



AMI – Status April 2011.

*Solveig Albrand
Jerome Fulachier
Fabian Lambert*





Summary

- Server problems.
- ORACLE problems.
- Security & Information Protection.
- Developments.
 - General
 - Real Data
 - MC
 - Other applications
- Plans.

07/04/2011



In brief

- Server problems. Some instability since the beginning of 2011. See SIT Tag Collector talk for details. *(extra slides)*
- Security & Information Protection. We are moving to VOMS for authentication (unless ATLAS management says "No"). Time scale to be fixed. No time to discuss here. See SIT Tag Collector talk for details.
- ORACLE.
 - "Back-up test" : I dropped one of the config tag tables table by accident ; our DBA@Lyon got it back again.
 - The underscore/case insensitive sorting incompatibility bug manifested itself again in a new form, following the latest ORACLE update (10.2.0.4 → 10.2.0.5) but once we spotted it we were able to get the behaviour we need. We used to get **unpredictable** results, now get the **opposite** of what we expected. *(see extra slides for more)*



Dev – Dataset General

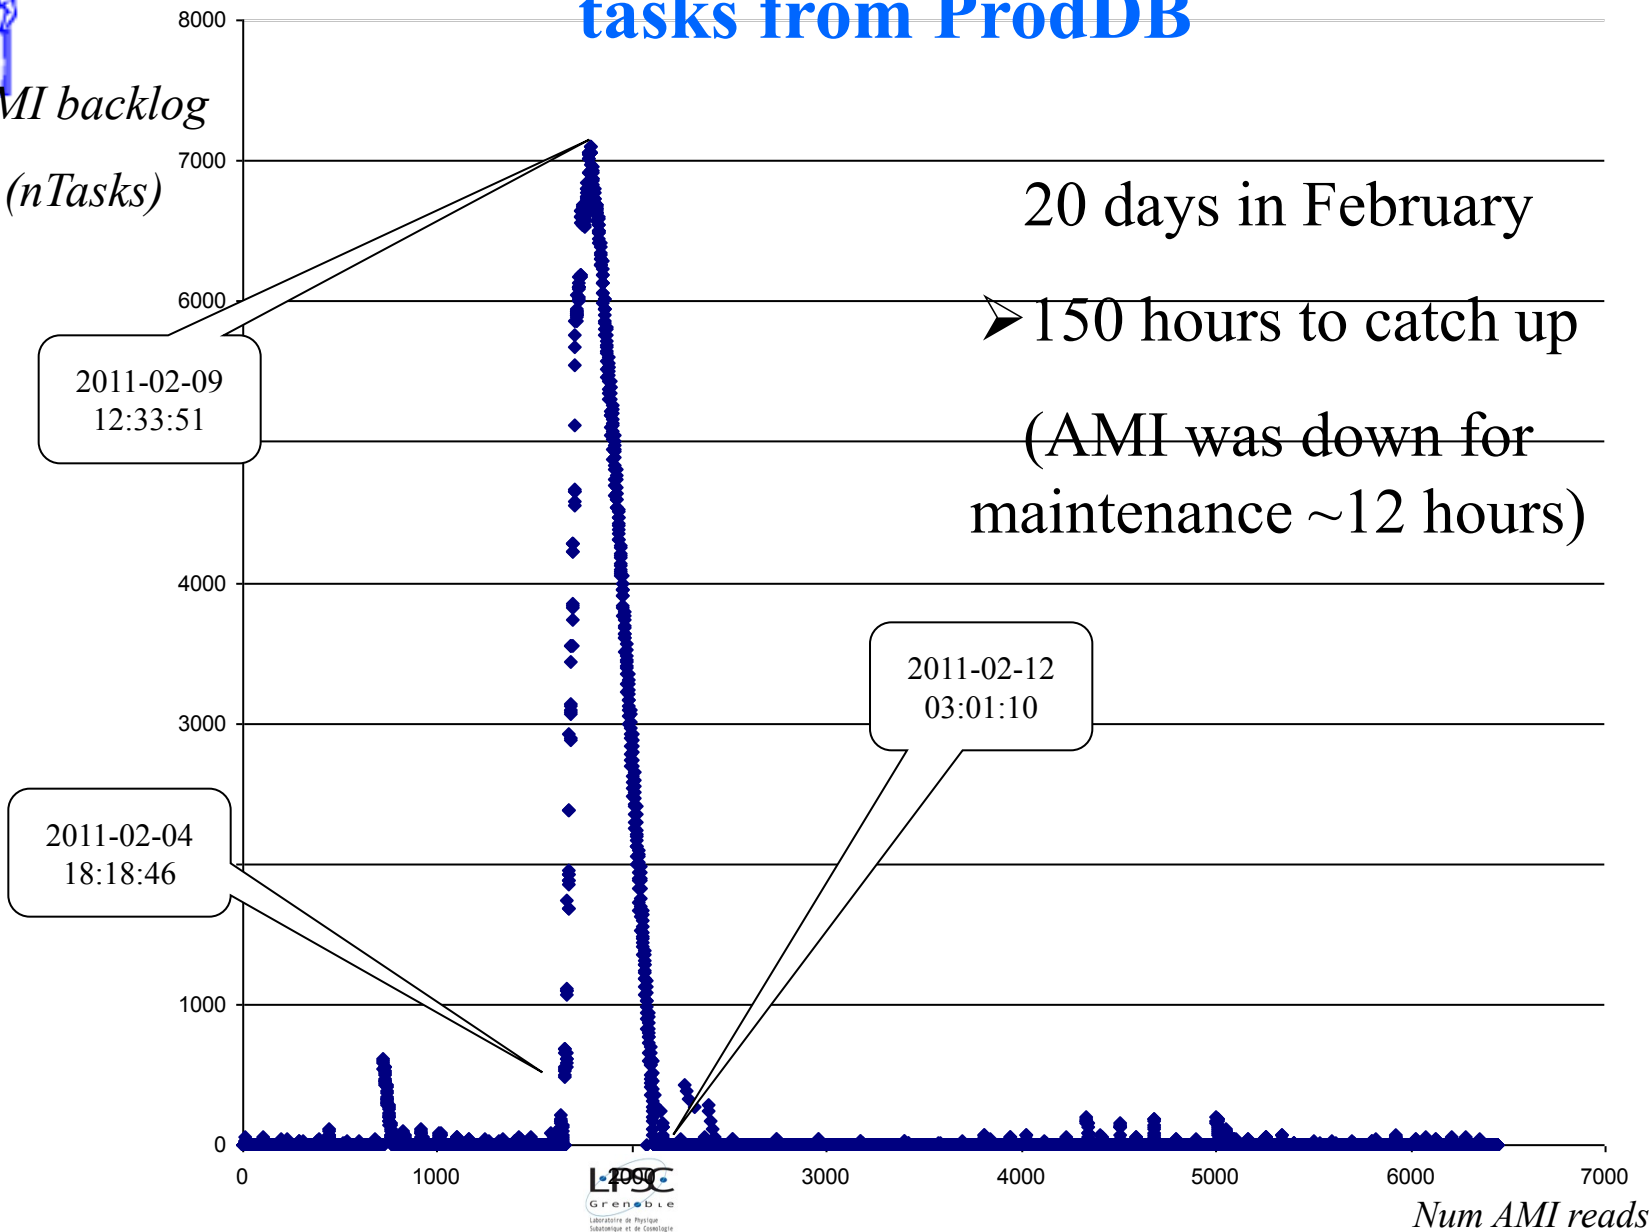
- A general view of metadata has been started. A document is in preparation (with metadata coordination). Will lead to some actions e.g. rework the AMI dataset state engine and remove panic-inducing states when data is deleted.
- Lost files - synchronized on DDM service. (see later)
- Scalability of reading prodDB (Reminder: We read metadata XML for all finished jobs for all finished tasks.)
 - Sequential since 2006. Knew it was not optimal, but that was not a problem up to now.
 - Had problem in February, so (at last) working on multi threaded reading of finished tasks. Not a panacea, because number of jobs in a task is not predictable, but ~ 50% improvement anticipated.
 - *WARNING – The graph on the next page has an "advertiser's" X axis (number of AMI reads). It doesn't mean anything much. The AMI task runs 300 seconds after it last finished – so not points are not evenly spaced in time in reality.*

Scalability of reading **FINISHED** tasks from ProdDB



AMI backlog

(nTasks)



20 days in February
 ➤ 150 hours to catch up
 (AMI was down for maintenance ~12 hours)

2011-02-09
 12:33:51

2011-02-12
 03:01:10

2011-02-04
 18:18:46



Real Data

- **Lost Luminosity Blocks.**
 - Lost files are marked once a week. (dq2 file consistency service)
 - **Lost files** are marked in **orange** in the file list, and removed from the event and file count. The dataset status is changed.
 - A comment is written to say when the file was lost.
 - All files in data10 and mc10 and up have been marked with their input file(s). Information is obtained from prodDB.ejobdefbig
 - The file to file provenance is traced recursively to obtain the lumi blocks which were in the lost file, and the information is stored.
 - The tracing is not 100% reliable:
 - ejobdefbig problems
 - with missing information,
 - Some surprises in the XML grammar ("inputESDFile=" but "inputTAGFile:",
 - badly formed XML,
 - deleted files mechanism in AMI. (this can be fixed !)
 - **What do I do now?** (*need guidance from data prep and/or luminosity group*) For example we could trace all file lumi blocks for data11 reprocessing.



Command Home Bookmark (this catalogue only)

dataset

1 - 1 / 1 order by modified - created dataset.created DESC

Help Options Edit Fields Advanced

Query : ((amiStatus=**VALID**' and logicalDatasetName like '%%') AND dataset.prodsysStatus=**EVENTS_AVAILABLE: MISSING LB**') AND dataset.dataType=**AOD**

additional Fields	logicalDatasetName	nFiles	totalEvents	dataType	prodsysStatus	period
details	data10_7TeV.00153565.physics_MinBias.merge.AOD.r1647_p306_p307 DQ2 - GANGA export - Provenance	350	21217326	AOD	EVENTS_AVAILABLE: MISSING LB LB_Lost	B1 All Runs - COMA - more Info

Command Home Bookmark (this catalogue only)

Command Home Bookmark (this catalogue only)

lost_lumi_block

1 - 3 / 3 order by lost_lumi_block.created DESC

Help Options Edit F

Query : dataset.logicalDatasetName='data10_7TeV.00153565.physics_MinBias.merge.AOD.r1647_p306_p307'

additional Fields	files.LFN	lumiBN	lastModified	modifiedBy	createdBy
details	AOD.198975__000297.pool.root.1	1022	2011-01-27 16:27:17	root	root
details	AOD.198975__000297.pool.root.1	1021	2011-01-27 16:27:16	root	root
details	AOD.198975__000297.pool.root.1	1020	2011-01-27 16:27:13	root	root

additional Fields	fileStatus	LFN	events	inputfile
details	LOST	AOD.198975__000297.pool.root.1	62682	AOD.194469__0003734,0003735,0003736,0003737,0003738,0003739,0003740,0003741,0003742,0003743,0003744,0003745].pool.root
details	VALID	AOD.198975__000039.pool.root.1	51743	AOD.194469__000636,000637,000638,000639,000640,000641,000642,000643,000644
details	VALID	AOD.198975__000040.pool.root.1	59146	AOD.194469__000648,000649,000650,000651,000652,000653,000654,000655,000656,000657
details	VALID	AOD.198975__000041.pool.root.1	54700	AOD.194469__000662,000663,000664,000665,000666,000667,000668,000669,000670,000671

file_comment

1 - 1 / 1 order by modified - created file_comment.crea

Query : files.identifier='10718332'

additional Fields	text
details	File F28D4E79-EBEF-DF11-ACEC-00A0D1E50B05 lost : From URL http://bourriat.cern.ch/dq2/media/monitoring/2010/49/lostFiles.txt



MC developments

Borut@ MD workshop


"Meta-data interface looks a bit technical for the end user"

- DONE
 - Transporting cross section values along the MC production chain (less clicks to get the values!) .
N.B. ~100 "physicsShorts" produce no value for cross section value.
 - Reworking the "dataset numbers" broker, and extending it to hold production requests in the future.
 - No longer reading the list of input parameters from Task Request (too many values are "NONE"). The reason is the hard coded argument list for job transforms. Get values only from metadata output of finished jobs, and the AMI tags.
- NOT DONE
 - Import of production requests from spreadsheet files; (we know how to do it but the input is too messy)
 - Pointers to job options files broken. (we lack a reliable way to do it)



Other Developments

- Data Periods :
 - Collaboration with COMA (Elizabeth G.) and Data Preparation (Beate).
 - Replaces text files



- Data is in the COMA database
- AMI "thinks" COMA is part of AMI
- Data Prep writes, several apps read



AMI interface

Home Searches Tools Bookmarks ?

Datasets Selection Account Validation

Hello Albrand, this page allows you to populate/add runs periods.

To edit an existing Period:

Period level "1" is selected
Project "data11_7TeV" is selected
Period Name "A1" is selected

Information about period/runs :

- Left column: runs that are part of the selected data period.
- Right column: runs that are not part of the selected data period.
- Middle column: Information about the selected data period.

Main Actions available:

- Lock/Unlock selected data period using checkbox in middle column.
- Filter on runs lists using top left and right column tools:
Type a pattern in text field, either hit return or click on go.
- Adding runs to selected (unlocked) data period:
Check wanted runs of right column, hit the middle < button.
- Removing runs from selected (unlocked) data period :
Check wanted runs of left column, hit the middle > button.

Create a new Period:

Click on the "Add new period link" in the middle column.

Key

= can't edit , = AMI info ,
 = delete , = save , = validate

Filter on runs already in period :

ALL
 177682
 177593
 177540
 177539
 177531

Edit an existing Period

Period Level: 1
Project: data11_7TeV
Period Name: A1

P_LEVEL	1
P_PROJECT	data11_7TeV (runs)
P_PERIOD	A1 (runs)
CREATE_USER	ate
P_DESC	First 2011 stable beams. Trigger being commissioned. No bunch trains: 2 nominal bunches 1 pilot bunch colliding in ATLAS
lock	<input checked="" type="checkbox"/>

Filter on runs to add in period :

ALL
 178109
 178108
 178047
 178044
 178026
 178024
 178023
 178021
 178020
 178019
 178017
 178012
 177990
 177986
 177966
 177965
 177964
 177960
 177929
 177928
 177924
 177912
 177911
 177910
 177904
 177823
 177759
 177648
 177622
 177591

Links to COMA

Runs loaded in COMA with selected project

See extra slides for more about COMA



Next steps for Data periods

- pyAMI commands for Data Period information (in beta testing)
 - GetDataPeriodsForRun
 - GetRunsForDataPeriod
 - GetDataPeriodTree
 - ListDataPeriods
- Document it all for users! (we advocate a written Period nomenclature)
- Extend to Physics Container creation.
- Other extensions in discussion.



Tracking of object sizes in reconstructed events.

- A new application in AMI
- In collaboration with SW dev. (Ilija Vukotic)
- Currently in test on Tier 0. If it works well we will find a way to extend it to Grid tasks.
- Has its own AMI/ORACLE resources
- Will lead to a new AMI graphics effort.




Other stuff.

- Fruits of the ADC retreat in Napoli
 - Can "inputfile peeker" mechanism be replaced by consulting AMI?
 - Can the configuration mechanism currently used by Tier 0 be extended to ProdDB tasks? See [Rod Walker's talk yesterday.](#)
- DA user survey – the comments on AMI are interesting but not directly helpful to us (we already knew not everyone likes our web interface). It would be better to complain directly – or better help us design a new interface!
 - "AMI web interface is awkward"
 - "AMI is also a bad tool, the web page is slow, too complicated for what it should offer - help on the mailing list is often difficult to get"

→ *We need a friendly user group to help complete redesign !
(During shutdown?)*



Number of Unique Visitors to AMI/TC





Dev – Partial "To Do Soon" list

- Synchronizing with DQ2 :
 - AMI client for DQ2 stomp Active MQ service has been working very well for several months.
 - We would like to extend this service to
 - Add/Remove primary datasets from dataset containers. **This is URGENT.**
 - File consistency. (not urgent since all ready have something working)
- *Borut* : 'No "automatic" way of marking datasets e.g. "September reprocessing"'. Have some ideas but don't see how it can be "automatic". Armin has a procedure to inform TAGS, and he has proposed to inform AMI at the same time.

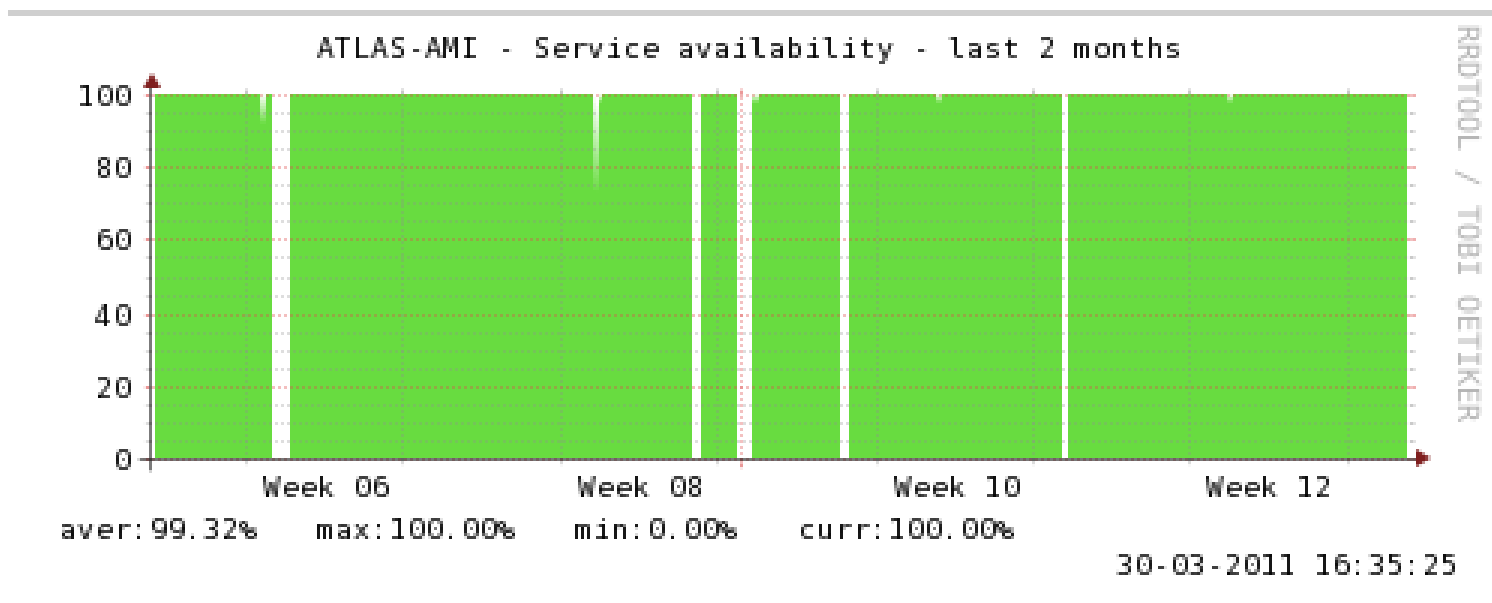


EXTRA SLIDES

- SLS + Load on AMI
- Information protection + security
- ORACLE & underscores
- COMA and Data periods



SLS for AMI




- Degradation since January.
 - We are not sure why exactly – it is not due to load. (see next two slides)
 - We suspect that the connection between the APACHE cluster and the Tomcat servers breaks.
 - The APACHE version changed in January.
 - We have treated the problem empirically (stronger watch dog) and we are planning an upgrade of Tomcat.




CCAMI02 – Number of commands per hour 10 Feb -> 28 Feb

Number of Commands per hour, 10 Feb -> 28 Feb 2011





Number of commands per hour CCAMI01 16th Feb -> 15th March





From Alex Undrus

- No nightlies are launched between 11:00 and 13:00 and between 13:30 and 20:00.
>>>>> The period between 21:00 and 23:00 is very "hot" in sense that the majority of nightly jobs are started during this period.



Security and Information Protection

- Following a security audit of the AMI web site at CERN we were asked to put the access to the AMI replica behind SSO and to clean up some rather ugly responses to error conditions or attempts to inject java script. This was done – but we had to take it away as SSO :-
 - Does not allow pyAMI through.
 - Does not protect any information from non-ATLAS members.
- The main site at Lyon remains world readable, and we cannot use SSO at Lyon.
- What we plan to do in the near future is to restrict world readable rights to the top page, and to permit only members of ATLAS VOMS to read AMI catalogues. (Waiting for management to agree)
- Everything is in place on the server side, some clients will need to adapt.



ORACLE behaviour

Which query treats "_" as a wild card?

```
ALTER SESSION SET NLS_COMP=LINGUISTIC NLS_SORT=BINARY_CI;
SELECT count(LOGICALDATASETNAME) FROM DATASET WHERE
LOGICALDATASETNAME LIKE '%data11_cos%';
```

```
ALTER SESSION SET succeeded.
COUNT(LOGICALDATASETNAME)
```

3103

```
ALTER SESSION SET NLS_COMP=LINGUISTIC NLS_SORT=BINARY_CI;
SELECT count(LOGICALDATASETNAME) FROM DATASET WHERE
LOGICALDATASETNAME LIKE '%data11\_cos%' ESCAPE '\';
```

```
ALTER SESSION SET succeeded.
COUNT(LOGICALDATASETNAME)
```

5286



COMA – complete presentation by Elizabeth Gallas

- <https://indico.cern.ch/materialDisplay.py?contribId=13&sessionId=2&materialId=slides&confId=130606>
<<https://indico.cern.ch/materialDisplay.py?contribId=13&sessionId=2&materialId=slides&confId=130606>>

Introduction: ATLAS Data Periods

- A Data Period is a set of ATLAS Runs grouped for a purpose
 - Defined by Data Preparation Coordinators
 - Used in ATLAS data processing, assessment, and selection ...
 - Each Period uniquely defined with a combination of
 - Project name (i.e. 'data10_7TeV')
 - Period name (i.e. 'C1', 'C2', 'C', 'AllYear' ...)
- Before 2011, Data Periods were
 - Described on TWiki page
 - <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/DataPeriods>
 - Stored in a file based system
 - Edited by hand by Data Prep Coordination (experts)
 - Structure evolved over last year with experience
 - This experience → valuable to decide/define long term solution
- New for 2011: Data Periods stored in the COMA DB
 - Thanks: Beate (DataPrep Coordinator), AMI team, DB experts.



Data Periods: Links to Reports and Services

The links/info below can be found on the revised TWiki page:
<https://twiki.cern.ch/twiki/bin/view/AtlasProtected/DataPeriods>

- **Interactive USERS → COMA Data Period Documentation Interface**

- https://atlas-tagservices.cern.ch/RBR/rBR_Period_Report.php

- Comments: hn-atlas-physicsMetadata@cern.ch.



Next slide

- **Programmatic USERS**

For systems needing period info: runQuery, beamspot, Data Quality, ..., “Data Period Services” provided via pyAMI:

- http://ami.in2p3.fr/opencms/opencms/AMI/www/Client/DataPeriods_pyAMI.pdf

- Comments: [AMI / Tag_Collector Team](#).

- **Data Preparation EXPERTS:**

Entry Interface:

- <https://ami.in2p3.fr/AMI/servlet/net.hep.atlas.Database.Bookkeeping.AMI.Servlet.Command?linkId=1479>

- Comments: [AMI / Tag_Collector Team](#).

COMA Period Documentation Menu

Click here for [Data Period Doc Help](#)



Period Documentation Menu

This interface generates Data Period descriptions (choose any link below). Periods are organized by Year, then by Project Name, then specific Periods (which are shown by clicking on the **./+/- highlighted text**).

- A Period is a set of Runs.
- A Period Group is a set of Periods.
- A Parent Group is a set of Period Groups.

Direct links to all Period Definitions

Domain:	DP -- Data Preparation		
by Year:	Describe all years Data Periods		
	2011		
	2010		
	2009		
by Project:	data11_7TeV	data10_hi	data09_900GeV
	data10_900GeV	data10_7TeV	data09_2TeV
by Period:	./+ data11_7TeV		
	./+ data10_hi		
	./+ data10_900GeV		
	./+ data10_7TeV		
	./+ data09_900GeV		
	./+ data09_2TeV		

Optional: textbox entry allows wildcards ('%')

Project Name (Filename_Tag)	<input type="text"/>	Example: data%_900% (all 900 GeV)
Period Name	<input type="text"/>	Example: %VDM% (all Van der Meer)

SUBMIT ! ; Clear form !

Purpose: Generate Period documentation for chosen input criteria

The report will include a description of all Periods

→ **By Year**

→ E.G. all '2010'

→ **By Project**

→ e.g. 'data10_7TeV'

→ **By specific Period or Group**

→ **Click on the project and then your Period of interest**

Wildcards can be entered in this optional section, then click on Submit button

https://atlas-tagservices.cern.ch/RBR/rBR_Period_Report.php

Example Report: All 2010 Data Period Descriptions


COMA Period Description Report








Project Name (fnt) : data10%



Input criteria: Shown in header

The Data Periods matching the input criteria (above) are described below.

- The  symbol indicates Periods which are locked (will not change Run members).
- Click on **+/- highlighted text** to expand the information in those sections.

Period	Run Links	Description	Dates	Runs	GRL
data10_7TeV AllYear  -/+A:1	COMA; runQuery	All 2010 pp 7 TeV collision runs	2010-Mar-30; 2010-Oct-29	152166-167844 (126)	DetStatus-v03-repro
data10_7TeV VdM  -/+VdM1:VdM3	COMA; runQuery	All 2010 pp 7 TeV collision runs with VdM scans	2010-Oct-05	166198 (6)	
VdM1  COMA; runQuery			2010-Apr-25; 2010-Apr-26	153599- 153599 (1)	
VdM2  COMA; runQuery			2010-May-08; 2010-May-09	154813- 154813 (1)	
VdM3  COMA; runQuery			2010-Oct-01; 2010-Oct-05	166052- 166198 (4)	
data10_900GeV AllYear  -/+B:B	COMA; runQuery	All 2010 pp 900 GeV collision runs	2010-May-02; 2010-May-03	154465-154471 (3)	
data10_hi J  -/+J1:J7	COMA; runQuery	First Heavy Ion running	2010-Nov-08; 2010-Dec-06	168665-170482 (44)	DetStatus-v03-pass1-analysis-201 DetStatu

-/+ highlighted links:
These sections expand to show period members

Members of data10_7TeV.VdM are VdM1, VdM2, VdM3

Links to COMA and runQuery multi-Run Reports for that Period