



"Become a star in just one hour"

thilo.lauer@hp.com

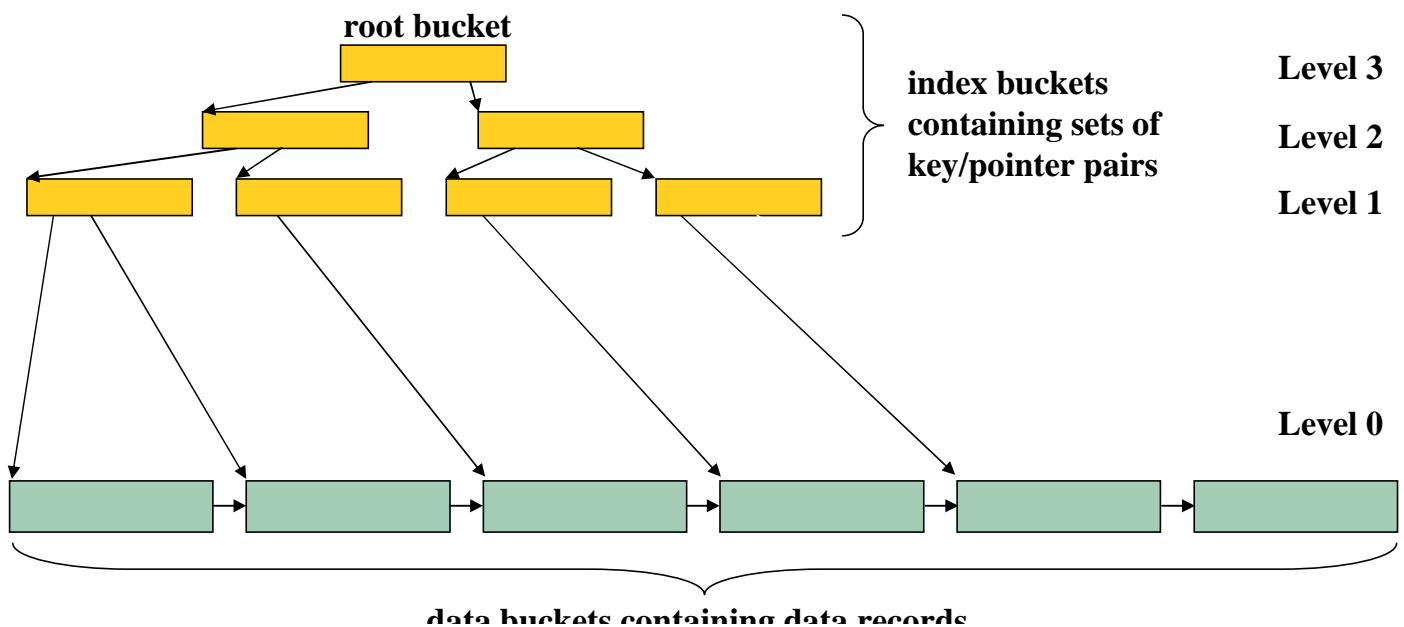
Topics



- Make up your mind!
- "What's in a name?" - some terminology
- Learn application behaviour
- Use the right tools
- The RMS Tuning Cookbook
 - things to try blindly
 - sophisticated techniques

- RMS doesn't use magic (yet)
- RMS behaviour is predictable
- most changes are transparent to applications
- many misconceptions floating around
- you don't even need to understand it - just do it!
- it's open-ended: you want more? Give us a call!

Some Terminology



- hot files
- read/write (update, delete) ratio
- file growth rate
- locality of new records
- key patterns
- shared access, number of users
- when does performance suffer?
- gather metrics

Use the right tools



- OpenVMS-supplied
 - \$MONITOR (/RMS) to analyze I/O behaviour
 - \$ANALYZE/RMS (/FDL, /STATISTICS) to analyze a file's permanent attributes
 - \$EDIT/FDL to change internal file attributes
- Other tools (from OpenVMS freeware CD)
 - RMS_STATS: similar to \$MONITOR, but displays raw numbers, more details
 - SIDR: shows count and values of duplicate keys

- Apply CONVERT regularly
- Increase bucket size
- Enable Global Buffers for shared files
- Evaluate compression settings
- Watch out for duplicates
- Adjust fill factor for growing files
- Consider removing secondary keys

Apply CONVERT regularly



`$CONVERT input-file output-file`

- reorganizes internal layout
- removes internal defragmentation
- (removes external defragmentation)
- removes levels of indirection (RRVs)
- removes deleted records

1 change bucket size parameter in FDL file

– let OpenVMS be smart:

```
$EDIT/FDL/NOINTERACTIVE fdlfile.FDL
```

– be smart yourself:

calculate better values by hand and edit the FDL-file manually

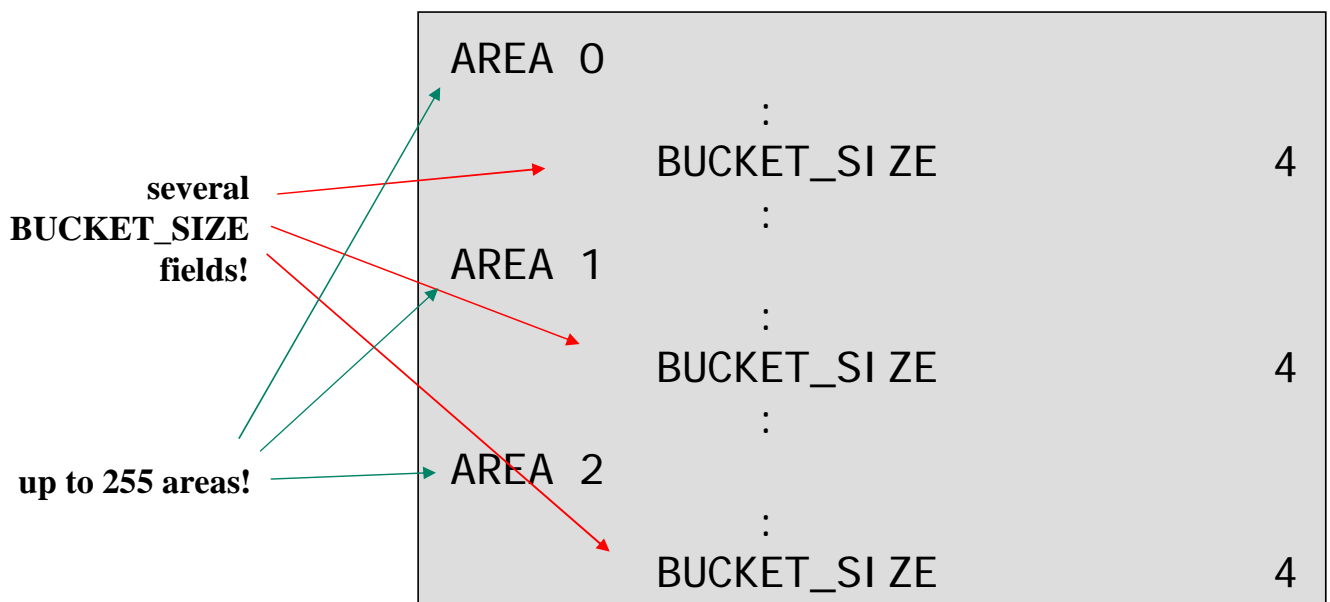
2 convert the file with the new FDL file

– `$CONVERT/FDL=fdlfile infile outfile`

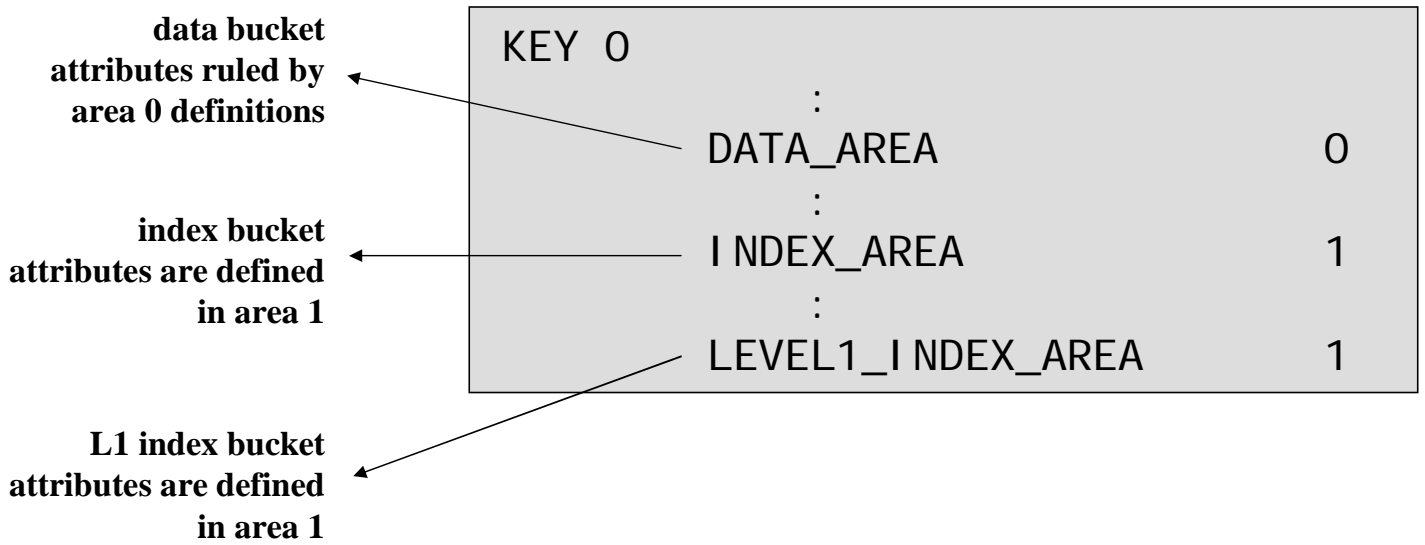
Increase bucket size - what to change



- which field should i change?




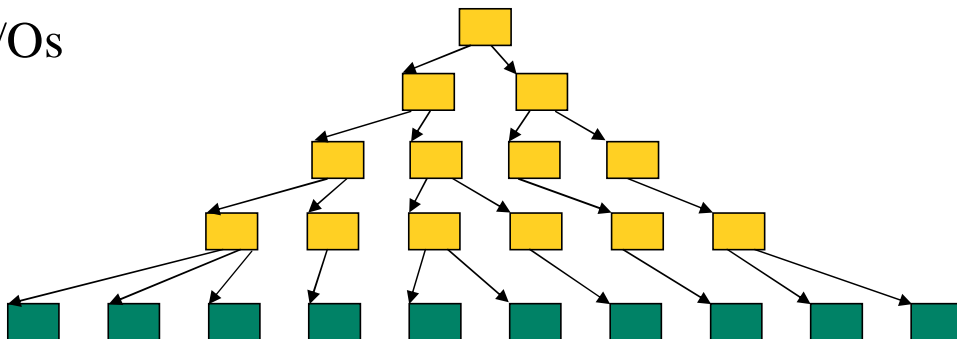
- get the area information from the key definition:




Increase bucket size - potential effect



old: 5 I/Os



new: 3 I/Os



To judge the effectiveness of your changes, examine the index depths of the resulting file:

```
$ pipe analyze/rms/statistics newindexfile | -  
_$_ search sys$input/window=(0,2) "STATISTICS FOR KEY"  
  
STATISTICS FOR KEY #0  
  
Number of Index Levels: 4
```

This number should have decreased compared to the original file!

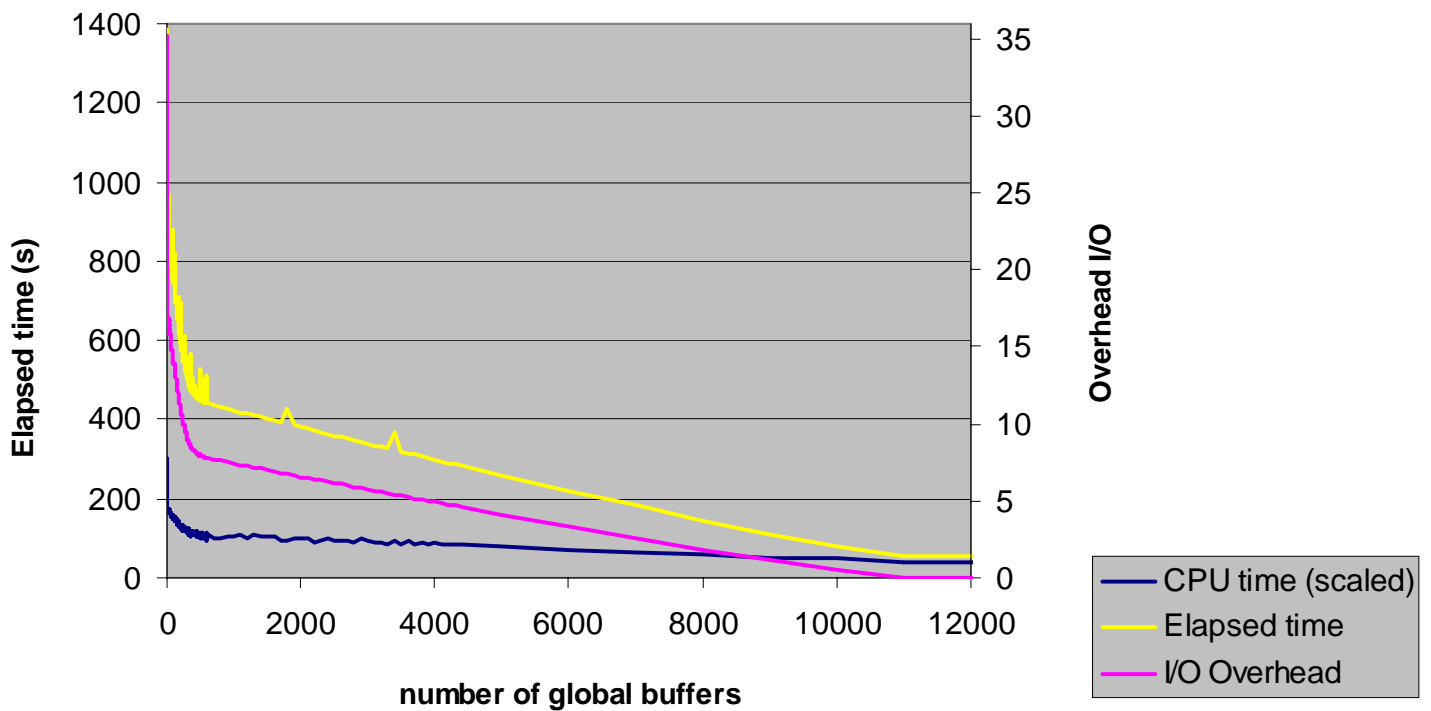
Enable Global Buffers for shared files



```
$SET FILE/GLOBAL_BUFFER=n filename
```

- probably the easiest RMS tuning method with the highest effects!
- forget all rumours about global buffers
- forget former pitfalls
- adjust GBLPAGES, GBLPAGFIL (dynamic since V7.1)
- adjust GBLSECTIONS

Effects of global buffers (100000 random reads, data+index=11051 buckets)




Evaluate compression settings



- concentrate on INDEX_COMPRESSION (FDL)
- default has been ON for historical reasons (disk space), changed in recent OpenVMS versions
- performance-wise, OFF would be better choice, but watch out for introducing additional index levels!
- easy as well:
 - change parameter in FDL file
 - perform CONVERT/FDL
 - enjoy!

traversal of index bucket

with compression: sequential search, performance: $N/2$



500 entries:
~250 lookups

without compression: binary search, performance: $\log_2(N)$



500 entries:
~9 lookups!

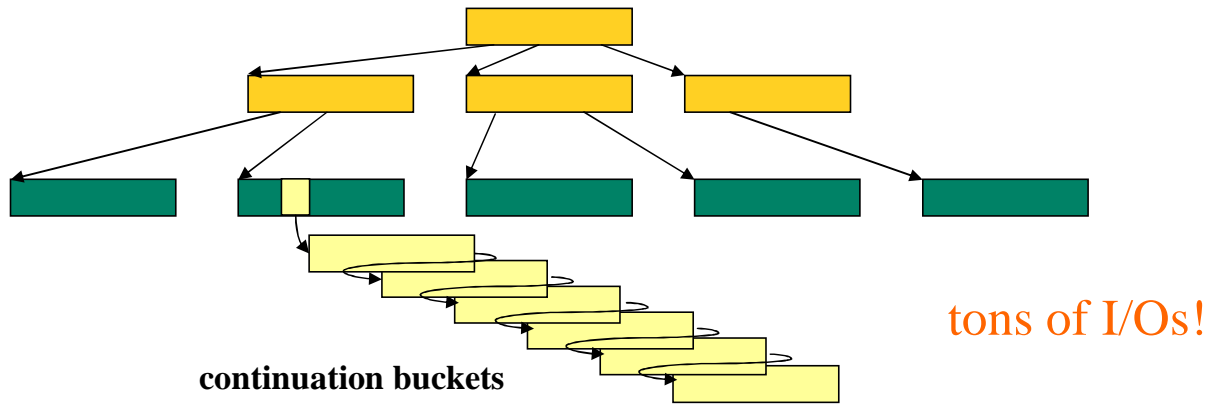
Result: less CPU overhead!

Evaluate compression settings



Recent performance tests have shown only marginal CPU-time improvements (<5%). Now, is this "improvement" worth thinking about index compression at all?

- Yes: for specific files with large index buckets (>30?), short keys (<10?), and a deep index (>2?).
- Having plain index entries might ease potential troubleshooting.
- Disabling index compression leads to increased space



- duplicate records are not maintained in the index
- sequential search through continuation buckets
- RMS preserves order of arrival: most recent record is at the end
- insertion of new duplicate is very expensive
- duplicates are often your most-frequent key value (default value?)

Watch out for duplicates - real-life example



Output of SIDR (file has ~46000 records):


Duplicate count	Buckets	Key value
17824	19	[REDACTED]
2982	4	OHNE
1055	1	ohne
173	1	-
139	1	ENTF. LLT
132	1	ENTFAELLT
127	1	ENTF. LLT
82	1	ENTF.
...

In FDL file: define SPACE character as NULL key value


KEY 3

CHANGES	yes
DATA_KEY_COMPRESSION	no
DATA_AREA	2
DATA_FILL	100
DUPLICATES	yes
:	:
LEVEL1_INDEX_AREA	2
NULL_KEY	yes
NULL_VALUE	' '
SEGO_LENGTH	13
SEGO_POSITION	40
TYPE	string

Watch out for duplicates - effect



- this particular key value is not maintained in the index
- records with this key value are not seen via indexed access




Adjust fill factor for growing files - solution




In FDL file: define Fill Factor for relevant areas

```

KEY 3      :
           DATA_FILL      : 50
           :
           LEVEL1_INDEX_AREA : 2
           NULL_KEY         : yes
           NULL_VALUE       : ' '
    
```



An RMS indexed file with 6 keys:

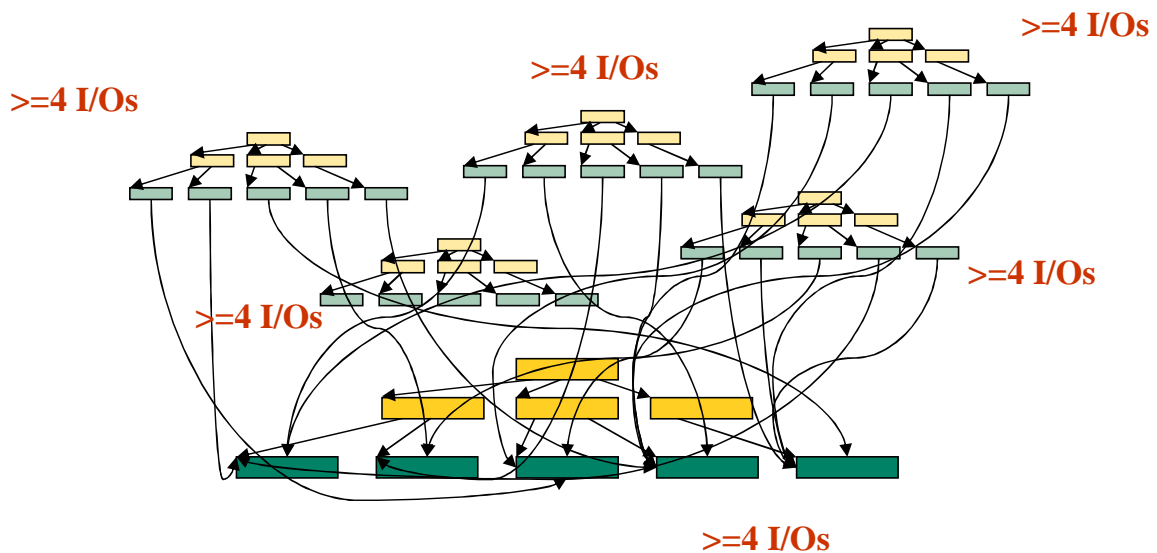


All this chaos works perfectly - but slowly!!

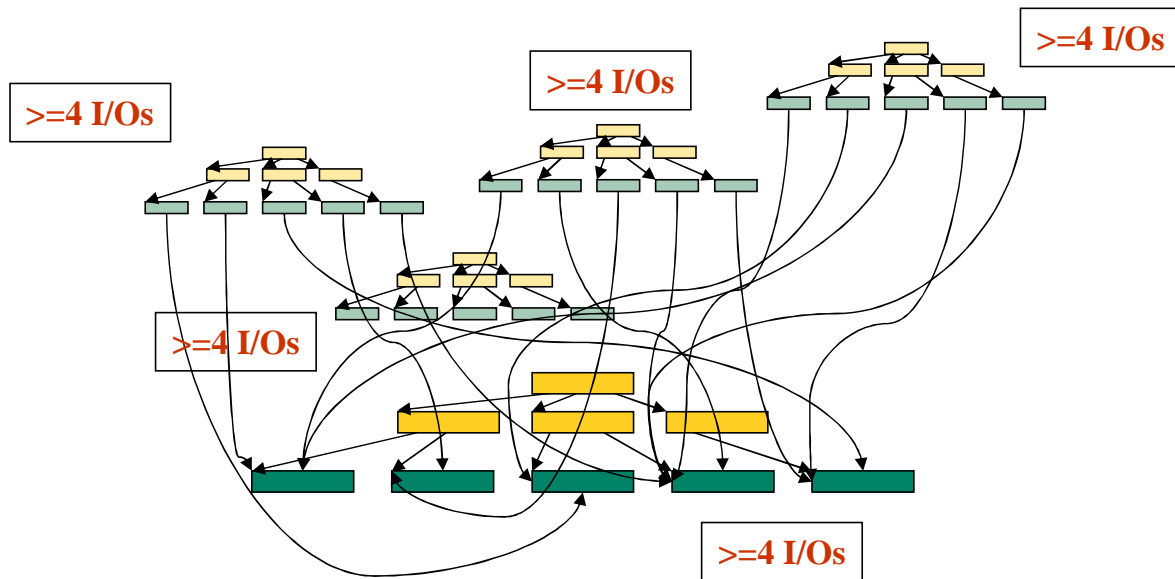
...Consider removing secondary keys



insert just one record:



same file, one key removed:



Removal of one index reduces administrative work (I/O+CPU)

...Consider removing secondary keys



Issues:

- application designers often implement keys 'just to have them handy'
- example: two identical key fields with different sorting order
- often secondary keys have poor 'duplicate behaviour'
- evaluate application impact!
- substitute keyed lookup by other methods
- You may need to change the application!

- Think about employing an RMS DB administrator
- Ask ~~Compaq~~: we sell more than just ~~PCs~~!

HP

Printers

- (RMS) Engineering is open to all sorts of problem reports, enhancement requests, discussions...



i n v e n t

Questions...

thilo.lauer@digital.com
thilo.lauer@compaq.com