

QEP ANALYSIS (COMMENTS IN BLUE)



1. Notice that none of the tables have statistics listed. These tables all need to be optimized with the command “optimizedb -zk-zr200 -zu200” to collect the statistics that will hopefully be the best possible.

1. help table accounting

```
Name:                accounting
Owner:               prod_dba
Created:             04-aug-2002 12:31:23
Location:            ingdata
Type:                user table
Version:             II2.6
Page size:           2048
Cache priority:      0
Alter table version: 0
Alter table totwidth: 77
Row width:           77
Number of rows:      1929165
Storage structure:   btree
Compression:         data
Duplicate Rows:      allowed
Number of pages:     157480
Overflow data pages: 0
Journaling:          enabled
Base table for view: yes
Optimizer statistics: none
```

→ It appears that there could be duplicate rows of data returned for a fully qualified keyed select. This could cause singleton select errors.

Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
uw_year	integer	2	no	yes	2
polno	integer	4	no	yes	1
layer	integer	1	no	yes	3
curr_pay	varchar	3	no	yes	4
tx_date	date		no	yes	
quarter	integer	1	no	yes	5
rev_acct_num	varchar	7	no	yes	
mliterr_code	integer	1	no	yes	
mlicurr_code	integer	1	no	yes	
ins_code	varchar	3	no	yes	
rev_locl_amt	money		no	yes	
rev_book_amt	money		no	yes	
acct_fed	date		no	yes	
finantxs_id	integer	4	no	yes	
legal_id	integer	4	no	yes	

Secondary indexes:

Index Name	Structure	Keyed On	
x_acct_ftxid	btree	finantxs_id	
x_acct_num	isam	rev_acct_num	→ “btree” better choice for this index.

2> help index x_acct_ftxid

Name: x_acct_ftxid
Owner: prod_dba
Created: 04-aug-2002 12:41:21
Location: ingdata
Type: secondary index on accounting
Version: II2.6
Page size: 2048
Cache priority: 0
Alter table version: 0
Alter table totwidth: 8
Row width: 8
Number of rows: 1929165
Storage structure: btree
Compression: none
Duplicate Rows: allowed
Number of pages: 20285
Overflow data pages: 0
Journaling: enabled if journaling on the base table is enabled
Base table for view: no
Optimizer statistics: will use any existing statistics on the base table

Index Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
finantxs_id	integer	4	no	yes	1
tidp	integer	4	no	no	2

3> help index x_acct_num

Name: x_acct_num
Owner: prod_dba
Created: 04-aug-2002 12:40:44
Location: ingdata
Type: secondary index on accounting
Version: II2.6
Page size: 2048
Cache priority: 0
Alter table version: 0
Alter table totwidth: 13
Row width: 13
Number of rows: 1929165
Storage structure: isam
Compression: none

3> help index x_acct_num (continued)

Duplicate Rows: allowed
Number of pages: 20905
Overflow data pages: 0
Journaling: enabled if journaling on the base table is enabled
Base table for view: no
Optimizer statistics: will use any existing statistics on the base table
Index Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
rev_acct_num	varchar	7	no	yes	1
tidp	integer	4	no	no	2

4> help table currency

Name: currency
Owner: prod_dba
Created: 04-aug-2002 12:13:13
Location: ingdata
Type: user table
Version: II2.6
Page size: 2048
Cache priority: 0
Alter table version: 0
Alter table totwidth: 71
Row width: 71
Number of rows: 69
Storage structure: btree with unique keys
Compression: none
Duplicate Rows: not allowed
Number of pages: 8
Overflow data pages: 0
Journaling: enabled
Base table for view: no
Optimizer statistics: none

Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
curr_abbr	char	3	no	yes	1
curr_name	char	30	no	yes	
exch_rate	float	8	no	yes	
inflate_fact	integer	4	no	yes	
mlicurr_code	char	2	no	yes	
restricted	char	1	no	yes	
rate_date	date		no	yes	
recip_rate	float	8	no	yes	
rskreg_n_abbr	char	3	no	yes	

Secondary indexes: none

5> help table contract

Name: contract
Owner: prod_dba
Created: 04-aug-2002 12:10:27
Location: ingdata
Type: user table
Version: II2.6
Page size: 2048
Cache priority: 0
Alter table version: 0
Alter table totwidth: 457
Row width: 457
Number of rows: 10003
Storage structure: btree with unique keys
Compression: none
Duplicate Rows: not allowed
Number of pages: 2710
Overflow data pages: 0
Journaling: enabled
Base table for view: yes
Optimizer statistics: none

Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
uw_year	integer	2	no	yes	2
polno	integer	4	no	yes	1
layer	integer	1	no	yes	3
curr_pay	varchar	3	no	yes	4
deal_desc	varchar	45	no	yes	
ref_num	varchar	20	no	yes	
rprr_cls1	char	3	no	yes	
line	varchar	3	no	yes	
nature_protn	char	3	no	yes	
treaty	varchar	3	no	yes	
war_risk	varchar	3	no	yes	
res	varchar	3	no	yes	
misterr_abbr	varchar	3	no	yes	
rsrve_srce	varchar	3	no	yes	
under_freq	varchar	1	no	yes	
loss_method	varchar	3	no	yes	
prem_type	varchar	3	no	yes	
depst_prem	varchar	3	no	yes	
curr_cov	varchar	3	no	yes	
prem_freq	varchar	1	no	yes	
status	varchar	3	no	yes	
ins_code	varchar	3	no	yes	
mgfee_type	varchar	3	no	yes	
mgfee_rate	float	4	no	yes	
commission	float	4	no	yes	
brkr_code	varchar	4	no	yes	
brkr_rate	float	4	no	yes	
brkr_type	varchar	3	no	yes	
profit_comm	float	4	no	yes	
effective	date		no	yes	
fst_effctive	date		no	yes	
next_anniv	date		no	yes	
cancel	date		no	yes	
risk_limit	money		no	yes	
retntion_dlr	money		no	yes	
agg_limit	money		no	yes	
max_limit	money		no	yes	
num_lvs_wrtty	integer	1	no	yes	

Column Information: (continued)

Column Name	Type	Length	Nulls	Defaults	Seq
retntion_pct	float	4	no	yes	
share_pct	float	4	no	yes	
canc_notice	integer	2	no	yes	
fac_reserved	varchar	1	no	yes	
treaty_exist	varchar	1	no	yes	
lst_fincl_tx	date		no	yes	
lst_contr_tx	date		no	yes	
accrual_done	varchar	1	no	yes	
pr_date	date		no	yes	
nxt_pay_due	date		no	yes	
tot_claims	money		no	yes	
asscomp_code	varchar	3	no	yes	
arc_block_id	integer	4	no	yes	
pre_merger_block_id	varchar	3	no	yes	
mgu_code	varchar	4	no	yes	
uw_assigned	varchar	12	no	yes	
last_possible_exp_date	date		no	yes	
lastposs_exp_status_code	char	8	no	yes	
ultimate_pric_prem_amt	money		no	yes	
ultimate_uw_prem_amt	money		no	yes	
net_risk_amt	money		no	yes	
retro_direct_type	char	6	no	yes	
pool_facility_type_code	char	9	no	yes	
orig_ced_comp_code	char	3	no	yes	
risk_region_code	char	3	no	yes	
uw_risk_type_code	char	10	no	yes	
syndicate_group_id	integer	4	no	yes	
risk_country_abbr	char	3	no	yes	

Secondary indexes: none

6> help table dates

```

Name:          dates
Owner:         prod_dba
Created:       04-aug-2002 12:13:10
Location:      ingdata
Type:          user table
Version:       II2.6
Page size:     2048
Cache priority: 0
Alter table version: 0
Alter table totwidth: 68
Row width:     68
Number of rows: 1
Storage structure: heap
Compression:   data
Duplicate Rows: allowed
Number of pages: 3
Overflow data pages: 0
Journaling:    enabled
Base table for view: no
Optimizer statistics: none
    
```

6> help table dates (continued)

Column Information:

Column Name	Type	Length	Nulls	Defaults	Key Seq
today_dt	date		no	yes	
monthend_dt	date		no	yes	
accrue_dt	date		no	yes	
mthend_st	varchar	18	no	yes	
arc_cutoff_date	date		no	yes	

Secondary indexes: none

7> SET QEP

8> set optimizeonly

```
9> SELECT a.uw_year
      ,a.polno
      ,a.layer
      ,a.curr_pay
      ,a.quarter
      ,a.legal_id
      ,prem_due =
DATE (c.effective)+ DATE (CONCAT (VARCHAR ((a.quarter - 1) * 3),'mos'))
      ,os_prem = SUM (a.rev_locl_amt)
      ,os_prem_book = SUM (a.rev_locl_amt * b.exch_rate)
      ,currency b
      ,contract c
      ,dates d
WHERE a.curr_pay = b.curr_abbr
      AND DATE (c.effective)+DATE (CONCAT (VARCHAR ((a.quarter - 1) * 3),'mos'))
      < d.monthend_dt - DATE ('180 days')
      AND a.rev_acct_num = '2112.01'
      AND a.uw_year = c.uw_year
      AND a.polno = c.polno
      AND a.layer = c.layer
      AND a.curr_pay = c.curr_pay
GROUP BY a.uw_year
      ,a.polno
      ,a.layer
      ,a.curr_pay
      ,a.quarter
      ,a.legal_id
      ,c.effective
HAVING ABS (SUM (a.rev_locl_amt)) > 1000
```

QUERY PLAN 3,4, **no timeout**, of function aggregate producing temporary table T4 → The optimizer found the plan that it felt was best.

aggregate expression -> → This is from the “HAVING” clause – just informational.

T4A9 =
sum
(agg expr)

aggregate expression -> → This is from the “GROUP BY” clause – just informational.

T4A8 =
sum
(rev_locl_amt)

by expression attribute -> T4A7
= effective
by expression attribute -> T4A6
= legal_id
by expression attribute -> T4A5
= quarter
by expression attribute -> T4A4
= curr_pay
by expression attribute -> T4A3
= layer
by expression attribute -> T4A2
= polno
by expression attribute -> T4A1
= uw_year

Sort Keep dups
Pages 3 Tups 111 → You will see these again!
D19728 C6160 Disk I/O count is high.

/
Cart-Prod → RED FLAG! Caused by a disjoint query.
Heap In this case it is the 'dates' table
Pages 3 Tups 111 only contains a single row of data.
D19728 C6149

/

	K Join (curr_abbrev) → Key join = good	Proj-rest
	Heap	Heap
	Pages 6 Tups 221	Pages 1 Tups 1
	D19727 C6144	D1 C0
/	FSM Join(polno, uw_year, layer, curr_pay) Partial(FA0 curr_pay) Pages 48 Tups 2214 D19723 C6102	\
	currency (b) B-Tree(curr_abbrev) Pages 8 Tups 69	/
		dates (d) cHeap Pages 3 Tups 1
/	T Join(tidp) Sort on(polno, uw_year, layer, curr_pay) Pages 288 Tups 19292 D18368 C388	
Proj-rest Sorted(polno, uw_year, layer, curr_pay) Pages 130 Tups 10003 D1355 C100		

→ Disk estimates are cumulative moving up the tree. We see that the biggest increase occurs at this node. That indicates that the index used might not have been the best choice. The best way to address this would be to add other supporting data columns to this index (adding curr_pay, polno, layer, quarter, uw_year, and rev_acct_num would add 22 bytes per row to the index, but then the index would satisfy most of the query and significantly reduce I/O, thereby improving performance.

→ Contract table is being scanned

```

/
contract
(c)
B-Tree (NU)
Pages 2710 Tups 10003

```

```

/
Proj-rest
Heap
Pages 144 Tups 19292
D211 C2

```

```

\
accounting
(a)
cB-Tree (NU)
Pages 157480 Tups 1929165

```

→ The 'accounting' table shows that its primary key is not being used (NU). This is OK since you can see that it earlier selected data from an index on this table (see T Join above).

```

/
x_acct_num

```

```

I(a)
Isam(rev_acct_num)
Pages 20905 Tups 1929165

```

→ Start here and work your way up. Once you get to the top then look at the query below. What you should see is the most restrictive part of the query "pushed down" to this level. The goal is to minimize the result set of data as quickly as possible.

This is showing us that 'x_acct_num' is an index (I) on the table alias "a" ('account', from the query above), that the storage structure of the index is isam, and that it has an estimated 1,929,165 rows of data on 20.905 Ingres pages (data and index).

QUERY PLAN 1,1, no timeout, of main query

```

Proj-rest
Sorted(T4A1 uw_year,
T4A2 polno,
T4A3 layer,
T4A4 curr_pay,
T4A5 quarter,
T4A6 legal_id,
T4A7 effective,
T4A9 rev_locl_amt)
Pages 2 Tups 55
D19729 C6161

```

→ Shows the sort order of the data set.

→ These are the estimated query. There was only one additional Disk I/O over the base query (above).

```

/
T4
Heap
Pages 3 Tups 111

```

→ This is the temporary table created by the first step (above). The default result structure is heap. The optimizer estimates that there will be 111 rows of data on 3 data pages.

Notes:

1. Make sure that the table that you feel is the most restrictive is pushed down to the bottom of the query.
2. Always look at the estimates at each level to see if they are reasonable. When an estimate is too high or too low that generally indicates an issue with statistics.
3. Look to see where I/O increases. This is often indicative of either a problem or an opportunity for improvement.
4. Look for Cartesian Product nodes (Cart-Prod). These are almost always bad.
5. Look for indexes / storage structures that are not used (NU). Unless there is an action on a related index prior to that node this indicates a table scan.
6. Remember that Ingres will only use a single access method per table. Because of this it is important to provide as much additional data as makes sense (e.g., for the 77 byte account table adding 22 bytes to an index would make sense, but adding 30+ bytes to an index would not make sense). This will maximize the probability of having good performing queries every time.

Recommendations:

1. Change the storage structure for the index x_acct_num from 'isam' to 'btree'.
2. Add the columns curr_pay, quarter, rev_acct_num, uw_year, polno, layer to the index x_acct_num.
3. Add statistics to the database using the optimizedb command (use the flags "-zk -zu200 -zr200").
4. After optimization run the "sysmod" command on the database.
5. If these four changes are made I would expect to see a decrease in Disk I/O of an order of magnitude (provided the estimates are reasonably correct). This should result in the query running significantly faster.