Summary of Atlas Database Meeting
2 September 1997

Agenda:

Introduction - RD Schaffer

Ongoing Atlas activities:

- Work for Tile test beam - Oleg Solovianov
- DAQ bookkeeping db - António Amorim
- Event storage - RD Schaffer

Report on activities in CDF - Simona Rolli

RD45 activities:

- Report from the July workshop - Jamie Shiers
- HepODBMS preview release - Dirk Duellmann
- Managing the type ids for multiple packages - Dirk Duellmann
- Replication and distribution studies - Eva Arderiu Ribera

RD Schaffer
RD.Schaffer@cern.ch
4 September 1997

Introduction - RD Schaffer

There are three domains of database activity within Atlas:

Event storage

Detector Description (for both simulation and reconstruction)

Calibration and alignment

Organization:

An Atlas database group was formed at the beginning of the year with myself as organizer.

The primary goal of this group is to gain experience with Objectivity by applying to both offline and test beam activities.

In June, people were assigned to be domain coordinators for the above:

RDS for events,

Katsuya Amako for det descr, and

Juergen Knobloch for calib and align.

RD Schaffer
RD.Schaffer@cern.ch
4 September 1997

State of progress:

Activity has concentrated on detector description and event storage.

Requirements for the domains:

event - 1st cycle complete and reviewed

detector description - 1st version still developing

Design and code - in a prototyping phase

Christian Arnault, et al., have developed ideas for storing and fetching objects according to a simple (hierarchical) description.

This can in particular be useful for accessing detector description information,

A first prototype (ADb) has been developed.

This is beginning to be tested on the Atlas muon database, and

the Tile Cal test beam (see Oleg Solovianov's talk)

RD Schaffer

RD.Schaffer@cern.ch

Other activities:

Objectivity server - a budget request has been submitted at CERN for a server machine and disk space for 1998.

This server will be interfaced to a general HPSS service (late) in 1998.

Test beams: for LArg and Tiles, discussions will soon be initiated with those responsible for the test beam event building.

RD Schaffer
RD.Schaffer@cern.ch

Design of Run Bookkeeper System for DAQ

Antonio Amorim and Helmut Wolters

Archives information about physics data recorded to permanent storage:

- DAQ and detector configuration
- tape management and logbook information
- beam characteristics

Will provide a Java applet (WWW) access to the information.

Have

- set of requirements
- an object design
- working on the implementation

They have also made some tests comparing remote access to Objectivity with and without using Corba server.

Progress report on Event Domain - RD Schaffer

Work has begun with prototype testing:

Concentrating on:

- how to organize the access to the information WITHIN and event

Has not yet begun the work on:

- how to organize events into collection and the collections themselves

The firsts tests have used Christian Arnault's hierarchical description prototype:

- Dice/Geant3 calorimetry data is now being stored

- Should be able to have some first tests for users wanting to read Geant 3 data sometime end October/ beginning November.

Work with Objy on HP is practically excluded until the new aCC compiler is supported - should be October.

Report on activities in CDF - Simona Rolli

CDF and D0 are both evaluating how they will treat their data management for run II

20 x the run I data volume:

250 TBytes/yr raw data

135 TBytes/yr reconstructed (DST)

80 TBytes/yr physics analysis (PAD)

Typical data access pattern for analysis:

- start from PAD -> ntuple
- analyse ntuple, but
- must go back to DST for some events

some analyses need

a small amount of information from many events

and others need

a large amount of information from a few events

New element in Run II: there will be little data replication

RD Schaffer
RD.Schaffer@cern.ch
4 September 1997

They are currently evaluating the use of Objectivity using Run I data.

Two approaches to putting data into Objectivity:

1. TRYBOS (bank manager) on top of Objectivity

   provides “backwards compatible analyses” access to banks stored in Objectivity

   but constrains the structure of the data and limits the use Objectivity features

2. Direct access to Objectivity

   using tag database classes from RD45 (see later)

   can also feed bank-style analyses if additional interface is written

Evaluation work will continue and a decision on the use of Objectivity will be taken this autumn or next spring.

RD45 Activities: Report form the July workshop - Jamie Shiers

Enhancement Requests to Objectivity

- scalability questions: HPSS interface, > 6 PB federations
- schema-related/multi-user/developer/partition issues

Next releases:

- aCC October 1997
- STL spring 1998

ODBMS+MSS Guide:

As part of next year's status report, they will provide a User's guide, which will be a living document

Beginning to consider the question of production services for 1998 for both groups (simulation/test beam) and user data (ntuples/histograms).

RD Schaffer
RD.Schaffer@cern.ch
4 September 1997

RD45 Activities: HepODBMS - Dirk Duellmann

Class library intended to:

- Insulate HEP applications from the ODBMS
- Provide facilities to simplify:

  database session and transaction management,

  clustering,

  event collections,

  selection predicates,

  tagDB access and

  indexing

A *preview* release is currently available and people are encouraged to use it and give feedback.

Software is available from (access from a cern node):

/afs/cern.ch/sw/lhcxx/specific/@sys/HepODBMS/0.0

html documentation is in HP branch.

RD Schaffer
RD.Schaffer@cern.ch

RD45 Activities: “Deployment” in HEP - Dirk Duellmann

Schemes for software development using Objectivity are currently being worked out.

Most likely, multiple federations will be needed, e.g.:

- development federation - full schema, no data, changes tested here before the production federation
- production federation - contains full schema and standard data
- user development/test federations - part of schema, some data

Type (or class) identification is done via a *unique type number*

This type ID is generated by the Objy preprocessor and stored in both:

1. the federation’s “schema” definition, and
2. in the class header files

*** This implies that one *must assure* that the same type IDs are being used in different federations so that

RD Schaffer
RD.Schaffer@cern.ch

compiled code (libraries) and data (db's) can be moved from federation to federation.

A scheme has been worked out for how to correctly preprocess .ddl files (see Dirk's transparencies)

This must be incorporated into out build facilities, i.e. SRT.

RD Schaffer
RD.Schaffer@cern.ch

RD45 Activities: Replication and distribution studies - Eva Arderiu Ribera

Eva gave us an update on the fault tolerance and data replication options.

She discussed two problems, one of which is solved:

1. **Re-synchronization when partitions come back online**:

   When taking a partition offline and then putting it back online, the re-synchronization was not automatic.

   This was solved with help from Objectivity, and a specific function must be called to re-synchronize.

2. **Replication via tape:**

   Replication of databases to a different autonomous partition has been tried successfully via network, which is the default means of transport.

   When trying to do the same while transporting via tape, Eva was able to do so only for the first replica/copy, and it is not currently possible to do so for the 2nd or greater replicas.

RD Schaffer
RD.Schaffer@cern.ch
4 September 1997