

Ting tar IKKE tid. Erfaringer med store databaser

Ingemar Jansson Haverstad

OUGN Vårseminar 8. april 2011

ingemar@oraklet.no
www.oraklet.no/foredrag

Store databaser

Det vil si:

Enkelte oppgaver tar
IKKE
særlig lang tid.

**Men de fleste oppgaver vil ta
svært
lang tid!**

Referanser - Datavarehus

- **Posten** Sybase IQ
- **Statistisk sentralbyrå** Oracle 7.3 og 8
- **Telenor** Oracle 9.2
- **Nets (tidligere BBS)** Oracle 10.1
- **Skattedirektoratet** Oracle (11.1) 11.2

Avgrensning
=
Partisjonering

Erfaring - Datavarehus

- **Posten**

Sybase IQ

Bitmap indekser

Erfaring - Datavarehus

- **Statistisk sentralbyrå** Oracle 7.3 og 8

Partisjonering
Explain plan...
Antall disk

Erfaringer - Datavarehus

- **Telenor**

Oracle 9.2

RAC – HP Superdome
Migrering fra Sybase IQ

Erfaringer - Datavarehus

- **Nets (tidligere BBS)** **Oracle 10.1**

ASM
Statistikk

Erfaringer - Datavarehus

- **Skattedirektoratet** **Oracle (11.1) 11.2**

Forenkling
Statistikk, 11gR1
SQL Monitoring
Statistikk, 11gR2

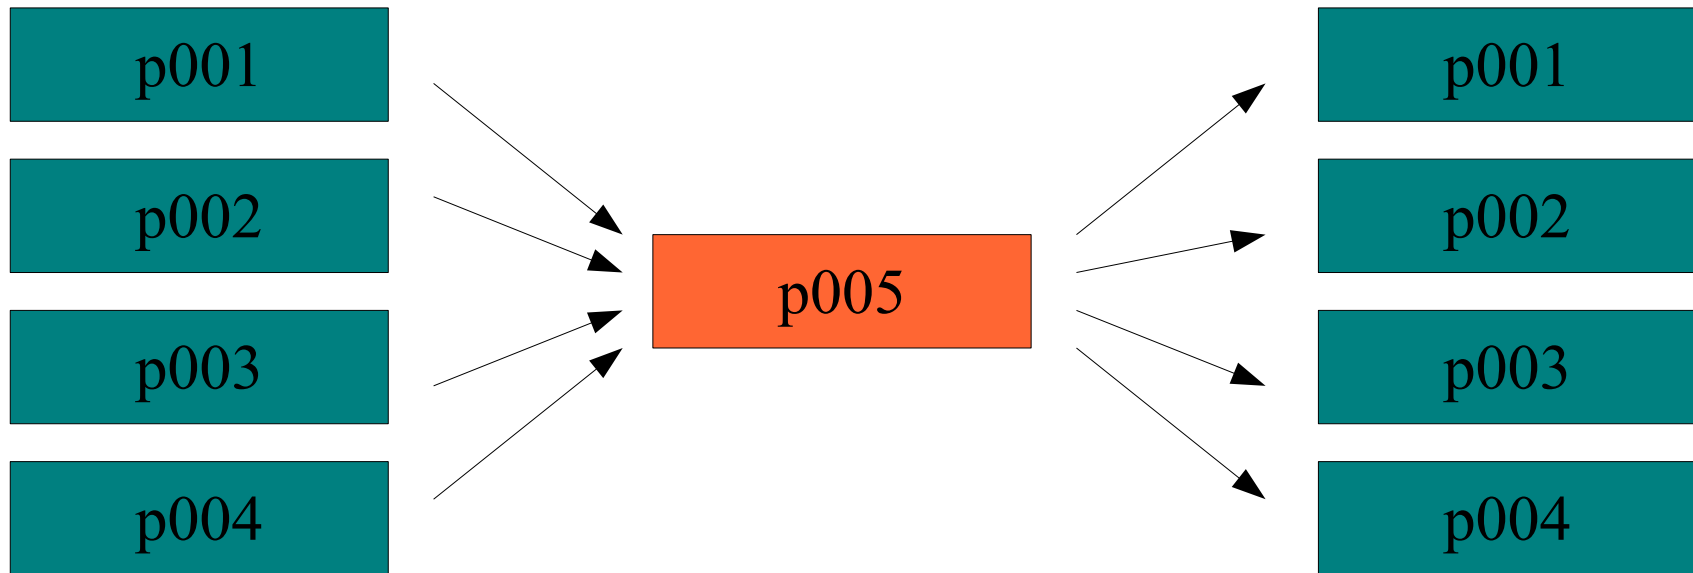
OPTIMIZER_FEATURES_ENABLE = 11.1.0.7

Mine erfaringer...

og

...mine forhåbninger!

Parallellisering Utfordring I




Parallell - Seriell

- Parallele operasjoner:


- `PARALLEL_TO_PARALLEL` P->P
- `PARALLEL_TO_SERIAL` P->S
- `PARALLEL_FROM_SERIAL`
- `PARALLEL_COMBINED_WITH_CHILD` PCWC
- `PARALLEL_COMBINED_WITH_PARENT` PCWP

0	SELECT STATEMENT	...		
1	PX COORDINATOR	...		
2	PX SEND QC (RANDOM)	Q1,00	P->S	QC (RAND)
3	PX BLOCK ITERATOR	Q1,00	PCWC	
4	INDEX FAST FULL SCAN	Q1,00	PCWP	



Parallellisering

Utfordring II



Parallelisering



PARALLEL_DEGREE_POLICY	MANUAL, AUTO eller LIMITED
PARALLEL_MIN_TIME_THRESHOLD	10 sekunder
PARALLEL_DEGREE_LIMIT	CPU, IO eller DEGREE (n)
PARALLEL_FORCE_LOCAL	Cluster parameter

Oracle Support Dokument 1264548.1

Parallellisering kan være vanskelig

```
MERGE /*+ aaPARALLEL ("RESKONTRO_FORFALL_HIST", 4) */  
INTO  
"RESKONTRO_FORFALL_HIST" "RESKONTRO_FORFALL_HIST"  
USING  
(SELECT /*+ PARALLEL ("W_MOT_AR_PAYMENT_SCHEDULES", 4)  
          PARALLEL ("RESKONTRO_FORFALL_HIST_KILDE", 4) */  
"RESKONTRO_FORFALL_HIST_KILDE"."RK_KRAV_ID" "KRAV_ID",  
"RESKONTRO_FORFALL_HIST_KILDE"."RK_KT_KRAV_KILDE" "KRAV_KILDE",  
"RESKONTRO_FORFALL_HIST_KILDE"."FORFALL_NR" "FORFALL_NR",  
  . . .
```

Kjetil Nordstrands ord:

Om man som her bruker parallell på samme tabell to ganger i samme spørring gir det:

```
«ORA-30926 unable to get a stable set of rows.»
```

Kjører man helt serielt tar det vinter og vår, men det går bra.

Fjerner man parallell hintet i «merge» delen går det hele bra & fort.

Merge - INSERT eller UPDATE

```
MERGE INTO sales s USING new_sales n
ON (s.sales_transaction_id = n.sales_transaction_id)
WHEN MATCHED THEN
UPDATE SET s.sales_quantity_sold =
           s.sales_quantity_sold + n.sales_quantity_sold,
           s.sales_dollar_amount =
           s.sales_dollar_amount + n.sales_dollar_amount
WHEN NOT MATCHED THEN INSERT
    (sales_transaction_id,
     sales_quantity_sold,
     sales_dollar_amount)
VALUES (n.sales_transaction_id,
        n.sales_quantity_sold,
        n.sales_dollar_amount);
```

- Hva er fordelingen mellom INSERT og UPDATE?

Jobben gikk fra mange timer
til noen minutter.

Hva skjer?

```
11gR2 sh@OUGN> ALTER SESSION ENABLE PARALLEL DML;

11gR2 sh@OUGN> INSERT /*+ APPEND PARALLEL(10, trg) */
                   INTO    sh.sales trg
                   SELECT /*+ PARALLEL(10, src) */ *
                   FROM    sh.sales_temp src
                   WHERE   aar = 2007;
```

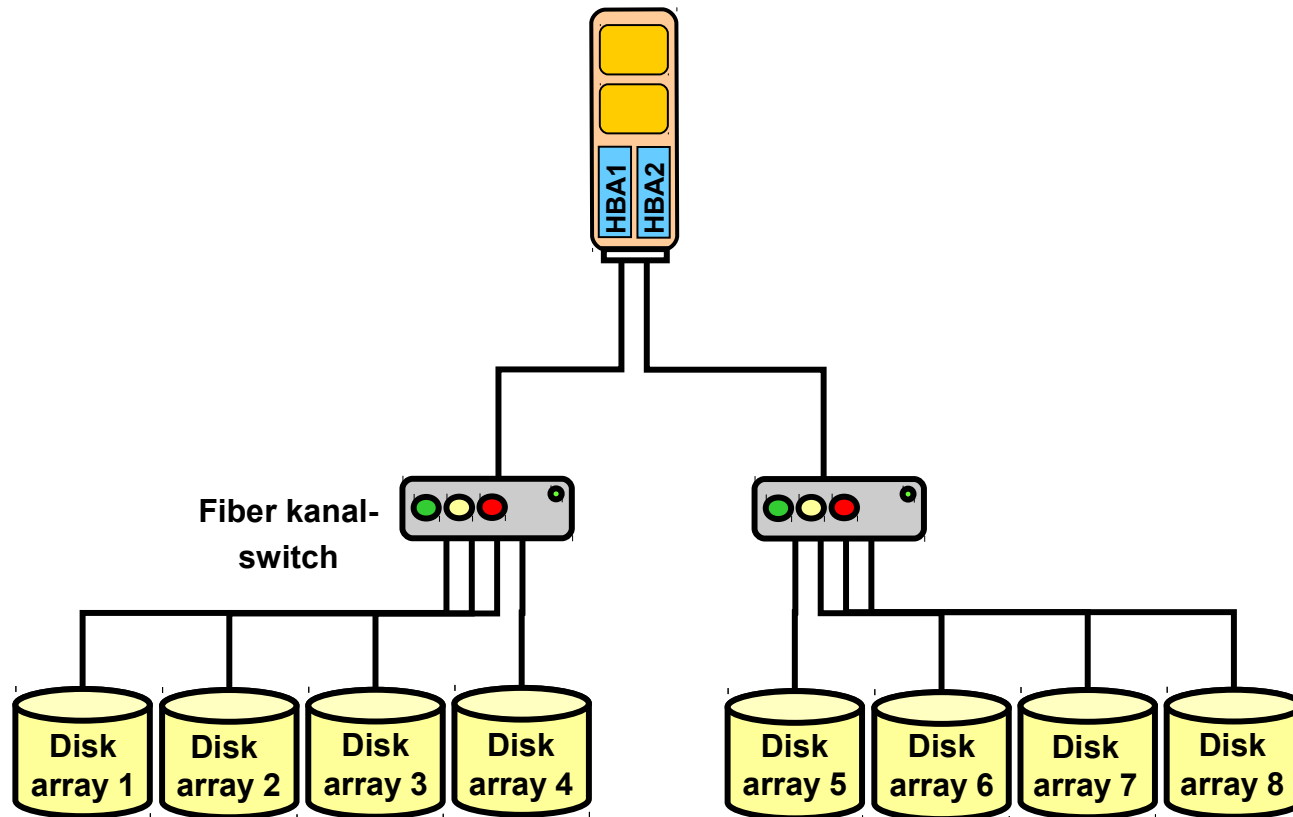


- Oppdatere data segmenter
- Oppdatere indeks segmenter
- Skrive til og lese fra undo segmenter
- Skrive til og lese fra temporære segmenter
- Skrive til og lese fra redo log
- Skrive til arkiv filer
- Auditing, FGAC og VPD

Server konfigurasjon

- **Faktorer som påvirker**
 - **CPU:** antall, type og frekvens
 - **Minne:** mengde og type
 - **I/O subsystem:** type, antall, oppsett, redundans og størrelse
 - **Nettverk:** type og redundans
- **Båndbredde (volum) og hastighet (latency)**
- **Virtualisering**

Et balansert system...



Det vil si,
systemet var balansert et øyeblikk.
Full harmoni rådde.

Virtualisering

På godt og vondt...

```

Topas CEC Monitor                      Interval:  10                      Sat Mar 19 16:38:47
2011
Partitions Memory (GB)                 Processors
Shr: 13      Mon: 140  InUse: 121  Shr:3.3  PSz: 12   Don: 0.0  Shr_PhysB  0.61
Ded:  0      Avl:    -           Ded:  0  APP: 11.4 Stl: 0.0  Ded_PhysB  0.00

```

Host	OS	M	Mem	InU	Lp	Us	Sy	Wa	Id	PhysB	Vcsw	Ent	%EntC	PhI
-----shared-----														
dbsrv1	A53	U	16	16	4	57	22	0	20	0.34	0	0.20	170.6	0
dbsrv4	A53	U	8.0	7.8	4	30	10	0	59	0.10	0	0.20	51.8	0
appsrv3	A53	U	8.0	3.8	4	6	10	0	83	0.04	525	0.20	21.5	1
crmsrv3	A53	U	8.0	7.8	4	5	4	0	90	0.03	0	0.20	13.7	0
crmsrv4	A53	U	4.0	3.9	4	4	2	0	93	0.02	0	0.20	9.6	0
dwh	A53	U	32	20	8	1	1	0	97	0.02	0	0.40	4.5	0
dbsrv2	A53	U	8.0	7.8	4	1	2	0	95	0.01	0	0.20	6.5	0
dwh_4	A53	U	16	16	8	0	0	0	98	0.01	0	0.40	2.4	0
crmsrv1	A53	U	4.0	3.9	4	0	1	0	97	0.01	271	0.20	4.2	0
crmsrv2	A53	U	4.0	3.8	4	0	1	0	97	0.01	372	0.20	4.0	0
appsrv	A53	U	16	16	4	0	1	0	98	0.01	322	0.20	3.1	0
dbsrv3	A53	U	8.0	7.9	4	0	0	0	99	0.01	162	0.50	1.0	0
dwh_1	A53	U	8.0	6.5	4	0	1	0	98	0.00	0	0.20	2.3	0

Virtualisering, mest vondt.

```

Topas CEC Monitor                      Interval: 10                      Sat Mar 19 16:38:47
2011
Partitions Memory (GB)                  Processors
Shr: 13      Mon: 140  InUse: 121  Shr:3.3  PSz: 12  Don: 0.0  Shr_PhysB  0.61
Ded:  0      Avl:    -              Ded:  0  APP: 11.4  Stl: 0.0  Ded_PhysB  0.00

Host      OS  M Mem InU Lp  Us Sy Wa Id  PhysB  Vcsw Ent  %EntC PhI
-----
-----shared-----
dbsrv1    A53 U   16  16  4  57 22  0 20  0.34   0  0.20 170.6  0
dbsrv4    A53 U   8.0 7.8  4  30 10  0 59  0.10   0  0.20  51.8  0
dwh_1     A53 U   8.0 6.5  4   0  1  0 98  0.00   0  0.20   2.3  0

```

An Oracle White Paper
November 2010

Oracle Real Application Clusters One Node:
Better Virtualization for Databases

«Virtual servers always suffer a performance handicap vs. physical hardware. This is because CPU instructions must be translated in a virtual server, adding to CPU overhead.»

Veldig vondt.

Den store diskgrøten...

- Hvor er de ulike segmentene lagret?




Server1



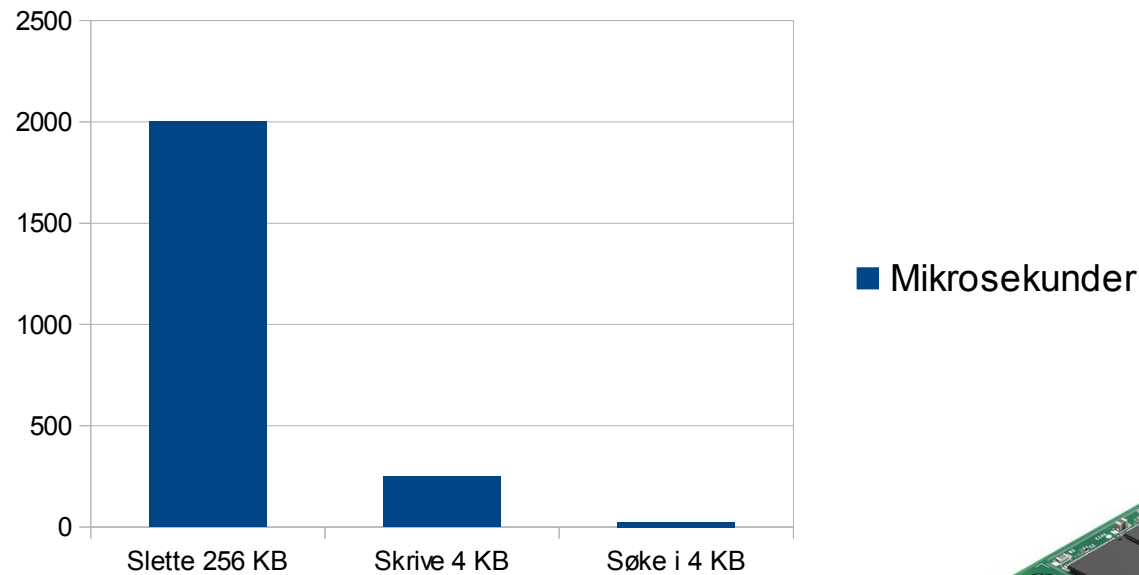
Server2



Server3



Solid State Disk - SSD



- Skriving og lesing 4 GB sider
- Ved sletting brukes 256 KB blokker
- 39 500 IO/s ved vilkårlig lesing
- 23 000 IO/s ved vilkårlig skriving
- 270 MB/s ved sekvensielle les
- 220 MB/s ved sekvensielle skriv



Solid State Disk - SSD

Alternativer:

- Plasser hele databasen på SSD.
- Plasser utvalgte segmenter på SSD.
 - `V$SEGMENT_STATISTICS`
- Benytt 11gR2 Smart Cache.
- Plasser temporær segmenter på SSD.
- Plasser Redo Log filene på SSD.

DB Smart Cache

- «DB Smart Cache» er en utvidelse av «Buffer Cache»
- Brukes i hovedsak for Solid State disk
- Skriv I/O må tilsvare les I/O
- Hovedsaklig for lese intensive OLTP systemer
- 2-10 ganger størrelsen av «Buffer Cache»

```
11gR2 SYS@OLTP SQL> SHOW PARAMETERS flash_cache
```

NAME	TYPE	VALUE
db_flash_cache_file	string	+FLASH_DG
db_flash_cache_size	big integer	5G

```
11gR2 SYS@OLTP SQL> CREATE TABLE postnummer(...)
                        STORAGE (FLASH_CACHE);
```

Heldigvis slipper vi å tenke på dette.

Consider adding the flash cache when all of the following are true:

- Your database is running on **Solaris** or **Oracle Linux**.
- ...
- ...
- ...

Uansett snakker vi om en stor svart boks!

Men hva har vi av muligheter?

For å kunne ha noe formening om noe som
helst...

Først av alt, få deg en Solid State Disk.

Og begynn å teste.

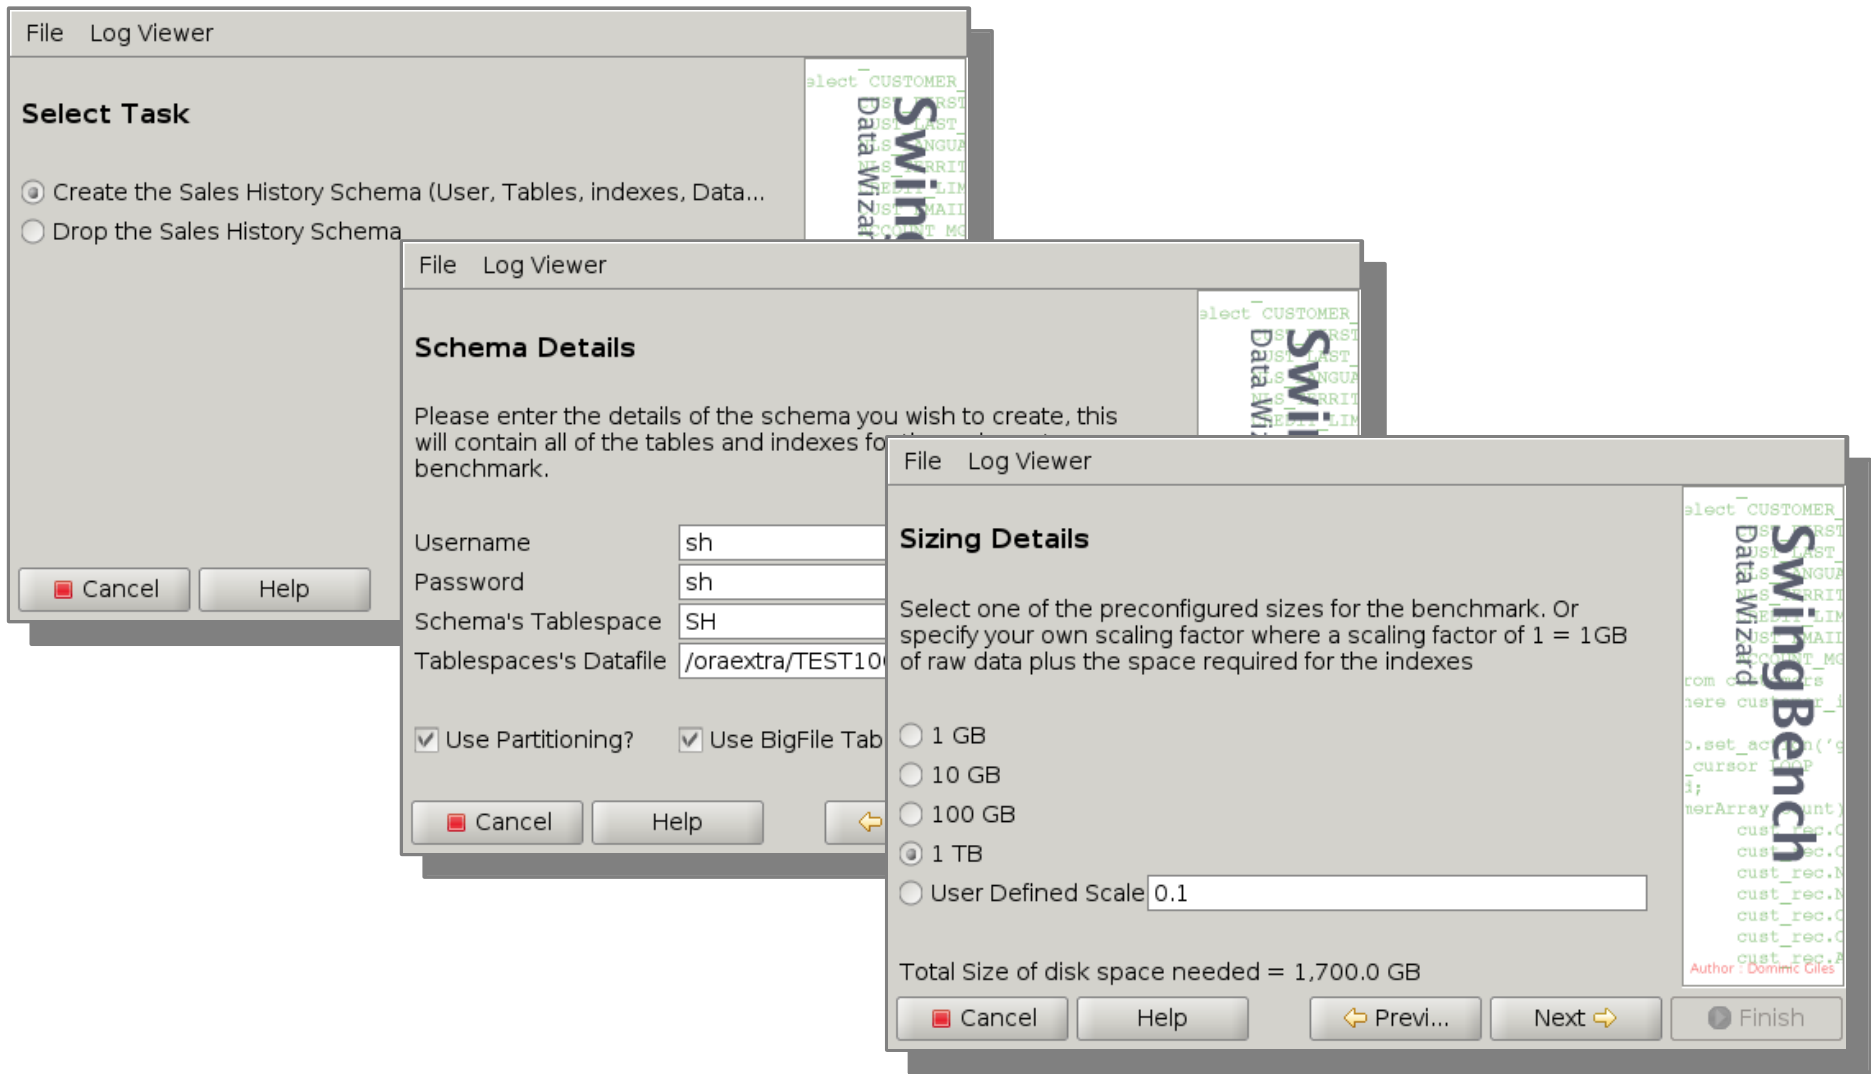
Swingbench

- Skrevet i Java av Dominic Giles, Oracle UK
- <http://ww.dominicgiles.com>
- Støtte for OLTP – Order Entry (SOE)
- Støtte for datavarehus – Sales History (SH)
- Telco: Calling Circle (CC)

```
Function getCustomerDetails(cust_id customers.customer_id%ty
CustomerArray customer_array_type := customer_array_type
cursor cust_cursor is select CUSTOMER_ID,
                             CUST_FIRST_NAME,
                             CUST_LAST_NAME,
                             NLS_LANGUAGE,
                             NLS_TERRITORY,
                             CREDIT_LIMIT,
                             CUST_EMAIL,
                             ACCOUNT_MGR_ID
                             from customers
                             where customer_id = cust_id;
begin
dbms_application_pool.begin_transaction('getCustomerDetails');
FOR cust_rec IN cust_cursor LOOP
CustomerArray.extend;
CustomerArray(CustomerArray.count) := customer_rec_t
                             cust_rec.CUST_FIRST_NAME,
                             cust_rec.CUST_LAST_NAME,
                             cust_rec.NLS_LANGUAGE,
                             cust_rec.NLS_TERRITORY,
                             cust_rec.CREDIT_LIMIT,
                             cust_rec.CUST_EMAIL,
                             cust_rec.ACCOUNT_MGR_ID;
end;
```

SwingBench
An Oracle Load Generation Tool
Version 2.4
Author : Dominic Giles, Oracle UK

Datagenerator



Testprogram

File Help

Configuration \ Preferences \ Output \ Events \

User Details \ Connection Pooling \ Properties \

Username: sh

Password: ●●

Connect String: 11z:1523:TEST11G

Driver Type: Oracle10g Type IV jdbc driver (thin)

Collect database statistics

Take AWR snapshots at start and end (10g/11g only)

System Username: _____

/ Load \ Environment Variables \ Distributed Controls \

Number of Users: 5

Min. Delay Between Transactions (ms): 250

Max. Delay Between Transactions (ms): 750

Logon Delay (milliseconds): 0

Logon Group: 1

Wait Till All Sessions Log On: true

Logoff Post Transaction: false

Benchmark Run Time (hh:min): 0 1

Record Statistics After (hh:min): 0 0

Stop Recording After (hh:min): 0 0

Transactions \ Jobs \

Id	Class Name	Short Name	Load Ratio	Activate ?
Sales Rollup by Month and ...	com.dom.benchmarking.swingbench.dsstransactions.S...	SRMC	100	<input checked="" type="checkbox"/>
Sales Cube by Month and C...	com.dom.benchmarking.swingbench.dsstransactions.S...	SCMC	100	<input checked="" type="checkbox"/>
Product Sales Cube and Ro...	com.dom.benchmarking.swingbench.dsstransactions.Pr...	PSCR	100	<input checked="" type="checkbox"/>
Sales Moving Average	com.dom.benchmarking.swingbench.dsstransactions.S...	SMA	100	<input checked="" type="checkbox"/>
Period to Period Sales Com...	com.dom.benchmarking.swingbench.dsstransactions.P...	PPSC	100	<input checked="" type="checkbox"/>
Top Sales by Quarter	com.dom.benchmarking.swingbench.dsstransactions.To...	TSQ	100	<input checked="" type="checkbox"/>
Sales within Quarter by Cou...	com.dom.benchmarking.swingbench.dsstransactions.S...	SQC	100	<input checked="" type="checkbox"/>

Chart Type: Transactions Per Minute Users: _____

Transactions Per Minute

Sales History

Transactions Per Minute : 29

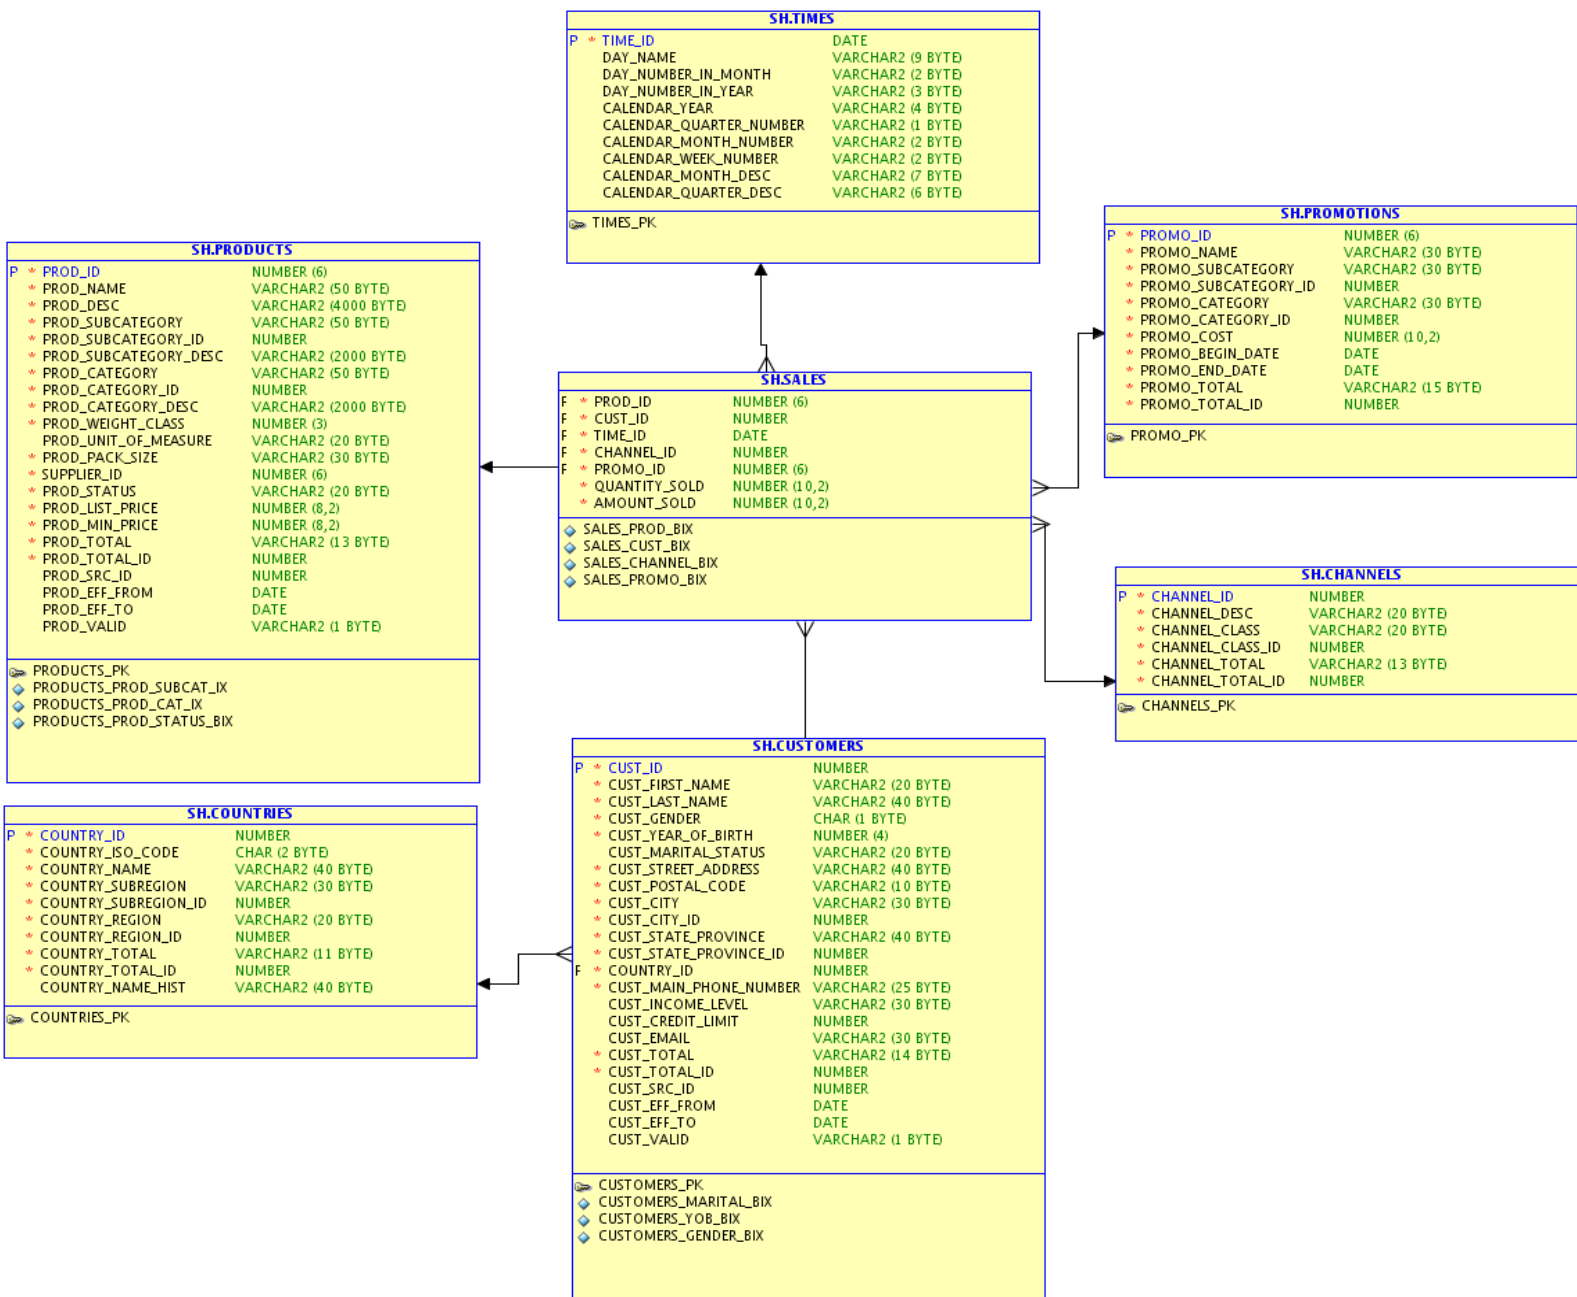
Users Logged On : 0

Generere data: *Sales*

	10 GB	100 GB
Opprette tabellplass	00:02:27	00:18:44
Generere data	00:15:56	02:47:39
Opprette indekser	00:18:15	02:15:36
Analysere skjema	00:09:16	00:22:47

- 4 CPU dual core
- Mid range SAN (400MB/sec I/O)
- Data 10 GB, indeks 5 GB

Sales model



ETL

- Ekstrahere – hente data fra kilde
- Transformer – konvertere data
- Laste inn – fylle opp tabeller

En ting som ikke kommer med...
... tid det tar å beregne statistikk.

ETL

ORACLE®
DATABASE

10^g

Schema Created

Completed schema successfully

Statistic	Value
Connection Time	0:00:03.185
Data Generation Time	0:41:48.786
DDL Creation Time	1:50:23.294
Total Run Time	2:32:15.281
Rows Inserted per sec	55 805
Data Generated (MB) per sec	4,2

The creation of the schema appears to have been successful.

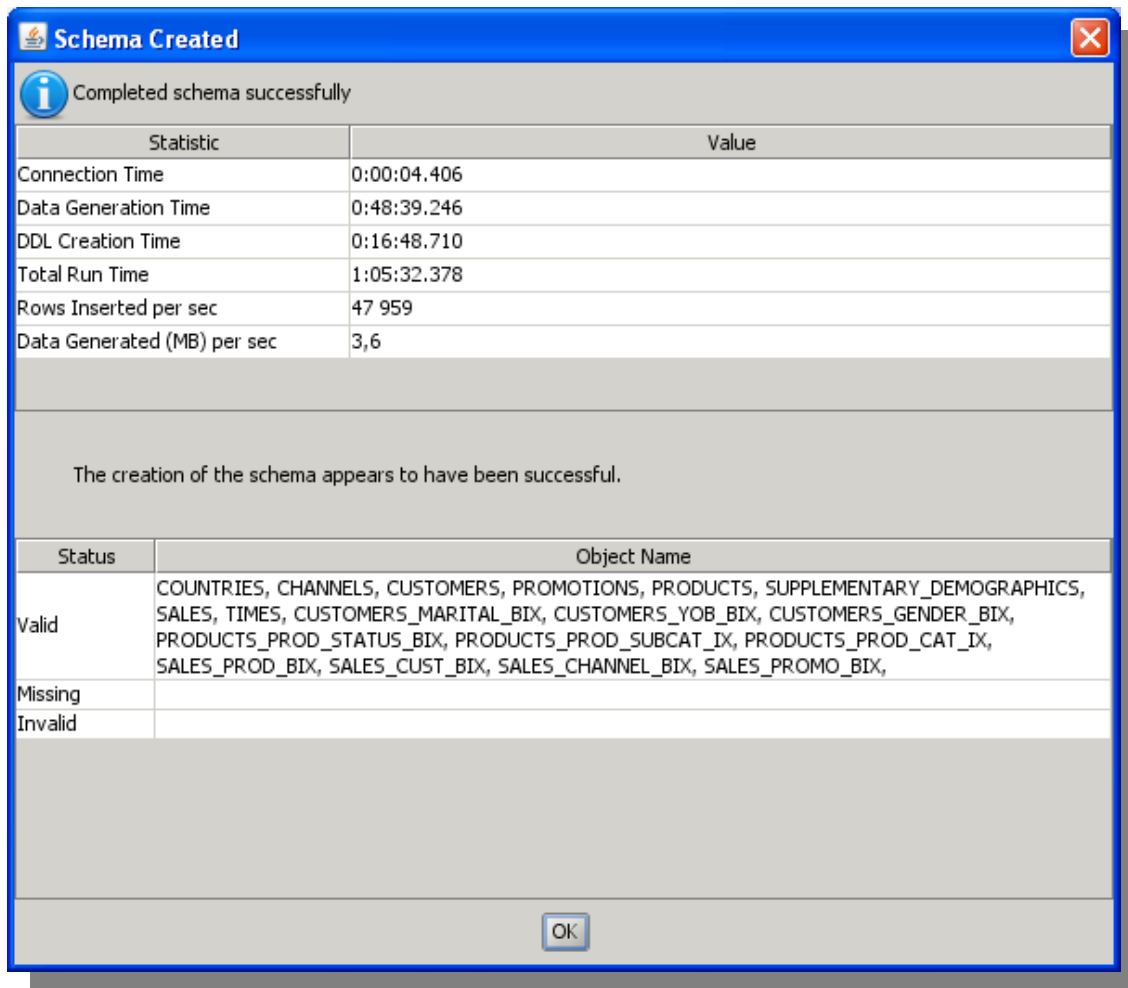
Status	Object Name
Valid	COUNTRIES, CHANNELS, CUSTOMERS, PROMOTIONS, PRODUCTS, SUPPLEMENTARY_DEMOGRAPHICS, SALES, TIMES, CUSTOMERS_MARITAL_BIX, CUSTOMERS_YOB_BIX, CUSTOMERS_GENDER_BIX, PRODUCTS_PROD_STATUS_BIX, PRODUCTS_PROD_SUBCAT_IX, PRODUCTS_PROD_CAT_IX, SALES_PROD_BIX, SALES_CUST_BIX, SALES_CHANNEL_BIX, SALES_PROMO_BIX,
Missing	
Invalid	

OK

Data Generation time 0:42
 DDL Creation time 1:50
 (statistikk 1:40)
Total Run Time 2:32

ETL

ORACLE®
 DATABASE 11g



Data Generation time 0:48
 DDL Creation time 0:16
 (statistikk 0:10)
Total Run Time 1:05

Statistikk

Oracle 10g	1:40
Oracle 11g	0:10

In 11g, using auto size for ESTIMATE_PERCENT defaults to 100% and therefore is as accurate as possible. Additionally, even though a 100% sample is collected, the gathering process is really fast **since a new hashing algorithm is used to compute the statistics rather than sorting** (in 9i and 10g the "slow" part was typically the sorting).

In 10g the default ESTIMATE_PERCENT sample size was extremely small which often resulted in poor statistics and is therefore not recommended.

Oracle Support Dokument 605439.1

Oppgradering til

ORACLE[®]

DATABASE **11gR1**

Vil gi deg mye bedre statistikk, hurtigere.

ORACLE[®]

DATABASE **11gR2**

Statistikk for SUBPARTISJONER 11gR2

Publisere statistikk

- Du kan nå velge når statistikk skal publiseres

```
11gR2 SYS@OLTP SQL> SELECT DBMS_STATS.GET_PREFS('PUBLISH')
                        FROM DUAL;
```

```
PUBLISH
```

```
-----
```

```
TRUE
```

- Kan unngå å gjøre statistikk tilgjengelig

```
11gR2 SYS@OLTP SQL> BEGIN
                        dbms_stats.set_table_prefs(
                        'SH', 'CUSTOMERS', 'PUBLISH', 'false');
                        END;
```

Benytte statistikk

- Teste ny statistikk

```
11gR2 SH@OUGN SQL> ALTER SESSION SET  
optimizer_use_pending_statistics = TRUE;  
...
```

- Gjøre statistikken tilgjengelig


```
BEGIN  
  dbms_stats.publish_pending_stats('SH', 'CUSTOMERS');  
END;
```

Nologging- The Gains and Pains

```
ORA-01578: ORACLE data block corrupted (file# 3, block#514)  
ORA-01110: data file 3: '/orassd/OUGN/OUGN_users_01.dbf'  
ORA-26040: Data block was loaded using the NOLOGGING option
```

Oracle Support Dokument 290161.1

Nologging

 Completed schema successfully

Statistic	Value
Connection Time	0:00:03.500
Data Generation Time	0:17:09.750
DDL Creation Time	0:08:58.098
Total Run Time	0:26:11.354
Rows Inserted per sec	135,959
Data Generated (MB) per ...	10.2

The creation of the schema appears to have been successful.


Status	Object Name
Valid	COUNTRIES, CHANNELS, CUSTOMERS, PROMOTIONS, PRODUCTS, SUPPLEMENTARY_DEMOGRAPHICS, SALES, TIMES, CUSTOMERS_MARITAL_BIX, CUSTOMERS_YOB_BIX, CUSTOMERS_GENDER_BIX, PRODUCTS_PROD_STATUS_BIX, PRODUCTS_PROD_SUBCAT_IX, PRODUCTS_PROD_CAT_IX, SALES_PROD_BIX, SALES_CUST_BIX, SALES_CHANNEL_BIX, SALES_PROMO_BIX,

OK

```
Data Generation time      0:17
DDL Creation time        0:09
Total Run Time          0:26

Rader generert/sekund 135.959
```

Logging

 Completed schema successfully

Statistic	Value
Connection Time	0:00:06.294
Data Generation Time	0:17:17.882
DDL Creation Time	0:11:02.814
Total Run Time	0:28:26.998
Rows Inserted per sec	134,894
Data Generated (MB) per ...	10.1

The creation of the schema appears to have been successful.

Status	Object Name
Valid	COUNTRIES, CHANNELS, CUSTOMERS, PROMOTIONS, PRODUCTS, SUPPLEMENTARY_DEMOGRAPHICS, SALES, TIMES, CUSTOMERS_MARITAL_BIX, CUSTOMERS_YOB_BIX, CUSTOMERS_GENDER_BIX, PRODUCTS_PROD_STATUS_BIX, PRODUCTS_PROD_SUBCAT_IX, PRODUCTS_PROD_CAT_IX, SALES_PROD_BIX, SALES_CUST_BIX, SALES_CHANNEL_BIX, SALES_PROMO_BIX,

OK

```
Data Generation time      0:17
DDL Creation time        0:11
Total Run Time          0:28

Rader generert/sekund 134.944
```

Resultat

```
Nologging      135.959
Logging        134.944
```

Nå kjører vi i produksjon med **LOGGING**.


The pain - Logging

```
SYS@OUGN SQL> ALTER DATABASE FORCE LOGGING;  
Database altered.
```

```
SYS@OUGN SQL> ALTER TABLESPACE users FORCE LOGGING;  
Tablespace altered.
```



Nocompress

 Completed schema successfully

Statistic	Value
Connection Time	0:00:03.590
Data Generation Time	0:17:48.876
DDL Creation Time	0:12:07.338
Total Run Time	0:29:59.814
Rows Inserted per sec	130,982
Data Generated (MB) per ...	9.8

The creation of the schema appears to have been successful.


Status	Object Name
Valid	COUNTRIES, CHANNELS, CUSTOMERS, PROMOTIONS, PRODUCTS, SUPPLEMENTARY_DEMOGRAPHICS, SALES, TIMES, CUSTOMERS_MARITAL_BIX, CUSTOMERS_YOB_BIX, CUSTOMERS_GENDER_BIX, PRODUCTS_PROD_STATUS_BIX, PRODUCTS_PROD_SUBCAT_IX, PRODUCTS_PROD_CAT_IX, SALES_PROD_BIX, SALES_CUST_BIX, SALES_CHANNEL_BIX, SALES_PROMO_BIX,

OK

```
Data Generation time      0:18
DDL Creation time        0:12
Total Run Time          0:30

Rader generert/sekund 130.982
```

Compress

 Completed schema successfully

Statistic	Value
Connection Time	0:00:03.613
Data Generation Time	0:18:33.213
DDL Creation Time	0:12:10.001
Total Run Time	0:30:46.835
Rows Inserted per sec	125,766
Data Generated (MB) per ...	9.4

The creation of the schema appears to have been successful.

Status	Object Name
Valid	COUNTRIES, CHANNELS, CUSTOMERS, PROMOTIONS, PRODUCTS, SUPPLEMENTARY_DEMOGRAPHICS, SALES, TIMES, CUSTOMERS_MARITAL_BIX, CUSTOMERS_YOB_BIX, CUSTOMERS_GENDER_BIX, PRODUCTS_PROD_STATUS_BIX, PRODUCTS_PROD_SUBCAT_IX, PRODUCTS_PROD_CAT_IX, SALES_PROD_BIX, SALES_CUST_BIX, SALES_CHANNEL_BIX, SALES_PROMO_BIX,

OK

```
Data Generation time      0:18
DDL Creation time        0:12
Total Run Time          0:30

Rader generert/sekund 125.766
```

Resultat

```
Nocompress      130.982
Compress        125.766
```

Compress / Nocompress

```
SH@OUGN SQL> SELECT table_name,
                    blocks * 8192 /1024 /1024 MB,
                    num_rows
                    FROM user_tables
                    WHERE table_name = 'SALES';
```

TABLE_NAME	MB	NUM_ROWS
SALES	3810.6	100000000

```
SH@OUGN SQL> SELECT table_name,
                    blocks * 8192 /1024 /1024 MB,
                    num_rows
                    FROM user_tables
                    WHERE table_name = 'SALES';
```

TABLE_NAME	MB	NUM_ROWS
SALES	3436.8	100000000

Pakke DBMS_COMPRESSION

```
SET SERVEROUTPUT ON
DECLARE
  blkcnt_comp PLS_INTEGER;
  blkcnt_uncomp PLS_INTEGER;
  row_comp PLS_INTEGER;
  row_uncomp PLS_INTEGER;
  comp_ratio PLS_INTEGER;
  comp_type VARCHAR2(30);
BEGIN
  dbms_compression.get_compression_ratio(
    SCRATCHTBSNAME => 'USERS',
    OWNNAME => 'SH',
    TABNAME => 'CUSTOMERS',
    PARTNAME => NULL,
    COMPTYPE => dbms_compression.comp_for_oltp,
    BLKCNT_CMP => blkcnt_comp,
    BLKCNT_UNCMP => blkcnt_uncomp,
    ROW_CMP => row_comp,
    ROW_UNCMP => row_uncomp,
    CMP_RATIO => comp_ratio,
    COMPTYPE_STR => comp_type);
  dbms_output.put_line('Block Count Compressed: ' || TO_CHAR(blkcnt_comp));
  dbms_output.put_line('Block Count UnCompressed: ' || TO_CHAR(blkcnt_uncomp));
  dbms_output.put_line('Row Count Compressed: ' || TO_CHAR(row_comp));
  dbms_output.put_line('Row Count UnCompressed: ' || TO_CHAR(row_uncomp));
  dbms_output.put_line('Block Count Compressed: ' || TO_CHAR(comp_ratio));
  dbms_output.put_line('Compression Type: ' || comp_type);
END;
```

ORACLE®
DATABASE 11gR2

DBMS_COMPRESSION

```
Block Count Compressed:      759
Block Count UnCompressed: 1444
Row Count Compressed:        73
Row Count UnCompressed:      38
Block Count Compressed:      2
Compression Type:            "Compress For OLTP"
```

ORACLE®
DATABASE **11gR2**

PL/SQL procedure successfully completed.

```
CREATE TABLE sales (
  prod_id      NUMBER      NOT NULL,
  cust_id      NUMBER      NOT NULL, ... )
PCTFREE 5 NOLOGGING NOCOMPRESS
PARTITION BY RANGE (time_id)
(partition sales_2008 VALUES LESS THAN (TO_DATE(...))
  COMPRESS BASIC,
 partition sales_2009 VALUES LESS THAN (MAXVALUE)
  COMPRESS FOR OLTP );
```

Komprimering Basic / OLTP*

- **BASIC**

- Komprimerer data ved bruk av «Insert direct path load».

- **OLTP**

- I hovedsak for OLTP database, håndterer «INSERT, UPDATE» og «DELETE».

```
2190,13770,25-NOV-00,S,9999,23,161
2225,15720,28-NOV-00,S,9999,25,1450
34005,120760,29-NOV-00,P,9999,44,2376
9425,4750,29-NOV-00,I,9999,11,979
1675,46750,29-NOV-00,S,9999,19,1121
```

```
2190,13770,25-NOV-00,S,%,23,161
2225,15720,28-NOV-00,S,%,25,1450
34005,120760,*,P,%,44,2376
9425,4750,*,I,%,11,979
1675,46750,*,S,%,19,1121
```

Symboltabell:

* 29-NOV-00

% 9999

*(Tidligere syntaks: COMPRESS FOR ALL OPERATIONS.)

Datapump

- Store datamengder
- Partisjonering
- Parallellisering

Tabeller

- **Tabell - Import via databaselink**

```
debian:DVH> impdp sh/sh parfile=SH-sales.par  
  
tables=SH.SALES  
job_name=SALES  
network_link=OUGN-dblink  
parallel=16  
logfile=SALES.log
```

Partisjonerte tabeller

- Kan ikke bruke nettverks-link for å laste inn partisjon:
Eksport -> dumpfil -> Import

```
debian:OUGN> expdp sh/sh parfile=SH-sales.par  
  
tables=SH.SALES:SALES_Q4_2009  
job_name=SALES_Q4_2009  
dumpfile=SALES_Q4_2009.dmp  
parallell=16  
logfile=SALES_Q4_2009-export.log  
compression=ALL
```

```
debian:DVH> impdp sh/sh parfile=SH-sales.par  
  
tables=SH.SALES  
job_name=SALES  
dumpfile=SALES_Q4_2009.dmp  
parallell=16  
logfile=SALES_Q4_2009-import.log
```

Seriell indeksering

```
debian:OUGN> impdp system attach=Sales_Job
```

```
Import: Release 11.2.0.2.0 - Production on Mar 15 18:56:24 2011  
...
```

```
Worker 1 Status:
```

```
Process Name: DW00
```

```
State: EXECUTING
```

```
Object Schema: SH
```

```
Object Name: SALES_CUST_BIX
```

```
Object Type:
```

```
SCHEMA_EXPORT/TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Completed Objects: 5
```

```
Worker Parallelism: 4
```

```
Worker 2 Status:
```

```
Process Name: DW01
```

```
State: WORK WAITING
```

```
Worker 3 Status:
```

```
Process Name: DW02
```

```
State: WORK WAITING
```

```
...
```

Import av *SH* skjema

```

$ time impdp system DUMPFILE=SH.dmp PARALLEL=4 LOGFILE=SH.log

Import: Release 11.2.0.2.0 - Production on Mar 15 18:56:24 2011
Processing object type SCHEMA_EXPORT/USER
Processing object type SCHEMA_EXPORT/SYSTEM_GRANT
Processing object type SCHEMA_EXPORT/ROLE_GRANT
Processing object type SCHEMA_EXPORT/DEFAULT_ROLE
Processing object type SCHEMA_EXPORT/TABLESPACE_QUOTA
Processing object type SCHEMA_EXPORT/PRE_SCHEMA/PROCACT_SCHEMA
Processing object type SCHEMA_EXPORT/TABLE/TABLE
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA

. . imported "SH"."SALES":"SALES_1996"          217.5 MB 6681012 rows
. . imported "SH"."SALES":"SAL"                216.9 MB 6662310 rows
. . imported "SH"."SALES":"SALES_H2_1997"      109.3 MB 3359310 rows
. . imported "SH"."SALES":"SALES_H1_1997"      107.5 MB 3304160 rows
. . imported "SH"."SALES":"SALES_Q2_2000"      54.15 MB 1662906 rows
...
Processing object type TABLE/INDEX/INDEX
Processing object type TABLE/CONSTRAINT/CONSTRAINT

Real: 19m 42s

```

Import av *SALES* tabell

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Sat Apr 2 18:50:59 2011  
Table "SH"."SALES" exists and has been truncated.  
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows  
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows  
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows  
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows  
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows  
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows  
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows  
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows  
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...  
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT  
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX  
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 2h 39m 42s
```

DEMO

Import av *SALES* tabell

Uten indekser

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Sun Apr 3 09:24:39 2011  
Table "SH"."SALES" exists and has been truncated.
```

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows  
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows  
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows  
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows  
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows  
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows  
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows  
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows  
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...
```

```
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT
```

```
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 1m 50s
```

Import av *SALES* tabell

Indeksering sist...

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Sun Apr 3 14:27:11 2011  
Table "SH"."SALES" exists and has been truncated.
```

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows  
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows  
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows  
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows  
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows  
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows  
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows  
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows  
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...
```

```
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT
```

```
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 7m 57s
```

SALES tabell

Eller indeksering etterpå...

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Sun Apr 3 14:27:11 2011
```

```
Table "SH"."SALES" exists and has been truncated.
```

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...
```

```
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT
```

```
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 1m 57s
```

Import av *SALES* tabell

```
Ora11gR2 SH@OUGN SQL> SELECT index_name,  
                                index_type,  
                                num_rows,  
                                leaf_blocks  
                                FROM user_indexes  
                                WHERE table_name = 'SALES';
```

INDEX_NAME	INDEX_TYPE	NUM_ROWS	LEAF_BLOCKS
SALES_PROMO_BIX	BITMAP	74882	34265
SALES_CHANNEL_BIX	BITMAP	20030	10024
SALES_CUST_BIX	BITMAP	97756282	357614
SALES_PROD_BIX	BITMAP	47947	23720

Indeks eksisterer før import: 2 timer og 39 minutter

Indeks eksisterer ikke før import: 2 minutter

Indeks opprettes etter import: 8 minutter

SALES tabell

Indeksering etterpå

```
Ora11gR2 SH@OUGN SQL> @SALES_PROMO_BIX  
Index created.
```

```
Elapsed: 00:00:27.27
```

```
Ora11gR2 SH@OUGN SQL> @SALES_PROD_BIX  
Index created.
```

```
Elapsed: 00:00:23.94
```

```
Ora11gR2 SH@OUGN SQL> @SALES_CHANNEL_BIX  
Index created.
```

```
Elapsed: 00:00:18.11
```

```
Ora11gR2 SH@OUGN SQL> @SALES_CUST_BIX  
Index created.
```

```
Elapsed: 00:07:50.75
```

```
Ora11gR2 SH@OUGN SQL> @SALES_CUST_IX  
Index created.
```

```
Elapsed: 00:02:28.22
```

Fremmednøkler

```
Ora11gR2 SH@OUGN SQL> SELECT constraint_name,  
                        status  
                        FROM user_constraints  
                        WHERE table_name = 'SALES'  
                        AND constraint_type = 'R';
```

CONSTRAINT_NAME	STATUS
SALES_TIME_FK	ENABLED
SALES_CUSTOMER_FK	ENABLED
SALES_CHANNEL_FK	ENABLED
SALES_PRODUCT_FK	ENABLED
SALES_PROMO_FK	ENABLED

```
Ora11gR2 SH@OUGN SQL> ALTER TABLE sales  
                        DISABLE CONSTRAINT sales_time_fk;
```

Table altered.

```
Ora11gR2 SH@OUGN SQL> ALTER TABLE sales  
                        DISABLE CONSTRAINT sales_customer_fk;
```

Table altered.

...

Import av *SALES* tabell

Uten aktive fremmednøkler

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Sun Apr 3 17:30:26 2011
```

```
Table "SH"."SALES" exists and has been truncated.
```

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...
```

```
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT
```

```
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 1m 47s
```

Import av *SALES* tabell

Med aktive fremmednøkler

```
$ time impdp sh DUMPFILE=SH.dmp PARALLEL=4 TABLES=SH.SALES
```

```
Import: Release 11.2.0.2.0 - Production on Mon Apr 4 19:20:43 2011
```

```
Table "SH"."SALES" exists and has been truncated.
```

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
```

```
. . imported "SH"."SALES": "SALES_Q2_2002"      54.07 MB 1660421 rows
. . imported "SH"."SALES": "SALES_Q3_2007"      54.72 MB 1680309 rows
. . imported "SH"."SALES": "SALES_Q2_2009"      54.06 MB 1660091 rows
. . imported "SH"."SALES": "SALES_Q1_2009"      53.48 MB 1642265 rows
. . imported "SH"."SALES": "SALES_Q4_2008"      54.72 MB 1680462 rows
. . imported "SH"."SALES": "SALES_Q3_2008"      54.73 MB 1680625 rows
. . imported "SH"."SALES": "SALES_Q2_2008"      54.07 MB 1660513 rows
. . imported "SH"."SALES": "SALES_Q1_2008"      54.09 MB 1661172 rows
. . imported "SH"."SALES": "SALES_Q4_2007"      54.70 MB 1679883 rows
```

```
...
```

```
Processing object type TABLE/CONSTRAINT/REF_CONSTRAINT
```

```
Processing object type TABLE/INDEX/FUNCTIONAL_AND_BITMAP/INDEX
```

```
Processing object type TABLE/STATISTICS/TABLE_STATISTICS
```

```
Real: 21m 8s
```

SALES tabell

Aktivering av fremmedenøkler

```
SH@OUGN SQL> alter table sales enable constraint SALES_TIME_FK;  
Table altered.
```

```
Elapsed: 00:01:35.08
```

```
SH@OUGN SQL> alter table sales enable constraint SALES_CUSTOMER_FK;  
Table altered.
```

```
Elapsed: 00:02:18.98
```

```
SH@OUGN SQL> alter table sales enable constraint SALES_CHANNEL_FK;  
Table altered.
```

```
Elapsed: 00:00:46.14
```

```
SH@OUGN SQL> alter table sales enable constraint SALES_PRODUCT_FK;  
Table altered.
```

```
Elapsed: 00:00:41.87
```

```
SH@OUGN SQL> alter table sales enable constraint SALES_PROMO_FK;  
Table altered.
```

```
Elapsed: 00:00:49.27
```

```
Real: 6m 10s
```

Alternativer? (P)CTAS

- PCTAS – Parallel Create Table As Select...

```
dvh@OUGN SQL> INSERT /*+ APPEND */  
                INTO dvh.oppgave_post dv1o_op  
                SELECT *  
                FROM dvh.oppgave_post@dv1p_datapump dv1p_op  
                WHERE aar = 2008;
```

```
dvh@OUGN SQL> INSERT /*+ APPEND */  
                INTO dvh.oppgave_post dv1o_op  
                SELECT *  
                FROM dvh.oppgave_post@dv1p_datapump dv1p_op  
                WHERE aar = 2009;
```

```
dvh@OUGN SQL> INSERT /*+ APPEND */  
                INTO dvh.oppgave_post dv1o_op  
                SELECT *  
                FROM dvh.oppgave_post@dv1p_datapump dv1p_op  
                WHERE aar = 2010;
```

Bigfile!

```
SYS@OUGN SQL> CREATE BIGFILE TEMPORARY TABLESPACE dvh_temp_01
                TEMPFILE '/u03/oradata/OUGN/dvh_temp_01.dbf'
                SIZE 10M AUTOEXTEND ON NEXT 10M MAXSIZE 1T
                TABLESPACE GROUP dvh_temp_group;
```

Tablespace created.

```
SYS@OUGN SQL> CREATE BIGFILE TEMPORARY TABLESPACE dvh_temp_02
                TEMPFILE '/u04/oradata/OUGN/dvh_temp_02.dbf'
                SIZE 10M AUTOEXTEND ON NEXT 10M MAXSIZE 1T
                TABLESPACE GROUP dvh_temp_group;
```

Tablespace created.

```
SYS@OUGN SQL> SELECT * FROM dba_tablespace_groups;
```

GROUP_NAME	TABLESPACE_NAME
DVH_TEMP_GROUP	DVH_TEMP_01
DVH_TEMP_GROUP	DVH_TEMP_02

```
SYS@OUGN SQL> ALTER DATABASE
                DEFAULT TEMPORARY TABLESPACE dvh_temp_group;
```

Database altered.

Krympe temporære områder

```
SYS@OUGN SQL> SELECT tablespace_size/1024/1024/1024 "GB",
                    allocated_space/1024/1024/1024 "Brukt GB",
                    free_space/1024/1024/1024 "Ledig GB"
                    FROM dba_temp_free_space;
```

```
      GB Brukt GB Ledig GB
-----
      2370      2370      2370
```

```
SYS@OUGN SQL> ALTER TABLESPACE temp SHRINK SPACE KEEP 1G;
Tablespace altered.
```

```
SYS@OUGN SQL> SELECT tablespace_size/1024/1024/1024 "GB",
                    allocated_space/1024/1024/1024 "Brukt GB",
                    free_space/1024/1024/1024 "Ledig GB"
                    FROM dba_temp_free_space;
```

```
      GB Brukt GB Ledig GB
-----
           1           1           1
```

Store filer på avveier...

```
11gR2 SYS@OUGN SQL> SELECT d.file_name,  
                             d.file_id,  
                             MAX(d.bytes)/1024/1024 TOTAL,  
                             MAX(block_id)*8192/1024/1024 MB,  
                             MAX(d.maxbytes)/1024/1024 MAKS  
FROM dba_data_files d, dba_extents e  
WHERE d.file_id = e.file_id  
      AND d.tablespace_name = 'USERS'  
      AND d.maxbytes > 0  
GROUP BY d.file_name, d.file_id  
ORDER BY d.file_id;
```

FILE_NAME	FILE_ID	TOTAL	MB	MAKS
/oradata/OUGN/OUGN_users_01.dbf	4	11460	97	32768

```
Or11gR2 SYS@OUGN SQL> PURGE TABLESPACE users;  
Tablespace purged.
```

```
Or11gR2 SYS@OUGN SQL> ALTER DATABASE DATAFILE 4 RESIZE 100M;  
Database altered.
```

Recover Manager

```
run {  
  allocate channel DISK1 device type DISK;  
  allocate channel DISK2 device type DISK;  
  allocate channel DISK3 device type DISK;  
  backup as compressed backupset  
    filesperset 1  
    incremental level 0  
    database  
    plus archivelog delete input;  
}
```

Recovery Manager

```
RMAN> restore datafile 1 from tag='TAG20110404T073445';
```

```
Starter gjenoppretting klokken 05/04/2011 06:41:37  
bruker kanalen ORA_DISK_1
```

```
gjenoppretter datafilen 00001 til  
/oradata/OUNG/OUNG_system_001.dbf  
kanal ORA_DISK_1: leser fra sikkerhetskopidelen  
/rman/OUNG/OUNG_84m8tn4m.bck  
kanal ORA_DISK_1: gjenopprettet sikkerhetskopidel 1  
delreferanse=/rman/OUNG/OUNG_84m8tn4m.bck  
kode=TAG20110404T073445  
kanal ORA_DISK_1: gjenoppretting fullført, medgått tid: 03:52:46
```

```
kanal ORA_DISK_1: gjenoppretting fullført, medgått tid: 00:01:23
```

Recovery Manager i full fart...

```
RMAN> sql 'ALTER TABLESPACE example OFFLINE';

RMAN> COPY DATAFILE 5 TO '/u07/oradata/UGN/UGN_example_01.dbf';

RMAN> SWITCH DATAFILE 5 TO COPY;
datafile 5 switched to datafile copy "/u07/UGN/UGN_example_01.dbf"

RMAN> sql 'ALTER TABLESPACE example ONLINE';

RMAN> DELETE COPY LIKE '/u06/UGN/UGN_example_01.dbf';

List of Datafile Copies
=====
Key          File S Completion Time          Ckp SCN   Ckp Time
-----
3           5      A 18.03.2011 20:52:28 26868693 18.03.2011 20:51:20
Name: /u6/UGN/UGN_example_01.dbf

Do you really want to delete the object (enter YES or NO)? Yes
deleted datafile copy
```

Oppsummering

De fleste oppgaver tar
IKKE
særlig lang tid.

Hvis korrekt statistikk er på plass.

**Og underliggende strukterer
IKKE
er for komplekse.**

Det var alt for denne gangen.

Takk for meg!