

Datasets for the Grid and for ATLAS

ATLAS Software Workshop

Database Session

CERN

David Adams

BNL

September 24, 2003



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Contents

Functional definition of dataset

Dataset properties

Dataset categories

Dataset category associations

Properties and categories

Implementation

Development plan

Functional definition of Dataset

Define dataset by the way it is used:

- Dataset is the unit of data with which users normally interact

There are two use cases

- User selects a dataset for processing
- User hands dataset to a “system” that distributes processing, gathers and merges results, and returns the result to the user
 - Result can be another dataset or
 - Summary analysis data such as histograms

What properties are datasets required to have to satisfy these use cases?

Dataset properties

0. Identity

- Dataset must have an unique index and/or name

1. Content

- Description of the type of data in the dataset
 - Event or non-event data
 - Simulation, reconstruction,
 - ESD, AOD, ...
 - Jets, tracks, electrons,...

2. Location

- Where to find the data
 - Logical files, physical files, site,...

3. Mapping

- Which content is at which location?



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Dataset properties (cont)

4. Provenance

- Prescription for creating the data
- E.g. input dataset and transformation

5. History

- Details of production beyond provenance
 - How task was split into jobs,
 - Processing node and time for each job, ...

6. Labels

- Assigned metadata outside other categories, e.g.
 - Integrated luminosity
 - Result of quality checks
 - Flag indicating ok for use in published analyses



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Dataset properties (cont)

7. Mutability

- May dataset be modified?
- Possible states: locked, unlocked, extensible, ...

8. Compositeness

- Dataset made up of other datasets.
- Two cases:
 - Construction: provenance is the list of sub-datasets
 - > E.g. the summer dataset is defined to be the union of the June, July and August datasets.
 - Assignment: factorization into sub-datasets
 - > Typically to reflect data placement
 - > E.g. a representation of a global dataset might include sub-datasets in New York, Paris and Moscow.



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Dataset categories

Categorize datasets according to the extent of their location information:

- Virtual
 - no location
- Logical
 - Collection of logical files
- Physical
 - Collection of physical files
 - Inferred from logical DS and file catalog (Magda, RLS, ...)
- Staged
 - Collection of “jobs”
 - > each sub-dataset matched to CPU/process
 - Not important for discussion here



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Datasets... ATLAS SW – DB session

September 24, 2003

7

Dataset category associations

One-to-many association as we move down these categories

- Virtual dataset may map to multiple logical datasets
 - Optimize file size for local mass store
 - Copy out only selected events (vs. all plus event list)
 - Move data into a DB at one site
 - Composite representation along placement boundaries
- Logical dataset maps to many physical datasets
 - Many combinations inferred from file catalog
 - No need to record all these datasets
 - But system (or user) might record LDS used to process one task and reuse it for the next request



David Adams

BROOKHAVEN
NATIONAL LABORATORY

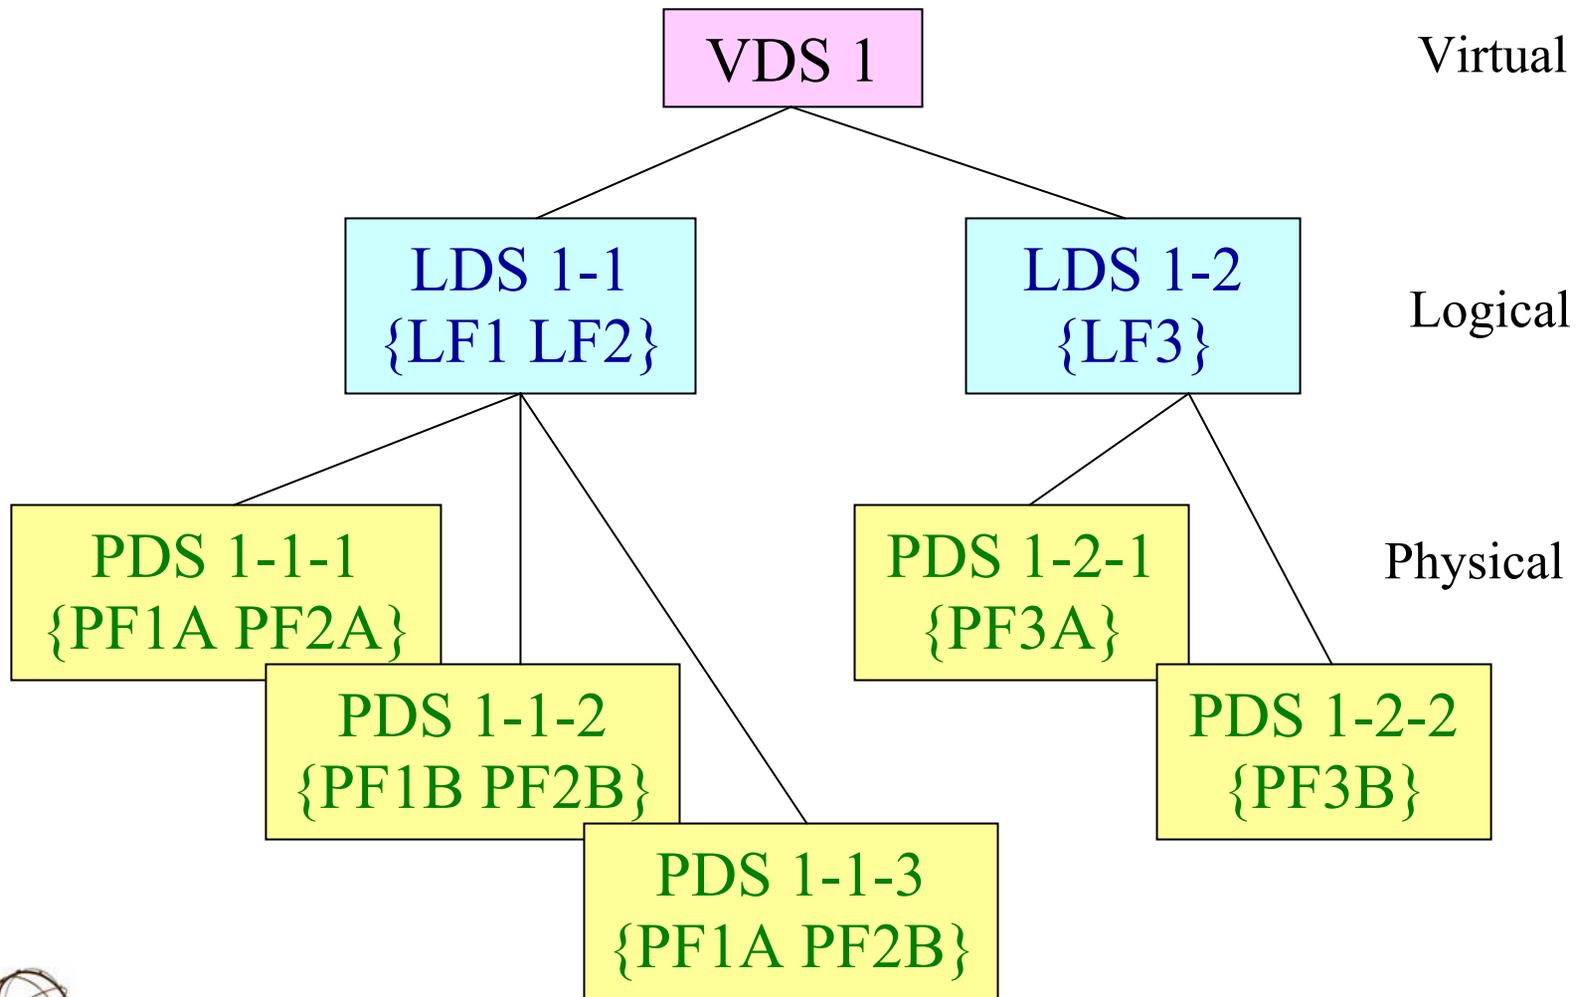


Datasets... ATLAS SW – DB session

September 24, 2003

8

Dataset category associations (example)



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Datasets... ATLAS SW – DB session

September 24, 2003

9

Dataset category associations (cont)

OO view

- LDS (logical dataset) “is a” VDS (virtual dataset)
- PDS (physical dataset) “is a” VDS
 - And “is a” LDS if files are in file catalog

Representation

- Use the word representation for this relationship
 - In the figure, LDS 1-2, LDS 1-2 are two different representations of VDS 1

Properties and categories

Virtual datasets have

- Content
 - Might have to have a representation to know all content
 - All representations have the same content
- Provenance
- Labels
 - Most are associated with the virtual view
 - Again may need representation to evaluate some labels
- Mutability
- Compositeness
 - Sub-datasets are also virtual
 - Composite because it was constructed by merging datasets

Properties and categories (cont)

Logical datasets add

- Location
 - Which logical files
 - Extend to add sites where LDS can be found?
- Mapping
- History
 - Perhaps stored in job tracking system
- Compositeness
 - Constituents are logical datasets
 - Typically they reflect data placement

Implementation

Dataset metadata catalog (DMC)

- Holds VDS properties
- User selects dataset based on these properties
- Receives a VDS name or ID
- DIAL wants programmatic interface for DMC

VDS class

- Programmatic interface to access VDS properties
 - E.g. what content is in a given VDS?
- Primary user is workload management system (WMS) rather than analysis users
 - E.g. to verify that a VDS has the content required for a given analysis task



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Implementation (cont)

Dataset replica catalog (DRC)

- Enables WMS to select a LDS representing a VDS
- Also get info about the sites where LDS can be found

LDS class

- Provide access to VDS, logical files (location) and sub-datasets (compositeness), if present
 - WMS uses this for splitting, matchmaking and staging (files)
- Means to select content or events to define an LDS with a subset of the data in the original LDS
 - New LDS has a set of files that is subset of original set
 - Used by WMS for splitting DS's and staging files



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Implementation (cont)

PDS/SDS catalogs and classes

- At a site, locate a PDS or SDS representing an LDS
- Not needed: WMS can use file replica catalog
- But PDS/SDS provide means to record the choices made for one task so they can be used for the next task using the same LDS
- However these pieces should wait to see what if (any) requirements come from WMS's

Following table shows where properties are recorded



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Implementation (cont)

Property		VDS	LDS	DMC	DRC
identity	virtual	X		X	X
	logical		X		X
content	# events	X	X	X	
	event ID's	X	X		
	event content	X	X	X	
location	logical files		X		
	site (with data)				X
mapping					
provenance	parent			X	
	xform			X	
history					
labels				X	
mutability		X	X	X	
compositeness	virtual	X		X	
	logical		X		X



Development plan

Define generic VDS and LDS interfaces

Implement VDS and LDS for CBNT files

- Existing CbntDataset is starting point

Define interface for DMC

Fill DMC for DC1 CBNT datasets

- Connect to AMI

Create trivial DRC for these datasets

- One LDS for each VDS

→ DIAL users can analyze all DC1 data

- Data must be staged at BNL and/or CERN
- Or file catalog does staging



David Adams

BROOKHAVEN
NATIONAL LABORATORY



Datasets... ATLAS SW – DB session

September 24, 2003 17