Grids: Why and How (you might use them)

J. Templon, NIKHEF

VLVνT Workshop

NIKHEF

06 October 2003

Information I intend to transfer

- Why are Grids interesting? Grids are *solutions* so I will spend some time talking about the *problem* and show how Grids are relevant. ***Solutions should solve a problem.***
- What are Computational Grids?
- How are we (high-energy physicists) using Grids? What tools are available?
- How might they be of interest to Neutrino Telescopists?

Our Problem

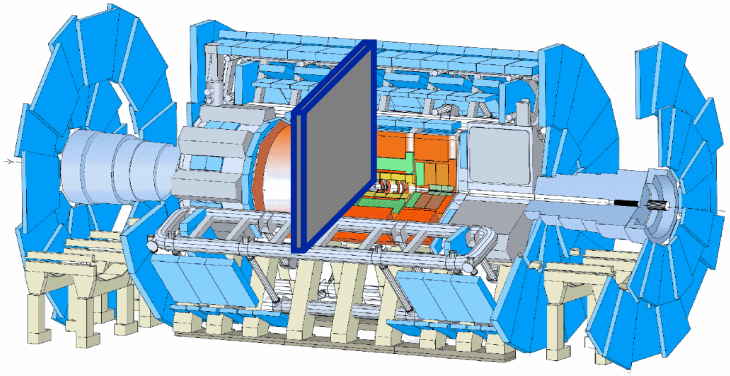

- Place event info on 3D map
- Trace trajectories through hits
- Assign type to each track
- Find particles you want
- Needle in a haystack!
- This is “relatively easy” case

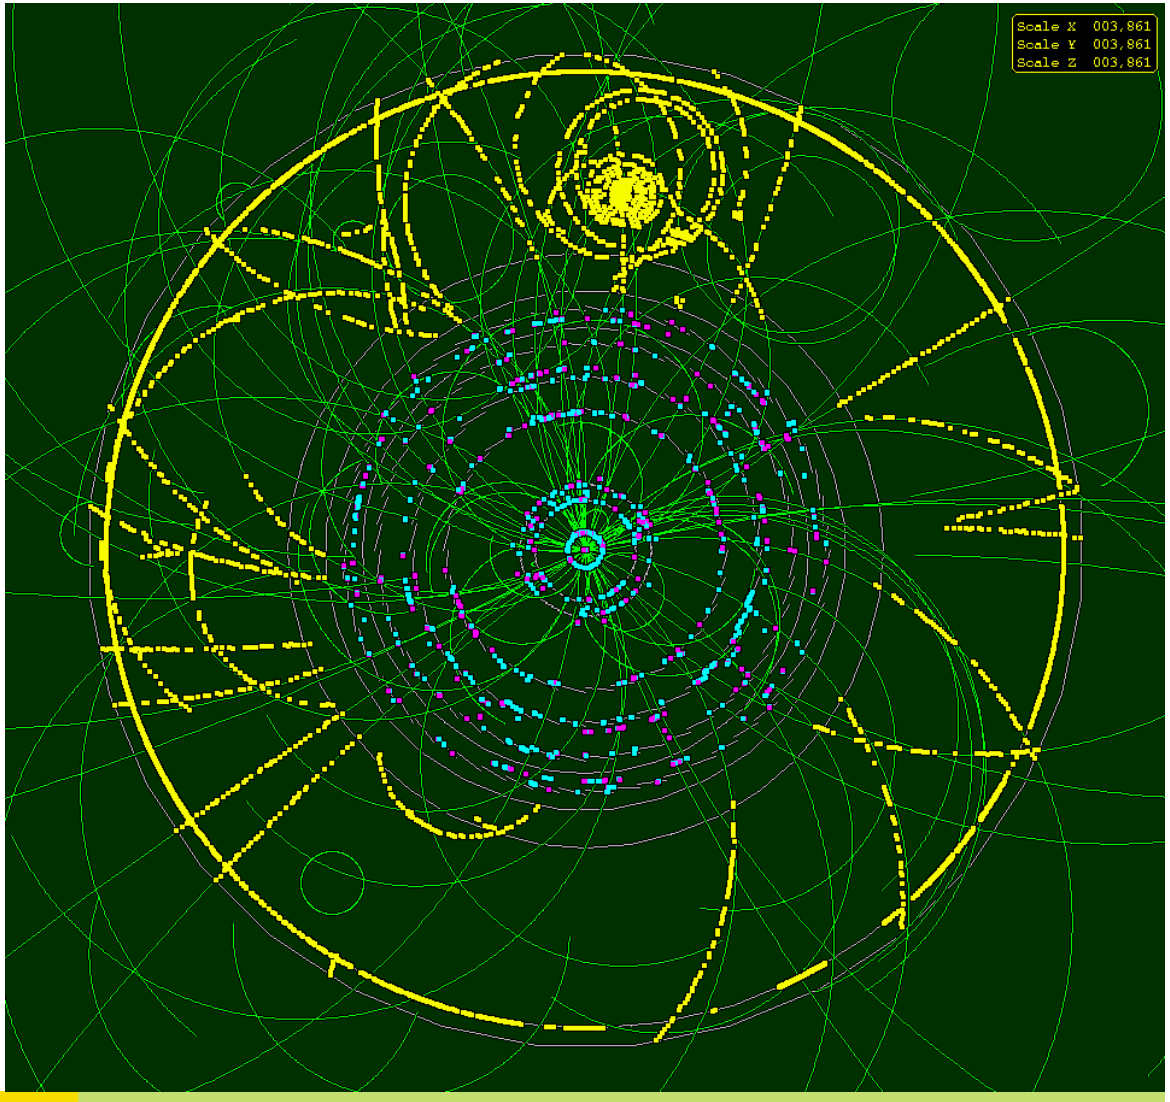

More complex example

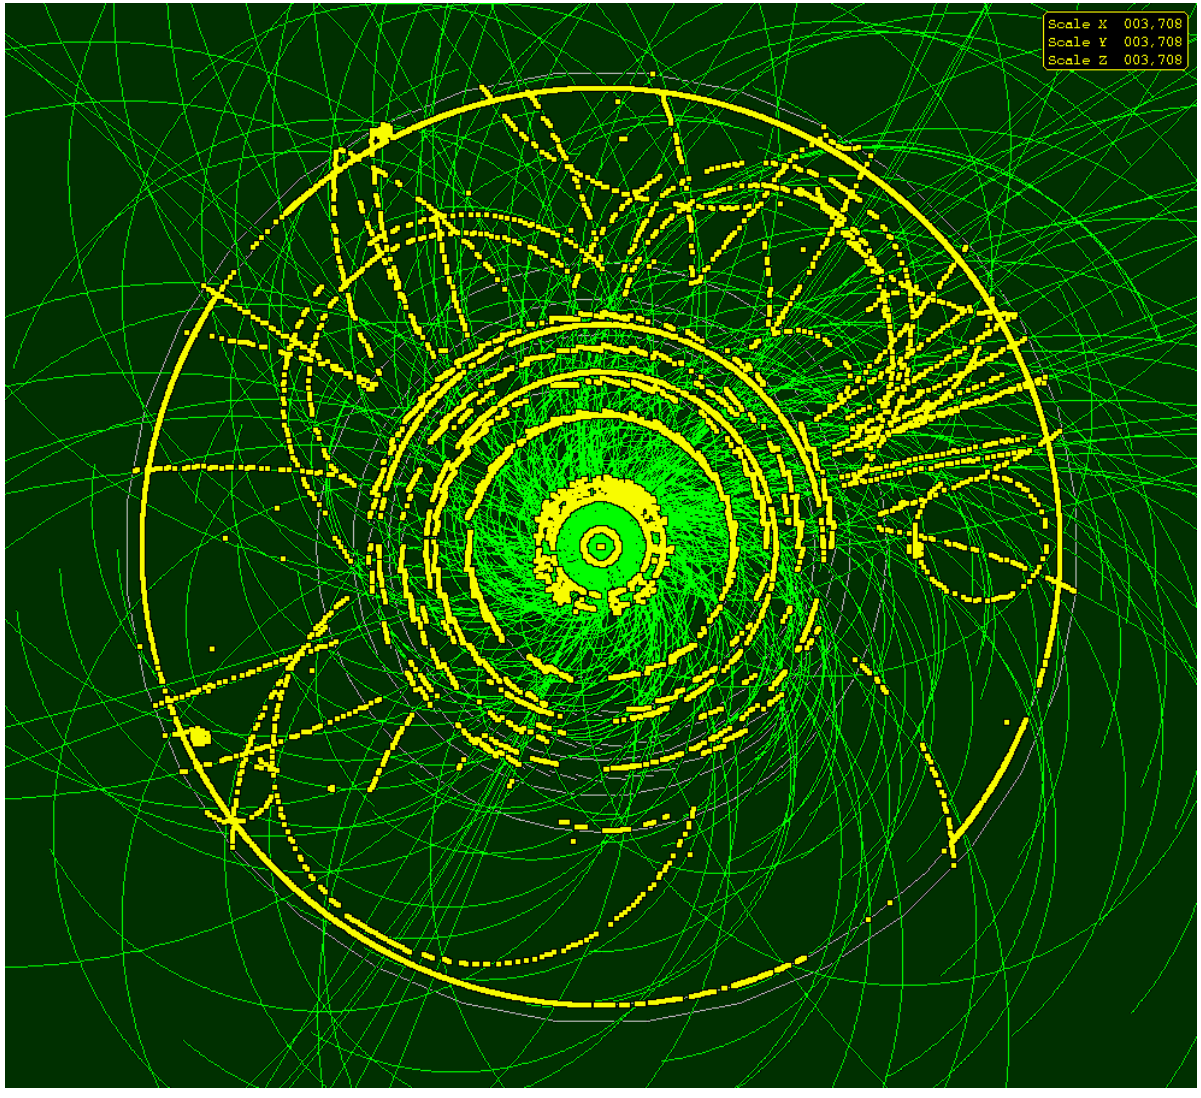

Data Handling and Computation for Physics Analysis

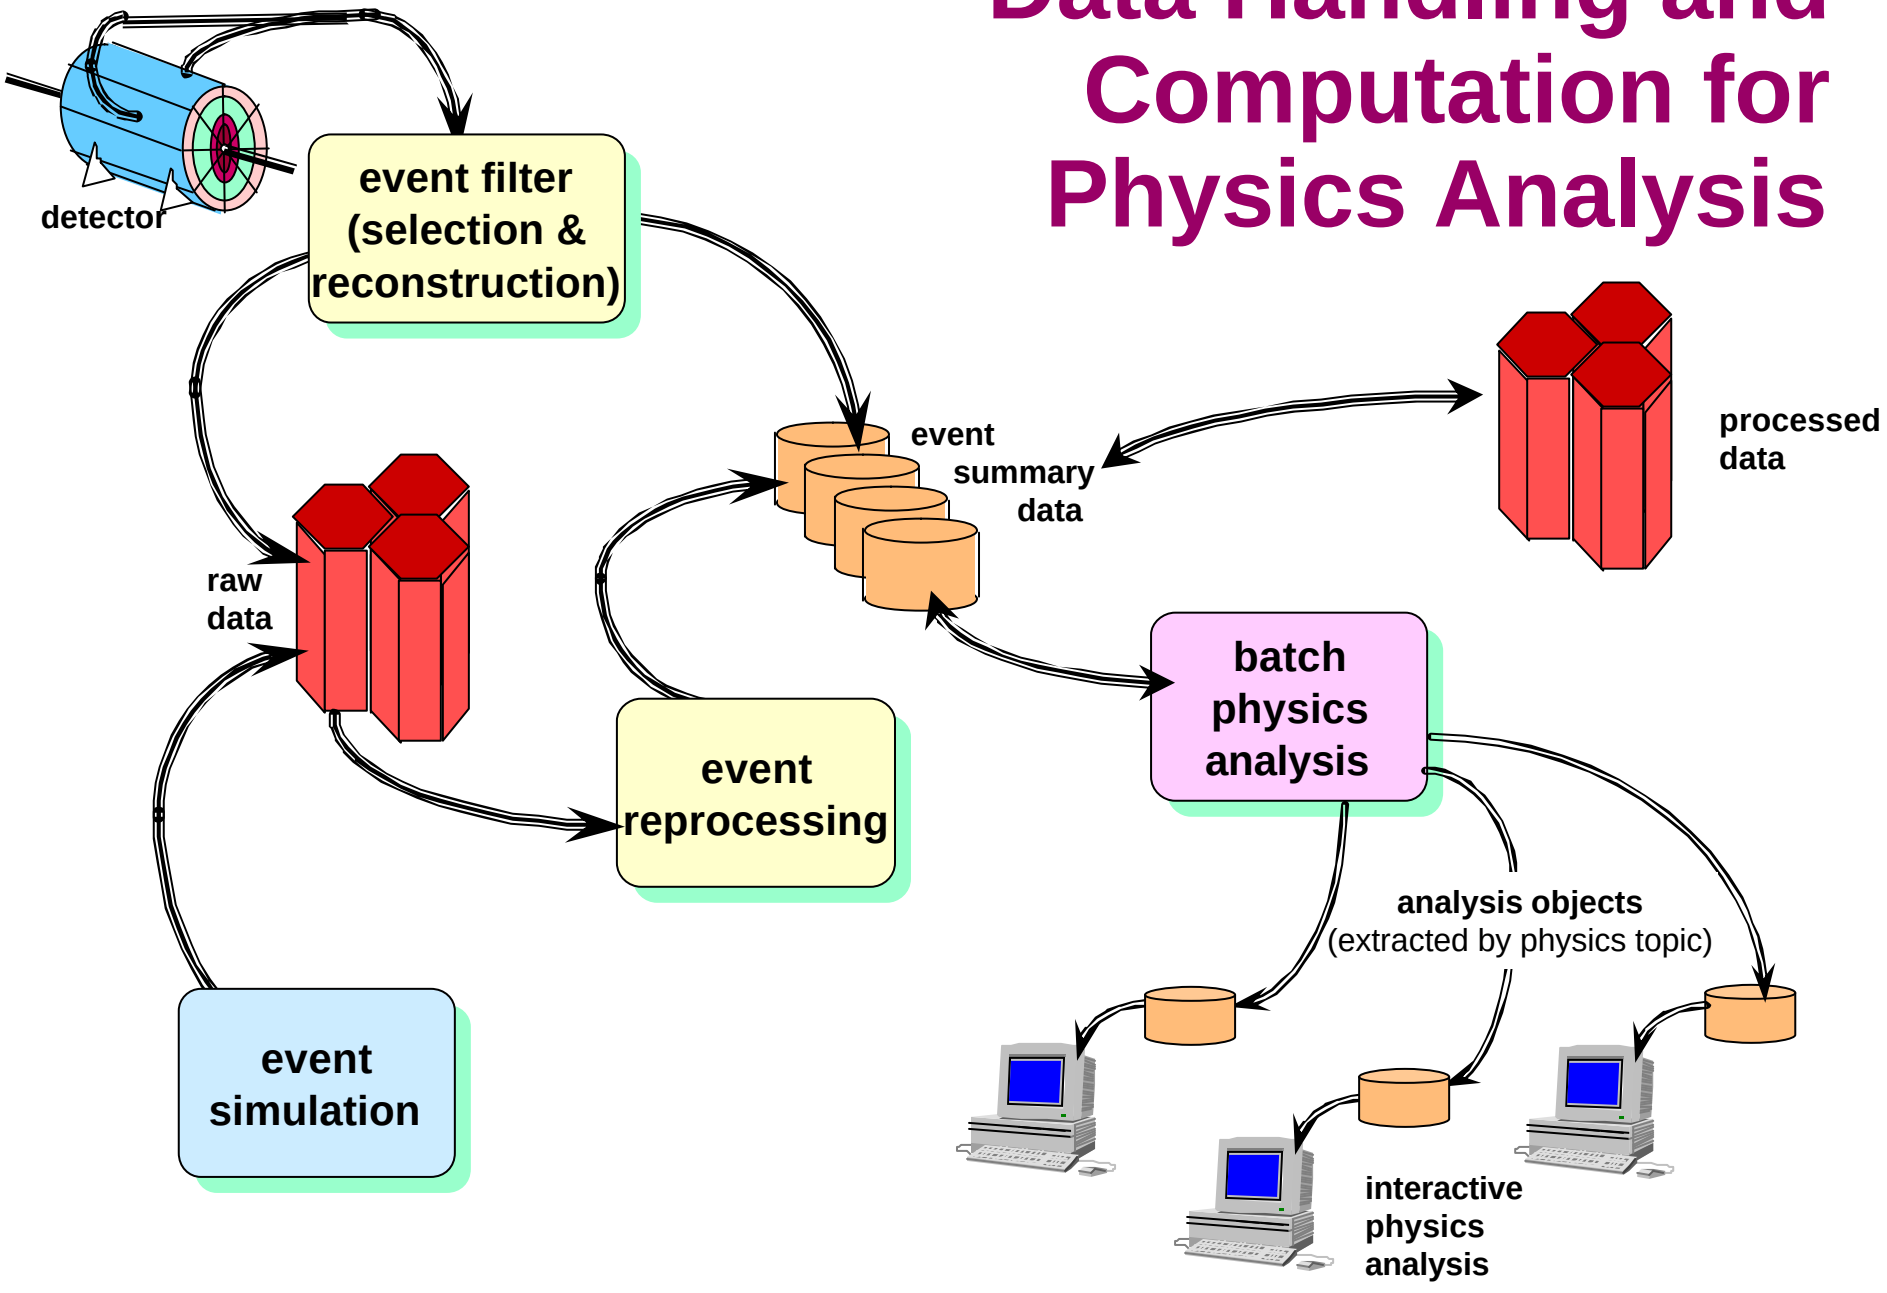

Computational Aspects

- To reconstruct and analyze 1 event takes about 90 seconds
- Most collisions recorded are not what we're looking for – maybe want as few as one out of a million. But we have to check them all!
- Analysis program needs lots of calibration; determined from inspecting results of first pass.
- ⇨Each event will be analyzed several times!

One of the four LHC detectors

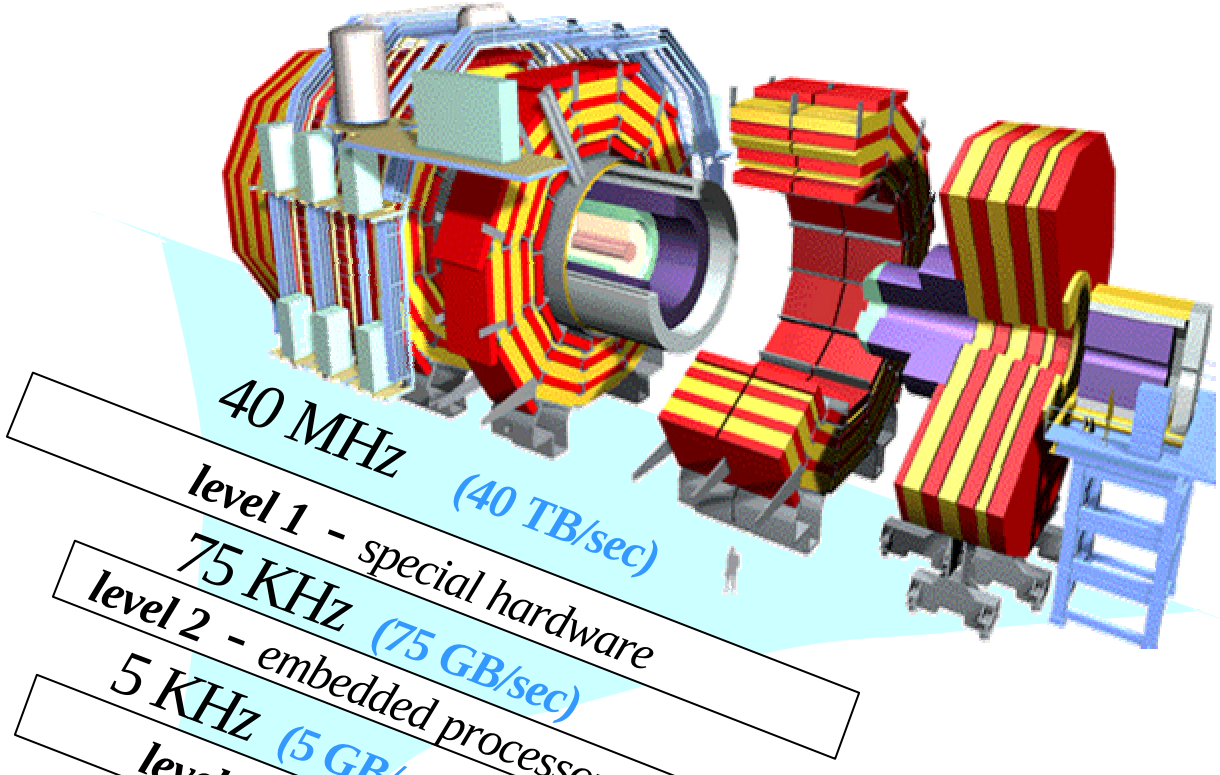

online system
multi-level trigger
filter out background
reduce data volume

Computational Implications (2)

- 90 seconds per event to reconstruct and analyze
- 100 incoming events per second
- To keep up, need either:
  - A computer that is *nine thousand times faster*, or
  - *nine thousand computers* working together
- Moore's Law: wait 20 years and computers will be 9000 times faster (we need them in 2006!)
- Four LHC experiments plus extra work: need >50k computers
- Grids: make large numbers of computers work together

So What are Grids Anyway??

A bunch of computers is not a Grid

- HEP has experience with a couple thousand computers in one place

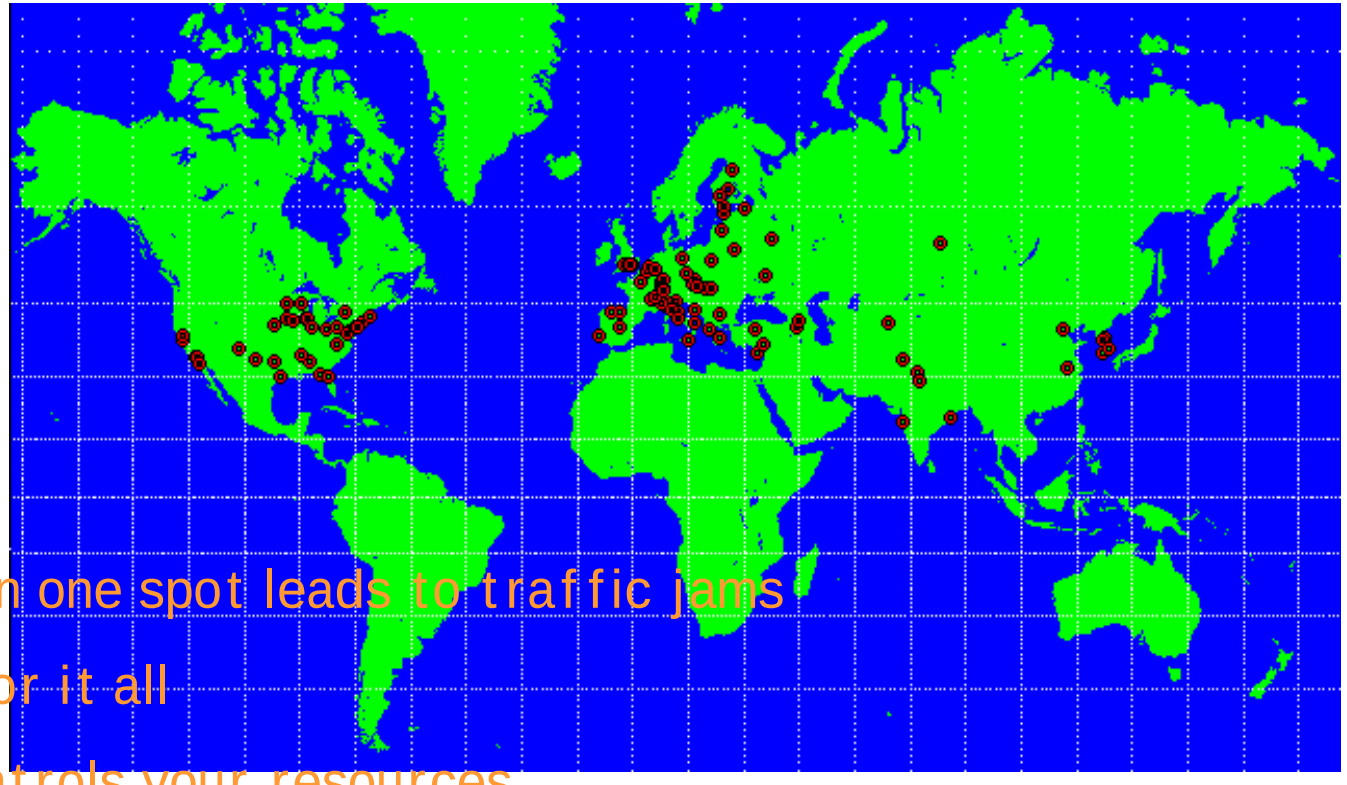

BUT

- Putting them all in one spot leads to traffic jams
- CERN can't pay for it all
- Someone else controls your resources
- Can you use them for other (non-CERN) work?

Distribute computers like users

- Most of computer power not at CERN
  - need to move users' jobs to available CPU
  - Better if jobs are sent to CPUs "close" to data they consume
- Need computing resource management
  - How to connect users with available power?
- Need data storage management
  - How to distribute?
  - What about copies? (Lots of people want access to same data)
- Need authorization & authentication for access to resources!

Grids are Information

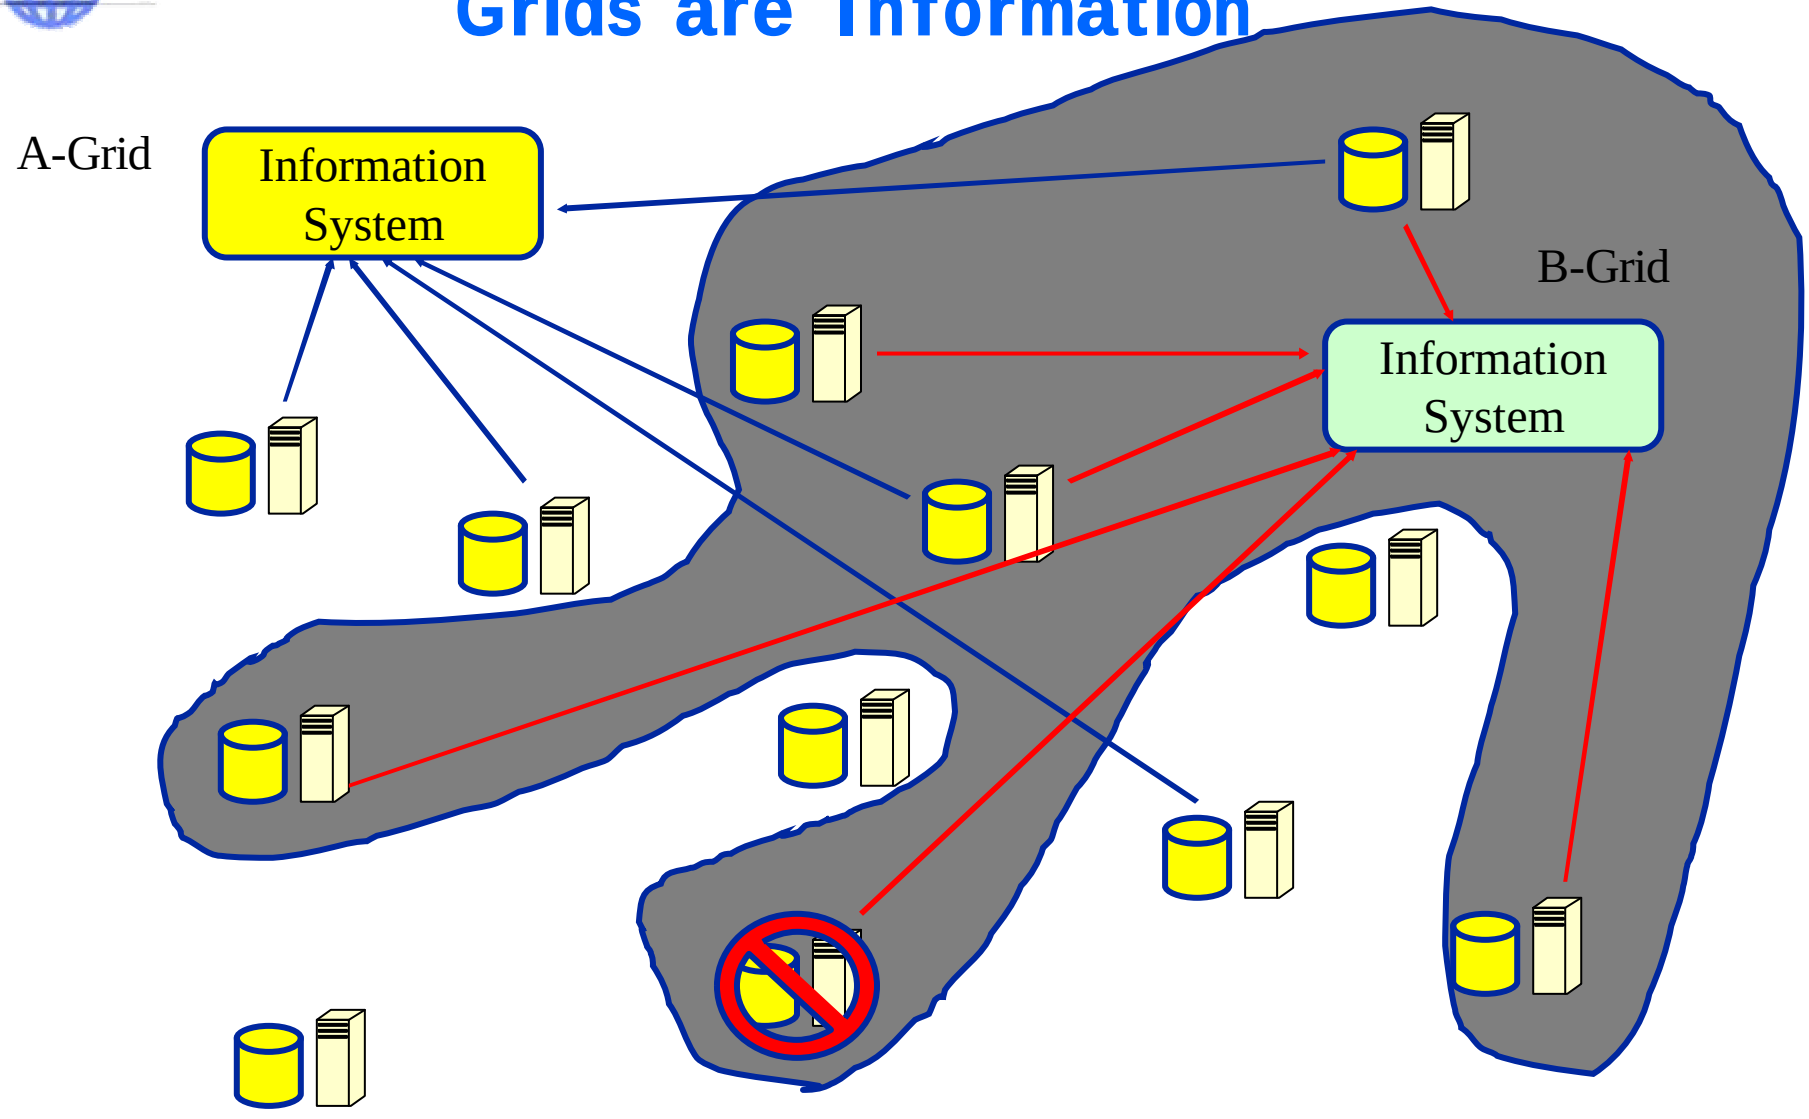

What does the Grid do for you?

- You submit your work, and the Grid
  - Finds convenient places for it to be run
  - Organises efficient access to your data
    - Caching, migration, replication
  - Deals with authentication to the different sites that you will be using
  - Interfaces to local site resource allocation mechanisms, policies
  - Runs your jobs
  - Monitors progress and recovers from problems
  - Tells you when your work is complete
- If your task allows, Grid can also decompose your work into convenient execution units based on available resources, data distribution

How it works

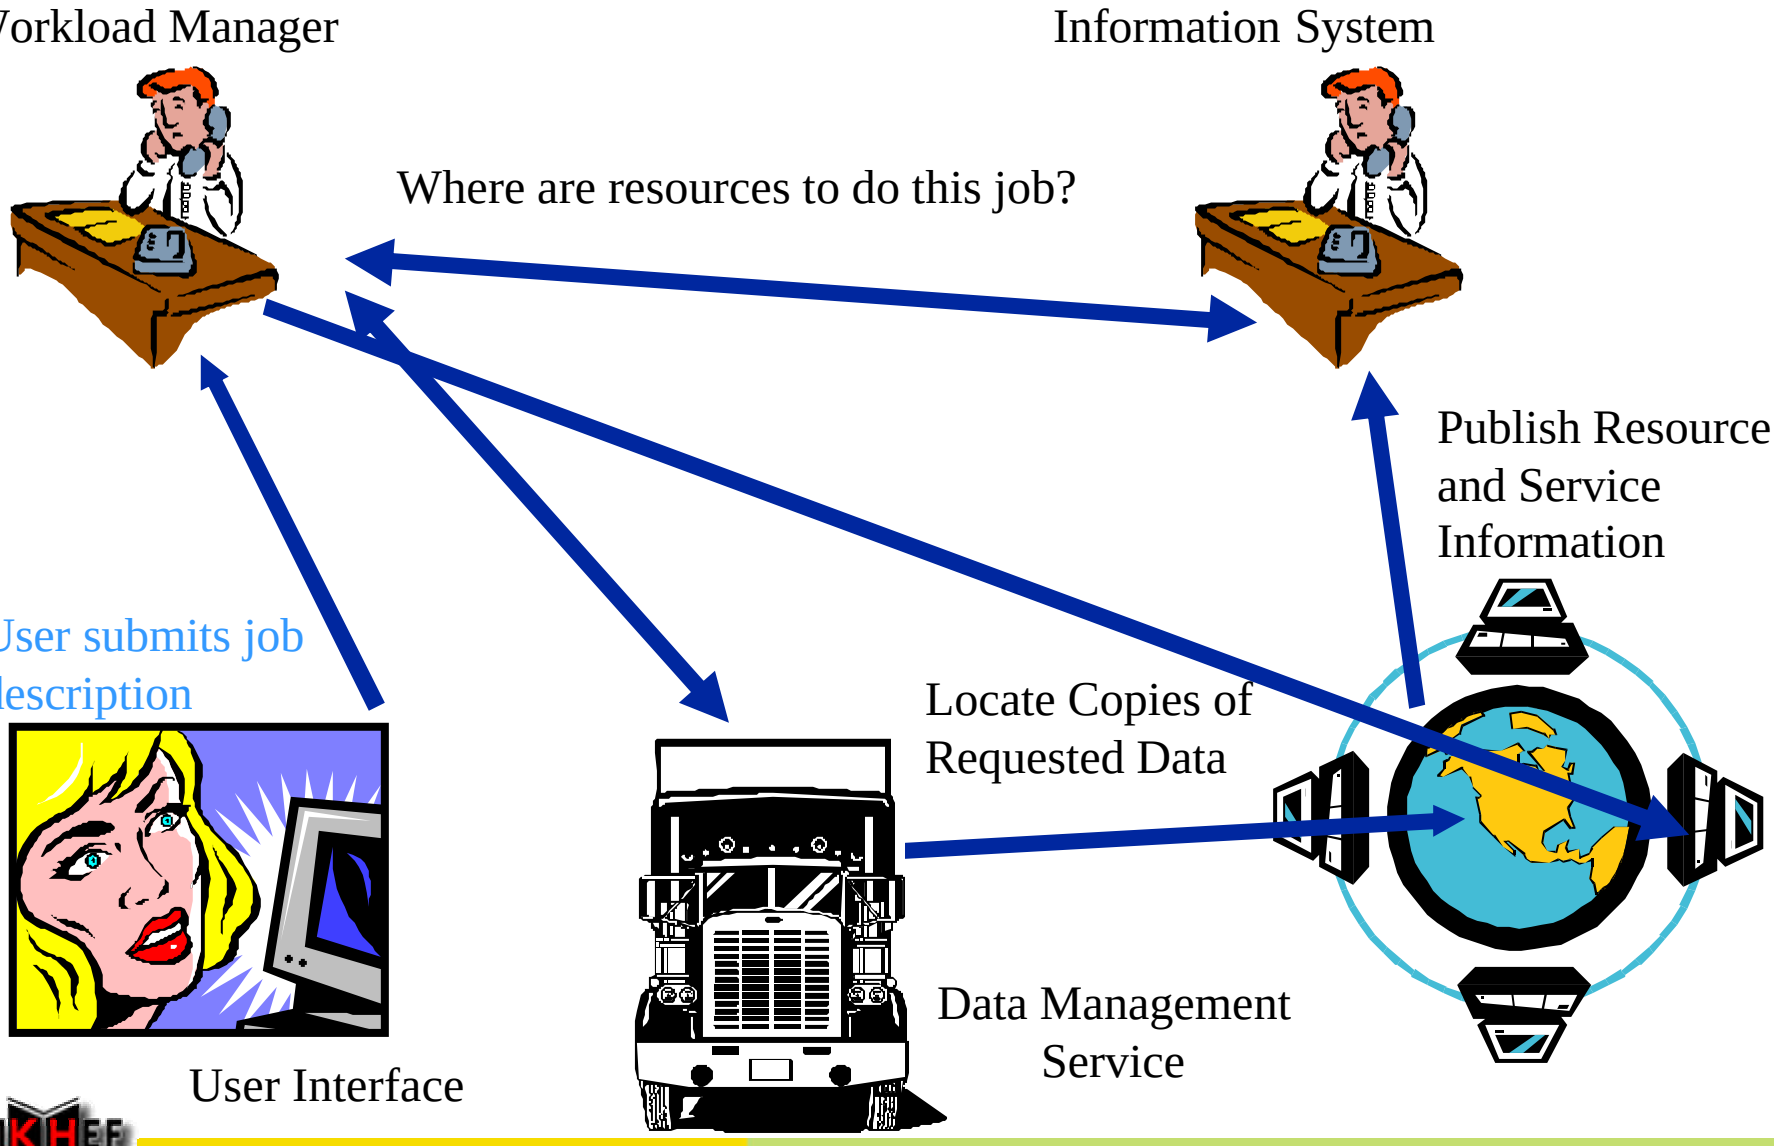
Workload Manager
Information System
Where are resources to do this job?
Publish Resource and Service Information
User submits job description
Locate Copies of Requested Data
User Interface
Data Management Service

What's There Now?

- Job Submission
  - Marriage of Globus, Condor-G, EDG Workload Manager
  - Latest version reliable at 99% level
- Information System
  - New System (R-GMA) – very good information model, implementation still evolving
  - Old System (MDS) – poor information model, poor architecture, but hacks allow 99% "uptime"
- Data Management
  - Replica Location Service – convenient and powerful system for locating and manipulating distributed data; mostly still user-driven (no heuristics)
- Data Storage
  - "Bare gridFTP server" – reliable but mostly suited to disk-only mass store
  - SRM – no mature implementations

VLVνT Reconstruction Model

Grid data model applicable, but maybe not computational model …

- Distributed Event Database?
- Auto Distributed Files?
- Single Mass Store + "Thermal Grid"?

Grid useful here – get a lot but only when you need it!

All connections through single pipe probably bad. Dedicated line to better-connected "redistribution center"?

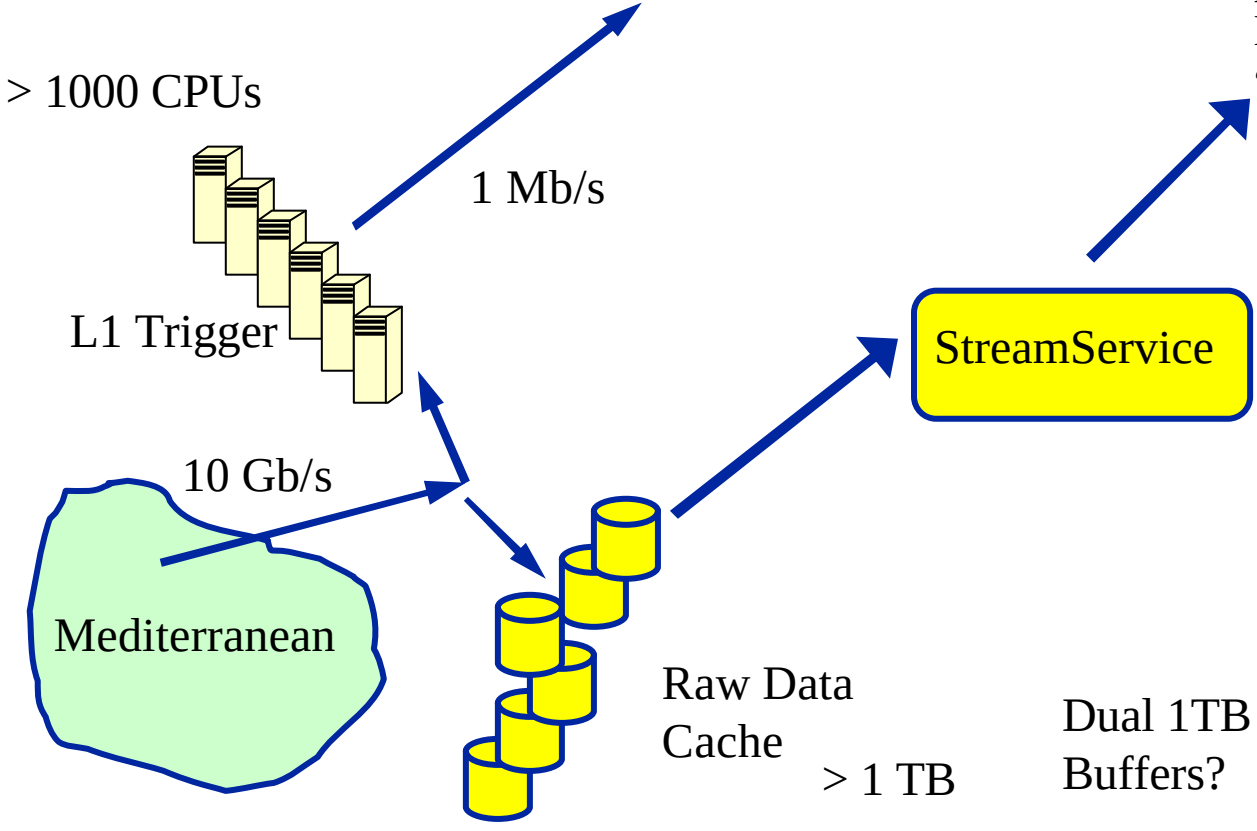

This needs work!! 2 Gbit/s is not a problem but you want many x 80 Gbit/s!

Dual 1TB Circular Buffers?

Networking Likely OK, Surely Welcome!

- History of Internet Land Speed Record
  - 930 Mbit/sec (NIKHEF/California) 1 year ago
  - 2.2 Gbit/sec (CERN) six months ago
  - 5 Gbit/sec today (tests last week from NIKHEF)
- This rapid advance results in network people looking for groups who can fill the pipes!!

Conclusions

- Grids are working *now* for workload management in HEP
- Data model of VLVνT matches Grid data model quite well
- Gamma-Ray burst scenario matches Grid computational paradigm well
- Network demands appear feasible and will be welcomed
- Sounds like a lot of fun!