

# ATLAS

## Plans for Distributed Analysis and Expectations for ARDA

ARDA workshop  
CERN

David Adams  
BNL  
January 21, 2004



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



# Contents

Scope

Current ATLAS projects

Strategy

AJDL

Scenario for the first release

Plan for the first release

Deliverables for the first release

Expectations for ARDA

Conclusions

References



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



# Scope

## Analysis (not necessarily distributed)

- Supports the **manipulation** and **extraction** of summary data (e.g. histograms) from any type of event data
  - AOD, ESD, ...
- Supports user-level **production** of event data
  - e.g. MC generation, simulation and reconstruction

## Distributed analysis

- Extends the **extraction** and **production** support to include distributed processing, data and users
- Natural extension of non-distributed analysis
- Easily invoked from any ATLAS analysis environment
  - including Python, ROOT, command line
  - easily ported to any other environments (e.g. JAS)



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 3

# Current ATLAS projects

A number of projects relevant to distributed analysis and production already exist in ATLAS

- DC2 production system
  - Supervisor/executor model allows grid-based production of data on different grid flavors
  - Includes DMS: file replica service combining all the ATLAS replica catalogs
- The AMI Database
  - Just introduced a web service interface to bookkeeping
- DIAL: Distributed Interactive Analysis of Large Datasets
  - Running analysis web service with ROOT-based client
- GANGA
  - User-friendly interface for distributed Athena/Gaudi processing

These constitute the starting point for evolution to ADA



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY

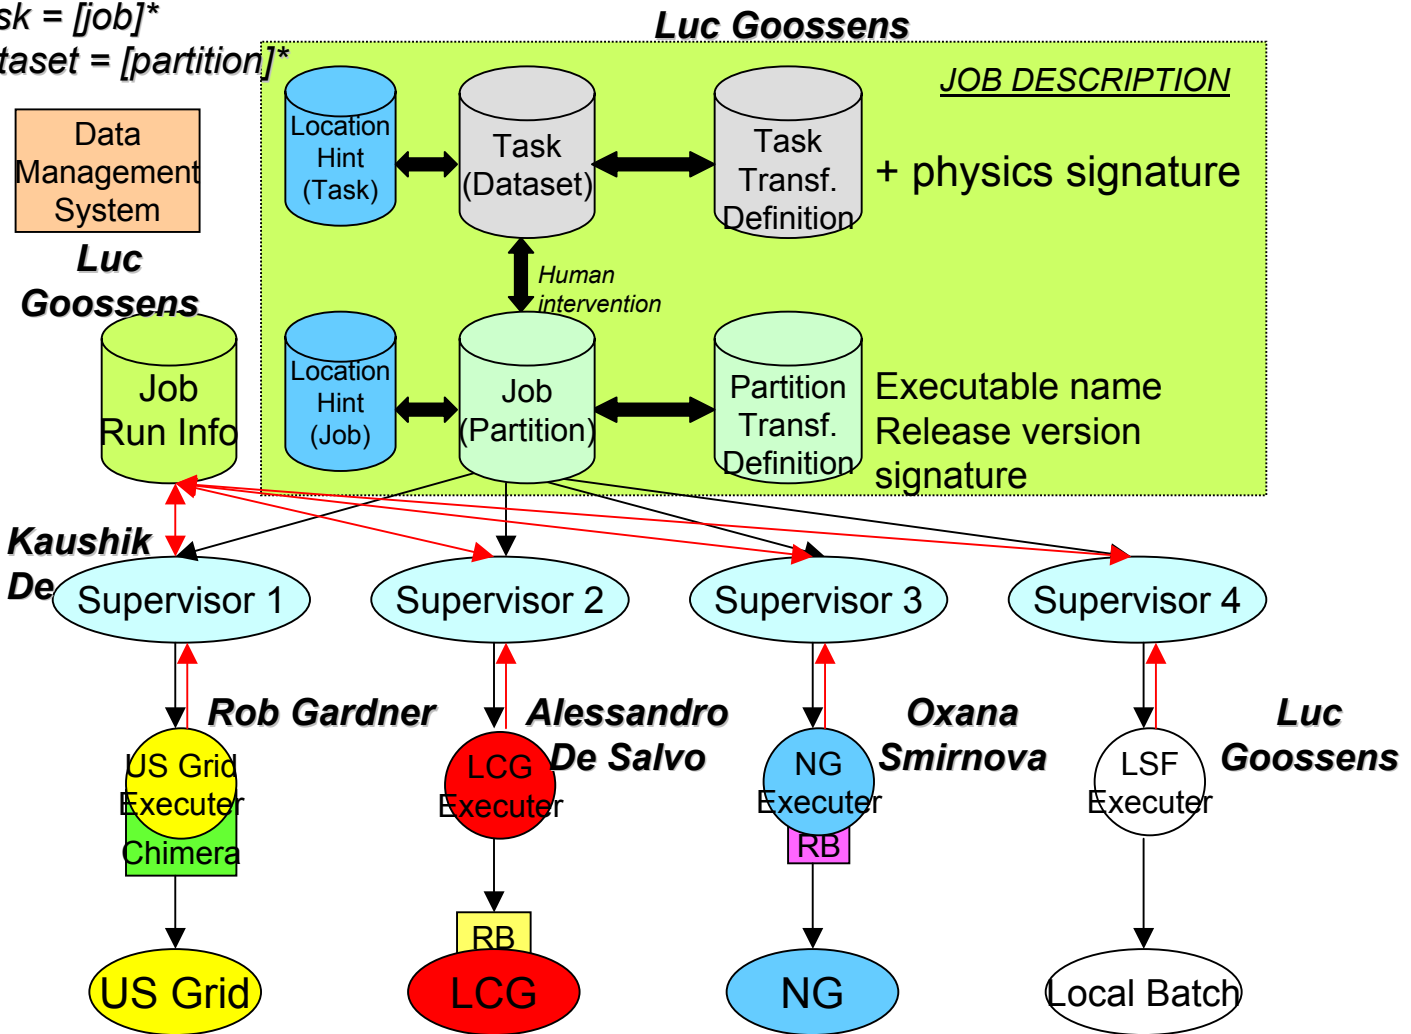


ATLAS Plans and Expectations ARDA workshop January 21, 2004 4

# ATLAS DC2 Production System

Task = [job]\*

Dataset = [partition]\*



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



# Strategy

Implement ADA as a collection of grid services

- As described in ARDA document
- Use ARDA components where possible
- Add missing and ATLAS-specific pieces

First identify the high-level services with which users interact directly

- Job submission (aka analysis service)
- Catalogs

Define language for interaction with these services

- AJDL: Analysis Job Description Language
- Objects exchanged between clients and services
- Service interfaces expressed in terms of these objects

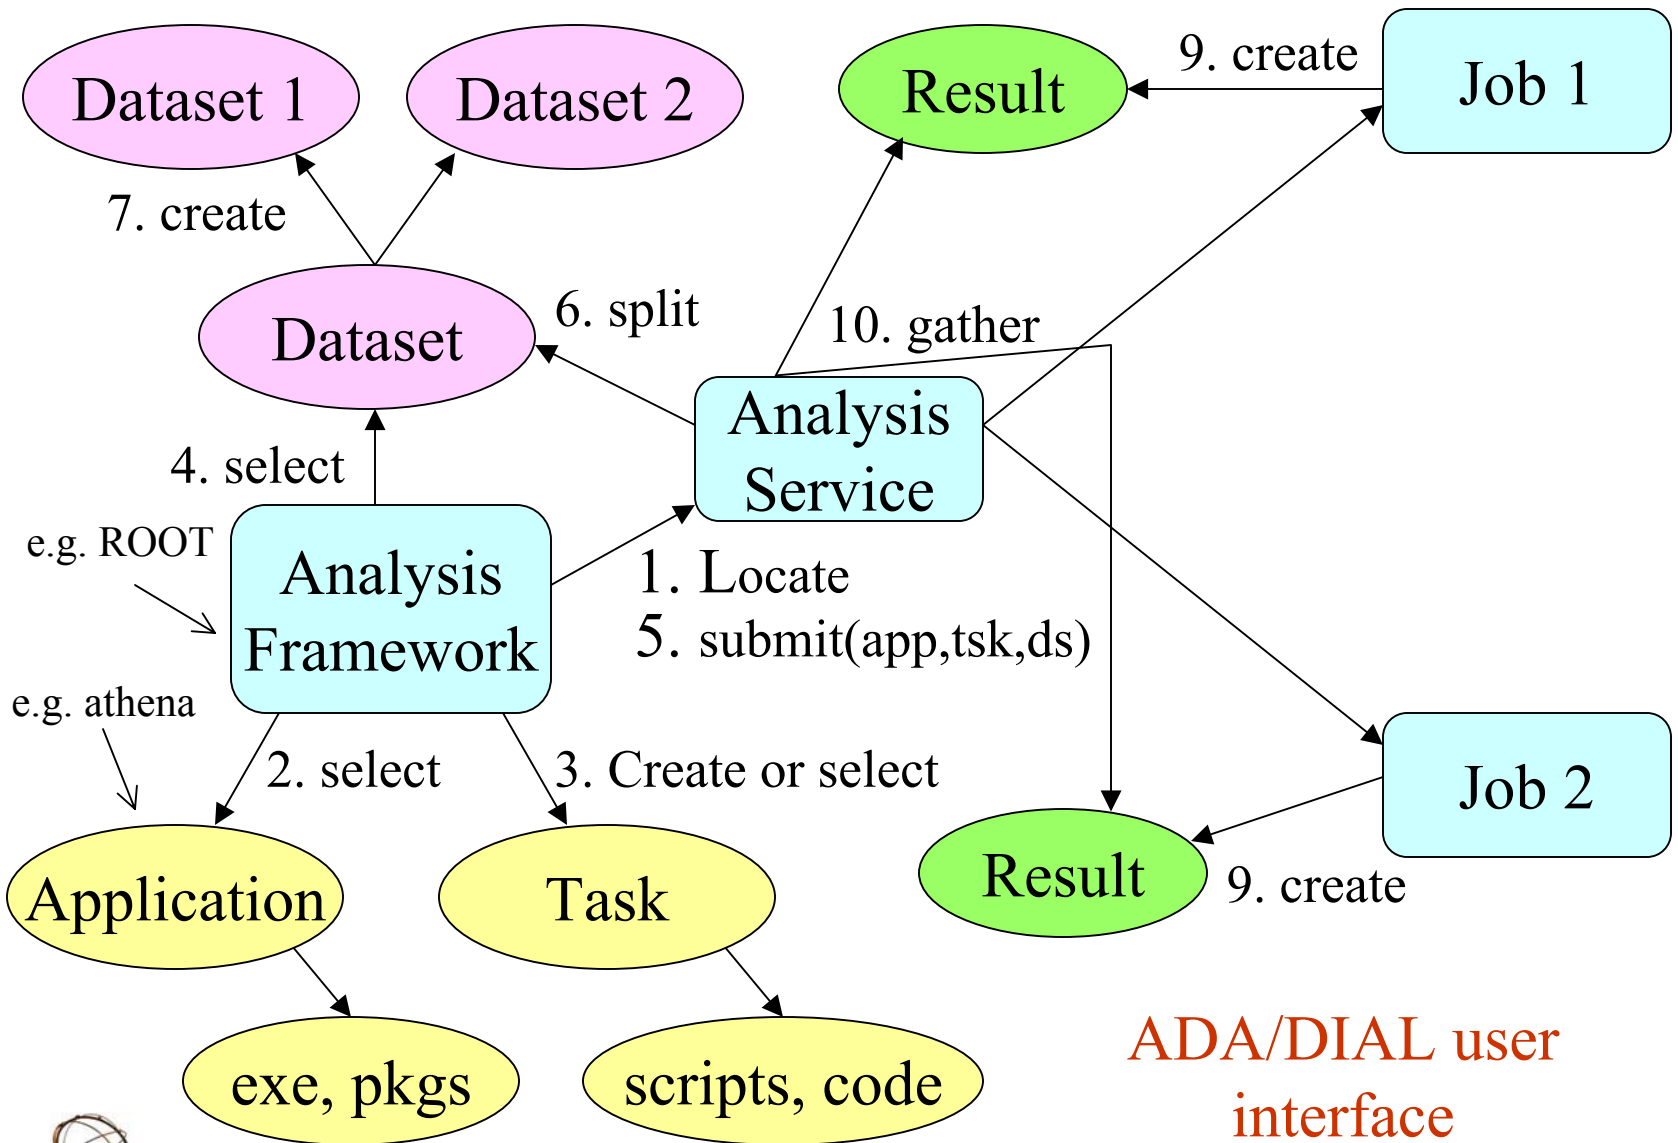


David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 6



ADA/DIAL user interface



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



# Strategy (cont)

## Implement high-level services

- Quickly to get user feedback
  - Reuse existing software where possible
- Re-implement as requirements become clearer and new middleware becomes available

## Provide clients for ATLAS analysis environments

- Python, ROOT, command line
- Also web portal and java

## Regular releases

- Perhaps for each SW week and ATLAS X.0
- Expand functionality with each release



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY

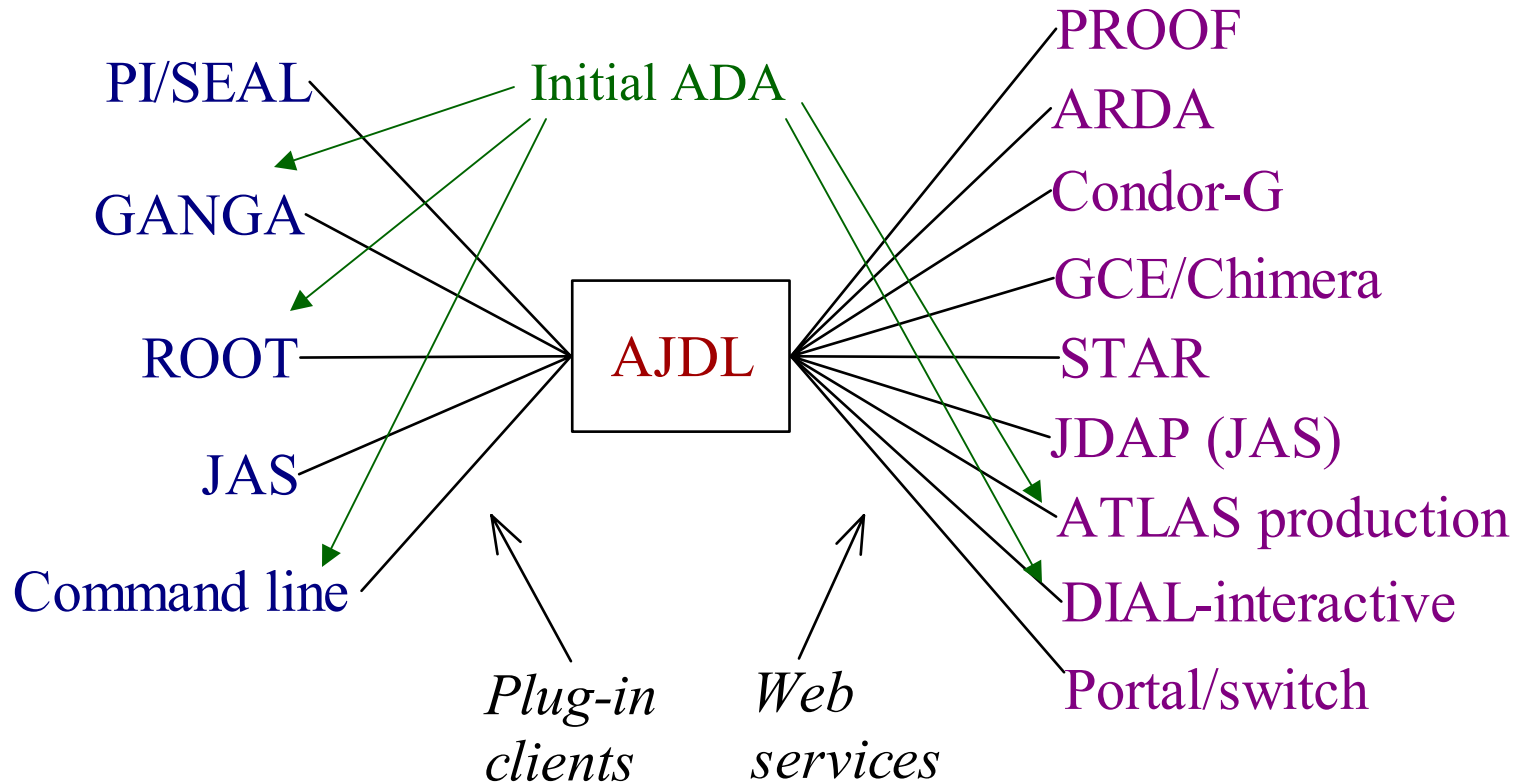


ATLAS Plans and Expectations ARDA workshop January 21, 2004 8



*Analysis environments*

*Analysis services*



## High level JDL as a bridge



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 9

# Strategy (cont)

Look to common projects for most of the pieces

- ARDA, GANGA, DIAL, ...
- Share as much as possible with ATLAS production
  - Also distributed
  - Similar interfaces and code for bulk and user-level production
- ADA must identify these pieces and tie them together

## Deployment

- ADA services must be deployed
  - Interactive service at one or two sites with data
- Provide testing and monitoring of these services
- Work with facilities to deploy and maintain
  - Also to develop facility-specific features
  - Looking for 1 or 2 initial sites for interactive service



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 10

# AJDL

## Base object types

- Application
  - Ultimately used to locate executable that processes data
  - Might be name plus version
  - Package management service used to install required packages and locate exe
  - Also provides mechanism to build (e.g. compile) task
- Task
  - Enables user to configure the application
  - Might be a collection of named text files
- Dataset
  - Specifies the data to be processed
  - Many possible implementations



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 11

# AJDL (cont)

## Base types (cont)

- JobConfiguration
  - Response time, budget, output data location, ...
- Job
  - Defined by application, task, dataset and configuration
  - Status (running, done, failed, ...)
  - Start, stop and update times
  - List of sub-jobs
  - Access to result (partial if job is not done)
  - Subtypes might add site, host machine, queue name, ...
- Result
  - Access to the data produced by the job
  - Probably this should be a dataset



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



# AJDL (cont)

## Analysis service

- Interface expressed in terms of these objects
  - has\_application(Application) : bool
  - install\_application(Application) : Status
  - has\_task(Application, Task) : bool
  - install\_task(Application, Task) : Status
  - submit(Application, Task, Dataset, JobConfiguration) : JobId
  - job(JobId) : Job
  - kill(JobId) : Status
  - and more ...

## Catalog services

- To store and select objects



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 13

# AJDL (cont)

## Nature of the objects

- Need subtypes to extend model for different applications, objectives and types of data
  - Inheritance and multiple inheritance
  - Intermediate types to allow common services
- XML representation for persistence and SOAP
- Also need methods
  - Associate classes with object types
    - > C++, python, java, ...
    - > Need bindings for all?
  - Services taking objects as arguments
  - Objects themselves could be service instances
  - Some combination of these



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 14

# AJDL (cont)

## Time scales

- Like to have broad participation so long term specification is deliberately vague at this point
- ATLAS would like to release a system soon so quick and temporary choices will have to be made over the next month
- DIAL has a C++ implementation
  - Similar to emerging specification
  - Has conflicts



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 15

# Scenario for the first release

Here is a scenario for user interaction with the first release of ADA

- Authenticate
  - Proxy from authentication service
- Choose application
  - E.g. PAW to process DC1 ntuples or
  - Athena to process DC2 AOD or
  - Athena reconstruction
- Define task
  - Analysis: provide code to define and fill histograms
  - Production: athena job options, maybe code
  - Perhaps select starting point from task catalog
- Select input dataset
  - From dataset (metadata) catalog service



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 16



# Scenario for the first release (cont)

- Create job configuration
  - Response time, role, optional splitter,...
- Locate processing service
- Submit job
  - Application, task, dataset, configuration
- While job is running
  - Query service for status and partial results
  - Examine partial results (e.g. histograms)
  - Kill job if results are bad
- When job is finished
  - Examine complete result
  - Modify task or select new dataset and repeat

- *Interactive* if job takes seconds or minutes



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 17

# Plan for the first release

## Schedule

- Implement and deploy by Spring 2004

## Building blocks

- Code and developers in GANGA and DIAL
- DC2 production system
- AMI cataloging services
- ARDA
  - Don't wait but stay closely coupled to that project
- Open to contributions (including effort) from others

## Deliverables in the following section

- These support the previously outlined scenario



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 18

# Deliverables for the first release

## Authentication service

- GSI based
- Support both EDG and US certificates

## AJDL – Analysis Job Description Language

- Start from current DIAL interface
- Incorporate ideas from PPDG, ARDA, ...
  - If available in time
- Goal: enable construction of a generic analysis service
  - Assume data made up of independently processed events
  - Few other assumptions
- Job-based model is also applicable to production
  - Rename analysis service to transformation service?



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 19

# Deliverables for the first release (cont)

## Interactive analysis service

- Goal is “interactive” response time
- Initial implementation distributes jobs over one site
- Make use of existing DIAL service
  - Modify DIAL to use AJDL

## Batch analysis service

- Batch-like response
- Processing distributed over grid or grids
- Could start from ATLAS supervisor/executor developed for DC2

## Clients for analysis environments

- Command line, GANGA and ROOT



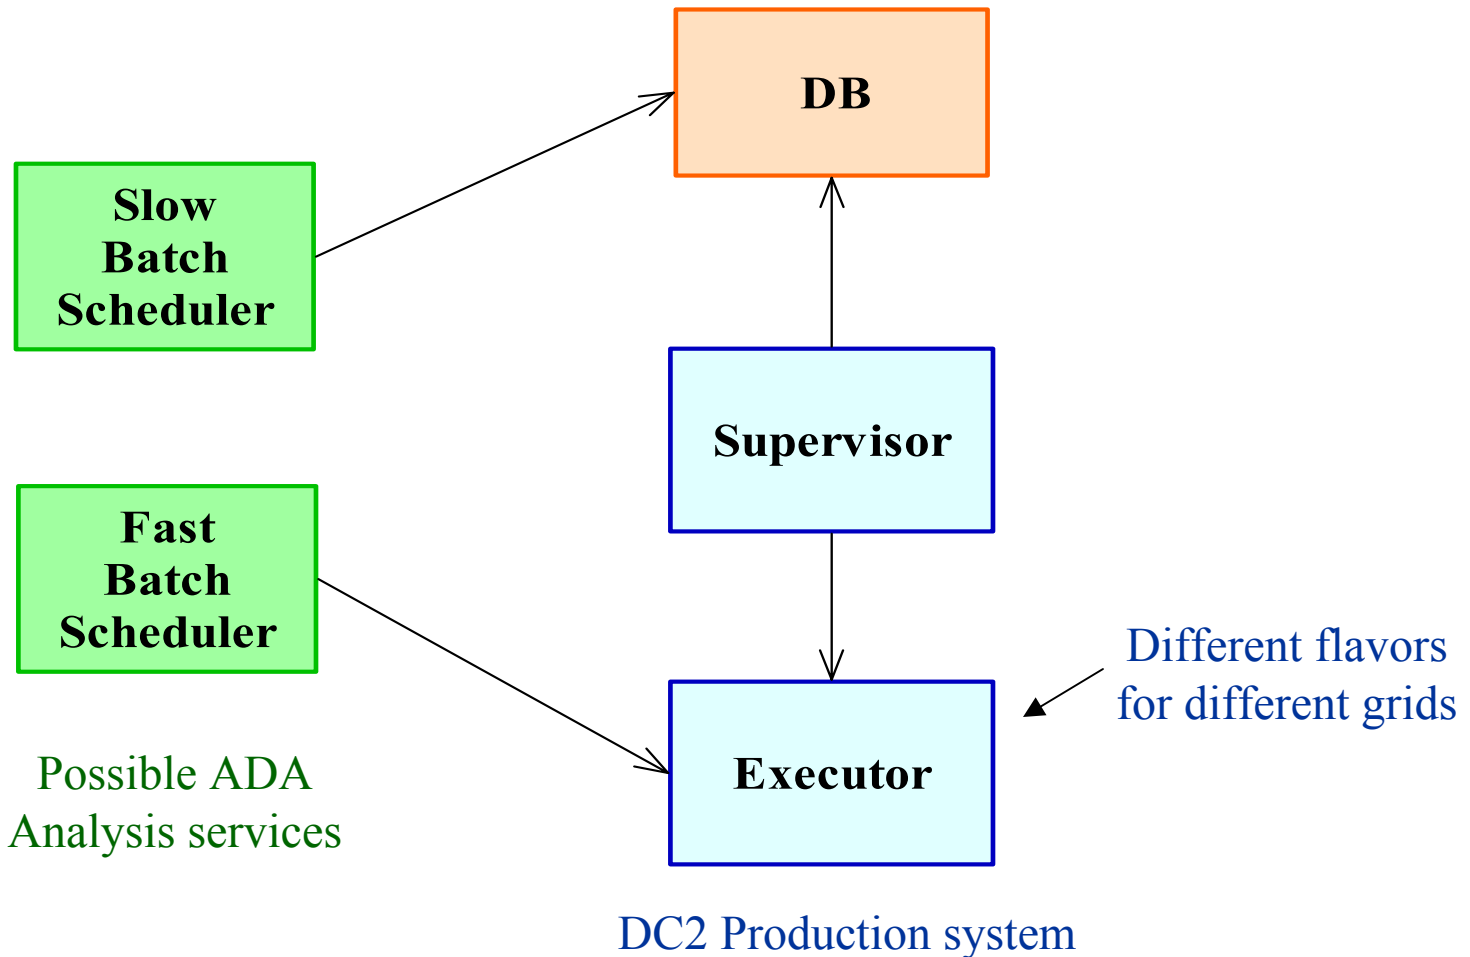
David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 20

# Possible connections to ATLAS production



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 21

# Deliverables for the first release (cont)

## ATLAS AJDL subtypes

- Enable useful system for ATLAS physicists
- Candidates summarized in the following table
  - DIAL already provides the first (PAW application)
  - Combined ntuple is current endpoint of ATLAS reconstruction
  - ATLAS is moving from combined ntuple to AOD

App.	Task	Dataset	Result
PAW	Kumac, fortran	CB ntuple (hbook)	Histograms (hbook)
Root	CINT macro	CB ntuple (root)	Histograms (Root)
Athena	C++ Algorithm	AOD	Histograms (AIDA?)
Athena	JobOptions	Raw data	ESD, AOD



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 22

# Deliverables for the first release (cont)

## Catalog services

- Repository
  - Repository to store object description indexed by ID
  - Applications, tasks, datasets, results and jobs
- Metadata (aka selection) catalogs
  - Enable users to make selections
  - Datasets and tasks
  - Includes virtual data catalog
    - > Prescriptions for creating a result (output dataset) from application, task and input dataset
- Replica catalogs
  - Enable user to find a concrete instance of a virtual object
  - Files and datasets



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 23

# Deliverables for first release (cont)

## Catalog service (cont)

- ATLAS bookkeeping is adopting a web service interface for DC2



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 24



# Deliverables for the first release (cont)

## Lower level services

- Eventually the analysis services will be constructed using lower level services
  - ARDA has done a good job identifying them and we will look there for delivery
- Initial releases will build on existing ATLAS components that are not service-based
- Likely will make use of the ATLAS services are being delivered
  - AMI catalog services
  - DMS file replica catalog



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 25

# Expectations for ARDA

## Service infrastructure

- Which service and service discovery frameworks should we be using?
  - Globus, Clarend, OGS::Lite, gsoap, ...
  - Answers for different time scales

## Authentication service

- Must support users around the globe
  - How to build list of ATLAS users?
- Support for roles
  - Ordinary user, physics coordinator, service administrator
- Much already in place: GSI, VOMS, CAS



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 26

# Expectations for ARDA (cont)

## AJDL

- Is this a useful interface?
- Are the object types best thought of as data objects, classes, services or some combination?
- What data and methods are associated with these types?
- What is the analysis service interface?

## What is a dataset?

- More than just a file or collection of files
- Connection to POOL collections
- Big part of AJDL



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 27

# Expectations for ARDA (cont)

## Deliver a generic analysis service

- With AJDL interface
  - Wrapper around ARDA prototype if ARDA does not embrace the AJDL abstraction
- Requires middleware outlined in ARDA report

## Deliver AJDL clients

- Python, CINT, java, command line, web, ...
- Simple user interfaces using these clients
  - Not too fancy; probably too difficult to satisfy all experiments



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 28

# Expectations for ARDA (cont)

## Catalog services

- Help identify types (repository, replica, metadata)
- Deliver these
- Includes virtual data catalog
- From S. Albrand:
  - “For Atlas Data Challenge 2, the AMI client is a web service. ARDA, or no ARDA this solution has a lot of advantages. The first version is ready for testing by “avant garde” users. This service is very simple.... What do we need to do to become an ARDA compliant service?”



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 29

# Expectations for ARDA (cont)

## Deployment of services

- Hardware and people

## Monitoring

- Infrastructure for monitoring services and data
- Agents to provide input to this infrastructure
- User tools for making sense of the data



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 30

# Conclusions

## ATLAS strategy

- Use grid service model
- Quickly define high-level service interfaces and implement services and clients
- Deliver end-to-end system to users
- Frequently re-design and re-implement based on user feedback

## ARDA collaboration

- Ideally we would come to consensus within ARDA on a high-level interface along the lines of AJDL and share end-to-end effort
- In any case, we will work closely with ARDA to define middleware services



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 31

# References

ADA (ATLAS Distribute Analysis) is described at

- <http://www.usatlas.bnl.gov/ADA>

Documents link on that page includes

- AJDL version 0.20
- “Datasets for the Grid” version 5



David Adams

**BROOKHAVEN**  
NATIONAL LABORATORY



ATLAS Plans and Expectations ARDA workshop January 21, 2004 32