




Distr. DB Operations workshop - November 2008



the **ATLAS Experiment**



The PVSS Oracle DB Archive in ATLAS (life cycle of the data)

Gancho Dimitrov (LBNL)



Outline



- Introduction to the PVSS system
- Nature of the data and organization
- The ATLAS PVSS schemas and EVENTHISTORY table description
- The need of having PVSS data replication from ATONR to ATLR
- Some optimizations
- Data volumes
- Conclusions



Introduction to the PVSS system and its use in ATLAS




PVSS (Prozessvisualisierungs- und Steuerungssystem) is a Control and Data Acquisition system, developed by the Austrian company ETM (now 100% owned by Siemens AG).

In year 2000, it has been chosen as a control system for the LHC experiments





Thousands of data point elements



The ATLAS detector

PVSS Oracle archive - keeps history of the detector status, e.g. high voltages, temperatures



The ATLAS 'online' Oracle DB



Nature of the data and organization



- The data representation in the PVSS Oracle archive is simple:

- data point element ID

- time stamp

- value

VALUE_NUMBER (BINARY_DOUBLE)

or

VALUE_STRING (VARHAR2(4000))

- plus few more additional columns

- The primary key is based on the ELEMENT ID and the TIMESTAMP



The ATLAS PVSS DB accounts and table desc.



- A database schema per subdetector (as total 14)

▶	ATLAS_PVSSCSC	▶	EVENTHISTORY_00000002
▶	ATLAS_PVSSCSC_W	▶	EVENTHISTORY_00000003
▶	ATLAS_PVSSDCS	▶	EVENTHISTORY_00000004
▶	ATLAS_PVSSDCS_W	▶	EVENTHISTORY_00000005
▶	ATLAS_PVSSDSS	▶	EVENTHISTORY_00000006
▶	ATLAS_PVSSDSS_W	▶	EVENTHISTORY_00000007
▶	ATLAS_PVSSIDE	▶	EVENTHISTORY_00000008
▶	ATLAS_PVSSIDE_W	▶	EVENTHISTORY_00000009
▶	ATLAS_PVSSLAR	▶	EVENTHISTORY_00000010
▶	ATLAS_PVSSLAR_W	▶	EVENTHISTORY_00000011
▶	ATLAS_PVSSLUC	▶	EVENTHISTORY_00000012
▶	ATLAS_PVSSLUC_W		ELEMENT_ID
▶	ATLAS_PVSSMDT		TS
▶	ATLAS_PVSSMDT_W		VALUE_NUMBER
▶	ATLAS_PVSSPIX		STATUS
▶	ATLAS_PVSSPIX_W		MANAGER
▶	ATLAS_PVSSRPC		TYPE_
▶	ATLAS_PVSSRPC_W		USER_
▶	ATLAS_PVSSSCT		SYS_ID
▶	ATLAS_PVSSSCT_W		BASE
▶	ATLAS_PVSSTDQ		TEXT
▶	ATLAS_PVSSTDQ_W		VALUE_STRING
▶	ATLAS_PVSSTGC		VALUE_TIMESTAMP
▶	ATLAS_PVSSTGC_W		CORRVALUE_STRING
▶	ATLAS_PVSSSTIL		CORRVALUE_NUMBER
▶	ATLAS_PVSSSTIL_W		CORRVALUE_TIMESTAMP
▶	ATLAS_PVSSSTRT		OLVALUE_STRING
▶	ATLAS_PVSSSTRT_W		OLVALUE_NUMBER
			OLVALUE_TIMESTAMP

Table is 'switched' when it reaches a certain size and a view is updated to keep them together for the application to access the data (the EVENTHISTORY view)

Data point elements, in the LAR case are about 4500

Not used from ATLAS, get NULL values, thus do not take occupy space

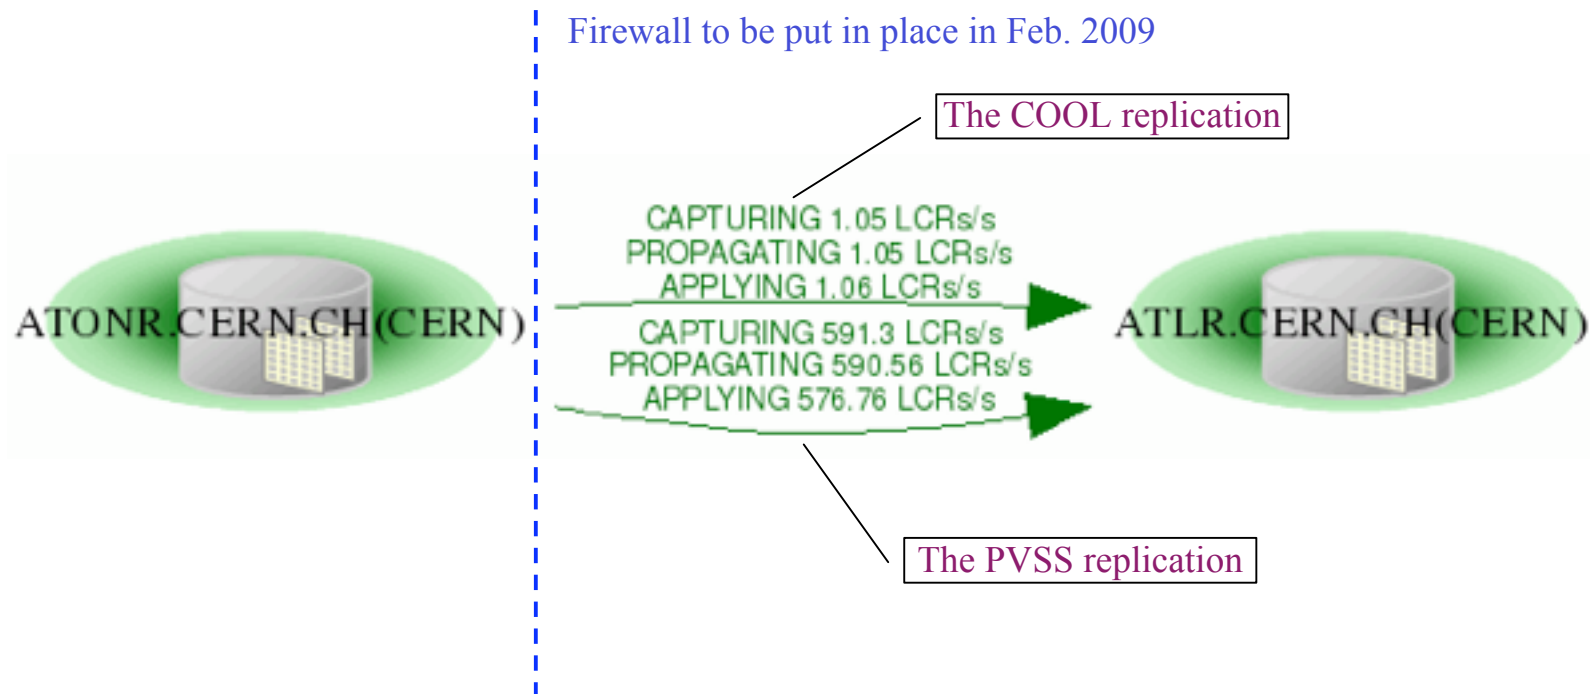
The row length is in the range 55-60 bytes



The need of having PVSS data replication from ATONR to ATRR ('online' => 'offline')



- In order to have the PVSS data accessible for the sub-detector expert analysis from the CERN public network and even from outside CERN a need for its replication showed up.
- After series of testing a production replication of the PVSS data was introduced in the beginning of July 2008.

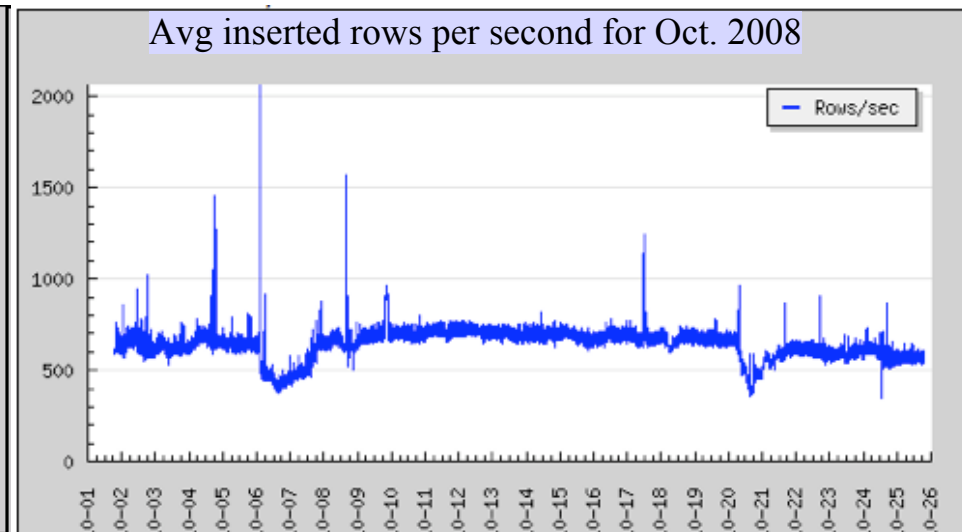
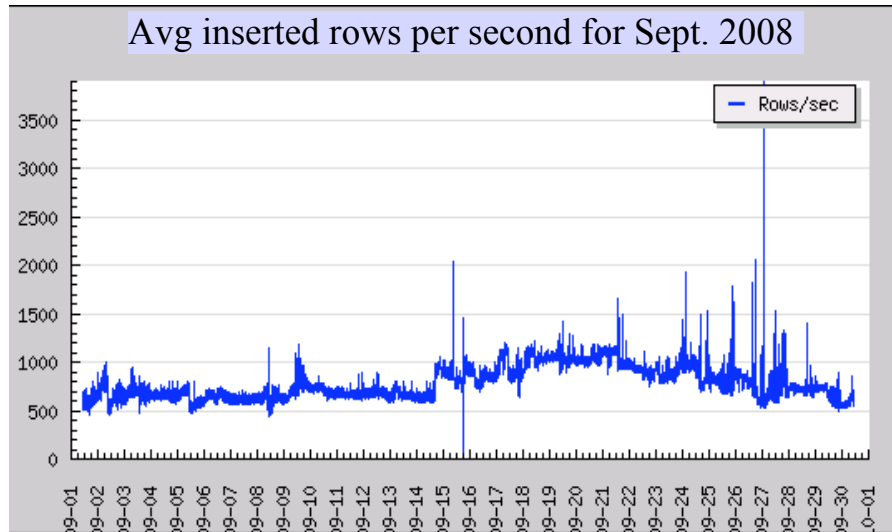




Replicating the PVSS Archive




- Load on the replication : avg. 4 GB per day of data
- Replication activity can be measured in LCR/sec (Logical Change Record)
- The limit on the APPLY on the destination DB is about 3000 LCR/sec
- Activity above 700-800 rows/sec is considered over the normal and we work with the sub-detectors to keep the rate within that range





Snapshot from the PVSS_APPLY in a moment of compensating the backlog





Optimizations on the ATLAS PVSS Archive



- Compression on the PVSS data – done on a regular basis on the old ‘switched history’ tables - (move/compress to a different TBS)
 - factor of 2 decrease in storage of the archived tables.
 - 35 % saved space from the index compressionThus the overall PVSS space saving is about 45%
- Intention to delete yearly window of data from the online server. Procedure still to be tested, as replication must block this operation on the offline.
- Usage offline: PL/SQL libraries are being developed to control and optimize the access to the offline archive, so that free-form SQL will not overuse resources.



PVSS data volumes



- Every day PVSS 'eats' about 8 GB of the disk space (table + index segments)
- Non-compressed PVSS data (estimated on 07.11.2008)

ATONR : 606 GB

ATLR : 593 GB

The difference comes from the fact that on the ATONR the indexes are not compressed on the 'current' tables (default for the PVSS)

- Compressed PVSS data -
about 1/2 TB on each of the DBs
- The policy is to keep the data for most recent 12 months on the online database ATONR (the older one will be deleted) and keep the replicated data on the ATLR forever.



Current status



- Currently several sub-detectors run in debug mode, thus sending more data to Oracle than is expected in stable running mode.
- On the other side in running more we may expect upto 30% increase because of the instability into the detector conditions
- With this in hand the maximum space we need would be 4TB. After applied compression, this requirements would 2,2 TB
- For 2009 we request 3 TB and we believe it would be sufficient



Conclusions



- We got almost a factor of 2 decrease in the disk space usage by compressing the PVSS data
- The PVSS replication worked quite well since it was put in production mode. However, we have to keep the insert rate in the reasonable range of 700-800 LCR/sec that we have now. Work with the sub-detector experts is ongoing.
- PVSS data retention on the ATONR agreed to be 12 months (sliding window). Procedure to be tested on the INTR database