

Method – GAPP

Used to Mine OEM 12c Repository and AWR Data

a Method to Diagnose and Predict Performance in Complex Architectures



Gerwin Hendriksen
Performance Architect
gerwin.hendriksen@gmail.com

Who I am...

- Worked for 1,5 year at Dedicate (Parity Solutions)
- Worked for 5 years at Oracle The Netherlands
 - First 1.5 year as RDBMS analyst
 - Rest worked as ARE EBS Performance and RAC
 - This included Content Lead of EBS with RAC world wide
- Worked for 2 years at IBM
 - Business Consultancy Services
- Worked for 3 years at AMIS Services BV
- Worked for 2.5 years at IT-eye
- Currently working for codecentric NL as Performance Architect

The Oracle logo, consisting of the word "ORACLE" in a bold, red, sans-serif font with a registered trademark symbol.The IBM logo, consisting of the letters "IBM" in a blue, sans-serif font with horizontal stripes through the letters and a registered trademark symbol.

Agenda...

- Real World Encountered Situations
- What is Method-GAPP
- End-User processes and infrastructure
- The Five D's from Method-GAPP
- OEM 12C data usage
- AWR data usage
- Conclusions
- Q/A

Real World Encountered Situations

Real World Encountered Situations...

- End user processes are sometimes too slow and the cause hard to be determined.
- The technical infrastructure has shared components like a SAN, Server, etc.
- The application is hard to be debugged, because it is a third party application.



- Virtualization
- Shared Components
- Many applications, sharing the same technical infra structure
 - 10's - 1000's
- Performance problems occur in production can't be all reproduced in test due to unknown inter-application influences

It would be great if...

- You could find the cause of such problems without
 - Changing code (hooks)
 - Technology dependency
- You could even find out the influence of other applications on our end user processes.
- You could predict what the performance impact would be before doing investments.

What is Method-GAPP

What is Method-GAPP...



Problem

DISALLOWED CHANGE TRACING

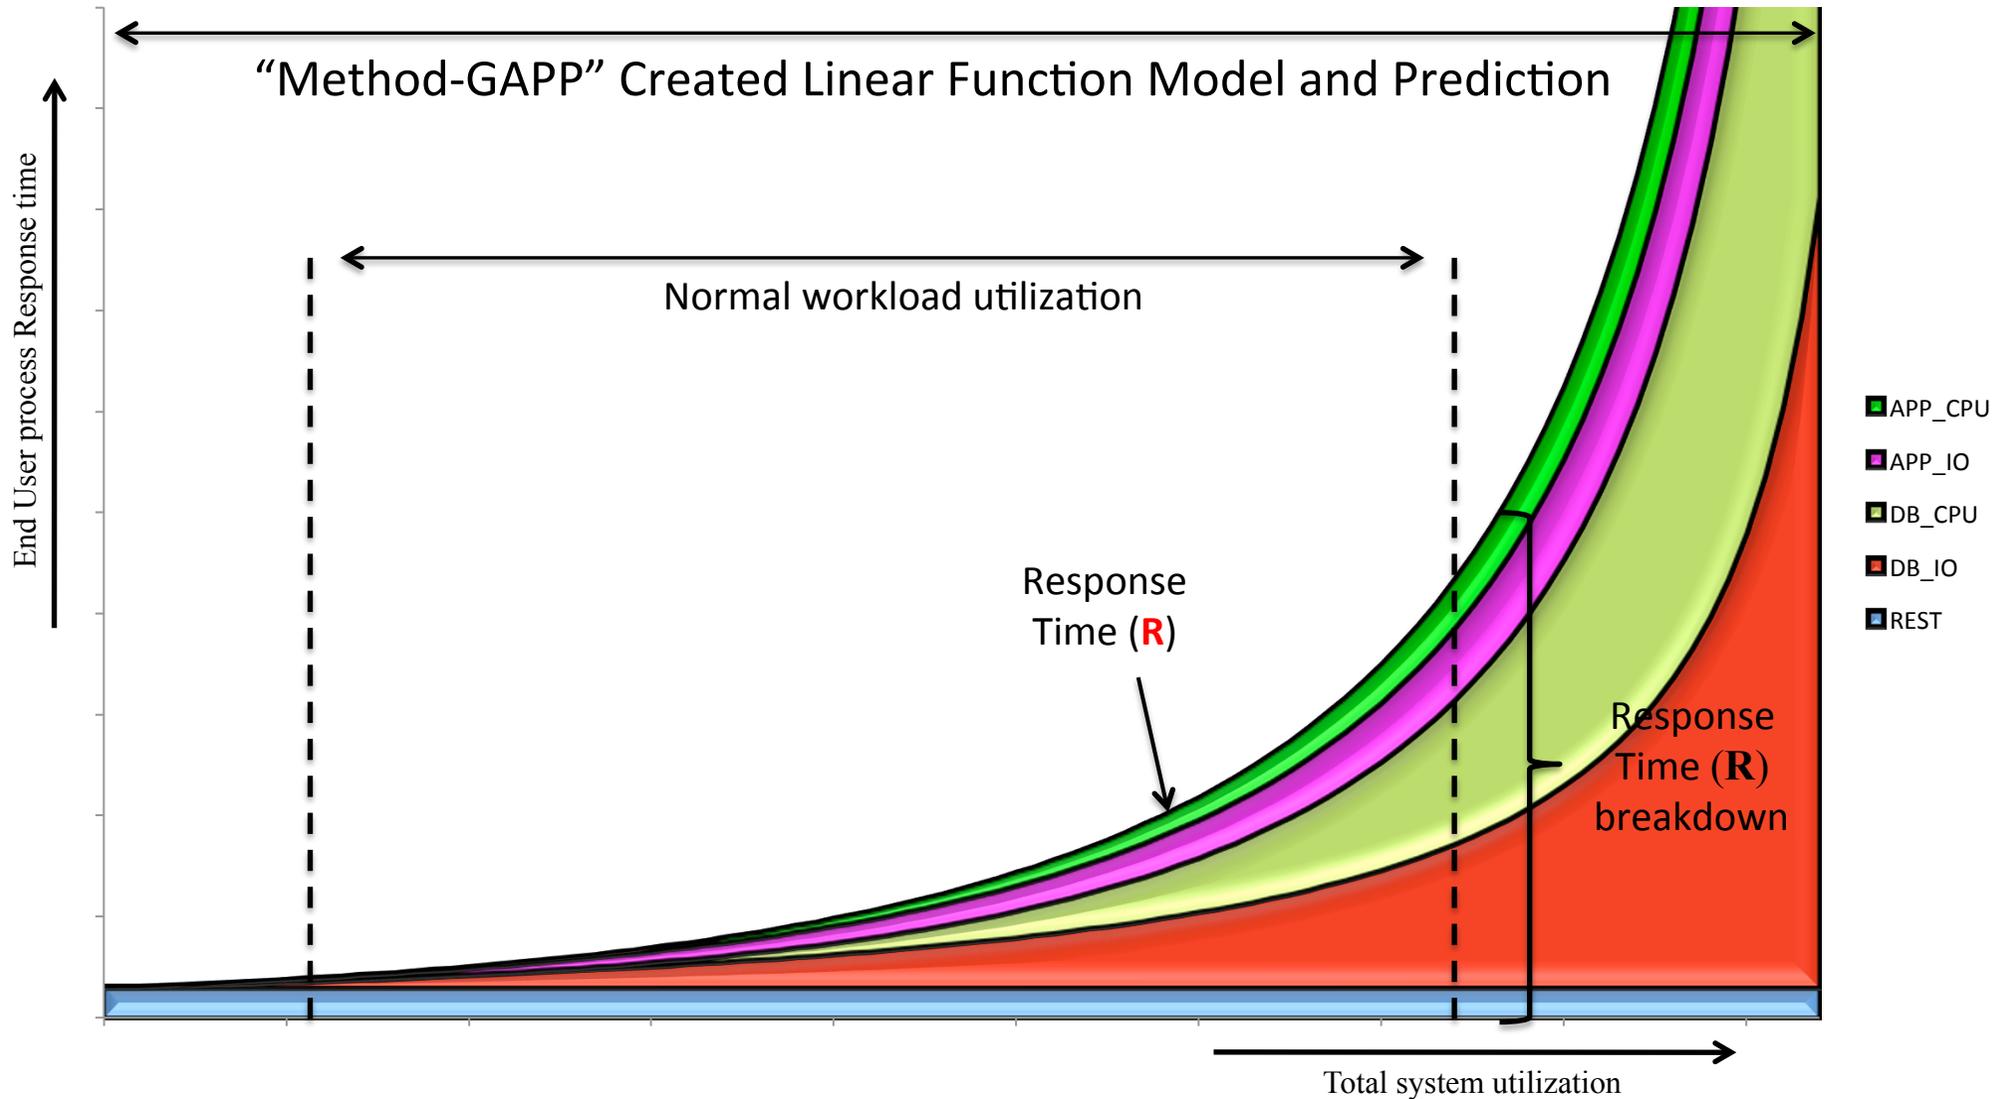
Solution

© G. Heintze 2008

Method-GAPP is a performance method which makes smart use of underlying queuing models and data mining to find bottlenecks in complex architectures, for specific end-user processes within an enterprise.

“GAPP” is an abbreviation of:
**“General Approach Performance
Profiling”**

Simplified capabilities Method-GAPP...



The Five D's from Method-GAPP...

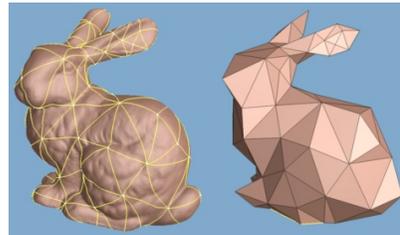
- Data Collection



- Data Synchronization



- Data Modeling



- Data Mining

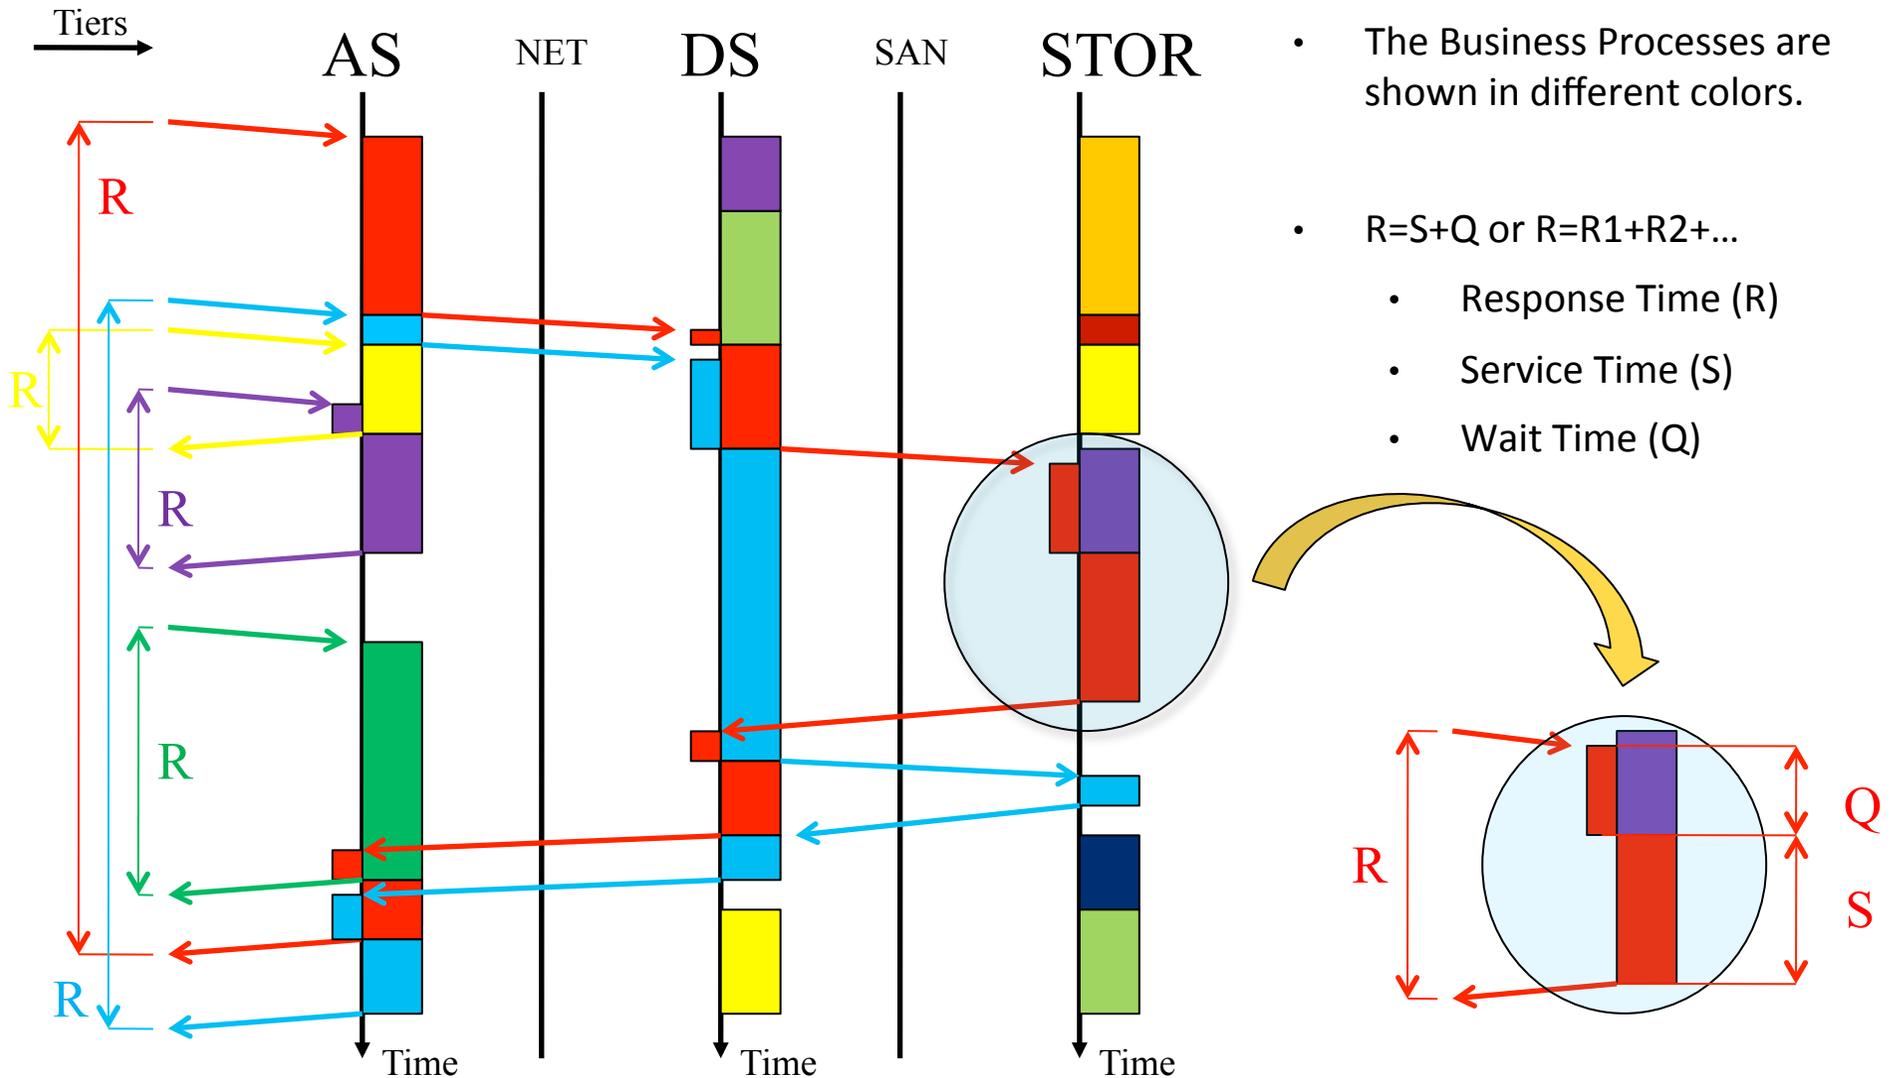


- Data Interpretation



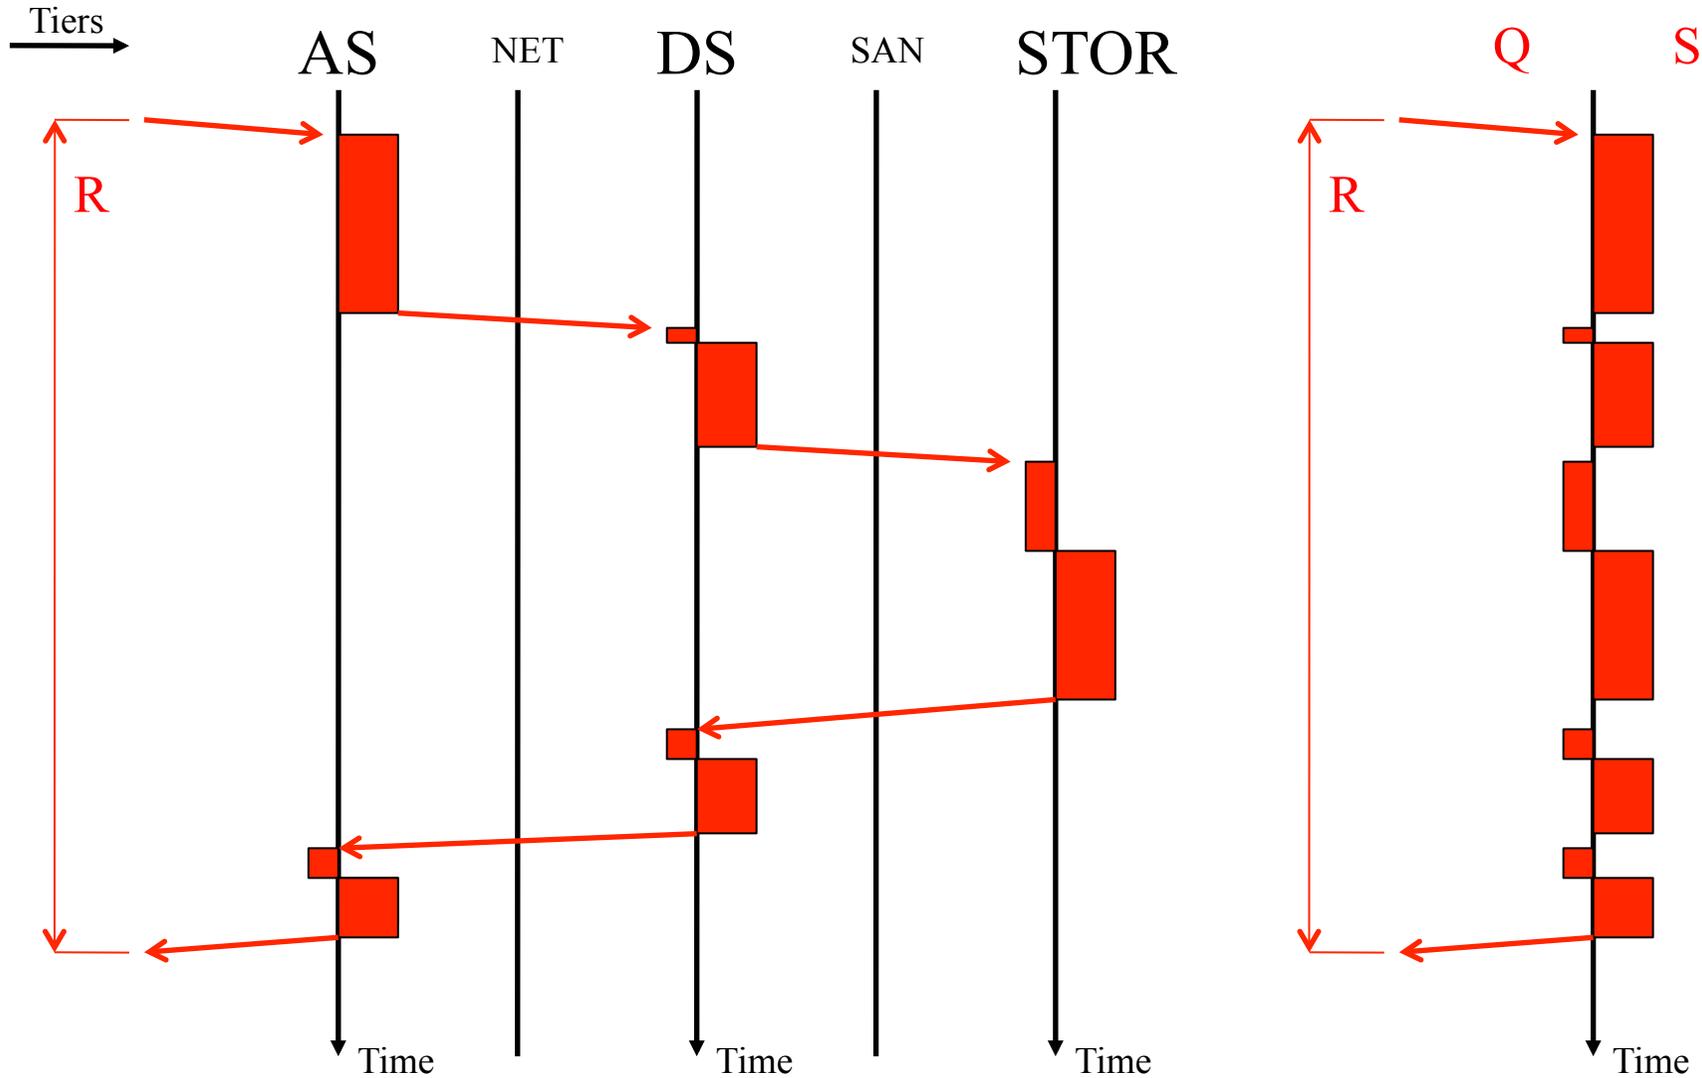
End-User processes and infrastructure

Method-R and the sequence diagram...



- The Business Processes are shown in different colors.
- $R=S+Q$ or $R=R1+R2+...$
 - Response Time (R)
 - Service Time (S)
 - Wait Time (Q)

Highlighted the "Red" process...



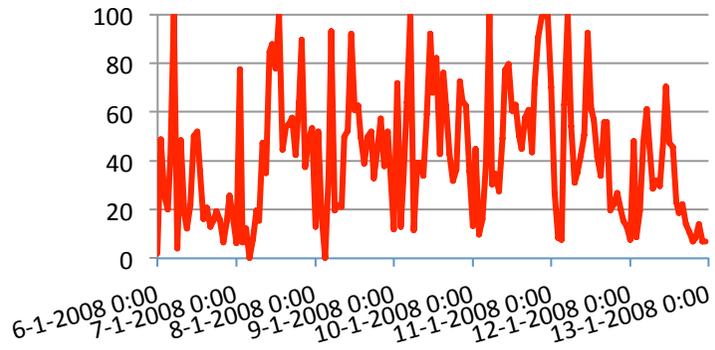
The First Two D's

Data Collection & Data Synchronization

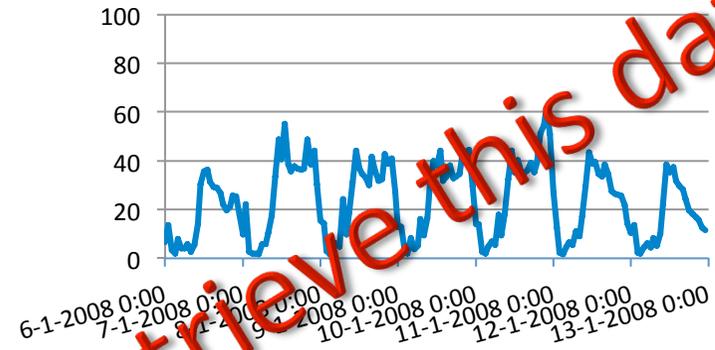
Primary and Secondary Components...

- Primary components are (System resources + Queues):
 - CPU
 - I/O (also network)
 - Memory
 - Queues (especially SOA environments)
- Secondary components are (rest) e.g.:
 - SQL response time
 - Events
 - Java threads
 - Virtual Machine CPU data
 - Etc.

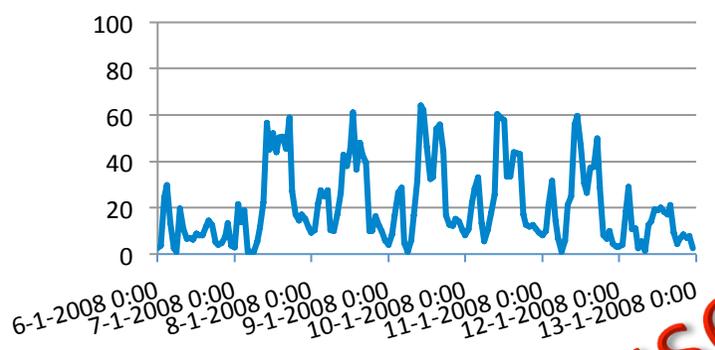
Data Collection and Data Synchronization...



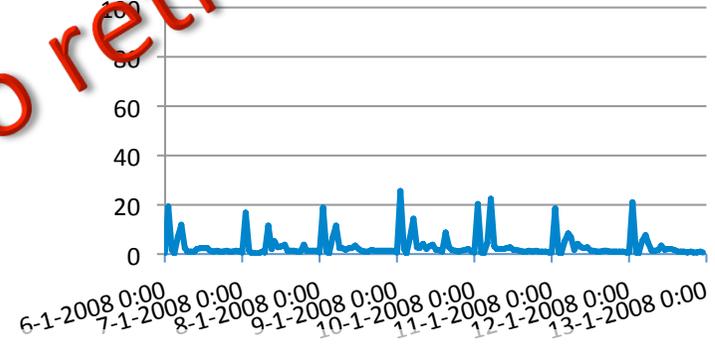
PCT_GT_5SEC



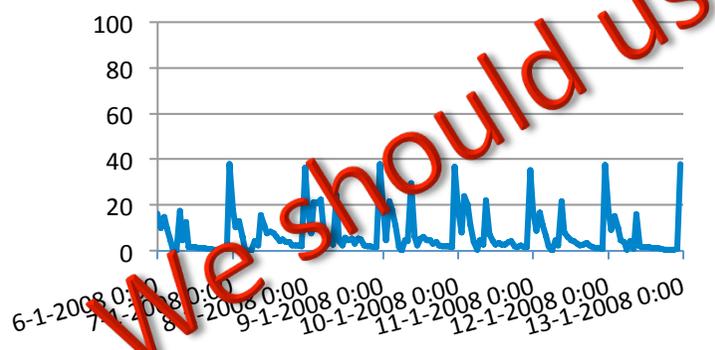
USER_E



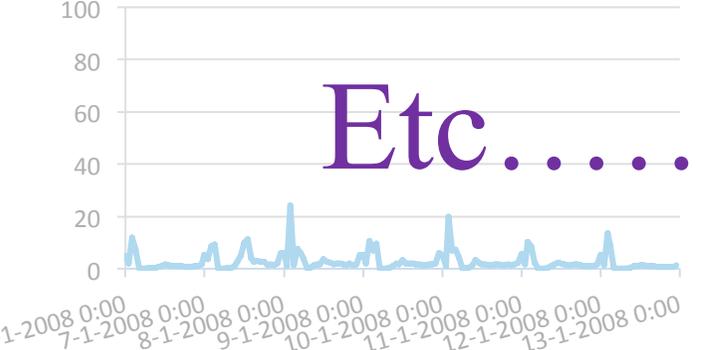
USER_Q



WAIT_E



WAIT_Q



HDISK5_Q

We should use OEM to retrieve this data!

Get the data from OEM 12C...

- When you want to use the data from the OEM12C repository make the sampling rate similar as AWR sampling.
- Use for the queues in the system, UDM's to have this information also available in OEM.
- For more information how data is gathered by the agent, you can checkout:
 - `$OH_AGENT12C/sysman/admin/scripts` (perl scripts)

Get the data from OEM 12C...

- 5 minute based (seven days) – “cpuUtil” example:

```
select to_char(collection_timestamp,'DD-MM-YYYY HH24:MI:SS') time
, target_name, metric_column, value
from sysman.mgmt$metric_details
where metric_column='cpuUtil'
and target_name like 'ghperf.gerwinhendriksen.local'
order by collection_timestamp;
```

- Hourly based – “cpuUtil” example:

```
select to_char(rollup_timestamp,'DD-MM-YYYY HH24:MI:SS') time, target_name
, metric_column, sample_count, average
, minimum, maximum, standard_deviation
from sysman.mgmt$metric_hourly
where metric_column='cpuUtil'
and target_name like 'ghperf.gerwinhendriksen.local'
order by rollup_timestamp;
```

All Metrics

Search

View

View By

- ▼ CPU Usage
 - CPU Idle Time (%)
 - CPU System Time (%)
 - CPU User Time (%)
 - CPU Utilization (%) of a cpu
- ▼ Disk Activity
 - Average Disk I/O Service Time
 - Average Outstanding Disk I/O I
 - Disk Block Writes (per second)
 - Disk Blocks Transferred (per se
 - Disk Device Busy (%)
 - Disk Reads (per second)
 - Disk Writes (per second)
- ▼ **Disk Activity Summary**
- ▼ File and Directory Monitoring
- ▼ Filesystems
- ▼ Host Storage Support
- ▼ Load
 - Active Logical Memory, Kilobyte
 - Active Memory, Kilobytes
 - CPU Utilization (%)
 - CPU in I/O Wait (%)
 - CPU in System Mode (%)
 - CPU in User Mode (%)
 - Free Memory (%)
 - Logical Free Memory (%)
 - Memory Page Scan Rate (per se
 - Memory Utilization (%)
 - Run Queue Length (1 minute a
 - Run Queue Length (15 minute
 - Run Queue Length (5 minute a
 - Swap Utilization (%)
 - Total Processes
 - Total Users

Data Synchronization...



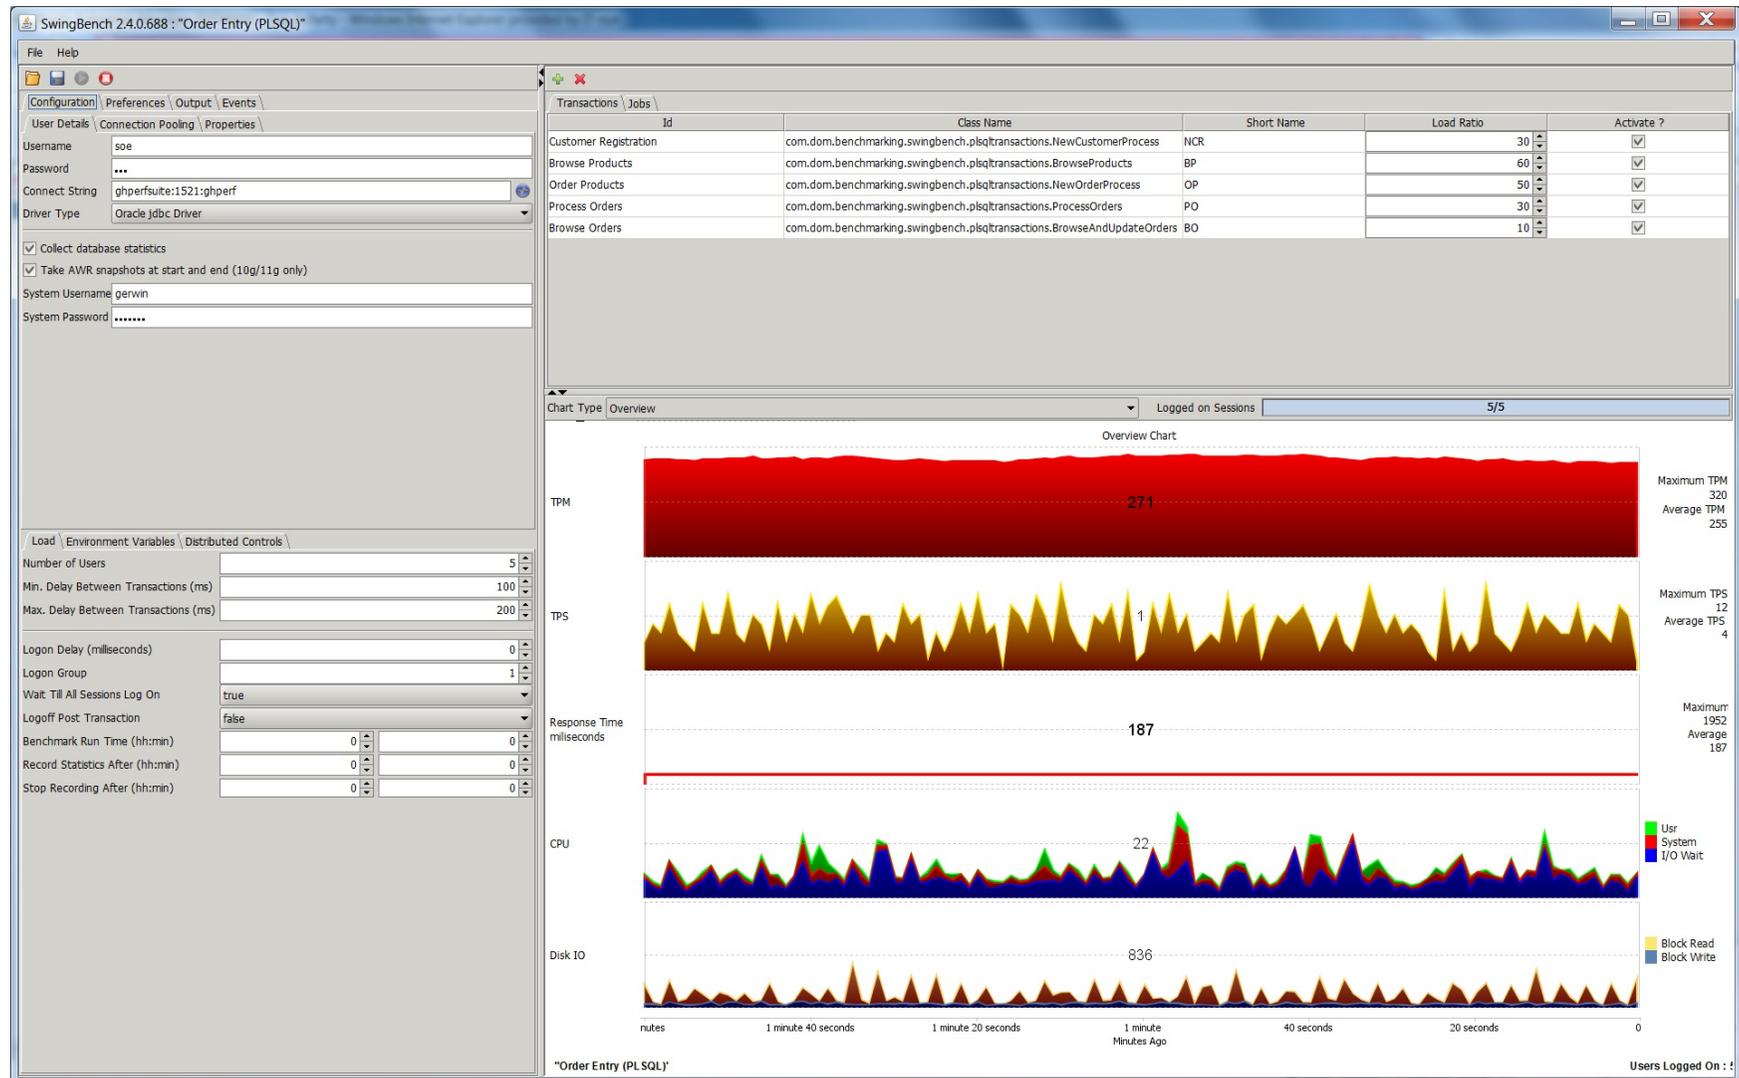
- Aggregate data on time stamps
- When using OEM this is already arranged
- Be aware of time differences
- Make sure that big workload changes are identified in your data
- Values examination
- Etc.

The Third / Fourth D

Data Mining – Explain

(factorial analyses or attribute importance)

Doing some tests in Swingbench...



Gathered data in the tests...

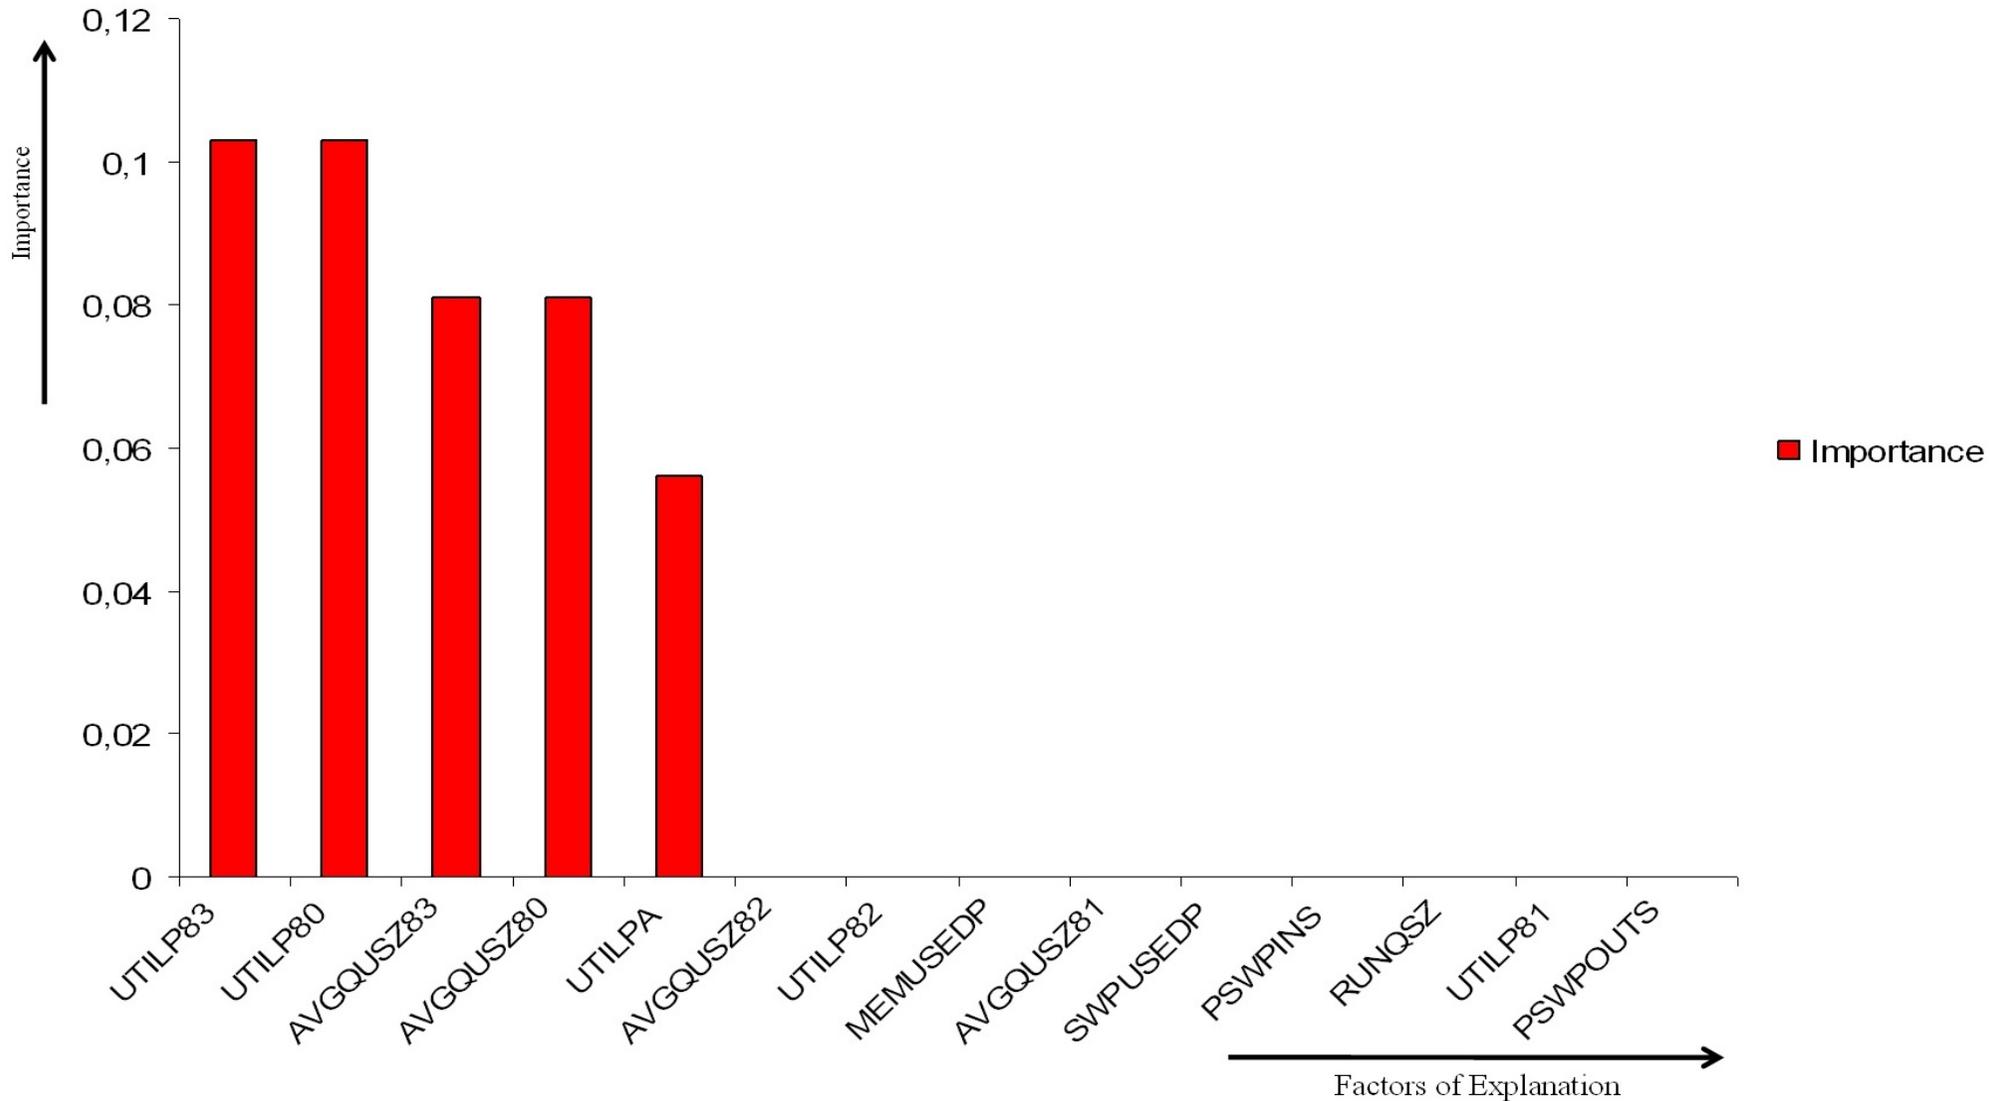


- Client data (Swingbench)
- SAR data on the linux virtual machine
- AWR data on the 11.2.0.3 database
- OEM 12C data

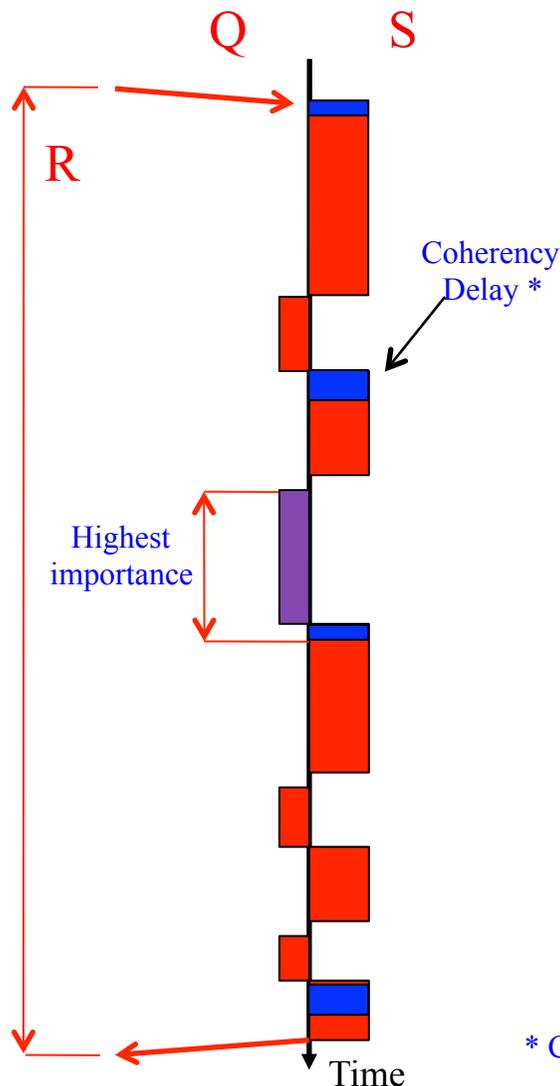
Do the data synchronization and...



Factorial Analysis on the SAR data...



The Importance...

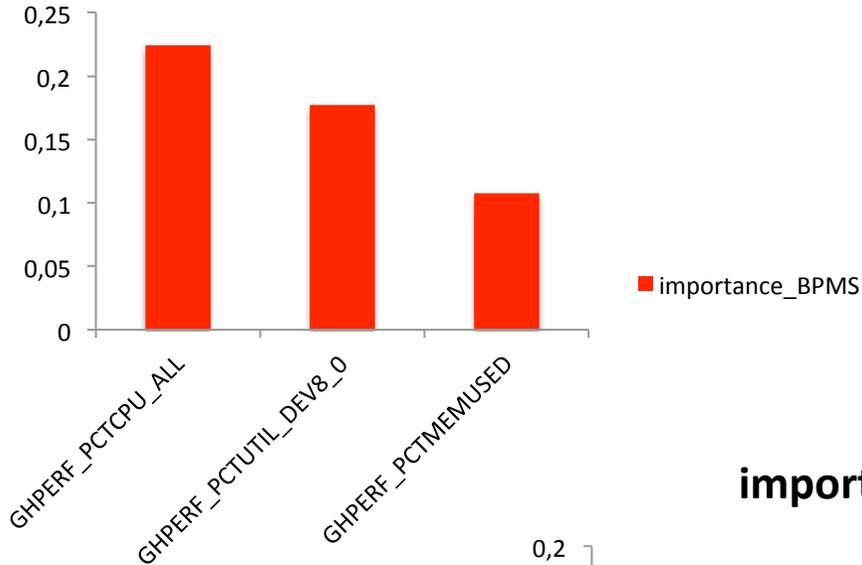


- Basically we search for the block with the biggest variance
- The factor with the biggest variance in the total of “Wait Time” + “Stretched Service Time”, will have the highest importance.
- So the higher the importance, the higher the impact on the total variance of R, and so worth investigation.

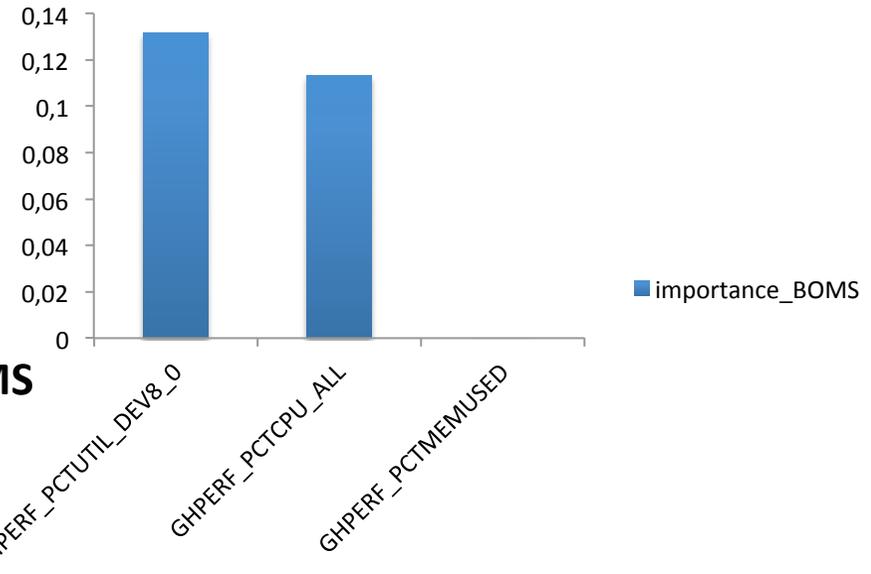
* Coherency Delay simplified

Different Transactions...

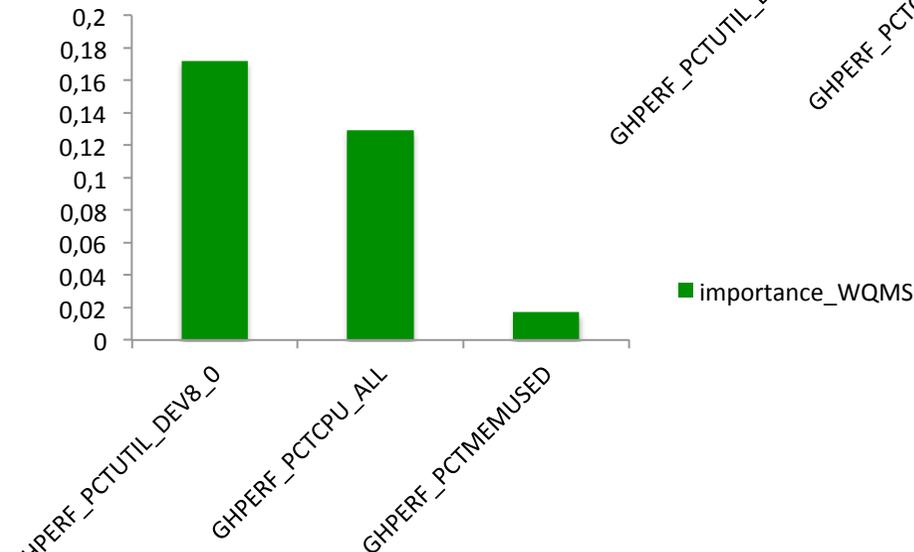
importance_BPMS



importance_BOMS



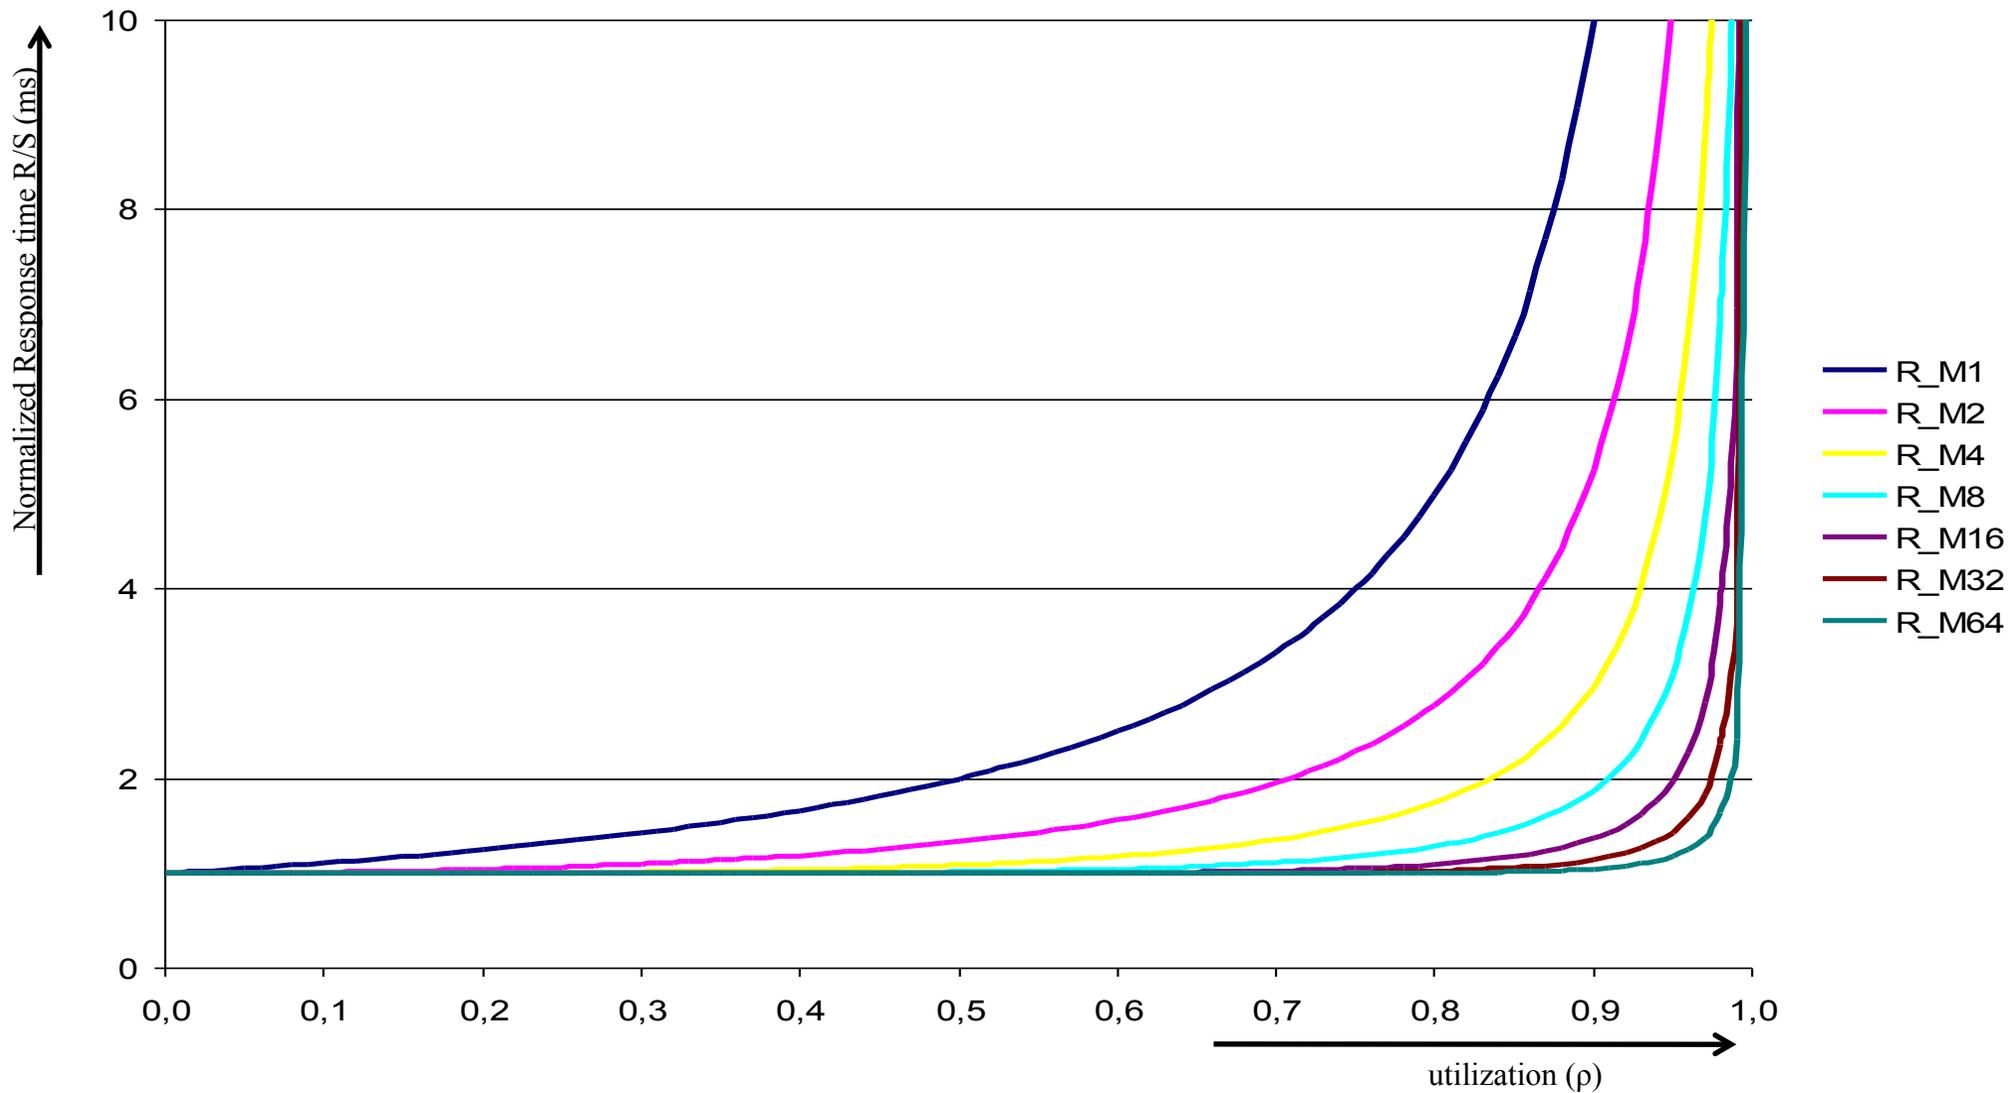
importance_WQMS



- “Browse Products”
- “Browse Orders”
- “Warehouse Query”

The Third D Data Modeling

Normalized R based on Erlang-C...



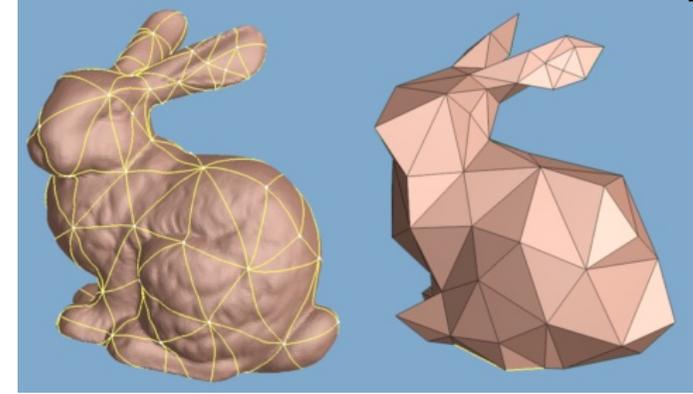
Erlang-C formula and R calculation...

$$C(m, \rho) = \frac{\frac{(m\rho)^m}{m!}}{(1-\rho) \sum_{n=0}^{m-1} \frac{(m\rho)^n}{n!} + \frac{(m\rho)^m}{m!}} \quad R = \frac{C(m, \rho)S}{m(1-\rho)} + S$$

Normalized R for M1 curve:

$$R = S \left(\frac{C(m, \rho)1}{m(1-\rho)} + 1 \right) \quad \xrightarrow[S]{/} \quad R_{M1} = 1 \left(\frac{C(1, \rho)1}{1(1-\rho)} + 1 \right)$$

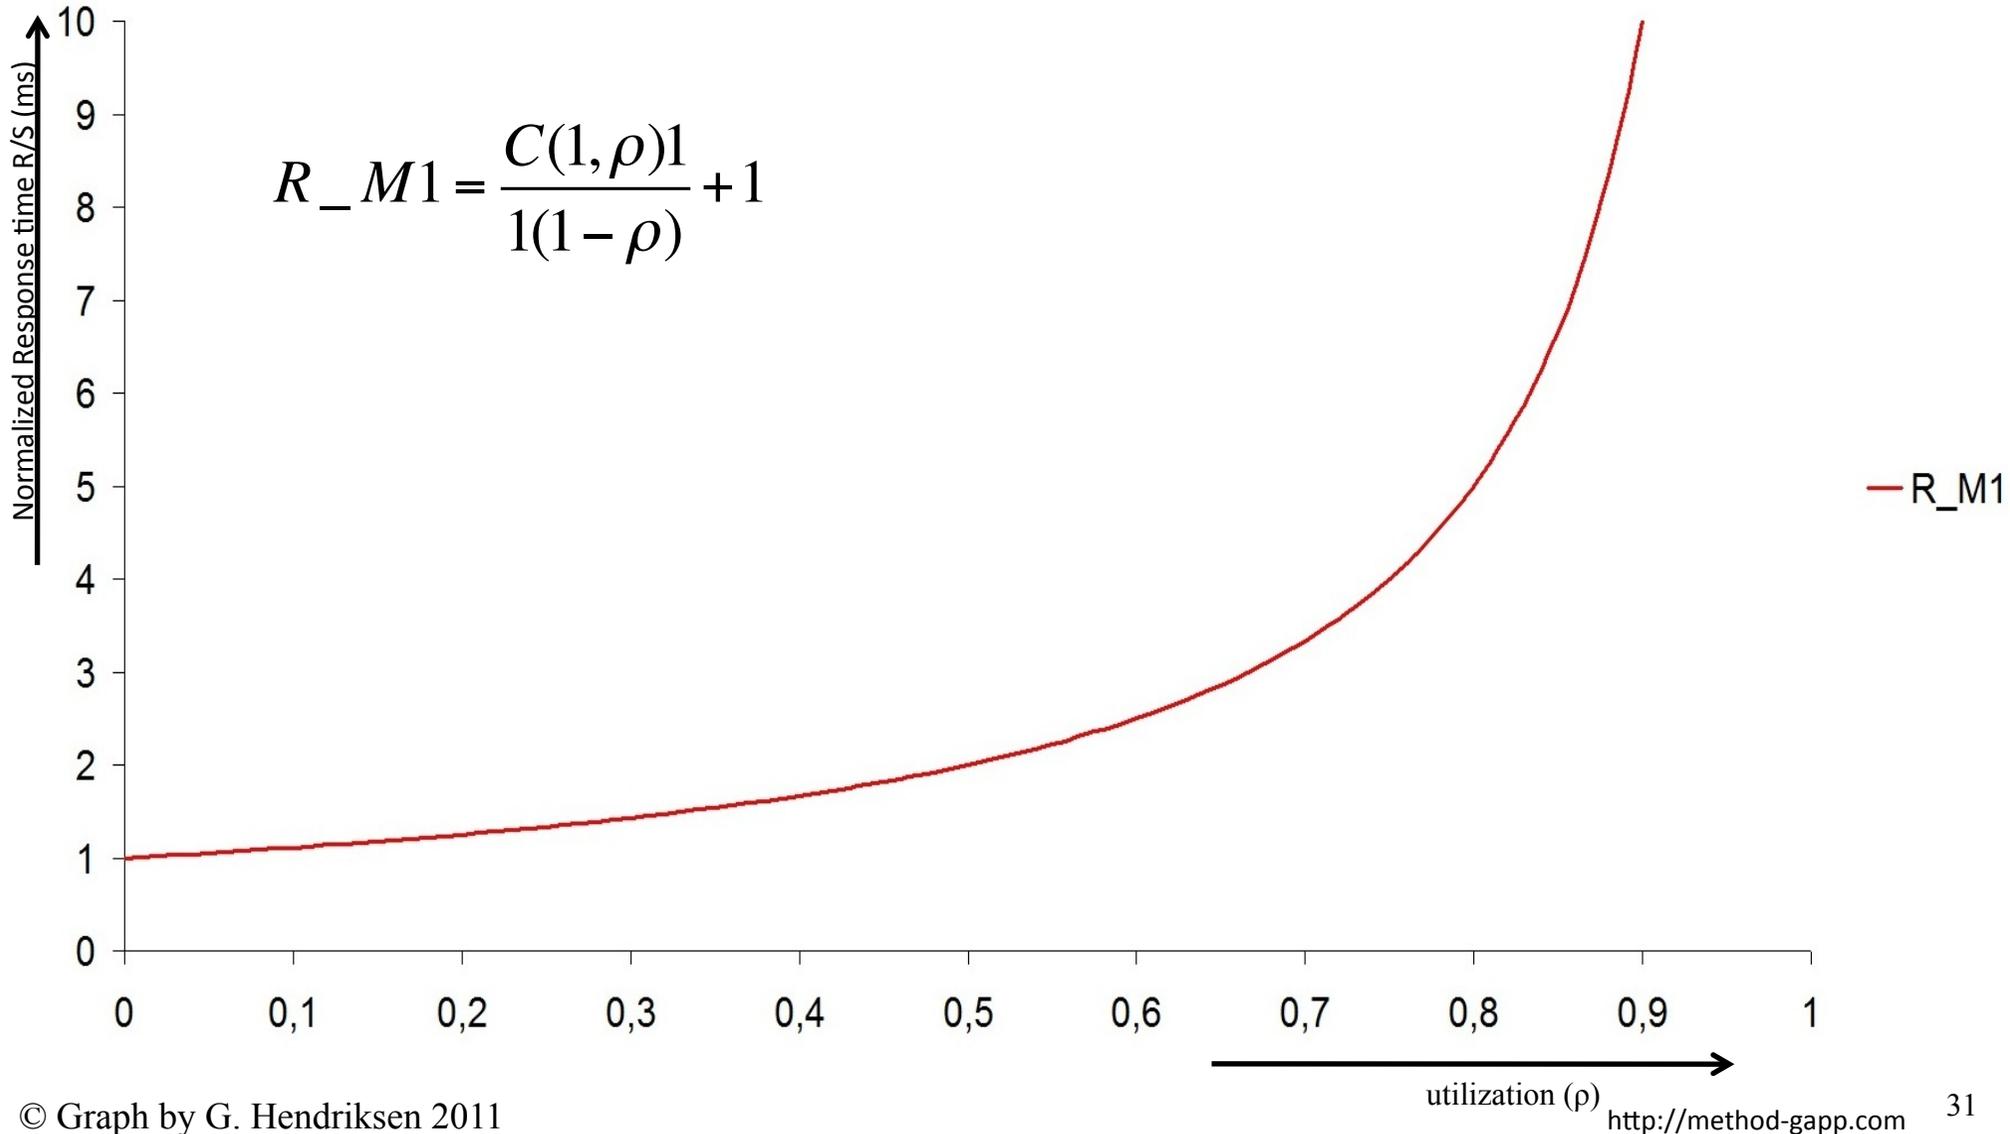
Why doing data modeling...



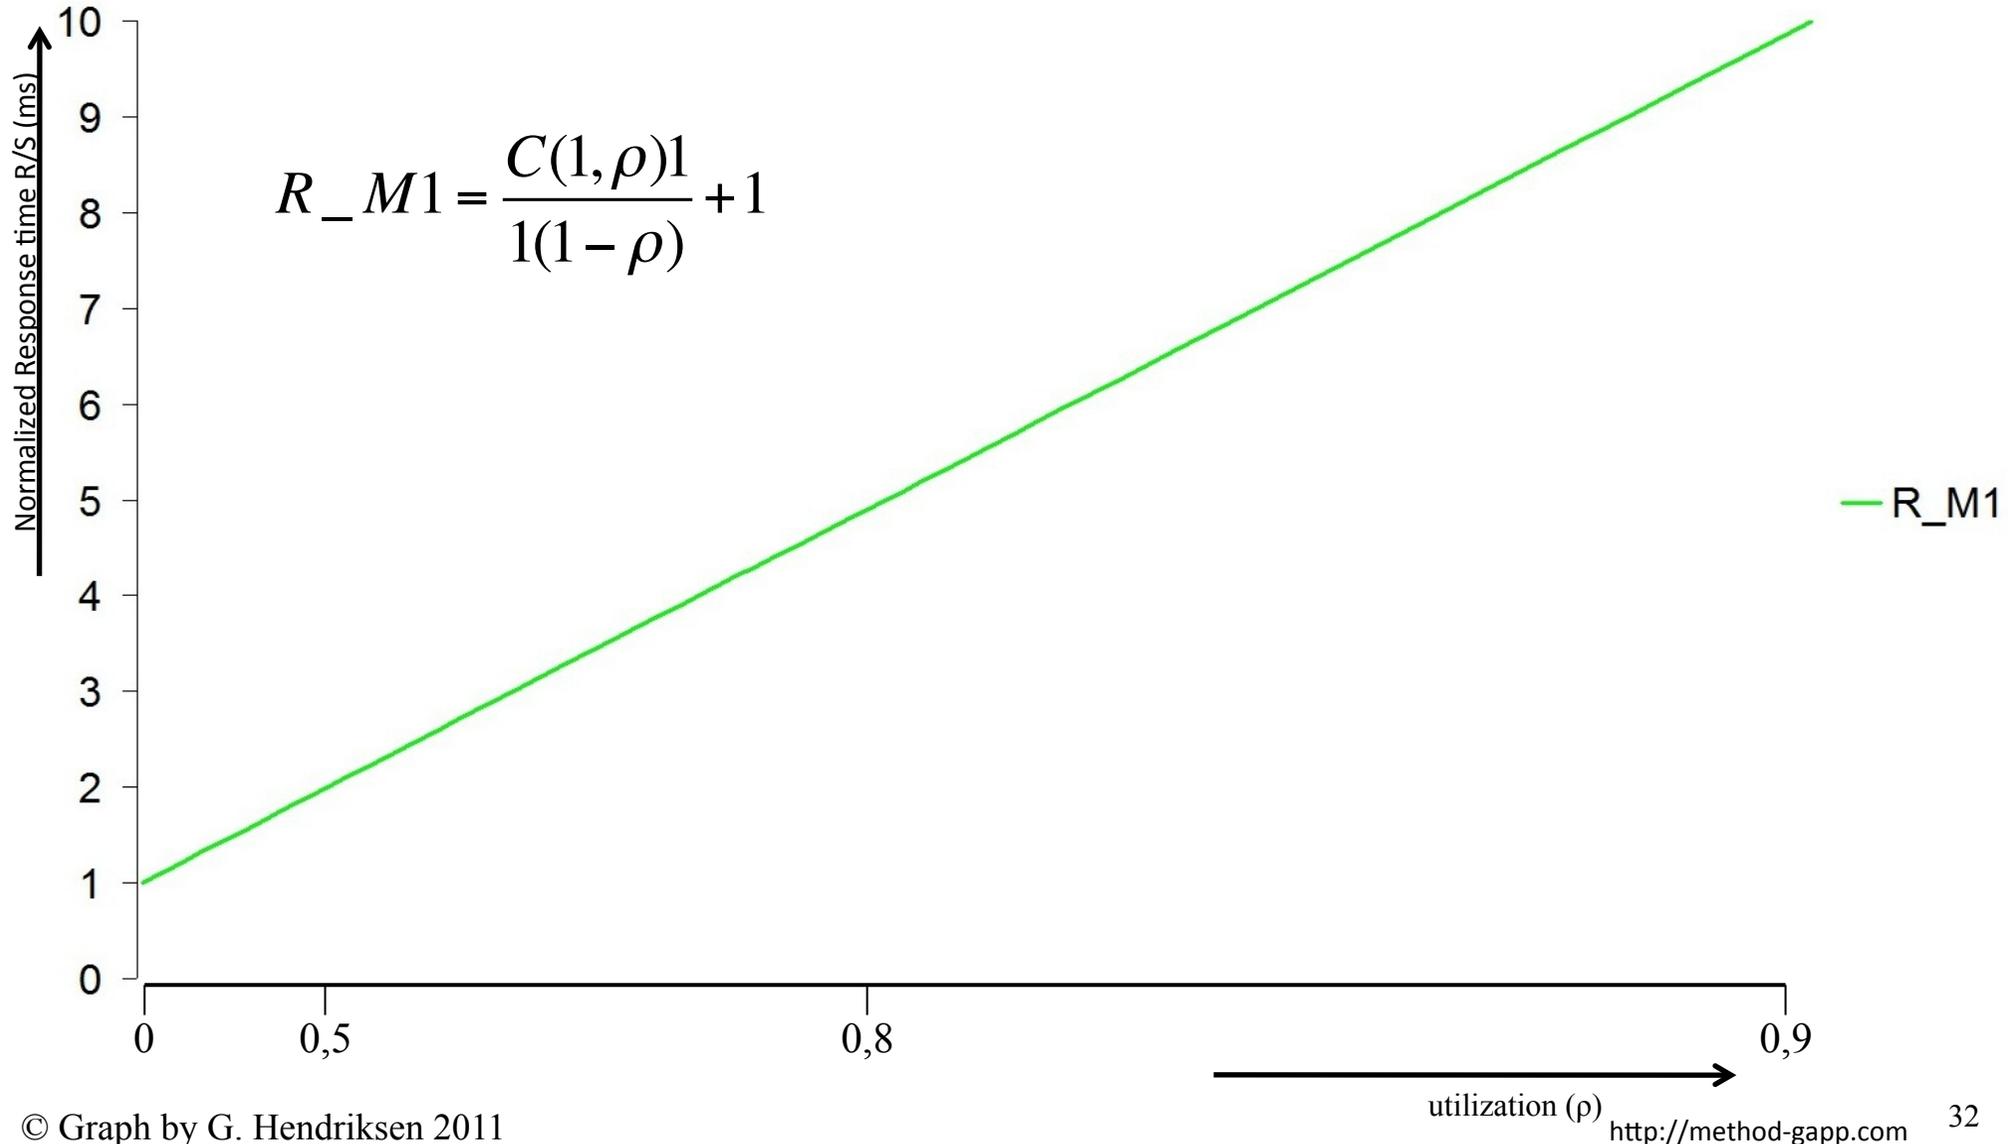
- The data modeling is done to be able to do:
 - Multi Linear Regression Data Mining
 - Able to create a linear function to describe the end-user response time (**R**)

M1 Curve - utilization linear scale...

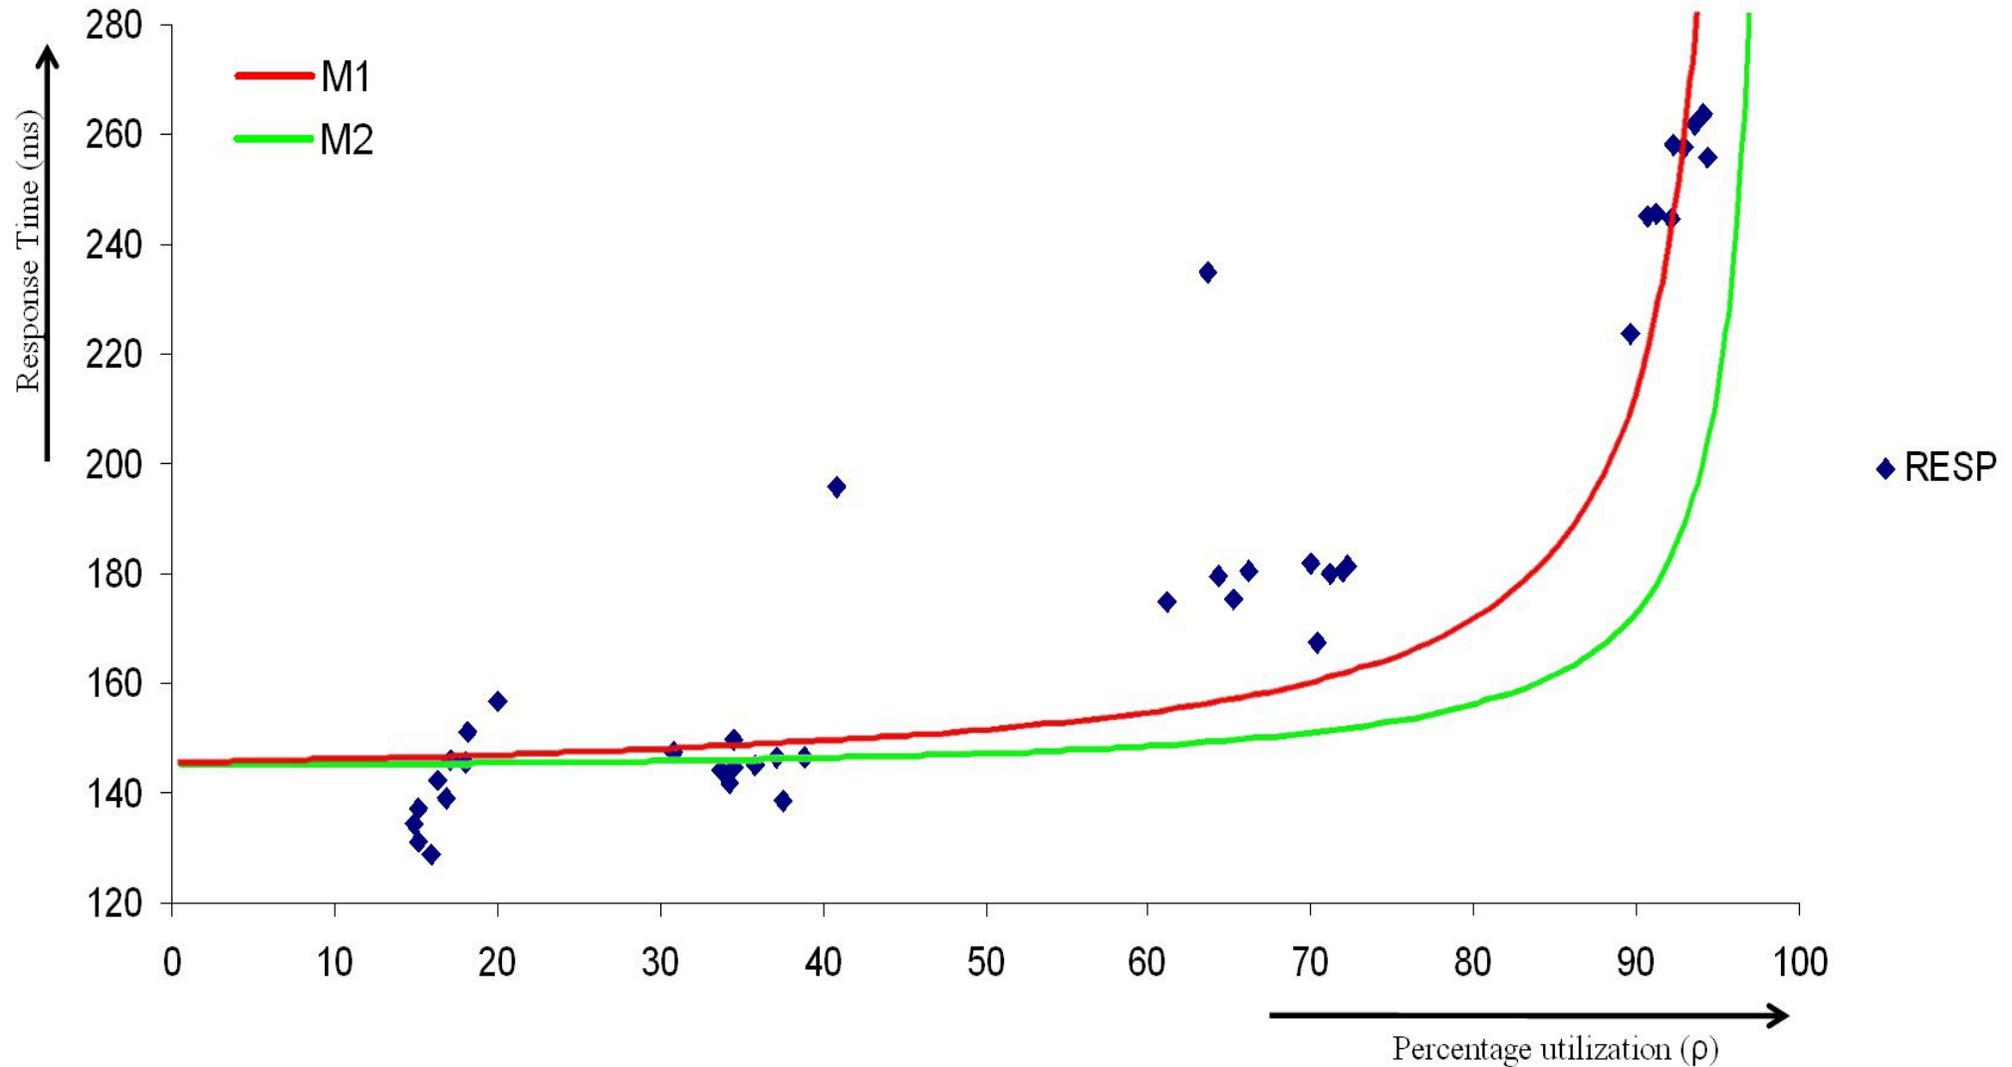
$$R_{M1} = \frac{C(1, \rho)1}{1(1 - \rho)} + 1$$



M1 Curve - utilization M1-curve scale...



Response time versus the I/O util...



The Fourth D

Data Mining – Model

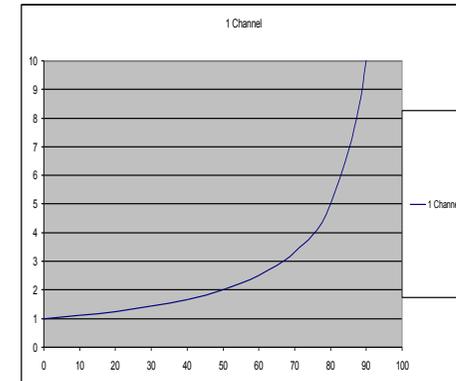
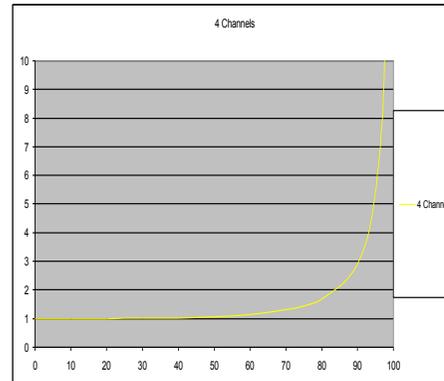
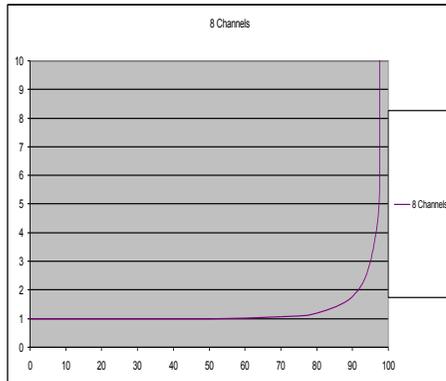
Data Mining and data modeling...



For Example:

Normalized R's for M curves !

$$R_{tot} = c_1 R_{1,n=8} + c_2 R_{2,n=4} + c_3 R_{3,n=1} + etc.$$



Gives the best fit, determined by data mining...

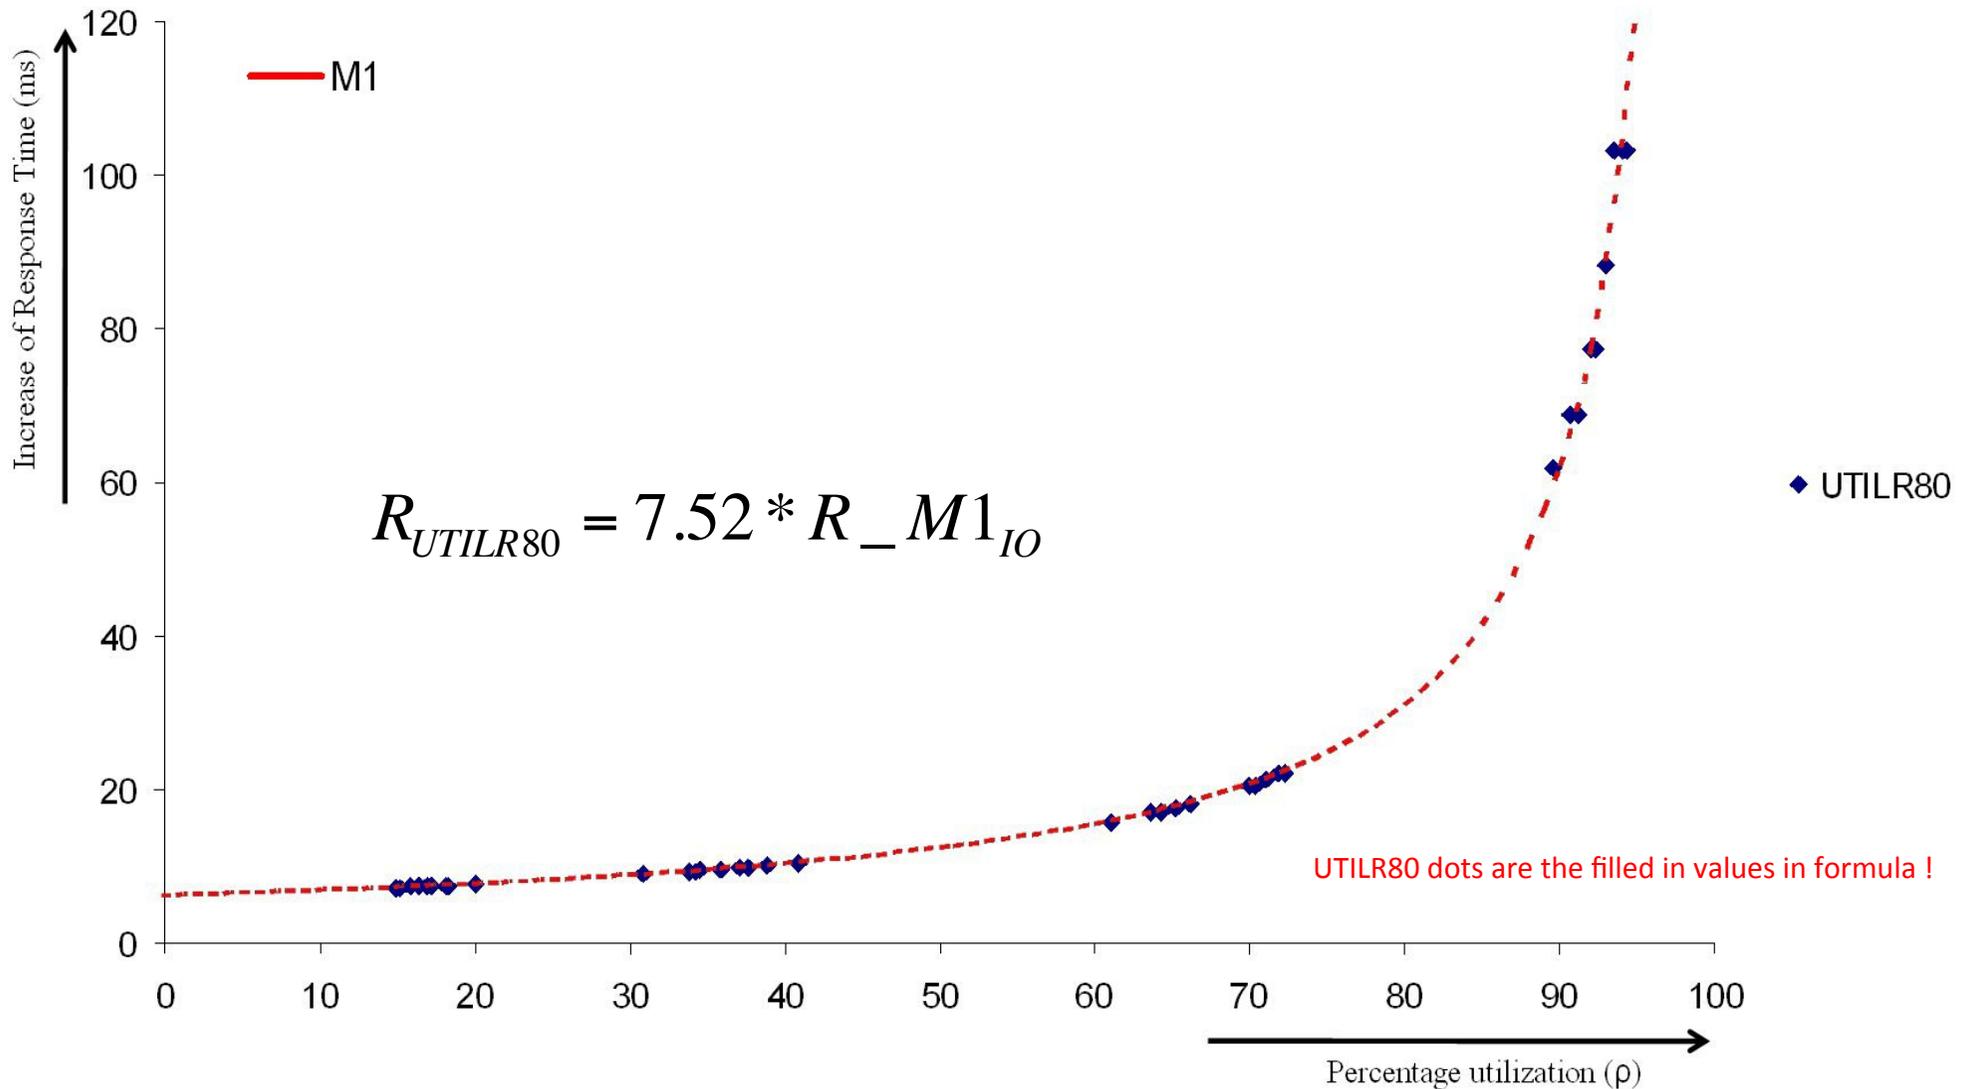
Creating the linear formula...

- Based on the coefficients of the calculated model and the “intercept row” of the model we can create the formula for the response time prediction. (below model only with I/O)

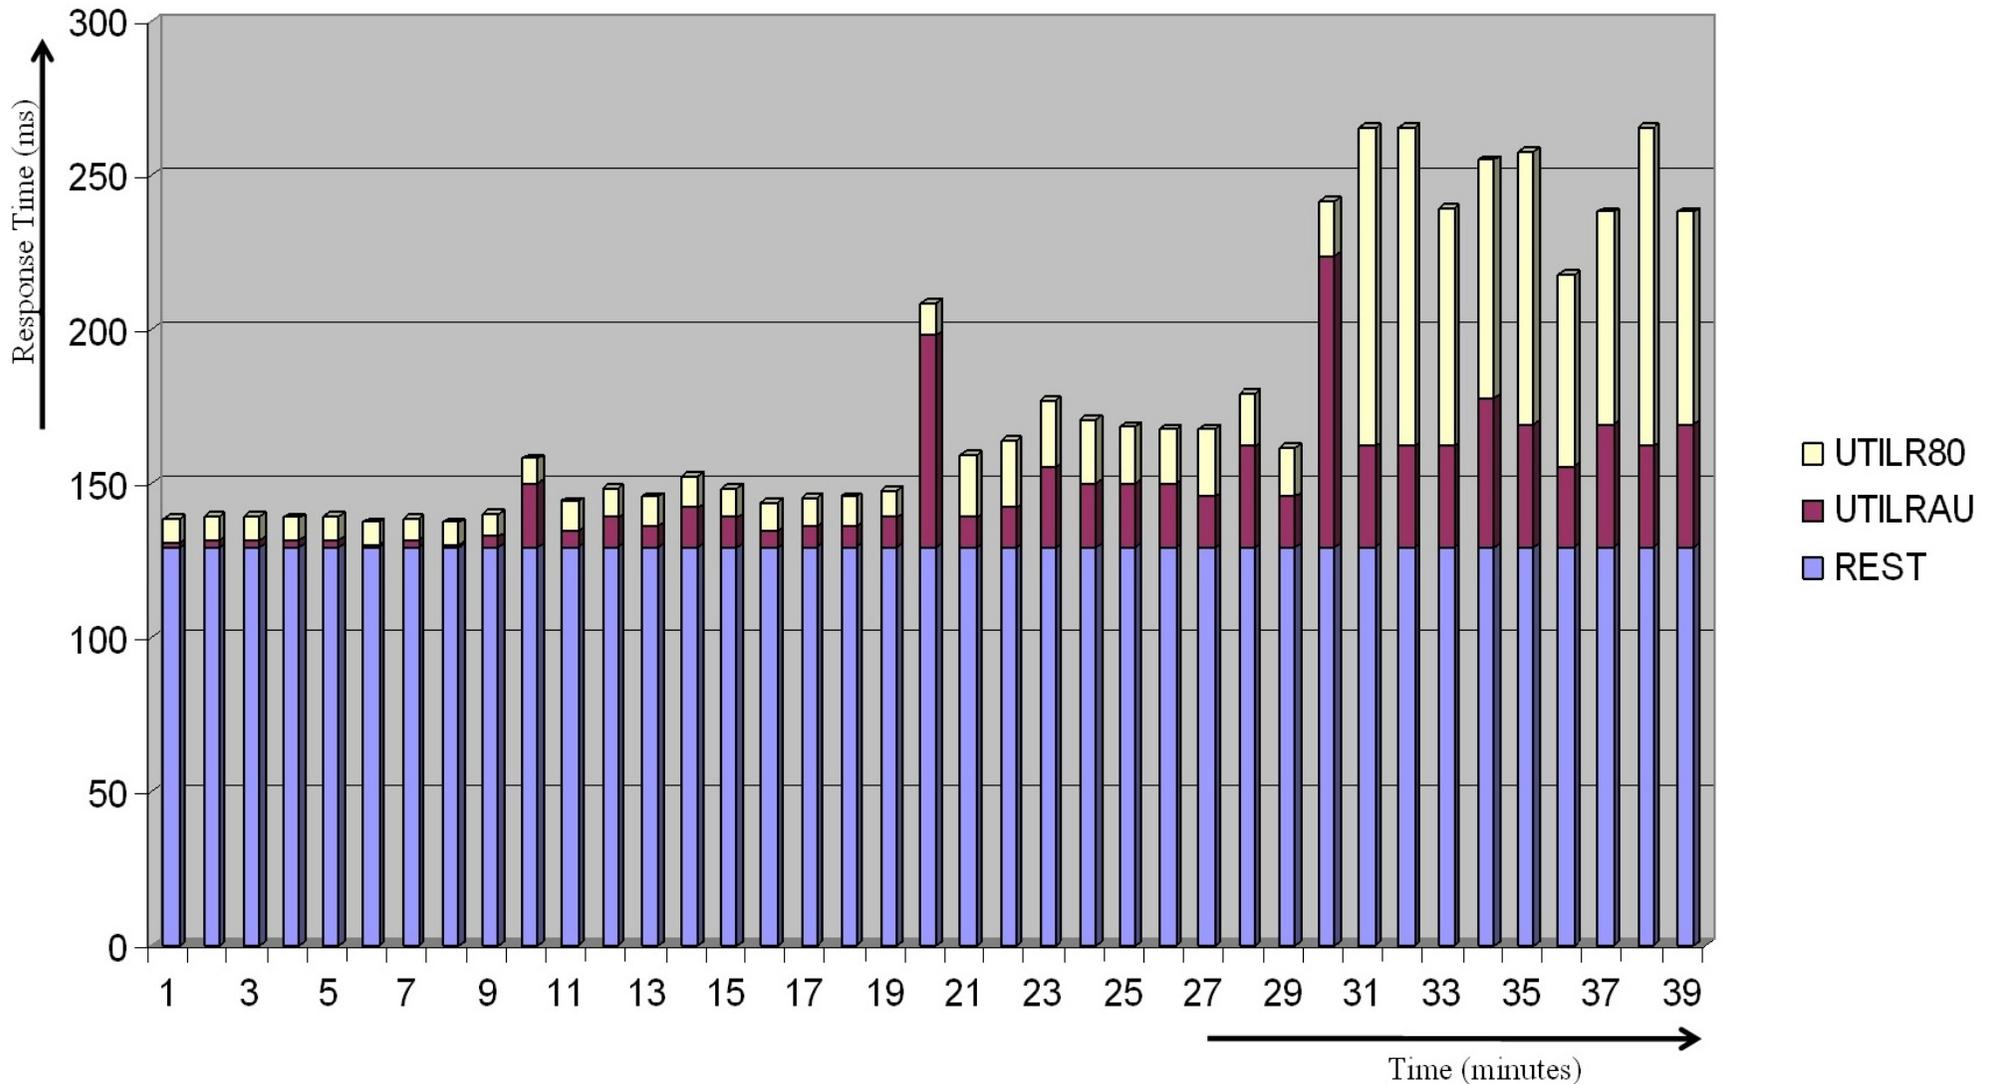
$$R_{RESP} \approx 7.52 \times M_{IO_{n=1}} + 145.41$$

$$R_{RESP} \approx R_{UTILR80} + R_{REST}$$

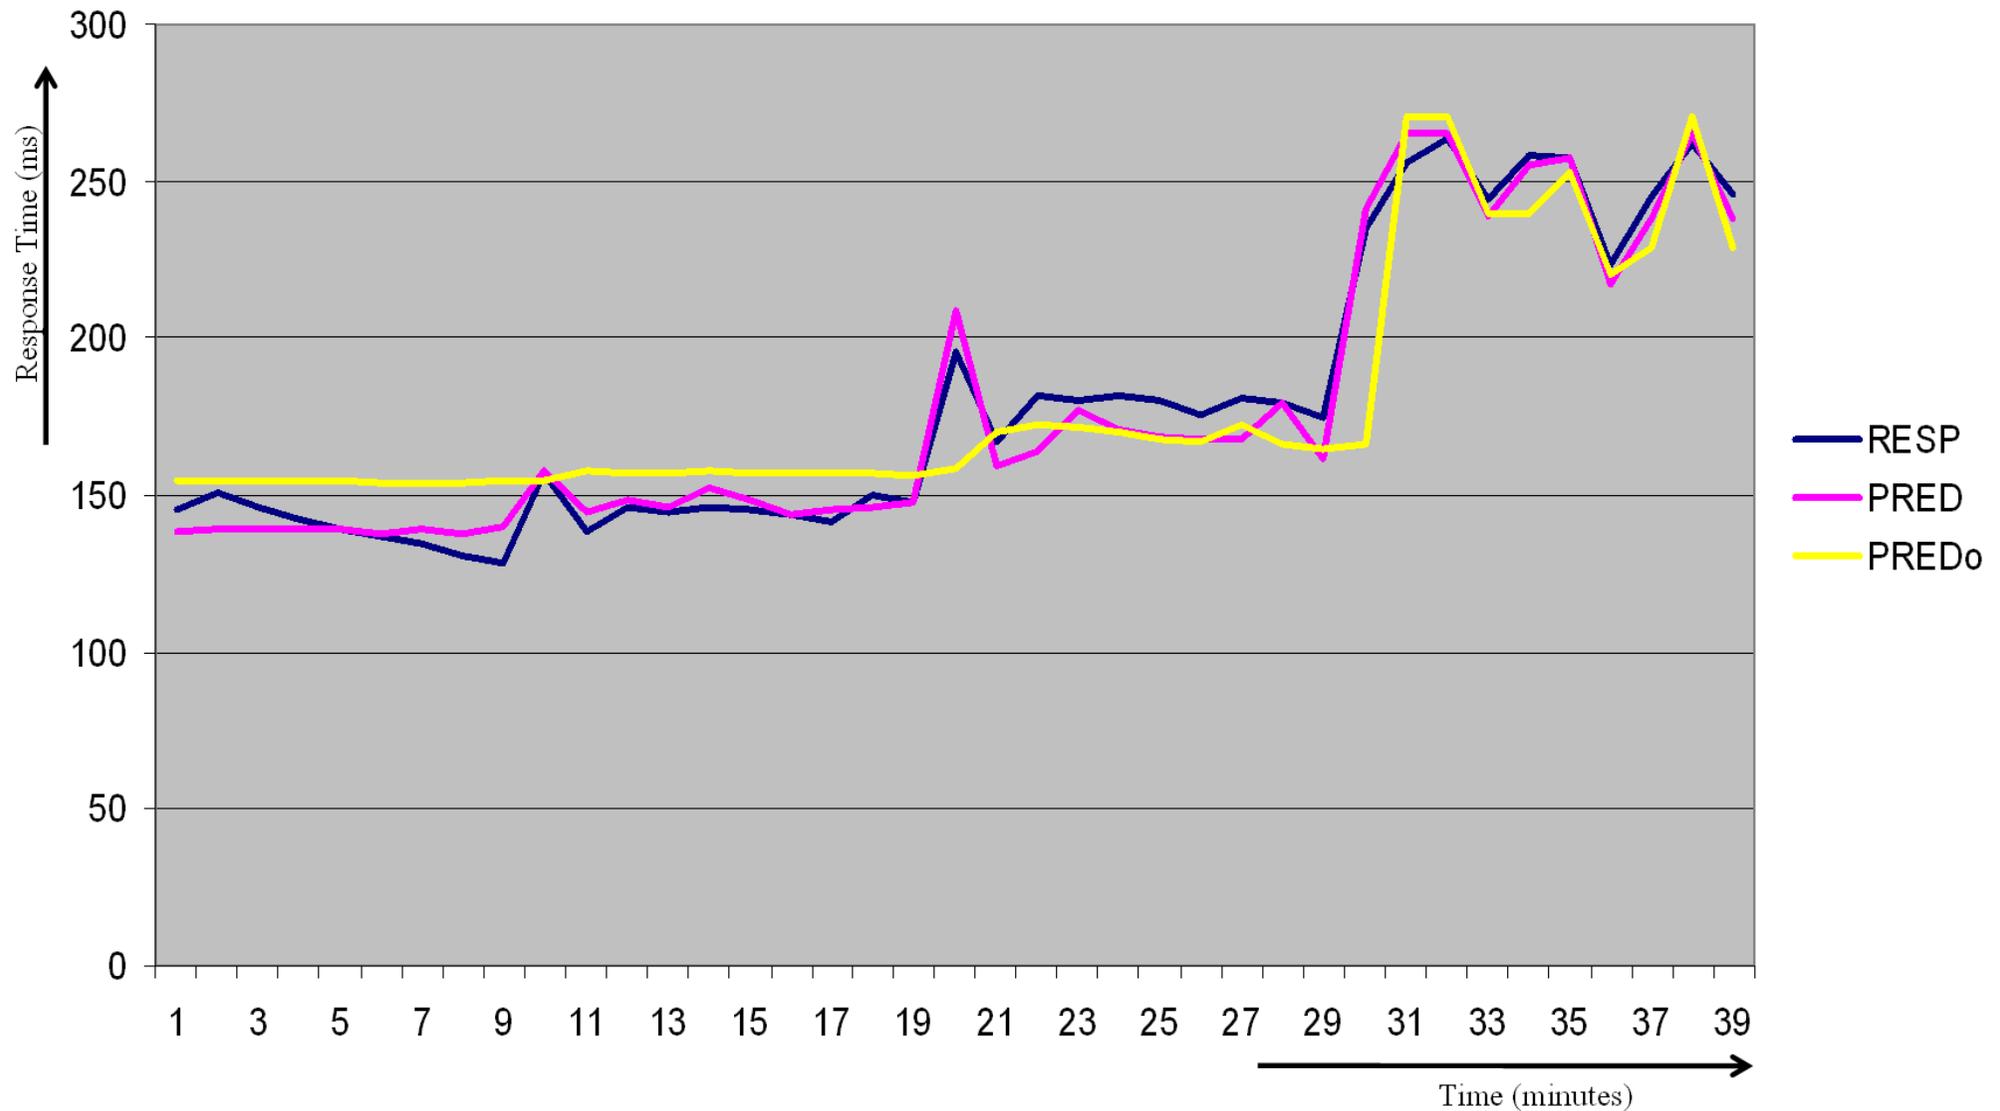
The modeled I/O utilization with M1...



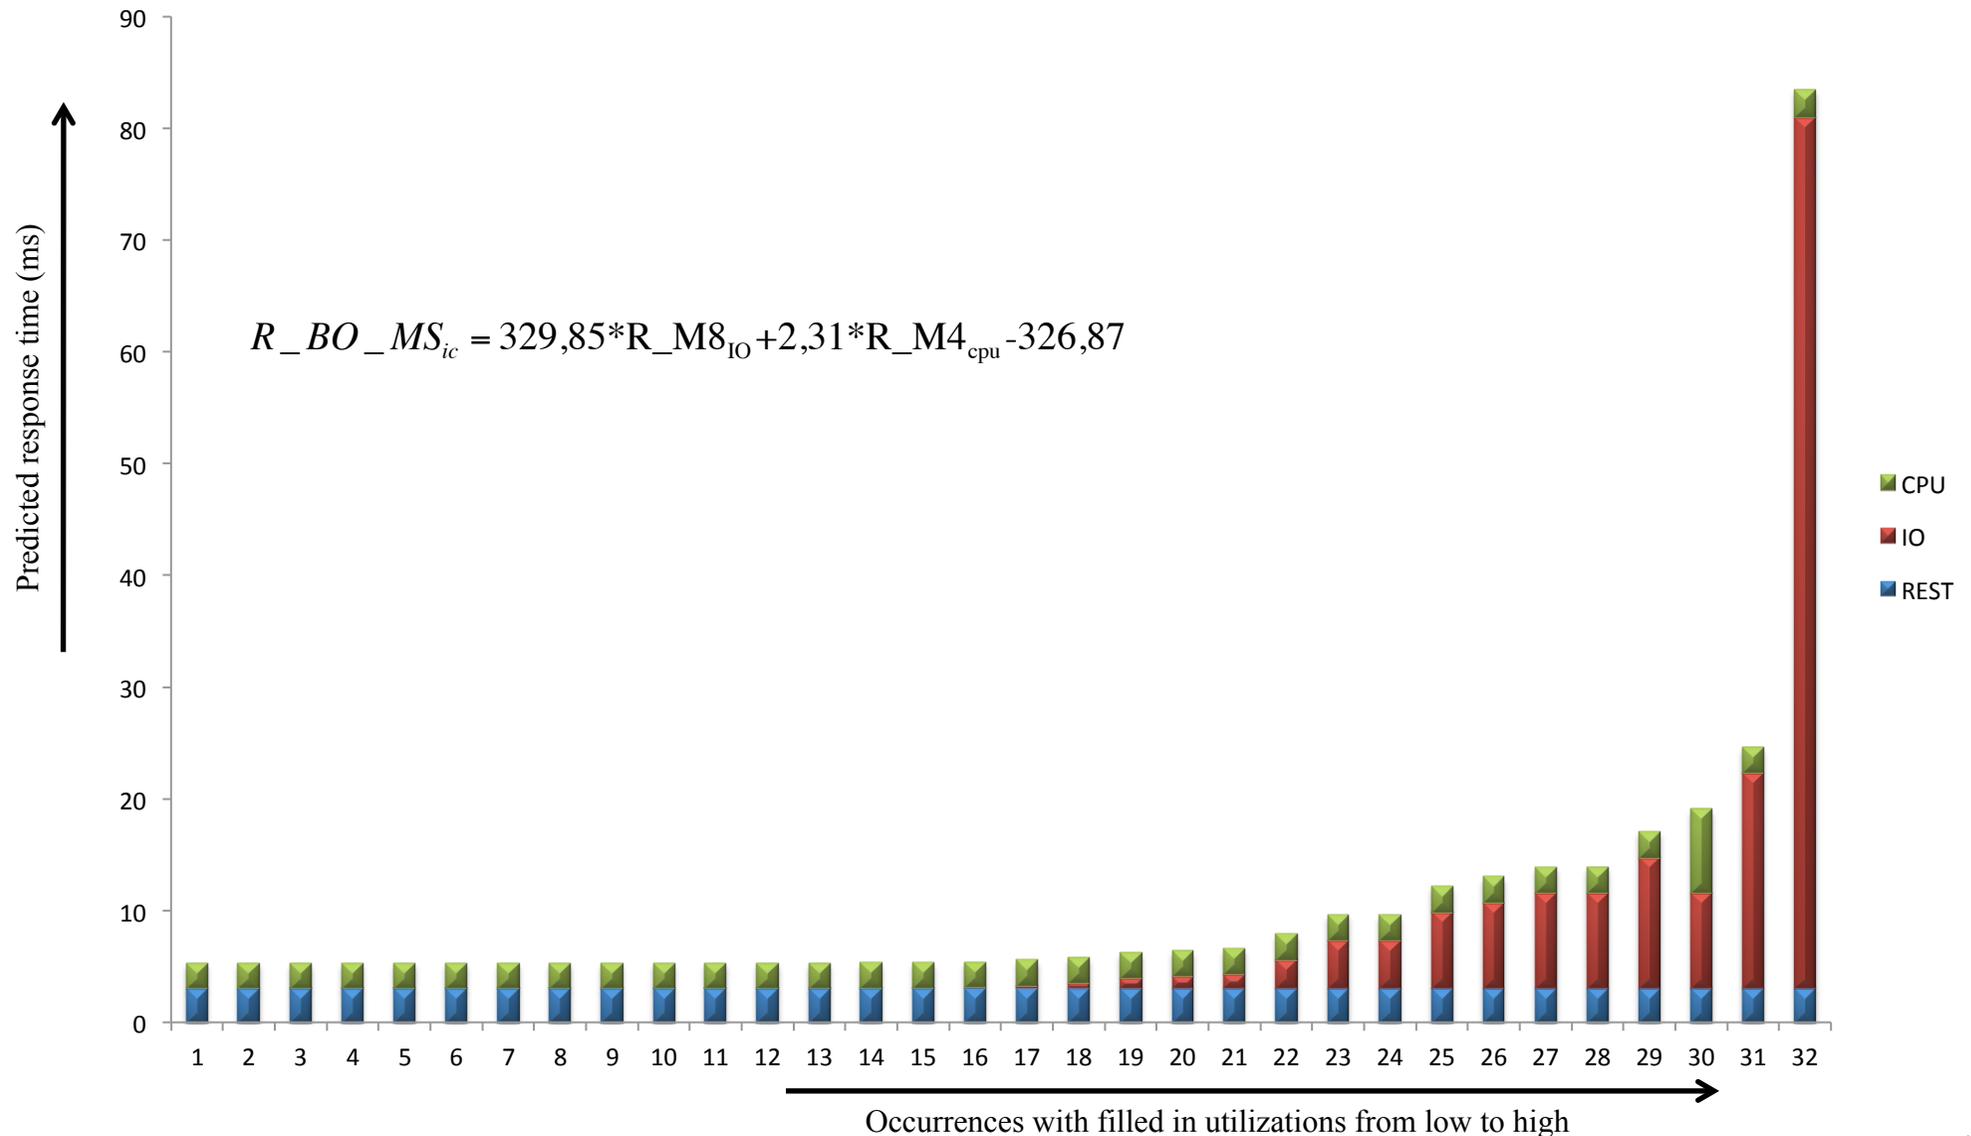
The model with the breakdown...



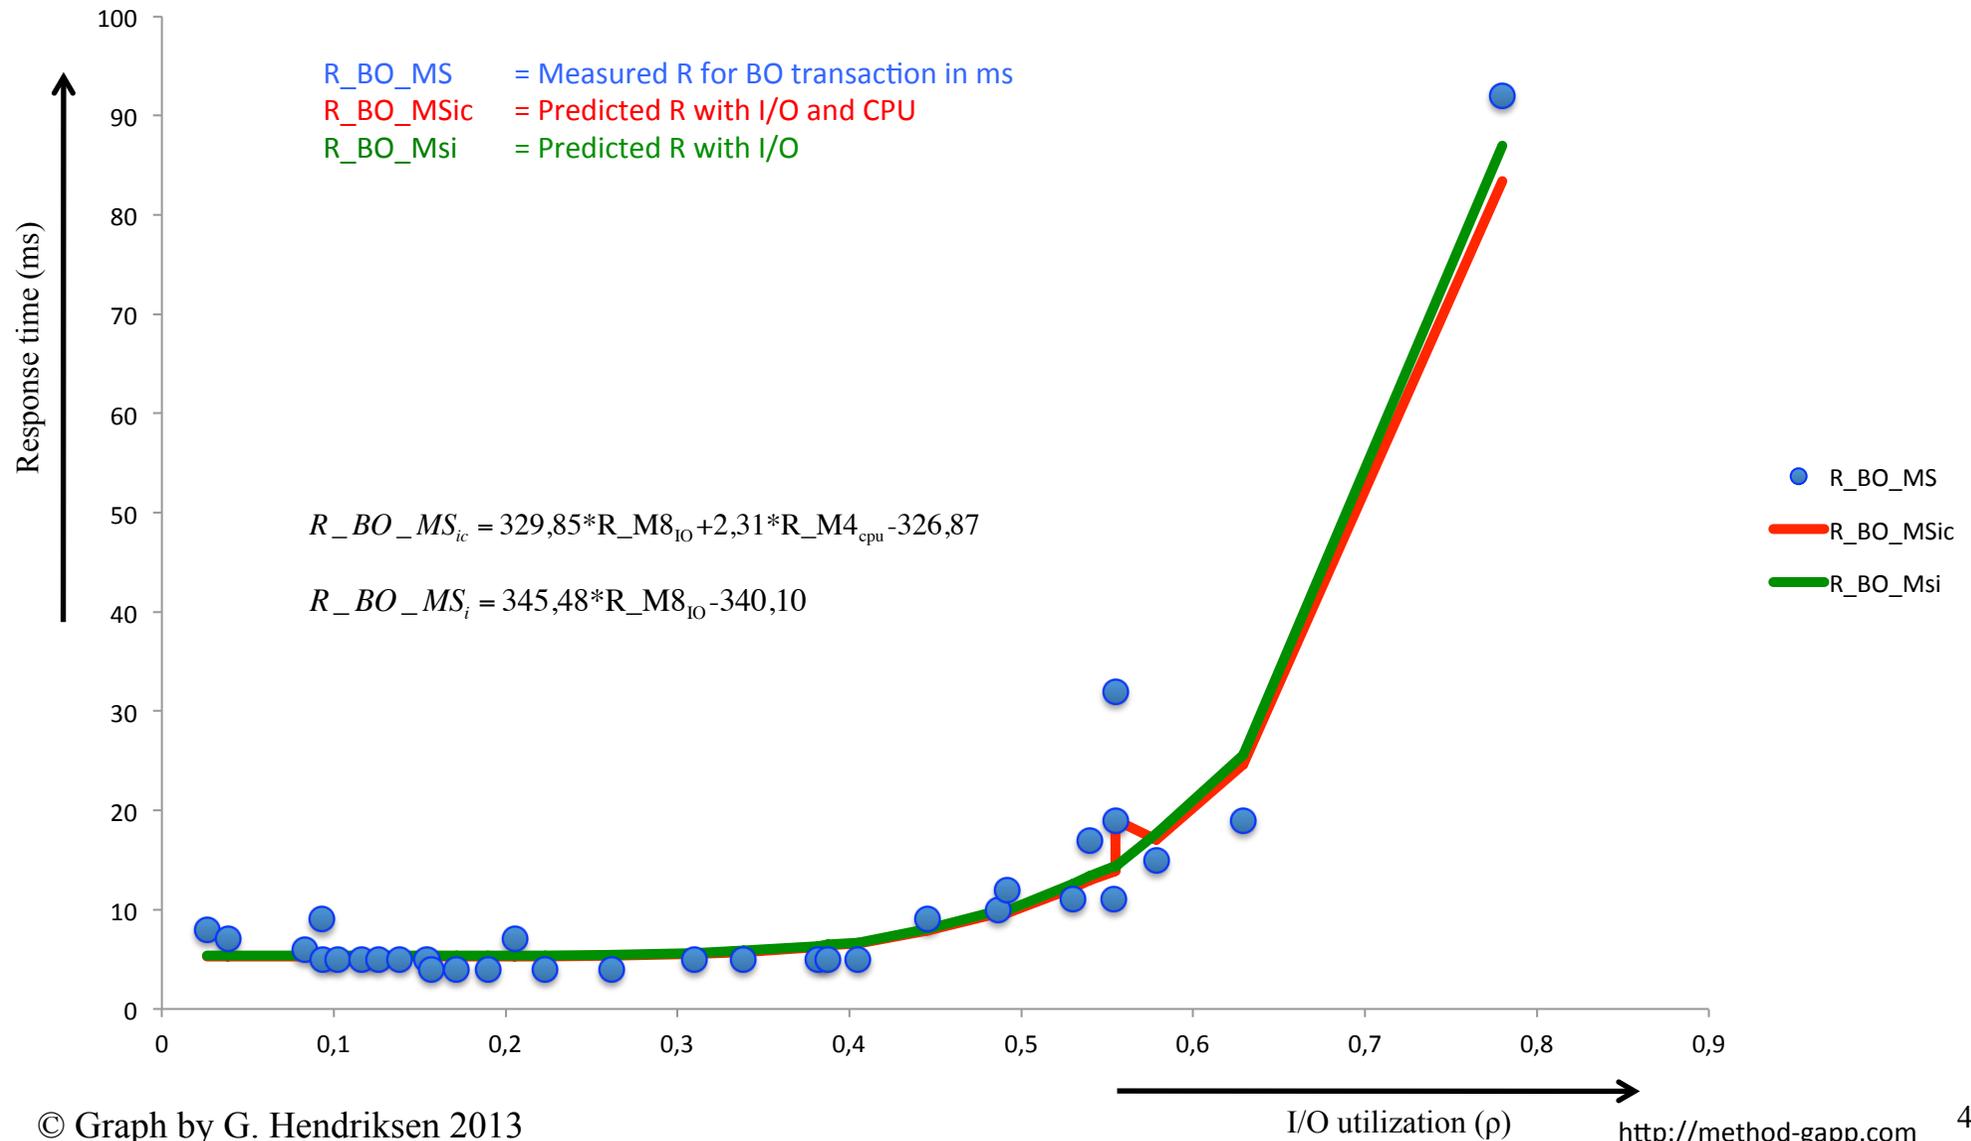
Models in compare with R measured...



Another model with the breakdown...



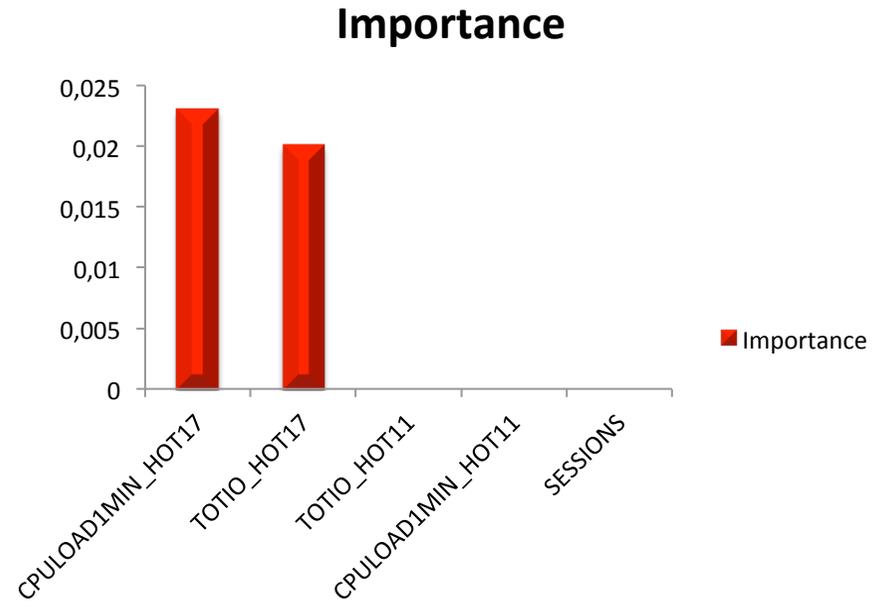
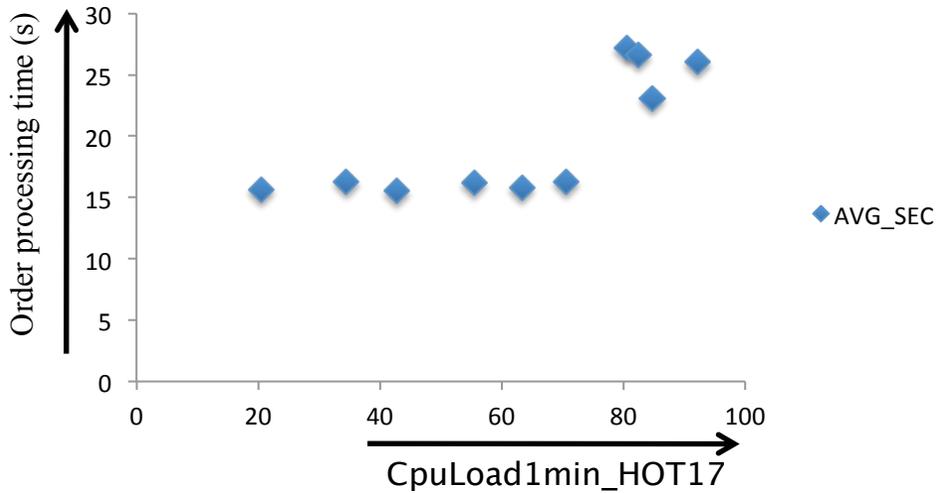
Models in compare with R measured...



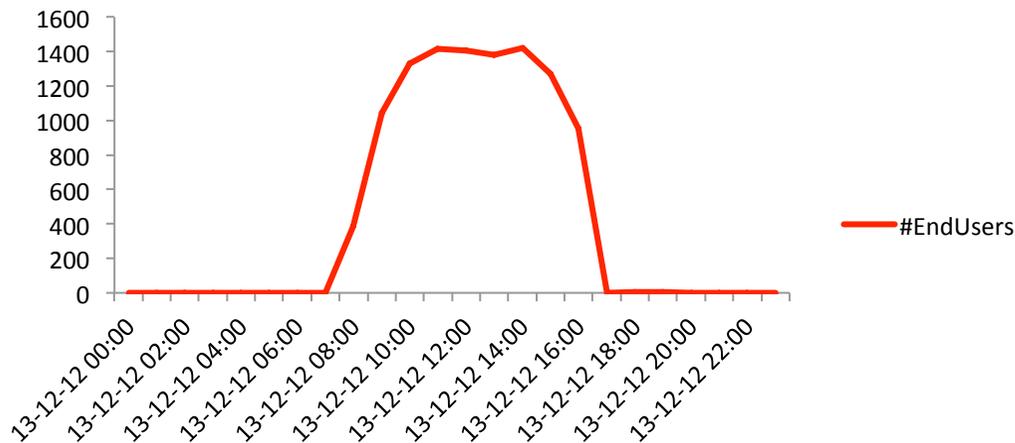
OEM Data usage

Fourth D Data Mining – Explain
using primary and secondary
component data

Data from OEM 12C...



#EndUsers



```
select * from mgmt$metric_hourly
where metric_column='CurrentTasks'
and target_name like 'PSCcObjMgr_HOT21'
order by target_name, rollup_timestamp asc
```

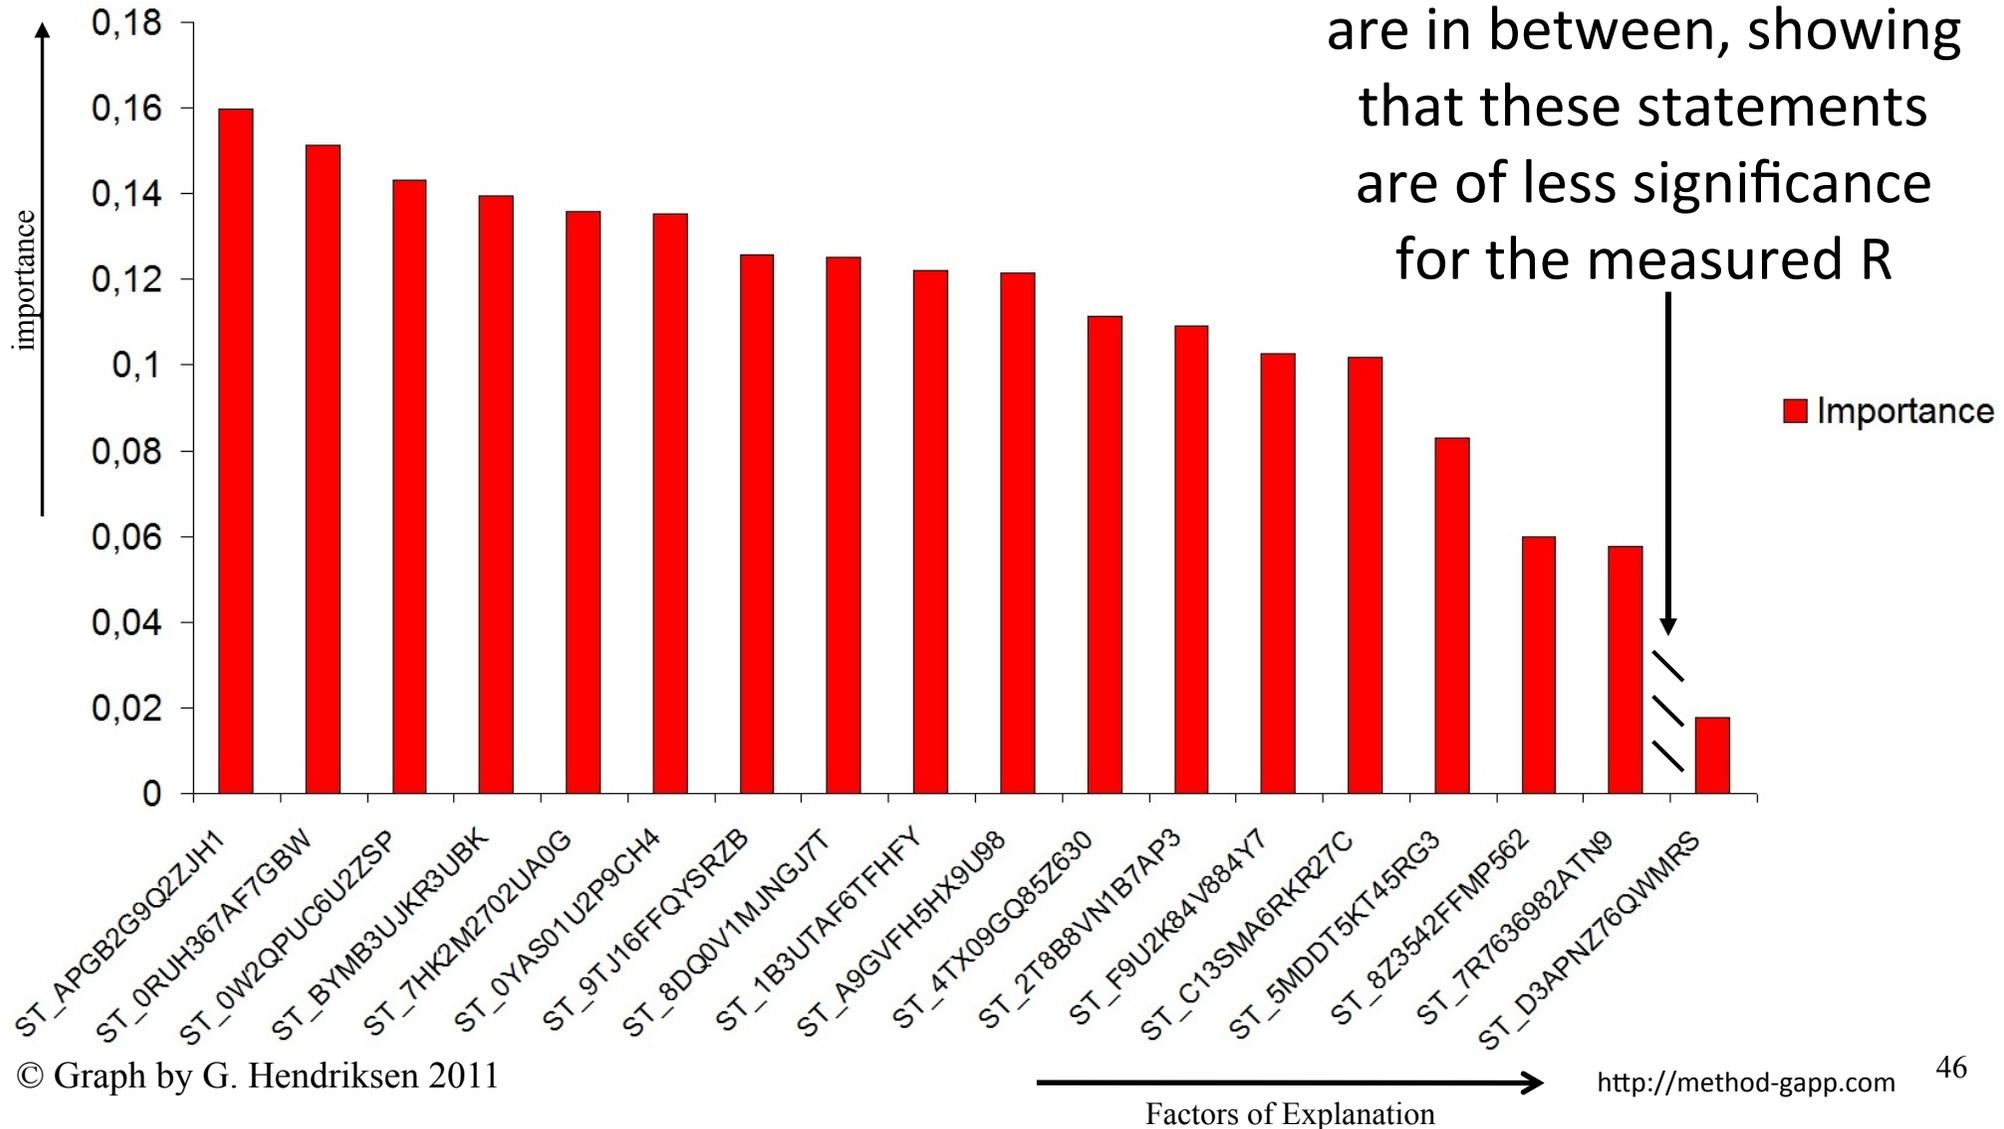
AWR Data usage

Fourth D Data Mining – Explain
using secondary component data

How to retrieve the used AWR data...

- To get the AWR SQL data response times:
 - $\text{elapsed_time_delta} / \text{executions_delta}$
 - From DBA_HIST_SQLSTAT
- To get the AWR events data:
 - TIME_WAITED_MICRO
 - From DBA_HIST_SYSTEM_EVENT
 - Calculate the delta between snapshots

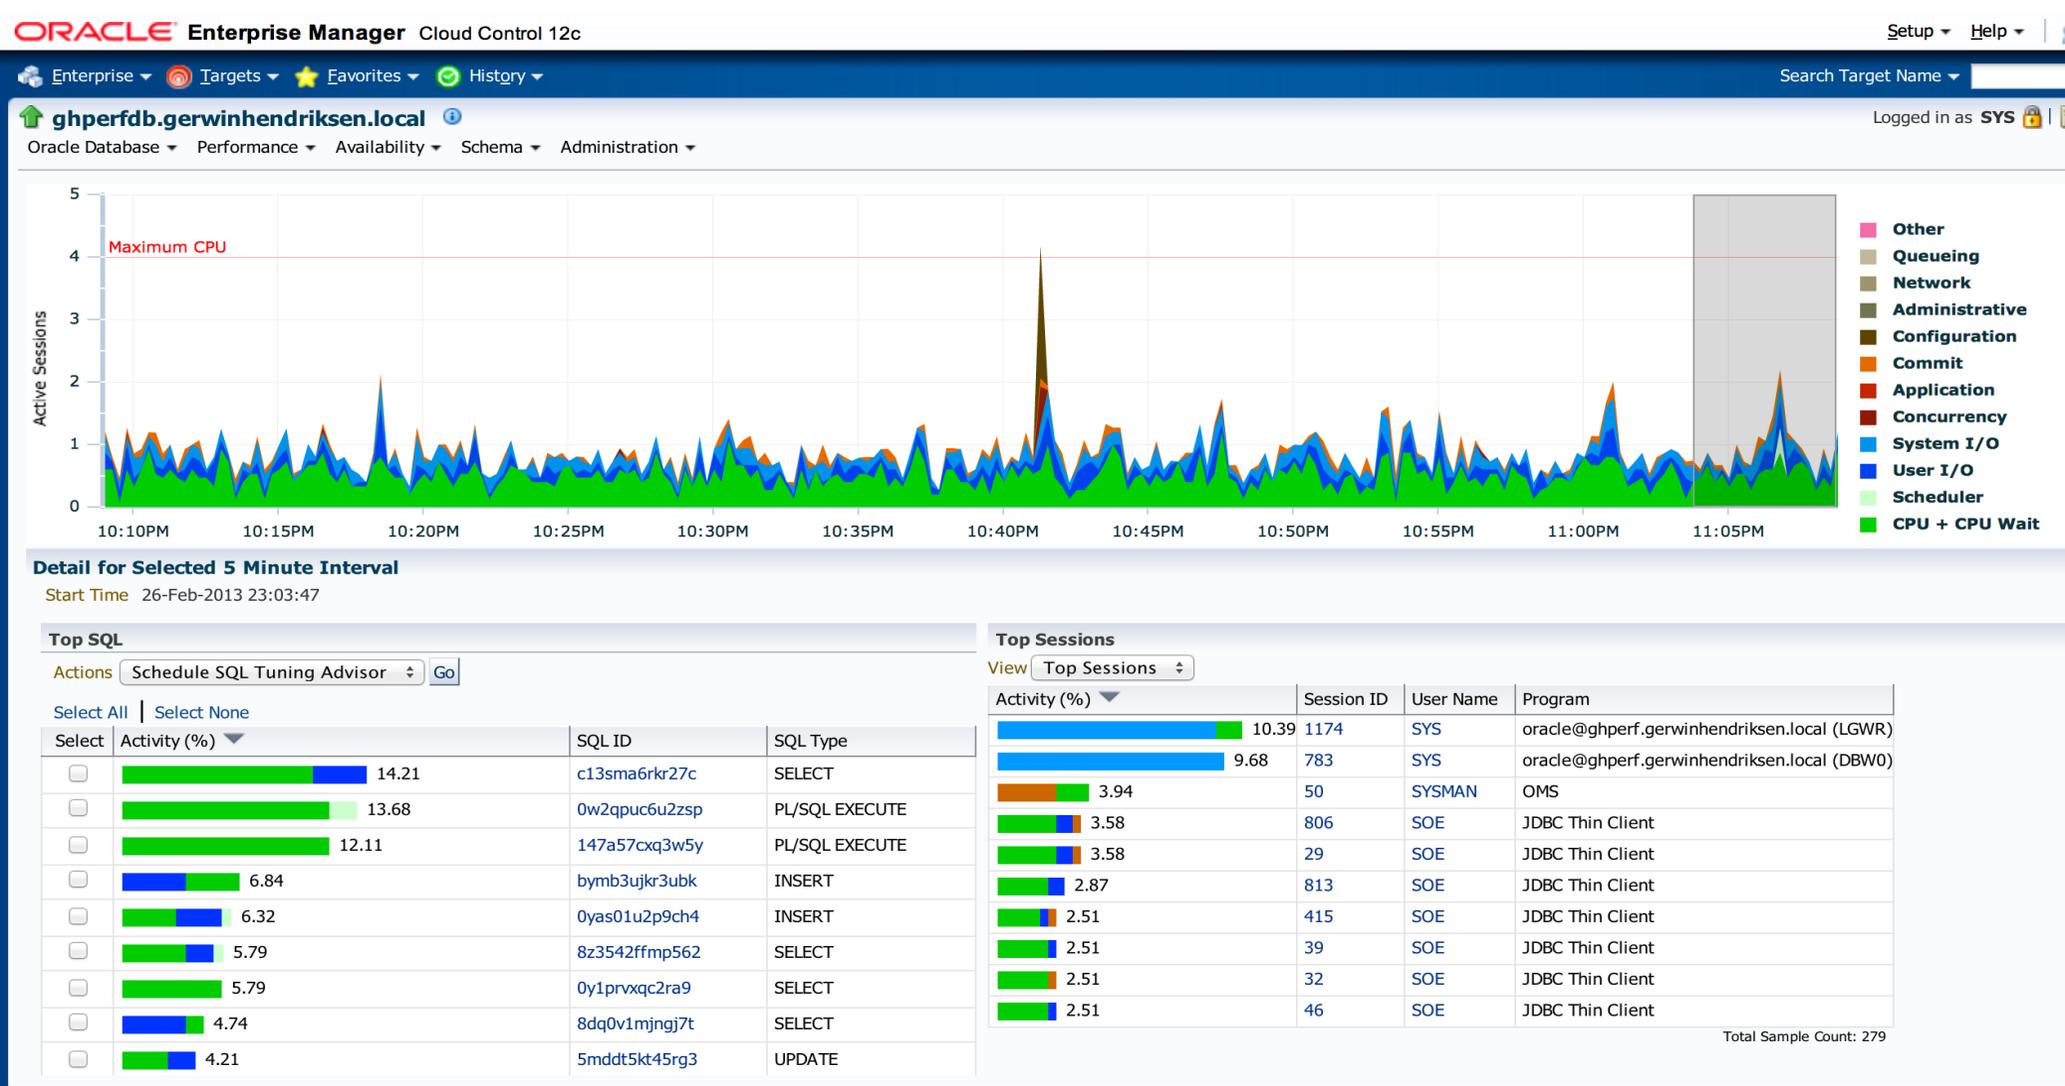
SQL Response Time Data from AWR...



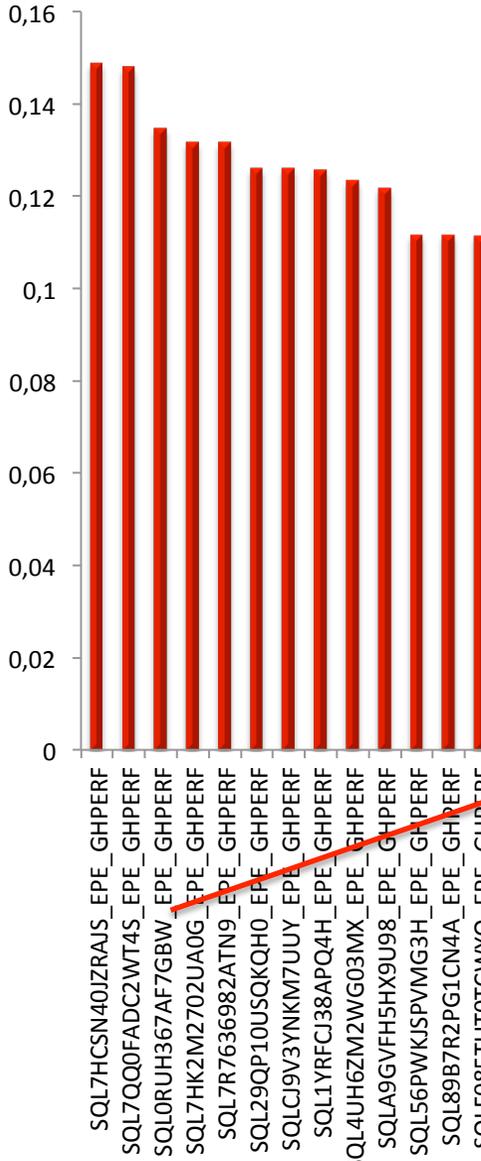
Involved important Statements...

- **'apgb2g9q2zjh1'**
 - BEGIN :1 := orderentry.browseandupdateorders(:2 ,:3 ,:4); END;
- **'0ruh367af7gbw'**
 - SELECT ORDER_ID, ORDER_MODE, CUSTOMER_ID, ORDER_STATUS, ORDER_TOTAL, SALES_REP_ID, PROMOTION_ID FROM ORDERS WHERE CUSTOMER_ID = :B2 AND ROWNUM < :B1
- **'0w2qpuc6u2zsp'**
 - BEGIN :1 := orderentry.neworder(:2 ,:3 ,:4); END;
- **'bymb3ujkr3ubk'**
 - INSERT INTO ORDERS(ORDER_ID, ORDER_DATE, CUSTOMER_ID, WAREHOUSE_ID) VALUES (ORDERS_SEQ.NEXTVAL + :B3 , SYSTIMESTAMP , :B2 , :B1) RETURNING ORDER_ID INTO :O0
- **'7hk2m2702ua0g'**
 - WITH NEED_TO_PROCESS AS (SELECT ORDER_ID, CUSTOMER_ID FROM ORDERS WHERE ORDER_STATUS <= 4 AND WAREHOUSE_ID = :B1 AND ROWNUM < 10) ...
- **'0yas01u2p9ch4'**
 - INSERT INTO ORDER_ITEMS(ORDER_ID, LINE_ITEM_ID, PRODUCT_ID, UNIT_PRICE, QUANTITY) VALUES (:B4 , :B3 , :B2 , :B1 , 1)
- **Etc.**
- **Not significant statement: 'd3apnz76qwmrs'**
 - select pol#, usr_name, usr_labels, package, privs from lbac\$user_logon where usr_name = :username

Top Activity screen in the new test...



Transactions – Browse Orders...



ORACLE Enterprise Manager Cloud Control 12c

Enterprise | Targets | Favorites | History | Search Target Name | Logged in as SYS

ghperfdb.gerwinhendriksen.local

Oracle Database | Performance | Availability | Schema | Administration

Active Sessions

Maximum CPU

Detail for Selected 5 Minute Interval

Start Time 26-Feb-2013 23:20:47

Top SQL

Actions Schedule SQL Tuning Advisor Go

Select All | Select None

Select	Activity (%)	SQL ID	SQL Type
<input type="checkbox"/>	32.99	apgb2g9qzjhl	PL/SQL EXECUTE
<input type="checkbox"/>	24.74	0ruh367af7gbw	SELECT
<input type="checkbox"/>	15.46	8dq0v1mjngj7t	SELECT
<input type="checkbox"/>	6.70	4tx09gq85z630	SELECT
<input type="checkbox"/>	3.61	d2krbdzrhvq	INSERT
<input type="checkbox"/>	3.61	c749bc43qqfz3	SELECT
<input type="checkbox"/>	3.61	3hatpjzrqvfn7	UPDATE
<input type="checkbox"/>	2.58	5raw2bxz227wp	INSERT
<input type="checkbox"/>	1.55	7hcsn40jzrajs	PL/SQL EXECUTE

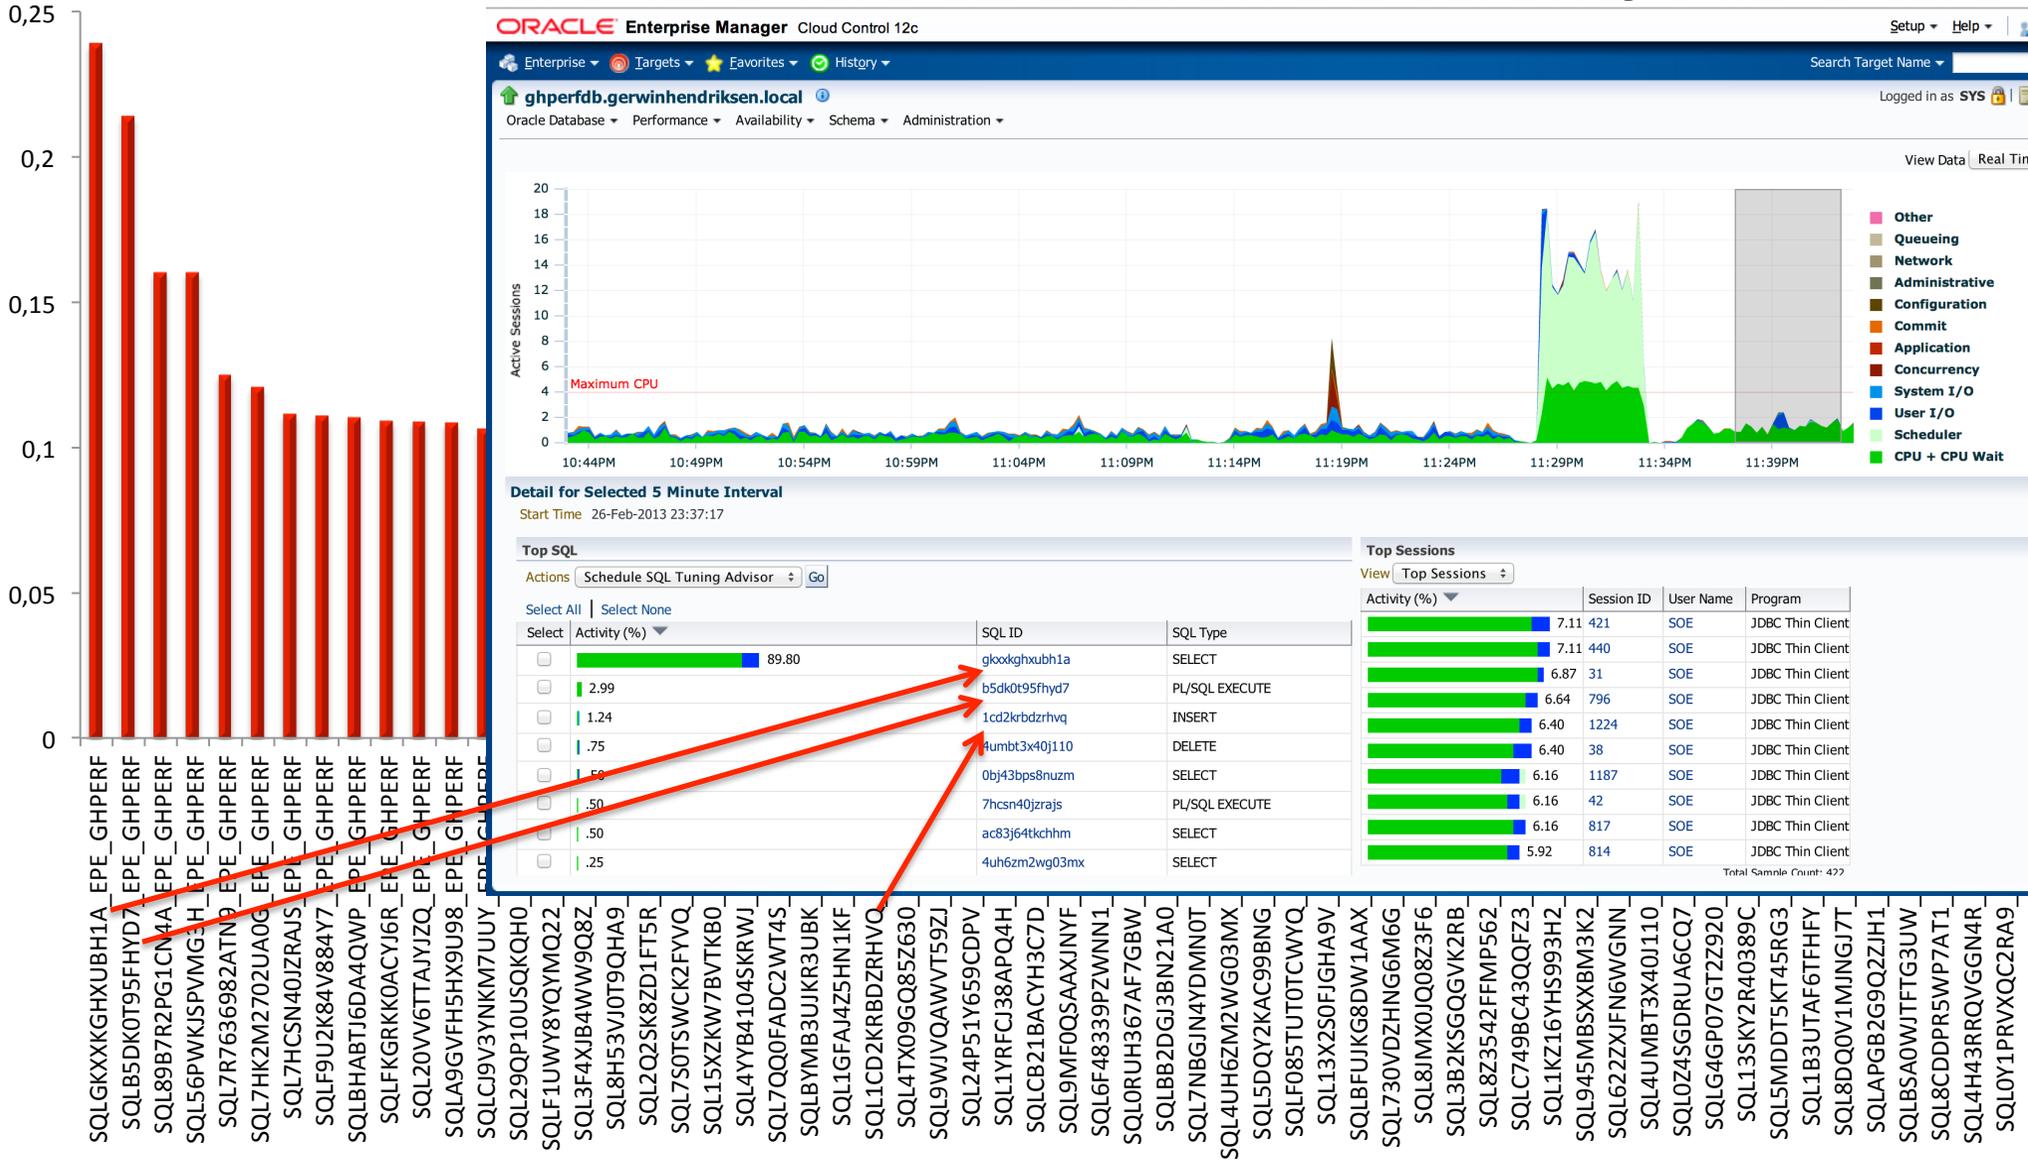
Top Sessions

View Top Sessions

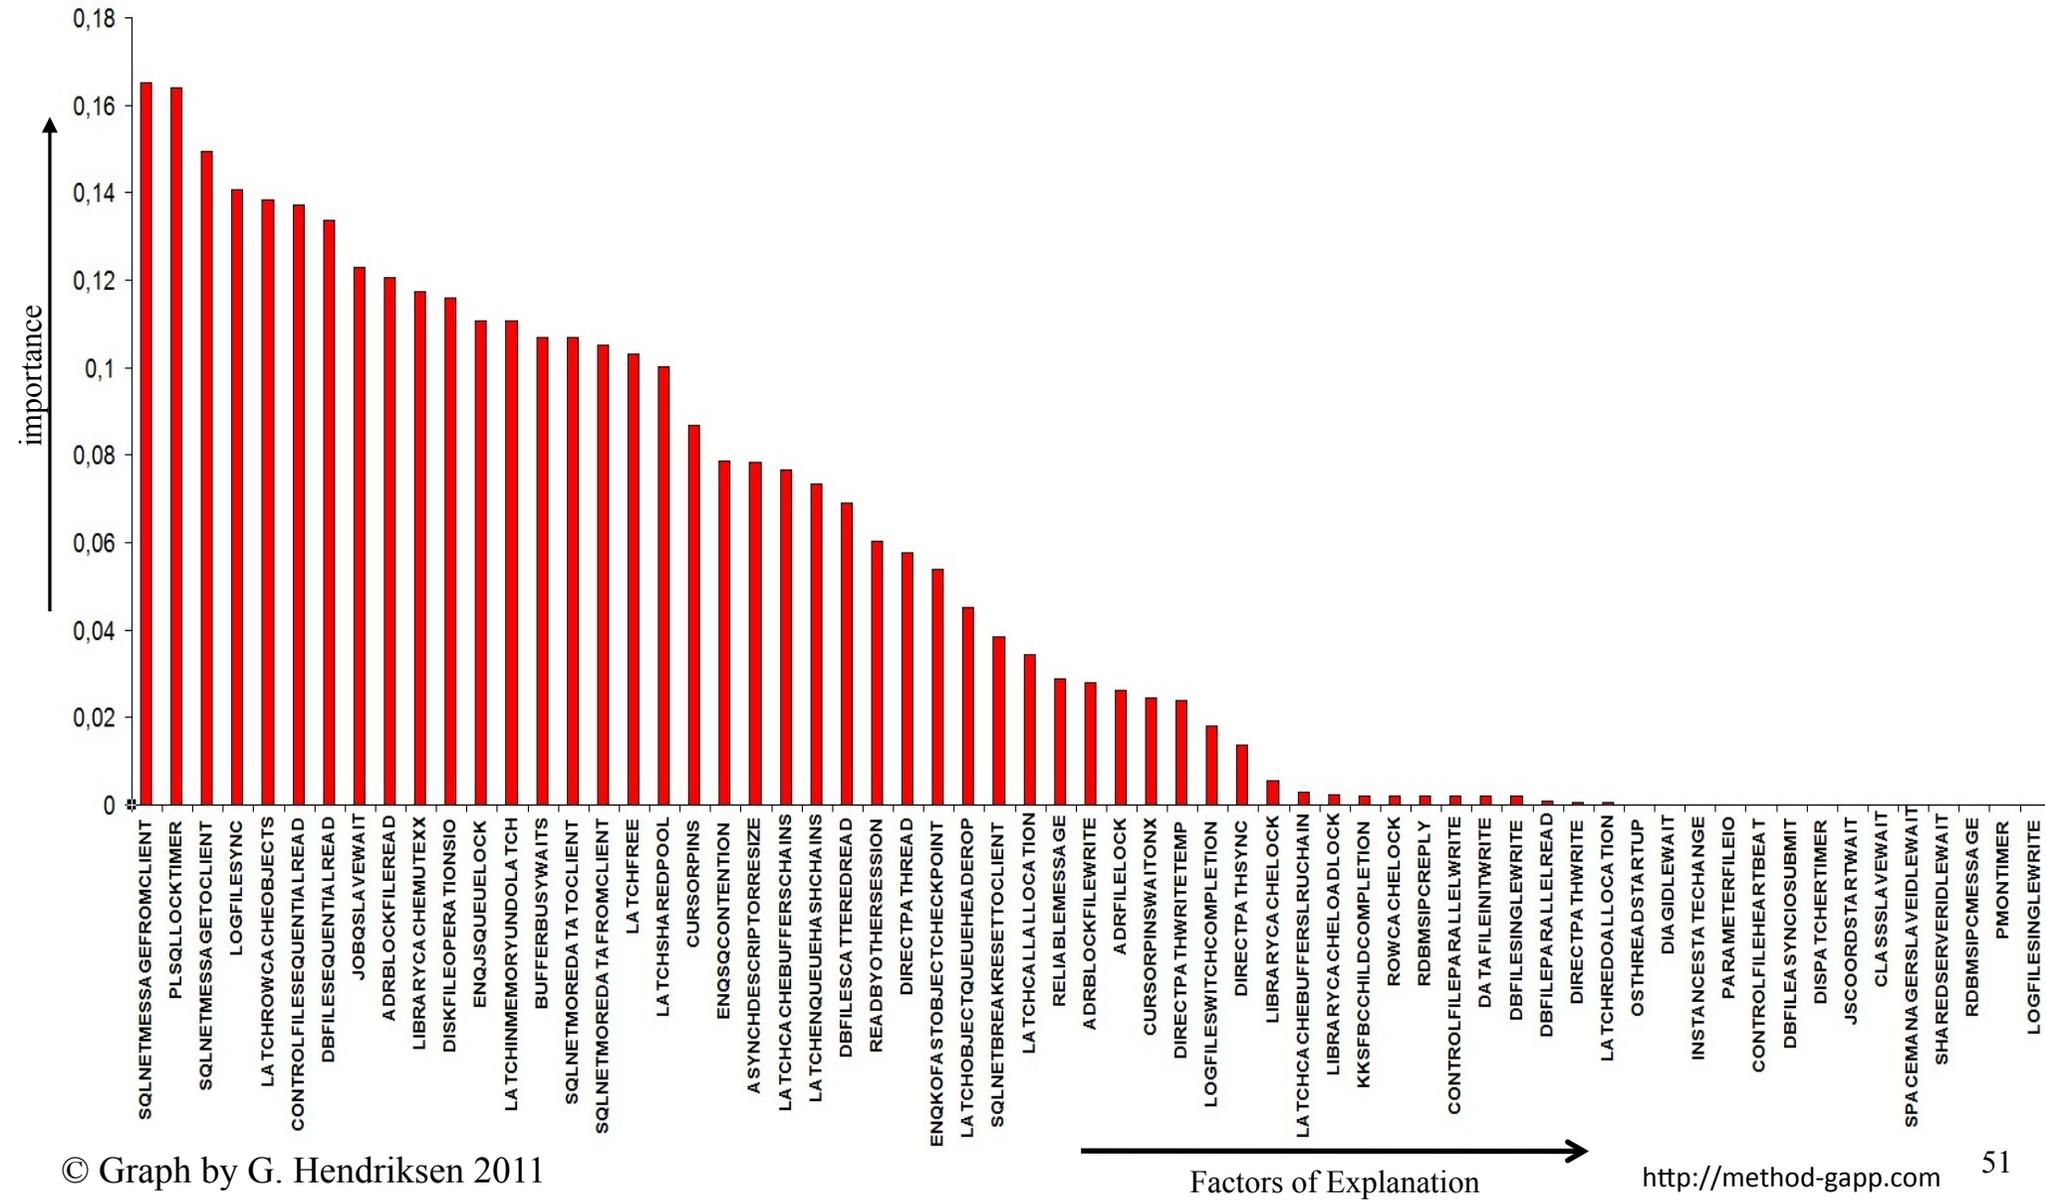
Activity (%)	Session ID	User Name	Program
15.36	1174	SYS	oracle@ghperf.gerwinhendriksen.local (LGWR)
5.36	783	SYS	oracle@ghperf.gerwinhendriksen.local (DBW0)
4.29	1171	SOE	JDBC Thin Client
3.21	1183	SYS	oracle@ghperf.gerwinhendriksen.local (M000)
3.21	439	SOE	JDBC Thin Client
3.21	442	SYSMAN	OMS
2.86	398	SOE	JDBC Thin Client
2.50	419	SOE	JDBC Thin Client
2.50	436	SOE	JDBC Thin Client
2.50	818	SOE	JDBC Thin Client

Total Sample Count: 280

Transactions – Warehouse Query



Events Data from AWR...



Involved wait events...

- SQL*Net message from client
- PL/SQL lock timer
- SQL*Net message to client
- **log file sync**
- latch: row cache objects
- control file sequential read
- **db file sequential read**
- Etc.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
log file sync	12,599	972	77	59.54	Commit
db file sequential read	29,465	379	13	23.19	User I/O
DB CPU		209		12.82	
resmgr:cpu quantum	2,475	15	6	0.91	Scheduler
log file switch (private strand flush incomplete)	28	14	509	0.87	Configuration

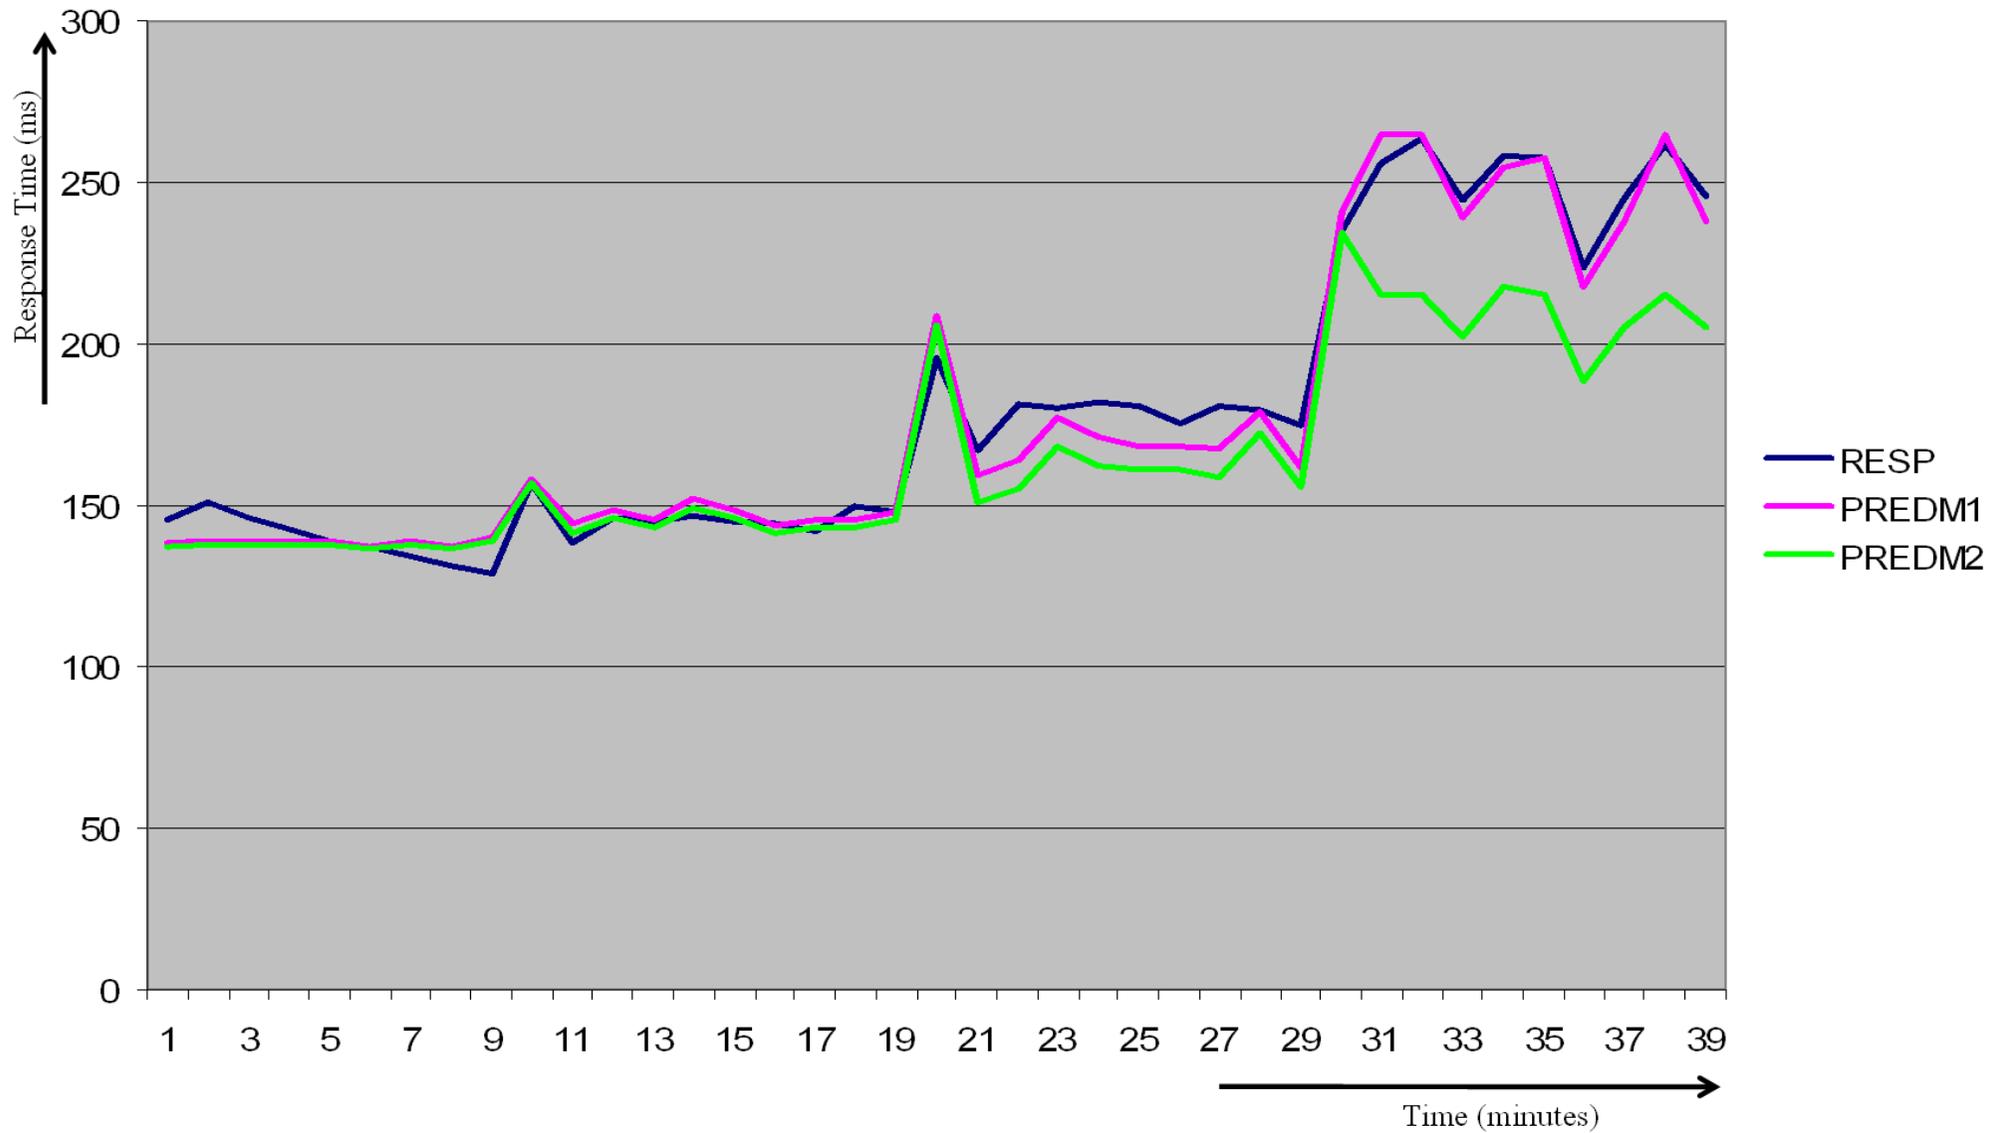
The data used in Method-GAPP was over a lot of hours the AWR report shown only over a 10 minute time slice.

Event	Waits	%Time -outs	Total Wait Time (s)	Avg wait (ms)	Waits /txn	% DB time
log file sync	12,599	0	972	77	0.43	59.54
db file sequential read	29,465	0	379	13	1.00	23.19
resmgr:cpu quantum	2,475	0	15	6	0.08	0.91
log file switch (private strand flush incomplete)	28	0	14	509	0.00	0.87
buffer busy waits	37	0	5	139	0.00	0.31
read by other session	10	0	0	33	0.00	0.02
SQL*Net message to client	19,922	0	0	0	0.68	0.02
library cache: mutex X	124	0	0	2	0.00	0.01
ADR block file read	4	0	0	38	0.00	0.01
control file sequential read	1,742	0	0	0	0.06	0.01
log file switch completion	2	0	0	30	0.00	0.00
latch free	11	0	0	4	0.00	0.00

The Fourth D

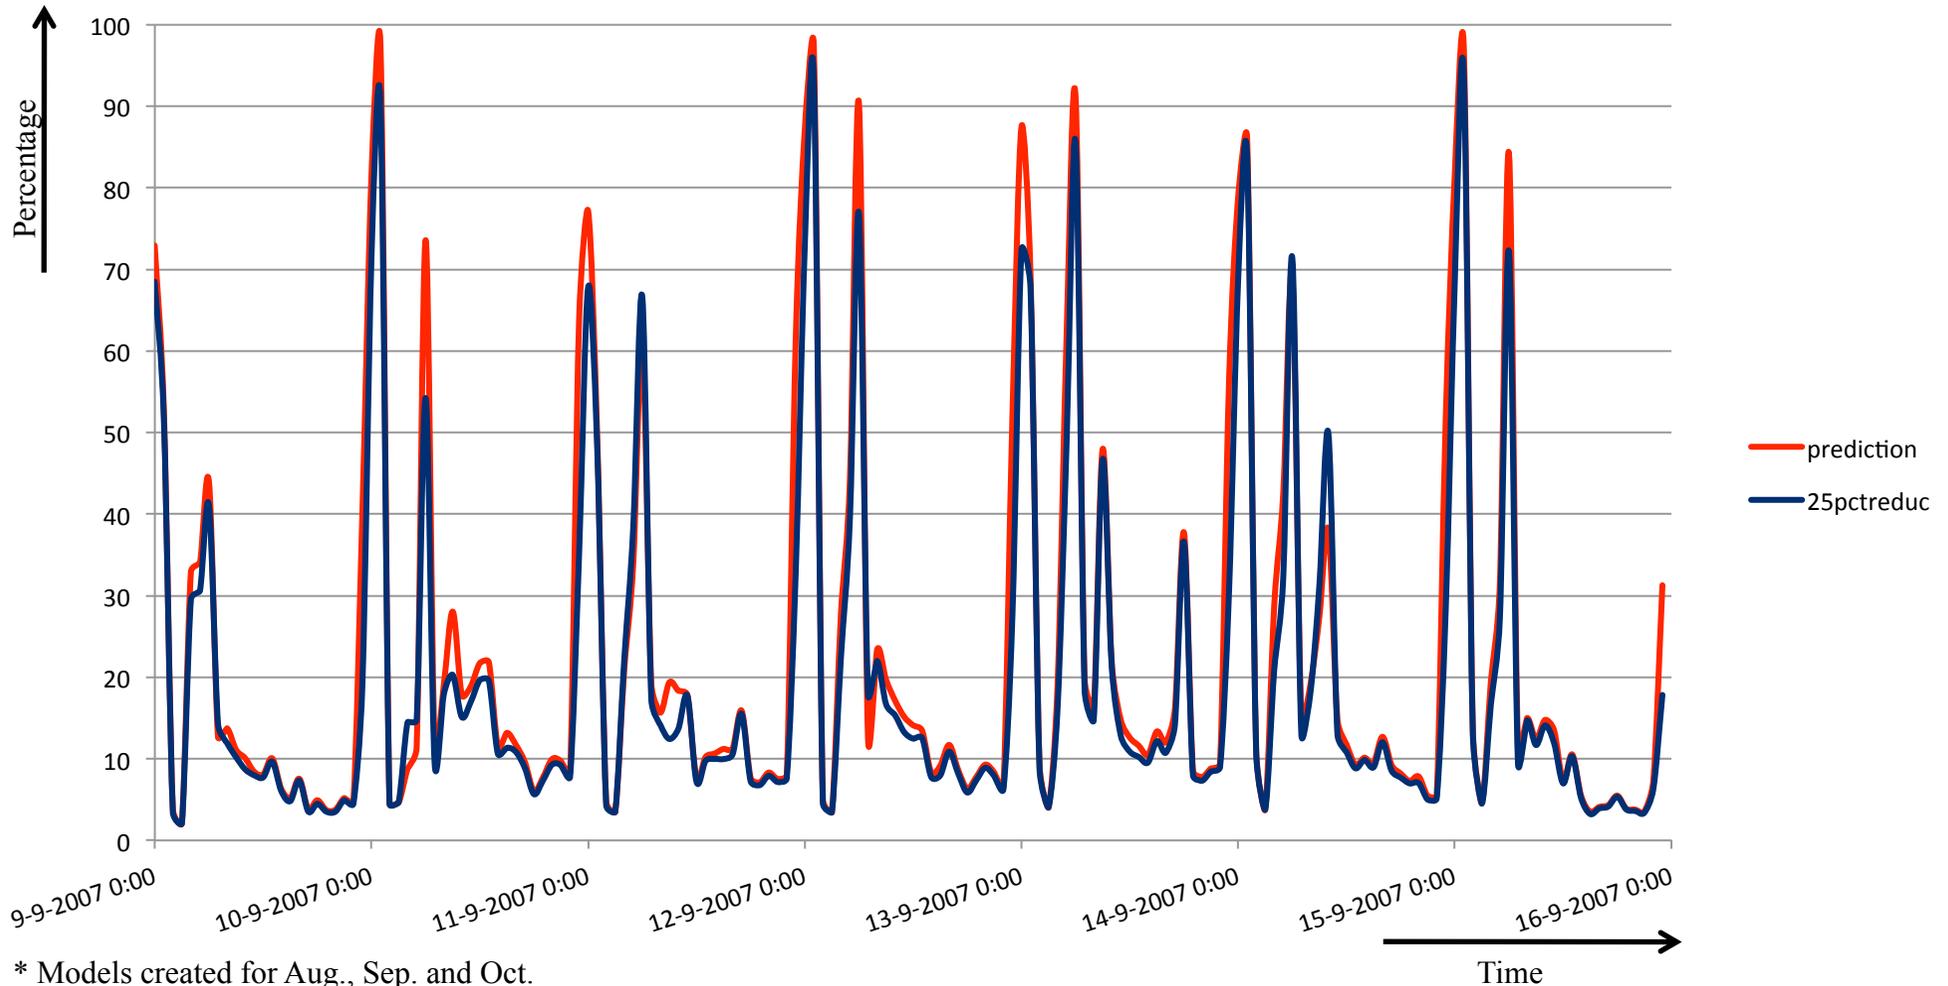
Data Mining - Predict

What will happen with 2 channels...



When 25% I/O Reduction...

Percentage logins (PCT_GT_5SEC) which took longer than 5 seconds:



* Models created for Aug., Sep. and Oct.

The Fifth D

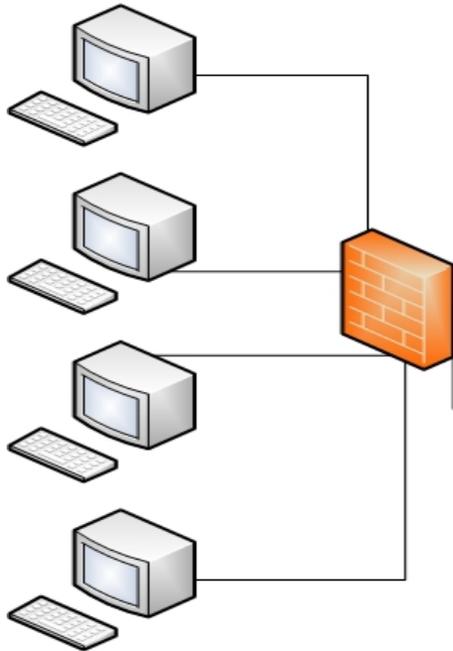
Data Interpretation

The Case...

- Response time of a very important business process in a time and labor system has sometimes very unpredictable performance
- The Used system is complex and the direct involved LPAR machines are:
 - HOT600 running several databases.
 - HOT720 running Portal and Oracle Internet Directory
 - HOT730 running Application Server for Time and Labor System

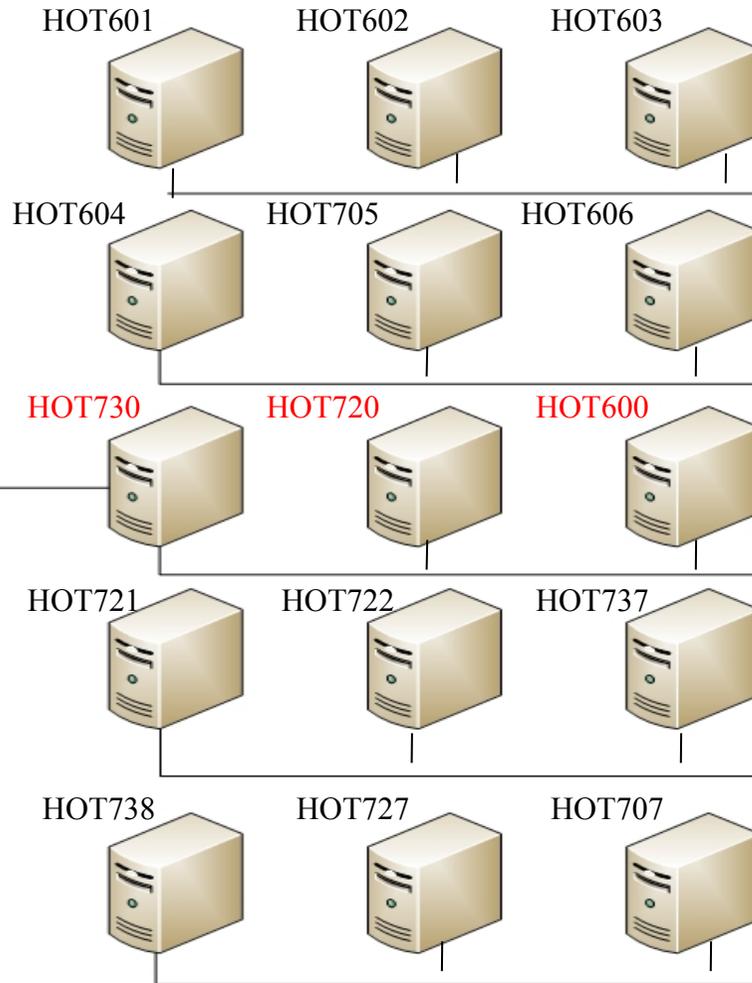
GAPP In A Complex Architecture...

Time and Labor
Users, via Web
Browser



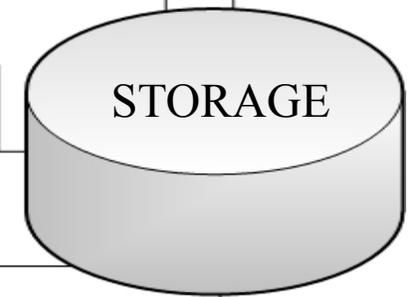
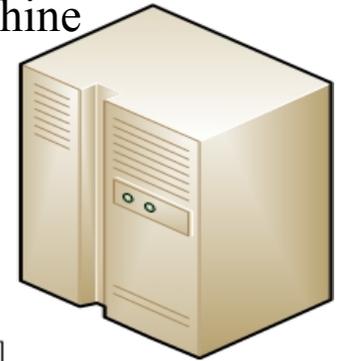
LPARS on the physical machine:

Direct involved LPAR's are in "RED" !!



Physical Machine

32 CPU's, 256Gb,
using AIX
Micropartitioning

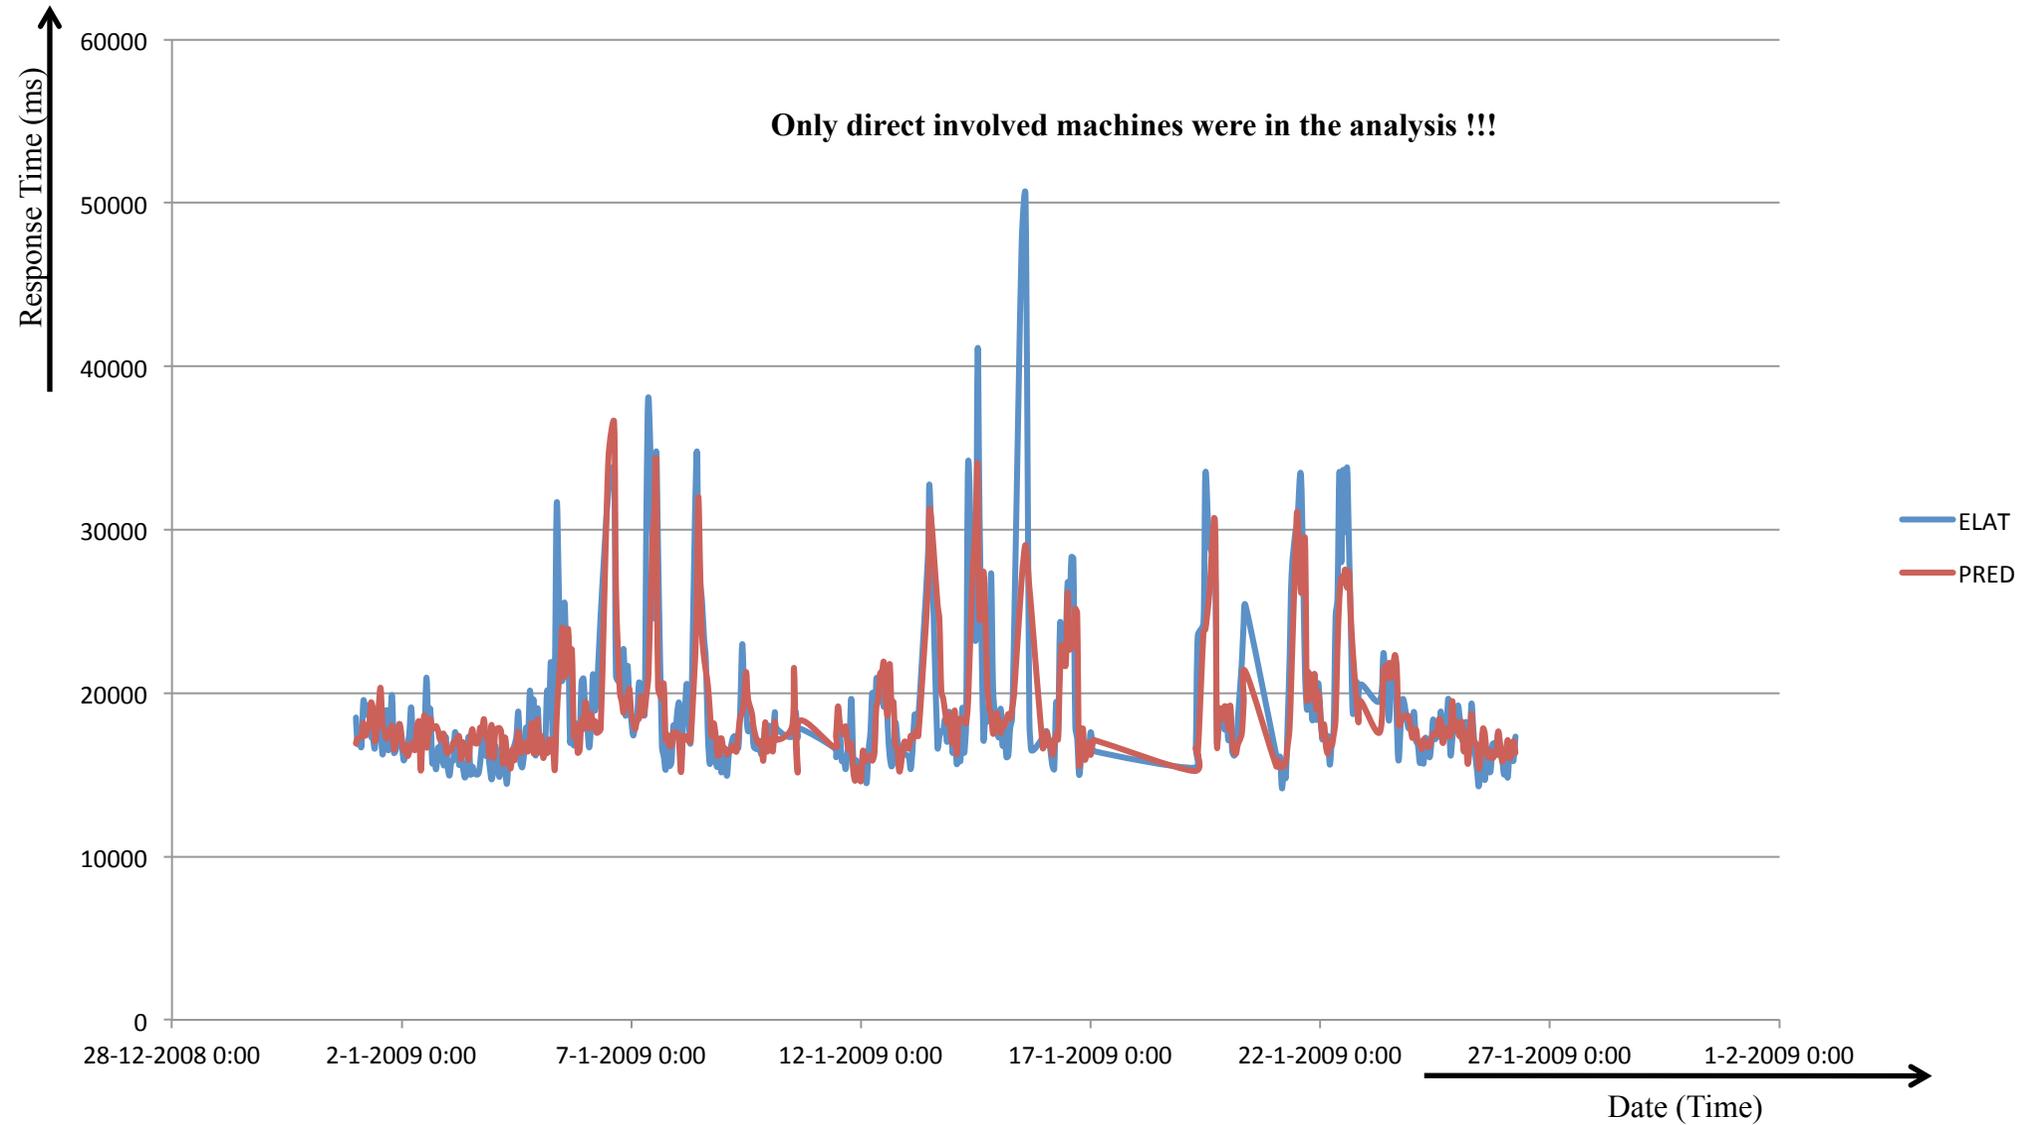


HOT730: Four Oracle Application Servers
for different enterprise applications.

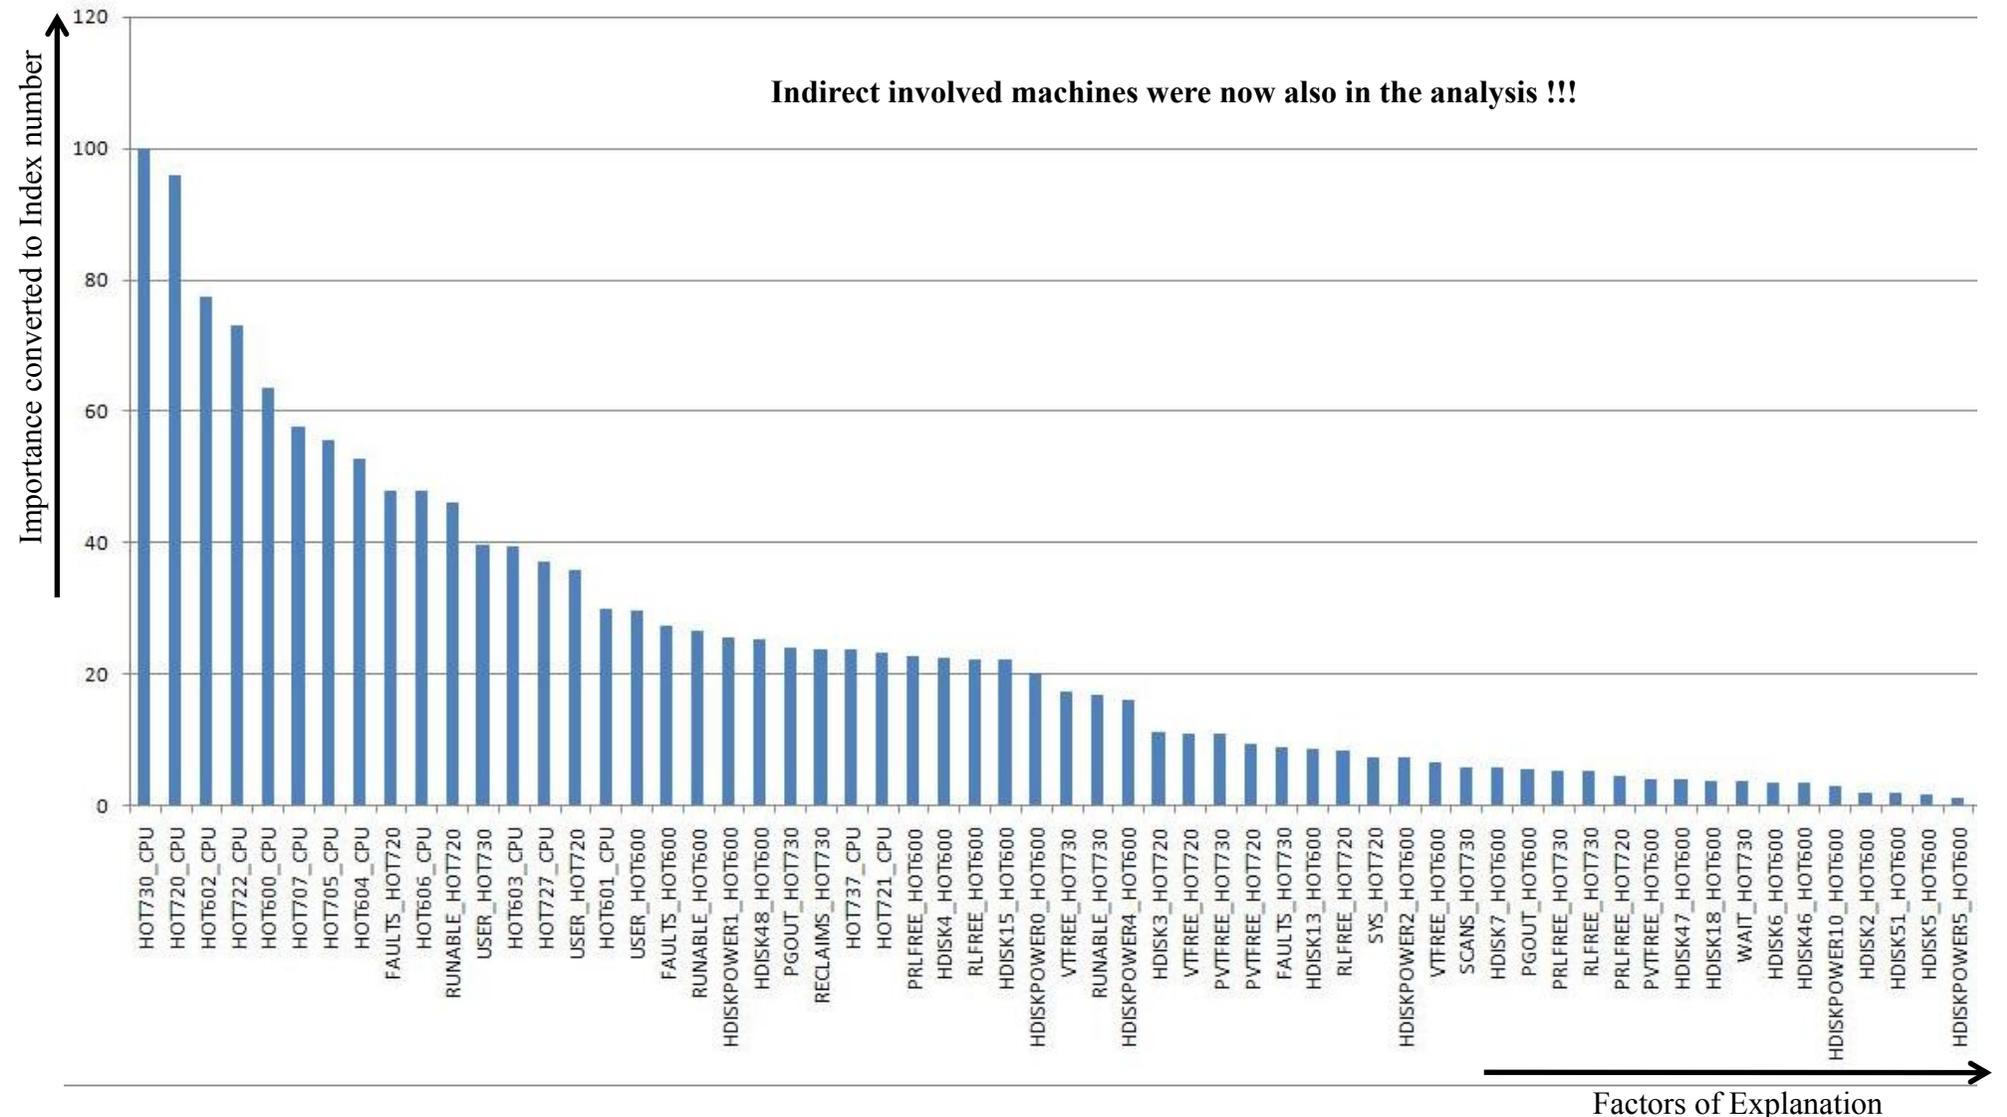
HOT720: Oracle Internet Directory (shared),
OID Database (shared), Portal Database
(shared), Legacy Application.

HOT600: Eight Oracle Databases

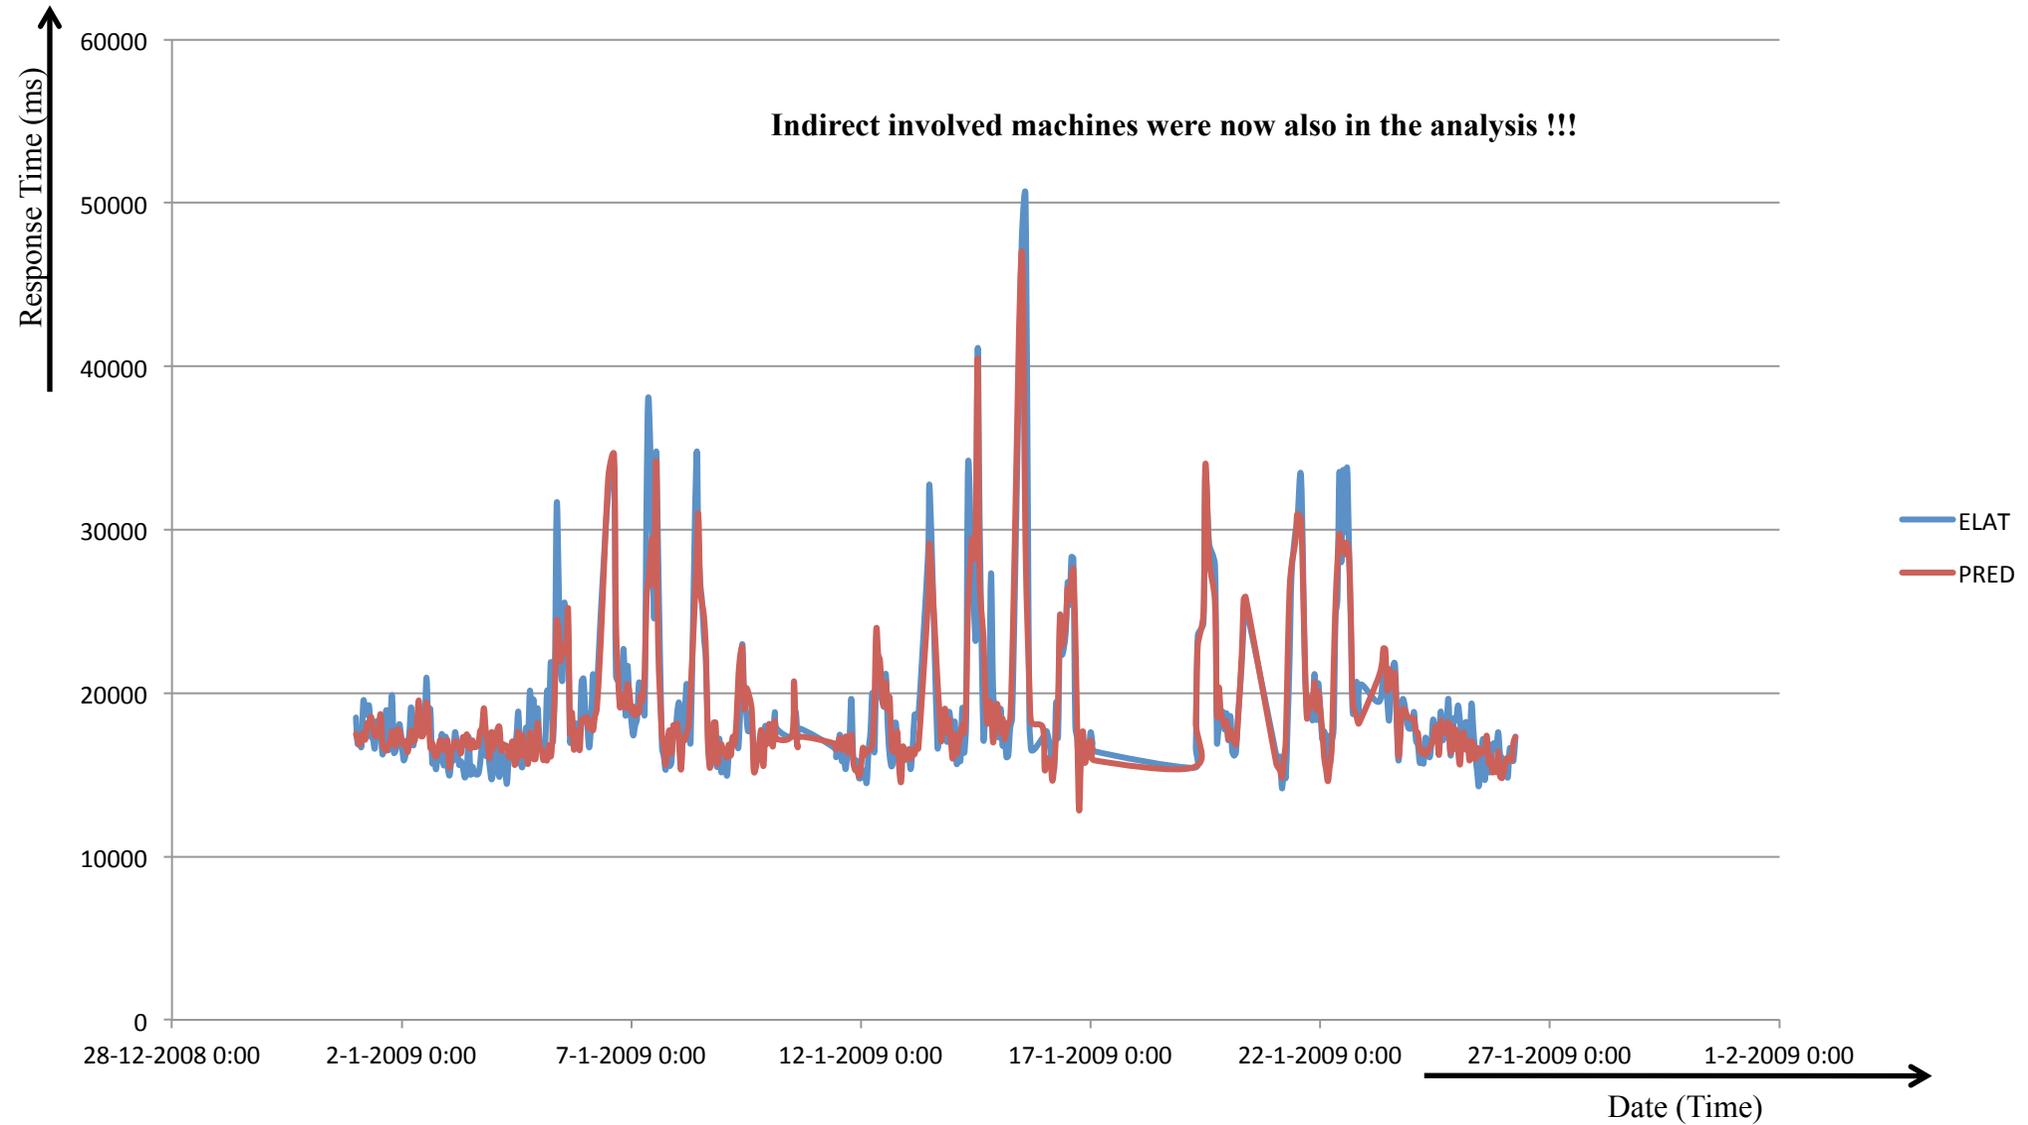
The Model (Predict)...



The Factorial Analyses (Explain)...



The Model (Predict)...



Interpretation...

- Most of the peaks are very well explained:
 - The performance of the involved business process was very depending on the machine's overall resources.
 - Beside the H0T730 and H0T720, the H0T602 and H0T722 had also much impact.
- The GAPP analyses showed that at the moment the physical machine gets short on physical Cpu all the different machines will have high impact on each other due to competition in hardware resources.
 - By Workload Scheduling the biggest problems could be solved

Conclusions

Conclusions...

- Using OEM 12C and AWR can help reduce significantly the time to gather and synchronize the data to use as input for the method.
- Although you didn't made "hooks" in the application you were able to find the problem, independent of technology.
- After your analyses was done you found out that also components in the infrastructure not part of the application were responsible for the encountered problem.
- The model you made, made it possible to do predictions of the impact of a possible investment on the end-user process performance.
- Adding secondary components data to the analyses can give a lot of extra detailed information regarding the faced performance problems.

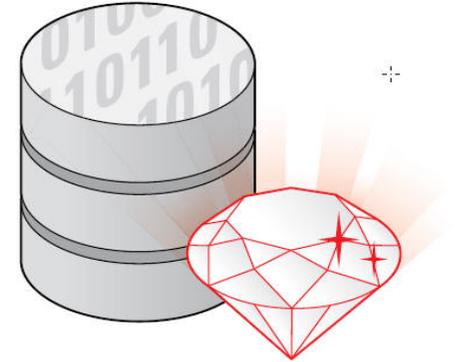
Considerations...

- Adding as secondary components “variance” of different components can be used to detect skew with Method-GAPP
- Adding as secondary components “Business Activity Metrics” can make Method-GAPP the way to connect business data with technical data.

What could you use for the mining...

- Oracle Advanced Analytics

- Cost 17K per CPU above EE
- Makes method-GAPP dependent on ORACLE
- Data can be stored in Oracle DB and be mined



- Project R:

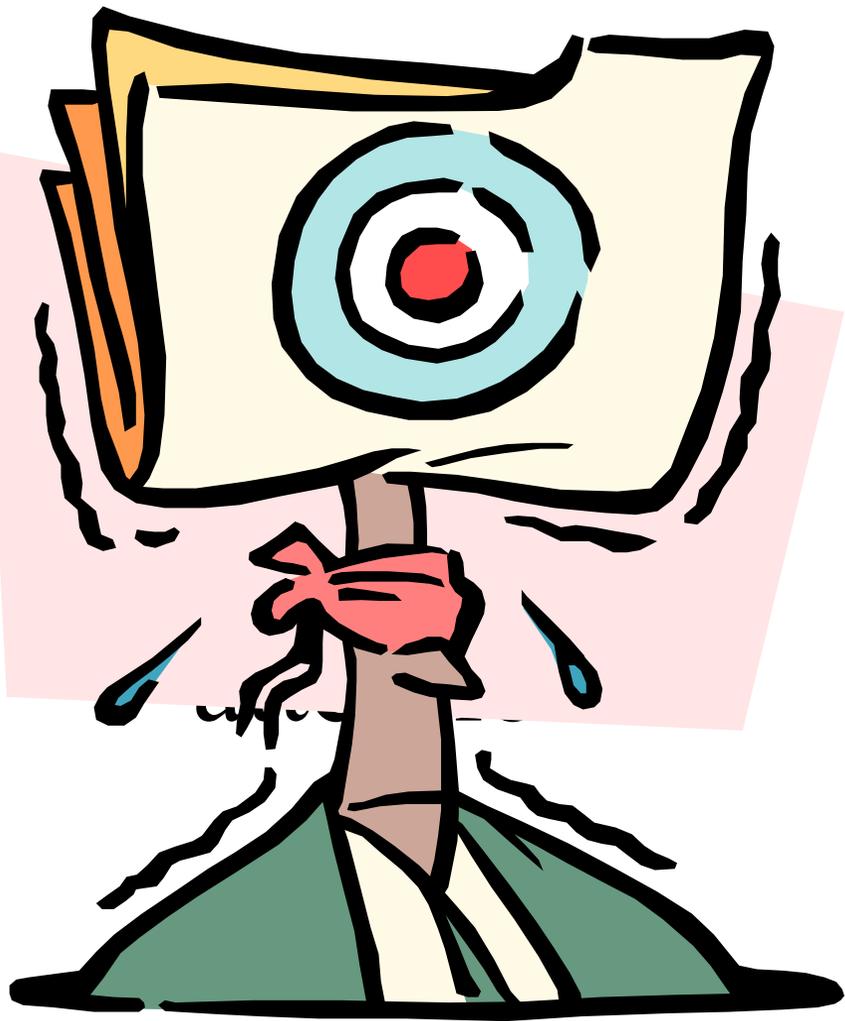
- Is Open Source
- Makes method-GAPP completely platform independent
- Needs Database (which could be open source)



Reference...

- Personal blog:
 - <http://blog.gerwinhendriksen.com>
 - <http://method-gapp.com>
- Method-R:
 - Oracle Performance by Cary Millsap (<http://method-r.com/>)
- Formula's:
 - Analyzing Computer Sys. Perf. By Dr. N.J. Gunther
- Oracle Data Mining:
 - <http://www.oracle.com/technology/products/bi/odm/index.html>
 - <http://www.oracle.com/technetwork/database/options/advanced-analytics/index.html>





Q/A