XXV Symposium on Nuclear Electronics and Computing - NEC'2015

The development of hybrid metadata storage for PanDA Workload Management System

M. Grigorieva, K. De, M. Golosova, A. Klimento, E. Ryabinkin
(National Research Center "Kurchatov Institute",
Brookheaven National Laboratory)

Meta-data, produced by large-scale scientific experiments

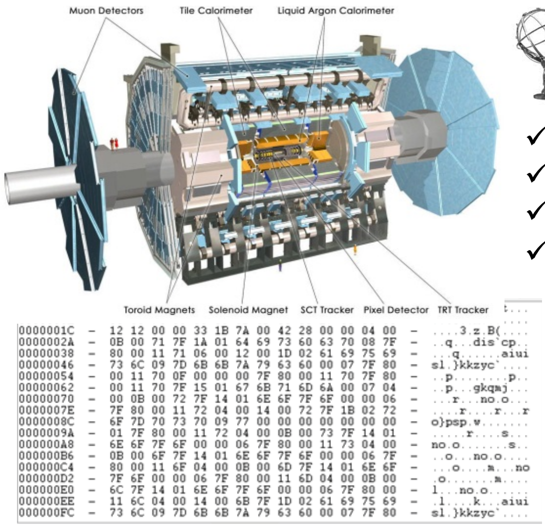

Raw data rate from LHC detector : **1PB/sec**

ATLAS EXPERIMENT

- ✓ Exabyte scale
- ✓ 3000 scientists
- ✓ 180 institutions
- ✓ 40 countries

- More than 30M jobs /month
- At more than 150 sites
- Running at 250 000 cores worldwide

- ✓ The ATLAS experiment uses **PanDA (The Production and Distributed Analysis) Workload Management System** for managing the workflow for all data processing on the WLCG.
- ✓ PanDA makes hundreds of distributed sites appeared as local

Processing and analysis of huge amounts of auxiliary meta-data, surrounding the life cycle of scientific experiments, is no less challenging task than the management and processing of experimental data files and results.

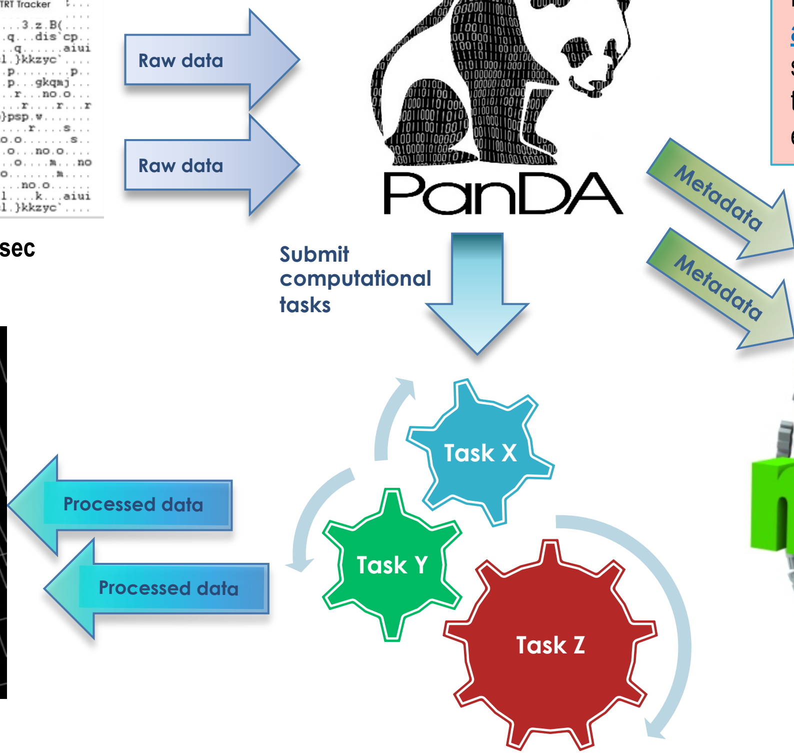

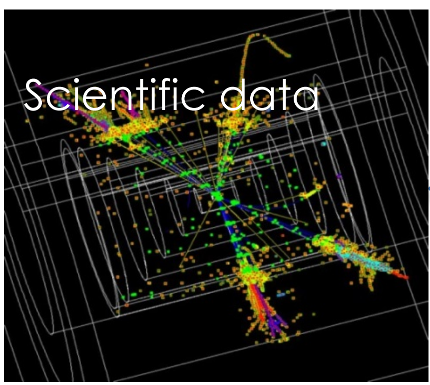

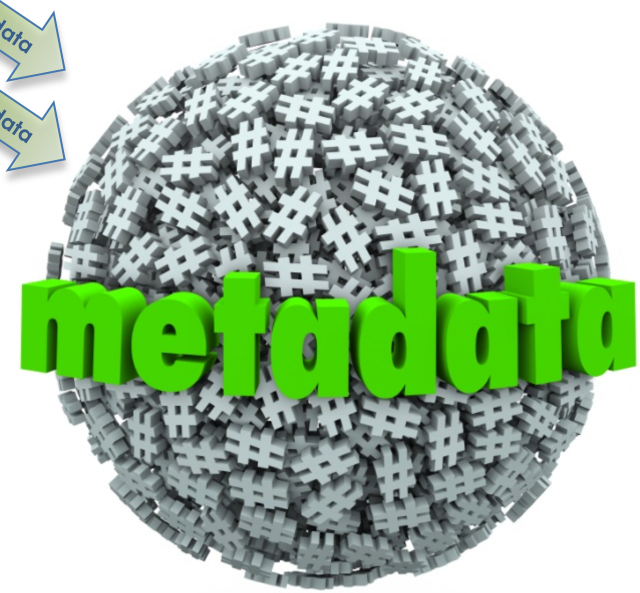

PanDA at Glance

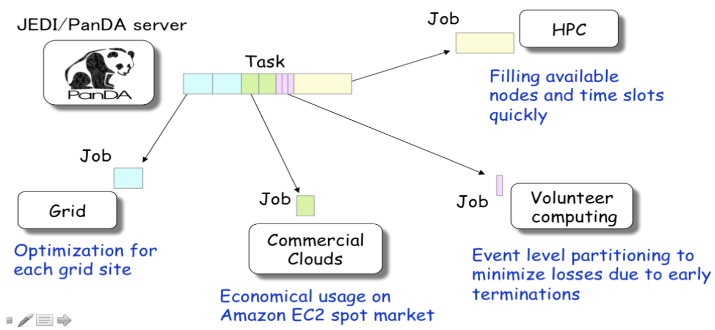

- ✓ Pilot-based WMS (pilots – “place holders” for payload)
- ✓ **PanDA Server** is the main component which provides a task queue managing all job information centrally
- ✓ Jobs are submitted to the PANDA server via a simple Python/ HTTP client interface
- ✓ Pilots retrieve jobs from the PANDA server to run the jobs as soon as CPU’s becomes available

ORACLE is used as the storage back-end, keeping all meta-information about computational tasks, jobs, datasets

BigPanDAMon

http://bigpanda.cern.ch/

Visualizing of the state of current and historical jobs, tasks, datasets, and performs run-time and retrospective analysis of failures on all used computing resources.

Job’s meta-data

- ✓ statistical analysis of recent workflows,
- ✓ detection of faulty resources,
- ✓ prediction of future usage patterns
- ✓ PanDA full archive now hosts information of over billion of records – all the jobs since the system started in 2006.

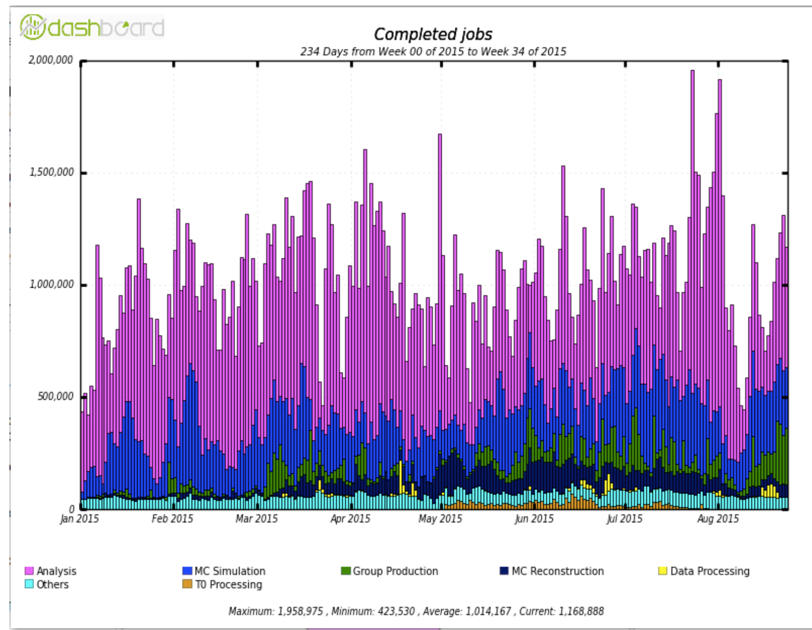

The number of finished jobs per day - currently it's up to 2 million jobs per day

PanDA meta-data storage technicalities

- ✓ *Actual DB clients*
  - ✓ *Real-time monitoring*
  - ✓ *queued/executing/paused jobs*

- ✓ *Archive DB clients*
  - ✓ *Analytical tasks*
  - ✓ *Finished jobs*

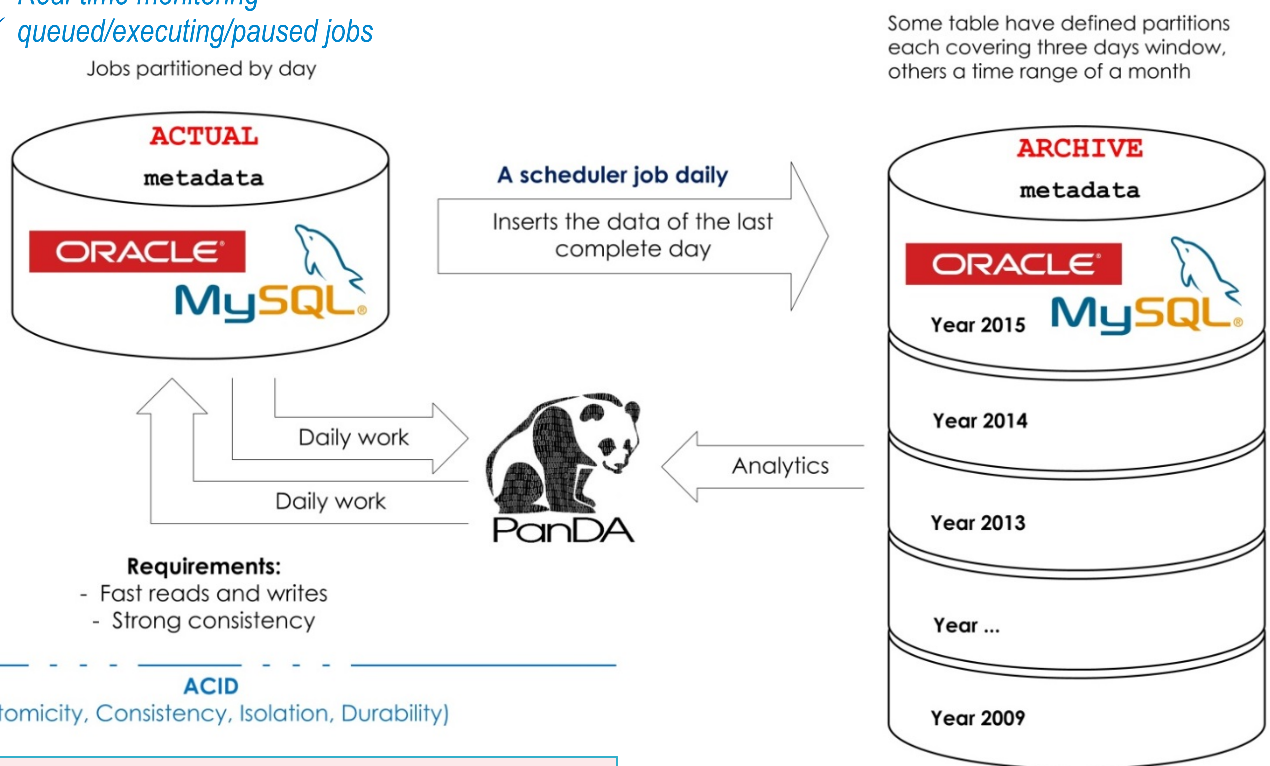

Requirements:
- Fast reads and writes
- Strong consistency

ACID
(Atomicity, Consistency, Isolation, Durability)

Requirements:
- Read-only access
- Bulk reads (long-term analysis)
- Fast reads (short-term analysis)

BASE
(Basic Availability, Soft-state, Eventual consistency)

Challenge

As the archived data volume grows, the underlying software and hardware stack encounters certain limits that negatively affect processing speed and the possibilities of meta-data analysis.

One of main paths:
Improve the scalability and performance of the **PanDA** monitor using "noSQL" solution for finalized, reference portion of the data where consistency is not essential, but sheer performance and scalability are.

PanDA Jobs Status/Error monitoring

- ✓ One of the primary goals of PanDA monitoring is the spotting **job failures.**
- ✓ Error messages play a key role in determining the potential problems that might arise.

BigPanDAMon http://bigpanda.cern.ch/errors/

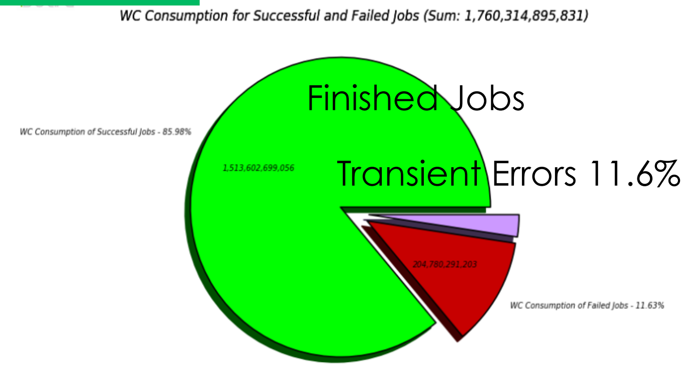

1. Errors Summary filtered by parameter

Number of Errors · Error Name : Code · Error Dialog · Job Statuses · Grid Sites

Site error summary		
AGLT2_MCORE	8	Total errors in 8 jobs. Finished: 108 Failed: 13 % failed: 10 Holding: 3 Cancelled: 0
	7	exe:65 Non-zero return code from EVNTtoHITS (65); Core dump at line 10313 (see jobReport for further details)
	1	exe:68 Fatal error in athena logfile: "G4 exception at line 17297 (see jobReport for further details)"
AGLT2_SL6	7	Total errors in 7 jobs. Finished: 1708 Failed: 11 % failed: 0 Holding: 70 Cancelled: 30
	5	ddm:100 Setupper._subscribeDistpatchDB() could not register location for prestage
	2	exe:69999 TRF_UNKNOWN | "!!! MADGRAPH TERMINATES NORMALLY: NO MORE EVENTS IN FILE !!!" | "../src. code 0: this-> |
ANALY_AGLT2_SL6	23	Total errors in 23 jobs. Finished: Failed: % failed: Holding: Cancelled:

2. Errors Count Timeline Histogram

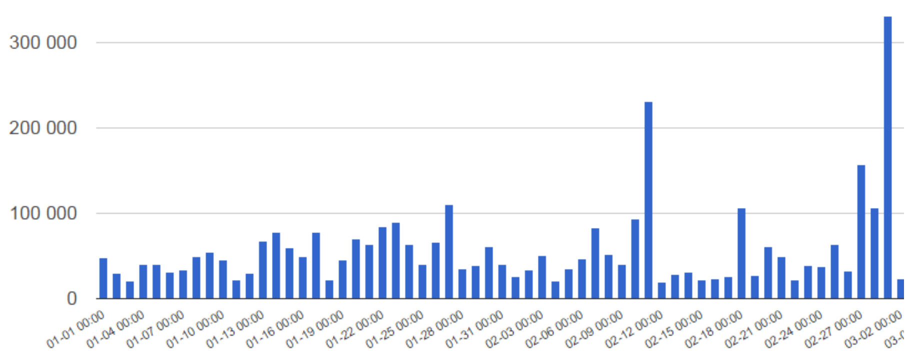

3. Overall error summary

Overall error summary

Category:code	Job list	Nerrors	Sample error description
transformation:11	jobs	7276	Error when handling transform output file
transformation:2	jobs	412	Payload core dump
transformation:1	jobs	395	Unspecified error, consult log file
jobdispatcher:100	jobs	301	lost heartbeat : 2015-09-14 06:54:23
taskbuffer:101	jobs	220	transfer timeout for HITS.06424728._096991.pool.root.1 log.0
pilot:1212	jobs	211	Payload ran out of memory
exe:65	jobs	191	Non-zero return code from EVNTtoHITS (65); Core dump at lir
pilot:1137	jobs	191	Put error: ^!^
transformation:6	jobs	184	TRF_SEGVIO - Segmentation violation
pilot:1201	jobs	164	Job killed by signal 15: Signal handler has set job result to FAI
taskbuffer:119	jobs	160	all event ranges failed
pilot:1213	jobs	119	Reached maximum batch system time limit
transformation:40	jobs	117	Athena crash - consult log file
transformation:220	jobs	106	Proot: An exception occurred in the user analysis code
pilot:1099	jobs	102	Get error: Staging input file failed

PanDA Monitor Workflow and Statistics

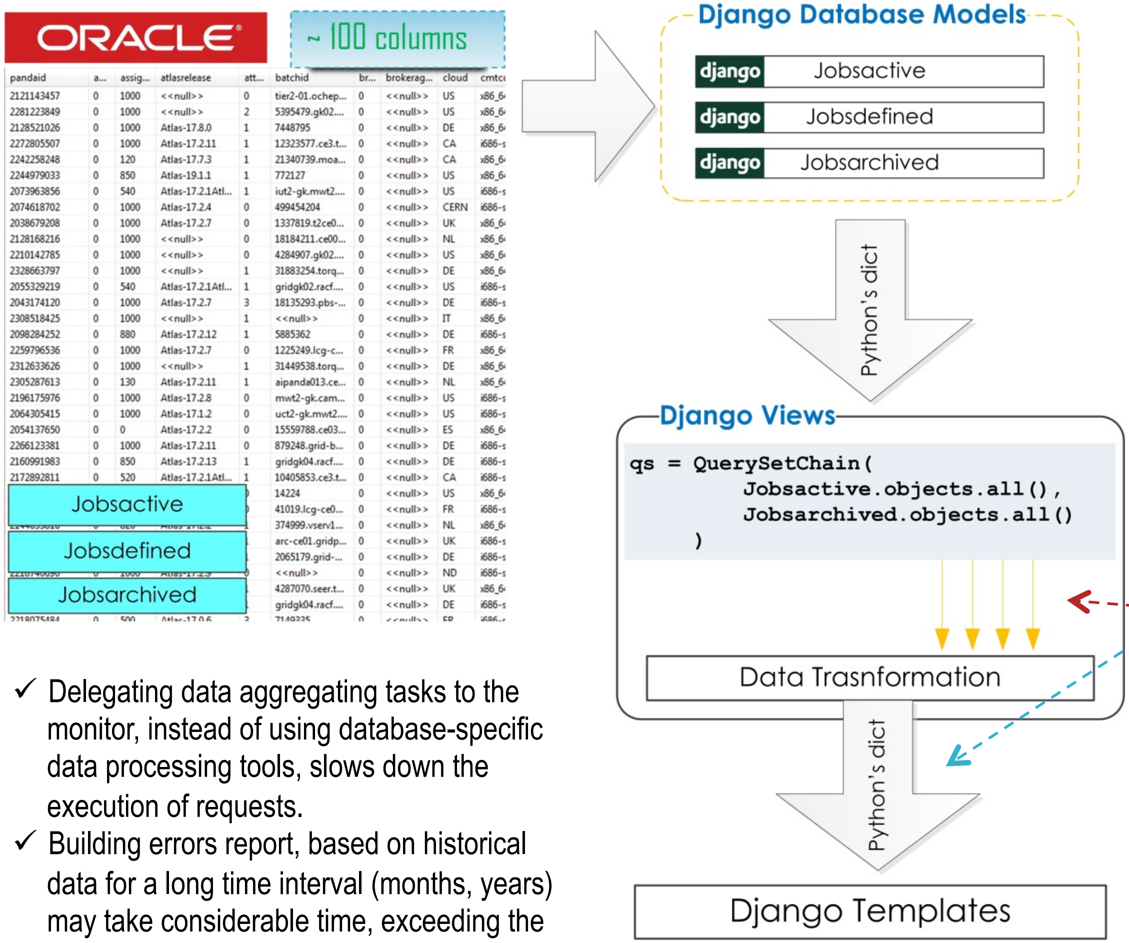

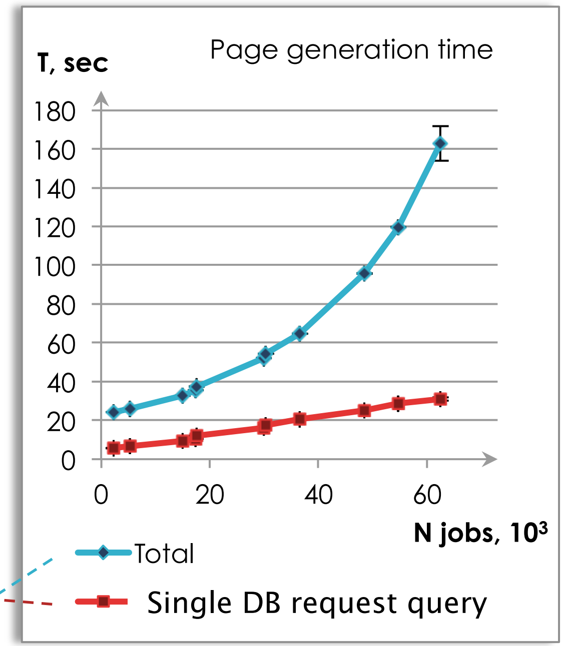

Total page generation time, including database request and aggregating obtained meta-data, dramatically decreases with the growth of the number of processing jobs

- ✓ Delegating data aggregating tasks to the monitor, instead of using database-specific data processing tools, slows down the execution of requests.
- ✓ Building errors report, based on historical data for a long time interval (months, years) may take considerable time, exceeding the reasonable time of web page generation.

Since we have already collected meta-data for more than 10 years and the amount of accumulated meta-data is constantly increasing, there is a need for long-term failures forecasting and analysis of system behavior under various conditions.

Hybrid Meta-data Storage Framework (HMSF) architecture

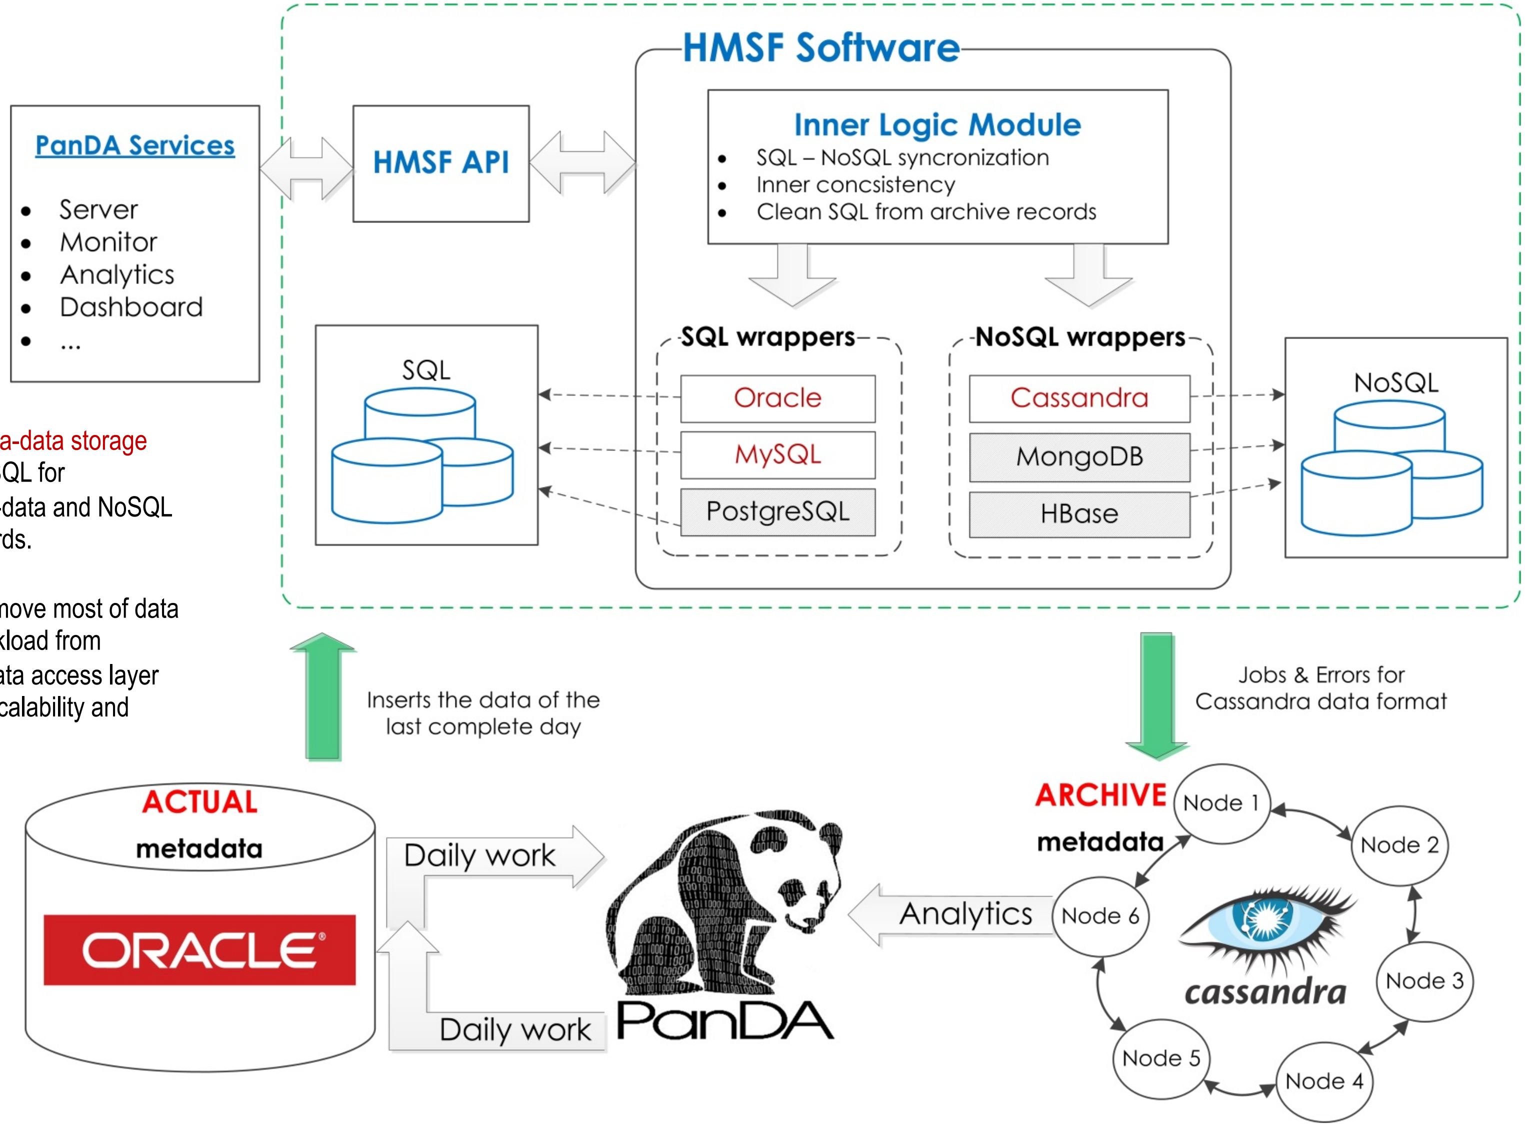

- ✓ Technique of meta-data storage segmentation – SQL for operational meta-data and NoSQL for archived records.
- ✓ This allows to remove most of data preparations workload from BigPanDAmon data access layer and improve its scalability and performance.

Cassandra as NoSQL database back-end

- ✓ A common data modelling strategy for NoSQL database systems is to store data in query-specific tables.
- ✓ No support of foreign key relationships, no JOINs of multiple CFs to satisfy a query.
- ✓ Cassandra performs best if all the data required for a given query is located in the same column family **(CF)**.

- ✓ Denormalize the data model so that a query can be served from the results from one row and query.
- ✓ All required data can be available in just one read which prevents multiple lookups.

Compound Primary Key

[**Partition Key** defines the physical location of the data in the cluster

+

Clustering Keys decides how data will be ordered/grouped in a partition]

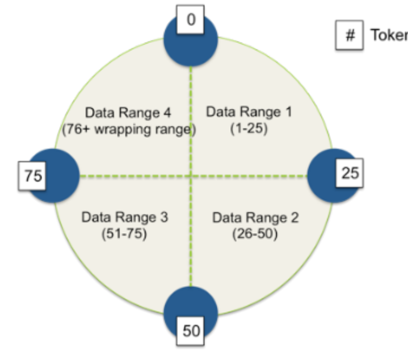

Node	Start range	End range	Primary key	Hash value
A	-9223372036854775	-46116860184273	3372	-672337285403678
B	-4611686018427387	-14894768874875	4976	-224546267672322
C	…	…	…	…
D	…	…	…	…

- ✓ The main Cassandra column family (table) – contain all job's meta-data (migrated from Oracle table "jobsarchived")
- ✓ The row of data is sent to nodes by the value of the JobID hash value
- ✓ The main table **contains the most information**, but it is **hard** for an application to work with **without any preprocessing**.

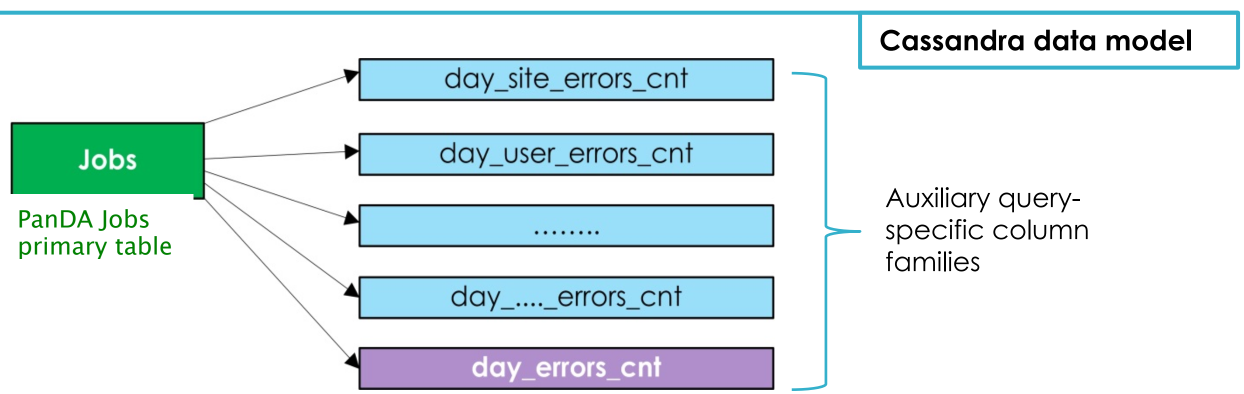

To improve BigPanDAmon performance, jobs meta-data aggregation logic was added to the HMSF.

Cassandra **time granularity** auxiliary table

GRADUALLY TIME-SERIES AGGREGATION

to maintain data at different levels of granularity ranging from fine-grained to coarse-grained data, where each level can be used for analysis and reporting purposes on different detalisation.

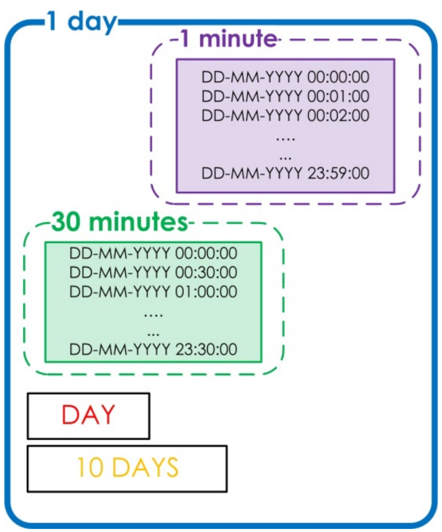

To get errors meta-data for some time period we need to make a set of requests to time granularity table. The number of requests is defined by time intervals.

- Errors metadata partitioned by days in Cassandra data model.
- For each day we have defined time intervals.
- To reduce the size of Cassandra partitions we use composite partition key.
- *Job errors meta-data are spread evenly across Cassandra cluster nodes according to the combinations of date and interval values.*

date	interval	Base_ mtime	param	errcode	Err_count	Job_count
2014-01-01	10 days [2014-01-01, 2014-01-11]	- - -		Pilot:1144 Transformation:2 Jobdispatcher:100 …	16773 988 736 …	16773 900 736 …
2014-01-01	1 day [2014-01-01, 2014-01-02]	- - -	AGLT2_SL6	Transformation:2 Jobdispatcher:100	5 3	3 1
2014-01-01	30 minutes	00:00:00	AGLT2_SL6	Transformation:2 Jobdispatcher:100	1 1	1 1
		00:30:00	AGLT2_SL6	Pilot:4476 Transformation:99	2 1	1 1
		…	AGLT2_SL6	…	…	…
		23:30:00	AGLT2_SL6	Pilot:100	1	1
2014-01-01	1 minute	00:00:00	AGLT2_SL6	Pilot:98 Transformation:98	1 1	1 1
		…	…	…	…	…
		00:01:00	AGLT2_SL6	Jobdispatcher:100	3	1
		23:59:00	AGLT2_SL6	Pilot:200	1	1

HMSF Performance Studies

Errors count timeline histogram for different internal intervals

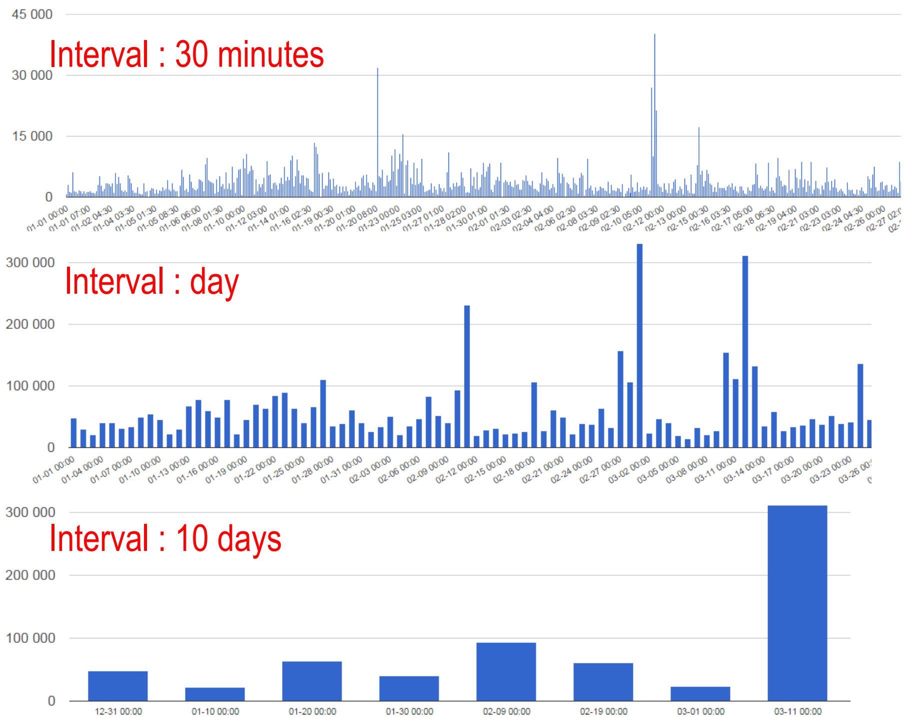

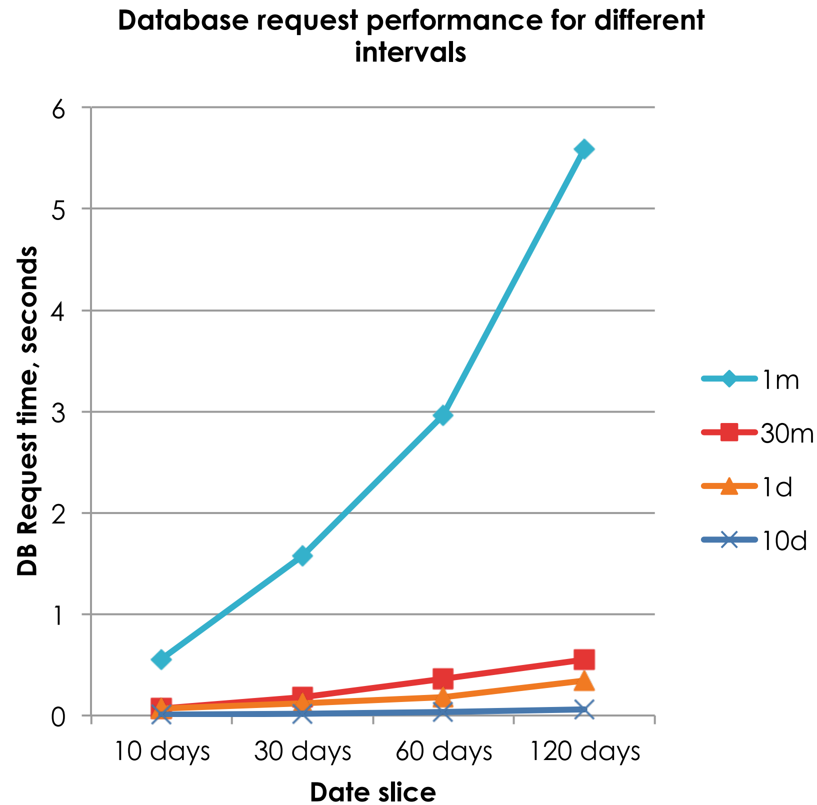

Number of rows, returned by database query

ROWS	1m	30m	1d	10d
10 days	14 306	432	10	1
30 days	42 906	1 392	30	3
60 days	85 819	2 832	60	6
120 days	170 074	5 661	120	12

Scalability test

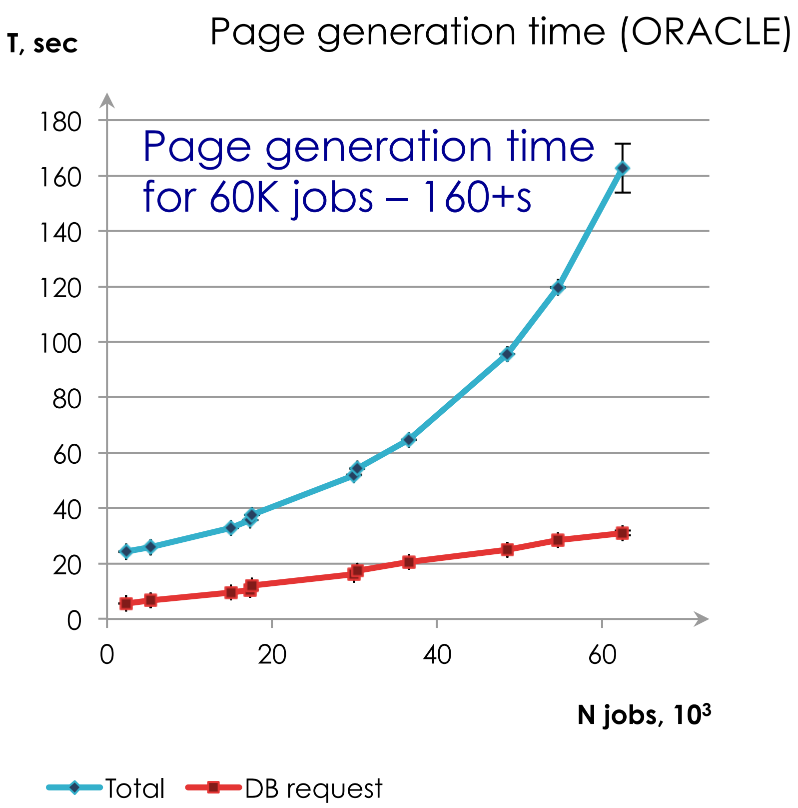

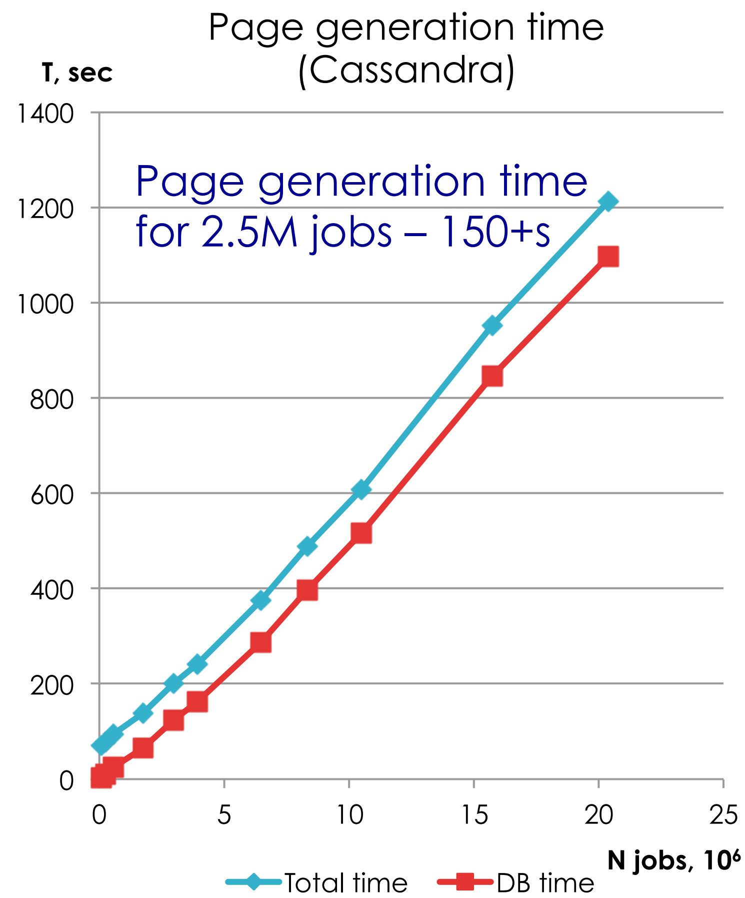

PanDA monitor adaptation to interact with HMSF

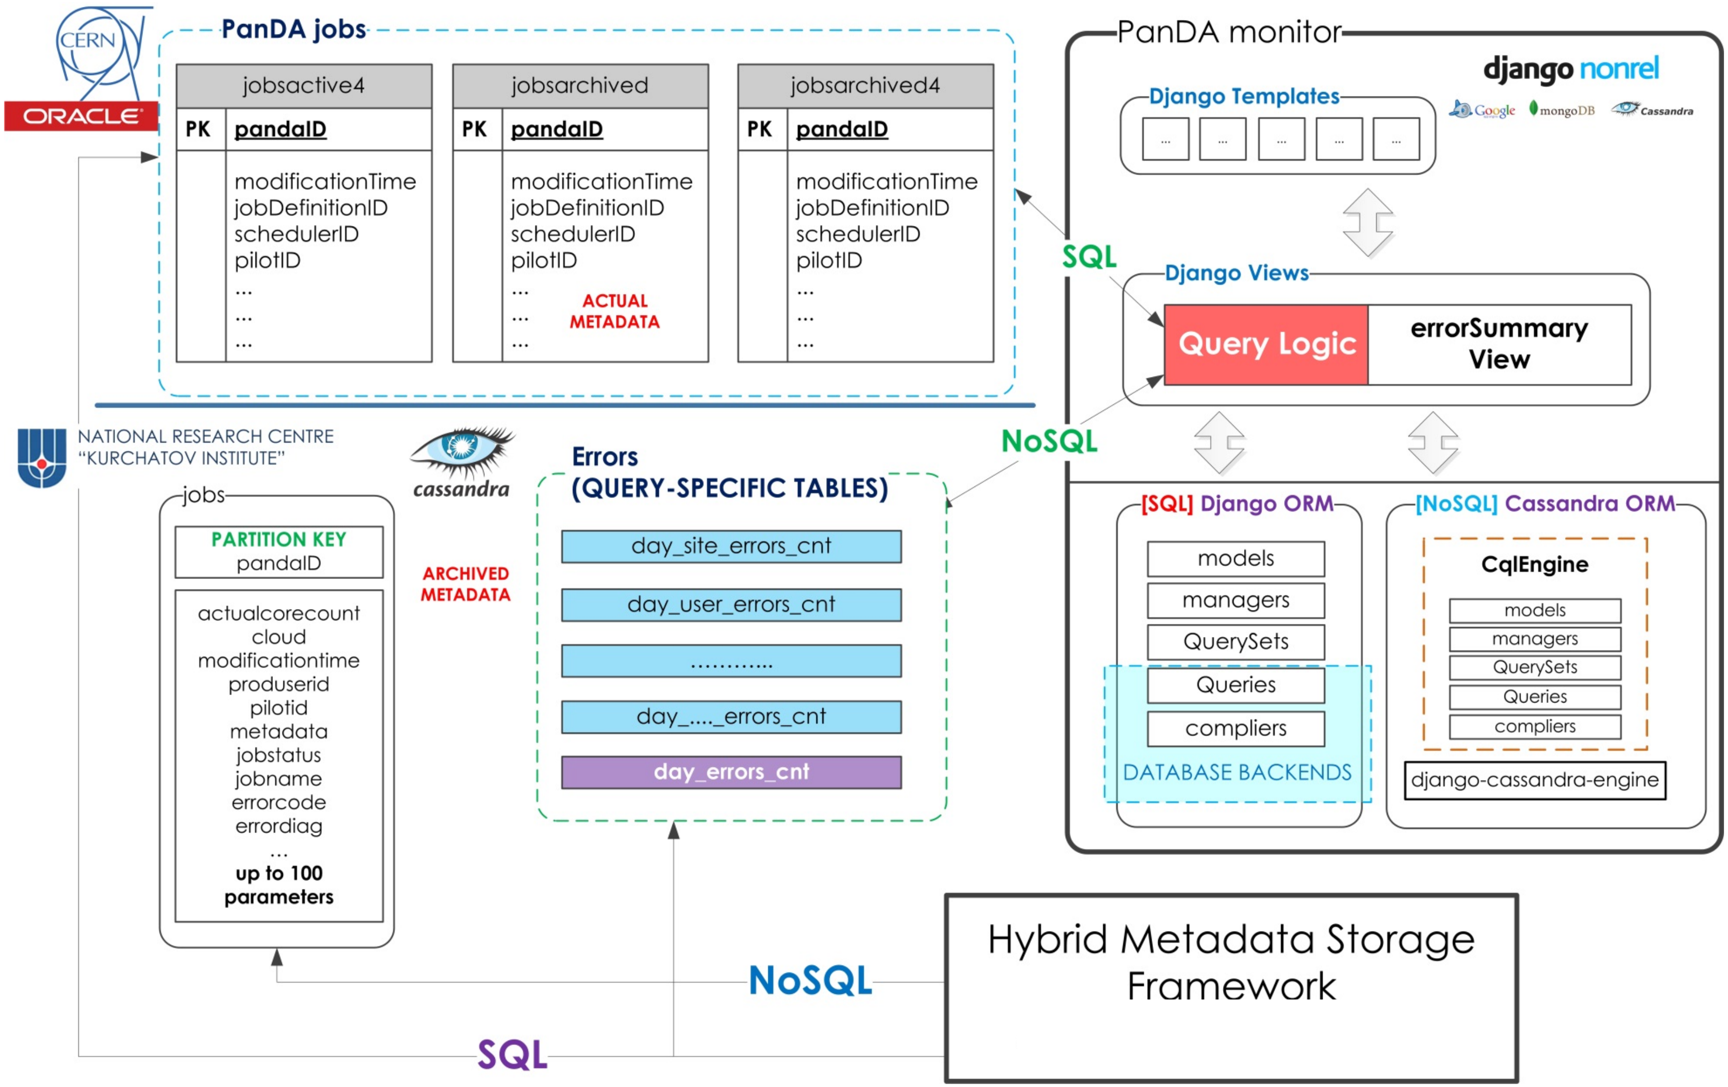

Summary

- It is hardly possible to perform long-term metadata analysis without any precalculation.
- Built NoSQL archive of metadata to improve availability of historical data.
- Prototype of Cassandra archive was created and tested on a 1-year slice of metadata from ATLAS PanDA Archive.
- Developed specific data structure for Cassandra: time granularity table.
- Near-term plans to conduct performance tests of time granularity table with Oracle.
- Adaptation of PanDA Monitor for work with NoSQL archive will be continued.

Acknowledgment

- This talk drew on presentations, discussions, comments, input from many our CERN, ATLAS, NRC-KI Colleagues, thanks to all people working in PanDA project

- This work was funded in part by the Russian Ministry of Science and Education under Contract N14.Z50.31.0024.