

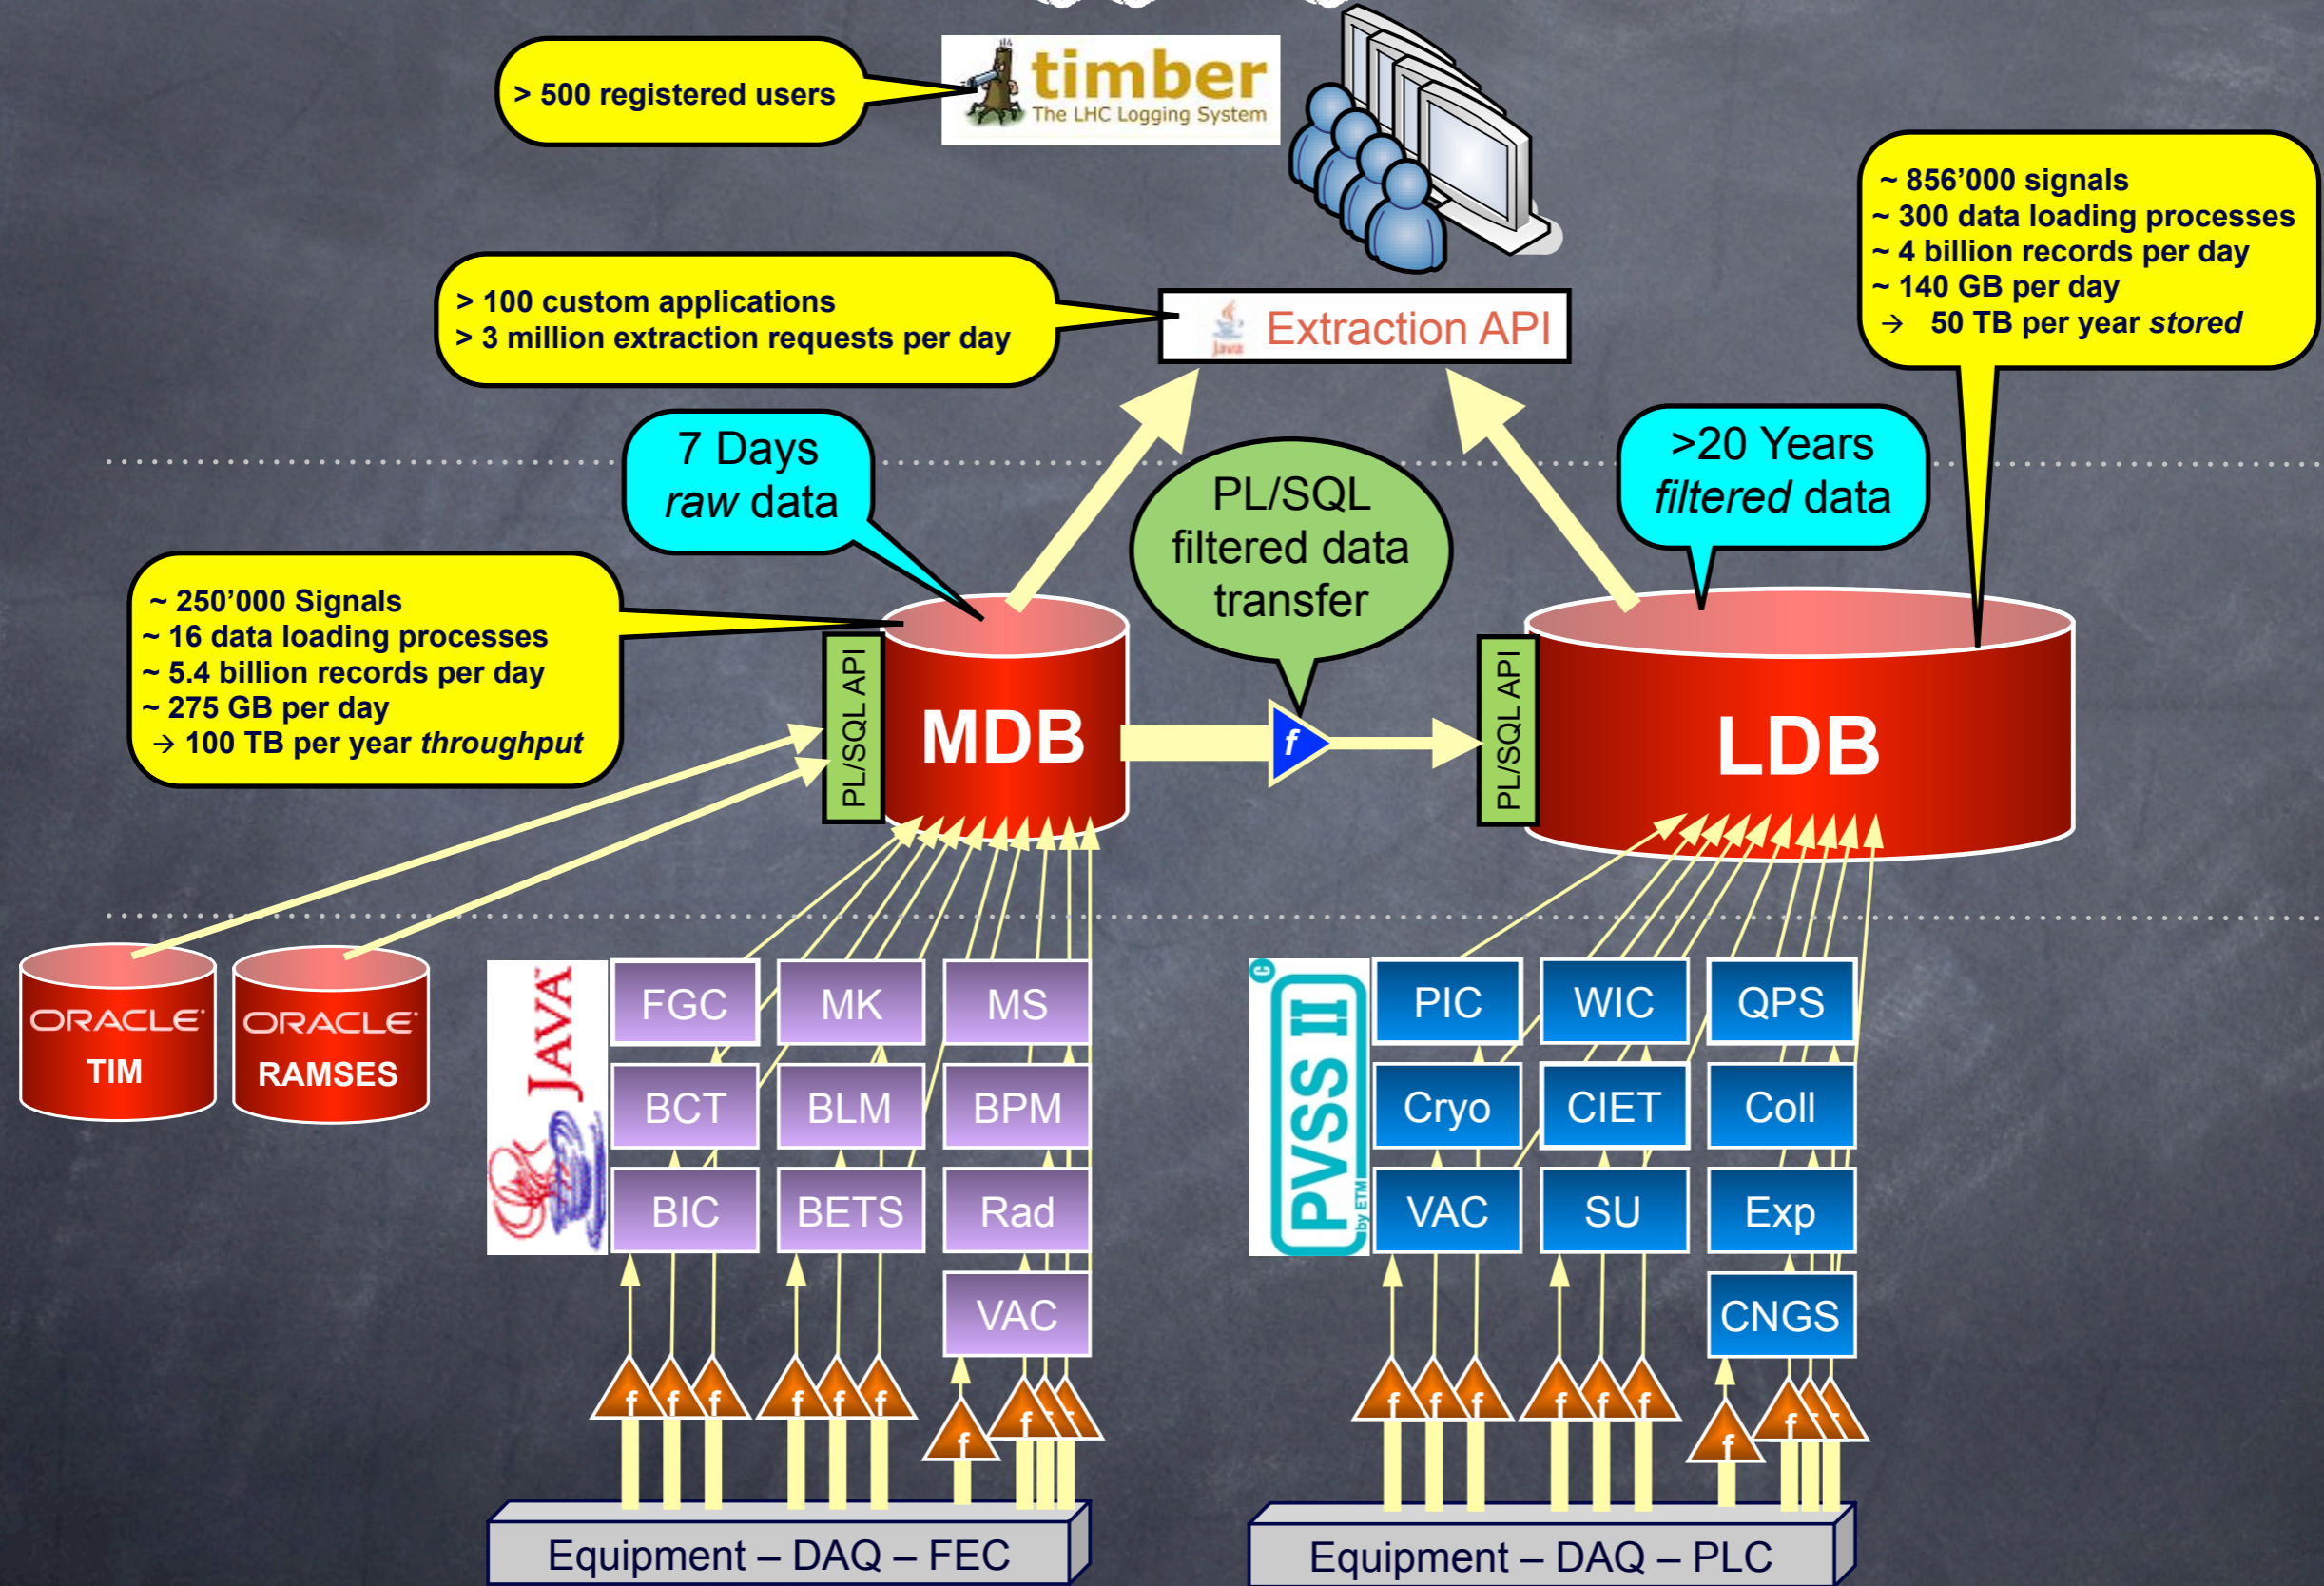
Instrumentation

An Essential Ingredient
for Performance Tuning

Chris Roderick, BE/CO

IT/DB Oracle Lectures, June 2012

Accelerator Logging Service Overview



Data Providers Data Persistence Data Consumers

Mission Critical Service !

Lots of Data !!

Lots of Users !!!

Be wary of the users!



Big difference between how:


You think users will read/write data

Users tell you they intend to read/write data

Users actually read/write data

Accelerator Logging Service

Lots of Data!



System misuse

can have severe

negative impact on stability

Instrument everything!

What is Instrumentation?

Capturing information about
system activity in real time
($\&$ over time)

Who?

What?

Where?

How?

How Long?

We **know** what the system is doing

We **know** how the system is performing

Throughput & Response times

Oracle is heavily instrumented out of the box

- **ASH** (using `v$sqlsession`, `v$sqlsession_wait` - real time)
- **AWR** (constant data capture over time)
- **ADDM** (analysis of AWR data)
- **Tracing** (on demand, per session, to file)
- **OEM** (exposing all of the above in a nice GUI)

Oracle instrumentation

does not include contextual information

- Already foreseen **MODULE, ACTION, CLIENT_INFO**

```
DBMS_APPLICATION_INFO.SET_MODULE (  
  module_name IN VARCHAR2,  
  action_name IN VARCHAR2);
```

```
DBMS_APPLICATION_INFO.SET_ACTION (  
  action_name IN VARCHAR2);
```

```
DBMS_APPLICATION_INFO.SET_CLIENT_INFO (  
  client_info IN VARCHAR2);
```

- Non-transactional

- Minimal performance impact (ms)

DBMS_APPLICATION_INFO in action

```
select osuser, machine, program, process, module, action, client_info
from v$session
where status = 'ACTIVE'
and action is not null
and type = 'USER';
```

Set via DBMS_APPLICATION_INFO

OSUSER	MACHINE	PROGRAM	PROCESS	MODULE	ACTION	CLIENT_INFO
copera	cs-ccr-apopl.cern.ch	JDBC Thin Client	1234	timber	QD: DN FF	bohl
copera	cs-ccr-logging0.cern.ch	JDBC Thin Client	1234	LHC-PRO-BI-BI_DEFAULT DW	loadNUMERICData	LHC-
copera	cs-ccr-logging0.cern.ch	JDBC Thin Client	1234	LHC-PRO-BI-BI_DEFAULT DW	loadTEXTUALData	LHC-
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP1 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP3 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP4 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP5 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP6 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP7 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	BLM-BLM_IP8 DW	loadNUMERICData	BLM:
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	LHC-BT-DEFAULT_1 DW	loadNUMERICData	LHC-
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	LHC-OP-TDC DW	loadNUMERICData	LHC-
copera	cs-ccr-logging1.cern.ch	JDBC Thin Client	1234	LHC-RF-DEFAULT_1 DW	loadNUMERICData	LHC-
oracle	dbsrva252.cern.ch	oracle@dbsrva252.cern.ch (J001)	20127	DATA_LOADING_5	LOG_LOGGING_DN	DL_GI
oracle	dbsrva252.cern.ch	oracle@dbsrva252.cern.ch (J002)	20138	DATA_LOADING_6	COLLECT_DN_CAD	DL_GI
oracle	dbsrva252.cern.ch	emagent@dbsrva252.cern.ch	22843	emagent_AQMetrics	DEQ	(nul
oracle	dbsrva252.cern.ch	oracle@dbsrva252.cern.ch (PZ99)	11092	emagent_SQL_rac_database	apply_queue_buffq	(nul

Accelerator Logging Service

- BigBrother Java package

- BIG_BROTHER

- DBMS_APPLI

- provides **transpa**

- audit trails (o

```
DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE(  
  service_name      IN VARCHAR2,  
  module_name      IN VARCHAR2,  
  action_name      IN VARCHAR2 DEFAULT ALL_ACTIONS);
```

```
DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(  
  service_name      IN VARCHAR2,  
  module_name      IN VARCHAR2 DEFAULT ANY_MODULE,  
  action_name      IN VARCHAR2 DEFAULT ANY_ACTION,  
  waits            IN BOOLEAN DEFAULT TRUE,  
  binds            IN BOOLEAN DEFAULT FALSE,  
  instance_name    IN VARCHAR2 DEFAULT NULL);
```

- session tracing (via DBMS_MONITOR)

- resource management (using Resource Manager)

- access control (e.g. VPD / RLS)

- Accurate instrumentation in 3-Tier environment

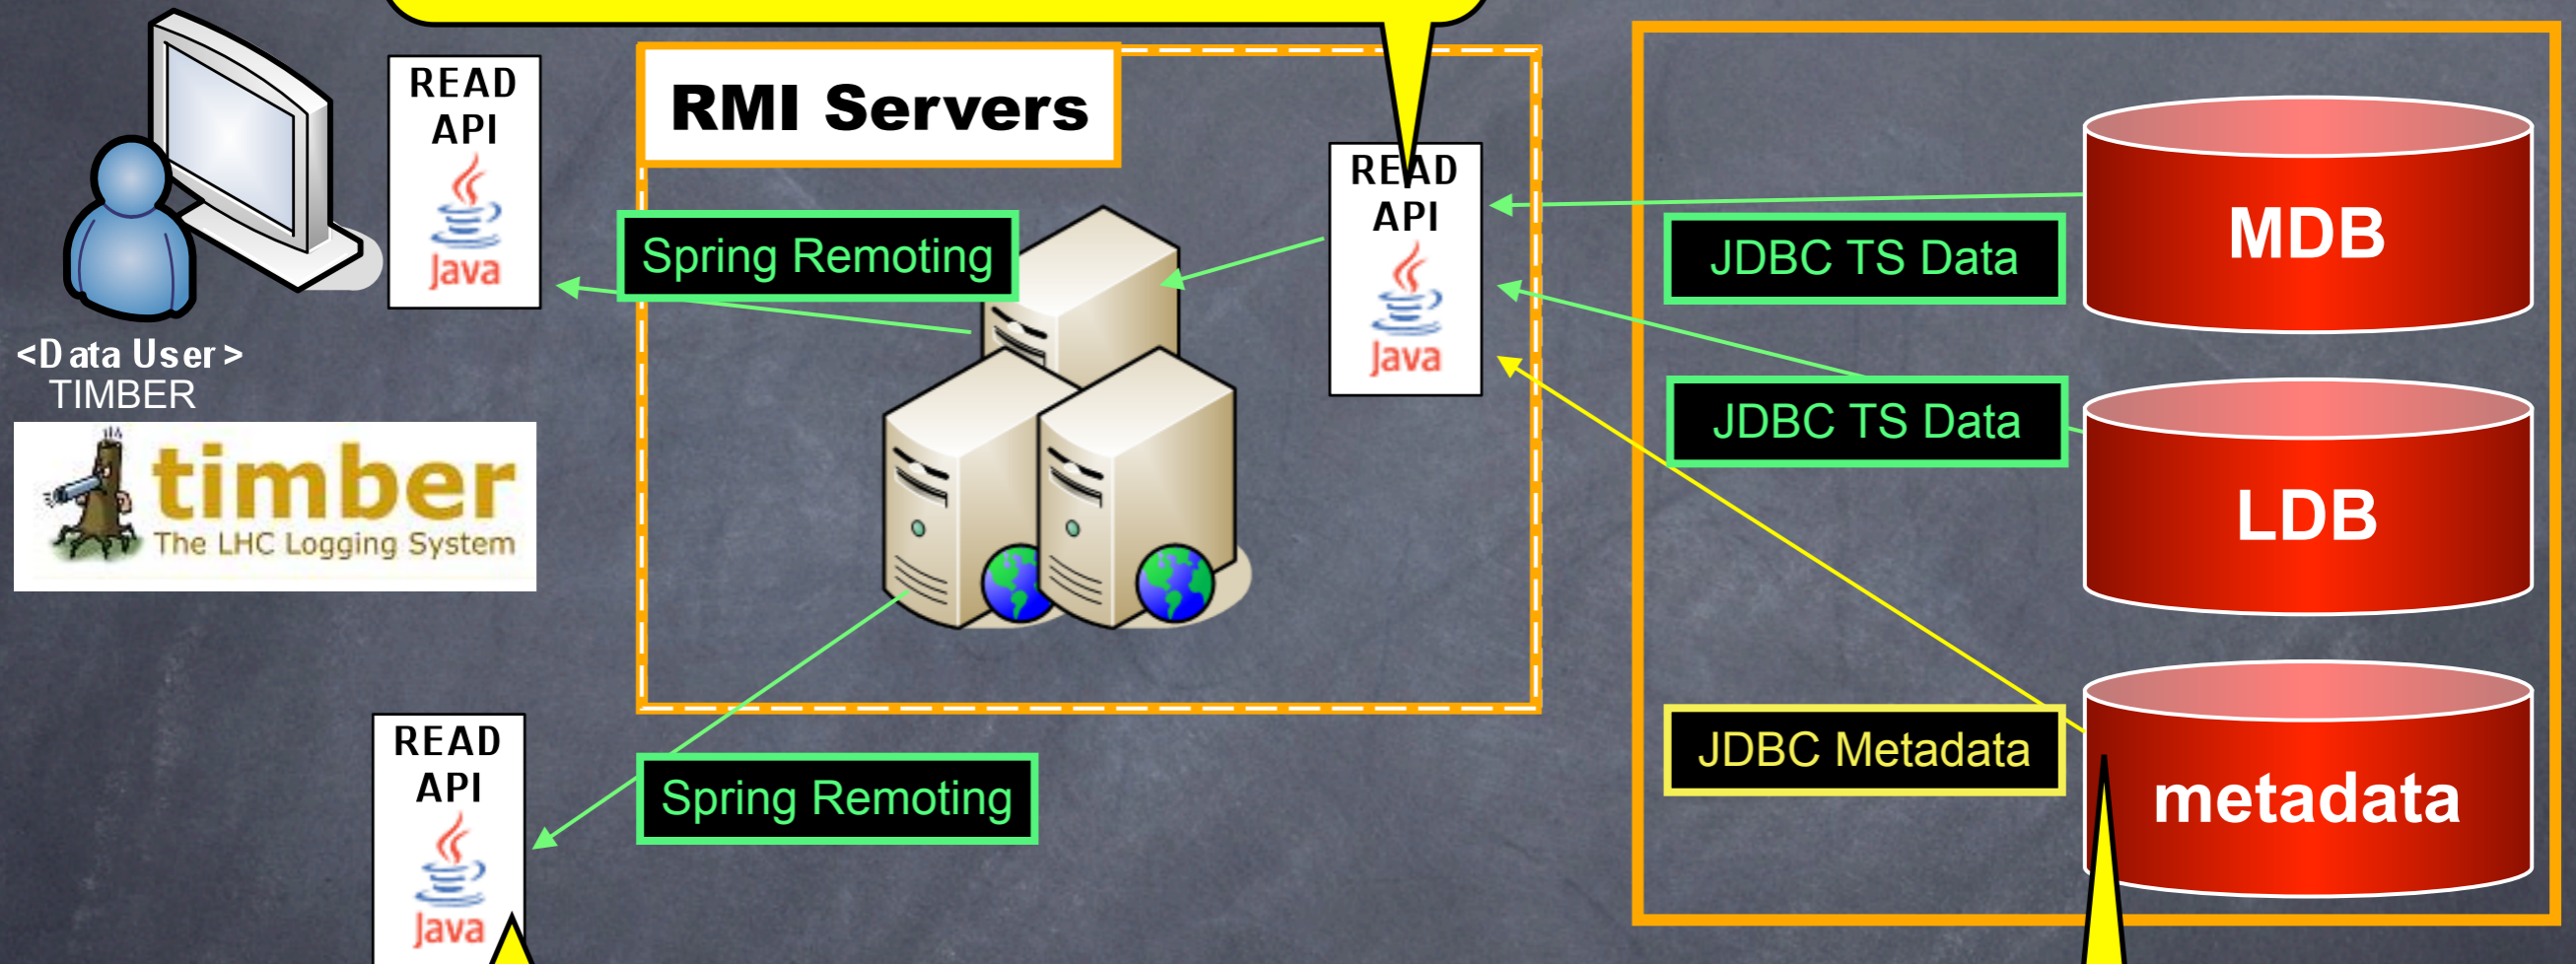
Accelerator Logging Service

- **BigBrother** Java package
 - **BIG_BROTHER** PL/SQL package
 - **DBMS_APPLICATION_INFO**
- provides **transparent hooks** to other uses:
 - audit trails (custom logging, profiling)
 - session tracing (via **DBMS_MONITOR**)
 - resource management (using Resource Manager)
 - access control (e.g. VPD / RLS)
- **Accurate instrumentation** in 3-Tier environment

Database Instrumentation Example

CERN Accelerator Logging Service

Server API instruments actions using BigBrother calls to BIG_BROTHER.



Client API obliges provision of username/application name.
Automatically picks up IP/host.

BIG_BROTHER package wraps calls to DBMS_APPLICATION_INFO.

Results

Session User Type DB User Type Kill Session Disable SQL Trace Enable SQL Trace

Select All | Select None

Select	SID	DB User	Program	Service	Module	Action	Client	Machine	OS User
<input type="checkbox"/>	373		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP7 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	404		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP1 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	405		oracle@ (J004)	ACCMEAS_JOBS_S	DATA_LOADING_8	COLLECT_DN_FD	DL_GRP_8@		
<input type="checkbox"/>	437		oracle@ (J001)	ACCMEAS_JOBS_S	DATA_LOADING_1	LOG_LOGGING_DVN	DL_GRP_1@		
<input type="checkbox"/>	442		oracle@ (J002)	ACCMEAS_JOBS_S	DATA_LOADING_3	COLLECT_DVN_GD	DL_GRP_3@		
<input type="checkbox"/>	449		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP5 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	453		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP3 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	492		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP2 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	505		oracle@ (J003)	ACCMEAS_JOBS_S	DATA_LOADING_2	COLLECT_DN_FD	DL_GRP_2@		
<input type="checkbox"/>	546		JDBC Thin Client	ACCMEAS_S	BLMS-BLM_IP8 DW	loadNUMERICData	BLMS:lhc-blm-meas@172.18.202.215		
<input type="checkbox"/>	560		oracle@ (J000)	ACCMEAS_JOBS_S	DATA_LOADING_4	LOG_LOGGING_DN	DL_GRP_4@		

Session User Type DB User Type Kill Session Disable SQL Trace Enable SQL Trace

Top Modules

View Top Modules

Activity (%)	Service	Module
18.67	ACCMEAS_S	DAQ_CHAIN_VALIDATION
14.50	SYS\$BACKGROUND	
6.53	ACCMEAS_JOBS_S	DATA_LOADING_3
6.29	ACCMEAS_S	LHC_HUMP_ANALYSIS
5.04	ACCMEAS_JOBS_S	DATA_LOADING_1
4.98	ACCMEAS_JOBS_S	DATA_LOADING_2
3.36	ACCMEAS_S	
3.11	ACCMEAS_JOBS_S	DATA_LOADING_4
2.80	ACCMEAS_JOBS_S	DATA_LOADING_6
2.74	ACCMEAS_JOBS_S	DATA_LOADING_8

Total Sample Count: 1,607

Module: DATA_LOADING_3

Actions Activity Statistics

Top Actions

Action	Percentage
COLLECT_DVN_GD	48.1%
LOG_LOGGING_DVN	27.3%
COLLECT_DN_FD	6.5%
COLLECT_DVN_CAD	5.2%
LOG_LOGGING_DN	5.2%
COLLECT_DN_GD	3.9%
COLLECT_DN_CAD	1.3%
COLLECT_DVN_IAD	1.3%
UPDATE_DN_VARIABLE_LOG_STATUS	1.3%

Active Actions

View Active Actions

Action	Activity (% for the last 5 minutes)	Aggregation Enabled	SQL Trace Enabled
COLLECT_DVN_GD	48.1	FALSE	FALSE
LOG_LOGGING_DVN	27.3	FALSE	FALSE
COLLECT_DN_FD	6.5	FALSE	FALSE
COLLECT_DVN_CAD	5.2	FALSE	FALSE
LOG_LOGGING_DN	5.2	FALSE	FALSE
COLLECT_DN_GD	3.9	FALSE	FALSE
COLLECT_DN_CAD	1.3	FALSE	FALSE
COLLECT_DVN_IAD	1.3	FALSE	FALSE
UPDATE_DN_VARIABLE_LOG_STATUS	1.3	FALSE	FALSE

What about specific Performance Instrumentation?

• good_performance =
((time_spent / work) <= user_expectation)



• Track context specific system workload vs time spent, over time:

• Measure work being done


• Time taken to do it

• Record the results (details & aggregates)

Accelerator Logging Service

- ◉ Who - username
- ◉ What - module, action
- ◉ Where - IP, host
- ◉ How - application, method, parameters
- ◉ How long - elapsed time for every action
- ◉ Contextual Information (dimensions affecting performance) e.g.
 - ◉ How many records
 - ◉ How many signals
 - ◉ Time window
 - ◉ data manipulations or filters
 - ◉ Operation type - INSERT or MERGE (4x slower)

Measurements to Logging Transfer




8x jobs running every 5'

PL/SQL

- bulk read data by category
- apply custom filters
- bulk transfer

MDB to LDB Transfer Instrumentation



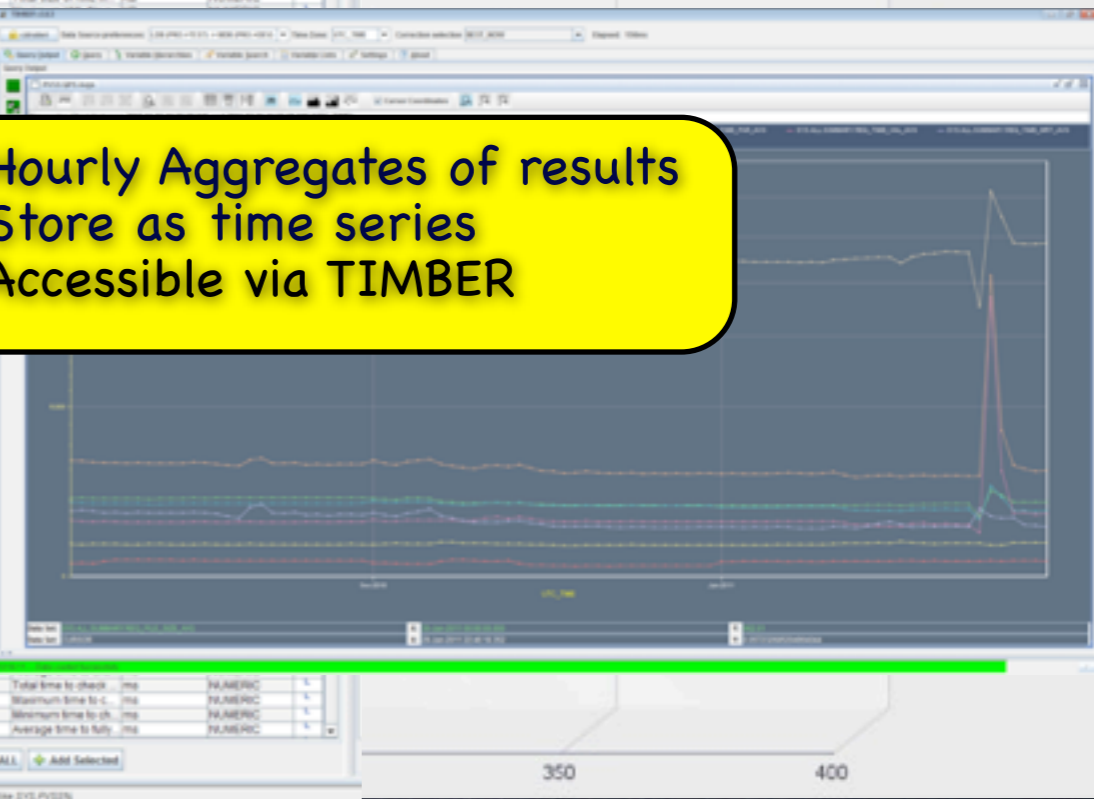
Capture details of every process execution
#records read, #records transferred, time spent, etc.
Keep details during 7 days

The screenshot shows the TIMBER application interface. On the left is a 'Hierarchy Variable Selection' tree with categories like ROOT, ADE, ATLAS, CMS, etc. The main area displays 'Search Results' for the variable 'SYS.PVSS%'. The results table includes columns for Variable Name, Description, Unit, and Category. Below the table are 'Add ALL' and 'Add Selected' buttons.

Variable Name	Description	Unit	Category
SYS.PVSS.PC.REQ.CNT.FAILED.SUBT	Number of times re-Requests	NUMERIC	
SYS.PVSS.PC.REQ.CNT.MERGED	No. of data loading Requests	NUMERIC	
SYS.PVSS.PC.REQ.FILE.SIZE.AVG	Average file size	KB	NUMERIC
SYS.PVSS.PC.REQ.FILE.SIZE.INT	Total size of file	KB	NUMERIC
SYS.PVSS.PC.REQ.FILE.SIZE.MAX			
SYS.PVSS.PC.REQ.FILE.SIZE.MIN			
SYS.PVSS.PC.REQ.FREQ.AVG			
SYS.PVSS.PC.REQ.OLDEST_STAMP			
SYS.PVSS.PC.REQ.REC.CNT.AVG			
SYS.PVSS.PC.REQ.REC.CNT.MAX			
SYS.PVSS.PC.REQ.REC.CNT.MIN			

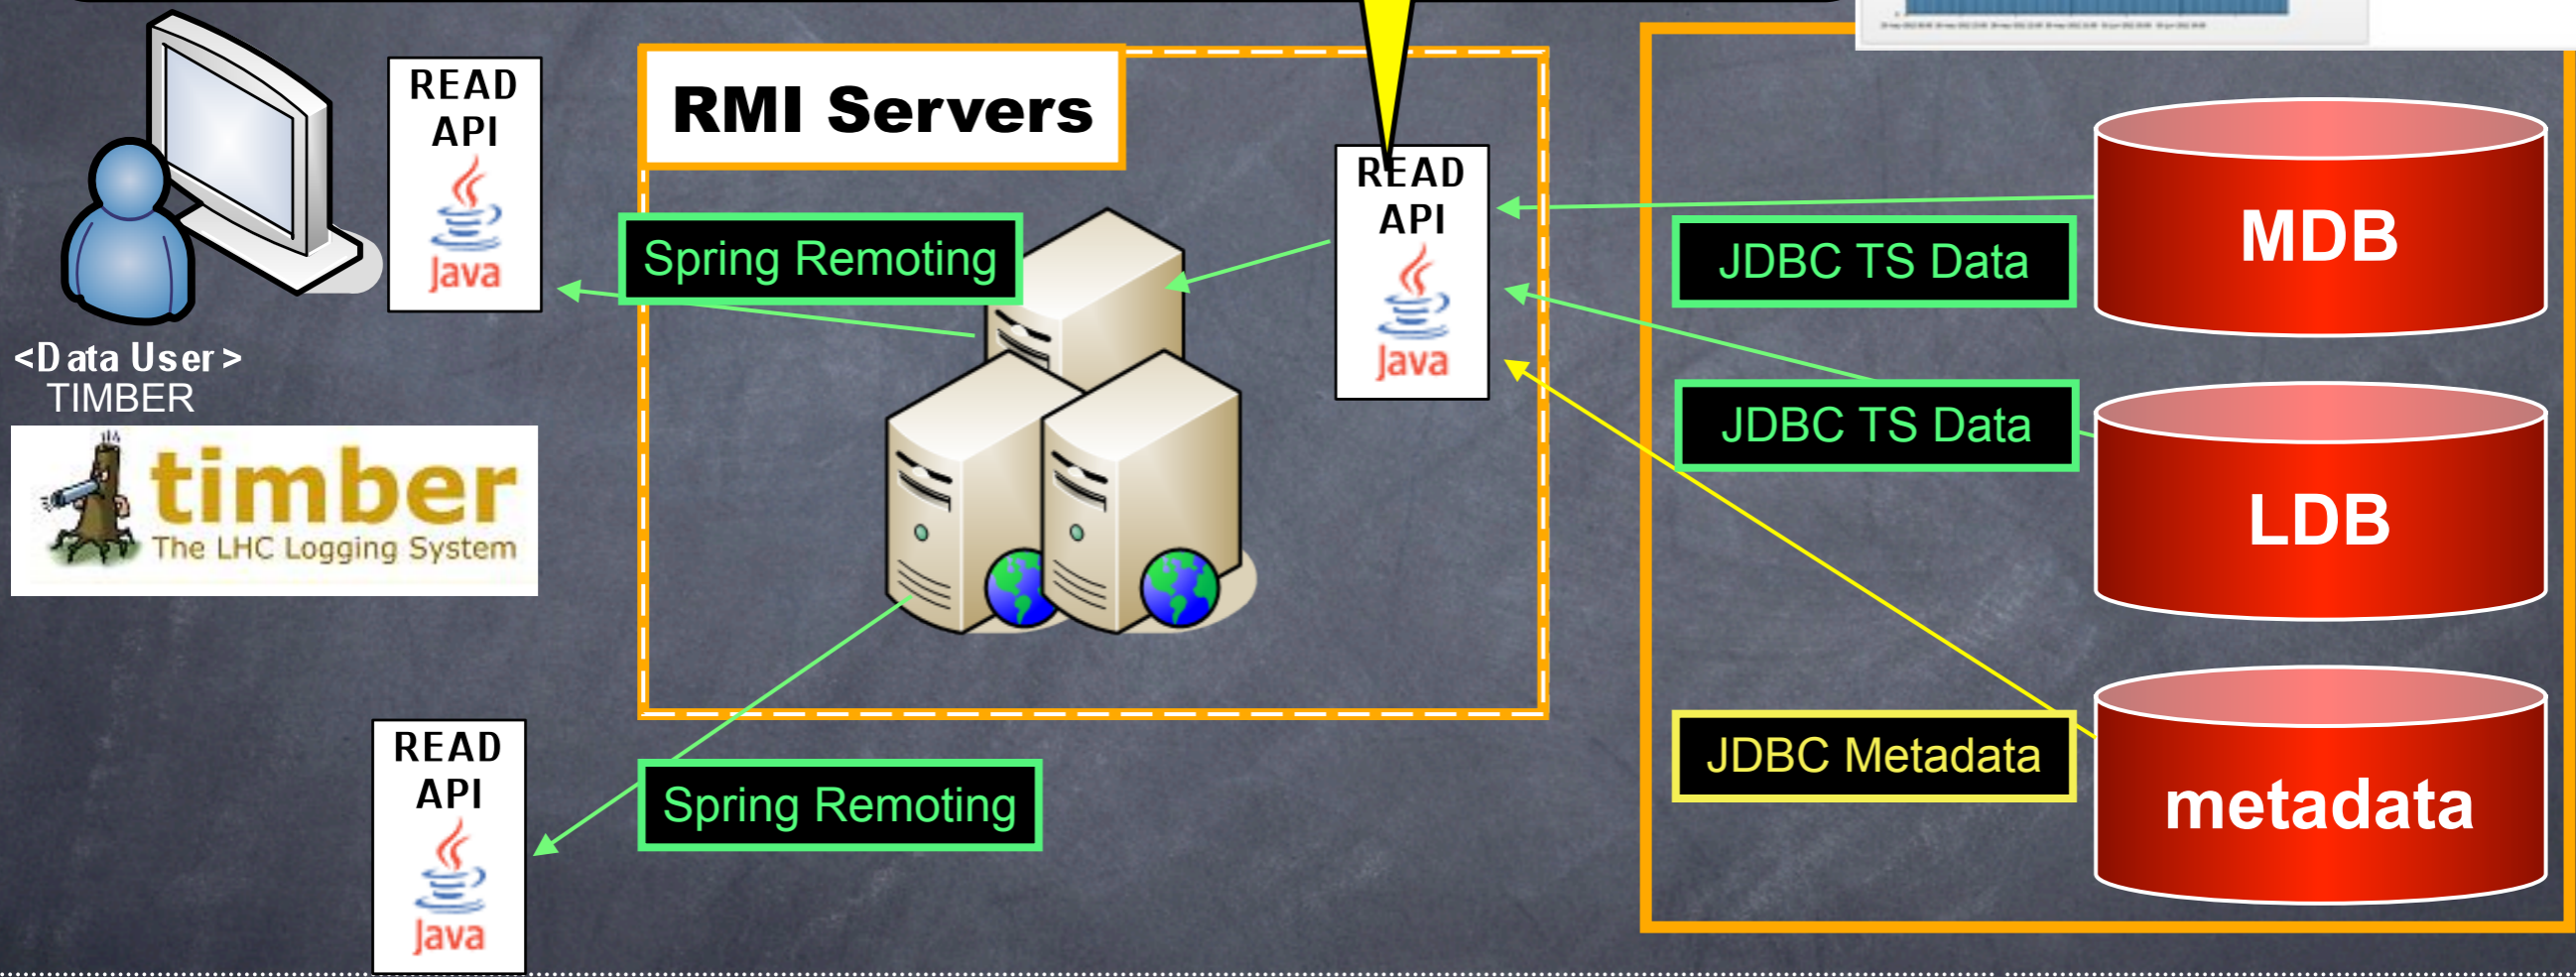
- DATA_LOADING_1
- DATA_LOADING_2
- DATA_LOADING_3
- DATA_LOADING_4
- DATA_LOADING_7
- DATA_LOADING_5
- DATA_LOADING_6
- DATA_LOADING_8

 Hourly Aggregates of results
Store as time series
Accessible via TIMBER



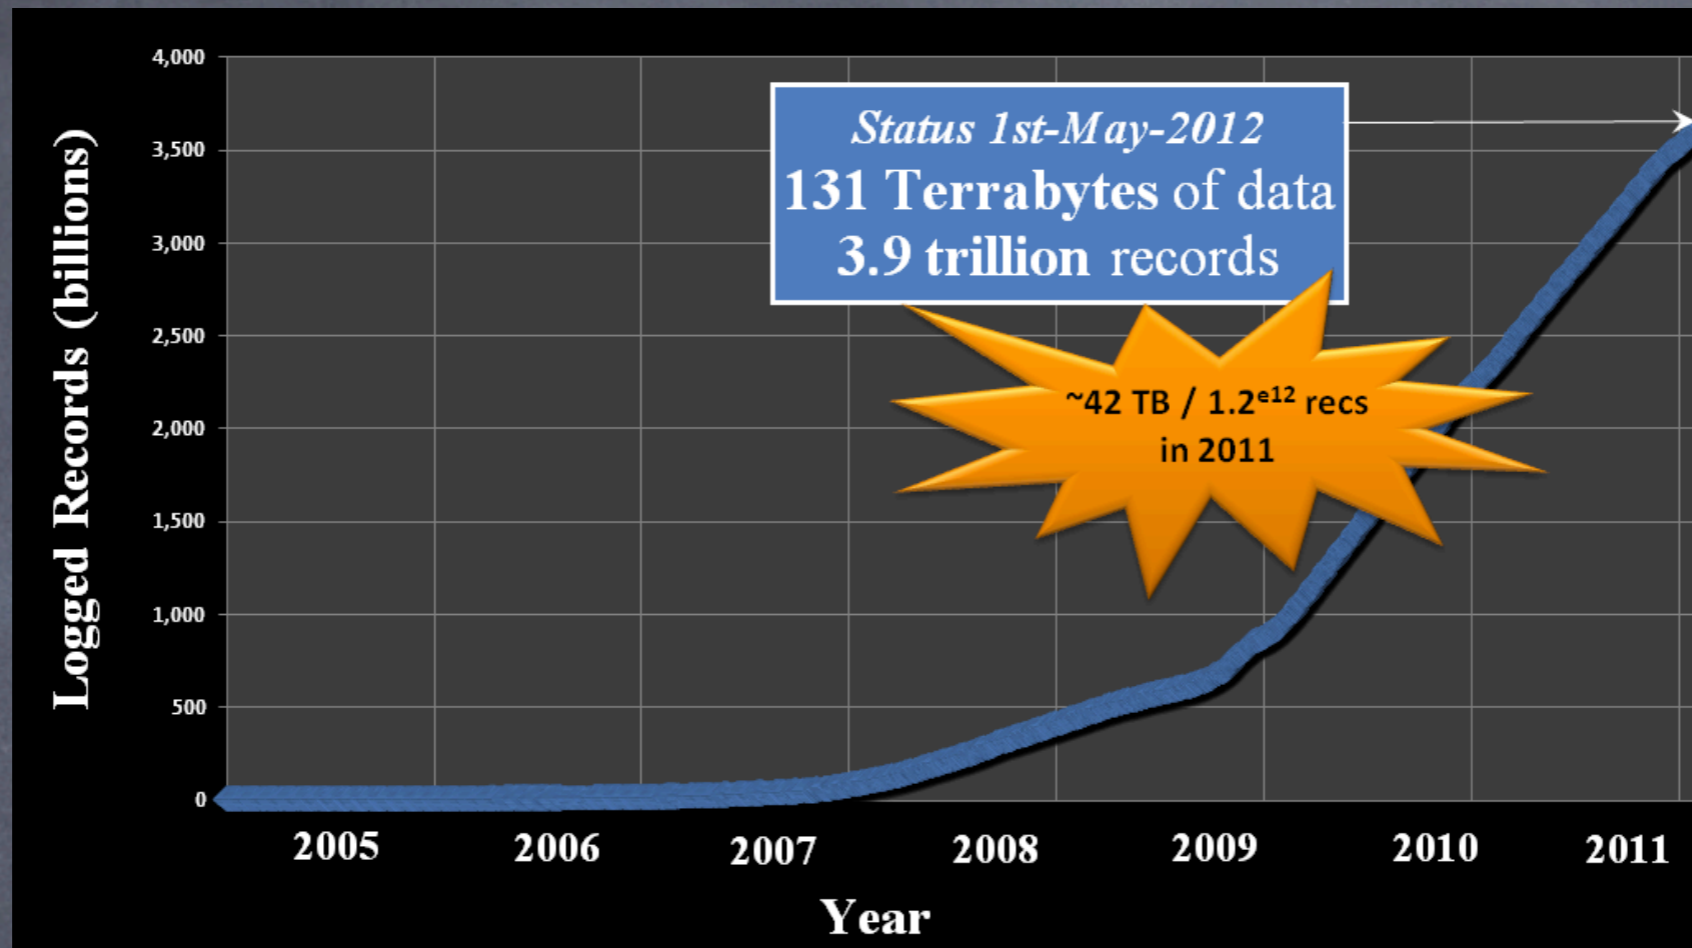
Data Extraction Instrumentation Example

Server API uses generic Tracing Component to instrument actions & measure performance.
Delayed batch writes of results to DB, analysis via APEX app.



Accelerator Logging Service

Many Performance Enhancements over the years



All based on understanding of system usage -
through instrumentation

Conclusions?

Good Instrumentation...

Facilitates Performance Tuning,
Scalability Planning, & Diagnostics

Enhances System Stability

Enables rapid reactive
& proactive Support

Questions?

Additional Use Cases

Code Examples

DBMS_APPLICATION_INFO wrapper

Generic Java Tracing Component

Chris.Roderick@cern.ch