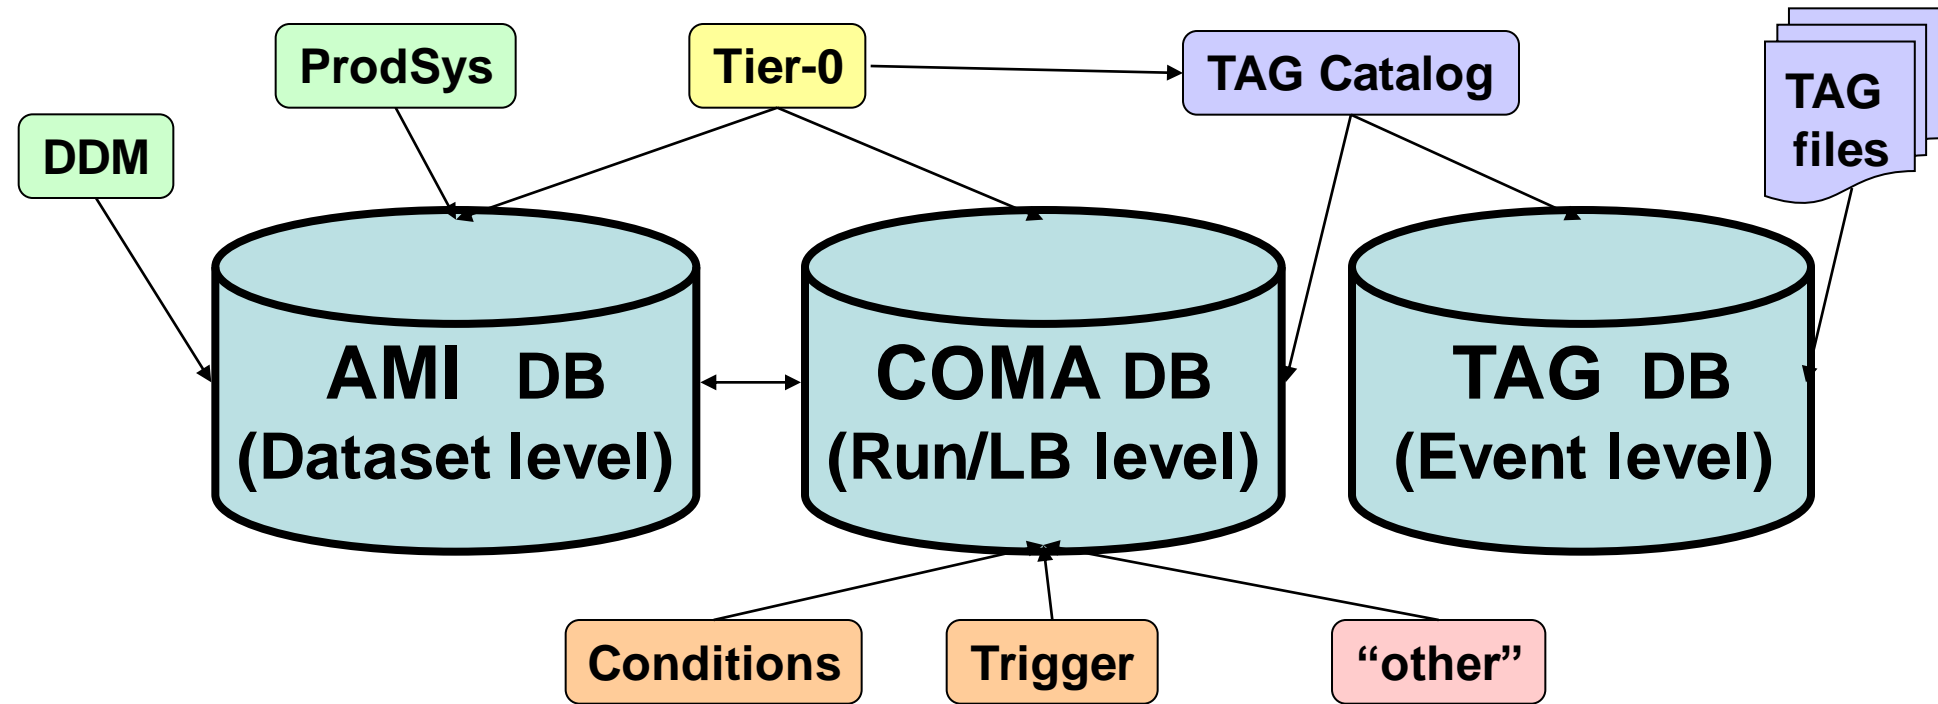


1. Introduction to COMA

- The COMA System (Conditions/Configuration Metadata for ATLAS), has been developed to make globally important run-level metadata more readily accessible. It is based on a relational database storing directly extracted, refined, reduced, and derived information from system-specific data sources as well as information from non-database sources.
- This information facilitates a variety of unique dynamic interfaces and provides information to enhance the functionality of other systems.
- COMA is one of the 3 dedicated metadata repositories in ATLAS, fitting nicely between the AMI and TAG databases, which store metadata at the dataset- and event-levels, respectively. The data sources of each of these metadata repositories is shown below:



- COMA data sources include:
 - Conditions database:** A wide variety of configuration information and measured conditions at the Run and Luminosity Block (LB) levels. An ATLAS Run is a interval of data taking (generally for many hours) with a fixed configuration, with selected configurations allowed to change at the sub-Run (or LB) level.
 - Trigger database:** Trigger specific information not readily accessible or available via the Conditions database for collision data and Monte Carlo (MC) datasets.
 - Tier-0 and TAG Catalog databases:** Information from these systems is used to identify Runs of offline “analysis” interest for upload into COMA (a subset of all ATLAS Runs). For those Runs, we collect a variety of information about the processing of those Runs. Via the TAG Catalog, COMA interfaces also have access to which processing versions are available in the TAG database.
 - AMI database:** AMI and COMA systems work symbiotically to store a variety of information about collections of Runs and their processing, making the data available to both systems.
 - Other:** Information from a variety of non-database sources: TWiki and other documentation, text and xml files, human entry.

3. COMA Interface Overview

The COMA project includes a number of interfaces for accessing, displaying, and browsing information in the COMA database.

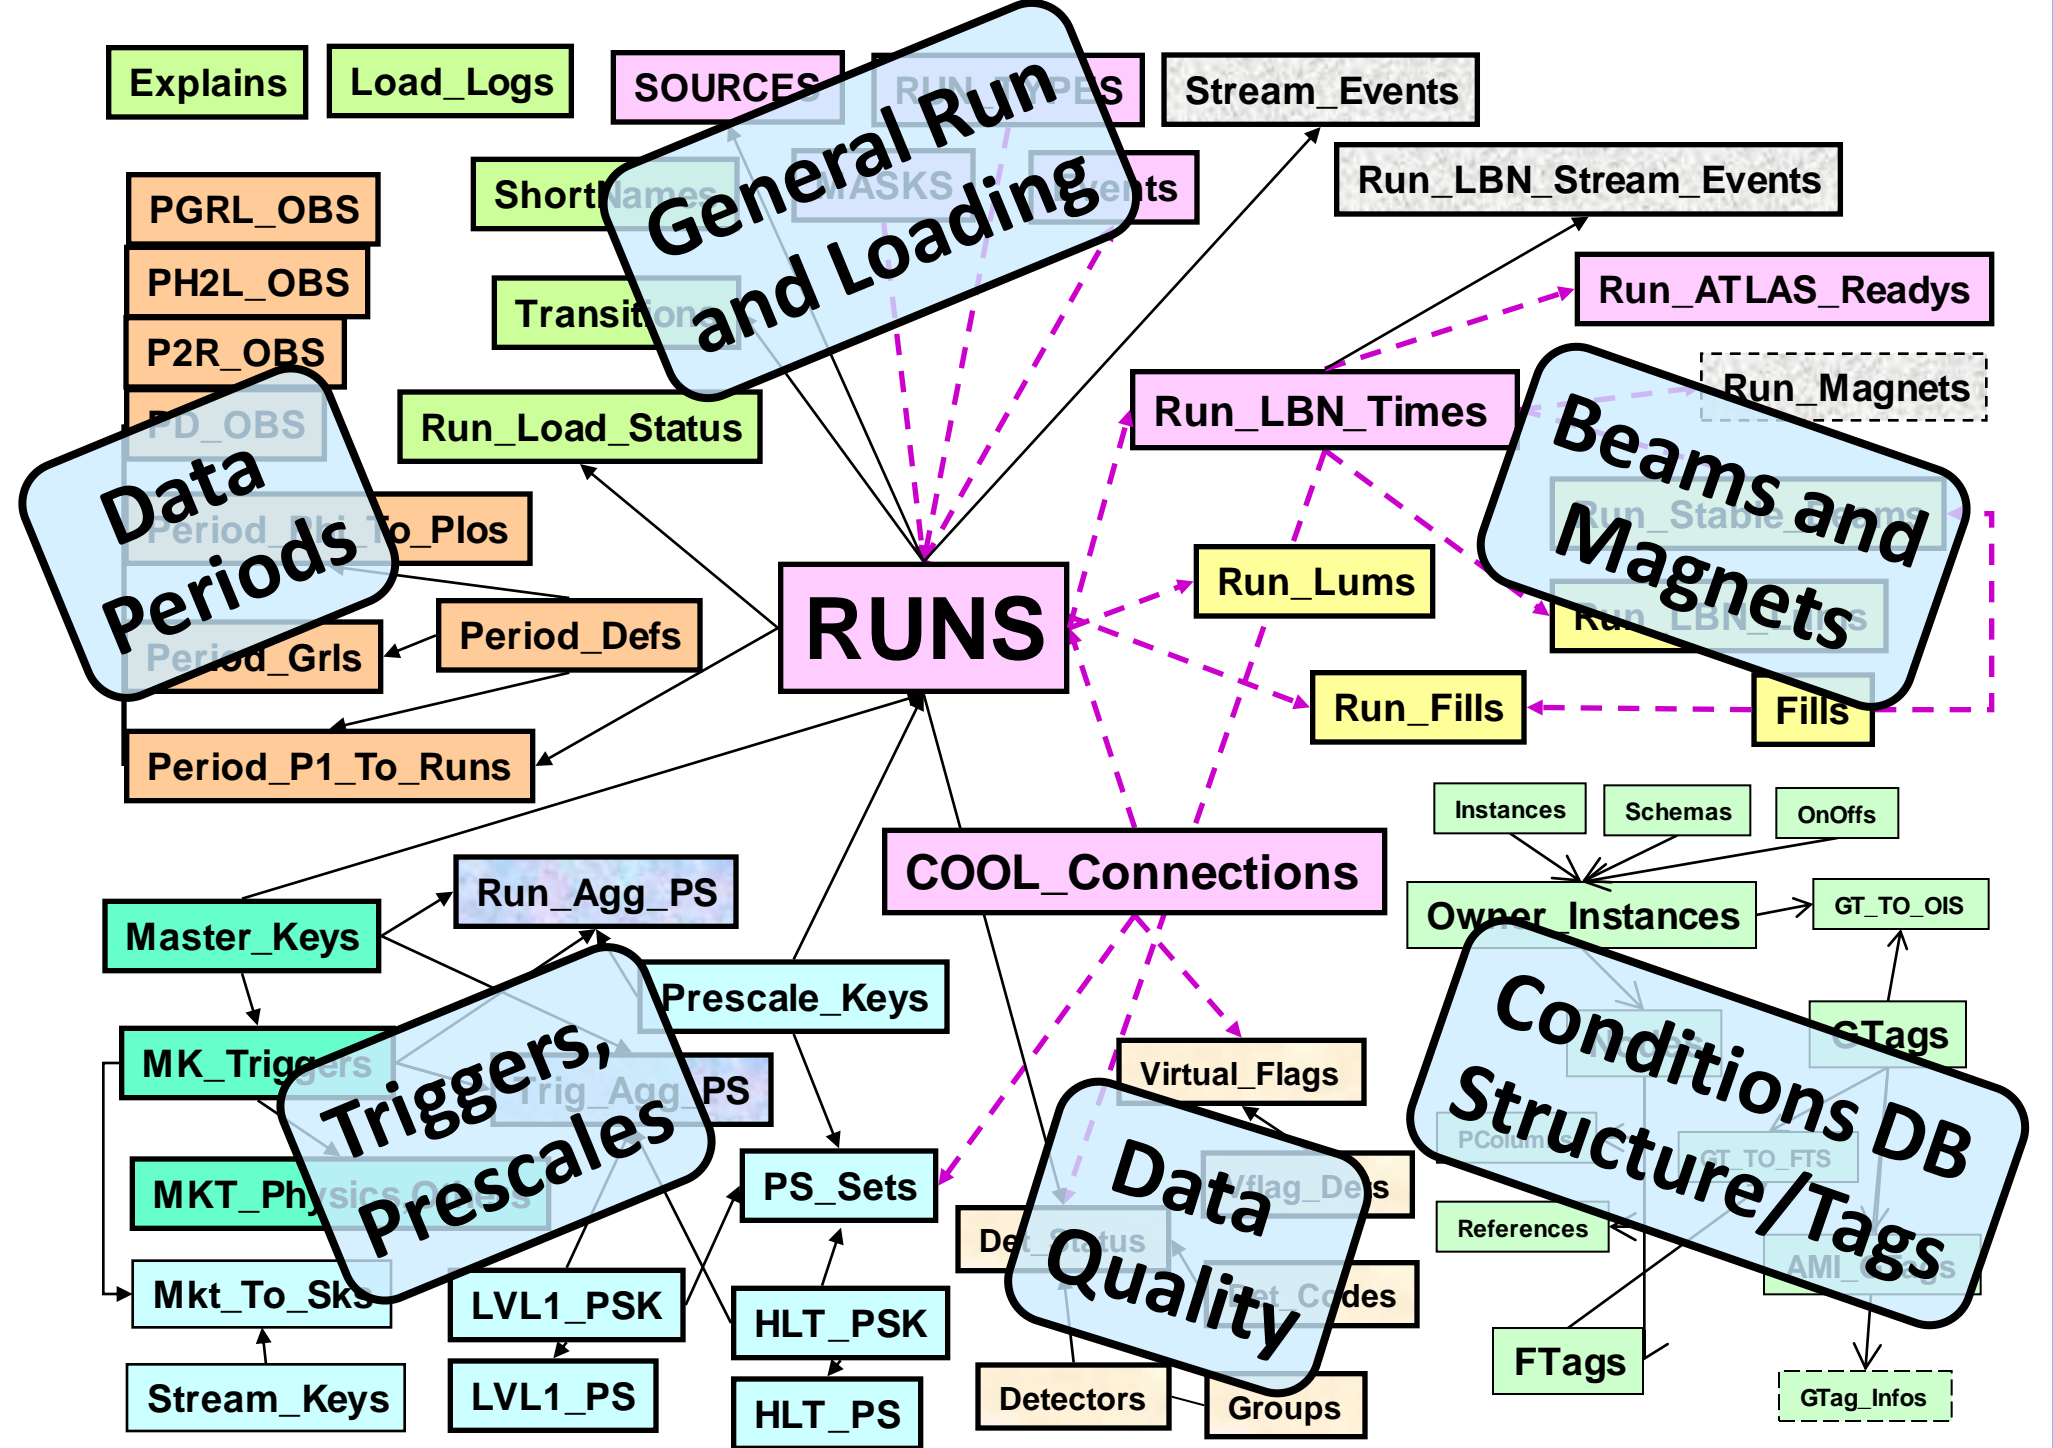
COMA Portal	
Interface (link)	Purpose
COMA Period Menu	Data Period Documentation Menu ... generates COMA Data Period Reports which describe the selected Data Period(s) and contains hyperlinks to COMA and other Run reports.
COMA Report Menu	Conditions Metadata Report Menu ... generates COMA Run, Stream, Master Key, Chain, and Prescale Reports . Displaying general properties, COMA derived properties, and related hyperlinks for the entities matching the input criteria.
COMA Browser	Browser of Conditions Metadata ... users can impose selection criteria in any order, see remaining runs and their metadata properties, iterating until the desired collection of Runs is achieved and, if needed, produce a Run-LB xml file.
COMA Timeline Menu	COMA and Trigger experts have joined forces to generate reports with information from COMA and Trigger Timelines schemas.
COMA Folder & Tag Browser	Browse Conditions database Folder and Tag metadata in COMA to find folders and display their properties. Create customized Global Tag reports.

- COMA Period Menu**
 - A dynamic menu of available Data Periods organized in a rational way (presented by Year, Project, ...). It provides links into the Data Period Reports which show metadata about the Runs contained in each Period and provide hyperlinks into the other COMA Reports and selected external reports.
- COMA Report Menu**
 - A portal page allowing a wide selection of criteria to be entered which is input to generate a wide variety of COMA reports: Single- or Multi-Run, Trigger Overview, Trigger Chain, and Trigger Prescale Reports.
- COMA Browser**
 - The browser of Conditions Metadata where users can impose selection criteria in any order, see remaining runs and their metadata properties, iterating until the desired collection of Runs is achieved and, if needed, produce a Run-LB xml file, a format recognized by downstream processing utilities. This browser was described in the CHEP 2010 conference.
- COMA Timeline Menu**
 - The result of collaboration between COMA and Trigger experts making available reports which combine information from the COMA and Trigger Timeline Schemas.
- COMA Folder / Tag Browser**
 - User is allowed to browse ATLAS Conditions database Folder and Tag metadata in COMA to find folders, display their properties and tags and create customized Global Tag reports.

5. COMA Interface Logging

- A logging system to collect details of the usage of some of the COMA web-based interfaces was recently put in place, collecting information such as:
 - General information like the report type generated, user input criteria, originator and the referring page of the request,
 - Metrics like query count, number of database rows retrieved and wall clock time passed in generating the output.
- Query logging gives a quantitative picture of usage: relaying report popularity and the variety of input criteria which is used (or not used).
- Based on the metrics in the logs, a number of optimizations have been undertaken: refining criteria, reducing the number of queries and revising the workflow of report generation. Report optimization is a multifaceted issue requiring a balancing of many aspects: from SQL complexity to php sorting and memory utilization.

2. Overview of the COMA Schema:



- The COMA database schema currently contains 68 tables and 12 views. It includes sets of tables dedicated to metadata related to data quality, trigger and prescales, data periods, loading status, general run properties, event counts, luminosity and beam related entries. In addition, COMA has a set of tables which offer an overview of the ATLAS Conditions database structure and versioning.
- The principles of COMA database design are tightly coupled to the structure of ATLAS systems and how they interact, the nature of the source data and its handling in deriving useful metadata quantities and how it is used in the client interfaces.
- Loading progresses in 8 distinct stages, each of which loads a group of logically related tables. These stages are decoupled to allow
 - Separation of exception handling and logging,
 - Flexibility in schema evolution and program development,
 - Decoupling of synchronization with their sources
- The implementation of each loading stage includes business rules which dictate how the data is managed, including what happens if some or all of related source data is absent.

4. COMA Report Overview and Example

- COMA Interfaces are based on PHP and JavaScript, using some features of the jQuery utility help users filter results.
- General Features of COMA Reports:
 - The scope of each report is a coherent and manageable collection of related information, providing hyperlinks to more detailed information from COMA or from external systems when it is possible and logical to do so.
 - Input criteria is read via GET methods (POST input also recognized) and is always shown in the report header. Input keywords are compact and mnemonic. Input criteria is flexible, allowing lists or ranges of values and is applied in a case-insensitive manner.
 - Yellow highlighted sections open to show more detailed information
 - Tabulated content is compact but offers users pop-up explanations of content and values which exploit self-documenting features built into the COMA schema.
- A COMA Single-Run Report is shown below to demonstrate some of the functionality of the interfaces and exhibit aspects of COMA content.

COMA Online Runs Report
Project Name (fnt) : data12_8TeV
Run Number (runs) : 201383

Online Run Number 201383, RunType Physics, Project data12_8TeV
Duration 31518 seconds => (8:45:18), Lumi Block Count [range] 525 [1-525]
Recording start/end: 2012-Apr-17 02:45:25 / 11:30:43
Events recorded [L1,L2,EF]: 13476100 [191383132, 97060459, 12512919], partition ATLAS.
HLT Release: 17.1.4.11.

COMA Reports: for Run 201383: [COMA \(self\)](#), [COMA_RL_201383](#), [RunLum](#), [RunLBLum](#), [Trigger SMK Report: 1375](#) (see Triggers section for COMA Chain, Prescale links).

External Links: [AMI](#), [LumiDS](#), [RunQuery\[short, All, DO, DOSum\]](#), [RL](#), [Trigger](#), [PS Evolution](#).

+ Data Periods (3):
+ LHC Fills (1):
+ LumiBlock Evolution (8):
+ Active Streams (26):
+ AMI Tags (for TAGs):
+ Triggers (688)

Yellow sections open to display detailed content

ATLAS Ready
The ATLAS Ready flag is a state declared by ATLAS that the detector is fully ready for stable colliding beams data taking and sub detectors have ramped up their voltage.

Table header links open pop up windows explaining table columns and values

LB_BegIn	LB_End	LVL1	PSKey	HLT PSKey	TriggerDB link	ATLAS Ready	Stable Beams	Fill
1	96	6423	5749	1375 / 6423 / 5749	NO	NO	2523	
97	103	6505	5786	1375 / 6505 / 5786	YES	NO	2523	
104	113	6505	5772	1375 / 6504 / 5772	NO	NO	2523	
114	218	6504	5772	1375 / 6507 / 5772	NO	NO	2523	
219	220	6507	5772	1375 / 6508 / 5772	NO	NO	2523	
221	513	6508	5749	1375 / 6423 / 5749	NO	NO	2523	
514	521	6423	5749	1375 / 6423 / 5749	NO	NO	2523	
522	525	6423	5749	1375 / 6423 / 5749	NO	NO	2523	

- The top section displays general information about this Run: when it was recorded, the duration and number of LB, event counts, etc. Hyperlinks are provided to related COMA and external reports.
- Highlighted yellow subsections show:
 - Inclusion of this Run in a Data Period,
 - LHC Fills coinciding with this run, including beam conditions and luminosity,
 - LB Evolution of this Run shows transitions in Fill number, readiness state of the ATLAS detector/TDAQ, declaration of Stable Beams, changes in the magnet states, changes in trigger prescales, and the integrated luminosity aggregated over each step in the LB evolution.
 - streams which were active during the Run and the LB ranges in which they started and ended activity,
 - datasets for this Run available in the TAG database, and
 - a large section dedicated to Trigger configuration over the course of the Run, showing which triggers were active and their range in prescales.