

HCC – A turbo seldom recognized as such



Nuremberg, 11/18/2015, C. Trieb, P. Böhme

Personal - Experienced - Efficient

- 60 employees at headquarters in Friedrichsdorf and Munich
- Turnover of €11 millions
- Industries: retail, logistics, public sector, pharma, banks
- Flat hierarchies, open communication, flexibility
- Comprehensive consulting with a focus on practice
- Individually optimized solutions
- We are a fair, professional partner on par with you
- Strong partner network



Specialized
Oracle Enterprise Manager 12c



Specialized
Oracle Real Application
Clusters 11g



Specialized
Oracle Linux 6

Oracle in Use

- ORACLE Real Application Clusters
- ORACLE Text
- ORACLE Data Guard
- ORACLE Partitioning
- ORACLE Enterprise Manager Cloud Control
- Oracle Exadata
- Oracle ZFS Appliance ZS3-2

Speakers

- Patrick Böhme
 - DBA
 - OCS

- Christian Trieb
 - DBA
 - OCS, OCA, OCP, OCE



Overview

Hugendubel

Die Welt der Bücher

 PARAGON DATA

Turnover: €11 M.

Employees: 60

Strong partner network



Fresh. Fair. Fullservice!

tolino  media 



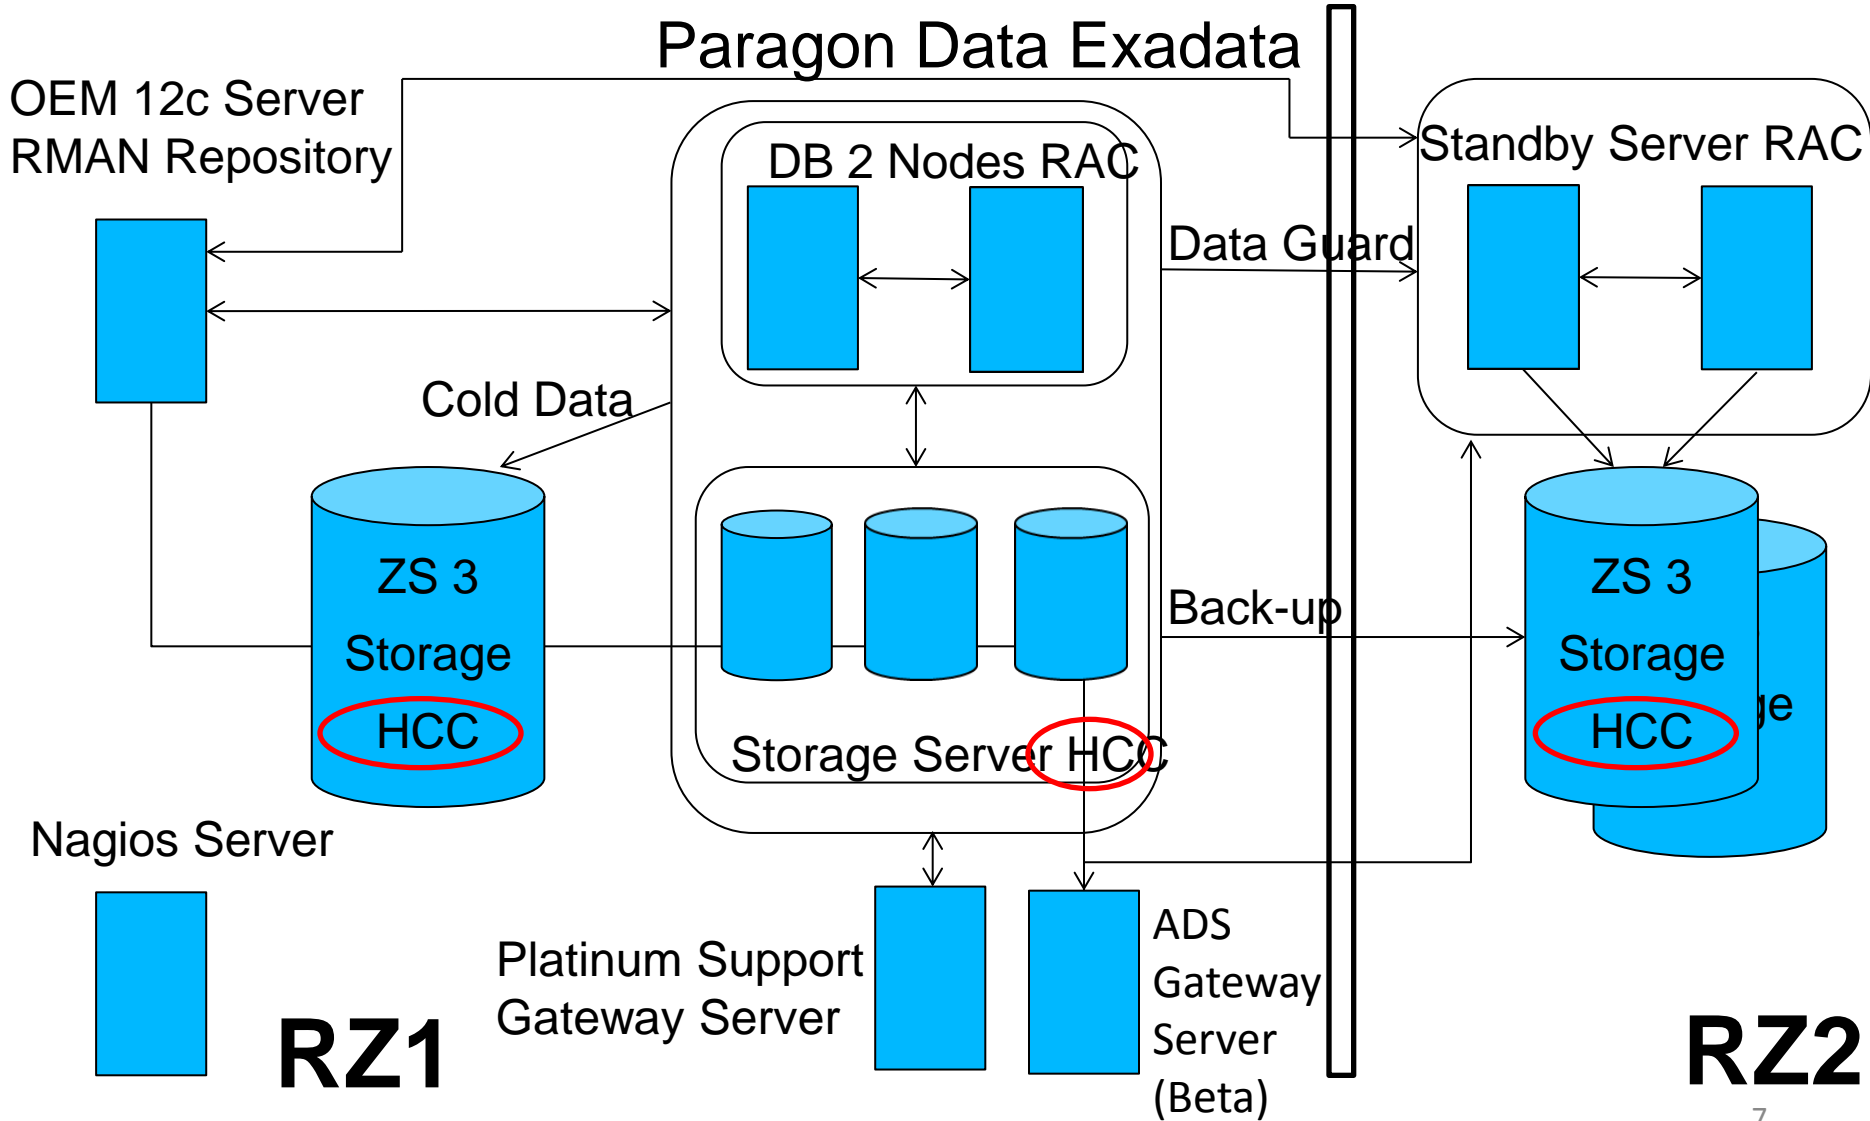
Weltbild *plus*



Overview

- Paragon's architecture
- Use of HCC
- HCC in a nutshell
- HCC Fields of application
- Summary

Engineered Architecture



Reasons for and Ways of Compression

Saving disk space

Improving performance through faster I/O

⇒ Reducing costs

Ways of compression:

- Index Key Compression
- BASIC
- OLTP (Advanced Compression Option)
- HCC Query Low
- HCC Query High
- HCC Archive Low
- HCC Archive High

Hybrid Columnar Compression – When and what for?

BASIC

-Only for bulk load operations

OLTP (Advanced Compression Option)

-To be paid extra

Warehouse Compression (optimized for performance):

HCC Query Low

HCC Query High

Archive Compression (optimized for storage savings):

HCC Archive Low

HCC Archive High

Combination of column- and row-based data storage Data sets are organized into logical compression groups and sorted into columns Repeated data is grouped in the same blocks

Hybrid Columnar Compression – When and what for

- Data with similar characteristics
- Bulk load operations
 - Parallel DML
 - CTAS
 - SQL*LDR
 - Insert with Append Hint
- Many read operations
- Archived data

Compression with HCC

HCC

- Works only with Exadata or ORACLE Storage
- Not possible with the data type “Long“!
- Great benefits / high compression ratio only with bulk load operations
- Lower compression ratio with INSERT / UPDATE
- Tables are organized into compression units, a compression unit is larger than an ORACLE block
- Within compression units, data is organized by column instead of by row => similar values together, enhancing compression

Müller	Müller	Schultz	Meier
Müller			
Müller			
Müller			
Schultz			
Schultz			
Schultz			
Meier			
Meier			

Hybrid Columnar Compression Syntax

HCC Query Low

- Create Table ... column store compress for query low ...;
- Alter Table ... column store compress for query low ...;

HCC Query High

- Create Table ... column store compress for query high ...;
- Alter Table ... column store compress for query high ...;

HCC Archive Low

- Create Table ... column store compress for archive low ...;
- Alter Table ... column store compress for archive low ...;

HCC Archive High

- Create Table ... column store compress for archive high ...;
- Alter Table ... column store compress for archive high ...;

Hybrid Columnar Compression Syntax

In that case the “Alter Table” command only applies to data sets that have been newly inserted afterwards.

Therefore the existing data needs to be compressed explicitly:

- Alter Table ... move compress for ...;

Undo:

- Alter table ... nocompress;

How do I know if a table was compressed?

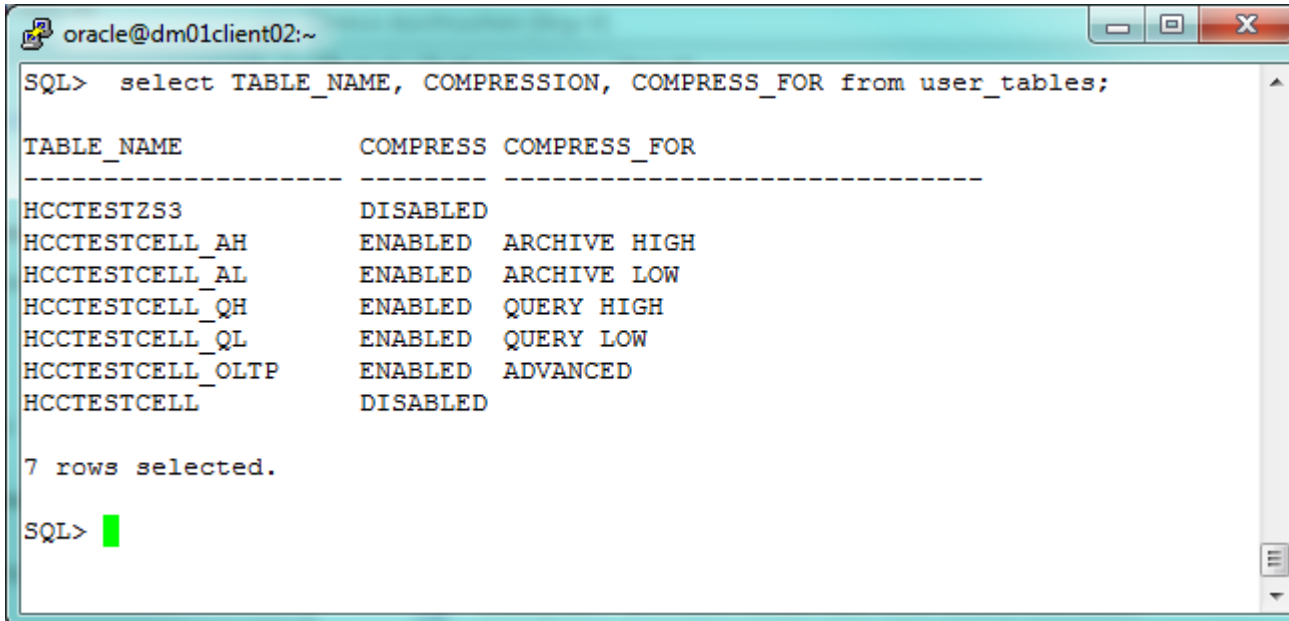
Check which type of compression there is:

select TABLE_NAME, COMPRESSION, COMPRESS_FOR from user_tables;
(or also dba_tables)

TABLE_NAME: Description of the object

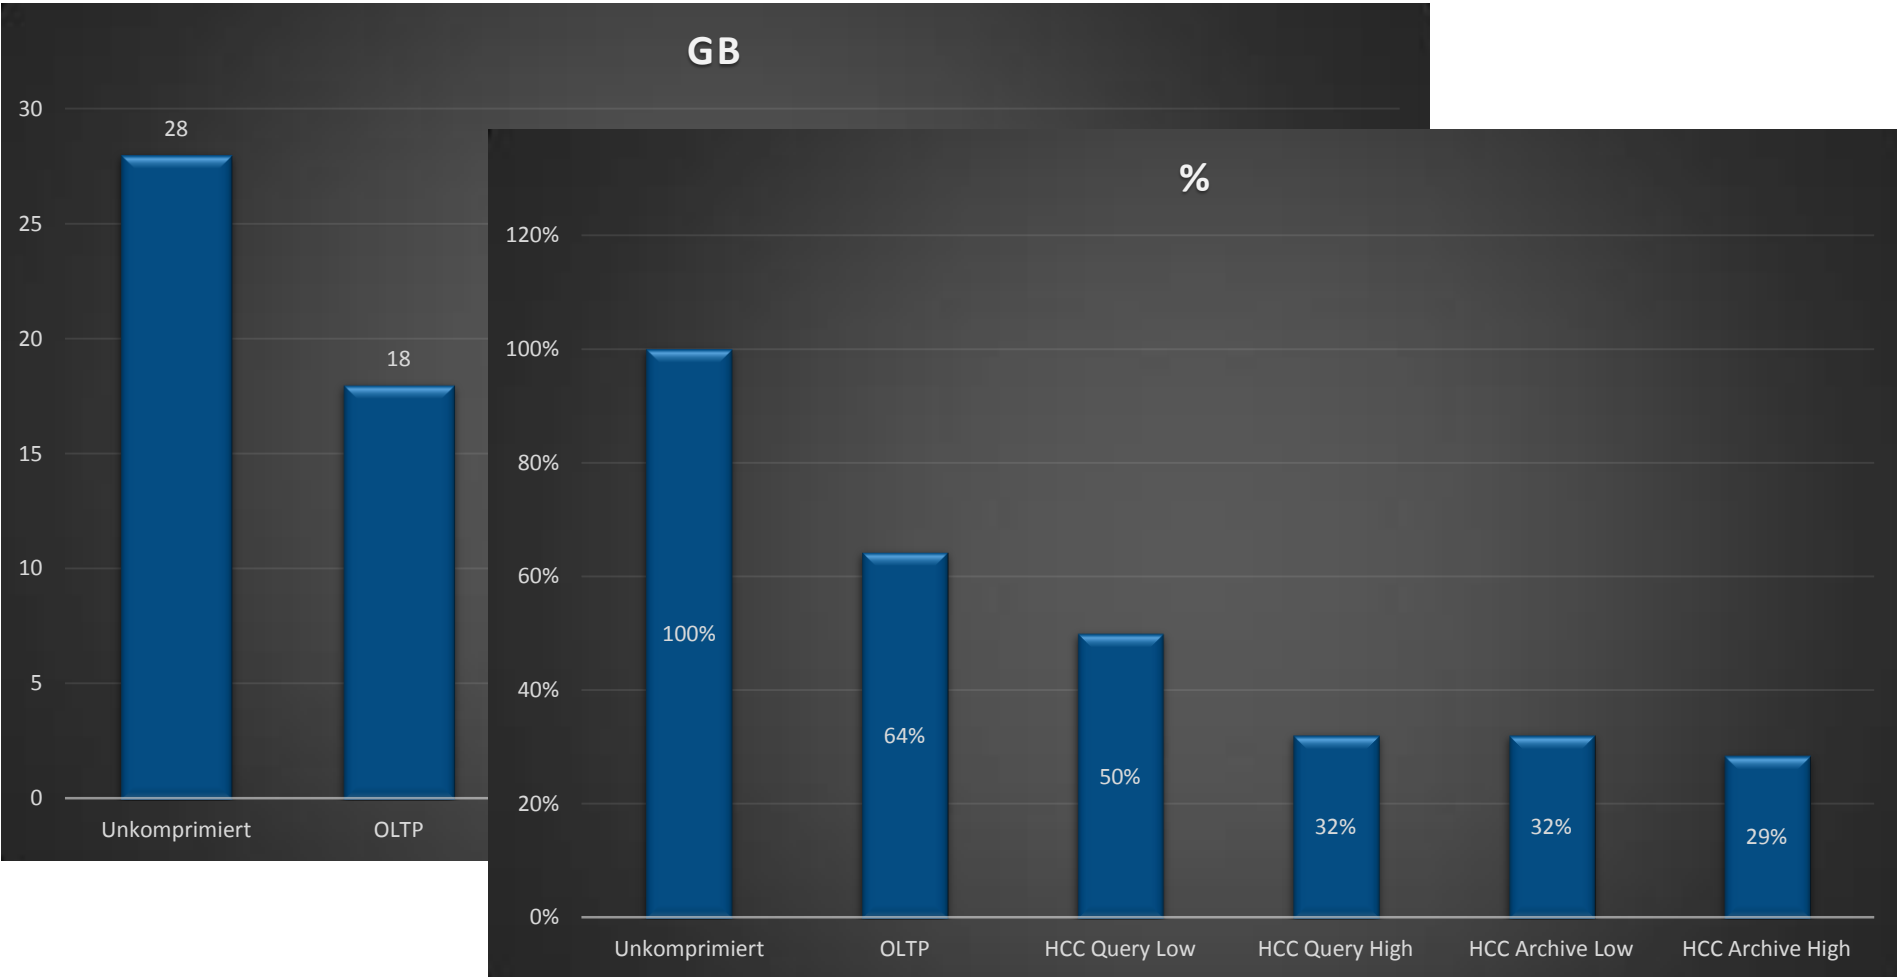
COMPRESSION: DISABLED / ENABLED

COMPRESS_FOR: Type of compression, e.g.: BASIC / ADVANCED / QUERY HIGH / ARCHIVE HIGH

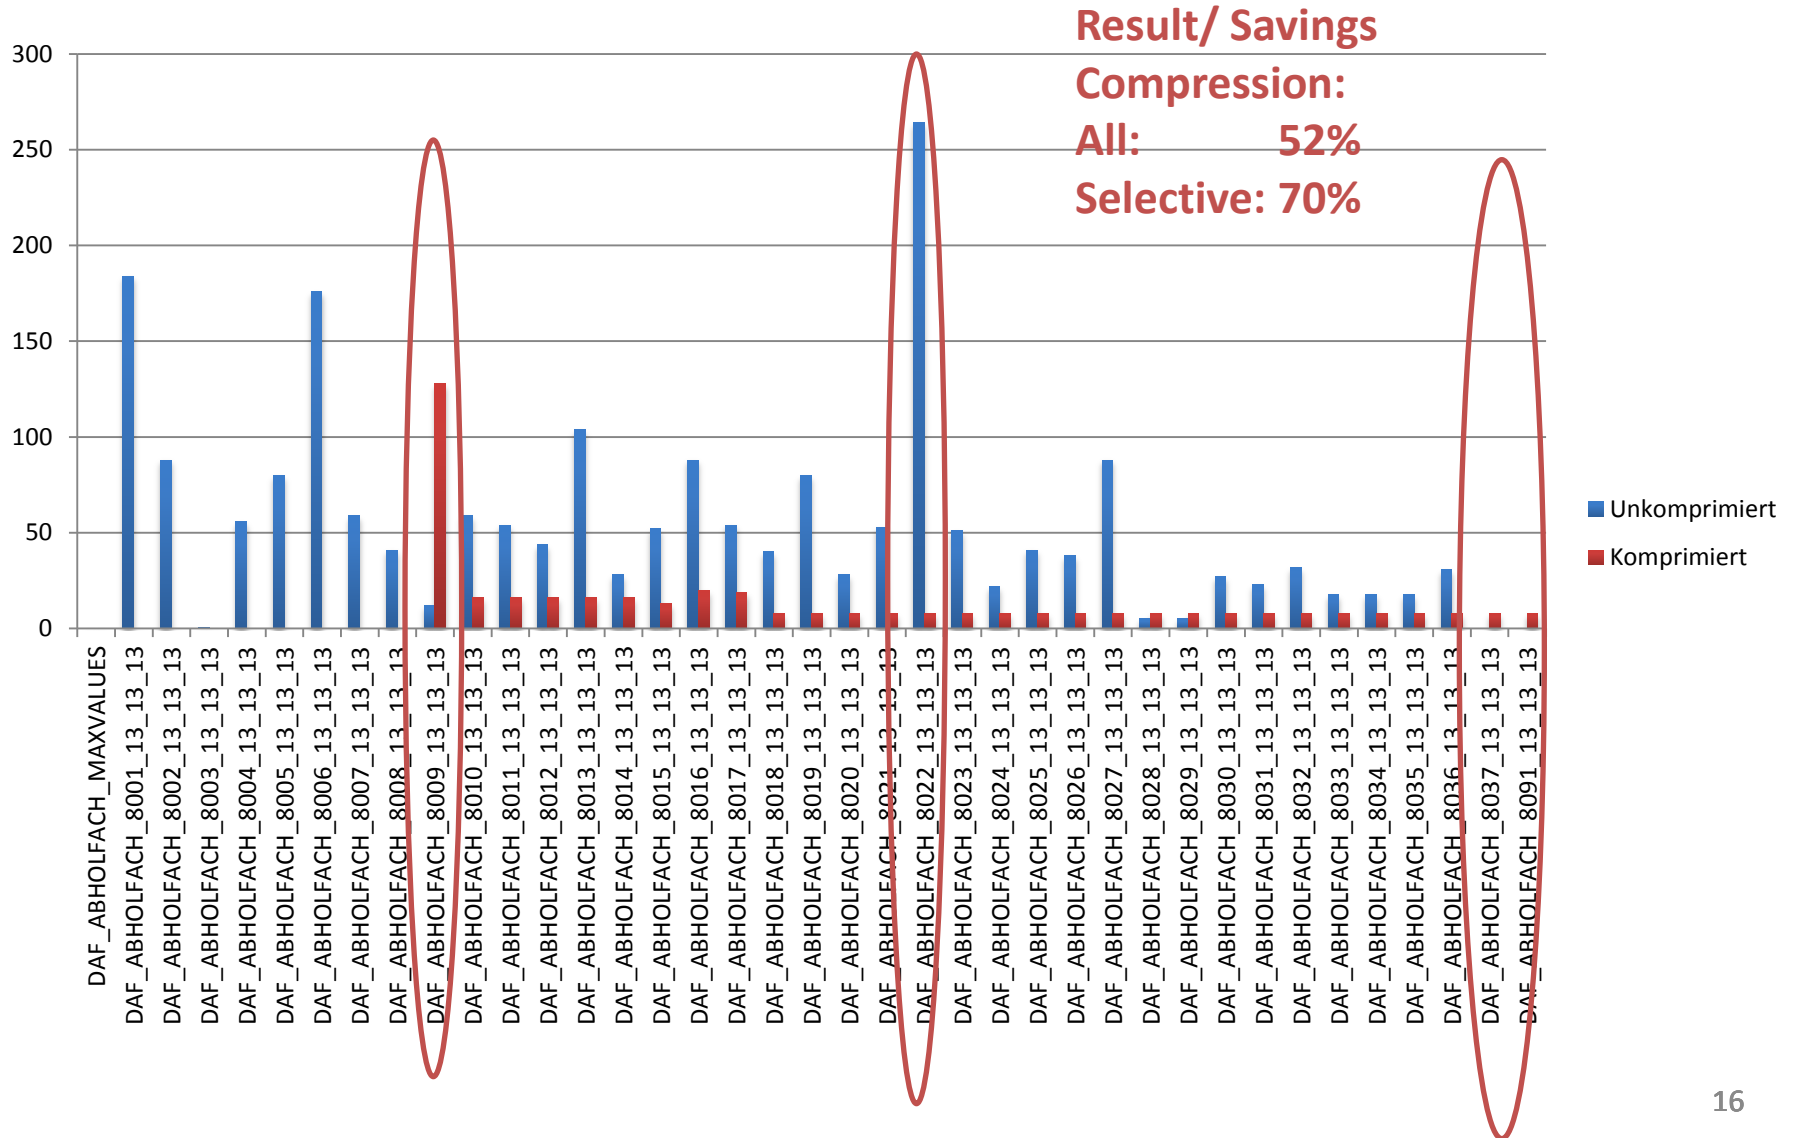


```
oracle@dm01client02:~  
SQL> select TABLE_NAME, COMPRESSION, COMPRESS_FOR from user_tables;  
  
TABLE_NAME          COMPRESS COMPRESS_FOR  
-----  
HCCTESTZS3          DISABLED  
HCCTESTCELL_AH      ENABLED  ARCHIVE HIGH  
HCCTESTCELL_AL      ENABLED  ARCHIVE LOW  
HCCTESTCELL_QH      ENABLED  QUERY HIGH  
HCCTESTCELL_QL      ENABLED  QUERY LOW  
HCCTESTCELL_OLTP    ENABLED  ADVANCED  
HCCTESTCELL         DISABLED  
  
7 rows selected.  
  
SQL> █
```

Compression Overview



HCC – Results of a Partitioned Table



HCC – The Success Story

Bibliography/ catalogue system:

1 master for processing

10 slaves for read applications

3500 users at the same time

Very good response times of the application expected

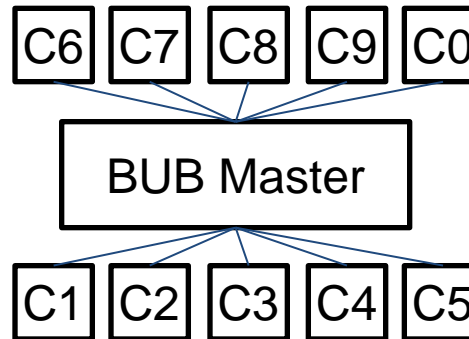
**Oracle Text full-text search + Oracle snapshots/
materialized views**

Consolidation BUB on Exadata?


**Oracle Text → almost only block IO
performance?**

Oracle HCC on Exadata → fewer reads

PoC



**High effort for
hardware and
administration!**




**Application without
HCC:
critical performance**

**with HCC:
Performance gain of 30%**

**Less hardware
but standby and HCC?
→ Oracle ZS3-2!**

What kind of hardware is needed for HCC?

Works only on Engineered Systems
and ORACLE Storage
But without additional
licensing fees there



ZFS Storage ZS3-2



Oracle FS1 Flash Storage System




Exadata Database Machine

What can HCC achieve?

Saving disk space

Improving performance

=> Reducing costs



Q & A

Thank you!