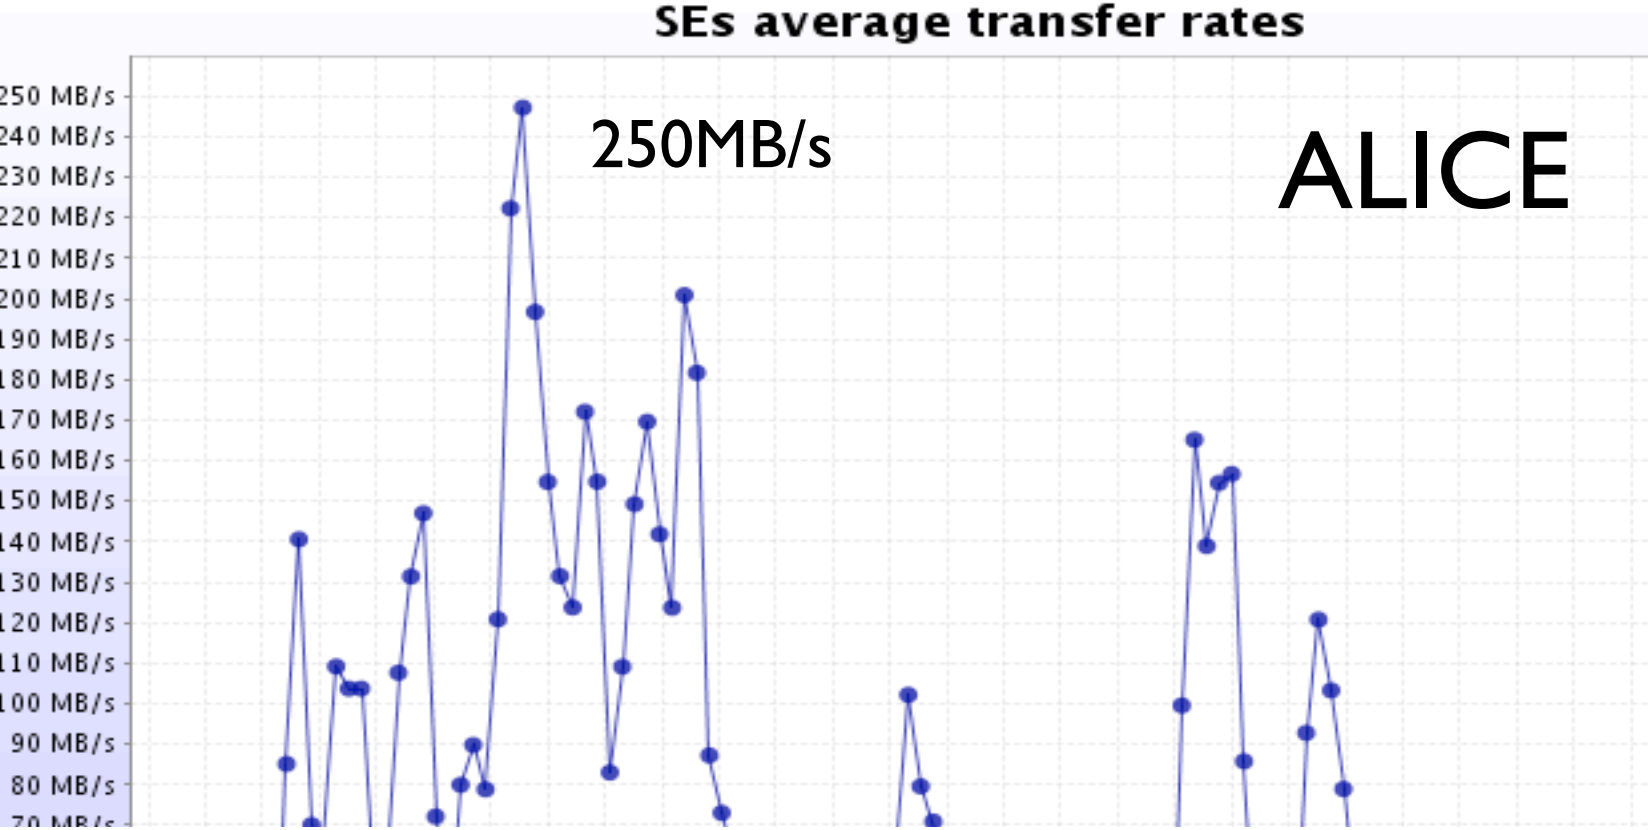

ects Tier-2 storage to be treated like a dynamic cache

r-2s can be updated with data from any Tier-1.

2008 data rates are expected in bursts of 50MB/s-500M

t below of data exported from FNAL to 21 Tier-2s

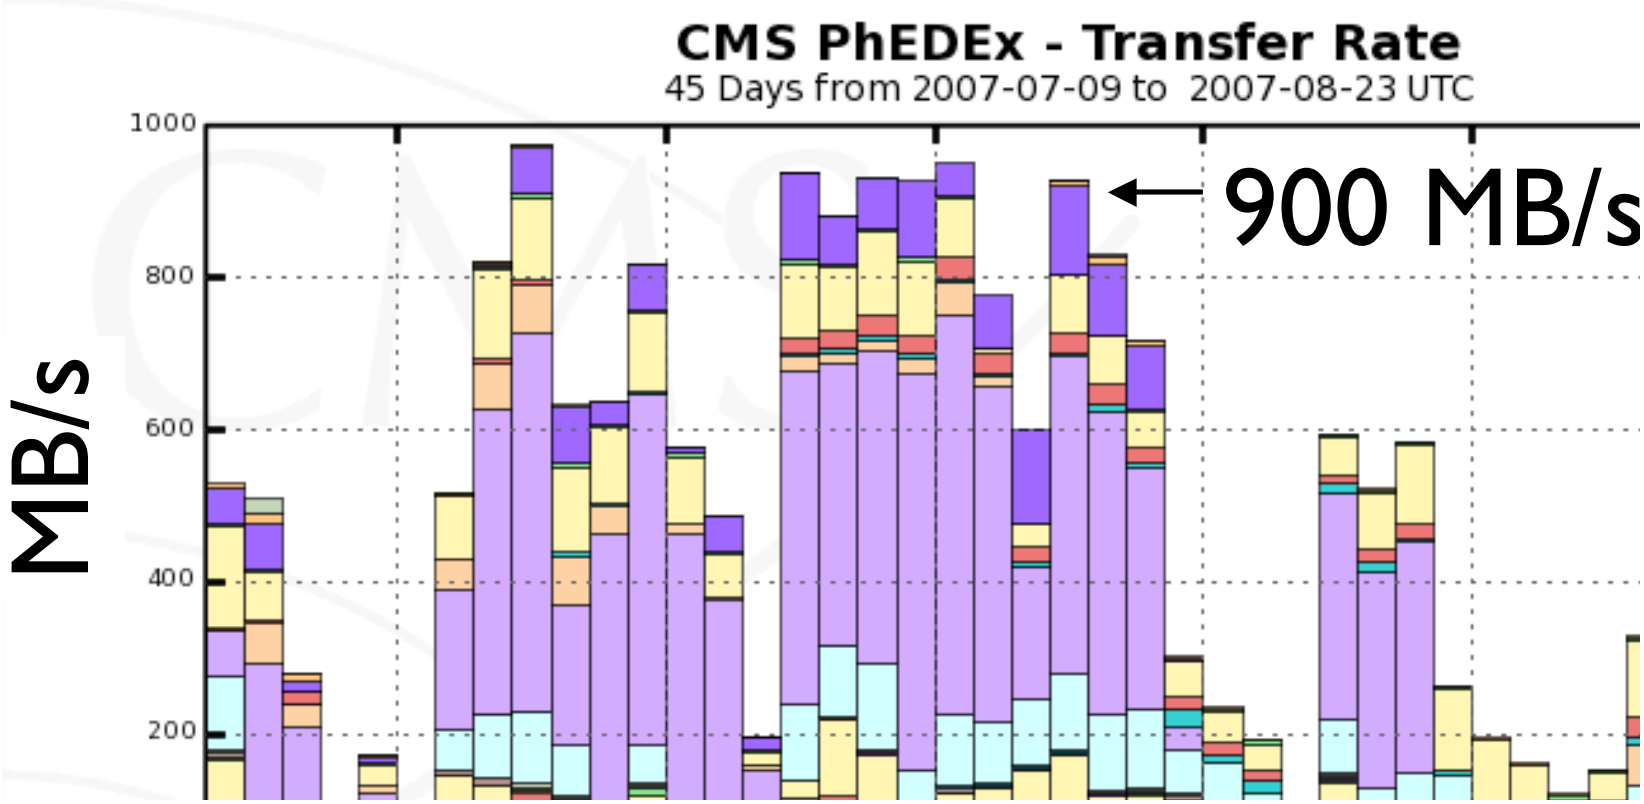

ess to applications has been a difficult area for LHC co
ioning

ge number of sites, CPUs, and large volume of data

rarchical mass storage

eed to be mindful of file size and rates of opening files

nplest solution, the mass storage system handles data s
ing to applications using an efficient local protocol (rfic

order to improve the performance for access and
ponsiveness ALICE has deployed PROOF and xrootd

Cb has implemented a pre-staging tool that will not rel
lications from the task queue until files are on disk

duction is an ideal candidate for distributed pro

rge output and CPU requirements but small inpu

plications. All four experiments are succeeding

CE

sites used

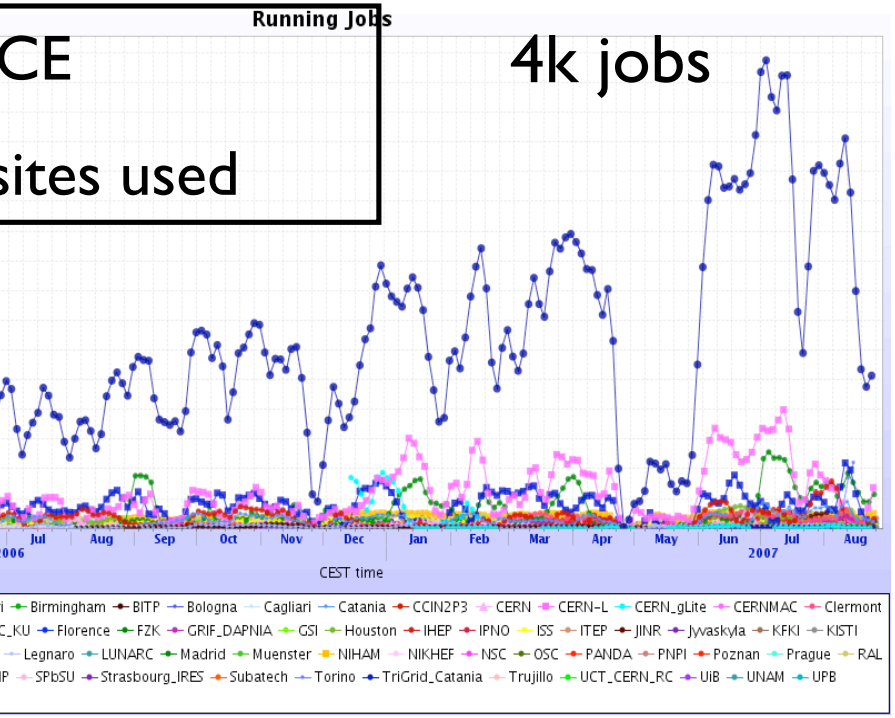

ATLAS

More than 100 sites used

CMS

170M Events

~180TB of data

40 sites used

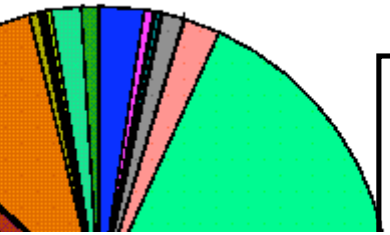

LHCb

375M Events

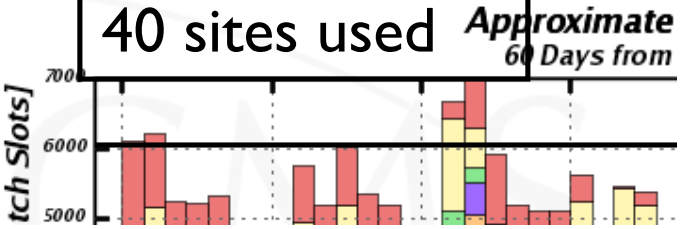

LICE and LHCb have developed pull based job su
n Production and Analysis

LAS uses pull for one of the work flow tools

entral
rvices

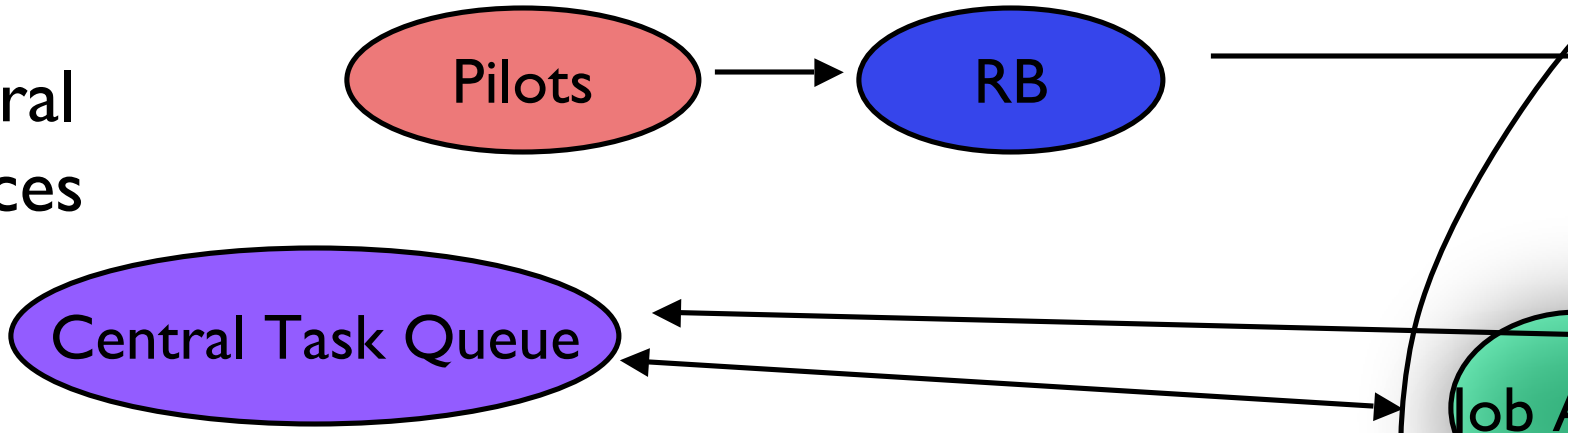

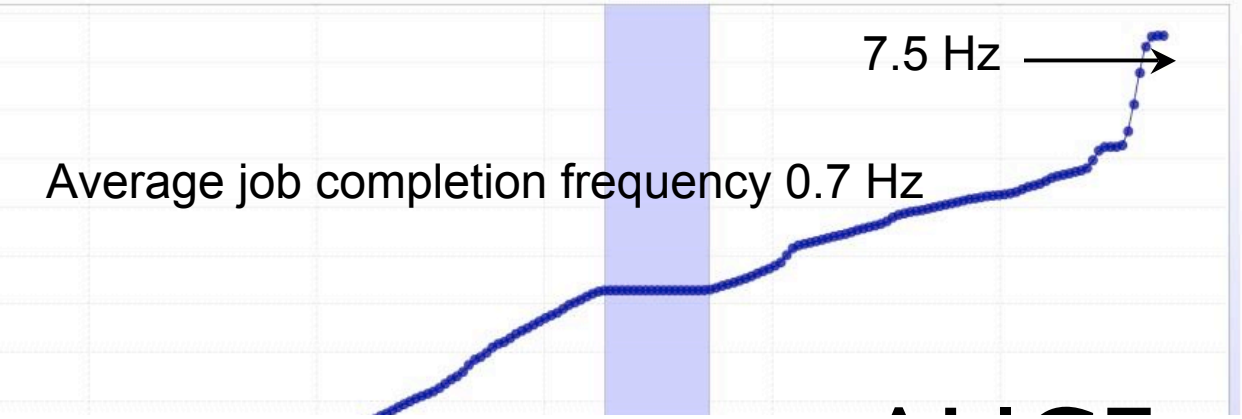

Average
The final
➡Cen
Central t

processing is more interesting need to match processir
s of data.

tems used in ATLAS and CMS are similar in the steps

nga and Panda in ATLAS and CRAB is CMS

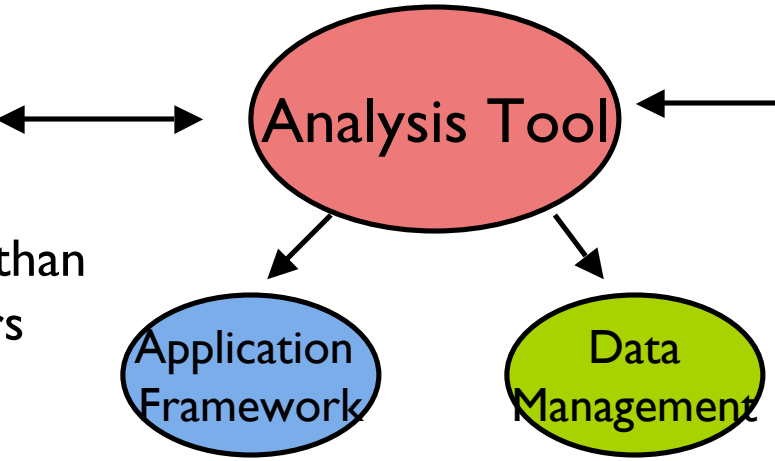

than
s

y CMS is averaging

per day

mate in 08 is > 100k

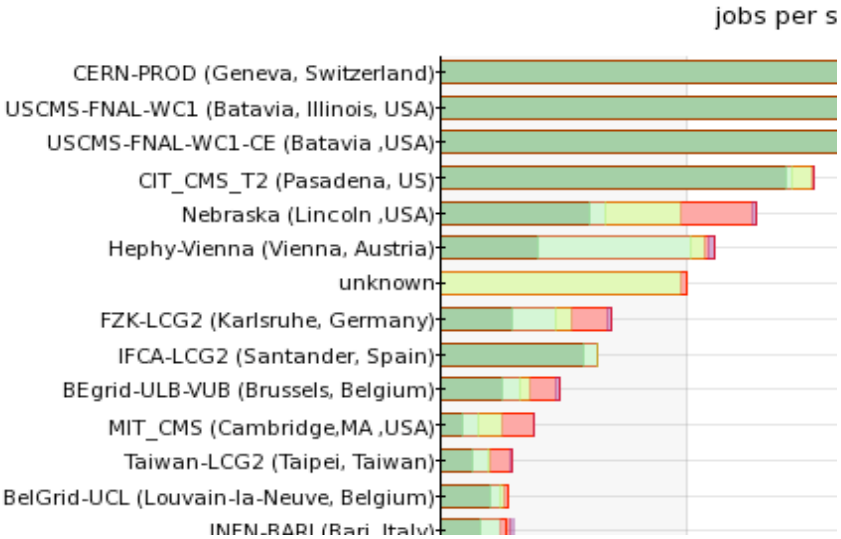

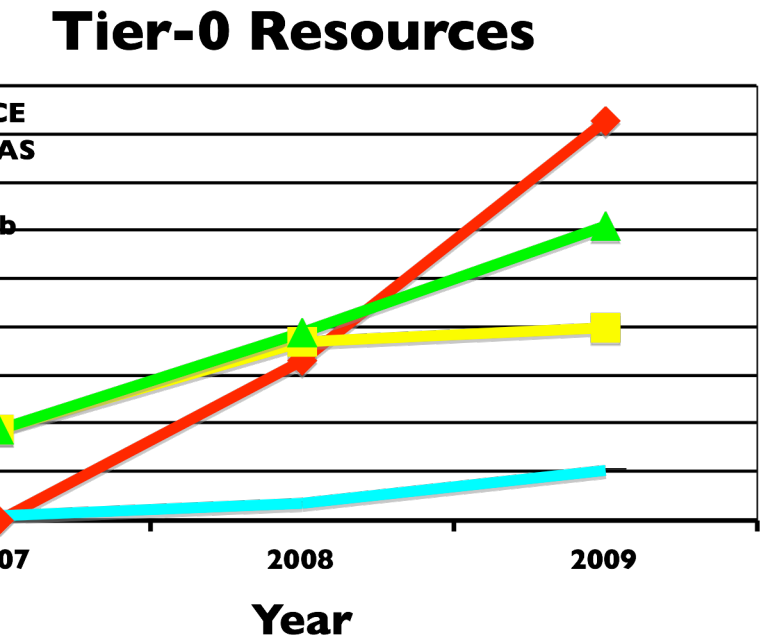

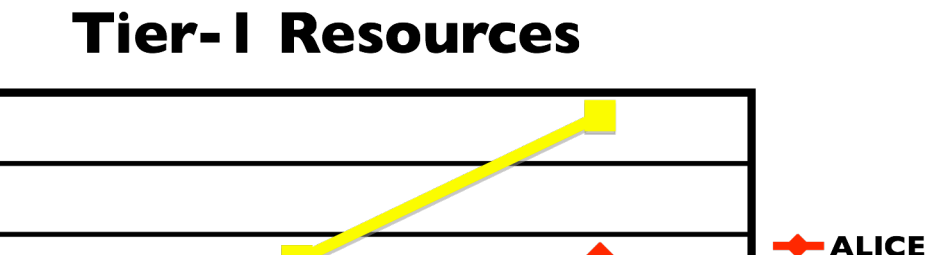

The total quantity of resources needs to the experiments ov

- Some of this ca Moore's law im
- While large the running farms t

In order to reach th of processes

- Node purchase and needed 2-3 them
- Node purchase

Tier-1 Disk Resources

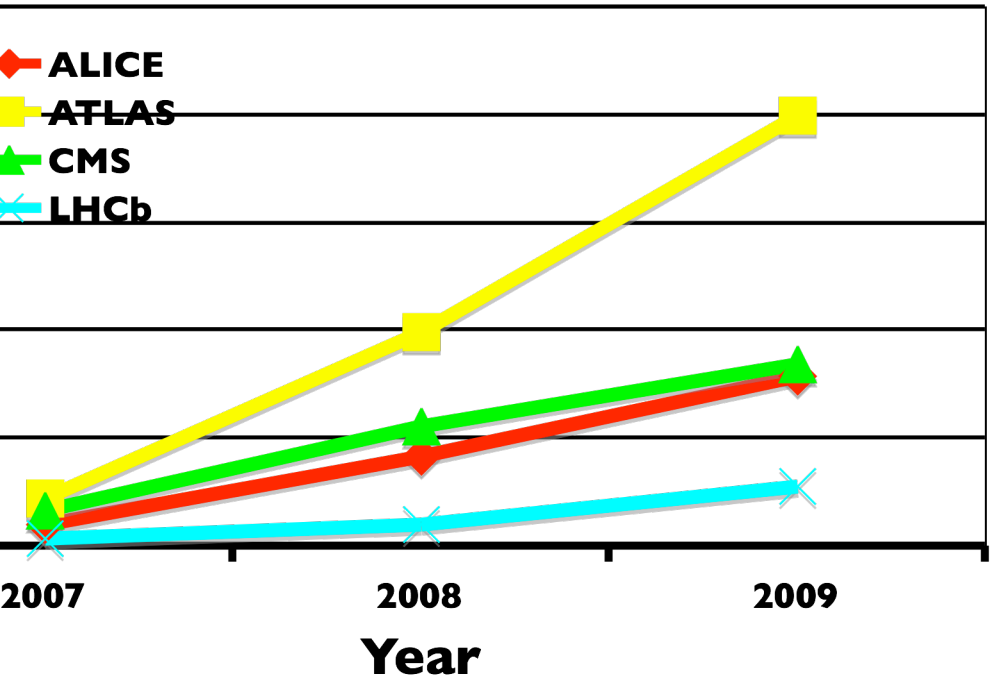

Disk ramp is a little

- The required accommodate improvement
  - There are a l examples of installations
  - Issues of facil scalability of

Tier-2 Disk Storage

Performance and st dependent on how

- Rely on exper

Tier-0 Tape Resources

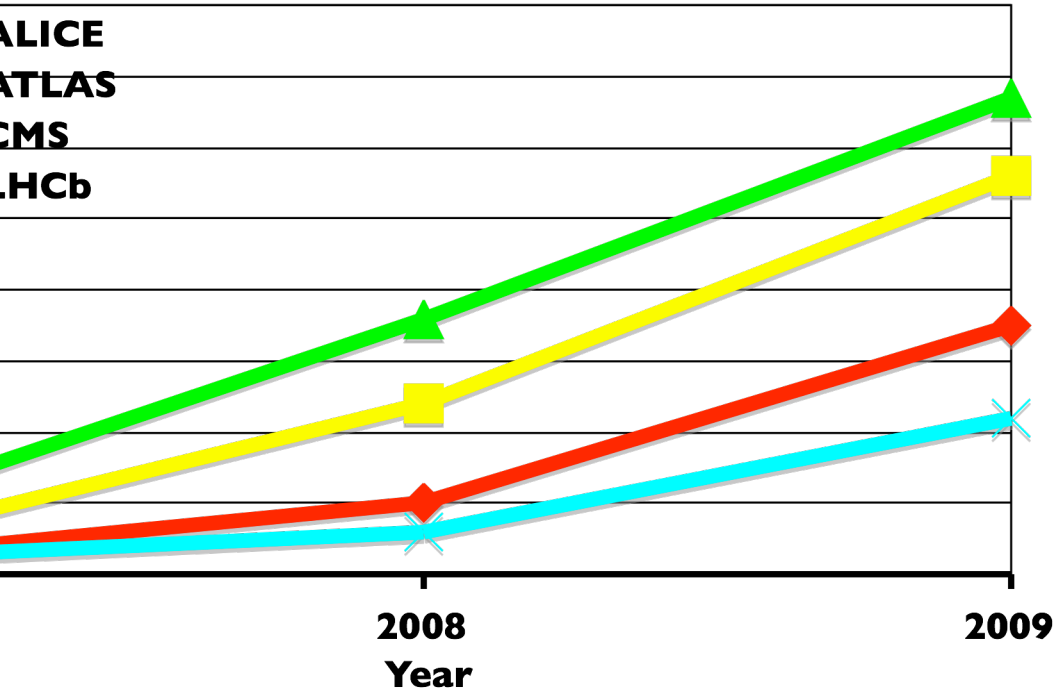

Tier-1 Tape Resources

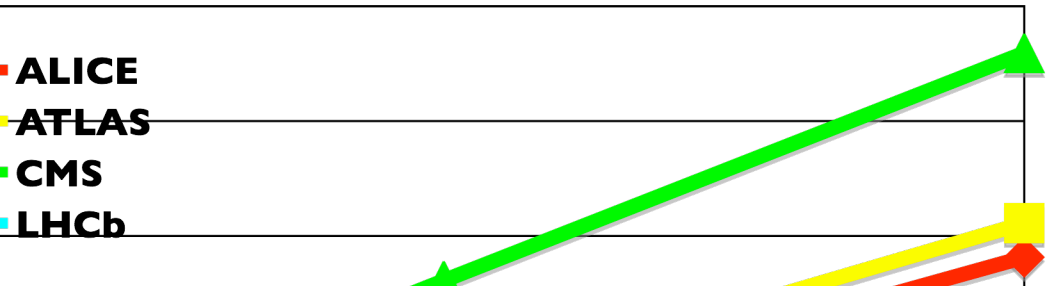

Tape resources
most scalable

- Robotic st
  handle larg

Also one of the
requires the lo
experience of

- Not all Tie
  experience

Most of the LH
to operate in t
many times reg

periments have begun demonstrating computing
xpected to be seen in running conditions

ansfers from CERN

sources utilized for simulated event production

work left in the final year of preparation

big increase in scale needed in facility infrastructu
e it routinely

er analysis access needs to ramp up

mplicated computing environment and we are st
erate it