Advanced Oracle Troubleshooting

No magic is needed,
systematic approach will do

Tanel Poder

http://www.tanelpoder.com

ORACLE | CERTIFIED MASTER

Introduction

- About me:
  - Occupation: DBA, engineer, researcher
  - Expertise: Oracle internals geek, *End-to-end* performance & scalability
  - Oracle experience: 10 years as DBA
  - Certification: OCM (2002) OCP (1999)
  - Professional affiliations: OakTable Network
  - Blog: http://blog.tanelpoder.com

ORACLE | CERTIFIED MASTER

Introduction

- About this presentation:
  - Systematic *approach*, rather than *methodology*
  - Use *right* tools for *right* problems
  - Break complex problems down to simple problems
  - Therefore, use simple tools for simple problems
  - In other words, use a *systematic approach* and life will be easier!

- All scripts used here are freely available:
  - http://www.tanelpoder.com

Simple (but common) question:

What the $#*&%! is that session doing?

demo1.sql

Non-systematic troubleshooting

- Check alert.log…
- Check for disk and tablespace free space…
- Check for locks…
- Check for xyz…

"We did a healthcheck and everything looks OK!"

?????!

Semi-systematic troubleshooting

- Quick check for *usual suspects*
  - System load, locks, etc…
- Look into Statspack…
- Enable sql_trace…

…then what?

Systematic Troubleshooting Demo

```
SQL> @sw 114

    SID STATE   EVENT                                     SEQ# SEC_IN_WAIT         P1         P2         P3 P1TRANSL
------- ------- ---------------------------------------- ---------- ----------- ---------- ---------- ---------- ----------------------
    114 WAITING enq: TX - row lock contention                  21           9 1415053318     131081       2381 0x54580006: TX mode 6

SQL> @sw &mysid

    SID STATE   EVENT                                     SEQ# SEC_IN_WAIT         P1         P2         P3 P1TRANSL
------- ------- ---------------------------------------- ---------- ----------- ---------- ---------- ---------- ----------------------
    107 WORKING On CPU / runqueue                              89           0 1413697536          1          0

SQL>
SQL> @sn 5 &mysid

-- Session Snapper v1.06 by Tanel Poder ( http://www.tanelpoder.com )

-------------------------------------------------------------------------------------------------------------------------------------------
HEAD,     SID, SNAPSHOT START   ,  SECONDS, TYPE, STATISTIC                               ,         DELTA,  DELTA/SEC,     HDELTA, HDELTA/SEC
-------------------------------------------------------------------------------------------------------------------------------------------
DATA,       9, 20080221 22:05:08,        5, STAT, recursive calls                       ,             1,          0,          1,         .2
DATA,       9, 20080221 22:05:08,        5, STAT, recursive cpu usage                   ,             1,          0,          1,         .2
DATA,       9, 20080221 22:05:08,        5, STAT, session pga memory max                ,         25292,       5058,     25.29k,       5.06k
DATA,       9, 20080221 22:05:08,        5, STAT, calls to get snapshot scn: kcmgss     ,             1,          0,          1,         .2
DATA,       9, 20080221 22:05:08,        5, STAT, workarea executions - optimal         ,            18,          4,         18,        3.6
DATA,       9, 20080221 22:05:08,        5, STAT, execute count                         ,             1,          0,          1,         .2
DATA,       9, 20080221 22:05:08,        5, STAT, sorts (memory)                        ,            11,          2,         11,        2.2
DATA,       9, 20080221 22:05:08,        5, STAT, sorts (rows)                          ,          1904,        381,       1.9k,      380.8
DATA,       9, 20080221 22:05:08,        5, WAIT, PL/SQL lock timer                     ,       4999649,     999930,         5s,   999.93ms
--  End of snap 1

PL/SQL procedure successfully completed.
```

Troubleshooting approaches

- How do you solve problems?

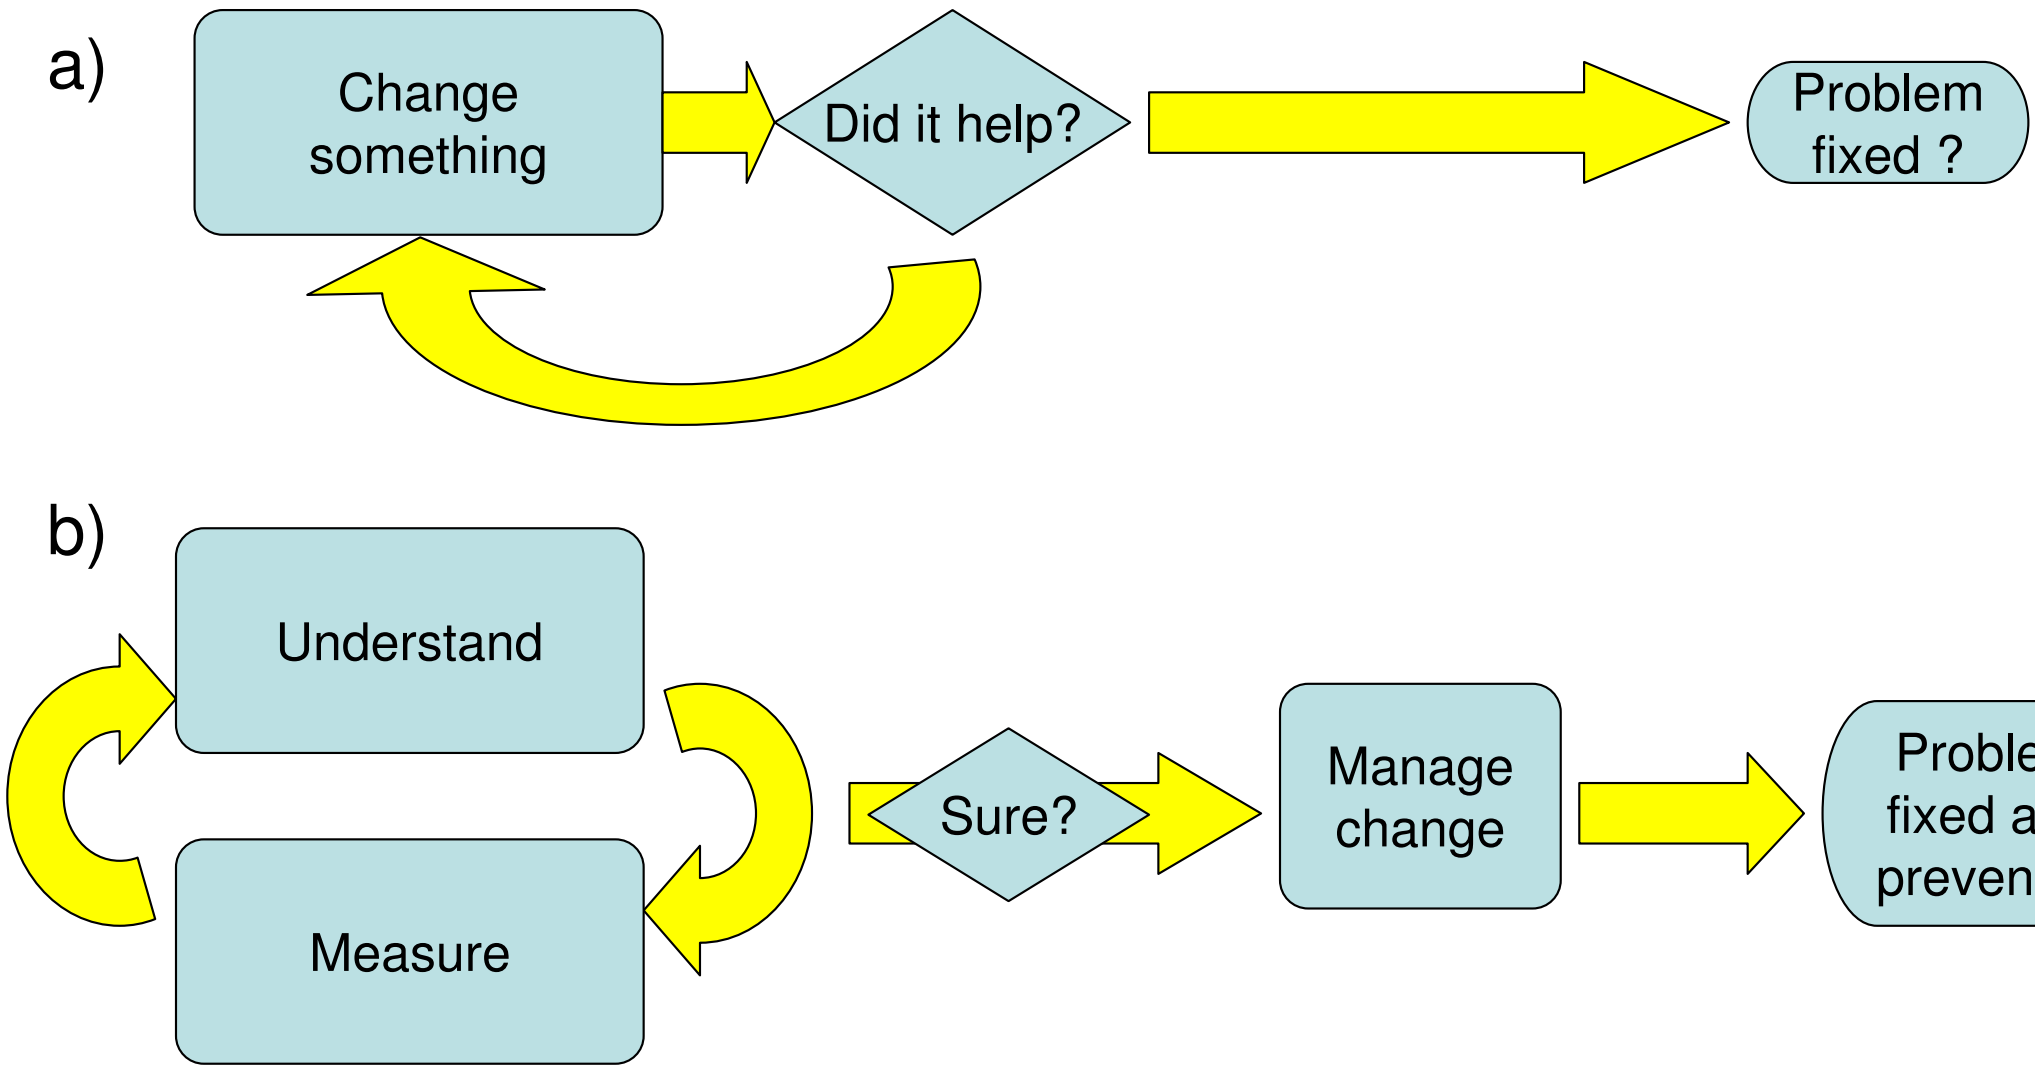

Systematic troubleshooting

- Understand the "flow" of a server process
- …and how to measure it
- …then measure it
- …step by step
- …using right tool at right step
- ...fix the problem *once you understand it*

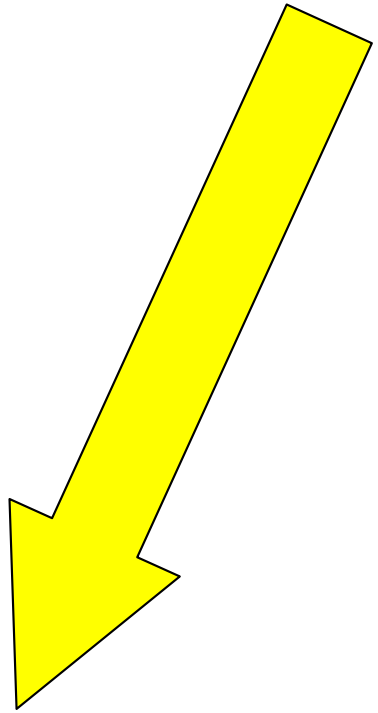

Simple (but common) question:

What the $#*&%! is that session doing?

demo2.sql

Understand the problem

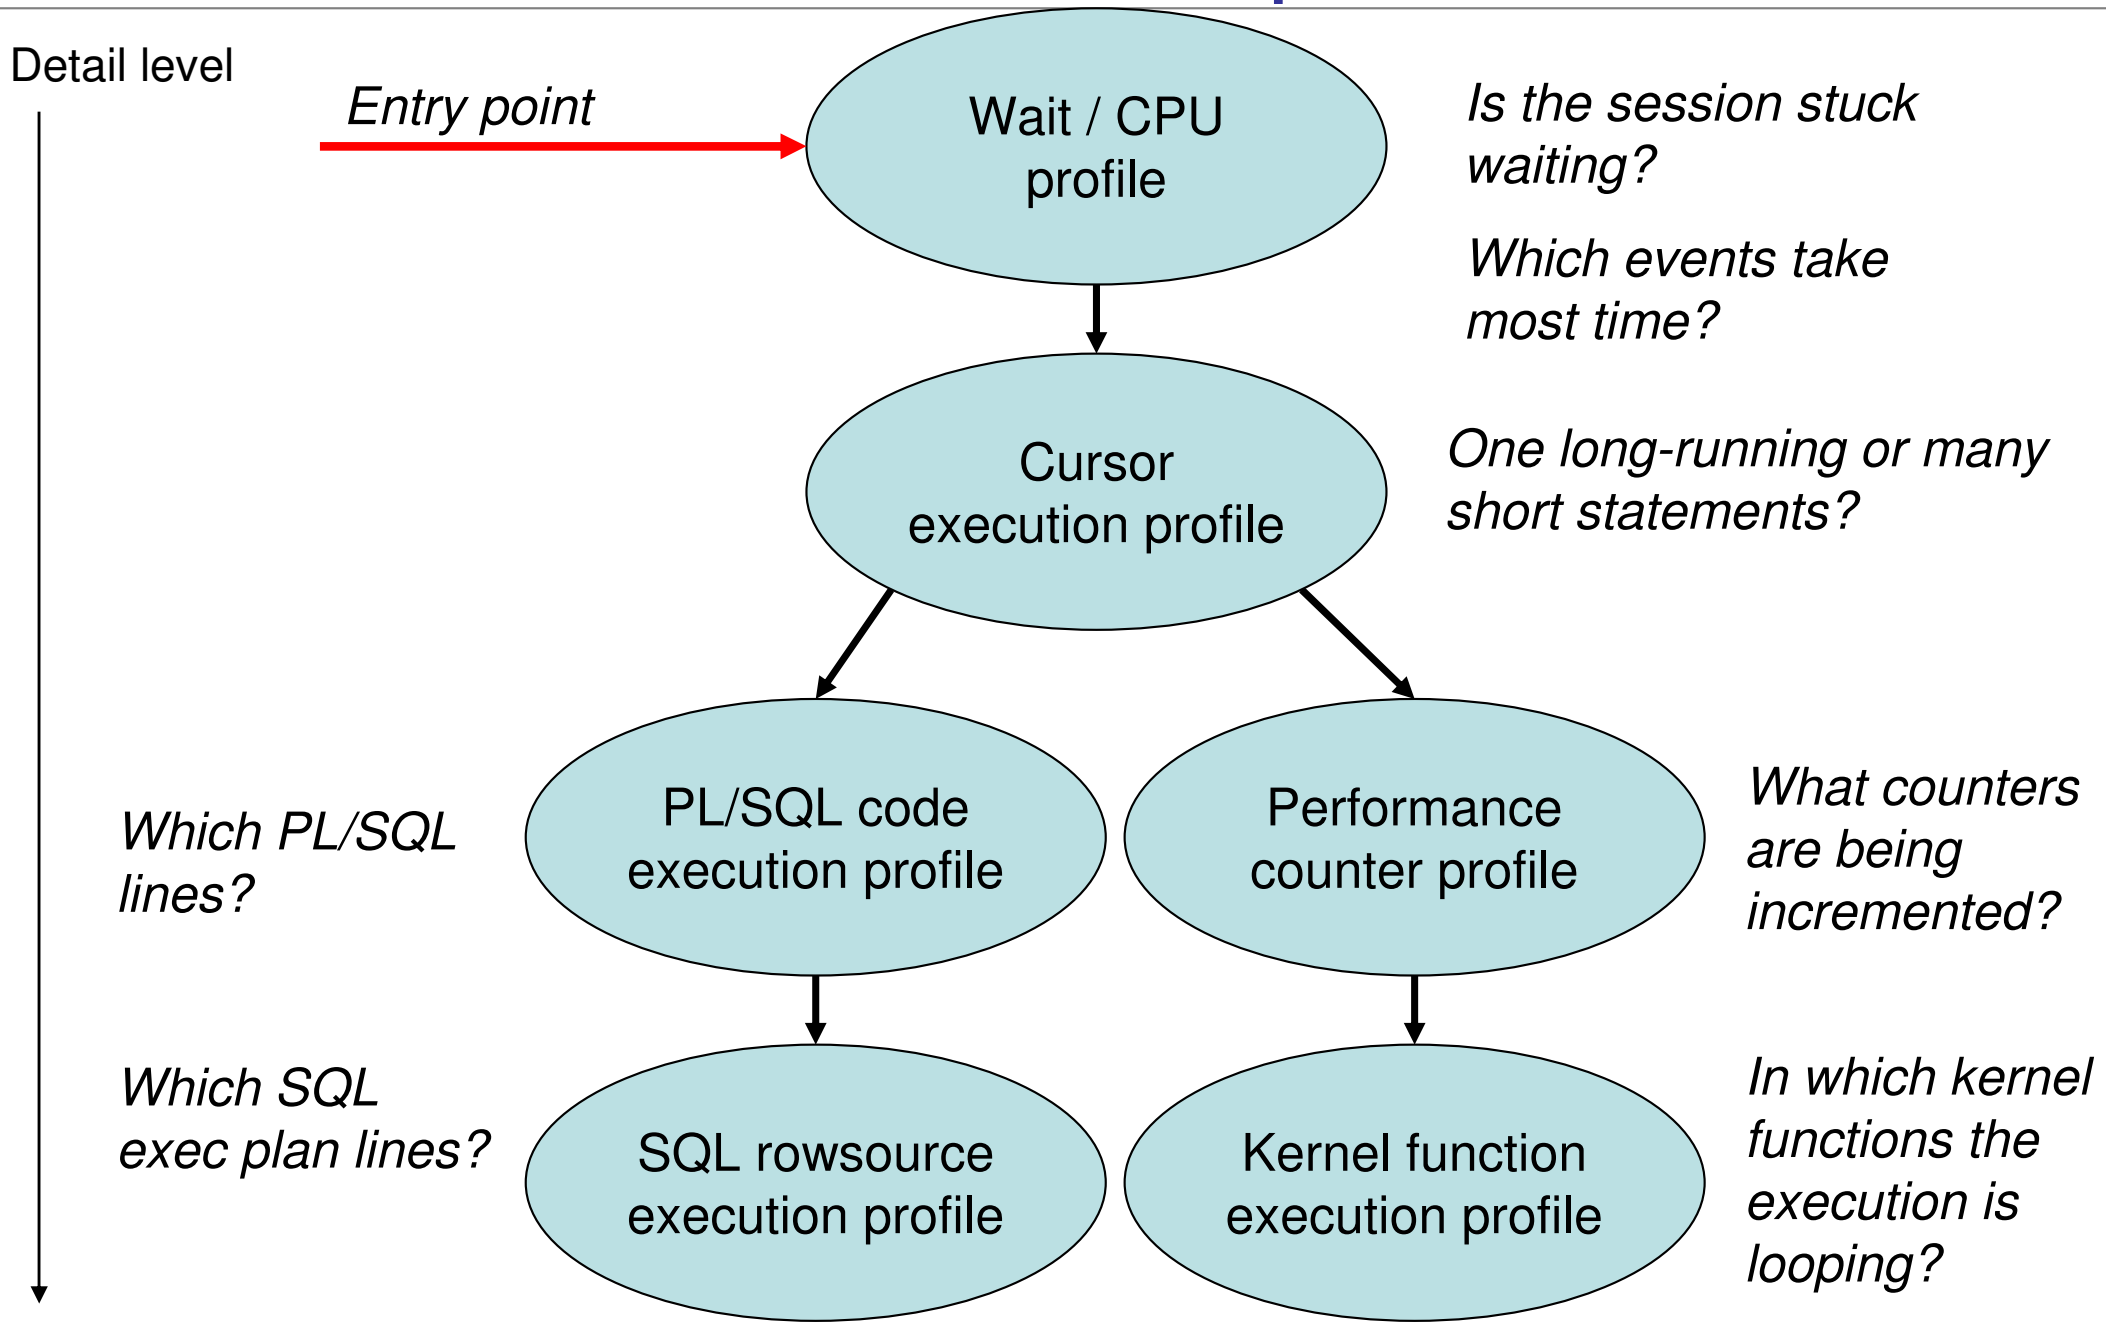

Right tools for *measuring* right problems

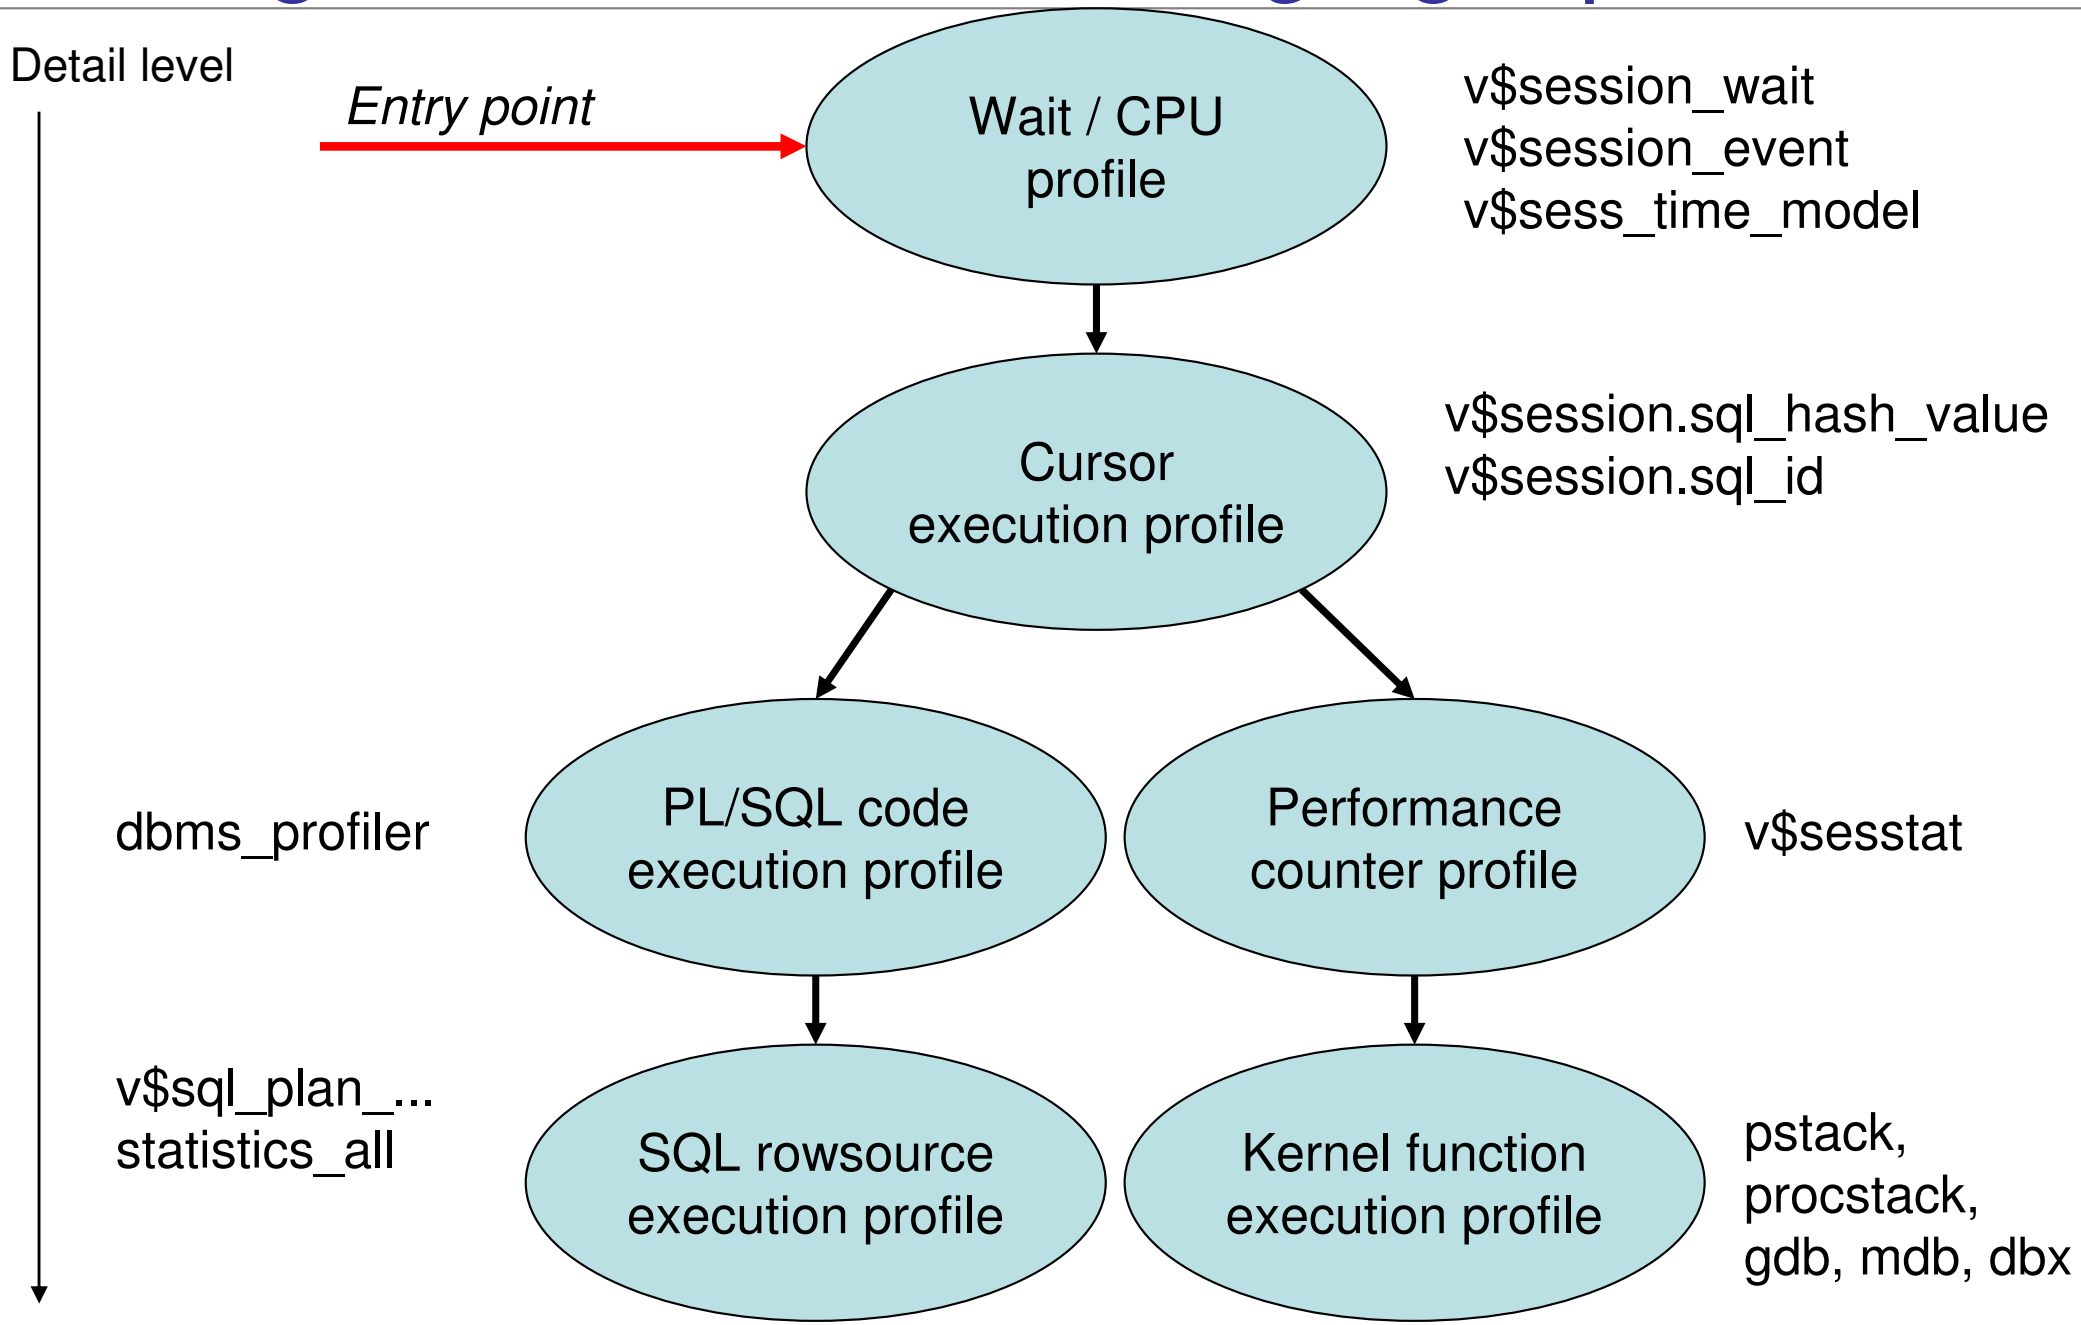

Right tools for *measuring* right problems

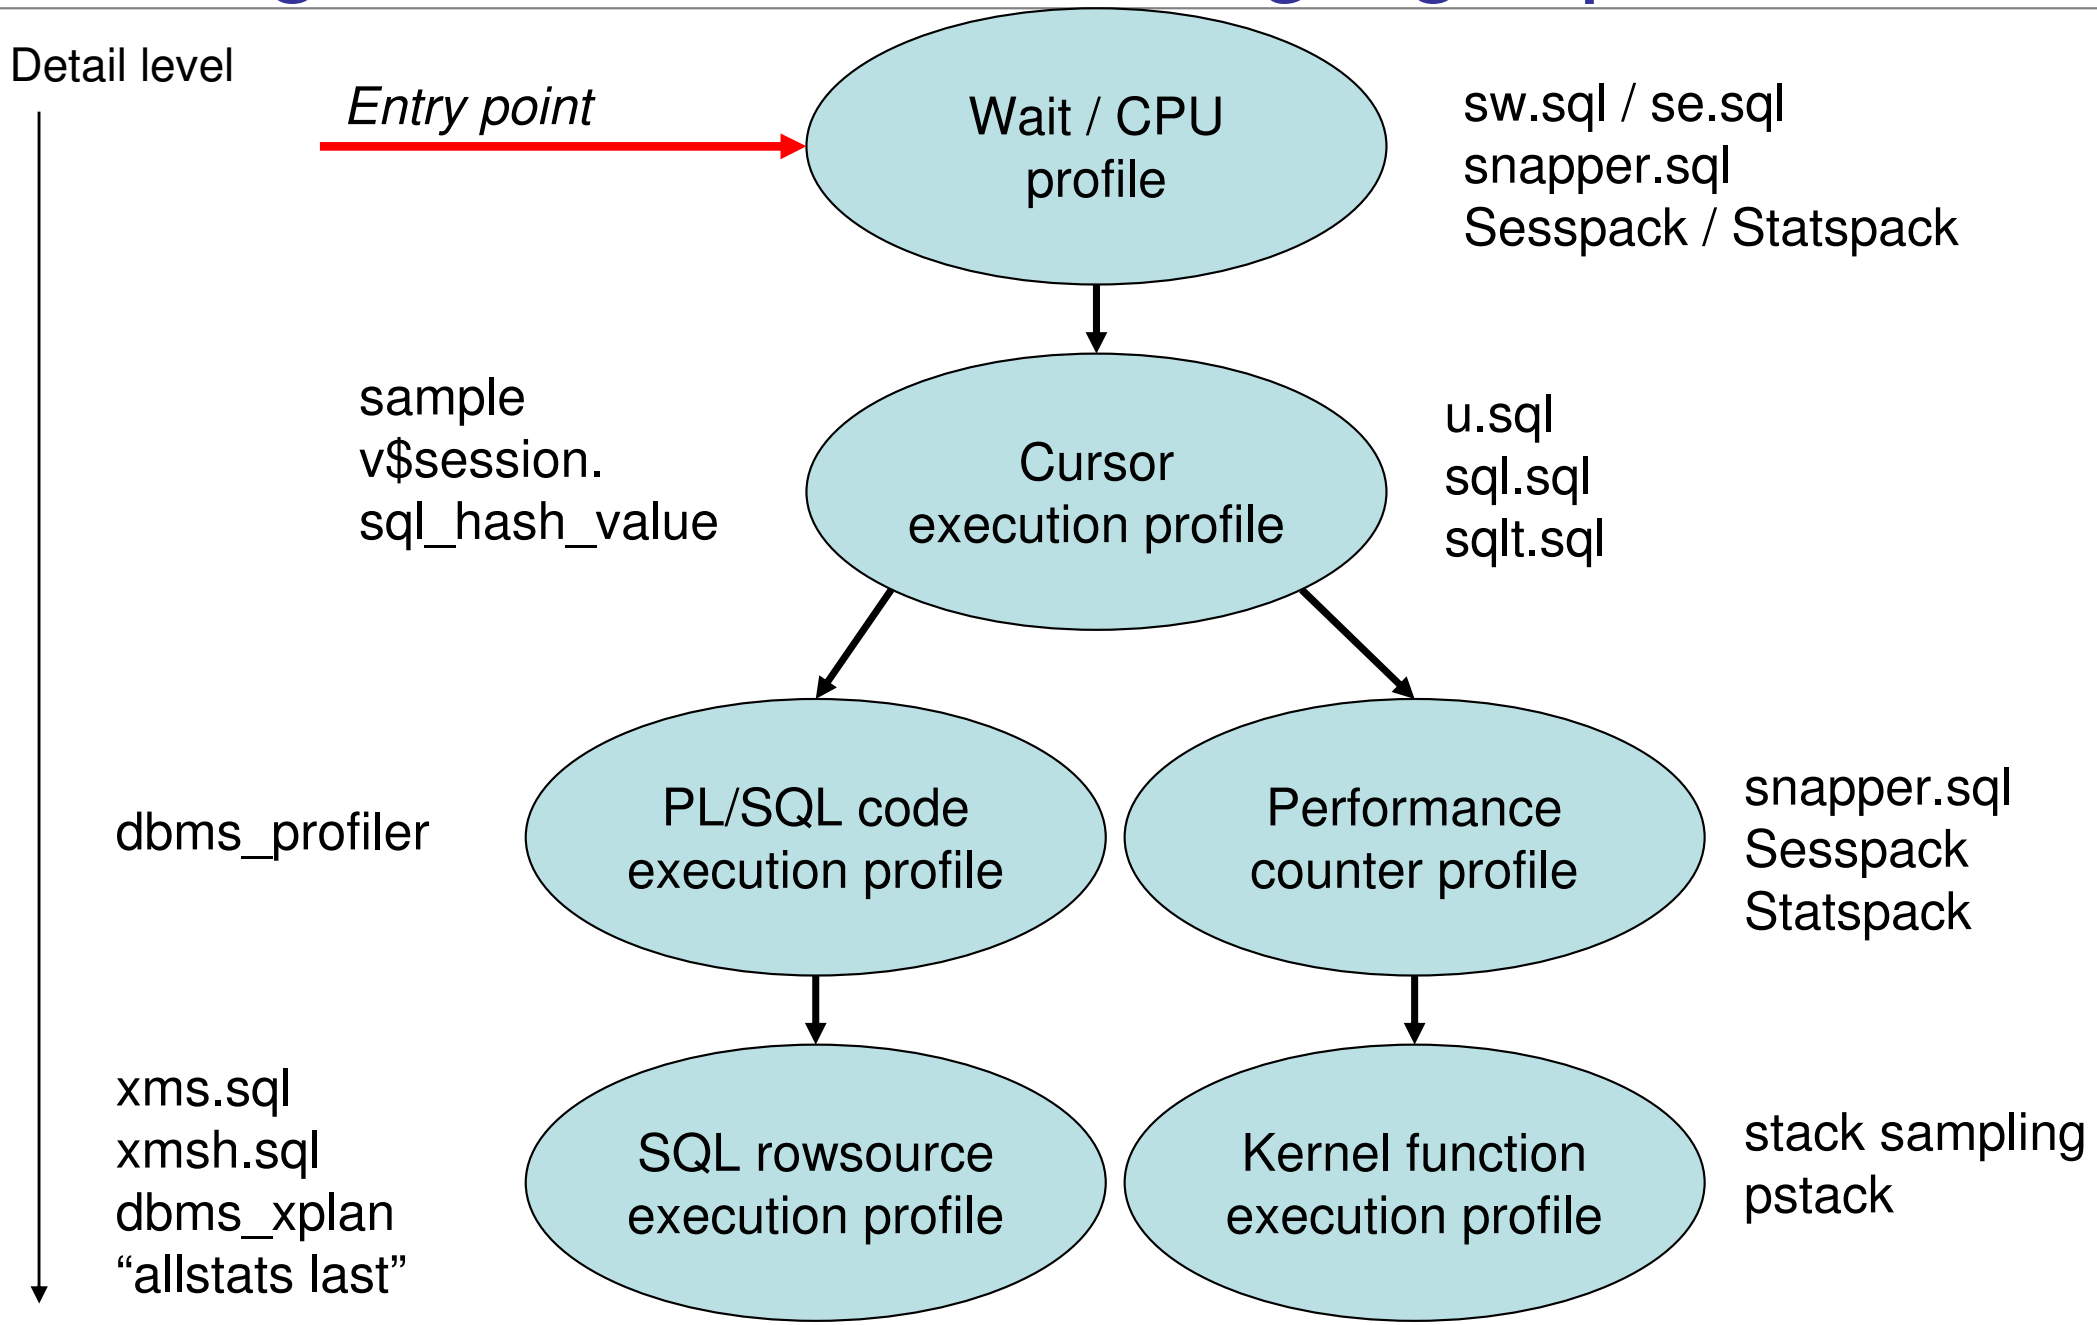

Simple (but common) question:

What the $#*&%! is that session doing?

demo3.sql

Understand the Oracle process flow…

- High level process flow explanation…

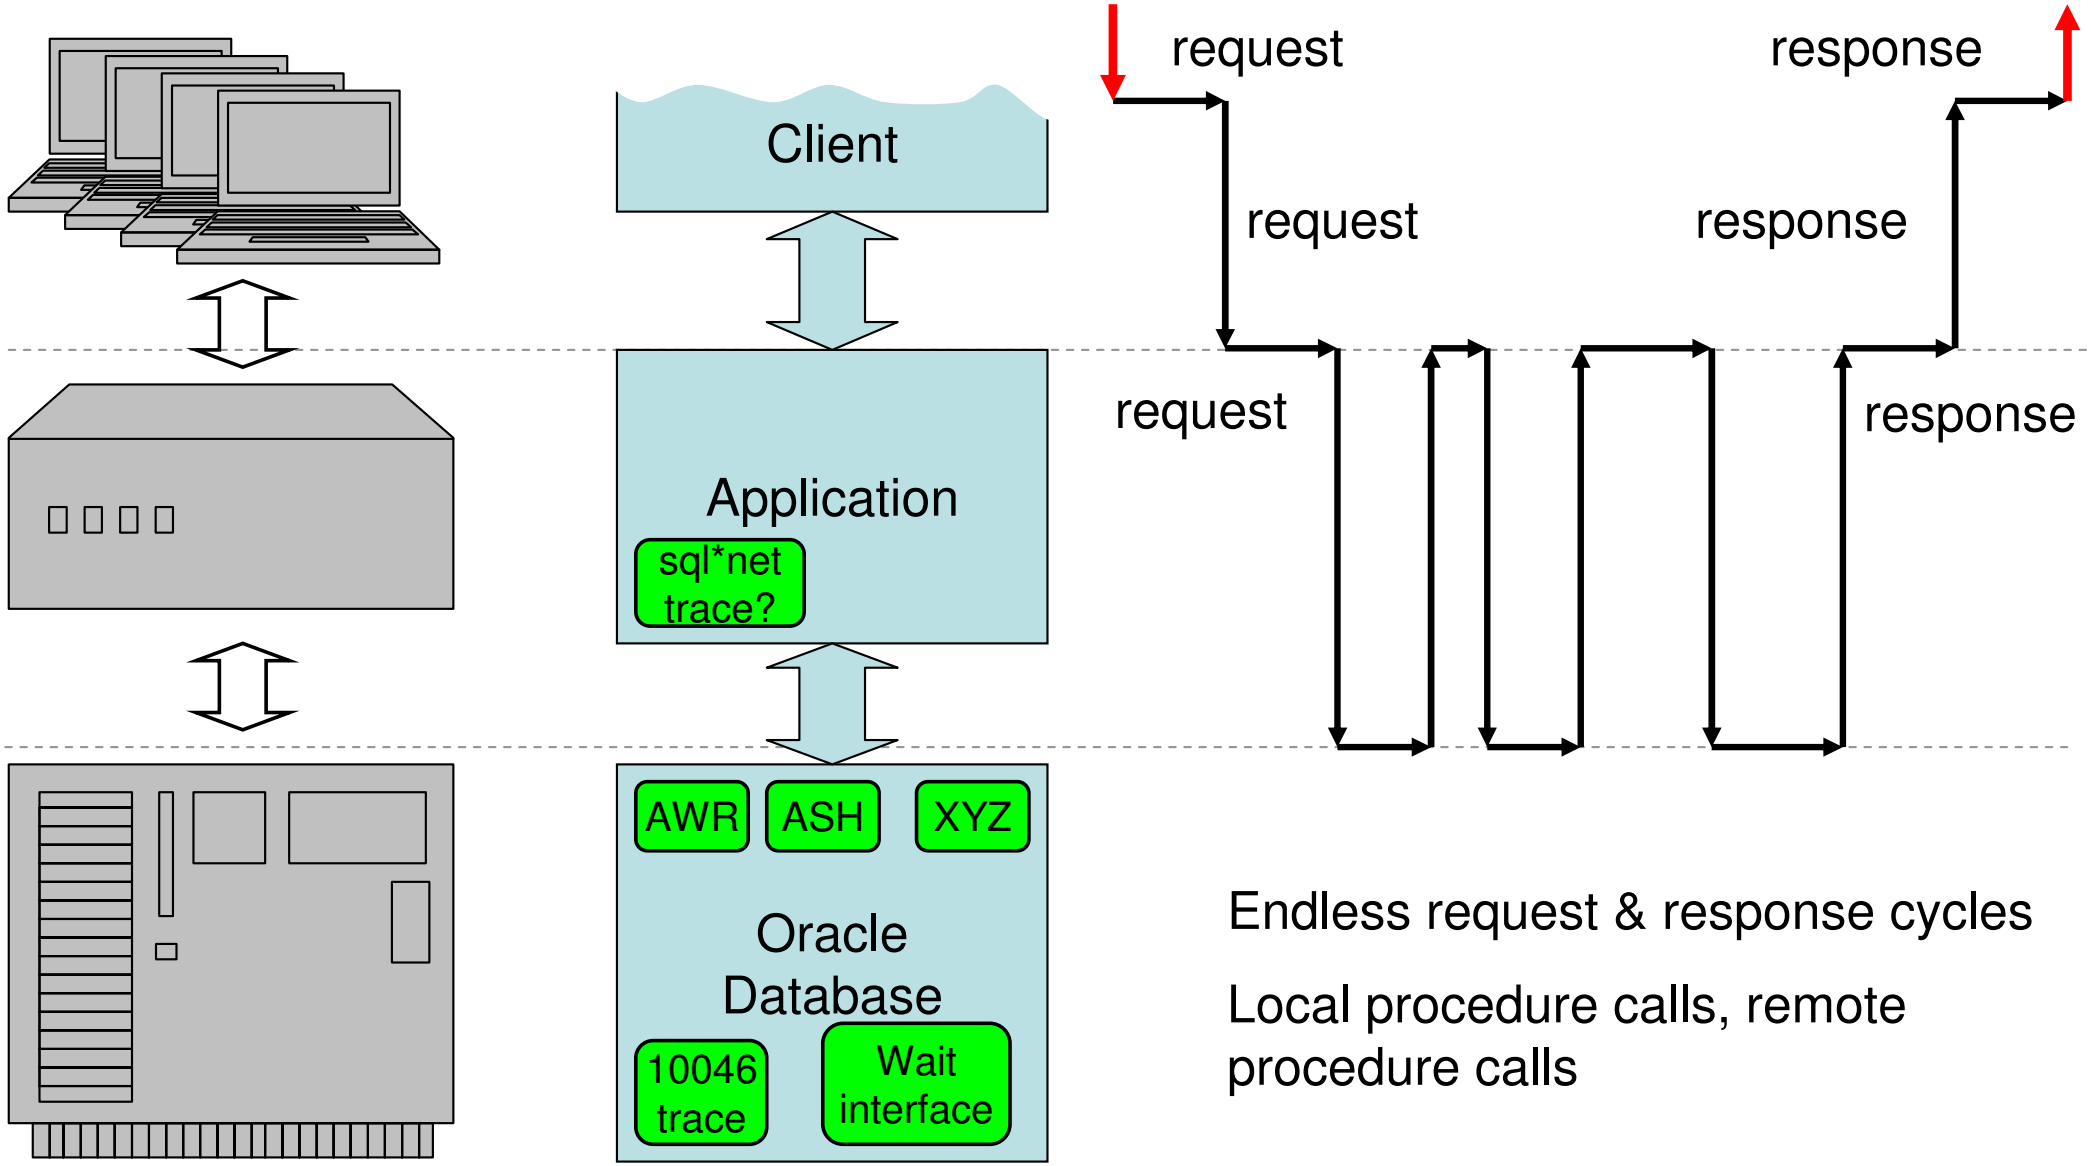

Endless request & response cycles

Local procedure calls, remote procedure calls

Understanding process flow

1. Application...
   a. ...waits for a request from a client
   b. ...issues SQL statements to a database and waits for result
   c. ...processes the SQL results
   d. ...returns processed results to client

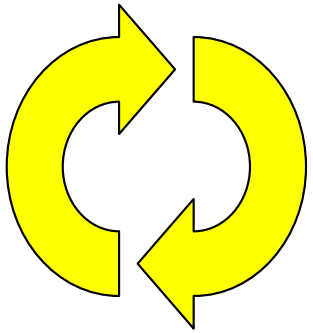

2. Database...
   a. ...waits for a request from an application
   b. ...issues physical IO calls to OS and waits for result
   c. ...processes the result data blocks
   d. ...returns processed results to application

Understanding process flow

3. OS...
   a. ...waits for a request from a database
   b. ...issues device driver calls to control hardware controller and waits for result
   c. ...processes the hardware access routine results
   d. ...returns processed results to database

4. Hardware controller...
   a. ...waits for a request from the OS
   b. ...sends (electric) signals to actual hardware and waits for result
   c. ...processes the result data
   d. ...returns processed results to OS

Oracle internal process flow

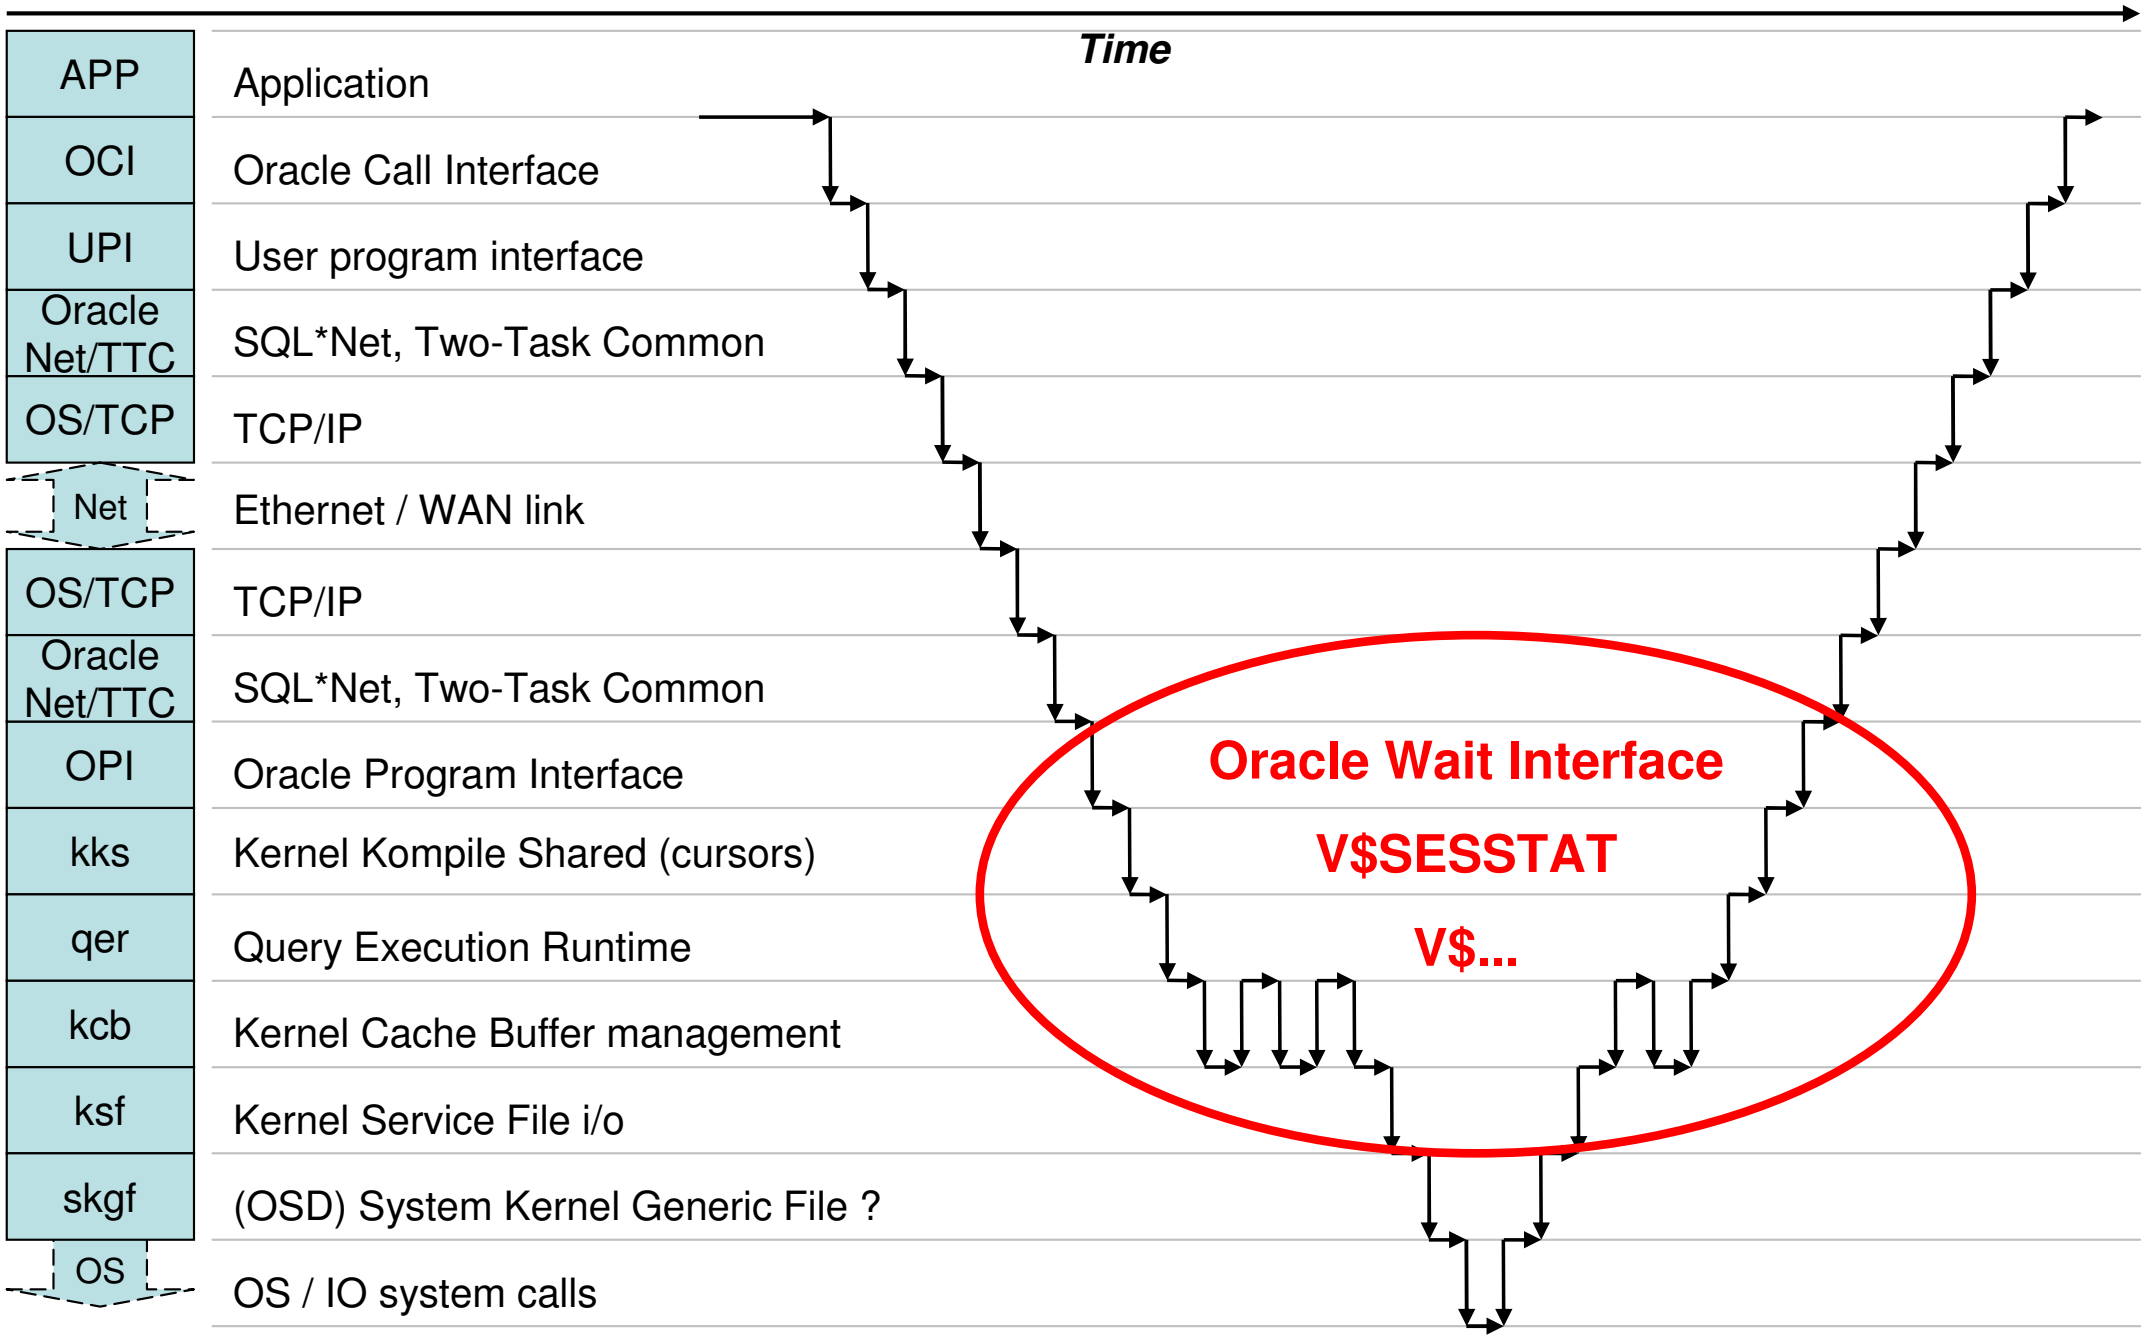

Oracle internal process flow

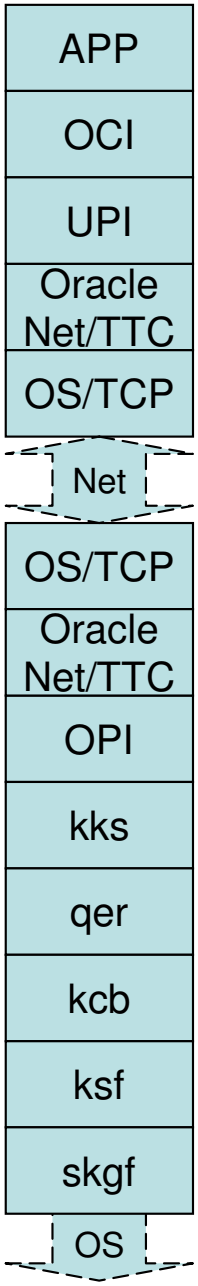

Application	***Application instrumentation***, ltrace, truss -u"libclntsh:*"
Oracle Call Interface	$OH/rdbms/demo/ociucb.mk, OCITrace
User program interface	-
SQL*NET, TNS, Two-Task Common	SQL*Net trace, Wireshark TNS protocol digester
TCP/IP	Wireshark TCP protocol digester
Ethernet / WAN link	snoop, tcpdump, Wireshark
TCP/IP	Wireshark TCP protocol digester
SQL*NET, TNS, Two-Task Common	SQL*Net trace, Wireshark, Event 10079
Oracle Program Interface	Event 10051
Kernel Kompile Shared (cursors)	sql_trace, Event 10046, 10270
Query Execution Runtime	v$sql_plan_statistics, v$sql_plan_statistics_all, sql_trace
Kernel Cache Buffer management	x$kcbsw, Event 10200,10298,10812, _trace_pin_time
Kernel Service File i/o	v$filestat, v$tempstat, v$session_wait, Event 10298
(OSD) System Kernel Generic File ?	-
OS / IO system calls	strace, truss, tusc, filemon.exe, procmon.exe

Process stack demos

```
$ pstack 5855
#0  0x00c29402 in __kernel_vsyscall ()
#1  0x005509e4 in semtimedop () from /lib/libc.so.6
#2  0x0e5769b7 in sskgpwwait ()
#3  0x0e575946 in skgpwwait ()
#4  0x0e2c3adc in ksliwat ()
#5  0x0e2c3449 in kslwaitctx. ()
#6  0x0b007261 in kjusuc ()
#7  0x0c8a7961 in ksipgetctx ()
#8  0x0e2d4dec in ksqcmi ()
#9  0x0e2ce9b8 in ksqgtlctx ()
#10 0x0e2cd214 in ksqgelctx. ()
#11 0x08754afa in ktcwit1 ()
#12 0x0e39b2a8 in kdddgb ()
#13 0x08930c80 in kdddel ()
#14 0x0892af0f in kaudel ()
#15 0x08c3d21a in delrow ()
#16 0x08e6ce16 in qerdlFetch ()
#17 0x08c403c5 in delexe ()
#18 0x0e3c3fa9 in opiexe ()
#19 0x08b54500 in kpoal8 ()
#20 0x0e3be673 in opiodr ()
#21 0x0e53628a in ttcpip ()
#22 0x089a87ab in opitsk ()
#23 0x089aaa00 in opiino ()
#24 0x0e3be673 in opiodr ()
#25 0x089a4e76 in opidrv ()
#26 0x08c1626f in sou2o ()
#27 0x08539aeb in opimai_real ()
#28 0x08c19a42 in ssthrdmain ()
#29 0x08539a68 in main ()
```

175982.1 ORA-600 Lookup Error Categories

453521.1 ORA-04031 “KSFQ Buffers” ksmlgpalloc

@d.sql - Report data dictionary & X$ tables

@pd.sql - Parameter descriptions

@la.sql - Latch by address

@lm.sql - Latch Misses by function location

@fv.sql - Fixed variable by name

@fva.sql - Fixed variable by address

Reading SQL plan execution stack

- os_explain script
- Uses pstack to get process execution stack
- Translates function names into execution plan step names
  - As an Oracle SQL plan execution means that just a bunch of row-source functions are executed in defined order
  - The order definition (in form of set of function pointers stored in library cache) is the execution plan
- Uses information from Metalink:
  - 175982.1 ORA-600 Lookup Error Categories
- Demo

What if my problem lies outside Oracle?

…Where to look next?

Oracle external process flow

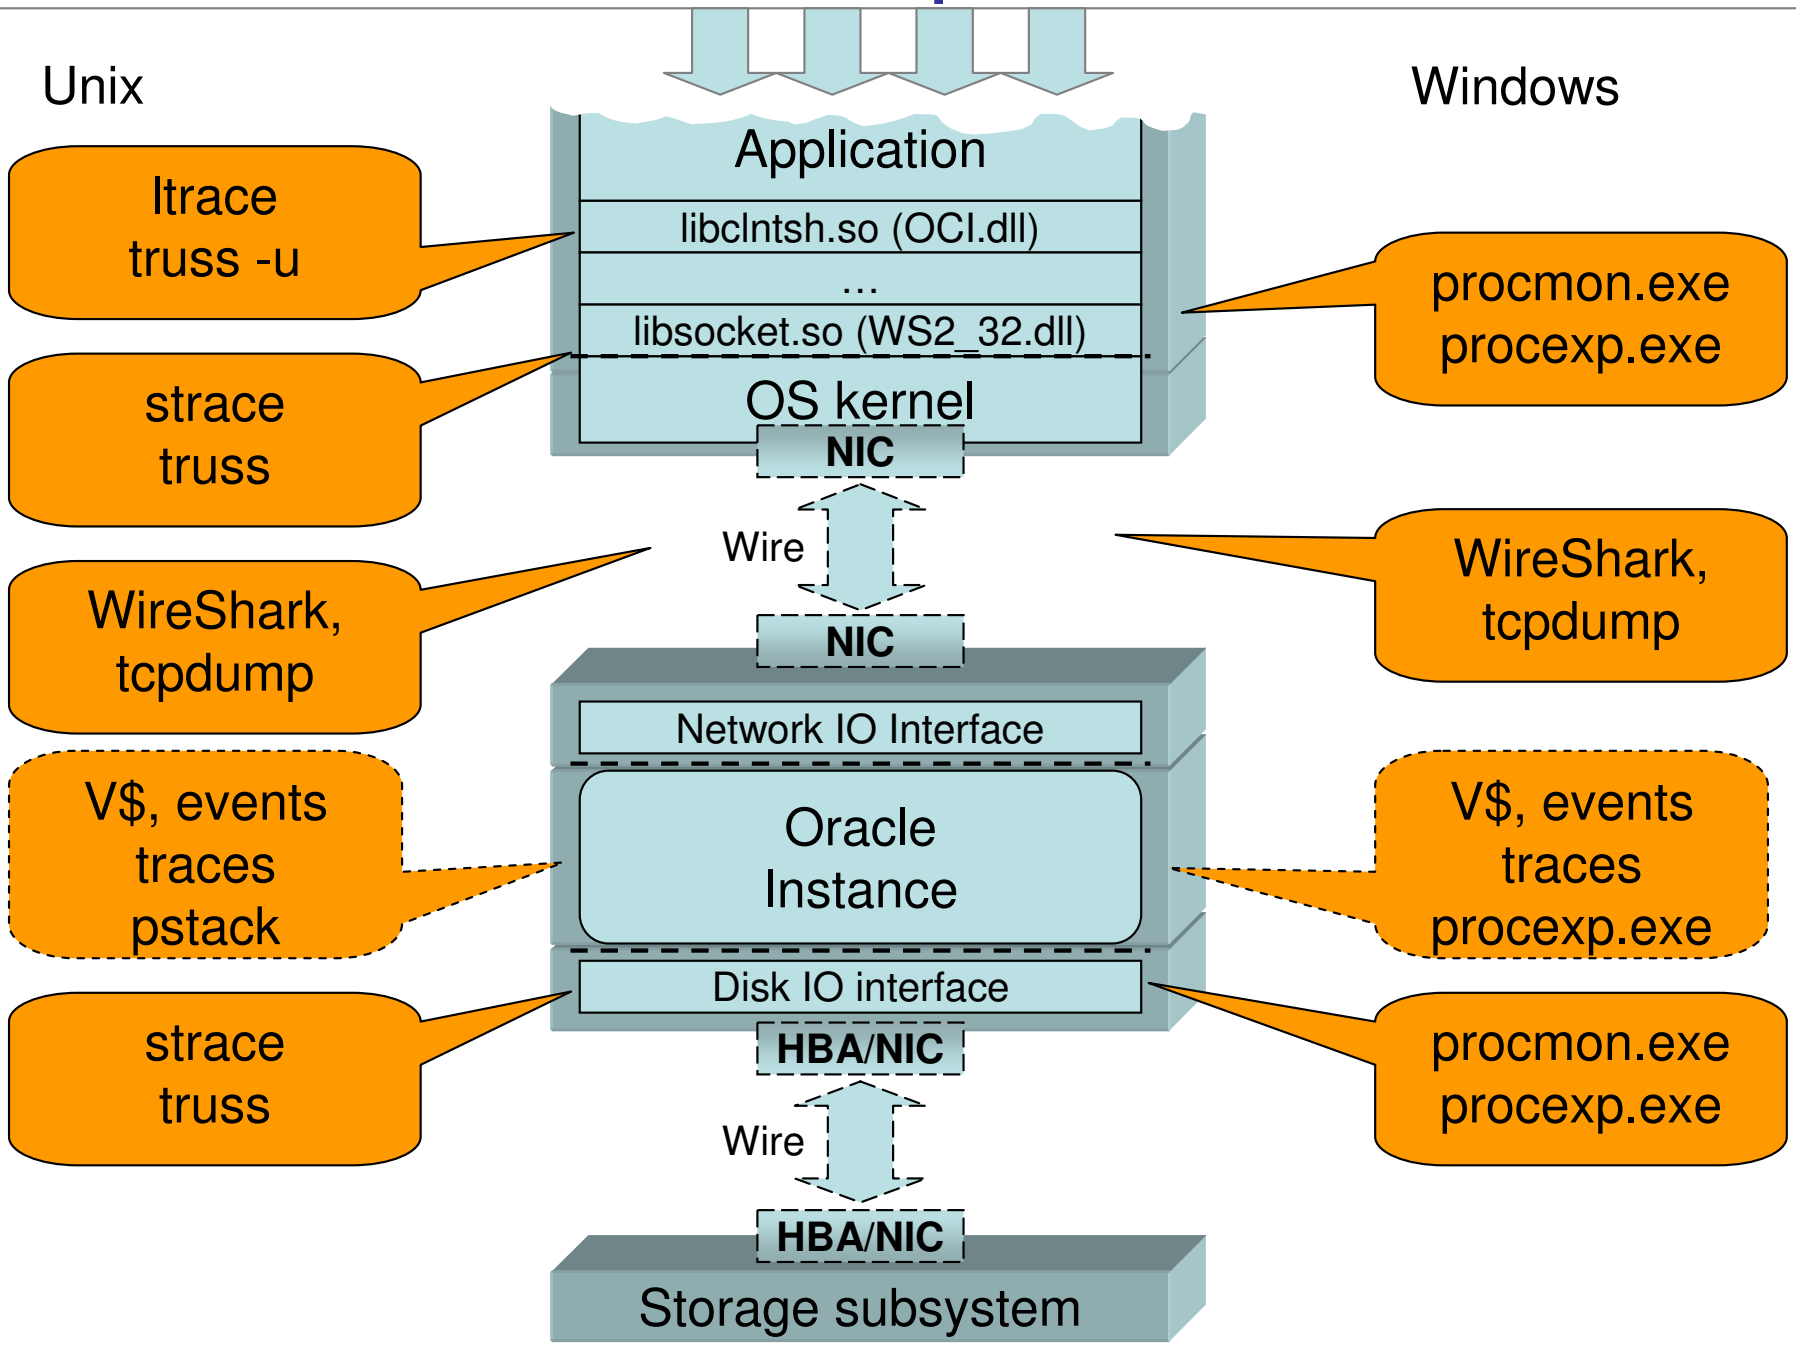

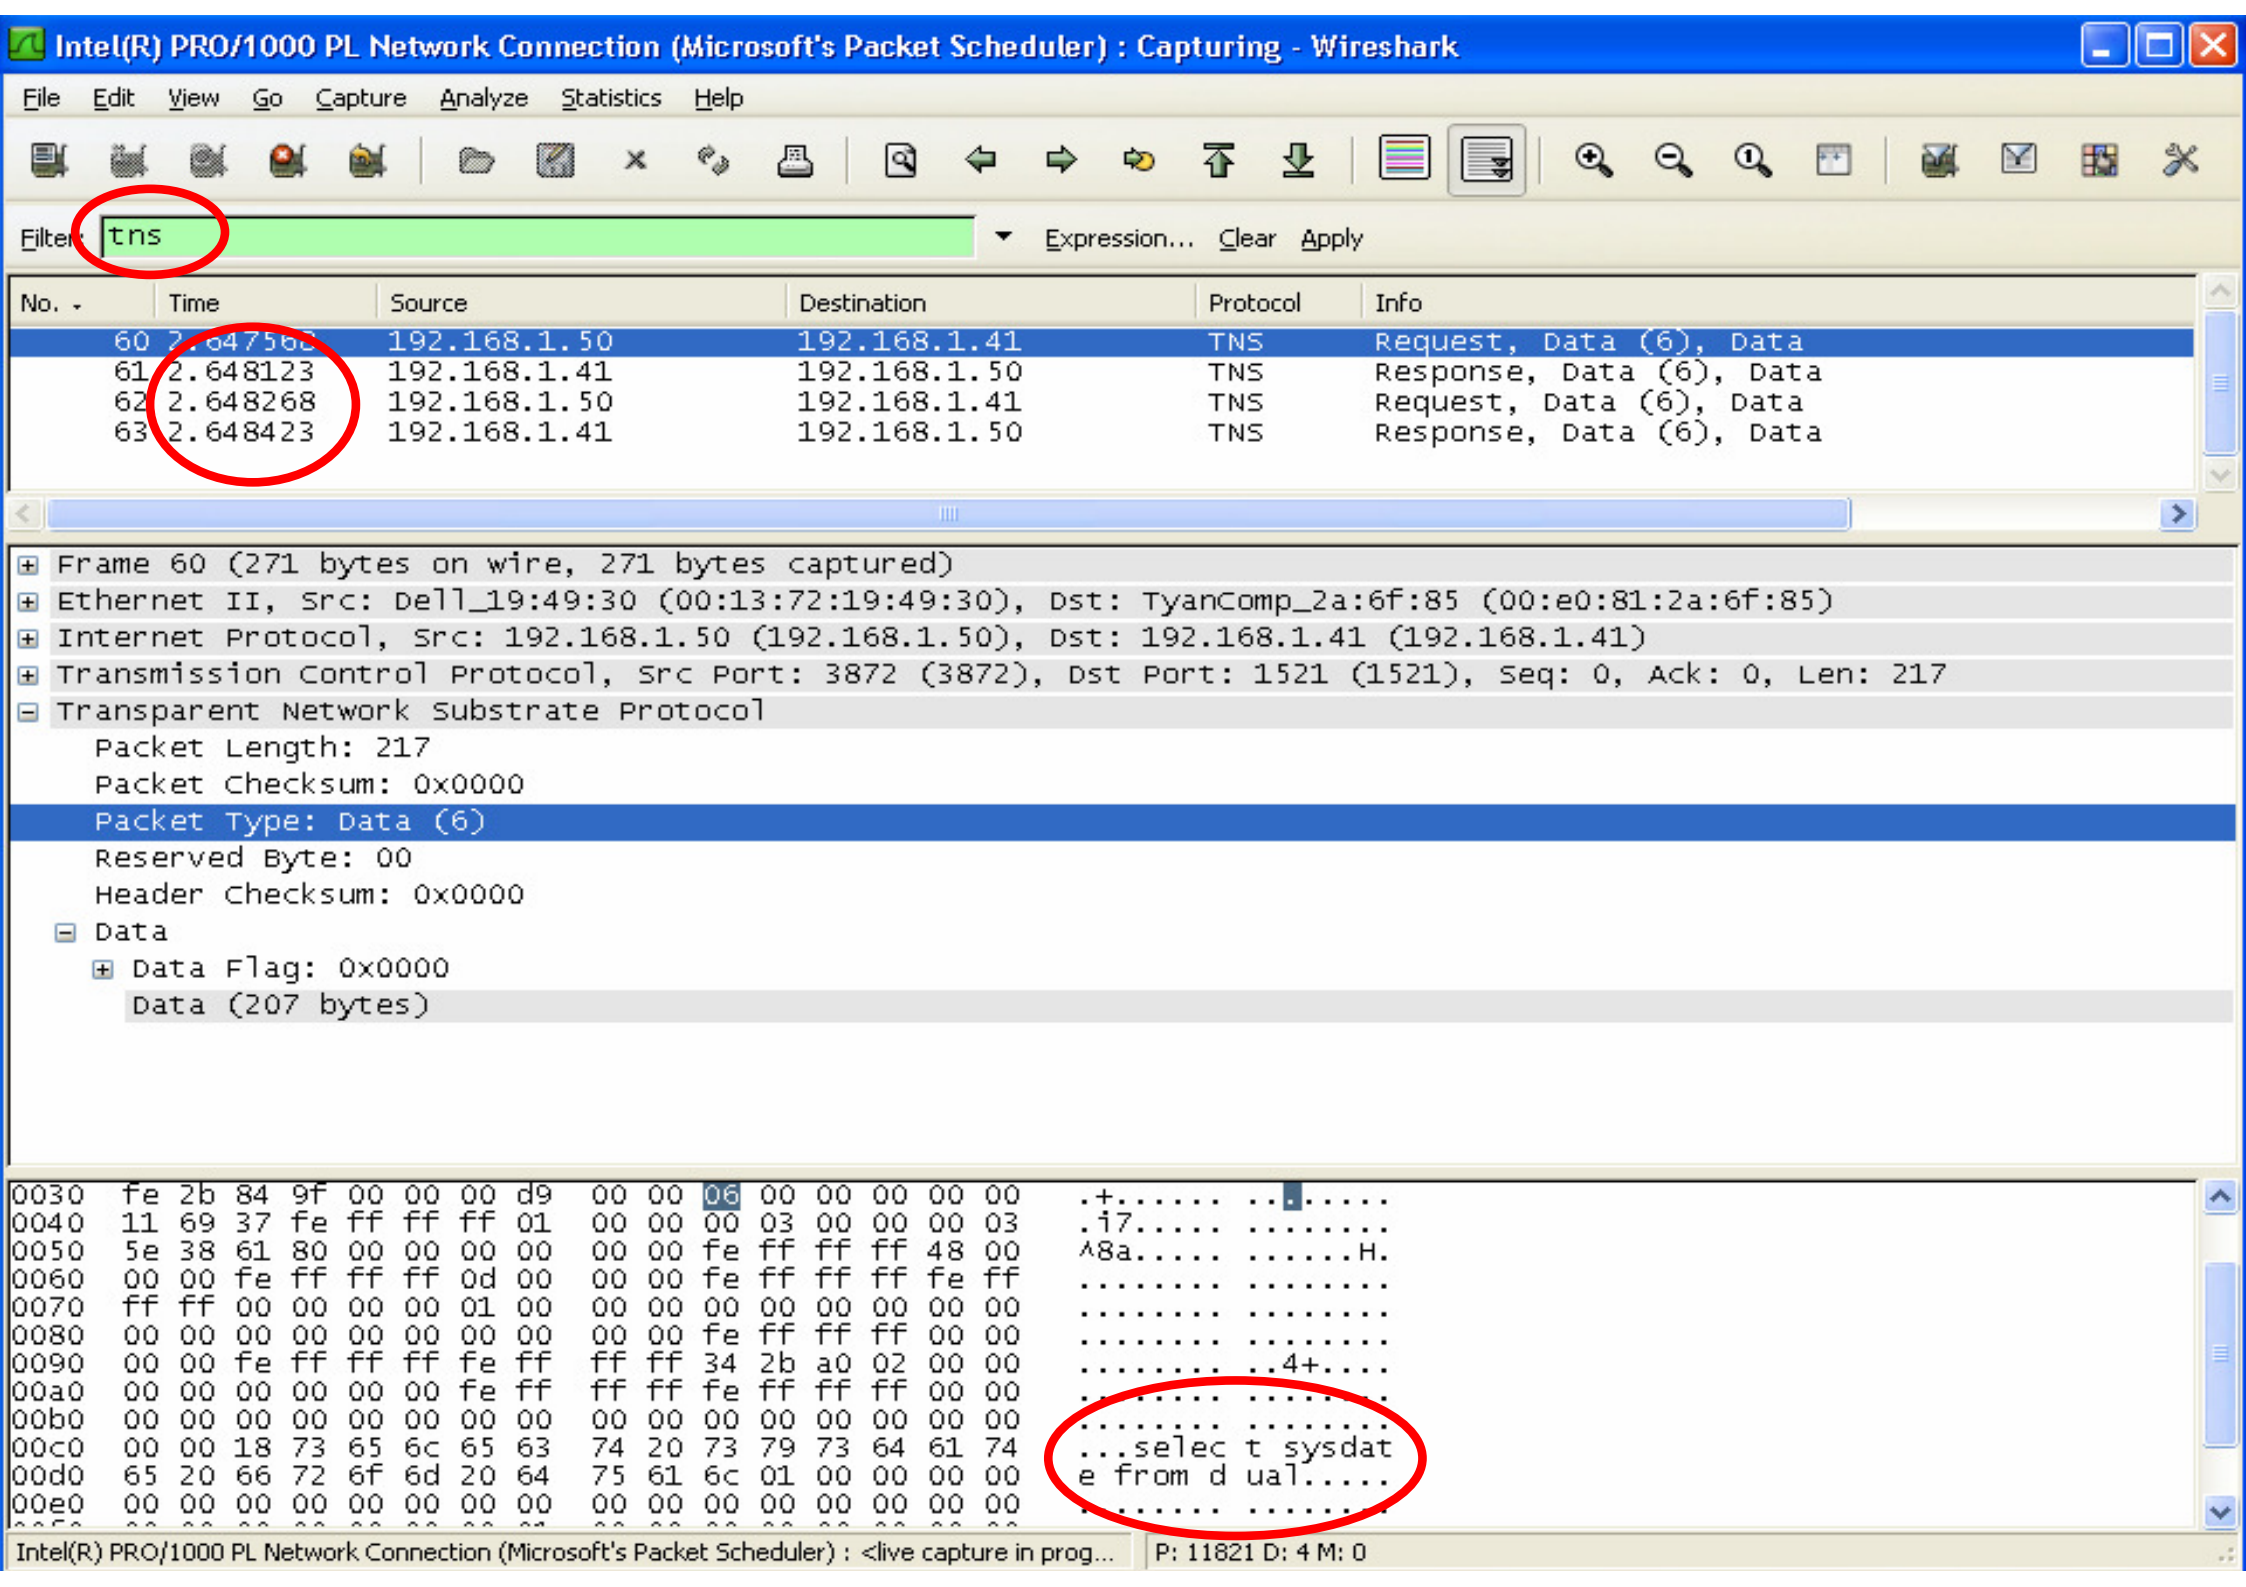
Intel(R) PRO/1000 PL Network Connection (Microsoft's Packet Scheduler) : Capturing - Wireshark
File Edit View Go Capture Analyze Statistics Help
Filter: tns
Expression... Clear Apply
No. Time Source Destination Protocol Info
60 2.647568 192.168.1.50 192.168.1.41 TNS Request, Data (6), Data
61 2.648123 192.168.1.41 192.168.1.50 TNS Response, Data (6), Data
62 2.648268 192.168.1.50 192.168.1.41 TNS Request, Data (6), Data
63 2.648423 192.168.1.41 192.168.1.50 TNS Response, Data (6), Data
Frame 60 (271 bytes on wire, 271 bytes captured)
Ethernet II, Src: Dell_19:49:30 (00:13:72:19:49:30), Dst: TyanComp_2a:6f:85 (00:e0:81:2a:6f:85)
Internet Protocol, Src: 192.168.1.50 (192.168.1.50), Dst: 192.168.1.41 (192.168.1.41)
Transmission Control Protocol, Src Port: 3872 (3872), Dst Port: 1521 (1521), Seq: 0, Ack: 0, Len: 217
Transparent Network Substrate Protocol
Packet Length: 217
Packet Checksum: 0x0000
Packet Type: Data (6)
Reserved Byte: 00
Header Checksum: 0x0000
Data
Data Flag: 0x0000
Data (207 bytes)
0030 fe 2b 84 9f 00 00 00 d9 00 00 06 00 00 00 00 00 .+...... ....
0040 11 69 37 fe ff ff ff 01 00 00 00 03 00 00 00 03 .i7..... ....
0050 5e 38 61 80 00 00 00 00 00 00 fe ff ff ff 48 00 ^8a..... ......H.
0060 00 00 fe ff ff ff 0d 00 00 00 fe ff ff ff fe ff ....
0070 ff ff 00 00 00 00 01 00 00 00 00 00 00 00 00 00 ....
0080 00 00 00 00 00 00 00 00 00 00 fe ff ff ff 00 00 ....
0090 00 00 fe ff ff ff fe ff ff ff 34 2b a0 02 00 00 ....4+....
00a0 00 00 00 00 00 00 fe ff ff ff fe ff ff ff 00 00 ....
00b0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ....
00c0 00 00 18 73 65 6c 65 63 74 20 73 79 73 64 61 74 ...selec t sysdat
00d0 65 20 66 72 6f 6d 20 64 75 61 6c 01 00 00 00 00 e from d ual.....
00e0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ....
Intel(R) PRO/1000 PL Network Connection (Microsoft's Packet Scheduler) : <live capture in prog...
P: 11821 D: 4 M: 0

What if I need to look further inside Oracle

...if standard Oracle instrumentation isn't detailed enough...

…OS tools don't understand Oracle internal workings

...only for experimental environments

IO tracing events

```
10200, 00000, "consistent read buffer status"
// *Cause:
// *Action:

alter session set "_trace_pin_time" = 1;
// trace how long a current pin is held

10812, 00000, "Trace Consistent Reads" ( Trace into X$TRACE )
// *Cause: N/A
// *Action: THIS IS NOT A USER ERROR NUMBER/MESSAGE.  THIS DOES NOT
// NEED TO BE TRANSLATED OR DOCUMENTED. IT IS USED ONLY FOR DEBUGGING.

10298, 00000, "ksfd i/o tracing"
// *Cause:
// *Action: If this event is set then ksfd module generates tracing
//          for each i/o request
```

Cursor usage tracing events

```
10270, 00000, "Debug shared cursors"
// *Cause: Enables debugging code in shared cursor management modules
// *Action:

10730, 00000, "trace row level security policy predicates"
// *Document: NO
// *Cause:
// *Action:
// *Comment:

10731, 00000, "dump SQL for CURSOR expressions"
// *Cause:
// *Action: set this event only under the supervision of Oracle development
// *Comment: traces SQL statements generated to execute CURSOR expressions
```

alter session set "_dump_qbc_tree" = 1; *(10.2+)*

```
// dump top level query parse tree to trace
```

Network / user call tracing events

```
10051, 00000, "trace OPI calls"
// *Cause:
// *Action:


10079, 00000, "trace data sent/received via SQL*Net"
// *Cause:
// *Action: level 1 - trace network ops to/from client
//          level 2 - in addition to level 1, dump data
//          level 4 - trace network ops to/from dblink
//          level 8 - in addition to level 4, dump data
```

Right tools for right problems

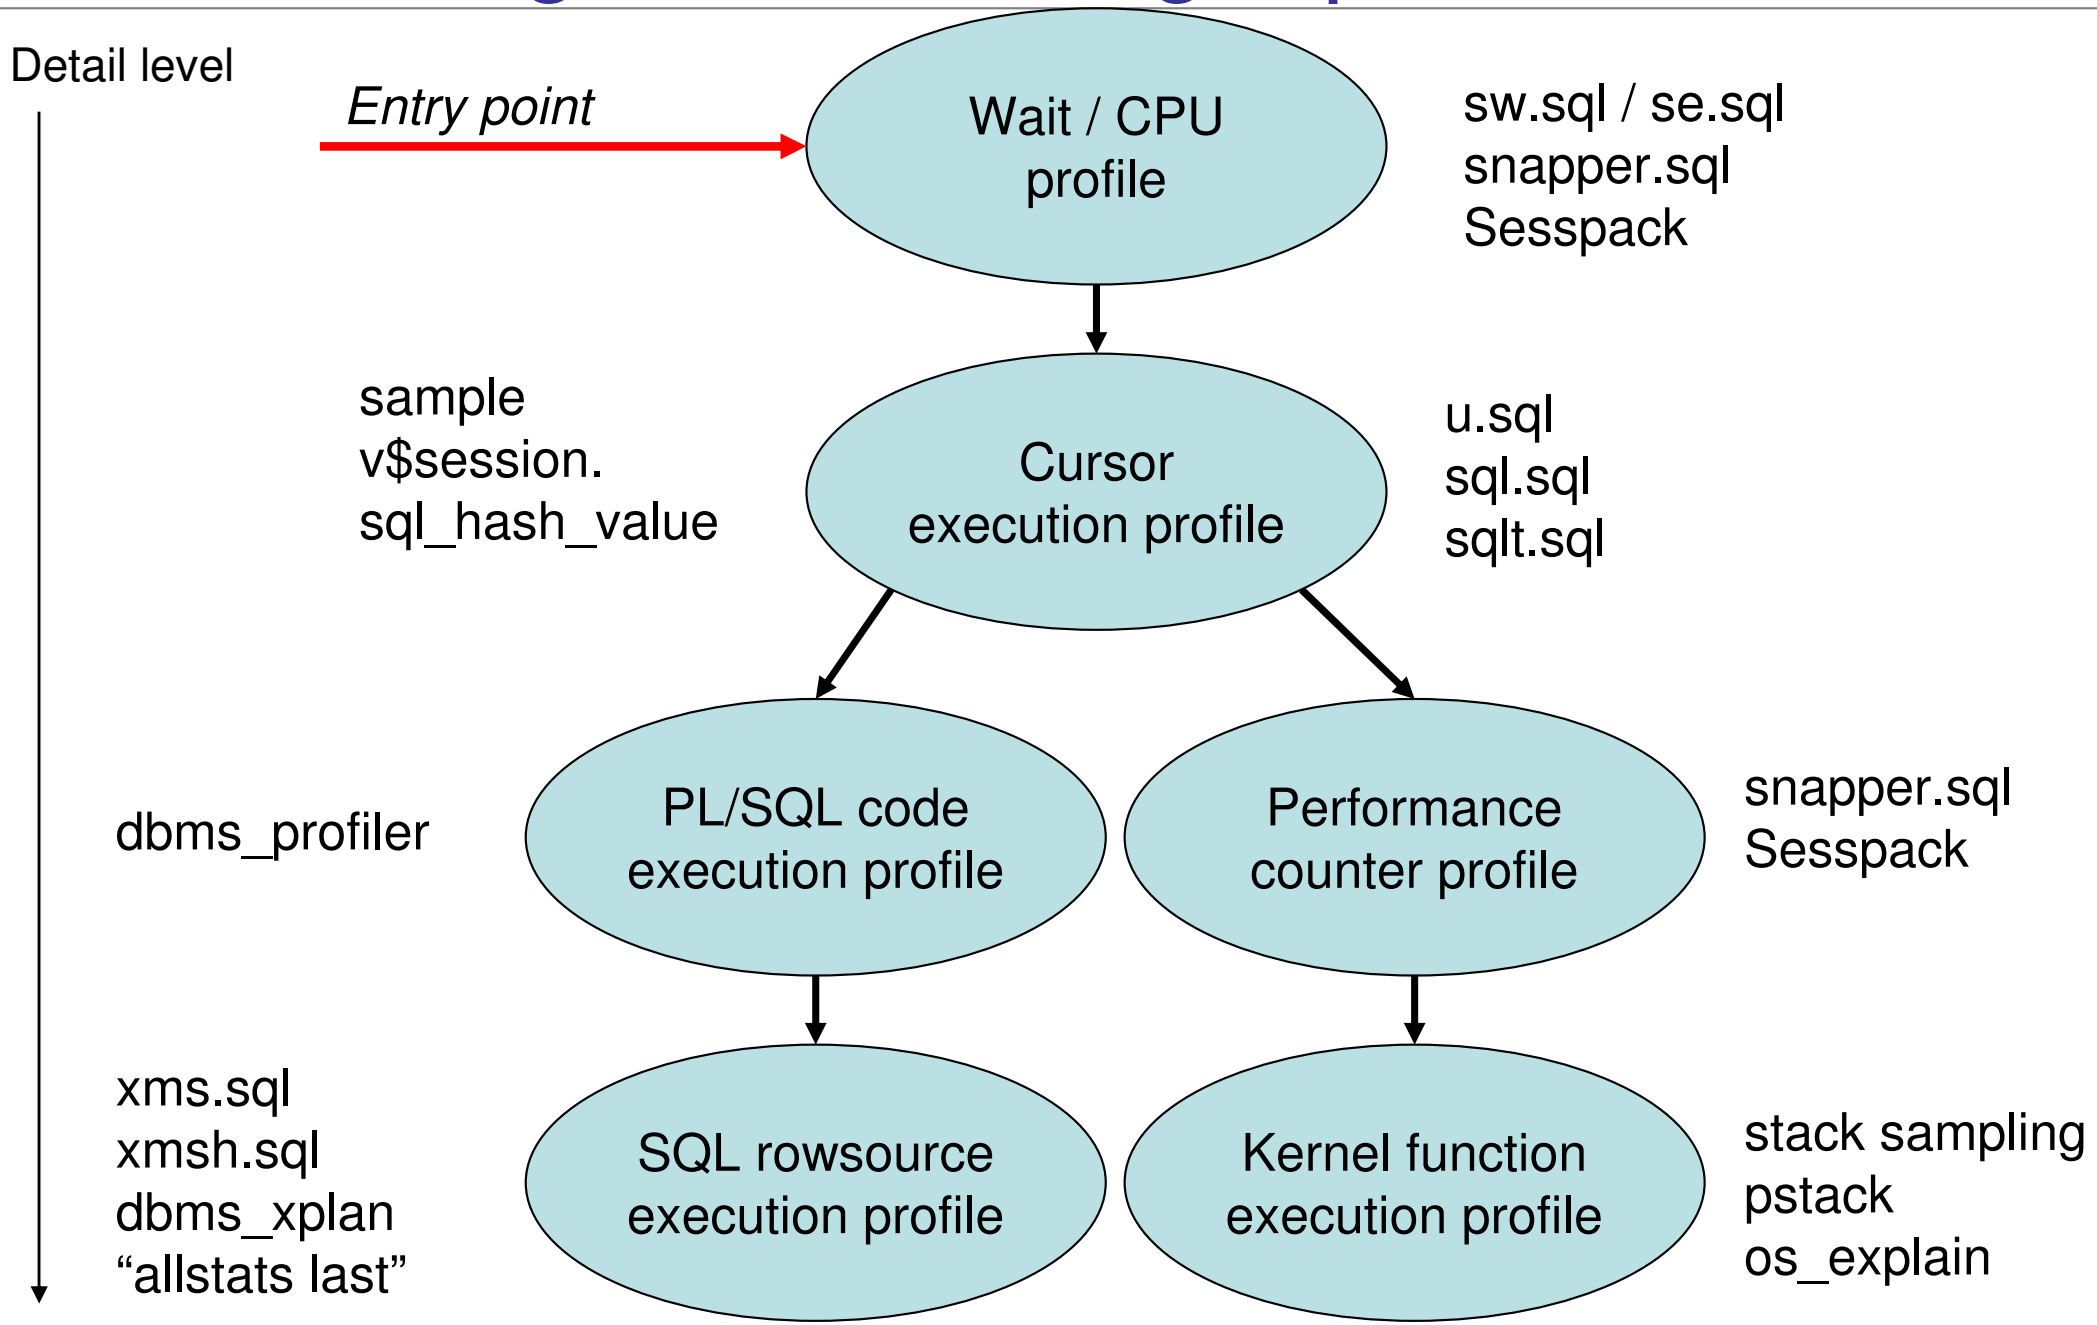

Questions?

Further questions welcome at
http://blog.tanelpoder.com

Thank you!

Tanel Põder

http://www.tanelpoder.com