

# Database & Client Monitoring

#### Radovan Chytracek, Giacomo Govi CERN IT Department





- Motivation
- What to monitor
- Database server side monitoring
- Client side monitoring
- Status
- Conclusions



- Constant fight on three front-lines
  - HW (CPU, network, storage)
  - Server side (OS, DB, schema design)
  - Client side (bugs, wrong practices, queries)
- HW can be improved by better "iron"
- SW should be safe by not making mistakes
  - New or upgraded apps have the same or new bugs
  - Good schema designs is often difficult
  - Following good practices seems to be tough job too
- DBAs are inevitable
  - Spit out & analyze the bad things & give advices 24/7



- Server side
  - Various levels, session tracing is the most used one
  - Must ship the server trace file back to user
  - Security issues, some development effort required
  - Supported by LCG SW (POOL Oracle plug-in)
- Client side
  - Required to make the whole picture complete
  - Does not exist out-of-the-box
  - Application code instrumentation needed
  - Often connected to monitoring systems
  - Support being built into LCG SW



- Allow DBAs and developers inspect the current state of a database instance in an easy way without a need for complex software
- Goal is to enable database & application level monitoring in coherent way with the existing OS level monitoring provided by LEMON
- Easy access via web interface to quantities and trends describing current database instance behavior with keeping their history and possibility to zoom in a given time period



- Considered OEM repository but requires OEM infrastructure in place and not all instances are in OEM
  - What if OEM is down?
  - Data kept only 1 month
- Source: instance's SYS.V\$... performance views
- The baseline DB metrics extracted from SYS.V\$SYSSTAT dynamic performance view
  - Recalculated exactly the same way as done in OEM
  - Examples: SQLNet in/out data rate, logical I/O, physical I/O, SQL per second...
- Application level monitored via SYS.V\$SESSION... views



### **DB sensor for LEMON version I**

- SQL script executed via SQLPlus
  - Connecting to the locally detected database
- Shell driver script executed by a simple Perl sensor in LEMON framework
  - Detects local DB settings from /etc/oratab file and names of local oracle daemons (pmon...)





### **DB sensor for LEMON version II**

- SQL queries still executed via SQLPlus
  - Connecting to the locally detected or remote database
- SQL\*Plus tool wrapped in Perl class module
- Allows to keep single permanent connection only
- DB instance & SQL\*Plus tool auto detection
- DDL & DML and queries API provided





- DB LEMON sensors tested on various systems
  - Single DB instances, Oracle 9i/10g
  - LEMON databases
  - RAC systems, Oracle 10g
- Web display & metrics deployed in LEMON development version
  - Little development needed
  - DB metadata read from OEM repository
  - Clicking a metrics graph in detailed view jumps to zoomable time period view similar to OEM
  - RAC cluster databases shown as computer cluster in LEMON



### **Database Instance Monitoring**

bu	Home Documentation	Alarms PCFinder	Metrics Help
ori	Db_service info: RAC for CMS applications       18 Oct 2005 Tue 04:20:53		
lit	Database Service/RAC Information		Logons and Cursors - last day
2	Database version(s):		
2	Number of CPUs:	0	
S.N.	Machine names:		8 0.5 k
ġ.	Select from sids:	None	
0	User stats	Tablespace stats	0.0 06:00 12:00 18:00 00:00
	User CPU stats	Wait events	<ul> <li>Curr logons aver: 177.34 max: 187.00 min: 171.00 curr: 177.62</li> <li>Curr cursors aver: 975.32 max: 1109.60 min: 896.65 curr: 1095.76</li> </ul>
			SQLNet utilization - last day 200 k 100 k 00 00 00 00 100 k 100 k 100 k 100 k 100 k 100 k 100 k 100 k 12:00 18:00 18:00 18:00 18:00 18:00 18:00 18:00 18:00 10:00 18:00 18:00 18:00 18:00 10:00 18:00 18:00 18:00 10:00 18:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 19:00 10:0
	Search host: Virtus	Il Clusters I Clusters I Racks I HW mo	dels   Databases  Virtual Organizations  PDUs  Power Locations

17 January 2006

Radovan Chytracek, CERN IT Department



### **Application Level Monitoring**



17 January 2006



- The monitoring of the WAIT events in progress
  - Performance tuning is difficult without having these
- Deployment of the new LEMON DB sensor on all physics databases
  - Currently running on selected instances and few RAC nodes
- Oracle installation procedures need to be updated to include proper monitoring settings



- CORAL code instrumentation to enable client activity tracing
- Information about DB client behavior
- Helps discovering changes in DB access pattern
- Security & Auditing
- Format easy to import/export
  - Databases, spreadsheets, data mining

#### Third party monitoring systems interoperability



## **Client Monitoring Status**

- User & Developer interfaces defined in CORAL
- Per session monitoring
- Basic MonitoringService implementation provided
- CSV & SQLite exports is being worked on
- Plans to enable interoperability with CMS Mona Lisa and Fermilab, UDP based, client monitoring system





- Performance tuning and testing are essential for the resource planning
- Proper monitoring helps to save some resources too
- DB Server monitoring done in collaboration within CERN IT Department with LEMON team
- Client side monitoring is in CORAL code and will be finalized in the coming weeks
- Collaboration in client side monitoring is open activity & third party support is foreseen