Mastering OpenEdge Query Performance

Best Practices and Tools for Optimizing Database Queries

White Star software

White Star Software

For over 30 years, we have been helping companies around the world simplify the job of managing and maintaining the world's best OpenEdge applications.

Our experts, combined with ProTop, the leading OpenEdge monitoring and alerting tool, deliver unparalleled peace of mind for your OpenEdge environments.

info@wss.com | wss.com

The speaker

Paul Koufalis (pk@wss.com)

- Speaker/teacher/OpenEdge DBA since 1994

Agenda

- ❑ First thoughts
- ❑ Tools to find inefficient code
- ❑ Improving queries
- ❑ Q&A

White Star
software

First thoughts

"The system is slow"

- We've all heard this from our end users
- Often code/query related. Sometimes not
- For this presentation we will assume that it is code
- …but more on this later

White Star software

Second thoughts…

- There are many ways to find the bad code
  - Some kind of monitoring and trending
  - Hopefully you attended our session on ProTop
- We will show ProTop RT screen captures
  - 100% free download

White Star
software

Our villains

- "churn.p": A program that reads the very small vacation table a very large number of times
- "paul_orders.p": a program that reads through all orders of all customers
- Both using a "big" sports2000 DB

- HOW DO WE FIND THEM?

Using ProTop RT

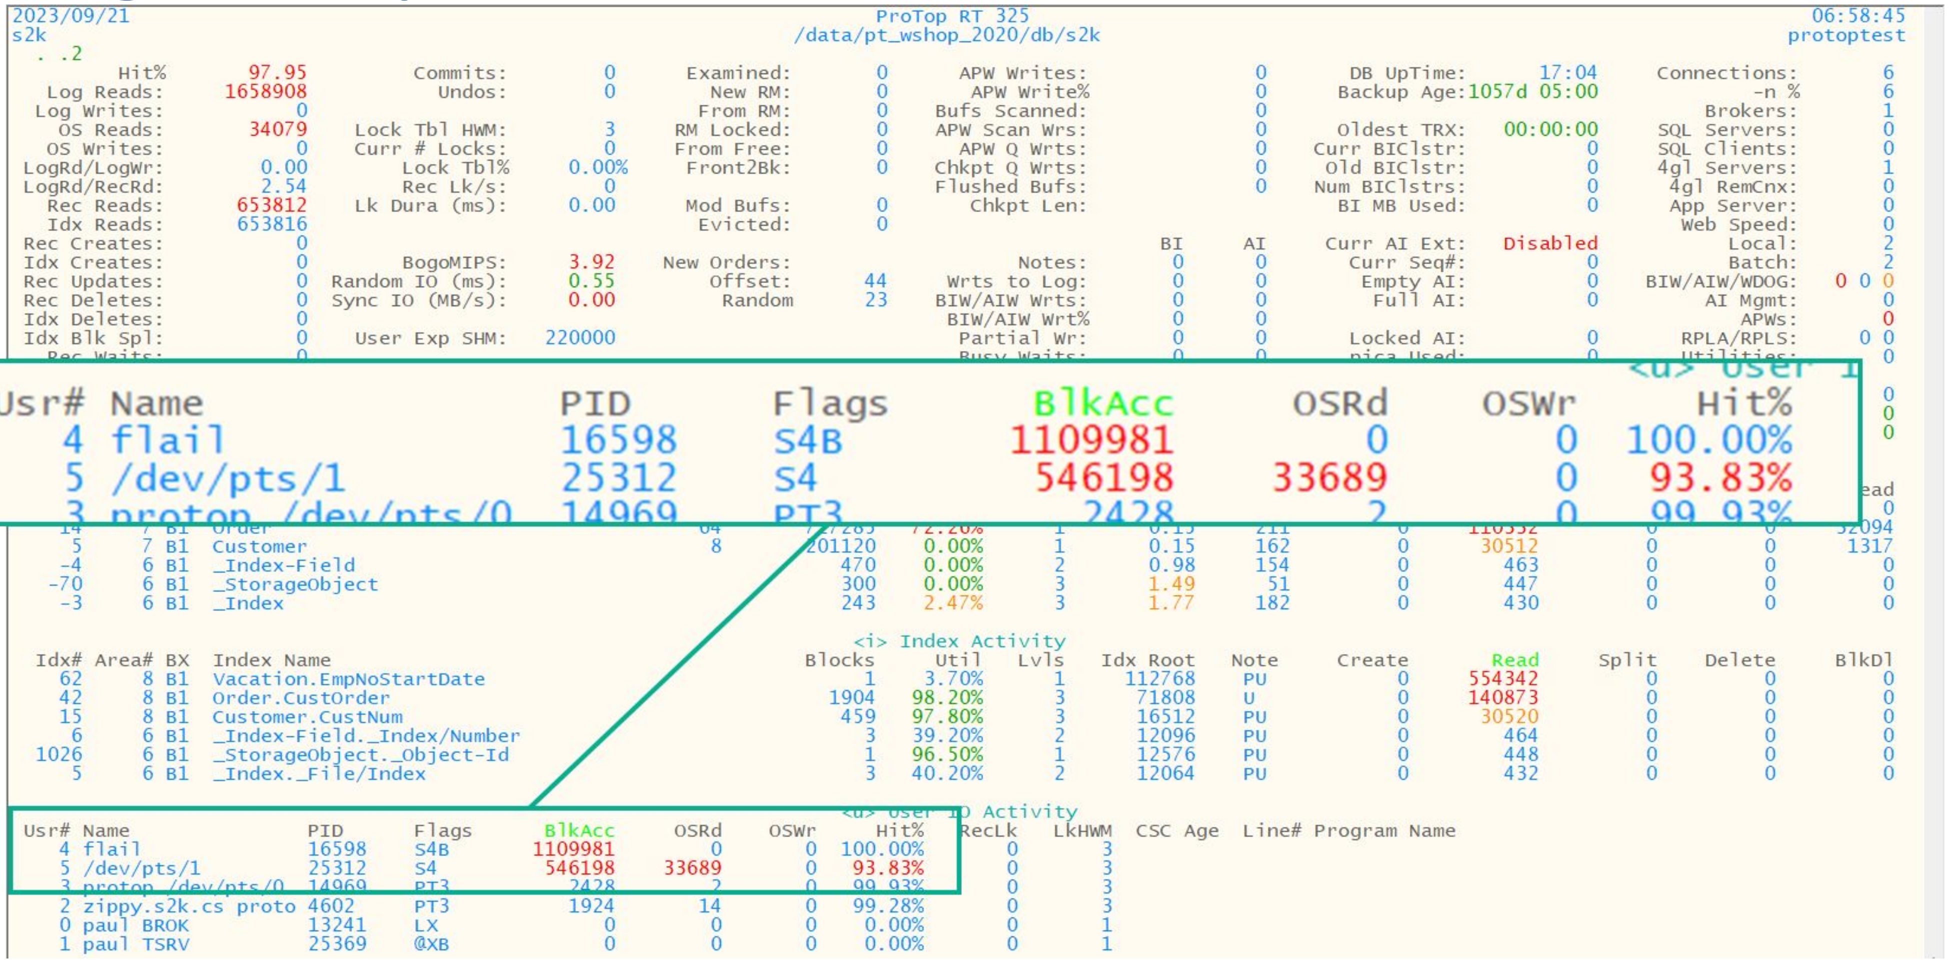

```
2023/09/21                                  ProTop RT 325                                         06:58:45
s2k                                  /data/pt_wshop_2020/db/s2k                                   protoptest
  . .2
          Hit%       97.95        Commits:        0       Examined:       0     APW Writes:              0     DB UpTime:       17:04     Connections:        6
     Log Reads:    1658908          Undos:        0         New RM:       0      APW Write%              0   Backup Age:1057d 05:00            -n %        6
    Log Writes:          0                                  From RM:       0   Bufs Scanned:              0                                   Brokers:        1
      OS Reads:      34079   Lock Tbl HWM:        3       RM Locked:       0   APW Scan Wrs:              0     Oldest TRX:    00:00:00    SQL Servers:        0
     OS Writes:          0   Curr # Locks:        0       From Free:       0     APW Q Wrts:              0    Curr BIClstr:          0    SQL Clients:        0
   LogRd/LogWr:       0.00      Lock Tbl%    0.00%        Front2Bk:       0    Chkpt Q Wrts:             0     Old BIClstr:          0    4gl Servers:        1
   LogRd/RecRd:       2.54      Rec Lk/s:         0                              Flushed Bufs:           0    Num BIClstrs:          0     4gl RemCnx:        0
     Rec Reads:     653812   Lk Dura (ms):     0.00       Mod Bufs:       0      Chkpt Len:                    BI MB Used:          0    App Server:         0
     Idx Reads:     653816                                 Evicted:       0                                                               Web Speed:        0
   Rec Creates:          0                                                                     BI       AI    Curr AI Ext:    Disabled          Local:        2
   Idx Creates:          0       BogoMIPS:     3.92    New Orders:                  Notes:      0        0     Curr Seq#:          0           Batch:        2
   Rec Updates:          0   Random IO (ms):   0.55        Offset:      44     Wrts to Log:      0        0      Empty AI:          0    BIW/AIW/WDOG:   0 0 0
   Rec Deletes:          0   Sync IO (MB/s):   0.00        Random       23   BIW/AIW Wrts:       0        0       Full AI:          0        AI Mgmt:        0
   Idx Deletes:          0                                                    BIW/AIW Wrt%       0        0                                      APWs:        0
   Idx Blk Spl:          0      User Exp SHM:   220000                          Partial Wr:      0        0     Locked AI:          0     RPLA/RPLS:      0 0
                                                                                                                                                             0
                                                                                                                                                             0
                                                                                                                                                             0
                                                                                                                                                             0
                                                                                                                                                           ead
                                                                                                                                                             0
                                                                                                                                                         32094
     5     7 B1  Customer                                   8  201120   0.00%     1      0.15      162        0      30512        0          0          1317
    -4     6 B1  _Index-Field                                     470   0.00%     2      0.98      154        0        463        0          0             0
   -70     6 B1  _StorageObject                                   300   0.00%     3      1.49       51        0        447        0          0             0
    -3     6 B1  _Index                                           243   2.47%     3      1.77      182        0        430        0          0             0

                                                       <i> Index Activity
  Idx# Area# BX  Index Name                              Blocks    Util  Lvls  Idx Root   Note    Create      Read     Split    Delete     BlkDl
    62     8 B1  Vacation.EmpNoStartDate                      1   3.70%     1    112768    PU          0    554342         0         0         0
    42     8 B1  Order.CustOrder                           1904  98.20%     3     71808    U           0    140873         0         0         0
    15     8 B1  Customer.CustNum                           459  97.80%     3     16512    PU          0     30520         0         0         0
     6     6 B1  _Index-Field._Index/Number                   3  39.20%     2     12096    PU          0       464         0         0         0
  1026     6 B1  _StorageObject._Object-Id                    1  96.50%     1     12576    PU          0       448         0         0         0
     5     6 B1  _Index._File/Index                           3  40.20%     2     12064    PU          0       432         0         0         0

                                                       <u> User IO Activity
 Usr# Name                  PID     Flags     BlkAcc     OSRd     OSWr      Hit%   RecLk   LkHWM  CSC Age  Line# Program Name
    4 flail                 16598   S4B      1109981        0        0   100.00%       0       3
    5 /dev/pts/1            25312   S4        546198    33689        0    93.83%       0       3
    3 protop /dev/pts/0     14969   PT3         2428        2        0    99.93%       0       3
    2 zippy.s2k.cs proto    4602    PT3         1924       14        0    99.28%       0       3
    0 paul BROK             13241   LX             0        0        0     0.00%       0       1
    1 paul TSRV             25369   @XB            0        0        0     0.00%       0       1
```

```
Usr# Name                PID      Flags     BlkAcc     OSRd     OSWr      Hit%
   4 flail               16598    S4B      1109981        0        0   100.00%
   5 /dev/pts/1          25312    S4        546198    33689        0    93.83%
   3 protop /dev/pts/0   14969    PT3         2428        2        0    99.93%
```

Where did these numbers come from?

- _UserIO VST
- You will probably also want to sample _UserTablestat and _UserIndexstat VST
- Set correct -basetable, -baseindex, -tablerangesize and -indexrangesize
  - Default only capture first 50 tables and 50 indezes

White Star software

What are they doing?

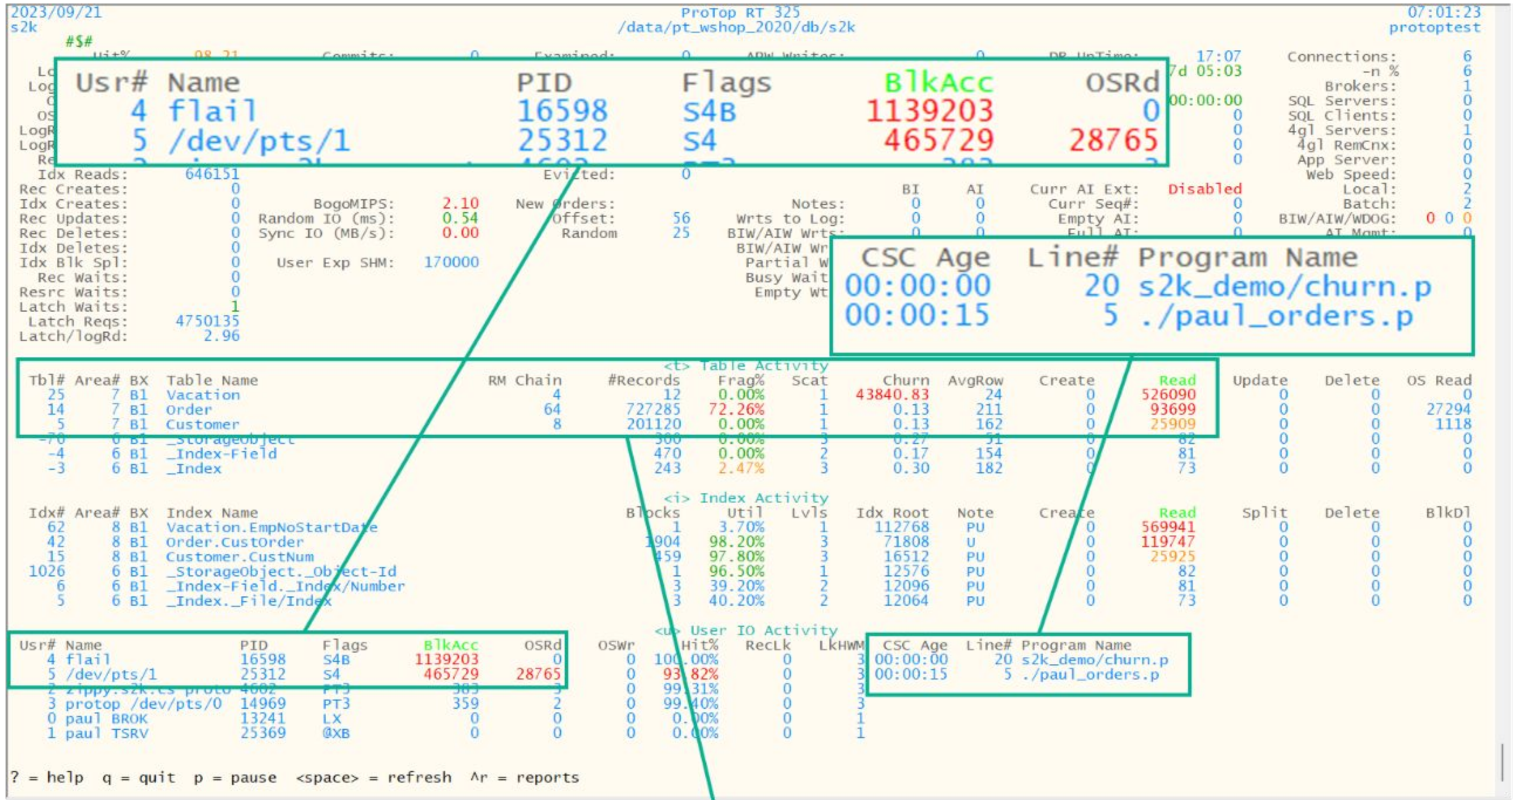

```
2023/09/21                                ProTop RT 325                                  07:01:23
s2k                                 /data/pt_wshop_2020/db/s2k                          protoptest
     #$#
  Idx Reads:     646151                                 Evicted:      0                                   Web Speed:     0
Rec Creates:          0                                                           BI    AI   Curr AI Ext:  Disabled      Local:     2
Idx Creates:          0        BogoMIPS:      2.10  New Orders:        Notes:     0     0    Curr Seq#:         0        Batch:     2
Rec Updates:          0  Random IO (ms):      0.54     Offset:    56  Wrts to Log:  0     0    Empty AI:          0  BIW/AIW/WDOG:  0 0 0
Rec Deletes:          0  Sync IO (MB/s):      0.00     Random     25  BIW/AIW Wrts: 0     0
Idx Deletes:          0                                               BIW/AIW Wr
Idx Blk Spl:          0    User Exp SHM:    170000                    Partial W
  Rec Waits:          0                                               Busy Wait
Resrc Waits:          0                                               Empty Wt
Latch Waits:          1
 Latch Reqs:    4750135
Latch/logRd:       2.96
                                                          <t> Table Activity
Tbl# Area# BX  Table Name                 RM Chain  #Records  Frag%  Scat     Churn AvgRow  Create    Read  Update  Delete  OS Read
  25     7 B1  Vacation                          4        12  0.00%     1  43840.83     24       0  526090       0       0        0
  14     7 B1  Order                            64    727285 72.26%     1      0.13    211       0   93699       0       0    27294
   5     7 B1  Customer                          8    201120  0.00%     1      0.13    162       0   25909       0       0     1118
 -70     6 B1  _StorageObject                            300  0.00%     3      0.27     51       0      82       0       0        0
  -4     6 B1  _Index-Field                              470  0.00%     2      0.17    154       0      81       0       0        0
  -3     6 B1  _Index                                    243  2.47%     3      0.30    182       0      73       0       0        0

                                                          <i> Index Activity
Idx# Area# BX  Index Name                            Blocks   Util  Lvls  Idx Root  Note  Create    Read   Split  Delete   BlkDl
  62     8 B1  Vacation.EmpNoStartDate                    1  3.70%     1    112768  PU         0  569941       0       0       0
  42     8 B1  Order.CustOrder                         1904 98.20%     3     71808  U          0  119747       0       0       0
  15     8 B1  Customer.CustNum                         459 97.80%     3     16512  PU         0   25925       0       0       0
1026     6 B1  _StorageObject._Object-Id                  1 96.50%     1     12576  PU         0      82       0       0       0
   6     6 B1  _Index-Field._Index/Number                 3 39.20%     2     12096  PU         0      81       0       0       0
   5     6 B1  _Index._File/Index                         3 40.20%     2     12064  PU         0      73       0       0       0

                                                          <u> User IO Activity
Usr# Name                  PID    Flags   BlkAcc   OSRd   OSWr    Hit%  RecLk  LkHWM  CSC Age   Line# Program Name
   4 flail                16598   S4B    1139203      0      0 100.00%      0      3  00:00:00     20 s2k_demo/churn.p
   5 /dev/pts/1           25312   S4      465729  28765      0  93.82%      0      3  00:00:15      5 ./paul_orders.p
   2 zippy.s2k.cs proto    4602   PT3        383      3      0  99.31%      0      3
   3 protop /dev/pts/0    14969   PT3        359      2      0  99.40%      0      3
   0 paul BROK            13241   LX           0      0      0   0.00%      0      1
   1 paul TSRV            25369   @XB          0      0      0   0.00%      0      1

? = help  q = quit  p = pause  <space> = refresh  ^r = reports
```

```
Usr# Name                  PID    Flags   BlkAcc   OSRd
   4 flail                16598   S4B    1139203      0
   5 /dev/pts/1           25312   S4      465729  28765
```

```
CSC Age   Line# Program Name
00:00:00     20 s2k_demo/churn.p
00:00:15      5 ./paul_orders.p
```

```
                                                   <t> Table Activity
Tbl# Area# BX  Table Name            RM Chain  #Records  Frag%  Scat     Churn AvgRow  Create    Read
  25     7 B1  Vacation                     4        12  0.00%     1  43840.83     24       0  526090
  14     7 B1  Order                       64    727285 72.26%     1      0.13    211       0   93699
   5     7 B1  Customer                     8    201120  0.00%     1      0.13    162       0   25909
```

How did ProTop get the program name?

- Enable Client Statement Cache
  - Use option “single”
- Can be enabled via promon or _connect VST

- NOTE:
  - CSC only records lines of code that touch the database
  - Only records new lines of code AFTER being enabled
  - Will not report on already running query: use proGetStack
- CAREFUL:
  - Could have performance impact, especially C/S

Tooling

Method #1: Use LOG-MANAGER

- Ex.: have an option in application to activate tracing
  - I.e. “Help - Trace App”, or a secret hotkey sequence

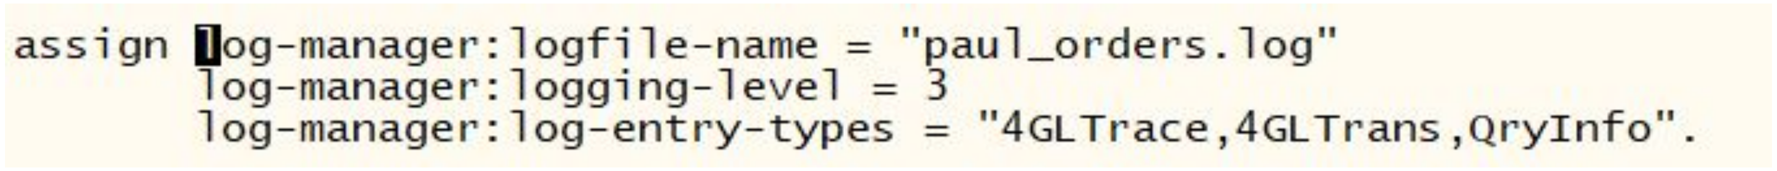

```
assign log-manager:logfile-name = "paul_orders.log"
       log-manager:logging-level = 3
       log-manager:log-entry-types = "4GLTrace,4GLTrans,QryInfo".
```

Output

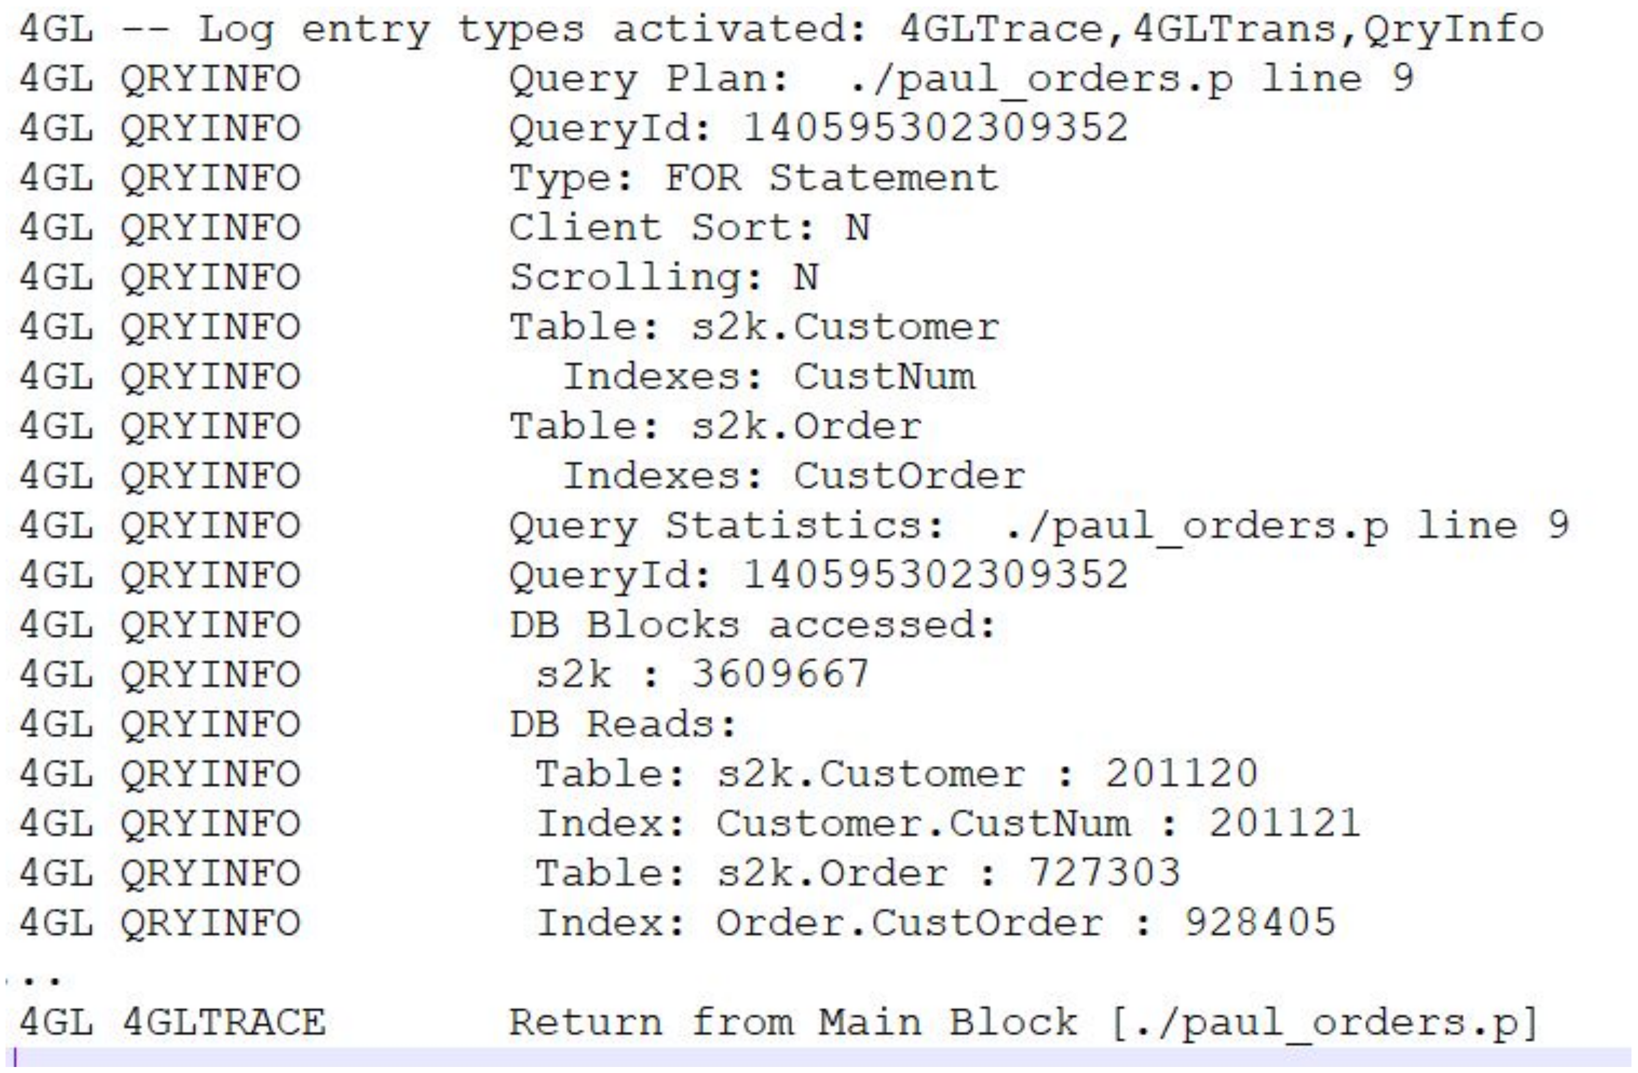

```
4GL -- Log entry types activated: 4GLTrace,4GLTrans,QryInfo
4GL QRYINFO        Query Plan:  ./paul_orders.p line 9
4GL QRYINFO        QueryId: 140595302309352
4GL QRYINFO        Type: FOR Statement
4GL QRYINFO        Client Sort: N
4GL QRYINFO        Scrolling: N
4GL QRYINFO        Table: s2k.Customer
4GL QRYINFO          Indexes: CustNum
4GL QRYINFO        Table: s2k.Order
4GL QRYINFO          Indexes: CustOrder
4GL QRYINFO        Query Statistics:  ./paul_orders.p line 9
4GL QRYINFO        QueryId: 140595302309352
4GL QRYINFO        DB Blocks accessed:
4GL QRYINFO         s2k : 3609667
4GL QRYINFO        DB Reads:
4GL QRYINFO         Table: s2k.Customer : 201120
4GL QRYINFO         Index: Customer.CustNum : 201121
4GL QRYINFO         Table: s2k.Order : 727303
4GL QRYINFO         Index: Order.CustOrder : 928405
..
4GL 4GLTRACE       Return from Main Block [./paul_orders.p]
```

Also cool…

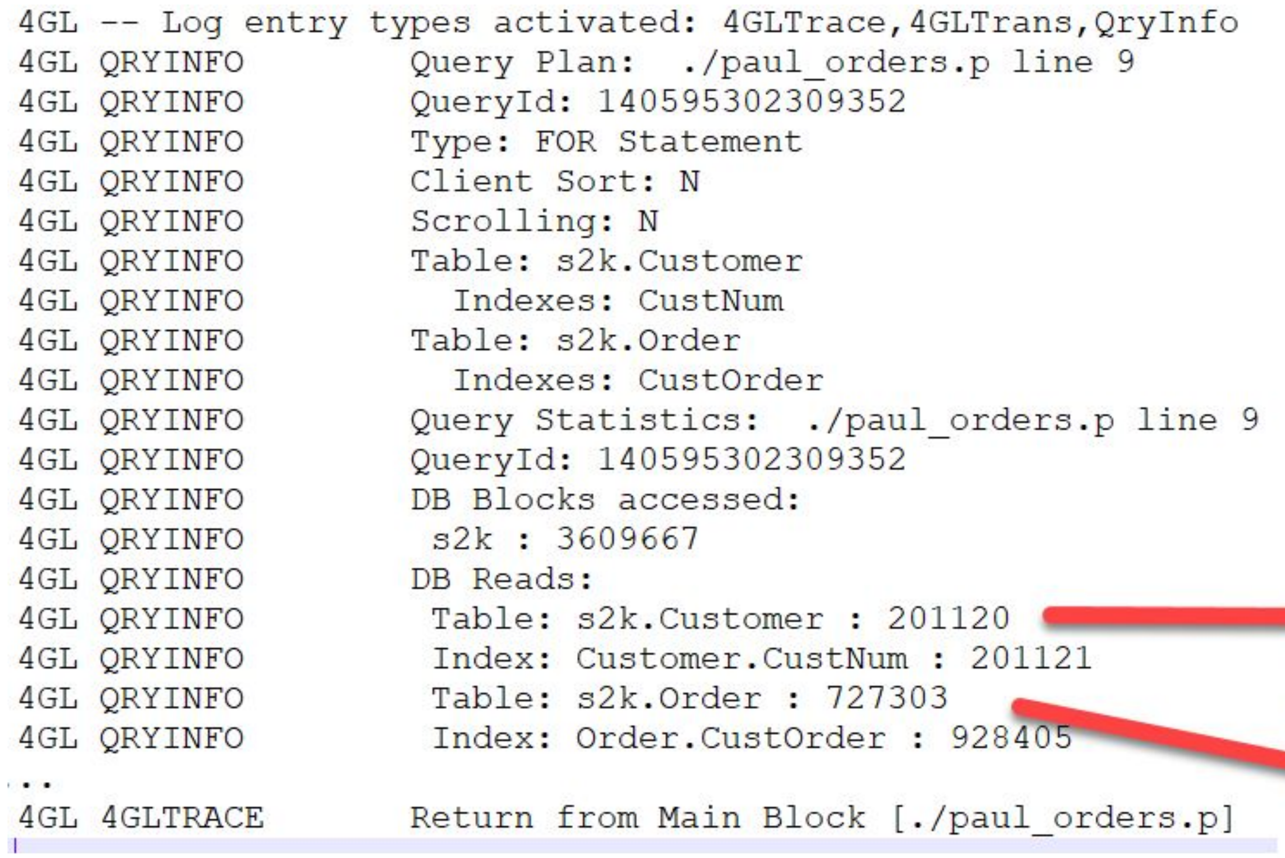

```
4GL -- Log entry types activated: 4GLTrace,4GLTrans,QryInfo
4GL QRYINFO        Query Plan:   ./paul_orders.p line 9
4GL QRYINFO        QueryId: 140595302309352
4GL QRYINFO        Type: FOR Statement
4GL QRYINFO        Client Sort: N
4GL QRYINFO        Scrolling: N
4GL QRYINFO        Table: s2k.Customer
4GL QRYINFO          Indexes: CustNum
4GL QRYINFO        Table: s2k.Order
4GL QRYINFO          Indexes: CustOrder
4GL QRYINFO        Query Statistics:   ./paul_orders.p line 9
4GL QRYINFO        QueryId: 140595302309352
4GL QRYINFO        DB Blocks accessed:
4GL QRYINFO         s2k : 3609667
4GL QRYINFO        DB Reads:
4GL QRYINFO         Table: s2k.Customer : 201120
4GL QRYINFO         Index: Customer.CustNum : 201121
4GL QRYINFO         Table: s2k.Order : 727303
4GL QRYINFO         Index: Order.CustOrder : 928405
...
4GL 4GLTRACE       Return from Main Block [./paul_orders.p]
```

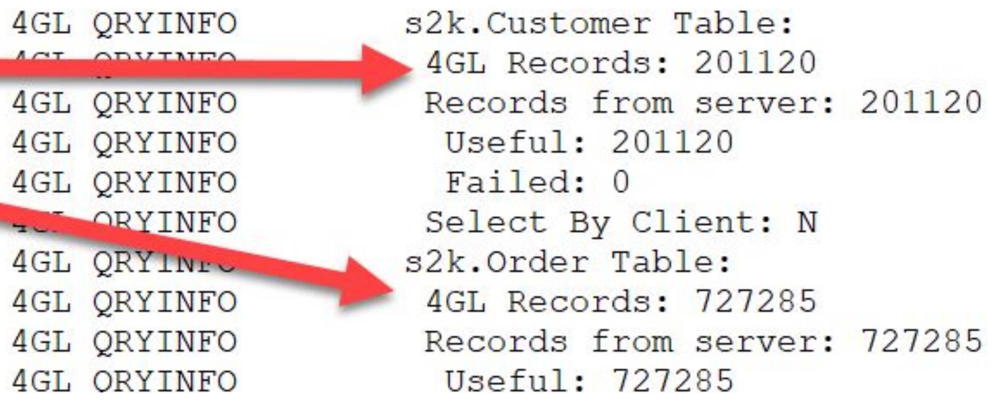

```
4GL QRYINFO        s2k.Customer Table:
4GL QRYINFO         4GL Records: 201120
4GL QRYINFO         Records from server: 201120
4GL QRYINFO          Useful: 201120
4GL QRYINFO          Failed: 0
4GL QRYINFO         Select By Client: N
4GL QRYINFO        s2k.Order Table:
4GL QRYINFO         4GL Records: 727285
4GL QRYINFO         Records from server: 727285
4GL QRYINFO          Useful: 727285
```

White Star software

Similar query…unindexed field

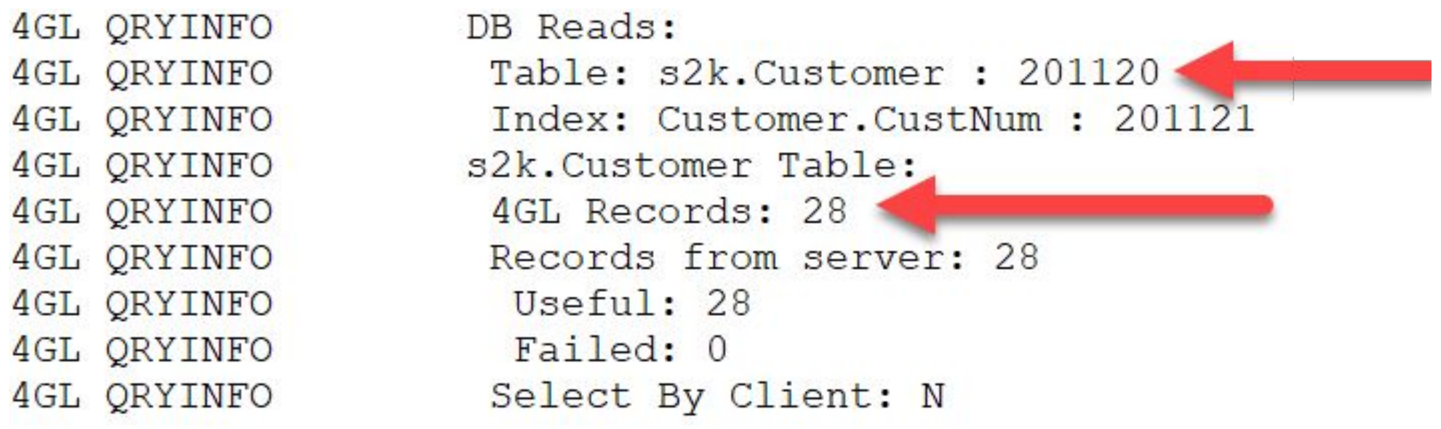

```
4GL QRYINFO         DB Reads:
4GL QRYINFO          Table: s2k.Customer : 201120
4GL QRYINFO          Index: Customer.CustNum : 201121
4GL QRYINFO         s2k.Customer Table:
4GL QRYINFO          4GL Records: 28
4GL QRYINFO          Records from server: 28
4GL QRYINFO           Useful: 28
4GL QRYINFO           Failed: 0
4GL QRYINFO          Select By Client: N
```

Tooling

Method #2: -zqil

- Unsupported/undocumented
- Do NOT use in production
- Writes query index usage information to the database log file

- Tells you which index is used and how deep into the index keys

-zqil

- Information is presented as upper and lower bounds
  - GT, GE are lower bounds
  - LT, LE are upper bounds
  - = is both an upper and lower bound

-zqil

```
for each customer no-lock where city = "Bellevue".
```

- There is no index that starts with “city”
  - Hence no upper no lower bound on index #15

```
ABL      3: (6135)  ==Compiled Query Resolution Method: Query No. 1==
ABL      3: (6157)  INDEX 15 0 0
ABL      3: (6136)  ==Server Query execution Method Query No. 1==
ABL      3: (6141)  INDEX 15
```

-zqil

```
for each customer no-lock where country > "A":
```

```
ABL      3: (6135)   ==Compiled Query Resolution Method: Query No. 1==
ABL      3: (6157)   INDEX 17 1 0
```

```
for each customer no-lock where country > "A" and PostalCode >
"100":
```

```
ABL      3: (6135)   ==Compiled Query Resolution Method: Query No. 1==
ABL      3: (6157)   INDEX 17 1 0
```

```
for each customer no-lock where country = "AT" and PostalCode >
"100":
```

```
ABL      3: (6135)   ==Compiled Query Resolution Method: Query No. 1==
ABL      3: (6157)   INDEX 17 2 1
```

What is index 17 ?

```
find _file where _file-name = "customer".
find _index of _file where _idx-num = 17.
displ _index-name.
```

```
Index-Name
--------------------------------
CountryPost
```

Combine QryInfo and -zqil

- QryInfo tells you how many record reads and how many useful records
- -zqil tells you how much of the index you are actually using!

Tooling

Method #3: COMPILE XREF

- Limited value:
  - Tells you WHICH indexes are used
  - Tells you if full table scan

```
custom.p 45 SEARCH s2k.Customer Name
custom.p 45 SEARCH s2k.Customer Salesrep

paul_orders.p 10 SEARCH s2k.Customer CustNum WHOLE-INDEX
```

White Star
software

Tooling

Method #4: Profiler

- Counts how often and how much time is spent in each line of code

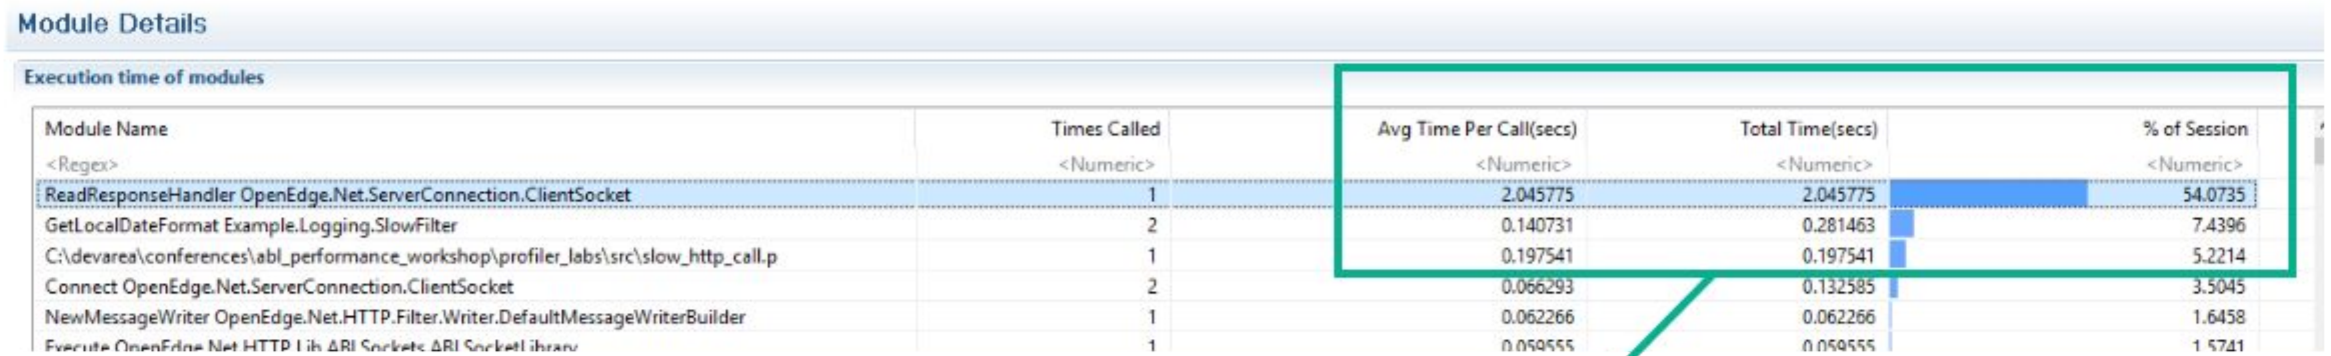

Module Details

Execution time of modules

Module Name	Times Called	Avg Time Per Call(secs)	Total Time(secs)	% of Session
<Regex>	<Numeric>	<Numeric>	<Numeric>	<Numeric>
ReadResponseHandler OpenEdge.Net.ServerConnection.ClientSocket	1	2.045775	2.045775	54.0735
GetLocalDateFormat Example.Logging.SlowFilter	2	0.140731	0.281463	7.4396
C:\devarea\conferences\abl_performance_workshop\profiler_labs\src\slow_http_call.p	1	0.197541	0.197541	5.2214
Connect OpenEdge.Net.ServerConnection.ClientSocket	2	0.066293	0.132585	3.5045
NewMessageWriter OpenEdge.Net.HTTP.Filter.Writer.DefaultMessageWriterBuilder	1	0.062266	0.062266	1.6458
Execute OpenEdge.Net.HTTP.Lib.ABLSockets.ABLSocketLibrary	1	0.059555	0.059555	1.5741

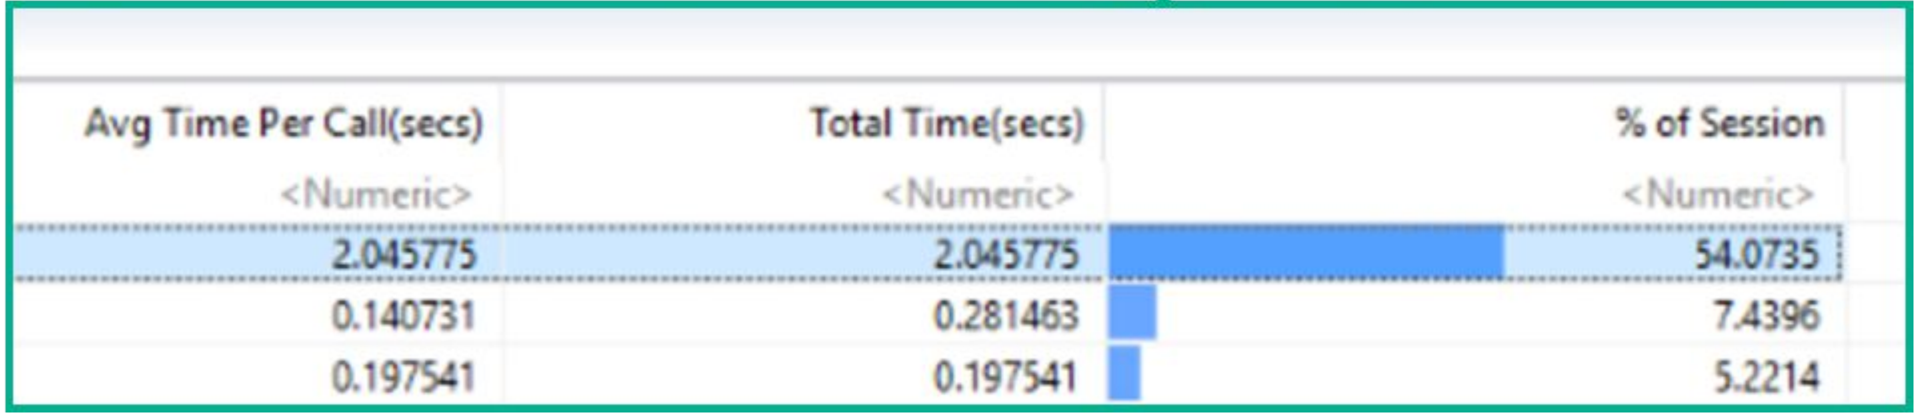

Avg Time Per Call(secs)	Total Time(secs)	% of Session
<Numeric>	<Numeric>	<Numeric>
2.045775	2.045775	54.0735
0.140731	0.281463	7.4396
0.197541	0.197541	5.2214

What’s next?

- You found the offending (and offensive) query
- Let’s improve it - i.e. read less records

Understanding index selection rules

- Only applies to ABL, not SQL
- Rules are applied in order, until only one index is left
    - IMPORTANT: Rules are NOT SELECTED, they are ELIMINATED
- Field match rules must be contiguous, from the first field in the index

Index selection rules

1. Pre-select only indexes with leading components in the where clause
2. If CONTAINS use word-index
3. Unique index with all components involved in the equality matches
4. Most active equality matches
   a. Sorta kinda…full matches trump partial matches
   b. But only if more than 1 field (sometimes)
5. Most active range matches
6. Most active sort matches

If you still have more than one index, or zero index, select one from

1. The primary index
2. First index alphabetically by name

White Star software

Example

```
for each order where orderNum = 12345.
for each order where orderNum = 12345 and CustNum = 5.
```

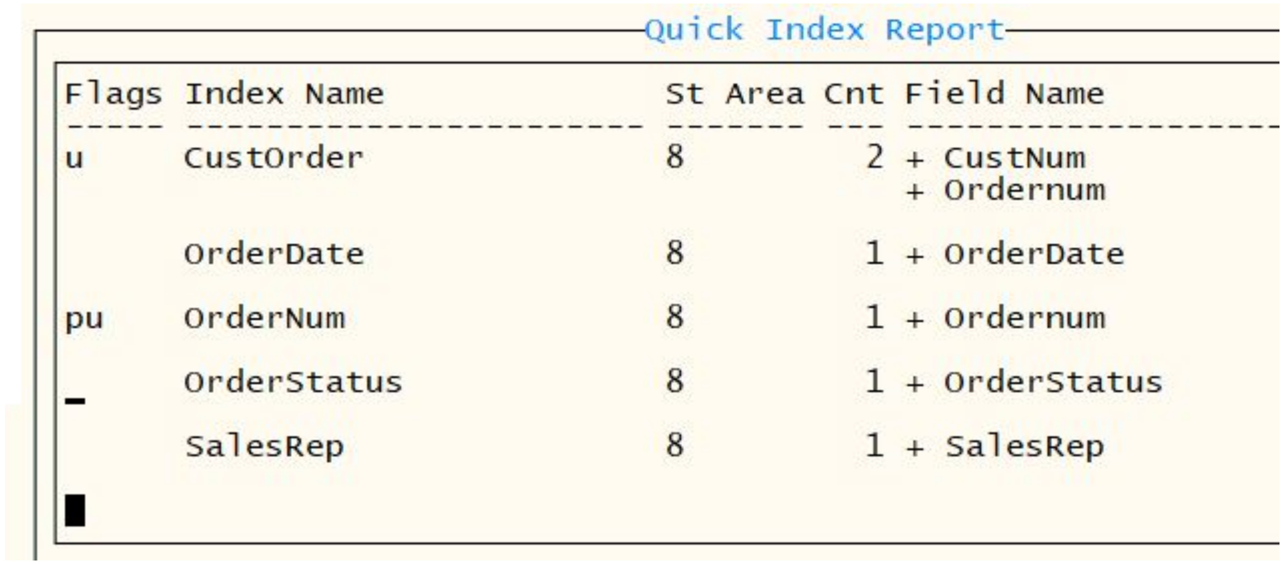

```
                              Quick Index Report
Flags Index Name                St Area Cnt Field Name
----- ------------------------- ------- --- --------------------
u     CustOrder                 8         2 + CustNum
                                            + Ordernum

      OrderDate                 8         1 + OrderDate

pu    OrderNum                  8         1 + Ordernum

      OrderStatus               8         1 + OrderStatus

      SalesRep                  8         1 + SalesRep
```

Tip: Use an elimination grid

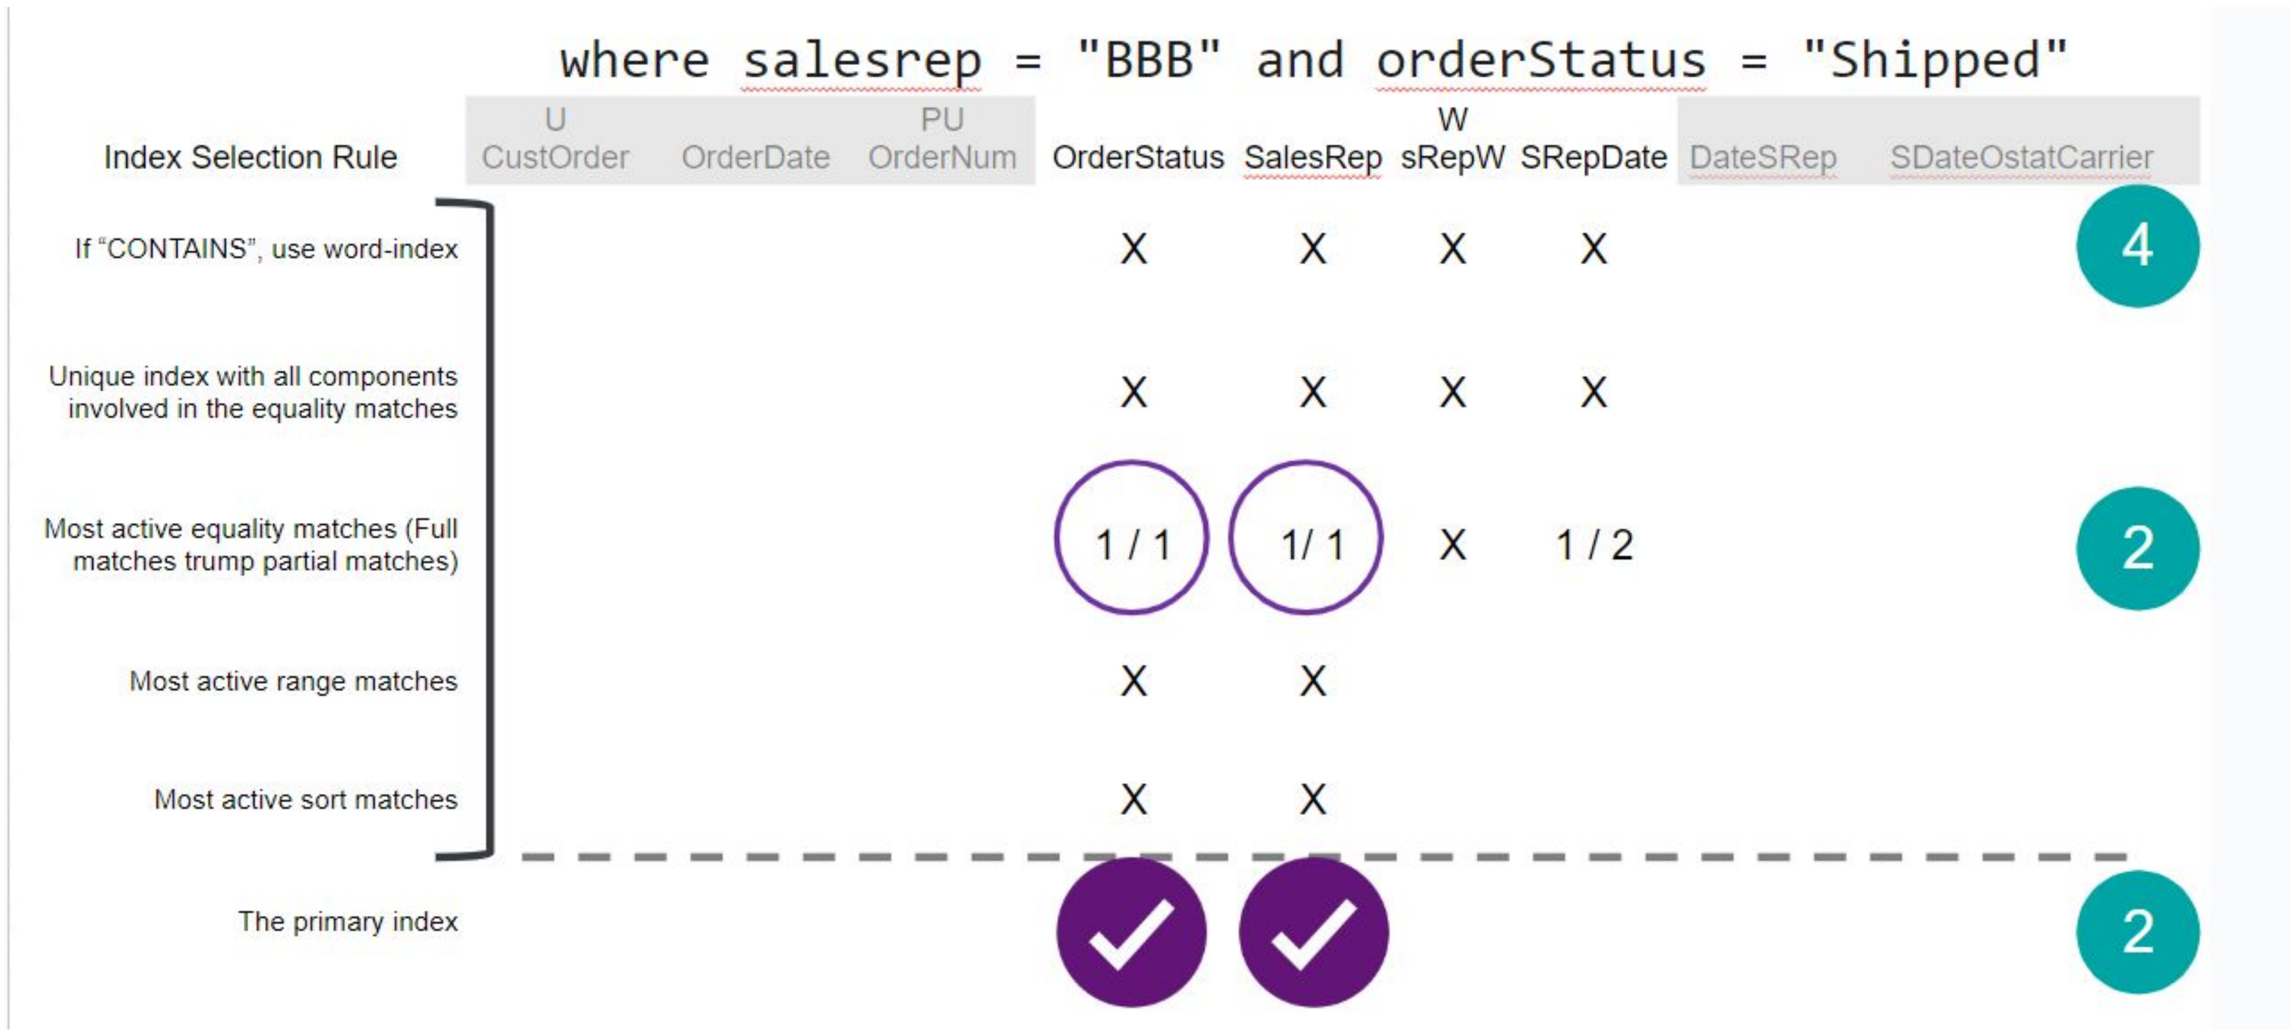

```
where salesrep = "BBB" and orderStatus = "Shipped"
```

	U		PU			W				
Index Selection Rule	CustOrder	OrderDate	OrderNum	OrderStatus	SalesRep	sRepW	SRepDate	DateSRep	SDateOstatCarrier	
If "CONTAINS", use word-index				X	X	X	X			4
Unique index with all components involved in the equality matches				X	X	X	X			
Most active equality matches (Full matches trump partial matches)				1 / 1	1/ 1	X	1 / 2			2
Most active range matches				X	X					
Most active sort matches				X	X					
The primary index				✓	✓					2

Multiple index use

- Where clause includes “AND”
  - ALL components of each index are involved in equality matches
  - No unique indexes are involved
- Where clause includes “OR”
  - Both sides of OR contain at least the lead component of an index
  - Either equality or range match
- CAREFUL: return order not guaranteed

Careful…

- Expressions break bracketing

```
for each order no-lock where month(orderDate) = 1 ...
```

- BEGINS does NOT break bracketing
  - Considered a range bracket

```
for each order no-lock where salesRep begins "D"
```

  - Uses the order.salesRep index
- MATCHES breaks bracketing
- Temp-table rules are subtly different

Special case: OR

- Each side of an OR is its own distinct index selection operation
  - Apply the rules to each side separately
  - Resulting records from both sides are then combined

Example

```
for each order no-lock where orderStatus = "Ordered" OR
                                         SalesRep = "BBB":
```

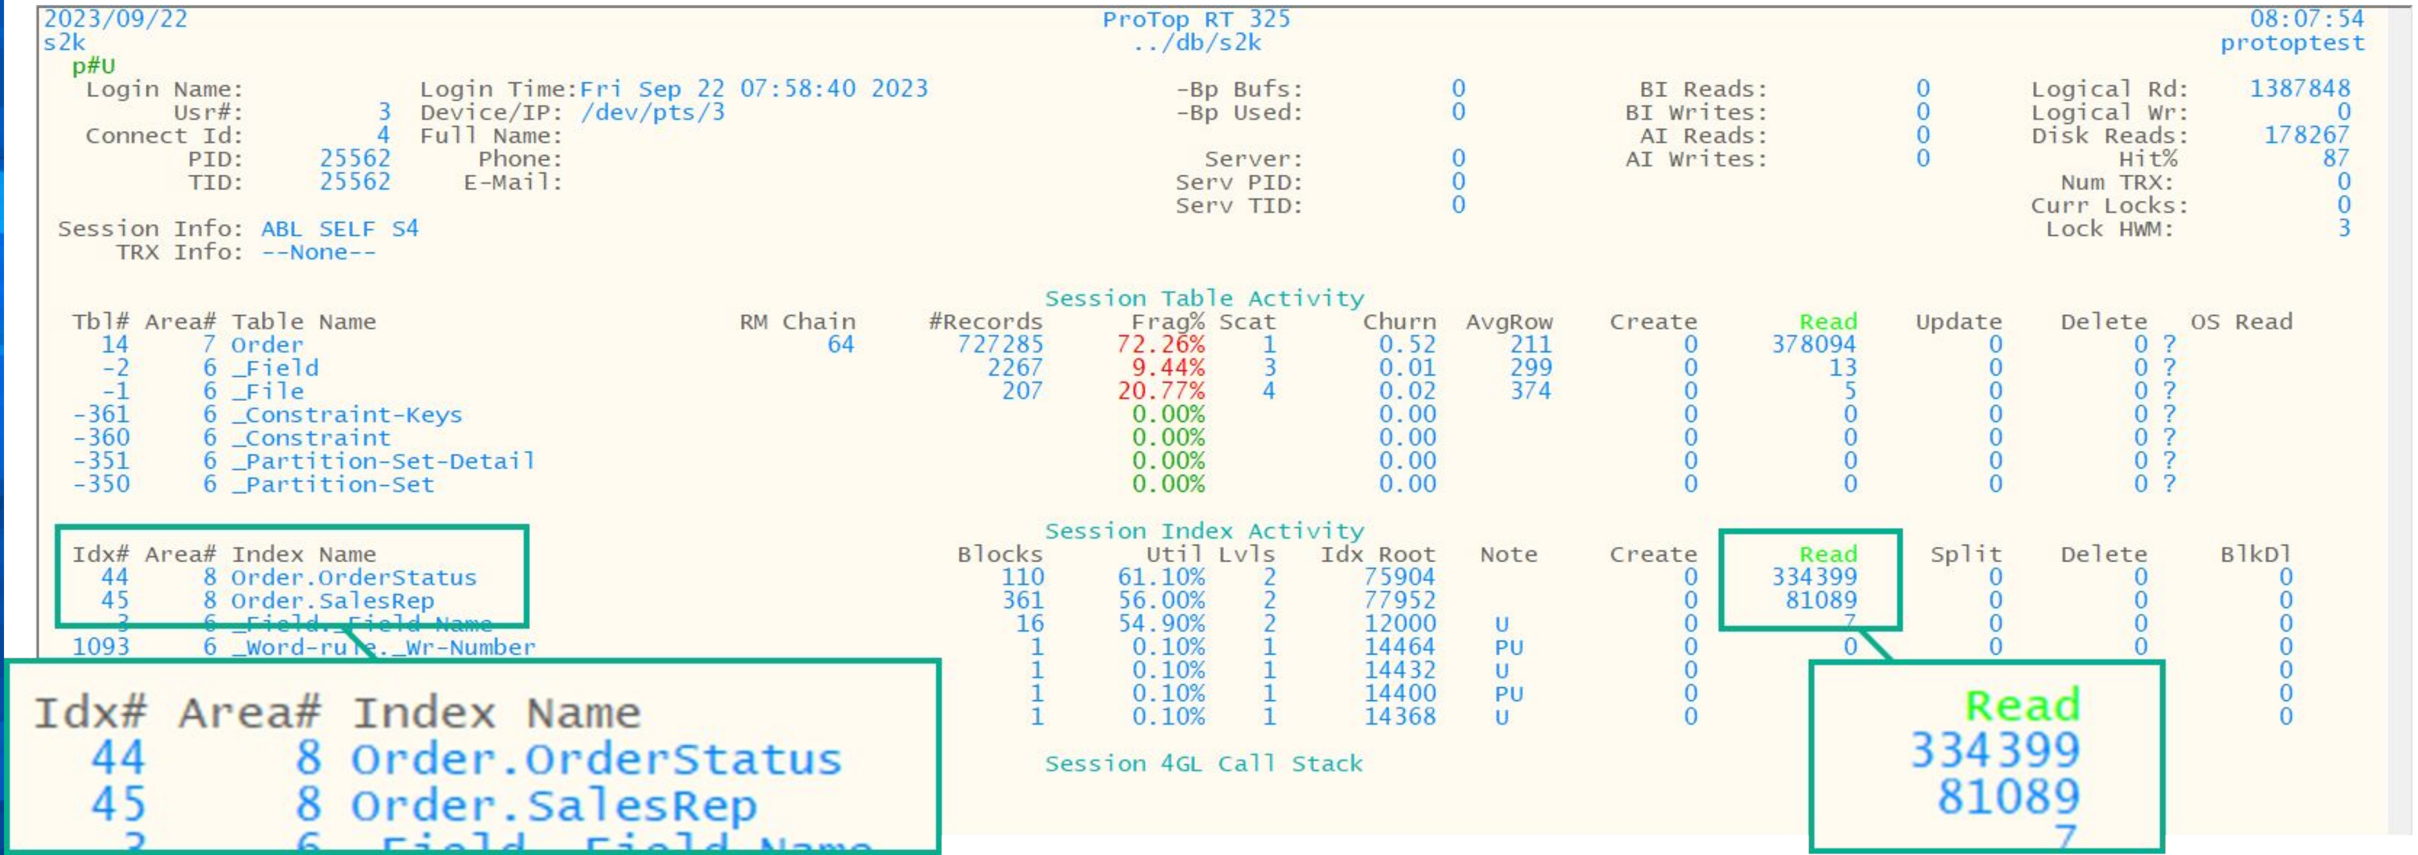

```
2023/09/22                                                       ProTop RT 325                                                      08:07:54
s2k                                                                ../db/s2k                                                     protoptest
 p#U
  Login Name:             Login Time:Fri Sep 22 07:58:40 2023           -Bp Bufs:          0          BI Reads:          0       Logical Rd:    1387848
        Usr#:          3  Device/IP: /dev/pts/3                          -Bp Used:          0          BI Writes:         0       Logical Wr:          0
  Connect Id:          4  Full Name:                                                                   AI Reads:          0       Disk Reads:     178267
         PID:      25562      Phone:                                      Server:          0          AI Writes:         0             Hit%         87
         TID:      25562     E-Mail:                                    Serv PID:          0                                          Num TRX:          0
                                                                        Serv TID:          0                                       Curr Locks:          0
 Session Info: ABL SELF S4                                                                                                           Lock HWM:          3
     TRX Info: --None--

                                                        Session Table Activity
 Tbl# Area# Table Name                  RM Chain     #Records     Frag% Scat     Churn  AvgRow     Create       Read     Update     Delete   OS Read
   14     7 Order                             64       727285    72.26%    1      0.52     211          0     378094          0          0 ?
   -2     6 _Field                                       2267     9.44%    3      0.01     299          0         13          0          0 ?
   -1     6 _File                                         207    20.77%    4      0.02     374          0          5          0          0 ?
 -361     6 _Constraint-Keys                                      0.00%           0.00                 0          0          0          0 ?
 -360     6 _Constraint                                           0.00%           0.00                 0          0          0          0 ?
 -351     6 _Partition-Set-Detail                                 0.00%           0.00                 0          0          0          0 ?
 -350     6 _Partition-Set                                        0.00%           0.00                 0          0          0          0 ?

                                                        Session Index Activity
 Idx# Area# Index Name                                   Blocks      Util Lvls  Idx Root   Note     Create       Read      Split     Delete     BlkDl
   44     8 Order.OrderStatus                               110    61.10%    2     75904               0     334399          0          0         0
   45     8 Order.SalesRep                                  361    56.00%    2     77952               0      81089          0          0         0
    3     6 _Field._Field-Name                               16    54.90%    2     12000    U          0          7          0          0         0
 1093     6 _Word-rule._Wr-Number                             1     0.10%    1     14464    PU         0          0          0          0         0
                                                              1     0.10%    1     14432    U          0                                          0
                                                              1     0.10%    1     14400    PU         0                                          0
                                                              1     0.10%    1     14368    U          0                                          0

                                                        Session 4GL Call Stack
```

```
Idx# Area# Index Name
  44     8 Order.OrderStatus
  45     8 Order.SalesRep
```

```
  Read
334399
 81089
```

Example

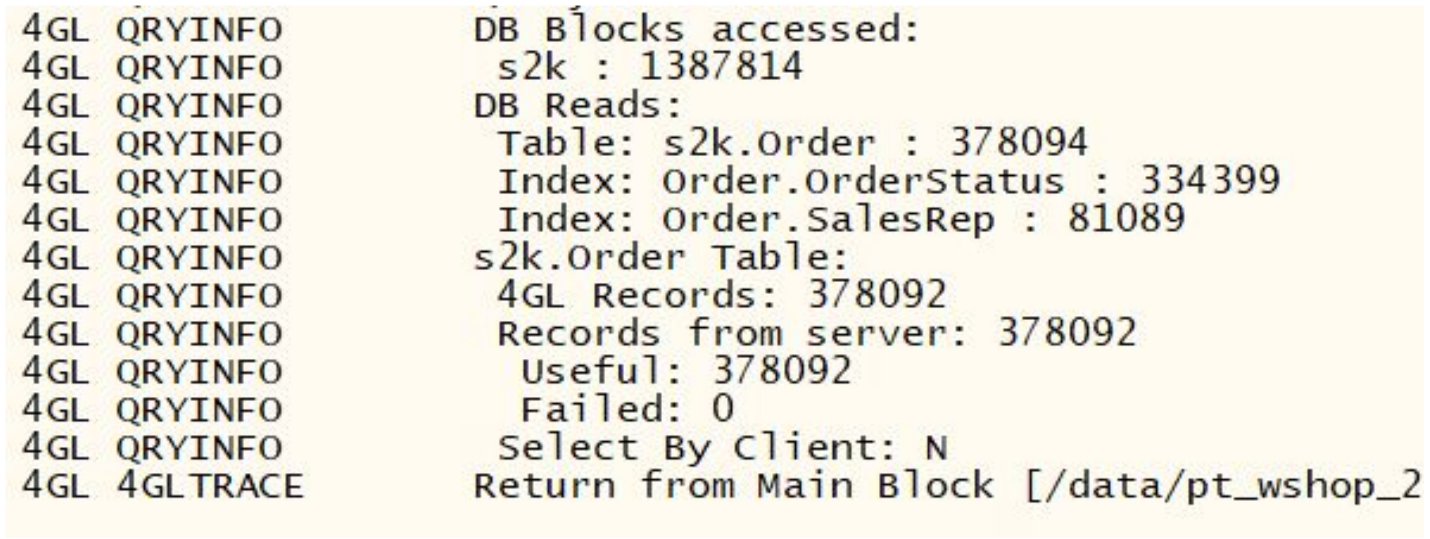

```
4GL QRYINFO          DB Blocks accessed:
4GL QRYINFO           s2k : 1387814
4GL QRYINFO          DB Reads:
4GL QRYINFO           Table: s2k.Order : 378094
4GL QRYINFO           Index: Order.OrderStatus : 334399
4GL QRYINFO           Index: Order.SalesRep : 81089
4GL QRYINFO          s2k.Order Table:
4GL QRYINFO           4GL Records: 378092
4GL QRYINFO           Records from server: 378092
4GL QRYINFO            Useful: 378092
4GL QRYINFO            Failed: 0
4GL QRYINFO           Select By Client: N
4GL 4GLTRACE         Return from Main Block [/data/pt_wshop_2
```

Table scans

```
for each order no-lock table-scan:
```

- If you expect to read more than ⅓ of the table, consider using the TABLE-SCAN option
  - Does not use any index
  - Returns data in “on-disk” order
  - Only if table is in a type 2 storage area

Client/Server Queries

- Client-server queries are going to be slower than shared memory queries
- Records are transported to client in "messages"
  - There is an OpenEdge message buffer size AND a TCP MTU (Maximum Transmission Unit)

- You can make it better with the following

White Star
software

Client/Server Queries

- Use NO-LOCK
  - Anything else will result in one record per OE message
- Use field lists
  - Don’t send the whole record if you only need one field
- Make the message buffer size (-Mm) bigger
  - Default is 1K
  - Use at least 8K or 16K
- Use -prefetch* parameters
  - No use having a big message unless you can fill it !!
- Server-side joins in OE 12

SQL

- SQL uses a cost-based optimizer
- Calculate cost statistics using UPDATE STATISTICS
- Repeat periodically or when texture of data changes
    - Ex: purge or mass load

Monitor OpenEdge. Anticipate Problems. Avert Disasters.

Prevent downtime, increase performance, and lower costs for cloud, on-premise, and hybrid environments with the only monitoring tool designed explicitly for OpenEdge.

```
otrx Old Transaction 3788 > 3600
[2022/07/28 11:59:05.584-04:00]
[server: naprod.acme.com]
[dbname: /db/prod/custprod]
[cfgname: custprod]
```

```
con_Block Blocked sessions 3 > 2
[2022/07/28 13:04:45.823-04:00]
[server: wx33euprd.xyz.com]
[dbname: g:\wx\prd\accntprd]
[cfgname: acctprd]
```

```
ai_Locked Locked AI Extents 7 > 2
[2022/07/28 22:38:11.987-04:00]
[server: naprod.acme.com]
[dbname: /db/prod/custprod]
[cfgname: custprod]
```

535	**1.5**	**4,500**	**325,000**
Customer Sites	Petabytes of data	Databases	Connected users