

Relational database evolution in ATLAS



**Elizabeth Gallas (Oxford),
Gancho Dimitrov (CERN),
For ATLAS DB Coordination**

**Database Futures Workshop
CERN**

May 29-30, 2017



- Relational databases:
 - Critical backend storage for many systems in ATLAS
 - Online, offline, and on the grid
 - For processing past and current data and supporting daily operations
- Relational storage is well suited for many of these systems
 - Relational technologies have evolved over time:
 - provisioned for established use cases
 - In functionality and in supporting increasing data volumes
 - While 'not trendy' now, relational storage is best for many applications:
 - Future application have chosen it: example MSW (muon small wheel)
- In this talk, a very brief overview:
 - Current schemas supporting applications and their volumes
 - It is logical to assume that their support must be continued
 - Note: Some smaller schemas are not listed
 - There is some investigation into, for example, 'Time series databases'
 - Which may be more appropriate

ADCR (ATLAS Distributed Computing) – 11 schemas

GROUP_NAME	Size_MB 2014	Size_MB 2017
LOCALGROUPDISK_MGT		149
PRODSYS	11259	9080
AGIS	7434	14696
GRISLI	19694	17254
PANDABROMON		20879
DEFT	22	41035
PANDABIGMON		69060
PANDAMETA	263859	626699
PANDA	415655	3203320
RUCIO	294382	3359724
PANDAARCH	6942185	9558537

ATLAS Relational DBs: ATONR

ATONR (ATLAS Online)

GROUP_NAME	Size_MB 2014	Size_MB 2017
ATLOG_CONF	1	1
RUN_NUMBER	37	43
DSS	73	98
GEOMETRY	197	174
DAQ_SHIFTER_ASSISTANT		190
COCA		386
MUBR_LVL1	486	392
DML_RATE	506	1130
ATLOG	4911	5835
TRIGGER	1453	6488
OKS_ARCHIVE	82567	76609
SFO	61891	104752
CONFIG	192713	160858
MDA		162609
LOG_MESSG		292210
MDT_DCS	442964	542550
COOL	1470184	1870027
PVSS	6029879	9582464

18 rows selected.

ATLAS Relational DBs: ATR (ATLAS Offline – not ADC)

GROUP_NAME	Size_MB 2014	Size_MB 2017
TODO_IDE	2	2
ATLOG_OF1	2	2
FRONTIER_TRACKMOD	3	3
ELISA_LOG		7
TILECAL	12	10
BEAMSPOT	15	19
RUN_NUMBER	35	41
GLANCE	49	48
DQ_RESULTS	22	52
COND_TOOLS	139	141
GEOMETRY	274	235
AUTHDB	170	341
OXS_TDAQ	3154	2544
MUON	2445	3751
ATLOG	4836	5756
TRIGGER	2314	8379
TILECOM	8792	9059
TAGS	9566	11998

GROUP_NAME	Size_MB 2014	Size_MB 2017
DBMON		18380
LTM_RTT	29777	24267
TO	41388	33377
NICOS	17392	44167
OXS_ARCHIVE	79610	74301
SFO	57639	100416
CONFIG	169207	141774
MUONCALIB	180575	219708
TRT	297618	267576
AMI	306910	450151
TZ		469170
MDT	982303	982100
EVENTINDEX		2406892
COOL	2052586	3237067
PVSS	15626505	15351423

Largest schema PVSS:
Breakdown by subsystem
In the next slide

ATLAS Relational DBs: ATRR PVSS by subsystem

Online and offline

Sub-detector	Billion rows
PVSSLAR	27.18
PVSSDCS	15.81
PVSSRPC	13.5
PVSSIDE	9.99
PVSSPIX	9.32
PVSSSTIL	7.6
PVSSMDT	7.23
PVSSSCT	6.6
PVSSSTRT	5.99
PVSSSTDQ	4.5
PVSSIS	2.22
PVSSSTGC	1.97
PVSSRPO	1.89
PVSSLUC	1.85
PVSSMUO	0.9
PVSSCSC	0.57
PVSSZDC	0.21
PVSSAFP	0.06
PVSSDSS	0.06
PVSSMMG	0

Online

Offline

Sub-detector	Billion rows
PVSSLAR	21.87
PVSSRPC	14.72
PVSSDCS	13.94
PVSSIDE	12.29
PVSSPIX	9.37
PVSSSTIL	9.33
PVSSSTRT	8.34
PVSSMDT	6.65
PVSSSTDQ	6.03
PVSSSCT	5.98
PVSSSTGC	4.25
PVSSIDESR	1.9
PVSSIS	1.34
PVSSLUC	1.33
PVSSPSR1	0.75
PVSSRPO	0.75
PVSSMUO	0.61
PVSSCSC	0.44
PVSSSTILPN	0.29
PVSSSTRTSR	0.23
PVSSZDC	0.17
PVSSDSS	0.04
PVSSSCTSR	0.02
PVSSMMG	0
PVSSAFPSPR	0

- Systems based on relational database storage
 - Are critical to current operations, processing
 - Store legacy data from the past
 - Many of these systems:
 - Refined over time based on this storage technology
- We assume we need to continue to support them
 - Projected data volumes for the future
 - I did not get time to estimate before this talk
 - Sometimes can be directly projected from previous growth
 - But may be application plan dependent
- Thanks to DBAs from ATLAS and CERN IT !