CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Overview of HEP software & LCG from the openlab perspective

Andreas Unterkircher, CERN openlab

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Contents

1. Opencluster overview
2. High Energy Physics (HEP) software basics
3. LCG overview (including an outlook to gLite)
4. SmartFrog

Part 1: Opencluster overview

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Opencluster hardware

- 4 Enterasys Matrix N7 Enterprise switches
- 2 Enterasys Matrix X-Series Core Router
- 96 HP dual RX2600 (Itanium 2 @ 1.5 GHz, 2-12 GB memory)
- 2 HP quad RX4640 (Itanium 2 @ 1.5 GHz, 8-16 GB)
- 8 ibm TotalStorage 200i systems (each system has 6 internal 73GB SCSI disks plus 3 extension units with 14 73GB SCSI disks each)
- 1 Voltaire ISR 9096 InifiniBand Switch Router

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Opencluster usage

- Porting (LCG, physics software)
- LCG IA64 testbed (incl. SmartFrog usage)
- Service challenges (gridFTP)
- Infiniband tests
- Benchmarking
- 10 GB NICs
- Virtualization (Xen)
- Numerical software (BeamX)
- High speed network tests
- StorageTank IA64 client

Part 2: HEP software basics

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Physics Software

Running a physics job on the grid is not just running one executable but starting an application within a framework.

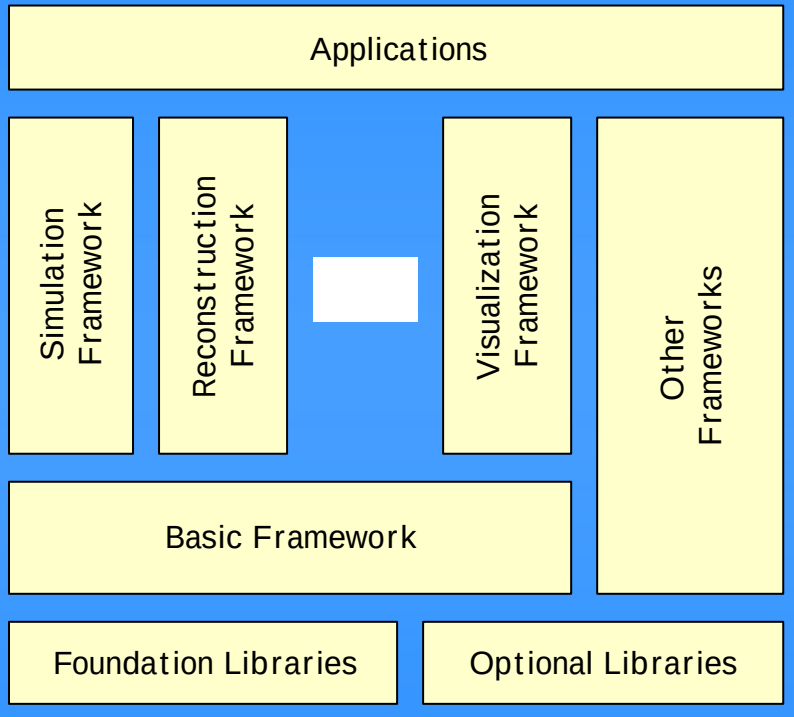

Domain decomposition

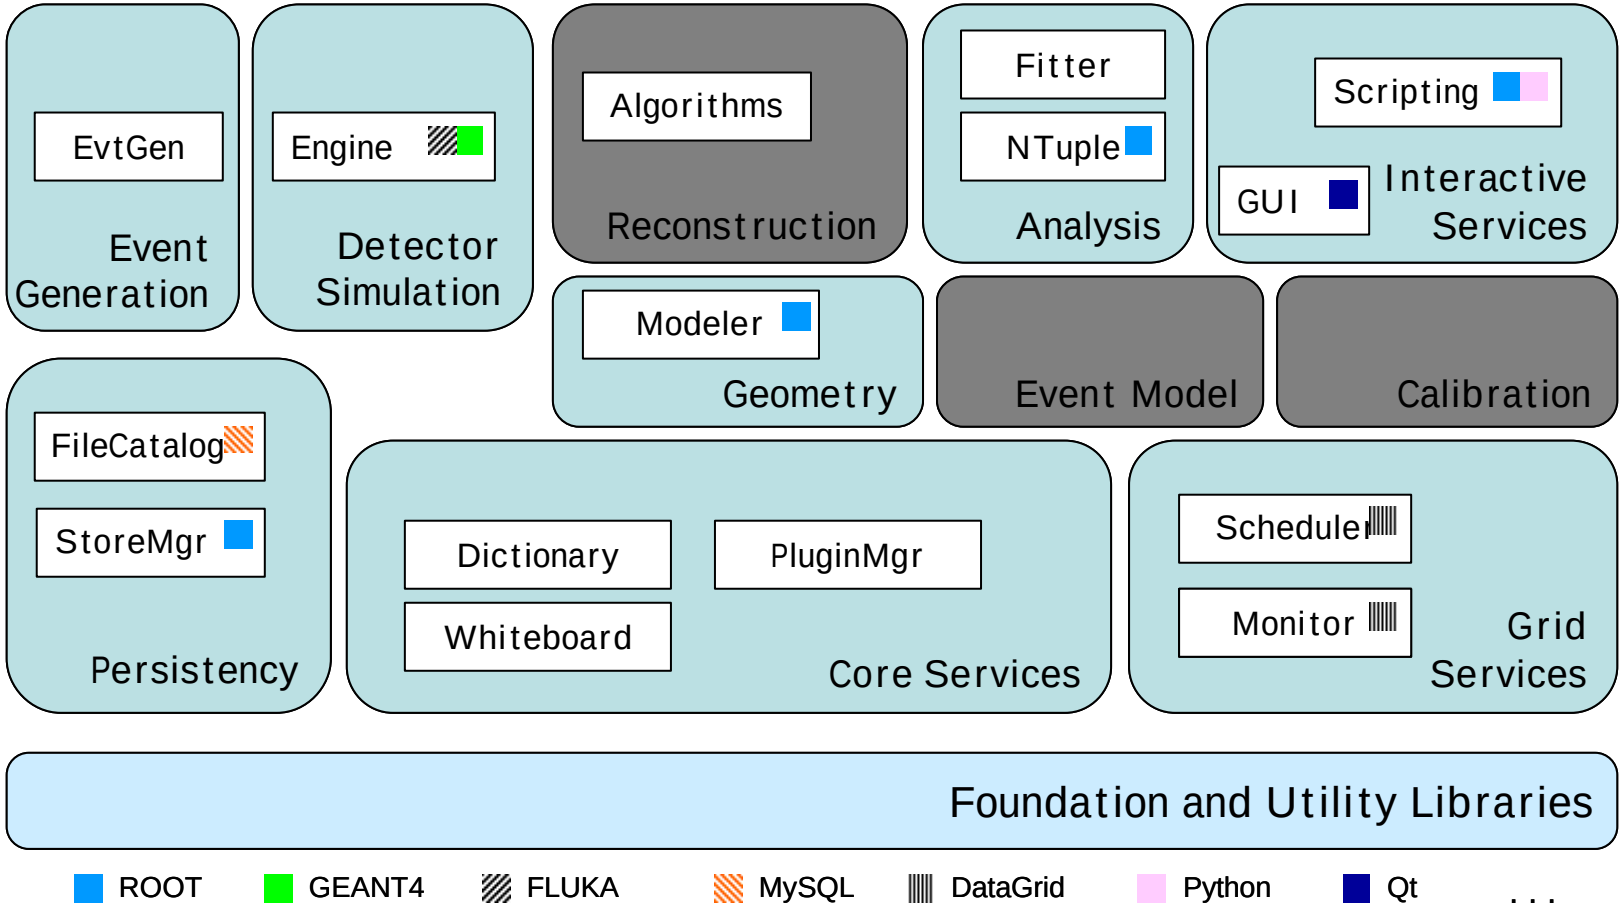

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Main components

- POOL: common persistency framework
- SEAL: core libraries (math, reflection,...)
- PI: physics infrastructure (abstract interface to analysis service)
- ROOT: I/O layer

These libraries are supposed to be used by the experiments, except for ALICE which relies only on ROOT (works on IA64).

Example: GAUDI

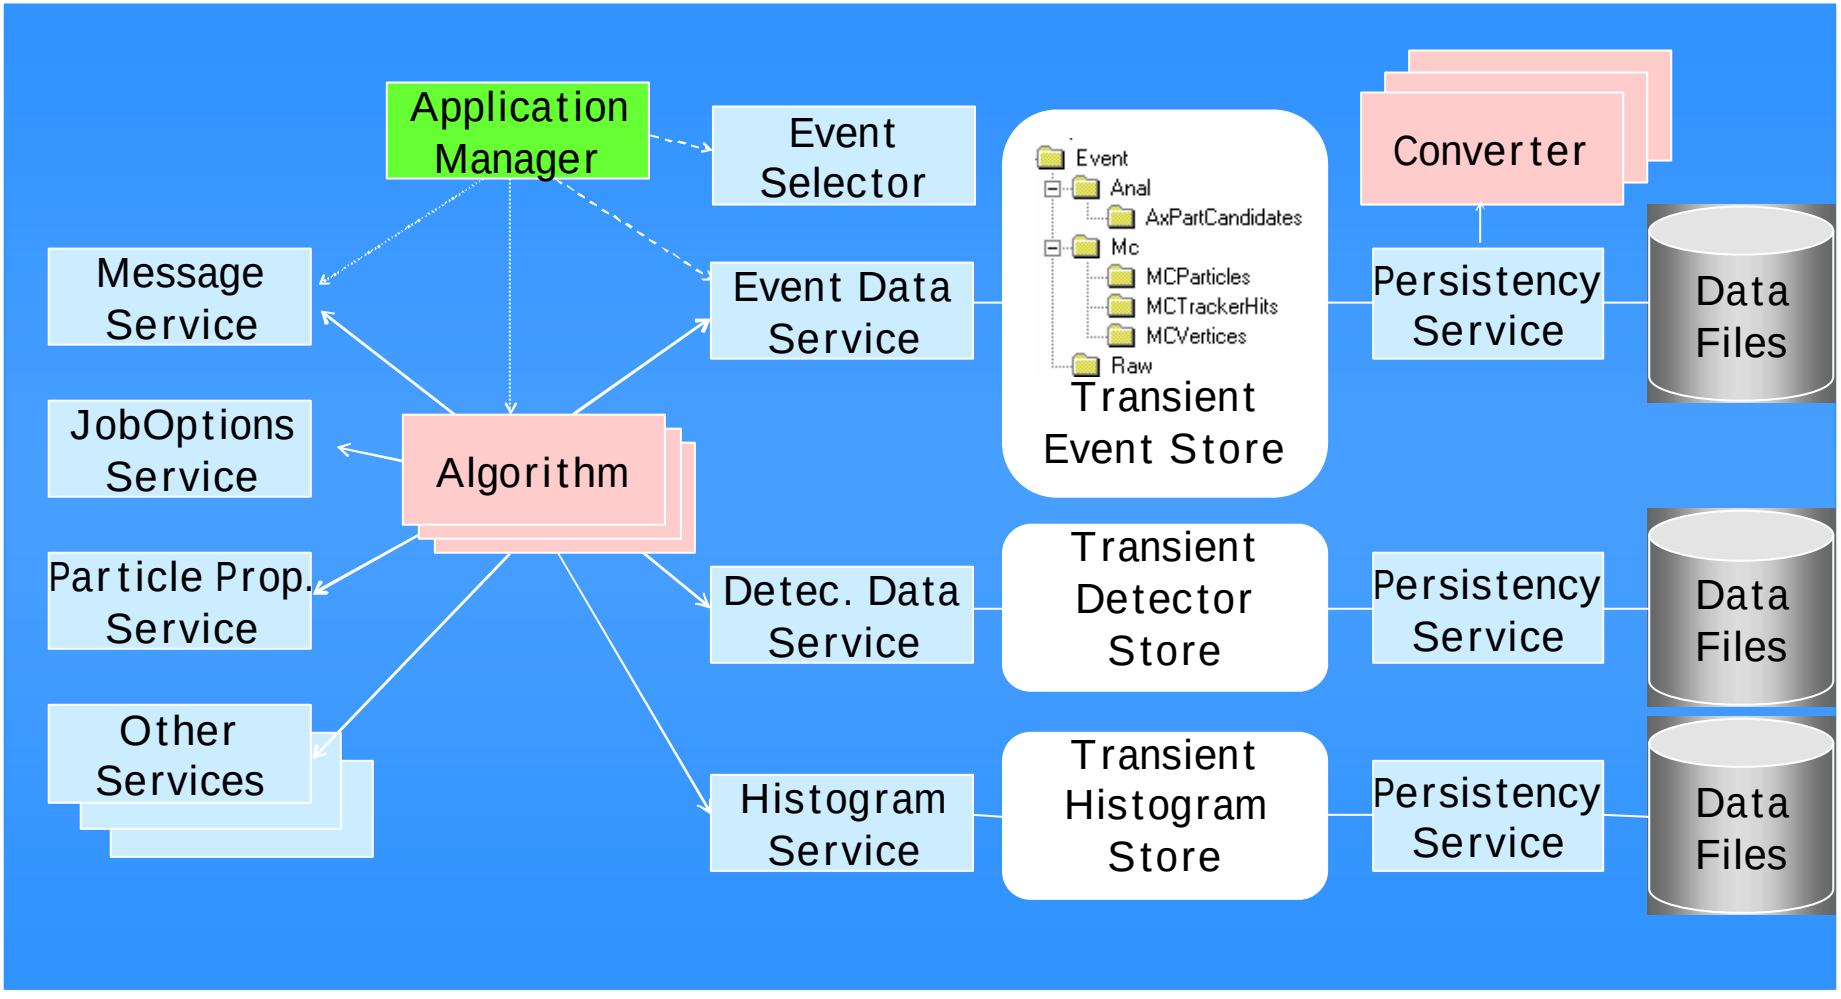

Example: GAUDI

Used by LHCb. Relies on SEAL, POOL and PI.
Also has interfaces to ROOT (data analysis) and GEANT4 (simulation).

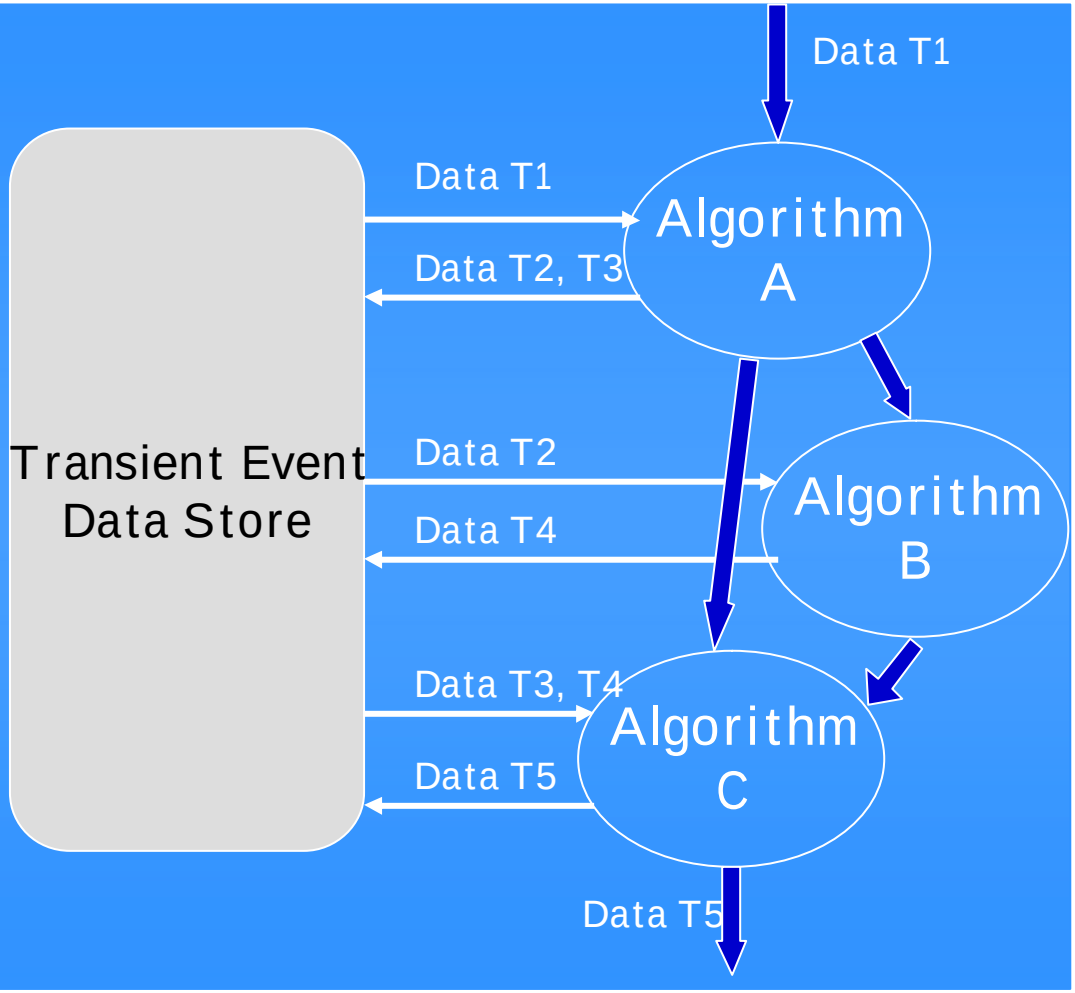

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Porting GAUDI to IA64

At openlab we aim at porting GAUDI to IA64 in order to be able to do an LHCb data challenge on IA64.

But first SEAL, POOL & PI have to be ported…

Started by a summer student in 04, now being followed up seriously.

Status: SEAL nearly finished. ROOT already runs on IA64.

Porting obstacles

- Frameworks have their own building mechanisms which are hard to understand.
- Only developed for very specific OS & compiler versions.
- It is essential to get the changes back into developers‘ cvs (i.e. to convince the developers to support IA64). Luckily now there is a general interest in 64bit Linux (AMD 64, EM64T).

Part 3: LCG overview

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Some LCG acronyms

CE	Computing Element	Gateway to computing resource (batch system)
UI	User Interface	Submit job & retrieve output
RB	Resource Broker	Handles job requests, selects CE, track jobs, handles I/O
SE	Storage Element	Gateway to storage
PX	Proxy Renewal	Extends lifetime of proxy certificate
WN	Worker Node	Computing resource controlled by a CE (batch node)
BDII	Information Index	Manages list of relevant information providers & some status information, gets data from MDS
MDS	Monitoring & Discovery Service	Hierachy of information providers (ldap)
GLUE	Grid Laboratory for a Uniform Environment	Data scheme to describe grid resources
RMC	Replica Metadata Catalog	Mapping between GUIDs and LFNs (aliases)
RLS	Replica Location Service	Information on physical location of replicas, composed of several local replica catalogs.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Architecture

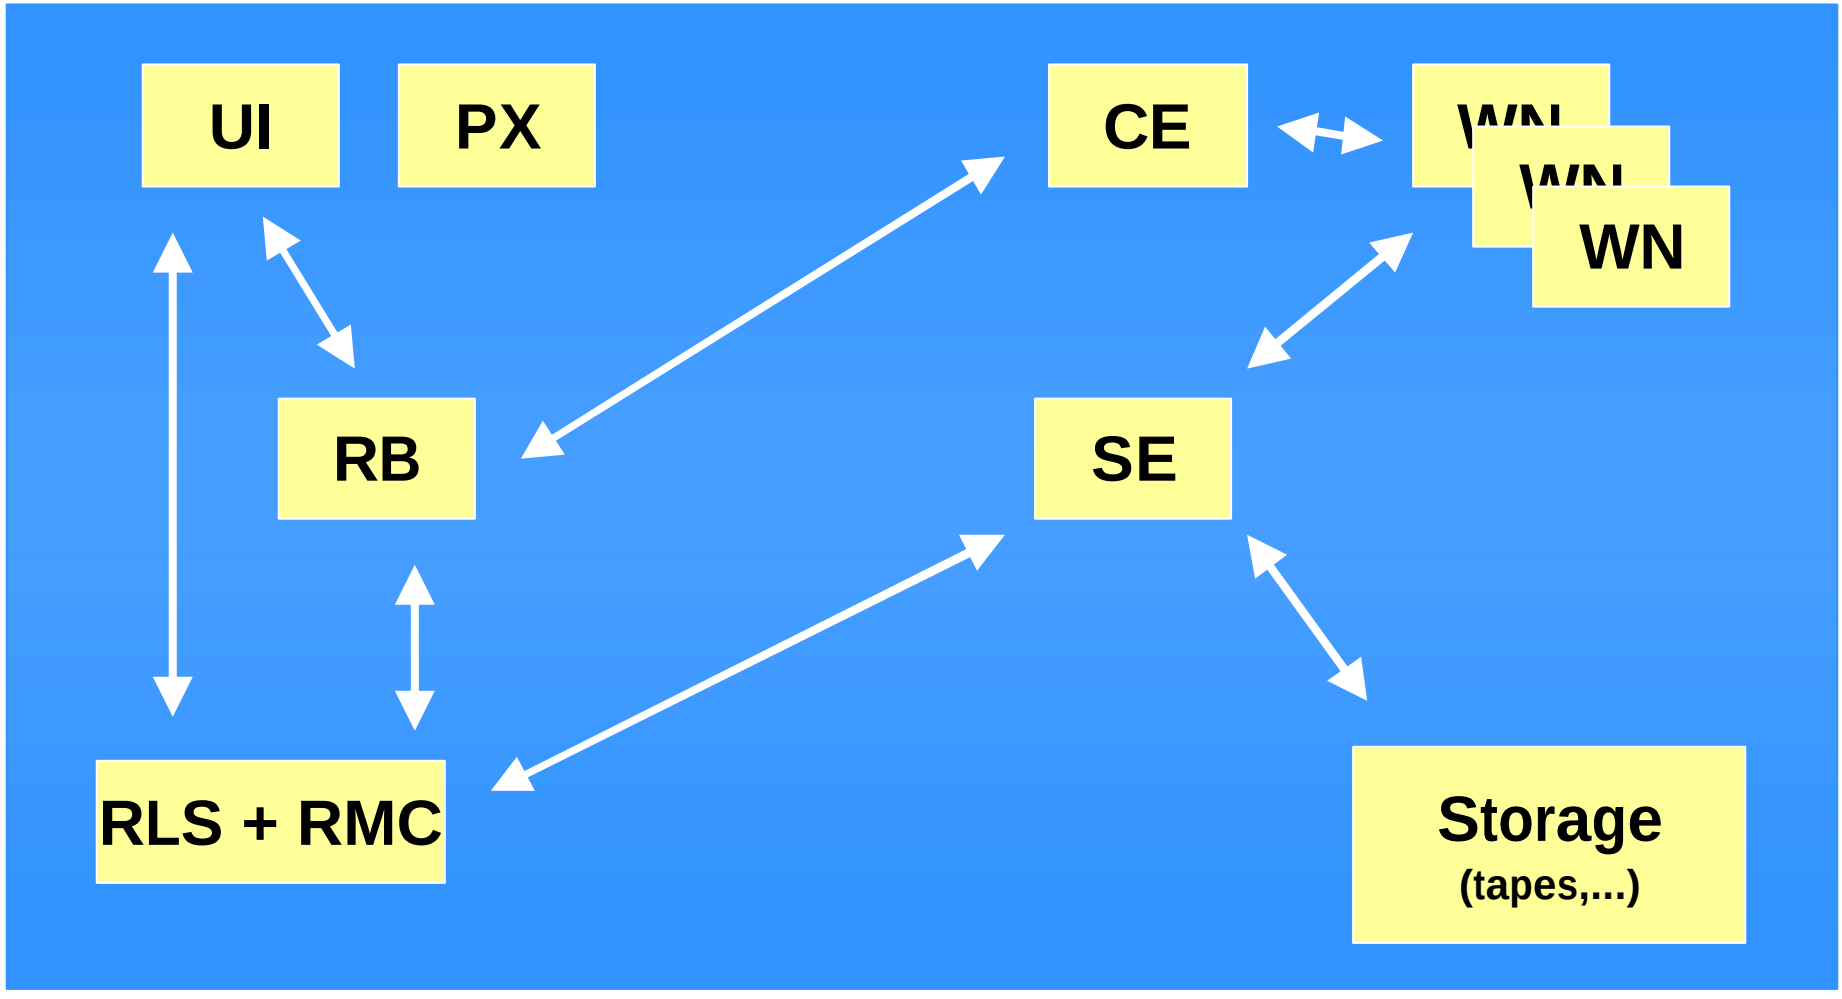

Information System

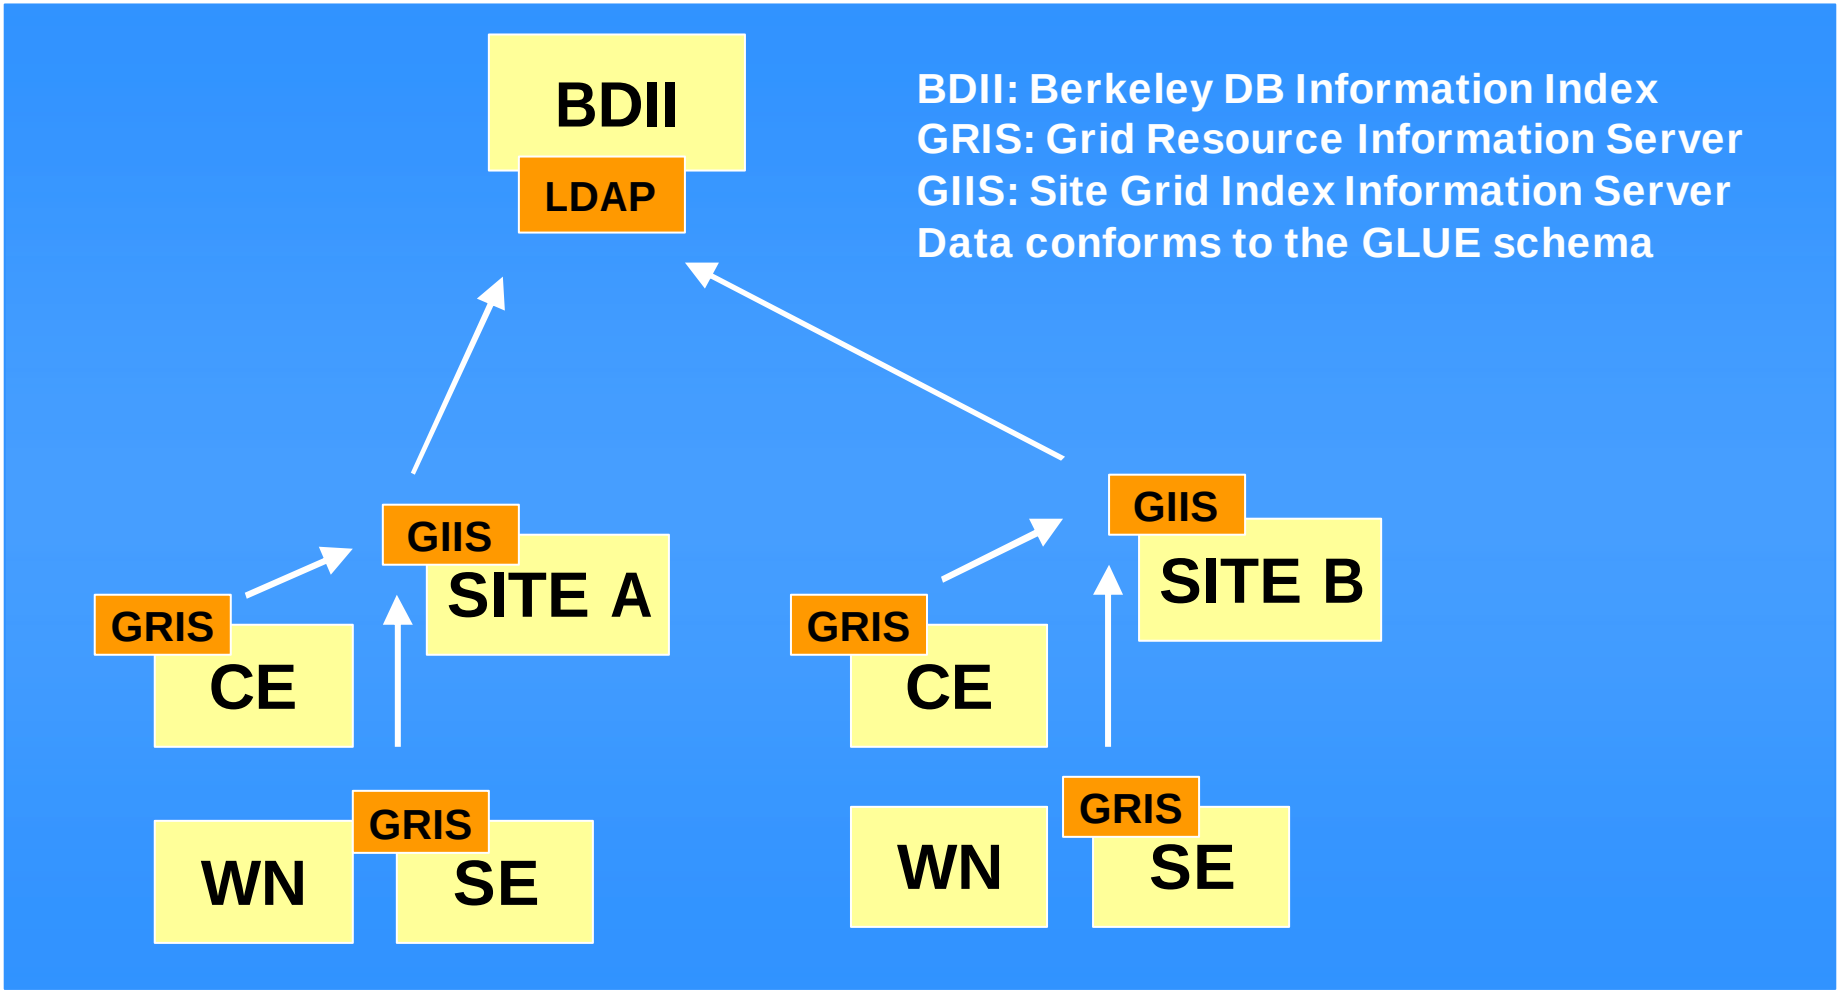

Monitoring

- R-GMA: implementation of the “Grid Monitoring Architecture“ as defined by the GGF. Gives the user the impression that all the information is contained in one large relational database. Currently in use with LCG.
- GridICE: uses the information system (GRIS,GIIS). Also distributed with LCG but only used by CMS (so far I know).

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

What is in LCG ?

- VDT Globus (globus 2.4 + patches)
- MyProxy
- Condor & Condor G
- EDG Workload Management
- EDG Logging and bookkeeping
- VOMS (virtual organization membership service)
- EDG Java security
- EDG Replica Manager and Replica Location Service
- EDG Information Service (R-GMA)
- GridICE
- LCAS (Local Center Authorization Service)
- LCMAPS (Local Credential Mapping Service)

- LCG Data Management Tools
- OpenPBS/Torque (batch system)
- CASTOR (hierachical storage management system)
- dCache (a kind of distributed filesystem)
- MySQL v4
- Several PERL modules
- Several Java modules

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Installation

- LCFGng: no longer supported for SL3.
- YAIM: apt-get + bash shell scripts.
- SmartFrog: only for IA64, under constant development at openlab. Does about 85% of installation work automatically. Some tasks still have to be done manually (adding users, crontabs, firewall, NTP...).

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

LCG on IA64

- Now ported & certified.
- Officially only announced for WNs but other nodes also work well.
- Download from openlab homepage.
- The code changes are in cvs thus we can follow up new releases quickly.
- Installation with SmartFrog. YAIM not yet working.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

LCG on IA64

- Next official LCG to be released soon (LCG_2-4-0).
- Current IA64 release is „between“ LCG_2-3-0 and LCG_2-4-0.
- We will synchronize with LCG_2-4-0.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

LCG on IA64 deployment status

Active deployment

- HP Puerto Rico (10 single CPU)
- PSNC Posznan (50+ nodes)
- HP Bristol (60 dual CPU machines)
- ISUFI/CACT Univ. of Lecce (some nodes)

Asked for IA64

- CYFRONET Cracow (20 HP rx2600)
- NGO Singapore (39 nodes)

gLite

Next generation grid middleware produced by EGEE.

Hardening & re-engineering of software that is already in LCG but also new components.

The plan is to deploy it in parallel with LCG and then move gradually to gLite (if it works well...).

In Feb. 05 LCG started to look at the first gLite release.

EGEE claims that gLite should work on IA64... we will have a look at it.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

gLite components

What‘s in gLite ?

I include information which has been presented by Frédéric Hemmer at the EGEE review in Feb. 05.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

gLite software stack

Computing Element
Gatekeeper (Globus)
Condor-C (Condor)
CE Monitor (EGEE)
Local batch system (PBS, LSF, Condor)

Workload Management
WMS (EDG)
Logging and bookkeeping (EDG)
Condor-C (Condor)

Storage Element
File Transfer/Placement (EGEE)
gLite I/O (AliEn)
GridFTP (Globus)
SRM: Castor, dCache

Catalog
File and Replica Catalog (EGEE)
Metadata Catalog

Information and Monitoring
R-GMA (EDG)

Security
VOMS (EDG)
GSI (Globus)
Authentication for C and Java based (web) services (EDG)

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Main differences to LCG

- Workload Management System works in push and pull mode
- Computing Element moving towards VO based scheduler guarding the jobs of the VO (reduces load on GRAM)
- Distributed and re-factored file & replica catalogs
- Secure catalogs (based on user DN, VOMS certificate being integrated)
- Scheduled data transfers
- SRM based storage
- R-GMA improved API and registry replication
- Move towards web services
- Some additional services as prototypes (Grid Access Service, Package Manager,...)

Part 4: SmartFrog

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

LCG installation with SF

1. Get list of rpms, download & install (resolve rpm conflicts manually)
2. Local site specific configuration file
3. Get templates for configuration files and write them to local disc with content according to local configuration file
4. Execute final scripts
5. Some things have to be set up manually (firewall, crontabs,…)

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

LCG installation with SF

You only need SF and the configuration file on the machine to be installed. Everything else can be downloaded.

Includes also undo components which delete all the new files.

Includes a component to add PBS nodes.

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

Next plans for SF

- Further develop PBS component (needs JSch)
- Install/uninstall whole WN
- Add „dynamics“: add/remove WNs based on demand (query batch system, RB, Tycoon,...); probably use SLP to discover free machines
- Eliminate the use of UNIX shell where possible (use Java methods instead)
- „Native“ rpm support (summer student)
- Components to handle configuration files like /etc/services (key/value pairs)

CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

SF thoughts

Rollback features, add an „undo“ method to the component

Better message sending between components (not just a string, maybe an object ?)