Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

Compression for Informix: What are the benefits and impacts?

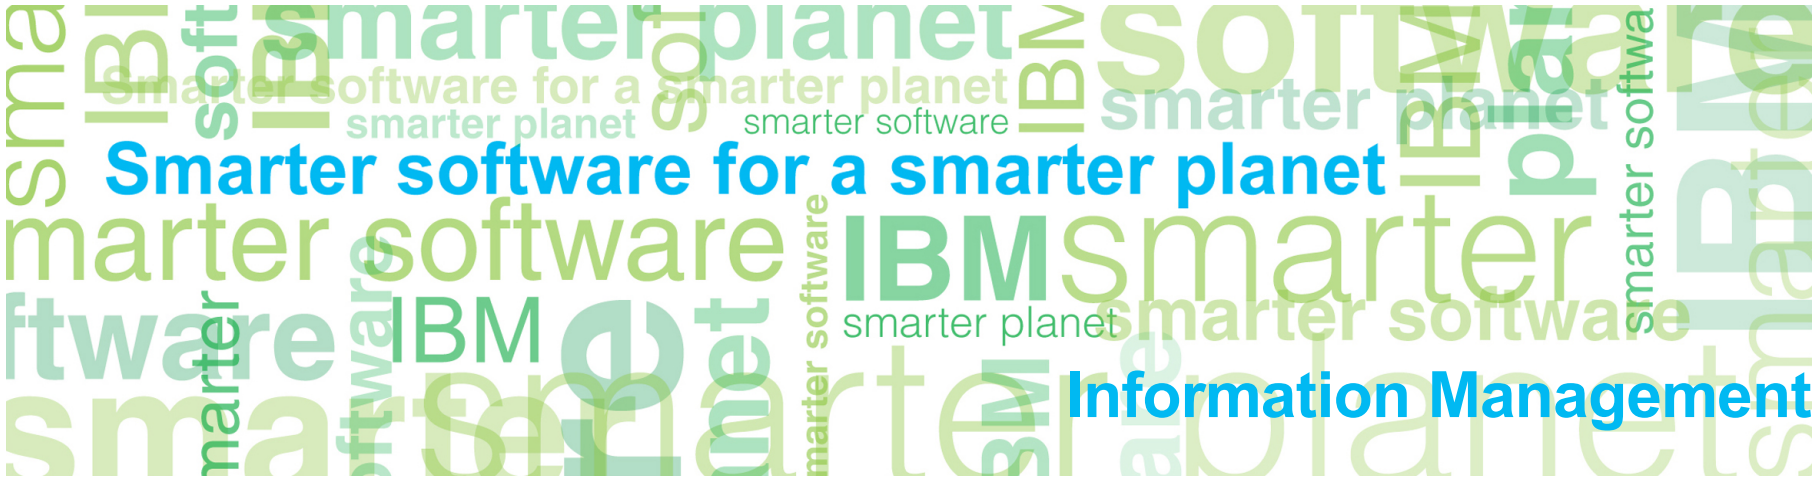

Agenda

- What is compression for Informix?
- What are the benefits of compression?
- How does compression work?
- Tooling for Compression
- How much space is saved?
- What is the performance impact?
- How is the impact of compression on other technologies and vice versa?
- Best practices
- Summary / Recommendations

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

What is Compression for Informix?

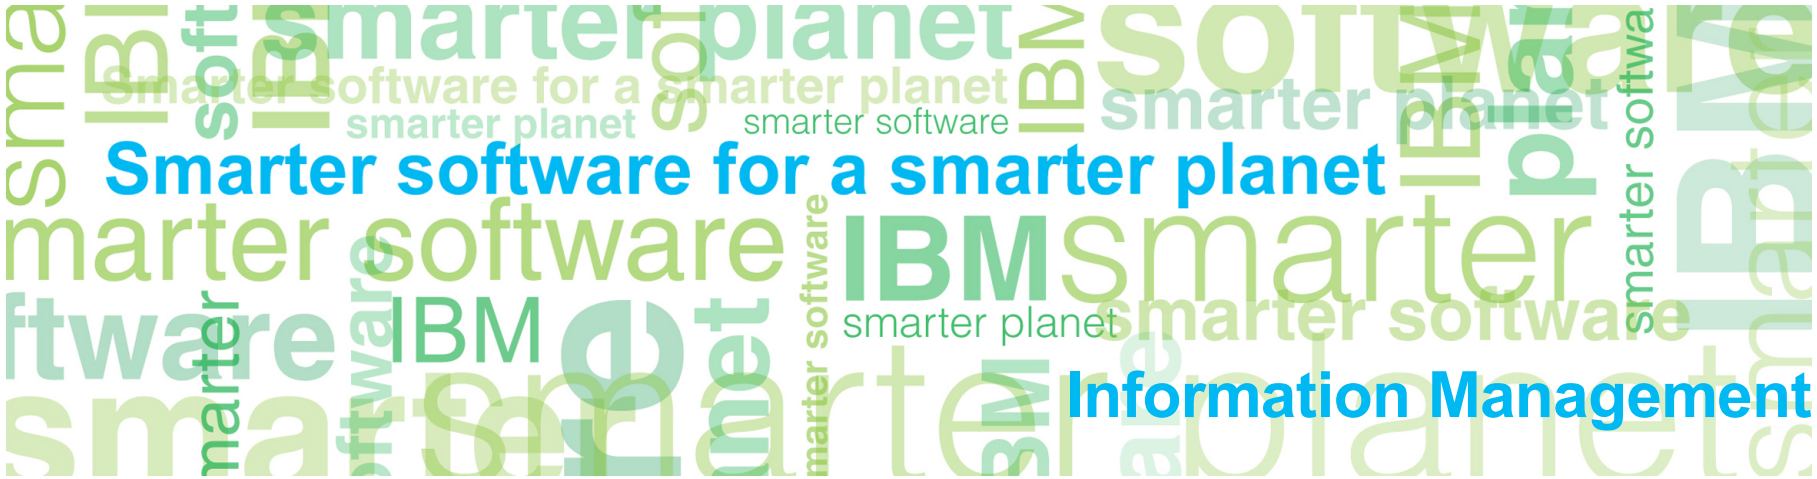

What Is Informix Compression?

- Ability to store data rows in compressed format on disk.
- Saves up to 90% of row storage space.
- Ability to estimate a possible compression ratio.
- Fits more data onto a page.
- Fits more data into the buffer pool.
- Reduces logical log usage.
- Saves lots of money by reducing the physical storage size of the data.
- Transparent for application

Availability

- Available as optional Storage Optimization Feature (SOF)
- Same license metric as server
- Since Informix Dynamic Server 11.50 xC4
- Since Extended Parallel Server 8.51 FC3
- Available on all supported platforms

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

What are the benefits of compression?

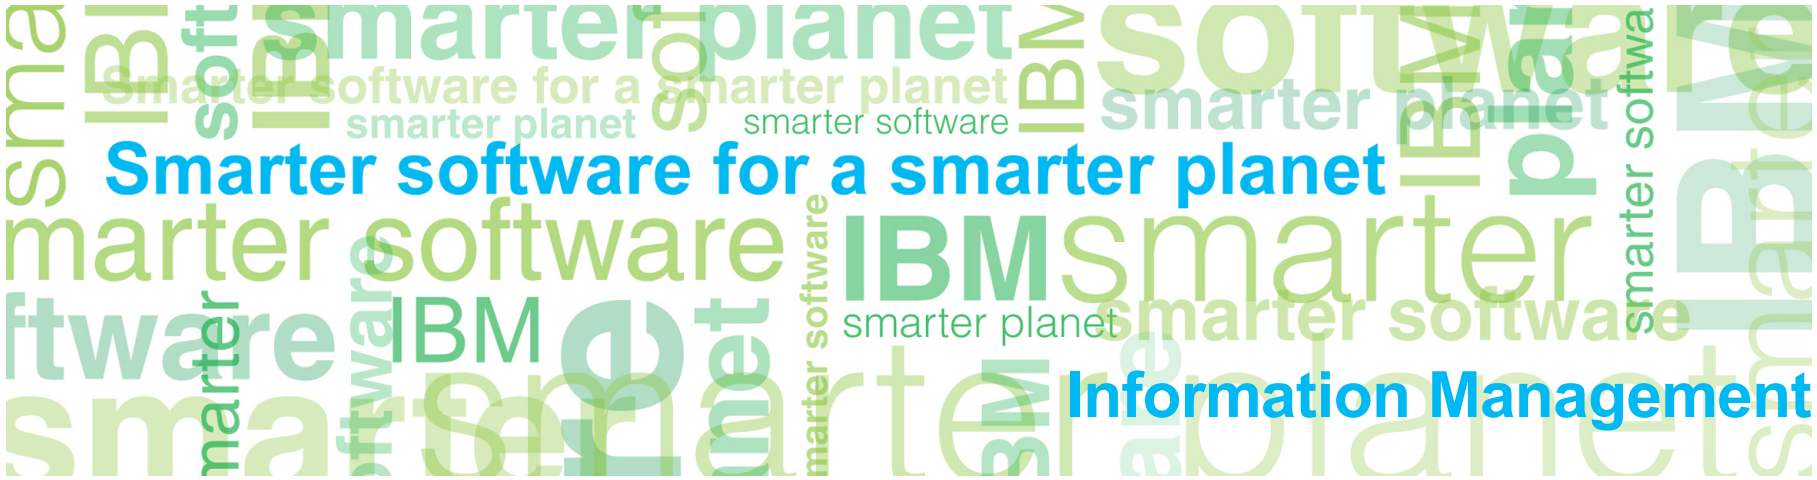

Benefits of Compression

- Cost Savings because of space savings
  - Memory savings
  - Disk savings
  - Tape savings
- Performance Improvements
  - OLTP environments:
    - More rows in bufferpool
    - faster technology like SSDs becomes affordable for critical data
  - Data Warehousing environments:
    - faster scans i.e. better use of IO bandwidth
  - Maintenance:
    - faster backup / restore, less overhead compared to backup compression

Memory Savings

- Bufferpool contains compressed pages
- Rows get uncompressed for processing
  - Data in Decision Support Memory (Sort, Hash Joins etc) not compressed
- Cost:
  - cost of acquisition
  - power consumption (roughly 40 kWh per GB memory per year; large variation based on memory chips used)
  - air conditioning

Disk Savings

- All table fragments with more than 2000 rows can be compressed
- More saving when many redundant copies of data are kept (e.g. HDR, RSS, RAID)
- Cost savings:
  - cost of acquisition
  - power consumption (roughly 60 kWh per drive per year)
  - air conditioning, floor space
  - administration

Tape Savings

- Same compression rate as for data on disk
- Applies to all copies (e.g. disk cache, all versions kept, etc.)
- No CPU overhead (especially compared to backup compression) since no processing of data during backup / restore
- Cost savings:
  - fewer drives
  - fewer media
  - fewer network bandwidth

Performance OLTP Environment

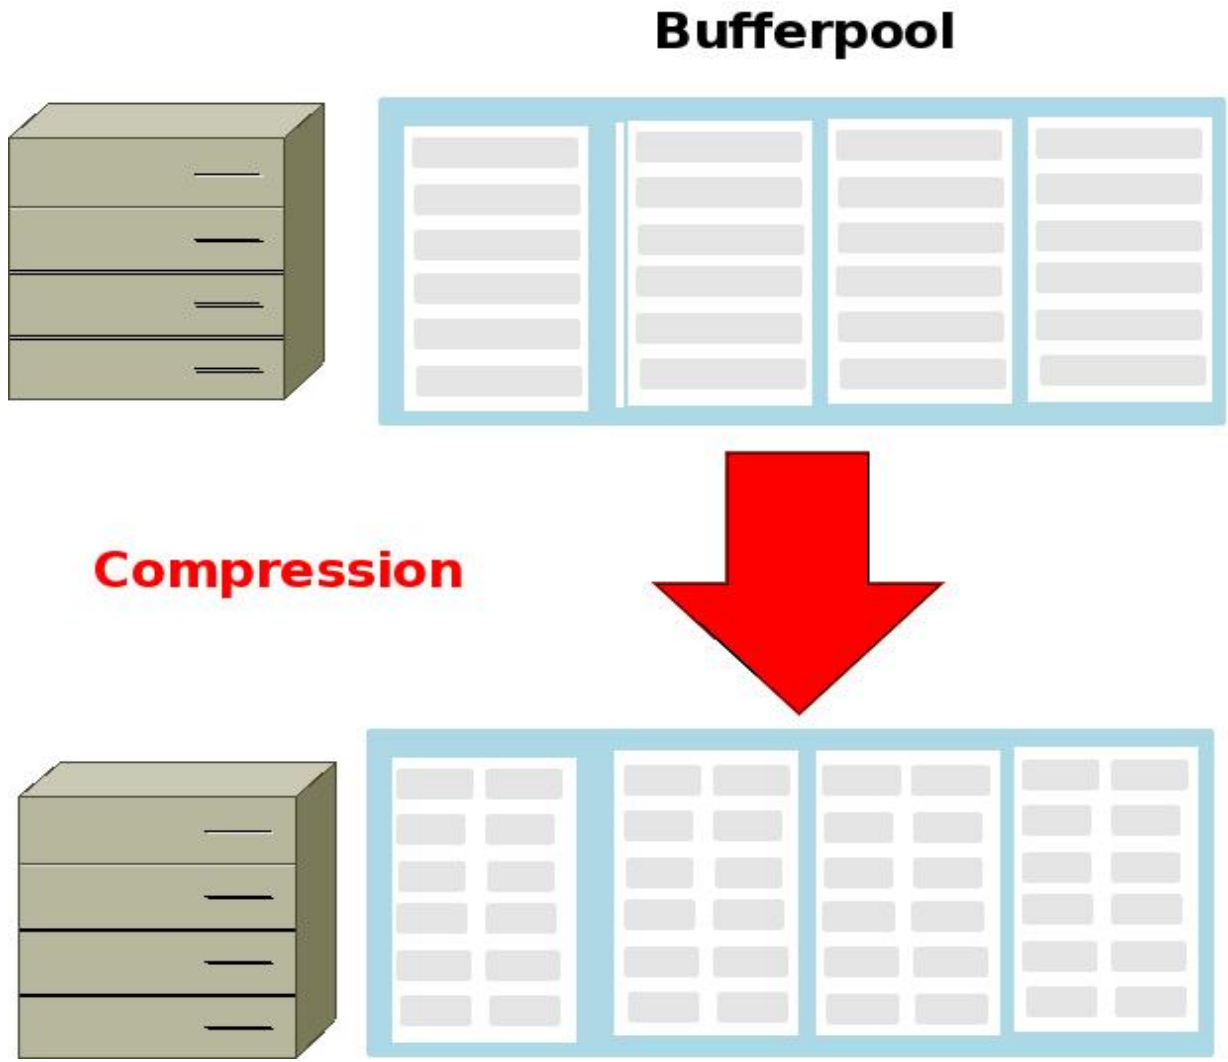

Performance Data Warehousing Environment

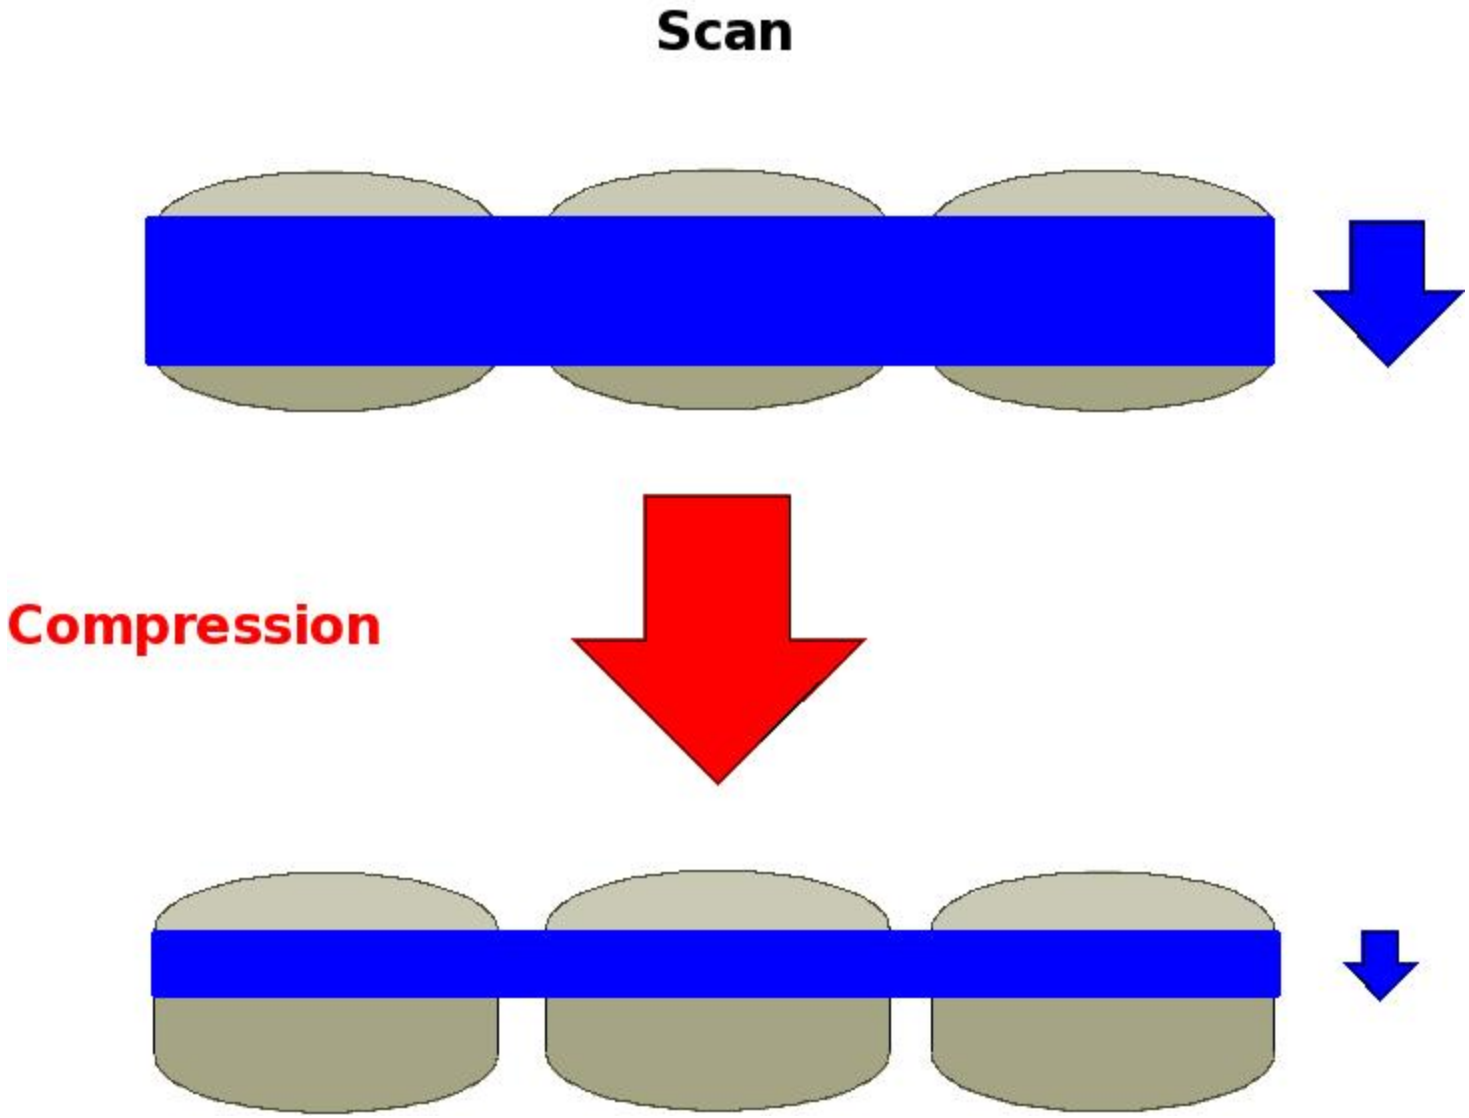

Backup / Restore

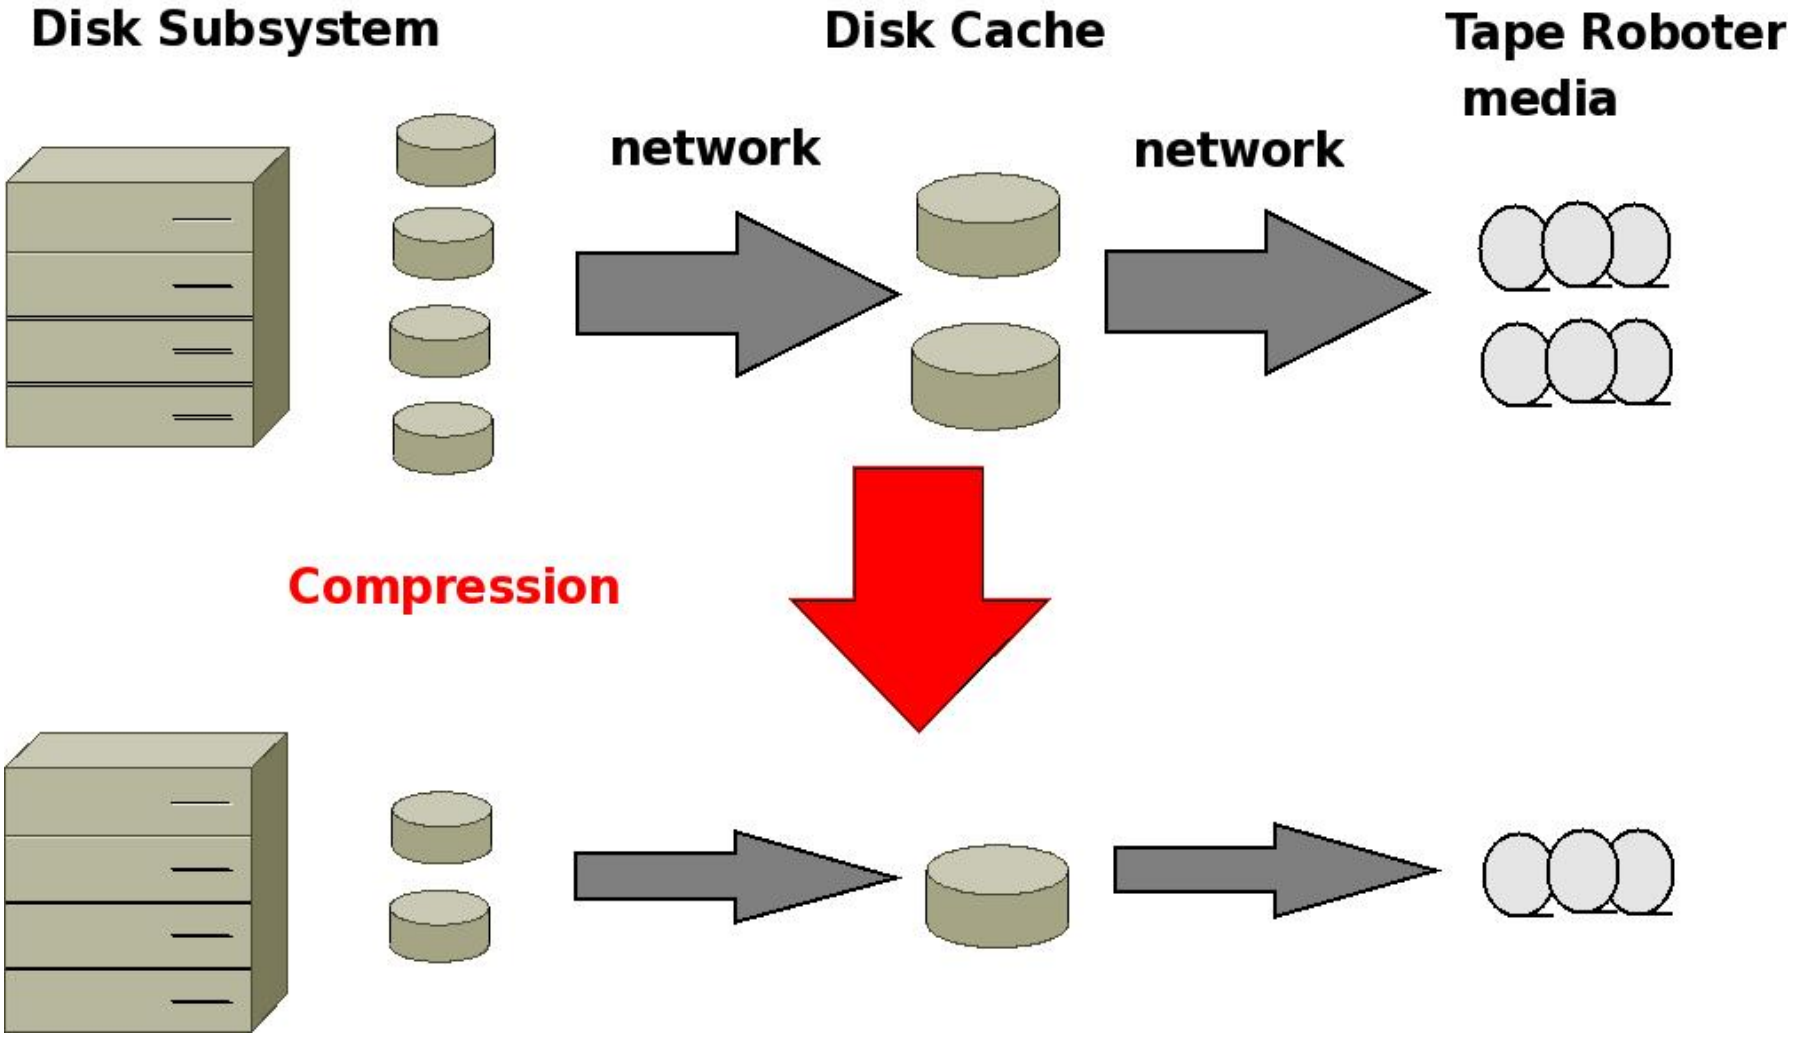

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

How does compression work?

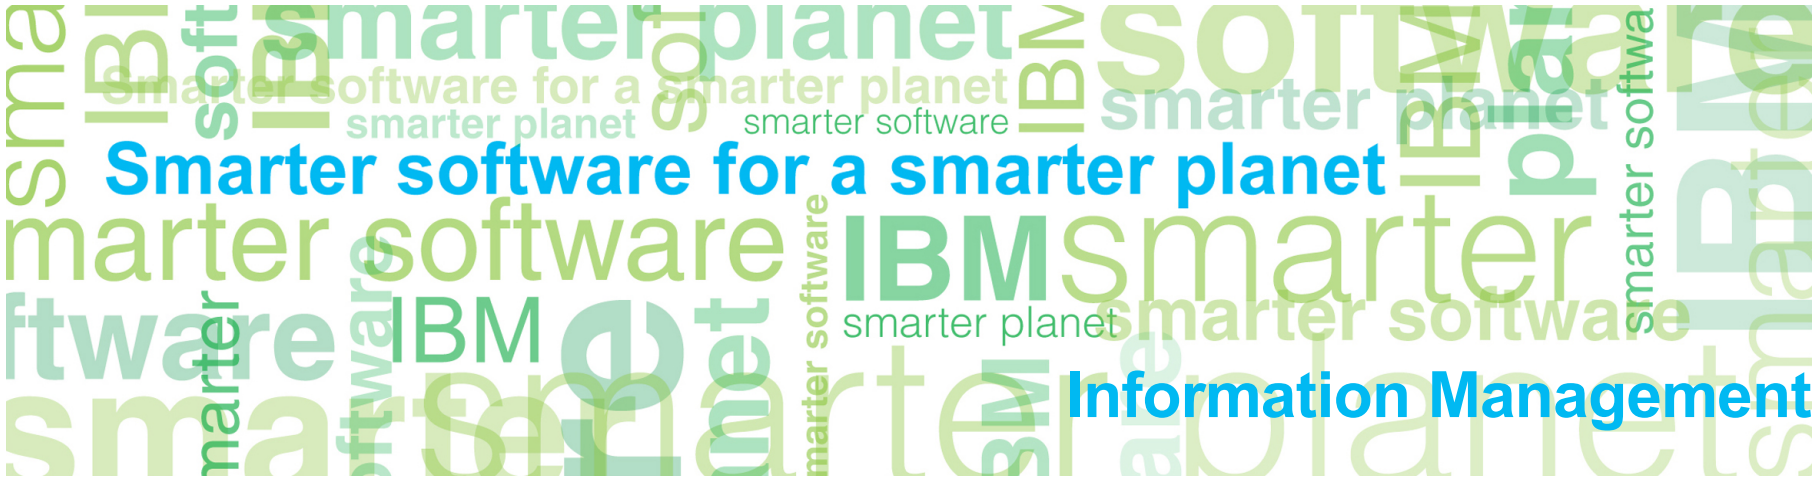

Compression Concepts

- Lempel-Ziv (LZ) based algorithm – static dictionary, built by random sampling.
- Frequently repeating patterns replaced with 12-bit symbol numbers.
- Dictionary tries to capture the “best” patterns (frequency x length).
- Any byte that does not match a pattern is also replaced with a 12-bit reserved symbol number.
- Patterns can be up to 15 bytes long.
- 12-bits means 4,096 symbols:
  - –256 reserved symbols for bytes that match no pattern.
  - –3,840 pattern symbols.

Compression and Storage Optimization

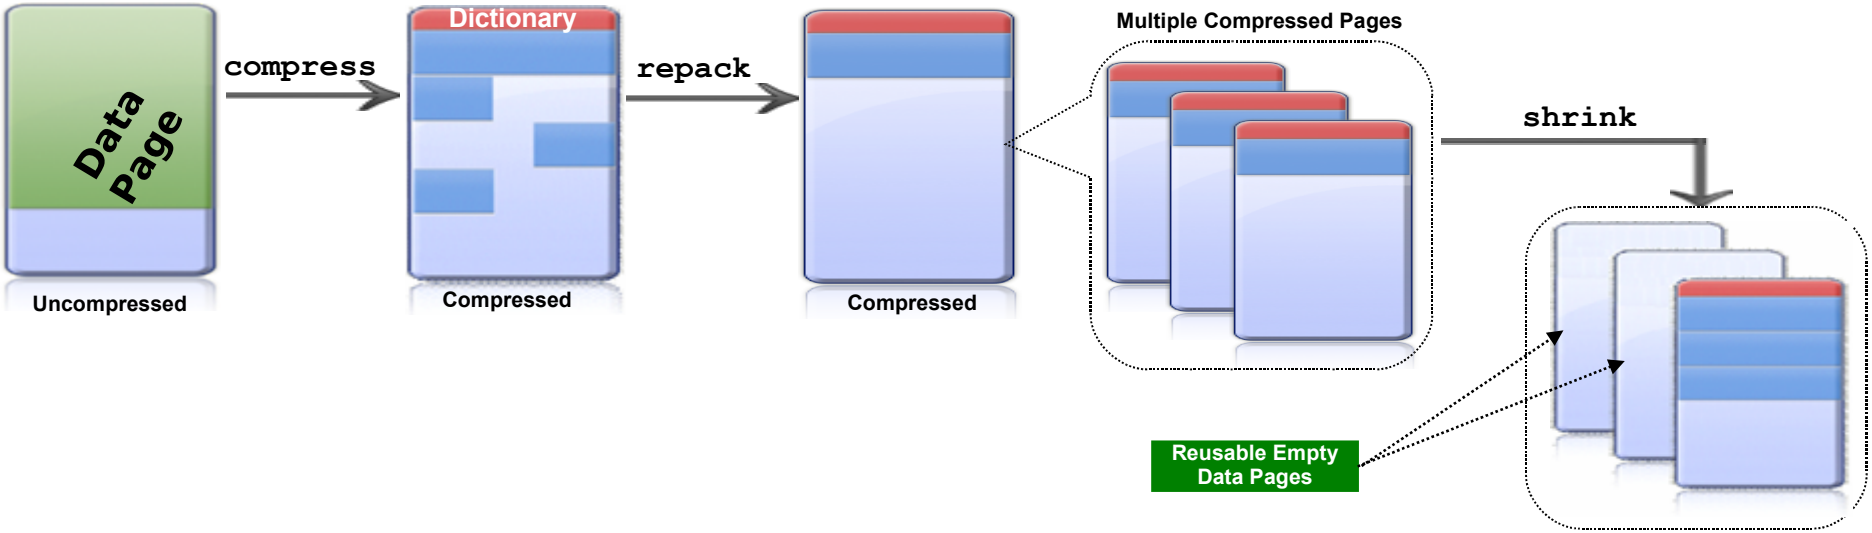

Shrink frees no longer used extents

For shortening first extent use

```
alter table ... modify extent size ...
```

Impact of Compression on System Maintenance

- As soon as a table is compressed all DML operations on that table will automatically produce compressed rows
- Compression is transparent to all accesses to the table
- Repack and shrink operations may be used for reorganizing (even non compressed) tables later if necessary
- Some maintenance operations (like backup if it is limited my network bandwidth) will be faster after compressions activated

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

Tooling for Compression

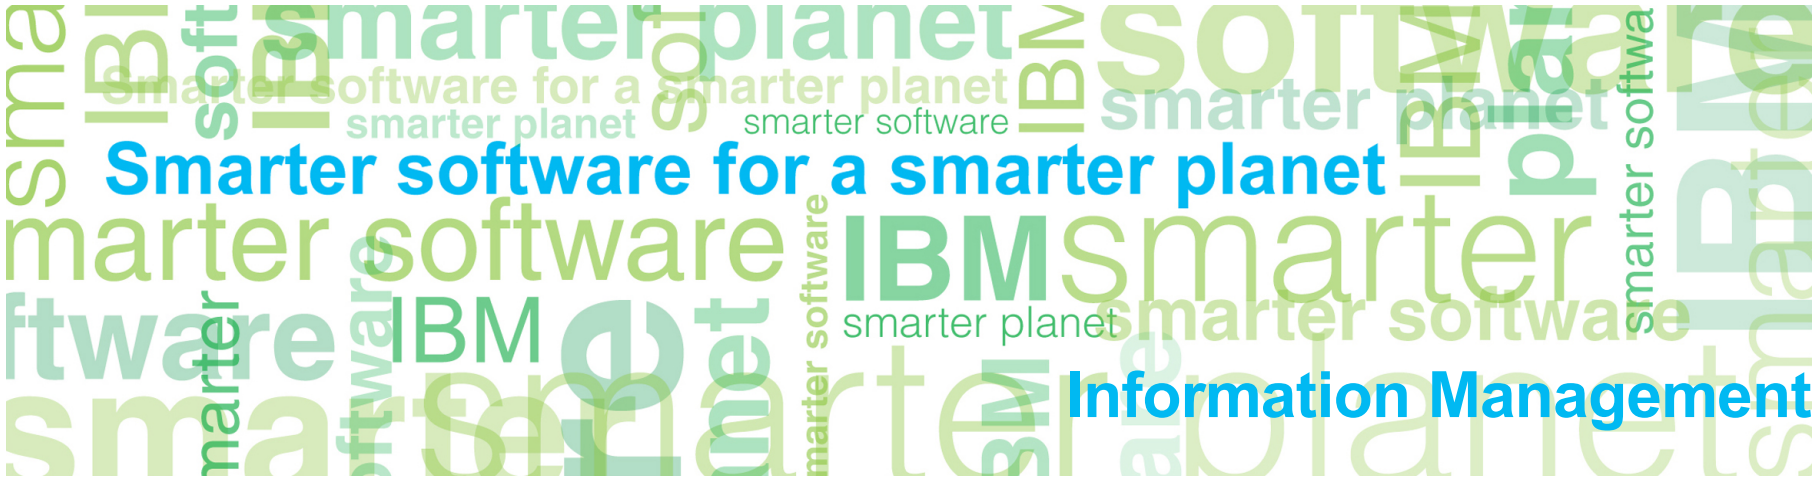

OAT Compression - General

Compression and Storage Optimization can be managed via the OAT graphical interface.

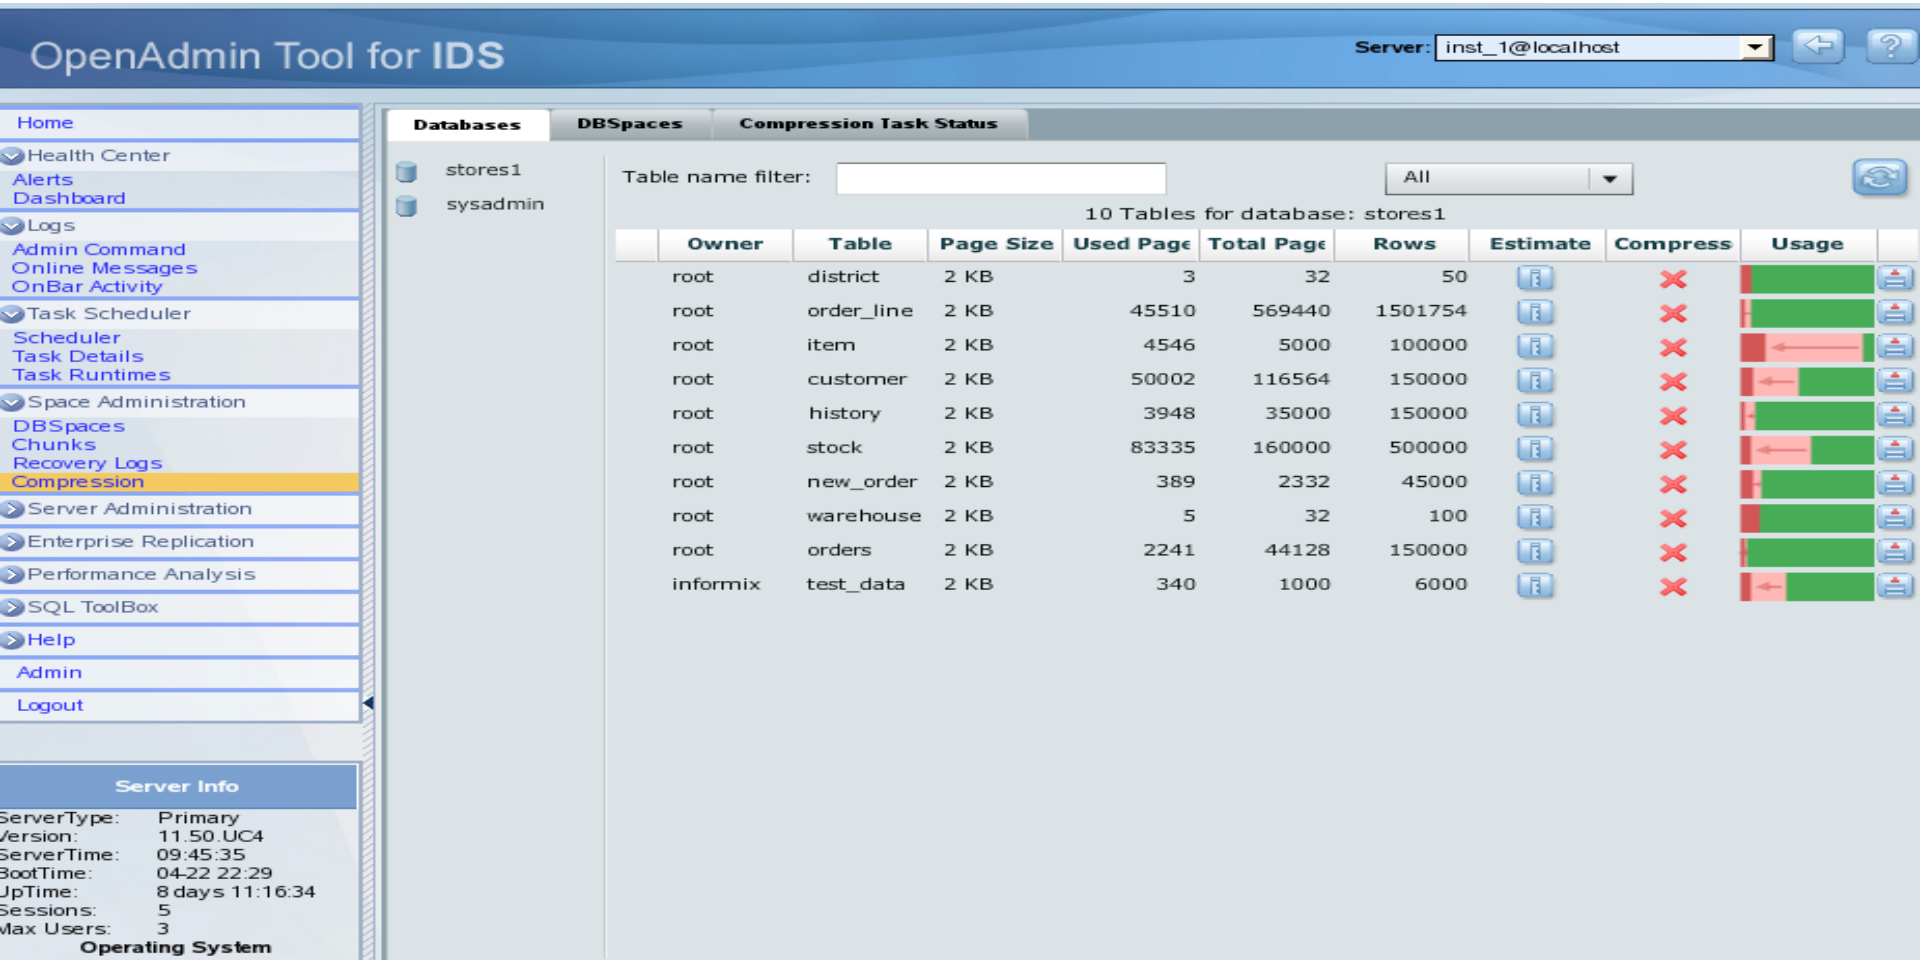

Compression Operations

- API Interface in Informix Dynamic Server:
  - –All compression and storage optimization operations are invoked via the IDS Admin API built-in UDRs:
    - execute function task(…);
    - execute function admin(…);
  - –Example:
    - execute function task(“table compress repack shrink”, “table_name”, “database_name”, “owner_name”);
  - –Enables remote execution (DBA does not need to log directly in to the target machine).
- Onutil Interface in XPS

Compression Estimation Tool

Schema based

http://www-01.ibm.com/software/sw-library/en_US/detail/L181272S36452U64.html

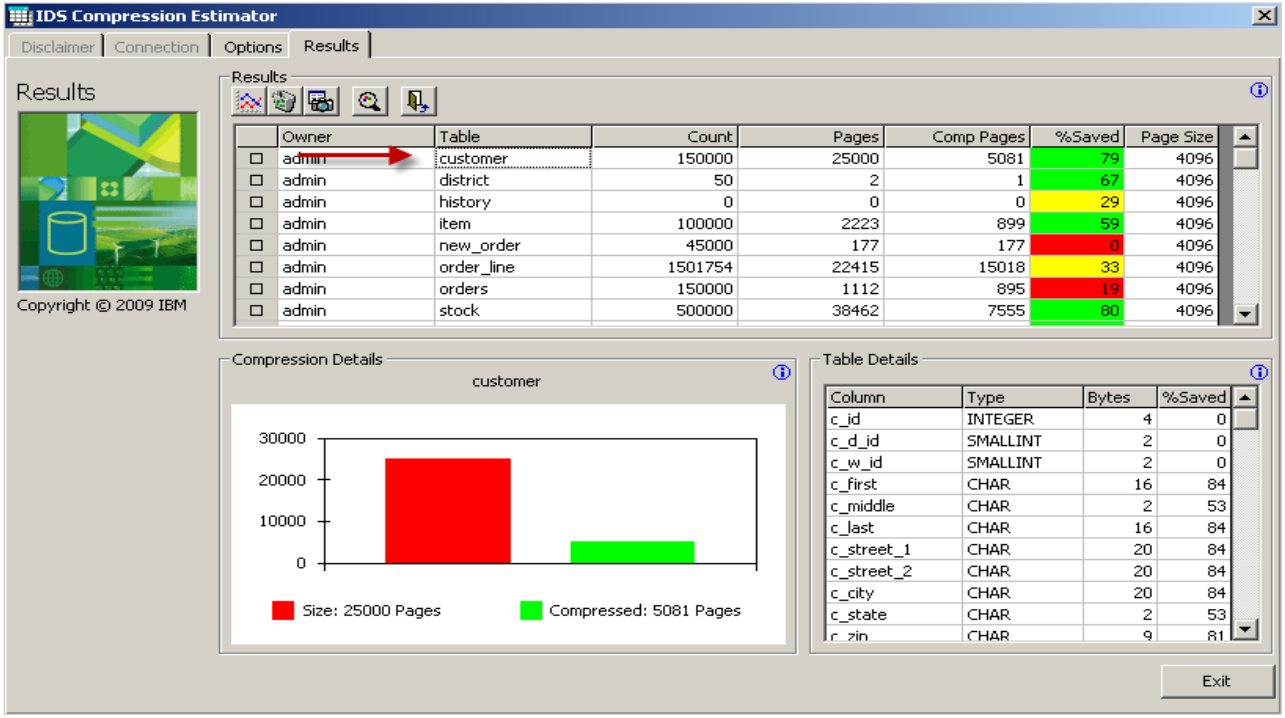

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

How much space is saved?

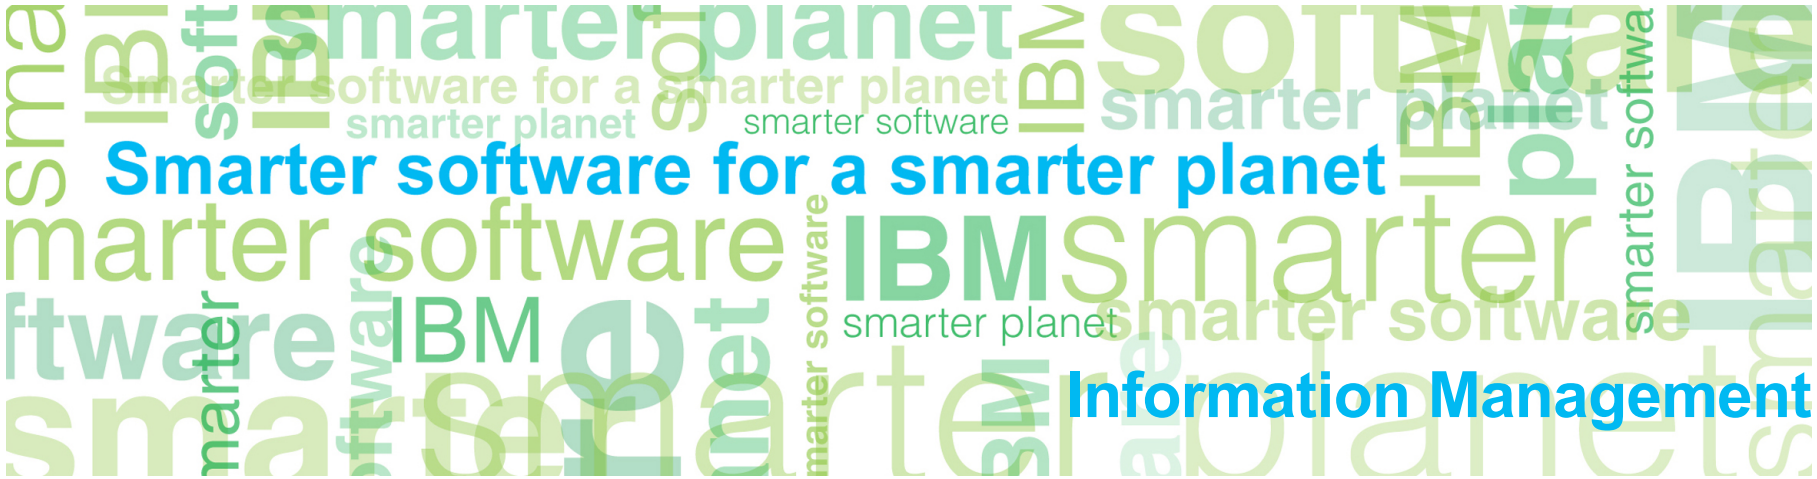

Customer Example: Space Used by Data Pages

Table	space before compression	space after compression	Savings in [%]
table 1	1,692,896	647,840	61.73
table 2	1,694,144	640,280	62.21
table 3	1,794,928	347,672	80.63
table 4	2,077,184	676,672	67.42
table 5	2,078,968	682,968	67.15
table 6	3,495,624	1,129,600	67.69
table 7	3,518,056	1,061,688	69.82
table 8	4,677,504	2,582,424	44.79
table 9	4,721,552	1,140,632	75.84
table 10	6,241,440	1,246,416	80.03
table 11	15,563,360	6,705,704	56.91
subtotal large tables:	47,555,656	16,861,896	64.54
total remaining tables	12,046,256	5,302,712	55.98
all tables:	59,601,912	22,164,608	62.81

space usage in [KB]

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

What is the performance impact?

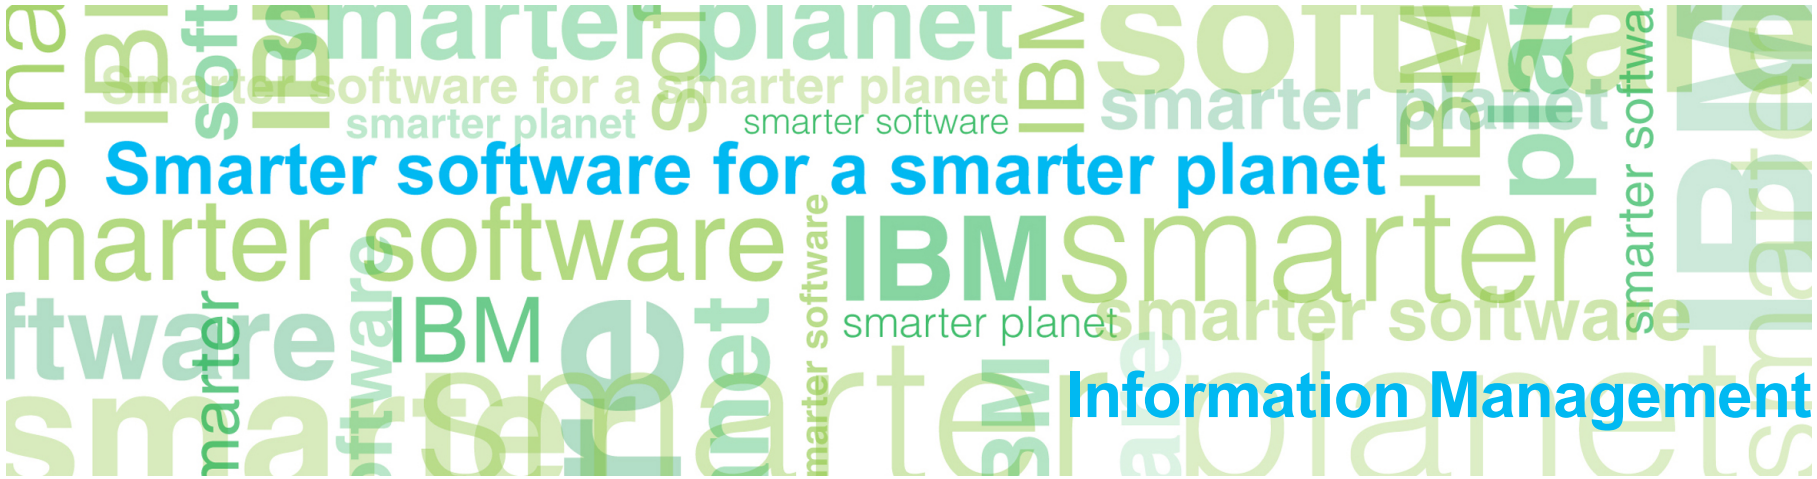

Performance Results in a Real Customer Benchmark
Batch Processing on OLTP System

Steps 1 to 3 mostly CPU bound

Job	Run time without compression	Run time with compression	Ratio: time with comp / time without comp	Time savings
Step 1	4564	4410	96.63 %	3.37 %
Step 2	8337	8461	101.49 %	-1.49 %
Step 3	5034	3748	74.45 %	25.55 %
Step 4	2537	1552	61.17 %	38.83 %
Step 5	1663	1266	76.13 %	23.87 %
Step 6	5279	4010	75.96 %	24.04 %
Total	27414	23447	85.53 %	14.47 %

All times in seconds

Performance Impact of Compression

- IO-bound workloads:
  - –Compression may improve performance by reducing IOs (both data page and logical log).
  - –More data fits on a page, so more in buffer pool.
  - –Log records are smaller. So there is less logging.
- For CPU-bound workloads:
  - –Additional CPU used to compress and expand rows.
  - –Should not be a large impact.
- Backups of compressed objects will take less time.

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

How is the impact of compression on other technologies and vice versa?

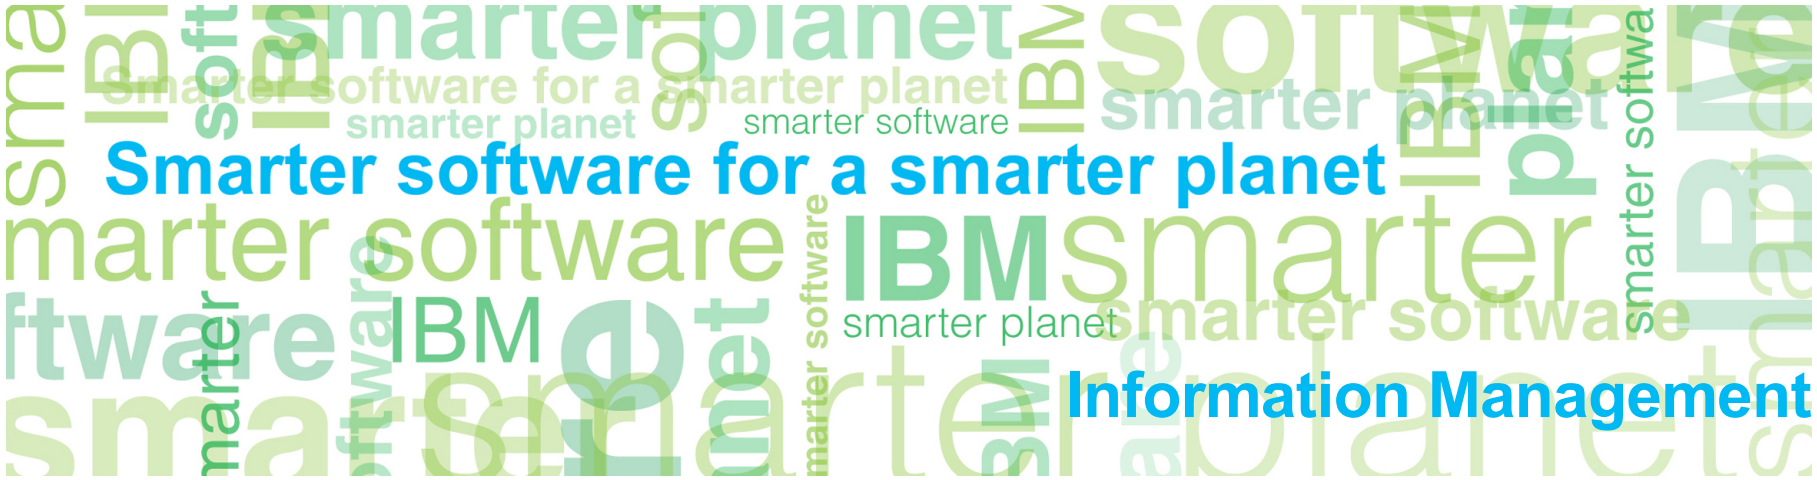

Compression and Solid State Disks (SSD)

- Characteristics of SSDs:
  - Very high number of IO/s:
    - e.g. Fusion-IO ioDrive Duo SSD
      > 100000 IO/s
      vs < 300 IO/s for traditional hard disk
  - High IO bandwidth
    - e.g. Fusion-IO ioDrive Duo SSD
    - > 1400 MB/s
    - vs < 200 MB/s for traditional hard disk
  - Capacity much more expensive

SSDs have
- Very good Price / IO/s ratio
- Bad Price / GB ratio

Compression and Solid State Disks (SSD) (cont)

- Compression can help to fit performance critical data on SSDs
- Additional storage tier with performance between bufferpool and regular dbspaces
- Example fragmentation by date

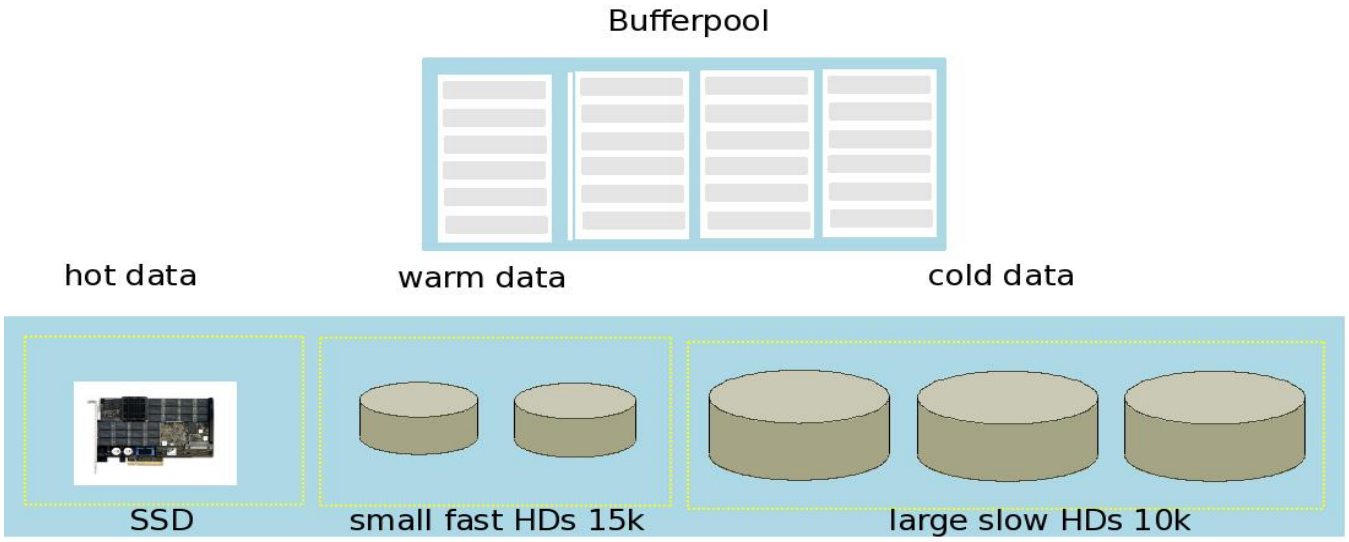

fragmentation by expression on date
compression makes sense across all storage tiers

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

Best practices

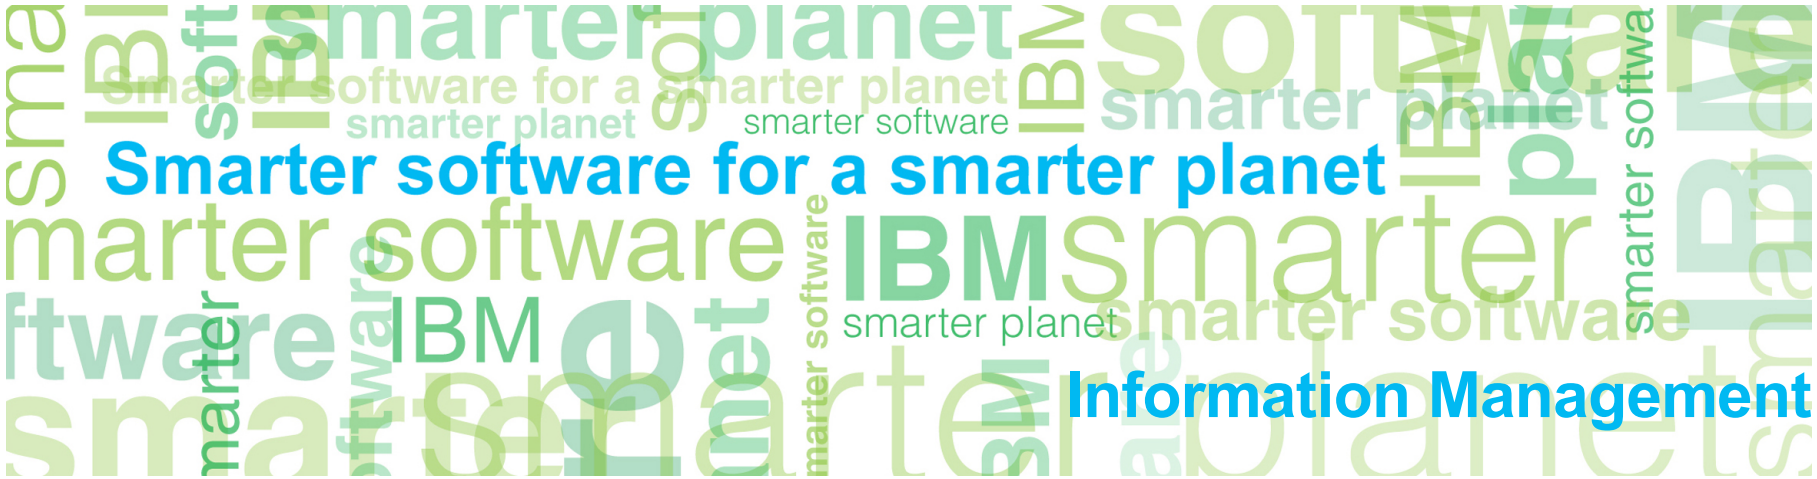

Which tables to compress?

- All tables with more than 2000 rows per fragment may be compressed
- Recommendation compression of all tables (ease of administration)
- Operation by fragment may make sense for online compress, repack, and uncompress to limit ressource usage

How to compress / uncompress?

- All compression operations are available online
- Offline version of operations usually not significantly faster
- Caveat: When tables are indexed it is recommended that table fragment and corresponding indexes fit into bufferpool for repack and uncompress operations for online operations (maintenance of index requires otherwise many random IOs

Compression and Rolling Windows

- Especially for data warehouses data are maintained with a “rolling window” method:
  - Data are fragmented by date column
  - Fragment for oldest time range is detached
  - Detached table is truncated
  - Detached fragment is loaded with newest data
  - Fragment is attached again
- Compression dictionary of detached fragment may be reused i.e. no compress, repack etc. operations necessary
- Experience shows that reused dictionary is usually good enough

What influences compression rates?

- Most important: Redundancy of data
- But: Data types influence also compression rate
- e.g. VARCHAR vs CHAR
  - without compression usually varchar more space efficient
  - with compression char has better compression ratio but is also more space efficient (!) after compression
- Other example: Decimal vs Integer

How to monitor / measure compression

- Testing performance of compression:
  - –Initial measurement of baseline for performance
  - –Compression/Repack/Shrink/Alter
  - –Test run on compressed data
- Tool for measuring disk space usage:
  - – oncheck -pT

oncheck -pT (Part 1)

```
TBLspace Report for dbs:usr.tab
...
                                                TBLspace use 4 bit bit-maps
                                                TBLspace is compressed
    Maximum row size                 1365
...
    Number of extents                1
...
    Pagesize (k)                     8
    First extent size                201500
    Next extent size                 80600
    Number of pages allocated        201500
    Number of pages used             4538
    Number of data pages             1372
    Number of rows                   36064
...

    Extents
         Logical Page     Physical Page        Size Physical Pages
                    0        14:8931772      201500     806000

    Compression Dictionary Identifiers
          rowid      loguniq       logpos
          5cf04        13927      65e8050
```

oncheck -pT (Part 2)

```
TBLspace Usage Report for dbs:usr.tab

    Type                  Pages      Empty  Semi-Full       Full  Very-Full
    --------------- ---------- ---------- ---------- ---------- ----------
    Free                200126
    Bit-Map                  2
    Index                    0
    Data (Home)           1372
    Data (Remainder)         0          0          0          0          0
                     ----------
    Total Pages         201500

    Unused Space Summary

        Unused data bytes in Home pages              1694281
        Unused data bytes in Remainder pages               0

    Home Data Page Version Summary

                Version                                Count

                     0 (current)                       1372

    Compressed Data Summary

        Number of  rows                          36064
        Number of compressed rows                36064
        Percentage of compressed rows             100.00
```

Andreas Weininger Andreas.Weininger@de.ibm.com
23 March 2011

Summary / Recommendations

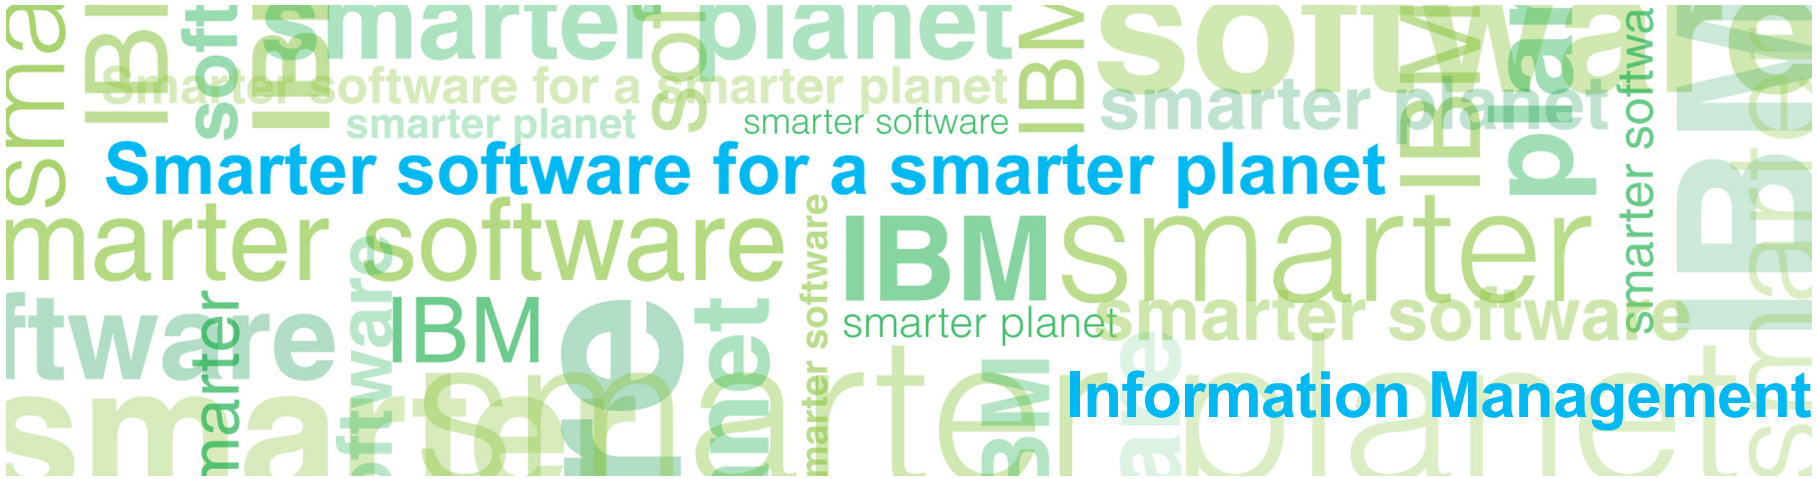

Summary Compression

- Compression and Storage Optimization can save disk space and thus €€€.
- For I/O-bound workloads, compression can also improve performance.
- Compression reduces logging.
- Compression fits more data into the buffer pool.
- Storage Optimization allows space saved by compression to be reclaimed from tables and fragments of tables.

Andreas Weininger Andreas.Weininger@de.ibm.com
9 November 2011

Questions?