



HAMBURG • ZEUTHEN

The Global Grid and the Local Analysis

Yves Kemp

DESY IT

DESY

GridKA School, 11.9.2008

Overview




- **“Global” and “globalization”: Some thoughts**
- **Anatomy of an analysis and the computing resources needed**
- **Boundary between “global” and “local”**
- **Some problems analysts and resource providers encounter**

Disclaimer: If you expect an “how-to” during this talk, you will be disappointed. The speaker does know not “The Right way to go” himself...

Wikipedia: “Globalization”

- **Globalization** in its literal sense is the process of transformation of local or regional things or phenomena into global ones. It can also be used to describe a process by which the people of the world are unified into a single society and function together. This process is a combination of economic, technological, sociocultural and political forces. (Sheila L. Croucher, 2004)
- **Globalization is NOT “Everyone can do everything without limits”**
 - Have a common set of rules
 - In Computer Science: This is called “defining standards”
 - The different actors need clear interfaces they can rely on

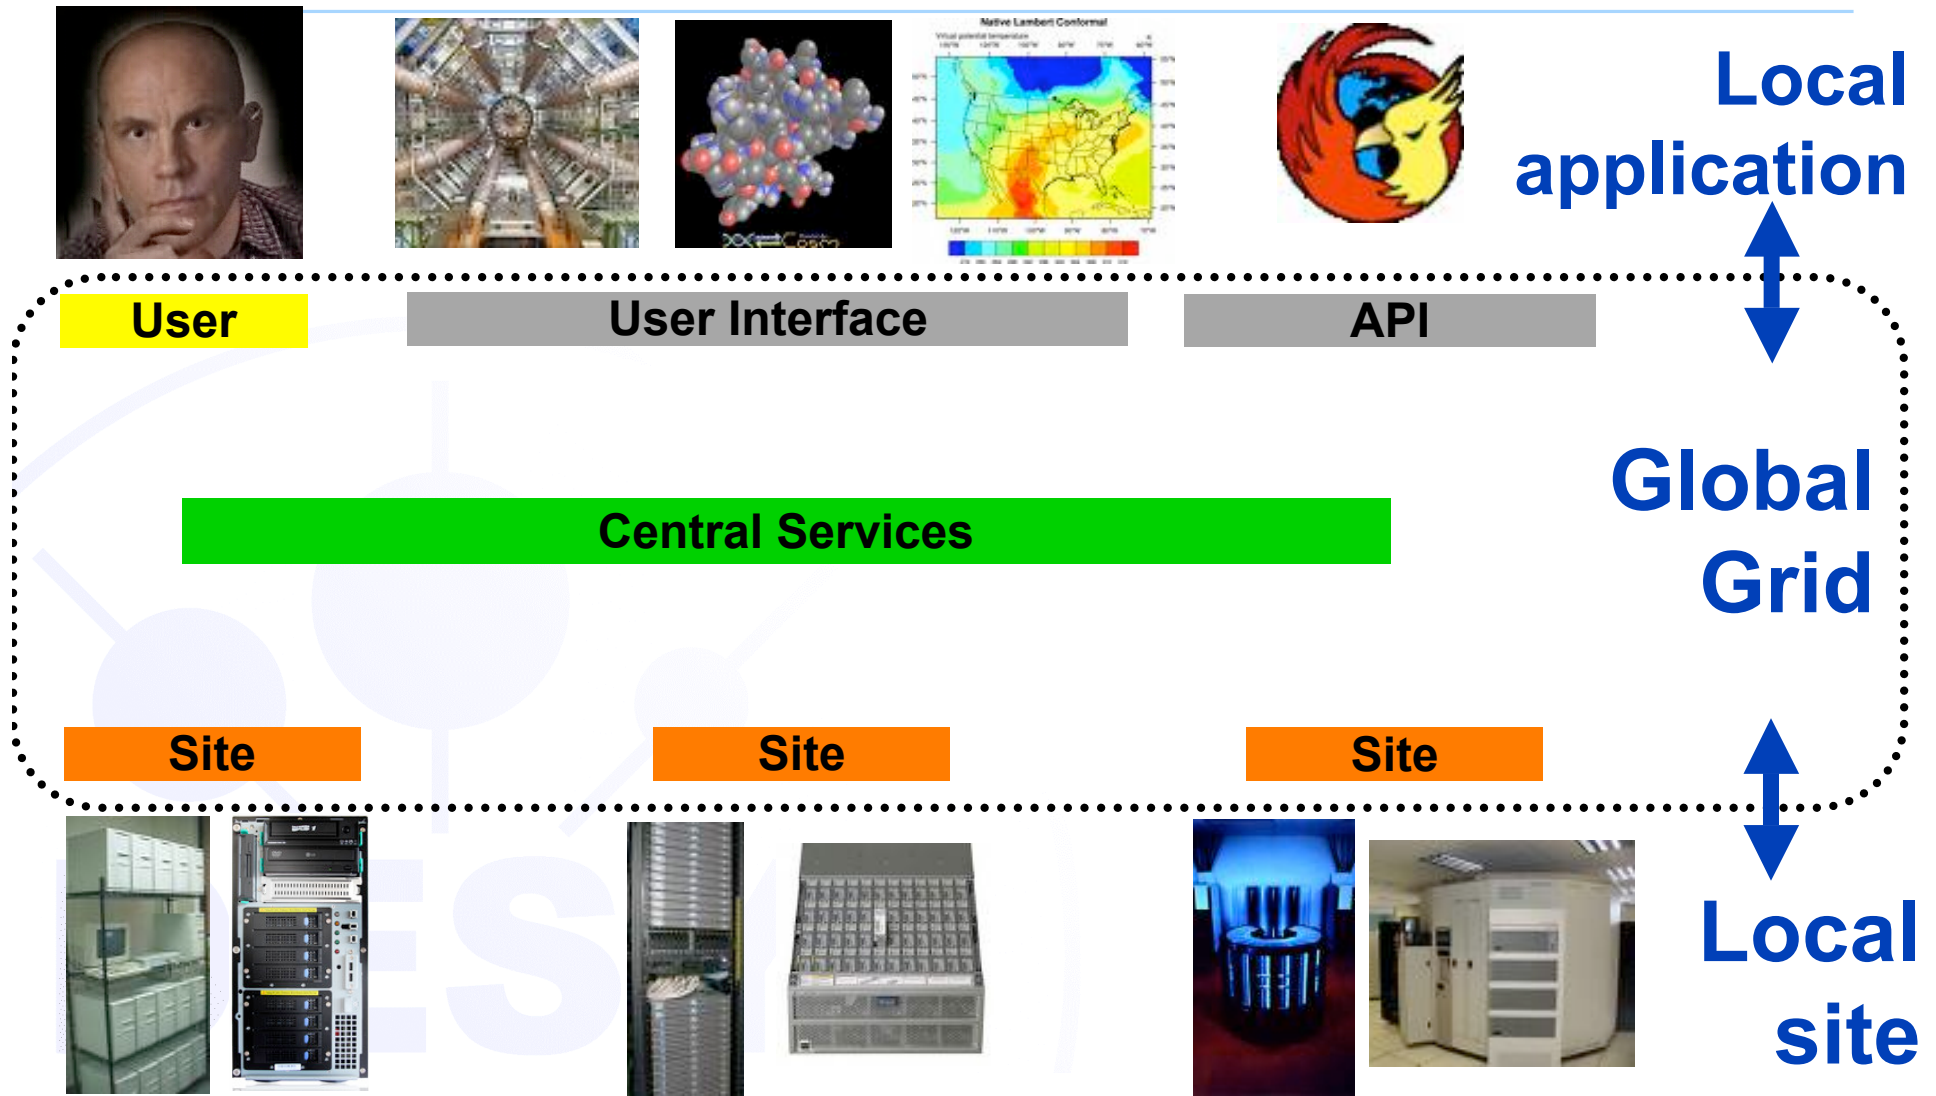


And local? What does that mean?


- **Every resource is local somewhere**
 - It is in a “local” building, with “local” admins, “local” users, ...
- **For users:**
 - You know the facility, the admins, the peculiarities, ...
 - The resource might even be near to you (geographically)
 - Eventually, you have a different (local) access



Anatomy of the Grid (very simplified)



Analysis facility (very simplified)



The central point
of data analysis
... is the data

→ Analysts need a
system where
storage is the central
part

→ Analysts want an
integrated system,
not a layered and
distributed system
like the Grid

**Analysis has its own
requirement!**

Different tasks: Different requirements



Coordinated &
global tasks

Example:
HEP model

■ MC Production

- Event Generation: no I; small O; large CPU
- Detector Simulation: small I; **large O & CPU**

■ Event Reconstruction/Reprocessing

- Reprocessing: **full I; full O; large CPU**
- Selections: large I; large O; large CPU

Uncoordinated,
unstructured
& local tasks

■ Analysis

- Usually: **large I; small O; little CPU**
- Performed by many users, many times!
- LHC StartUp phase: Short turn-around

Performance criteria for analysis



Real-time response

- Get the answer after a definite time

Fast

- Well, as fast as possible
- Rapid analysis (like Rapid Application Development)

Interactive


- You can react instantaneously and steer your analysis

From Global to Local:



Data
collected
with an
experiment

MC
simulations

- 
- Reprocessing & converting/stripping data
 - Creating group or personal data skim
 - Applying own & modified algorithms
 - ...

(Like the “Waterfall model”)

Higgs mass
is ... GeV/c²

Earth warming
is ... °C / year

Where to put the boundary?

Boundary: Different approaches

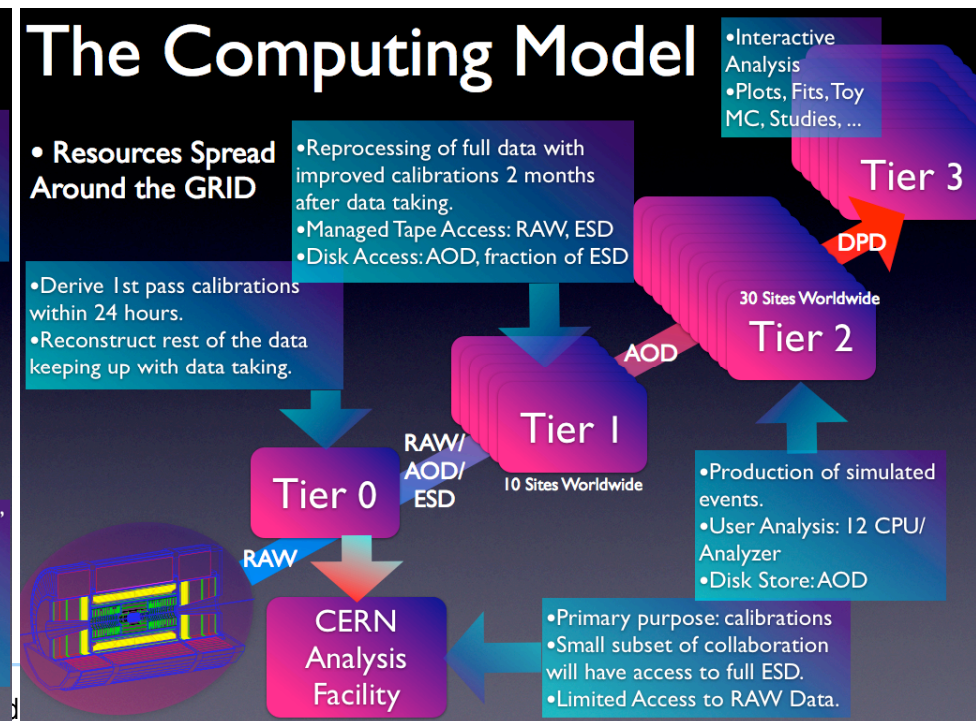
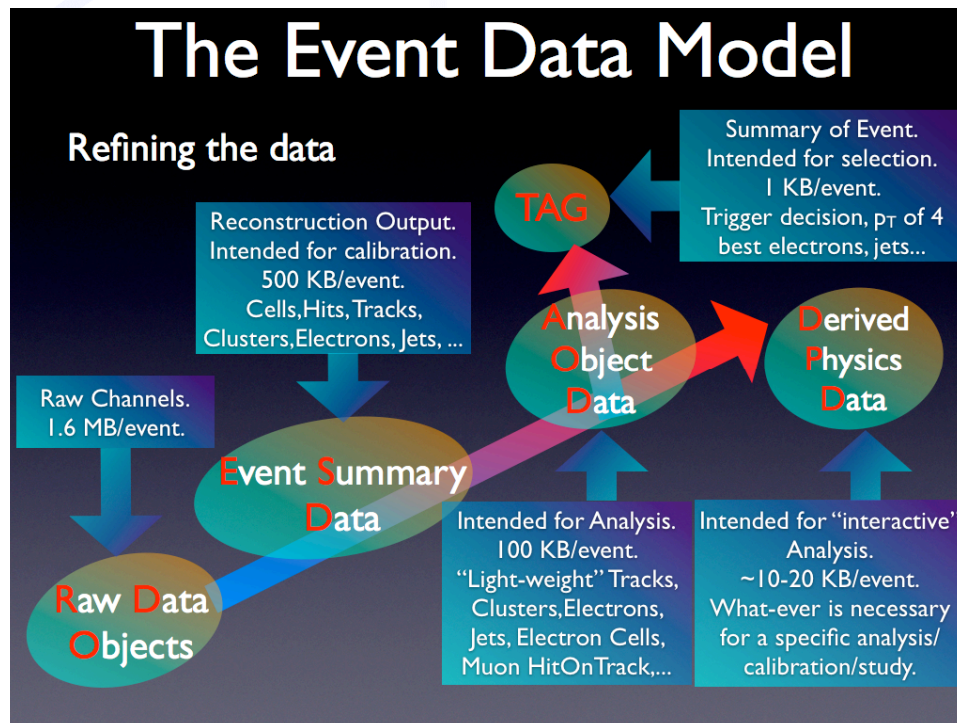


- In the Computing Model
 - E.g. LHC Tier Model: Each layer has different roles, T3 “the most local one”
- Adding layer of abstraction!
 - E.g. Submission frameworks hide underlying resources
- Integrate into existing facilities, logical separation
 - E.g. NAF @ DESY (additional fairshare and storage,...)
- Additional facilities, only partly coupled to the Grid
 - E.g. NAF @ DESY (WGS, local batch queue, local filesystem...)
- Enhancing Grid:
 - E.g. Interactive Grid: Having the instantaneous response users expect from a local machine (remember the previous talk)

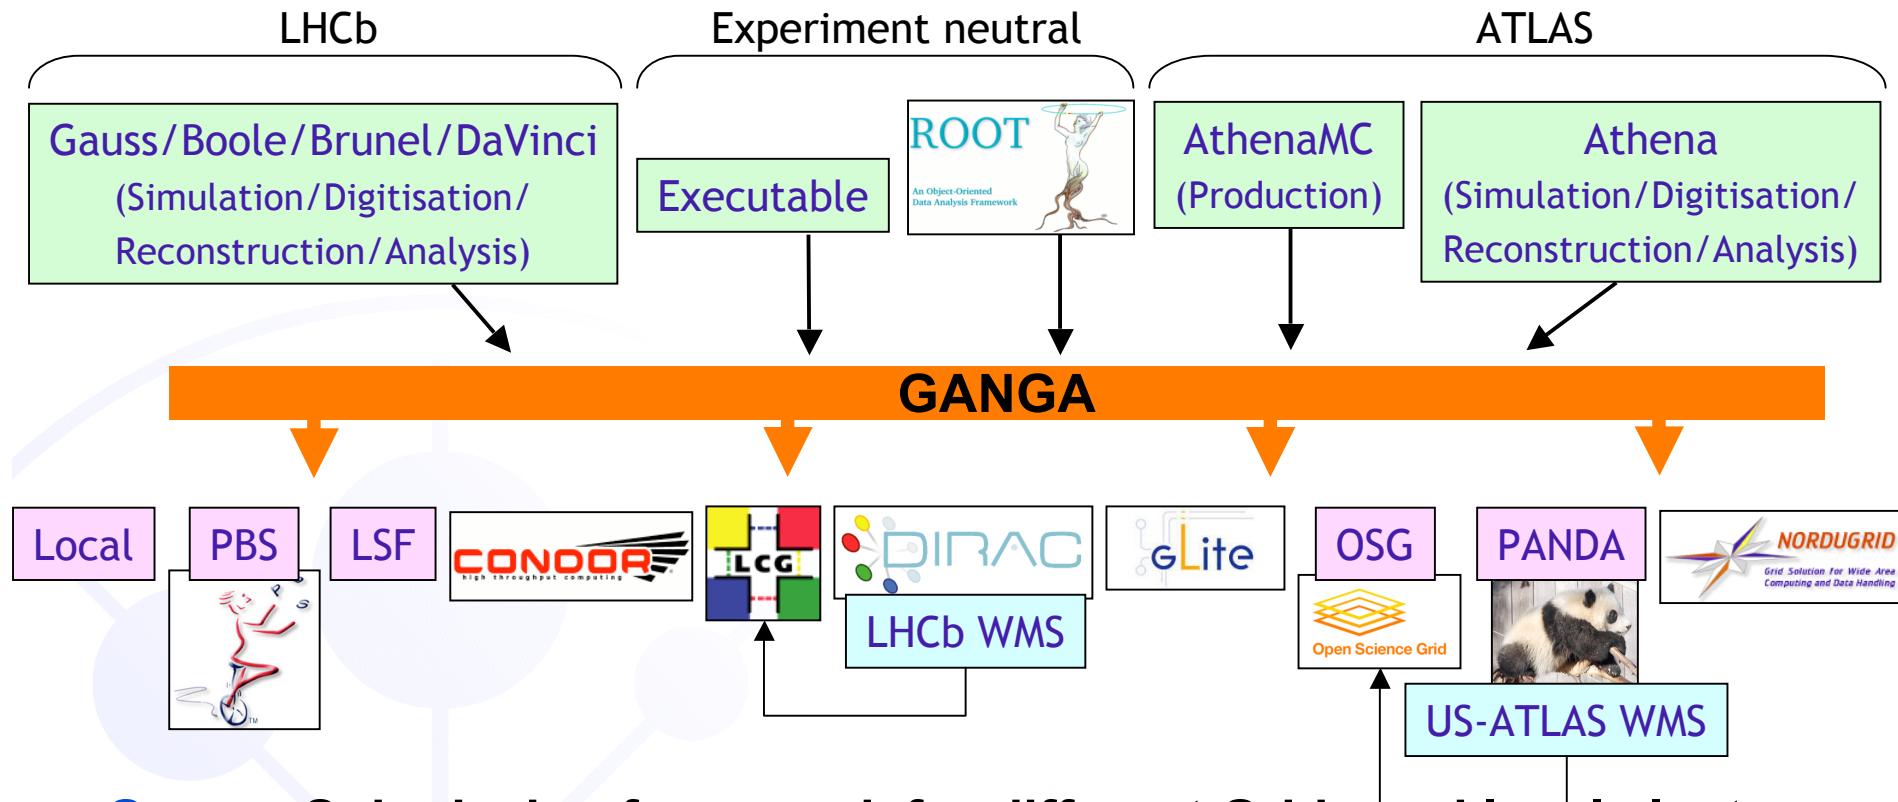
→ **No single road to success!**

Boundary: Computing Model

- Different tasks require different setups
 - Break-up with uniformity and homogeneity of an “Ideal Computing Grid”
 - Have different hierarchical layers for the different tasks
 - Each layer is specialized for one set of tasks, dataflow among layers
 - Example: ATLAS Data Model and Computing model (slides from Amir Farbin)



Boundary: Additional layer of abstraction (e.g. Ganga)



- **Ganga:** Submission framework for different Grids and local clusters and local machines!
- **CRAB** (submission tool developed by CMS)
- **PROOF** and **glitePROOF** (ALICE,...)

Ganga & PROOF:
See tutorials!

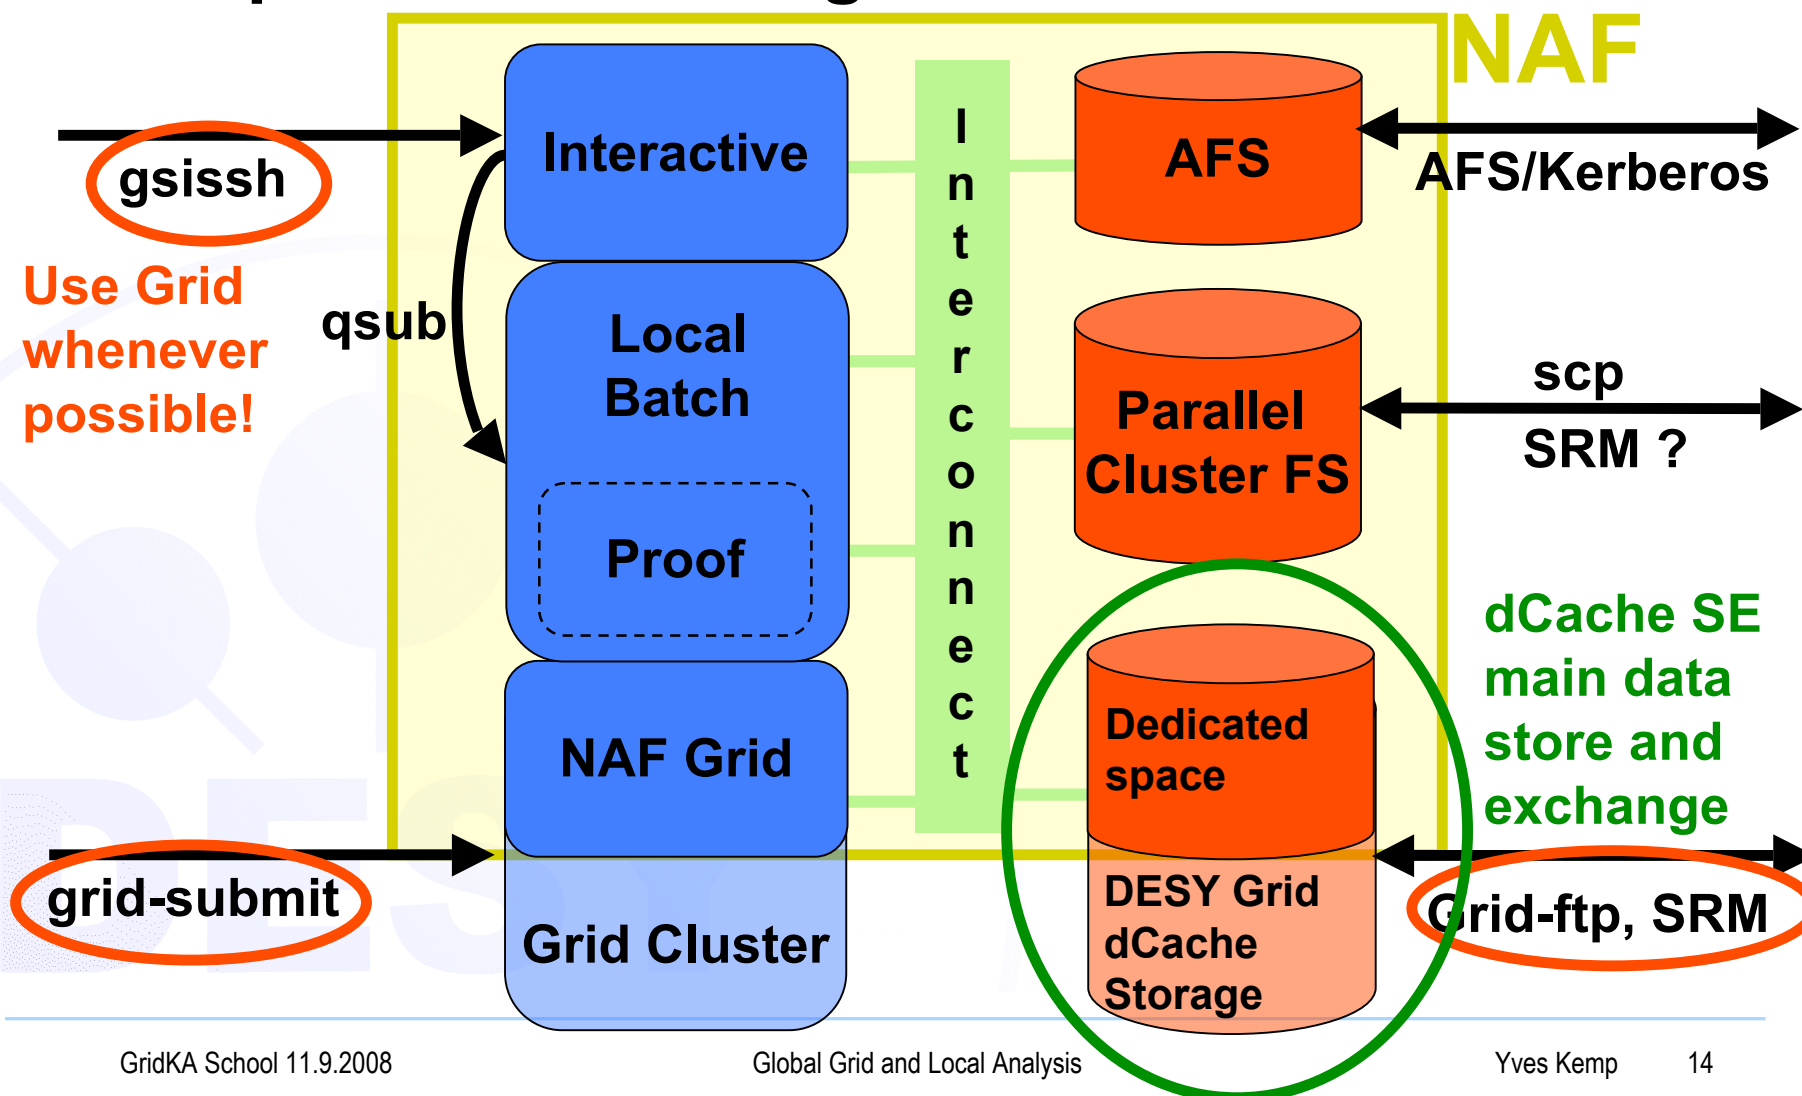
Boundary: Integration into existing




- **Example NAF @ DESY:**
 - Use **VOMS-Groups** to differentiate local users
 - CE: Scheduler gives prioritization and fairshare
 - SE: additional space only writable by /.../de
- **Problems:**
 - Info system is per-VO (ignores VOMS groups)
 - Education of users: Use of VOMS-extension!
 - Potential danger of split into too many groups

Boundary: Build something new

Example NAF: building blocks:





? **Some concerns still remain in the areas:**

- **Storage technique and management**
- **Support**
- **Computing services**
- **Administrators**

(following slides show only some aspects)

Storage technique



- **Global file catalog and name space vs. local one.**
 - Are the two in sync? Which one is the authoritative?
 - Lots of effort going into *manually* synchronizing the two
 - Problem known as **Dark Data** (see G. Cowan, Monday)
- **Access to Storage**
 - Local access protocols to Global storage
 - Global: e.g. gridftp, local: could be mounted file system
 - File system: Has Users (UID/GID) / VOMS: has Groups/Roles
 - Some operations are just not possible that one expects from the simple laptop HD (read-modify-write somewhat difficult with tapes...)


Storage management



- **Who can write where?**
 - Experiments usually have “Role=production” to write data
 - And there are different groups of users
 - Quotas and ACLs: How to consistently set them up?
 - “Home” directory in the Grid
 - ➔ Lots of communication needed between partners to set up space for different groups
- **Orphaned files somewhere in the Grid**
 - No owner (in the sense of: is using them)
 - On SE, but not in LFC (recall: Dark Data)
 - ➔ Need *people* taking care of storage

Do not forget support!

- The more users, the more support groups exist.
 - How to choose the right one?
 - People want “personal” support
 - does not scale at a Grid level
 - Information flow to your local working group is (relatively) easy.
 - When and how to inform a globalized community?
 - Debugging:
 - Large latencies in a distributed system
- ➔ Probably time will show how support will organize itself
- And organization will change with time



- **Again: How to handle (many small?) subgroups?**
 - Is fairshare enough? Will it scale and is it manageable?
 - They might want to install their own software in the Grid
 - This is one reason for separate clusters!
- **Different jobs profile**
 - Can the (meta-) schedulers handle this?
 - Again: One reason for separate clusters
- **Accounting**
 - Different supported subgroups might induce different financing bodies with different accounting wishes ... on the same resources

From a site admins perspective

- **Also for the admin, many services are global**
 - He/she has no control over them
 - Cannot easily debug things
 - Examples: VOMS server, FTS, WMS, LFC...
 - ➔ Needs inter-site communication and global support channels
 - ➔ Keep services as generic as possible
- **The site resources are part of a bigger thing**
 - WLCG/EGEE sets standards (e.g. Software installation, data access,...)
 - Admin has little leeway, and gives away some control
- **Everything is larger ... and more anonymous**
 - E.g. for support
- **Admin have a hard time explaining this to the users**

Summary and Outlook



- **Analysis has special requirements**
- **Most analysis is data-centric**
- **Finding the right boundary between Grid and local facilities is important for success of the final analysis**
- **Many ways to success, but still some open concerns and work to be done**
- **Future will tell which way(s) will be chosen**