



Of course it will take  
the right index!  
Oh... Why did it take that one?

Simon L. Prinsloo  
simon@vidisolve.com



Simon has been working with the Progress 4GL and the database for the last 18 years. Keeping up with the developments in the 4GL and the database has always been his passion.

Simon has been involved in CASE tools since his first introduction to the 4GL and worked on various other systems, including the Namibian Tax System, large ERP systems in Southern Africa and a number of smaller systems.

Progress Software South Africa used Simon at various times to present Progress training courses for version 9 and early version 10. He also presented sessions at three regional Progress conferences in South Africa.

He created Vidisolve in 2011 and started to focus much more actively on the modernisation efforts of his clients. He is currently also busy with his Masters degree in Information Technology at the University of Pretoria.

# Introduction

Simon Prinsloo

Vidisolve in Pretoria

Working with Progress since v.7 in 1996

Worked on various commercial systems

Mostly focused on CASE tools and implementing new functionality in legacy projects



## Why do we care?

- Hardware is getting bigger and faster
- The compiler selects the index
- It uses a fantastic set of rules
- It will combine indexes



## Why do we care?

Knowing the importance of creating indexes to support common data access patterns is a big step toward efficient design. However, to make your query code and indexes work together effectively, you must understand how ABL chooses indexes to satisfy a particular query.

OpenEdge Web Paper: ABL Database Triggers and Indexes



## Why do we care?

- Data bases grow faster than ever
  - This negates the effect of faster hardware
- The compiler lacks key information
  - It has no way to predict data distribution
  - It has no way to resolve certain constructs



The index selection is rule based. That enables it to be done at compile time and it results in predictable, reproducible index selection. But applying good rules on an abnormal data distribution can have unwanted consequences.

For example, a pipeline of records that needs to be processed can have a logical field indicating whether the record was processed or not. As long as we read the unprocessed records, the set will be very small, but if we read processed records, the set will very quickly be very large.

# Terminology

## WHERE

WHERE *searchExpr* [ BY *field* ]

WHERE *searchExpr* AND *searchExpr* [ BY *field* ]

WHERE *searchExpr* OR *searchExpr* [ BY *field* ]

## Bracketing

WHOLE-INDEX

## Sort

SORT-ACCESS



## Using multiple index brackets

Multi-bracket queries can be used by

- FOR EACH
- PRESELECT
- QUERY objects

Different brackets can be for the same index or for other indexes.



## Using single index bracket

A single index bracket is selected when

- FIND is used
- USE-INDEX is specified
- A condition is encountered that forces a WHOLE-INDEX bracket.



## Prime directive

- The programmer's wish, the compiler's command.
  - ✓ TABLE-SCAN (Type II Storage only)
  - ✓ USE-INDEX

With great power comes great responsibility.



If you do not have a Type II storage area and you use TABLE-SCAN, the DBMS will simply fall back to a full index scan of your primary index.

## Rules for selecting a single index

1. CONTAINS – Use the word index
2. Use a unique index – 0 or 1 record
  - All components must be used in equality matches
3. Use the index with the most active equality matches
  - Must apply to successive, leading components
  - Must be joined with AND



Of course, if there is a USE-INDEX or a TABLE-SCAN, the compiler will not select an index, it will do what it is told.

These rules will also apply top to bottom until only one index is left. In essence, each rule eliminates all indexes that does not fit the rule, but if it would eliminate all indexes, the rule is skipped. If more than one index is left after a rule was applied, we move on to the next rule.

## Selecting a single index

4. Use the index with the most active range matches
  - The component is the first or only component
  - All preceding components in the index have active equality matches – joined with AND
5. Use index with the most sort matches
6. Use index that comes first alphabetically
7. Use the primary index



## Selecting a single index

4. Use the index with the most active range matches
  - The component is the first or only component
  - All preceding components in the index have active equality matches – joined with AND
5. Use index with the most sort matches
6. Use index that comes first alphabetically
7. Use the primary index



If none of the first four rules are applied to reduce the index set, you are most likely in trouble.

Rule 5 can cause a selection or act as a tie breaker between two or more indexes, in which case it is beneficial. But if it is the sole reason for index selection, it means the whole table will be scanned.

Rule 6 and 7 are essentially mutually exclusive. If two or more suitable indexes remains even after applying rule 5, rule 6 will apply. But if none of the indexes are in the list after rule 5, rule 7 will apply. It is however possible that rule 7 may outrank rule 6 in a case where the Primary index are one of the tied indexes. I did not test this (yet).

## When multiple indexes can be combined

### WHERE with AND

- All components of each candidate index are involved in equality matches
- None of the candidate indexes are unique

### WHERE with OR

- Each side of the OR can be used to select an index that does not result in a WHOLE-INDEX scan



The index selection is rule based. That enables it to be done at compile time and it results in predictable, reproducible index selection. But applying good rules on an abnormal data distribution can have unwanted consequences.

For example, a pipeline of records that needs to be processed can have a logical field indicating whether the record was processed or not. As long as we read the unprocessed records, the set will be very small, but if we read processed records, the set will very quickly be very large.

## How to find out which index is used

### ➤ COMPILE ... XREF

- SEARCH
- SORT-ACCESS
- WHOLE-INDEX

### ➤ QUERY-HANDLE:INDEX-INFORMATION (n)

- Comma separated list of index names, one for each bracket
- First entry can be WHOLE-INDEX



## Single-field indexes

glTransaction

transactionID

transactionDate

transactionType

transactionDate + transactionDate

\* transactionID + transactionID

transactionType + transactionType



1

glTransaction

transactionID  
transactionDate  
transactionType

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FOR EACH glTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015:
  ...
END.
```

**Which Index?**

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

# 1

glTransaction

transactionID  
transactionDate  
transactionType

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FOR EACH glTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015:
  ""
END.
```

**Which Index?**

➤ transactionType

**Why?**

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

1

glTransaction

transactionID  
transactionDate  
transactionType

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FOR EACH glTransaction NO-LOCK
  WHERE transactionType = "INV"
     AND transactionDate >= 05/01/2015
     AND transactionDate <= 05/31/2015:
  ...
END.
```

**Which Index?**

➤ transactionType

**Why?**

➤ Rule 3: Use the index with the most active equality matches.

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

transactionDate has range matches and will not be considered for the combination of indexes when used as part of an AND.

2

glTransaction

transactionID  
transactionDate  
transactionType

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

**Which Index?**



This is a rather awkward structure, but will it give us some better performance if we need it?

## 2

### glTransaction

transactionID  
transactionDate  
transactionType

### Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

### Which Index?

- transactionType
- TransactionDate

### Why?



## 2

### glTransaction

transactionID  
transactionDate  
transactionType

### Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

### Which Index?

- transactionType
- TransactionDate

### Why?

- Rule 3 renders two options
- FOR EACH can combine indexes



All index components of the two candidate indexes are used in equality matches. For that reason, they can be combined by the FOR EACH.

# 3

## glTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FIND LAST glTransaction NO-LOCK  
WHERE transactionType = "INV"  
AND transactionDate = TODAY  
NO-ERROR.
```

## Which Index?



# 3

## glTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FIND LAST glTransaction NO-LOCK  
WHERE transactionType = "INV"  
AND transactionDate = TODAY  
NO-ERROR.
```

## Which Index?

➤ transactionDate

## Why?



# 3

## glTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

transactionType  
+ transactionType

```
FIND LAST glTransaction NO-LOCK  
WHERE transactionType = "INV"  
AND transactionDate = TODAY  
NO-ERROR.
```

### Which Index?

➤ transactionDate

### Why?

- Rule 3 renders two options
- FIND cannot combine indexes
- Rule 6: Take the index that comes first alphabetically.



In this case we are lucky. It is pure coincidence that the date index should work better. If it was the other way round, we could have been stuck with reading many more records and dropping most of it.

## Compound indexes

### glrTransaction

transactionID

transactionDate

transactionType

transactionDate + transactionDate  
+ transactionType

\* transactionID + transactionID

transactionType + transactionType  
+ transactionDate



# 4

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK  
  WHERE transactionType = "INV"  
    AND transactionDate >= 05/01/2015  
    AND transactionDate <= 05/31/2015:  
  ...  
END.
```

## Which Index?



# 4

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?

➤ transactionType

## Why?



# 4

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
     AND transactionDate >= 05/01/2015
     AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?

➤ transactionType

## Why?

➤ Rule 3: Use the index with the most active equality matches.



There is essentially no difference between selection rules for this example and example one, but we will get significantly better performance in this case, as we not only have an equality bracket on the first component of the index, but also have a range match on the second component.

# 5

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glrTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

## Which Index?



This one is the same example as we had in example two, where both indexes was selected by rule 3 and then combined.

# 5

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glrTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
  END.  
END.
```

## Which Index?

➤ transactionDate

## Why?



# 5

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH glrTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
  END.  
END.
```

## Which Index?

➤ transactionDate

## Why?

- Rule 3 yields two possibilities
- Both indexes contains the same fields and one is selected.



I cannot find a reference to this behaviour in the documentation, but it makes sense, as both indexes would render the exact same list of rowids, albeit in a different sequence. However, this construct is awkward and gains us nothing over example 4.

# 6

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE (transactionType = "INV" OR
         transactionType = "CRN")
         AND transactionDate >= 05/01/2015
         AND transactionDate <= 05/31/2015:
    ...
END.
```

## Which Index?



This is basically the same as example 4 – Equality on the type and range on the date, but here we have an OR

# 6

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK  
  WHERE (transactionType = "INV" OR  
         transactionType = "CRN")  
         AND transactionDate >= 05/01/2015  
         AND transactionDate <= 05/31/2015:
```

...

END.

## Which Index?

> **transactionDate**

## Why?



Not a good idea, as it will not leverage the second (transaction type) level of the index.

# 6

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE (transactionType = "INV" OR
         transactionType = "CRN")
         AND transactionDate >= 05/01/2015
         AND transactionDate <= 05/31/2015:
    ...
END.
```

## Which Index?

➤ **transactionDate**

## Why?

➤ Rule 4: Use the index with the most active range matches



Because the expression inside the ( ) must resolve first, we end up with WHERE clause that essentially boils down to a simple WHERE with AND, but the first one of the three expressions bound by AND is not a straight equality or range match.

# 6

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE (transactionType = "INV" OR
         transactionType = "CRN")
         AND transactionDate >= 05/01/2015
         AND transactionDate <= 05/31/2015:
    ...
END.
```

### Which Index?

➤ **transactionDate**

### Why?

➤ Rule 4: Use the index with the most active range matches

Can we fix this?



With the expected data distribution this is not too terrible, but we do utilize the index in an optimal way.

# 6

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

## Yes, we can fix it

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
    OR transactionType = "CRN"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015:
  ...
END.
```

## Index:

- transactionType (INV + Date range)
- transactionType (CRN + Date range)



We now have a WHERE with an OR and each OR fragment can be treated as an independent WHERE with AND.

This leads to the selection of the better index. Two brackets are made on the index and both levels of the index are leveraged.

# 7

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
    OR transactionType = "CRN"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
  USE-INDEX transactionType:
  ...
END.
```

## Which Index?



This is the same as before, is it not? After all, the compiler would have selected the index that we indicate in any case.

# 7

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
  OR transactionType = "CRN"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
  USE-INDEX transactionType:
  ...
END.
```

## Which Index?

➤ transactionType WHOLE-INDEX

## Why?



# 7

## glrTransaction

transactionID  
transactionDate  
transactionType

## Indexes

transactionDate  
+ transactionDate  
+ transactionType

transactionID  
+ transactionID

transactionType  
+ transactionType  
+ transactionDate

```
FOR EACH glrTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
    OR transactionType = "CRN"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015
  USE-INDEX transactionType:
  ...
END.
```

### Which Index?

➤ transactionType WHOLE-INDEX

### Why?

- Prime directive
- Only one bracket can be used



USE-INDEX allows for a single bracket. The only bracket that enclose all the options encompass the entire index. If we did not have the OR, i.e. if we had only the first or the last part, it would still bracket properly and the WHOLE-INDEX will go away.

## Mixing single-field and compound indexes

### glrTransaction

transactionID

transactionDate

transactionType

documentNo

transactionDate + transactionDate

\* transactionID + transactionID

typeDocNo + transactionType

+ documentNo



This set is fairly similar than the first one, except that the index with transactionType has a second component. We will now see how that change some rules.

# 8

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK
  WHERE transactionType = "INV"
    AND transactionDate >= 05/01/2015
    AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?

➤ typeDocNo

## Why?

➤ Rule 3: Use the index with the most active equality matches.



This is the same as in example 1.

# 9

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH gldTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

## Which Index?



This is essentially the same as the second case, where two single indexes was selected and combined.

# 9

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH gldTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

## Which Index?

➤ transactionDate

## Why?



# 9

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH gldTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

## Which Index?

➤ transactionDate

## Why?

➤ With AND, indexes do not combine unless all fields are used in equality matches.



The compound index is not combined with the single index, because the whole index must be used in equality matches.

# 10

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
DO vDate = 05/01/2015 TO 05/31/2015:  
  FOR EACH gldTransaction NO-LOCK  
    WHERE transactionType = "INV"  
      AND documentNo      = "I2300121"  
      AND transactionDate = vDate:  
      ...  
    END.  
  END.
```

## Which Index?

- transactionDate
- typeDocNo

## Why?

- All fields are used in equality matches, neither index is unique.



If the full index participated in the equality match, it would indeed combine with the other index, as happened in example 2.

# 11

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK
  WHERE transactionType = "INV"
     OR transactionDate >= 05/01/2015
     AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?



Note that this (rather unlikely) query contains an OR.

# 11

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK
  WHERE transactionType = "INV"
     OR transactionDate >= 05/01/2015
     AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?

- typeDocNo
- transactionDate

## Why?



# 11

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK
  WHERE transactionType = "INV"
     OR transactionDate >= 05/01/2015
     AND transactionDate <= 05/31/2015:
  ...
END.
```

## Which Index?

- typeDocNo
- transactionDate

## Why?

- Rule 3 for the first part
- Rule 4 for the second part



Each side of the OR is evaluated as a separate query and gets its own index bracket. Note that if any one of these resulted in a full index scan of any index, only that index would be used, as all the records will be accessed in any case.

# 12

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FIND FIRST gldTransaction NO-LOCK  
WHERE transactionType = "INV"  
OR transactionDate > 05/01/2015  
NO-ERROR.
```

## Which Index?



What will happen when we use a FIND?

# 12

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FIND FIRST gldTransaction NO-LOCK  
WHERE transactionType = "INV"  
OR transactionDate > 05/01/2015  
NO-ERROR.
```

## Which Index?

➤ transactionID – WHOLE-INDEX

## Why?



# 12

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FIND FIRST gldTransaction NO-LOCK  
WHERE transactionType = "INV"  
OR transactionDate > 05/01/2015  
NO-ERROR.
```

## Which Index?

➤ transactionID – WHOLE-INDEX

## Why?

➤ Rule 7 Use primary



FIND allows for the use of a single index bracket only. Of the two candidate indexes that was used in example 11, neither one will satisfy the other side of the OR, hence both gets eliminated early on and the compiler falls back to rule 7.

# 13

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK  
  WHERE documentNo = "123001214":  
  ...  
END.
```

## Which Index?

➤ transactionID – WHOLE-INDEX

## Why?

➤ Rule 7 Use primary



Once again, all indexes are eliminated early on, and the compiler falls back to its default of scanning the primary index.

# 14

gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK  
  WHERE transactionType > ""  
        AND documentNo    = "I23001214":  
  
  ...  
END.
```

## Which Index?

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

I often see this previous problem solved using code as above. Will it work?

# 14

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK  
  WHERE transactionType > ""  
        AND documentNo      = "I23001214":  
  
  ...  
END.
```

## Which Index?

➤ typeDocNo

## Why?



# 14

## gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

## Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

```
FOR EACH gldTransaction NO-LOCK  
  WHERE transactionType > ""  
        AND documentNo    = "I23001214":  
  ...  
END.
```

## Which Index?

➤ typeDocNo

## Why?

➤ Rule 4 Use index with most active range matches.

We made it worse but hides  
he fact!

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

The index with the most range matches was selected. This means that components after the range matching component (transactionType) are ignored.

Since anything is > "", we have a range here that spans the whole index, but the compiler does not realize this, as it has a "value" for the lower boundary of the bracket, so the WHOLE-INDEX is missing from the cross reference. But this index has two fields. Thus the index tree consists of all the rowids in the table plus all the transaction types plus all the document numbers, which could in total actually occupy more disk space than the primary index, which consists of all the rowids and all the transactionIDs, in which case we are now reading more blocks from disk and still access the entire table. Note however that in this design we expect the transactionID to be large, random values that will compress poorly, while the document type and number should lead to fairly good index compression, so this might still be the better option.

15

gldTransaction

transactionID  
transactionDate  
transactionType  
documentNo

Indexes

transactionDate  
+ transactionDate

transactionID  
+ transactionID

typeDocNo  
+ transactionType  
+ documentNo

### Will the compound index solve this?

```
FOR EACH gldTransaction NO-LOCK
  WHERE (transactionType = "INV" OR
         transactionType = "CRN")
         AND transactionDate >= 05/01/2015
         AND transactionDate <= 05/01/2015:
    ...
END.
```

No

➤ transactionDate

**PUG**  
**CHALLENGE**  
**EXCHANGE**  
**AMERICAS**

Once again, we have a range match bound with AND to an OR function, hence the index with the most range matches is selected.

Questions?



Feedback and more examples welcome

Simon Prinsloo

[simon@vidisolve.com](mailto:simon@vidisolve.com)

