SQL Server 2000 Performance Tuning

Decus 2004

Juergen.Krautner@hp.com HP-Services

Whats Up ?

- Traceflags
- Profiler
- Traces
- Indizes
- Executionplans

- Wait types
- others

- Need Sysadmin Rights
- dbcc traceon(xxx, ... , -1)
  - Setting Traceflags impacts all connections which issued a traceon !
- dbcc tracsestatus(-1)

- 8602 - turn off all *index* hints for the current connection.
- 8755 - turn off all *locking* hints for the current connection.
- 8722 - turn off all *other* hints (other than index and locking)

- 2520 – DBCC extended help
- 3604 - DBCC Sends trace output to the client
- 4022 - Bypasses automatically started procedures
- 8687 - Used to disable query parallelism.
- 8721 - Dumps info into the error log when AutoStats has been run.

Traceflags Example

- **set statistics io on**
- **set Statistics time on**

- **select * from [order details]**
  --Table 'Order Details'. Scan count 1, logical reads 10, physical reads 0
  -- elapsed time = 467 ms.

- **select * from [order details] order by discount desc**
  --Table 'Order Details'. Scan count 1, logical reads 10, physical reads 0,

- **select * from [order details] with (index(DiscountIdx)) order by discount desc**
  --Table 'Order Details'. Scan count 1, logical reads 4315, physical reads
  --elapsed time = 566 ms.
- **dbcc traceon (8602)** **-- turn off indexhints**

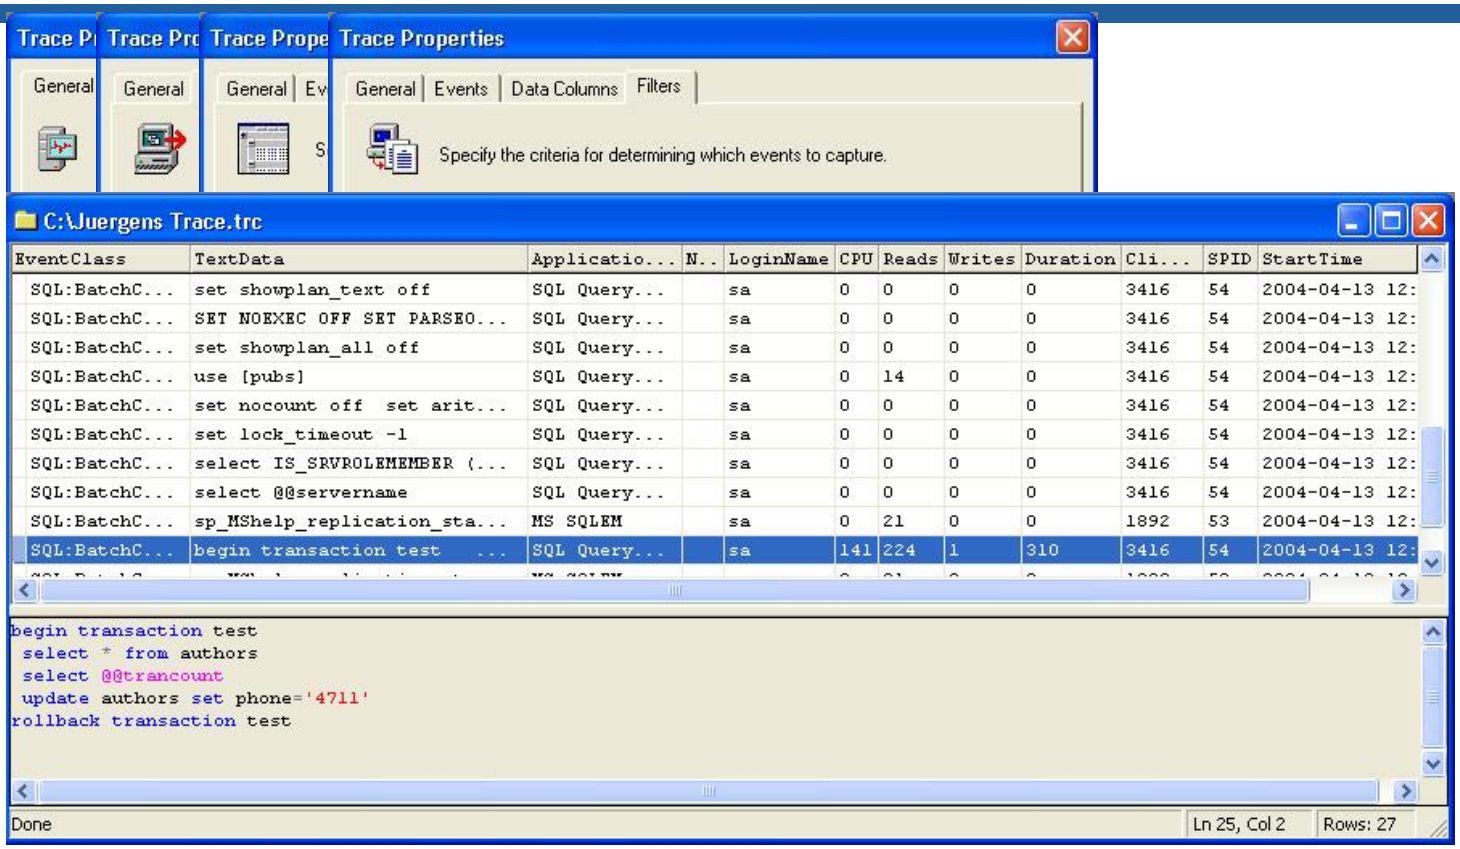

- Sort via properties/datacolums/grouping (up/down columns)
- http://support.microsoft.com/default.aspx?kbid=224587

Profiler ...

- Events of interest:
  - Database - DataFile & Logfile AutoShrink
  - Errors - Missing Column Statistics
  - Locks - Deadlock/ Deadlockchain/ Timeout
  - Objects - Autostats / Opened
  - Performance - Execution Plan / Show plan statistics
  - Stored Proc - cache miss / recompile / starting /stmtstarting / completed
  - TSQL - batch starting / stmtstarting / completed
  - Session - existing connections (settings)

- Quick Tracefile analysis
  - SELECT * Into #temptab
    FROM ::fn_trace_gettable('c:\itwinput.trc',default)

- **Declare @tid int**
- **EXEC sp_trace_create @traceid = @tid OUTPUT ,**
  @options = 0 , @tracefile = N'c:\autotrace.trc'
- **EXEC sp_trace_setevent @traceid = @tid,**
  @eventid = 34, @columnid = 6, @on = @bon
- **EXEC sp_trace_setfilter @traceid = @tid,**
  @columnid = 6, @logical_operator = 1,
  @comparison_operator = 6, @value = N'sa'
- **EXEC sp_trace_setstatus @traceid = @tid, @status = 1**
  -- 0=Stop,1=Start,2=Del
- **SELECT * FROM :: fn_trace_getinfo(default)** **-- Info**
- **sp_trace_setstatus @traceid =2,@status = 0** **-- Stop**

Blackbox & Stackdump

- Usefull on Crashes and Hangs
- Results (...\data\blackbox.trc) analyzed with Profiler
- Contains SP/TSQL Execution, Errors/Warnings
- Setup
  - **USE master**
    CREATE PROC startblackbox
    AS
    declare @TraceID int
    exec sp_trace_create @TraceID output, 8 (8=TRACE_PRODUCE_BLACKBOX)
    exec sp_trace_setstatus @traceID, 1 (1=start)
    RETURN
  - Mark the black box to start at startup of SQL Server
    exec sp_procoption startblackbox, 'startup', 'on',

- SEM
  - Table info
  - Manage indexes
  - Script
  - sp_helpindex ,tablename'

- TSQL:
  - exec sp_indexes_rowset 'titles'
  - exec sp_indexes_rowset ; 2 ,all tables in a DB

TABLE_NAME PRIMARY_KEY UNIQUE **CLUSTERED** TYPE **FILL_FACTOR**
INITIAL_SIZE **NULLS** SORT_BOOKMARKS **AUTO_UPDATE**
COLLATION ORDINAL_POSITION **COLUMN_NAME**
COLLATION CARDINALITY **PAGES** FILTER_CONDITION INTEGRATED

Index statistics

- Histogram / Density
- DBCC SHOW_STATISTICS (mytable, idxName)

Statistics for INDEX 'idxName'.

Updated	Rows	Rows Sampled	Steps	Density	Average key length
Apr 13 2004 3:00PM	1024	1024	4	0.0	30.0

All density	Average Length	Columns
0.25	30.0	lname

RANGE_HI_KEY	RANGE_ROWS	EQ_ROWS	DISTINCT_RANGE_ROWS	AVG_RANGE_ROWS
Juergen	0.0	256.0	0	0.0
Paul	0.0	256.0	0	0.0
Peter	0.0	256.0	0	0.0
Thomas	0.0	256.0	0	0.0

- Use Profiler
  - collect Execution Plans
  - Collect statements

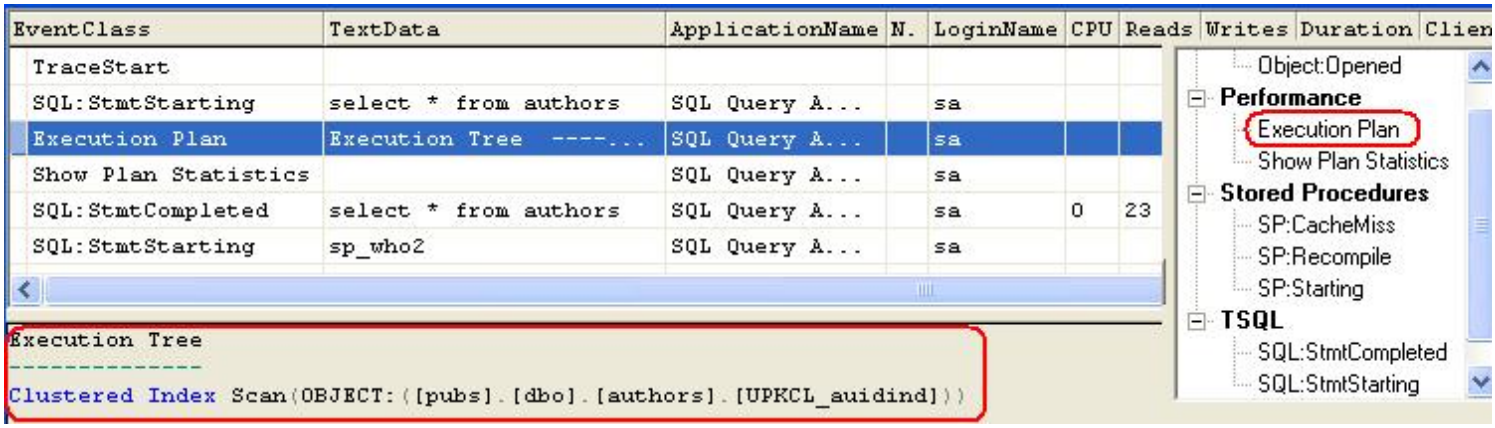

- **Use Index Tuning Wizard**
  - Load profiler trace
  - Select tables
  - Look at recommendations and INDEX USAGE REPORT
  - (based on IUR you might distribute index to different disk)

Execution plans

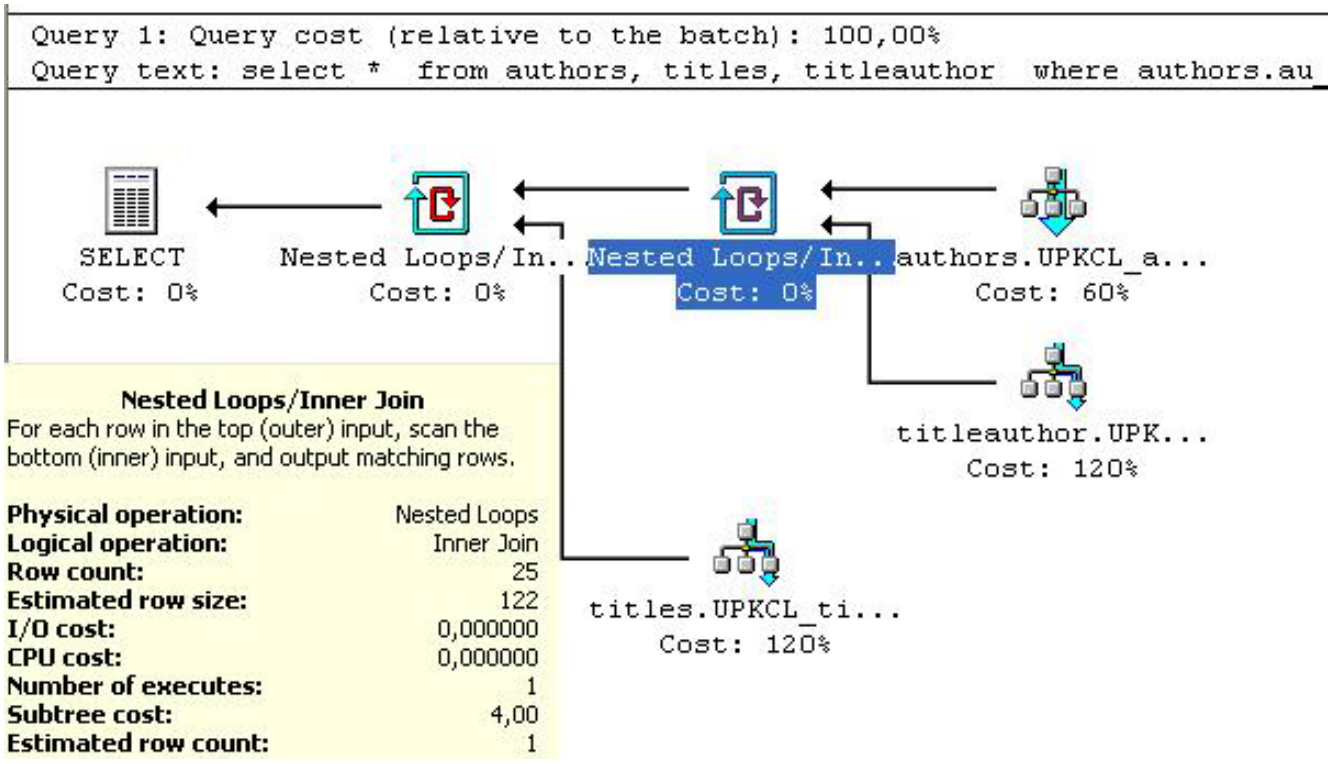

- SET STATISTICS PROFILE ON.

```
Rows  Executes StmtText ..EstimateRows ..EstimateIO ..EstimateCPU TotalSubtreeCost
------  --------- ----------------------------------------------------------------------------------------------------
25      1         select *  from authors, titles, titleauthor
                            where authors.au_id = titleauthor.au_id
                            AND titles.title_id = titleauthor.title_id
25      1         |--Nested Loops(Inner Join, OUTER REFERENCES:([titleauthor].[title_id..
25      1            |--Nested Loops(Inner Join, OUTER REFERENCES:([authors].[au_id])).
23      1            |    |--Clustered Index Scan(OBJECT:([pubs].[dbo].[authors].[UPKC..
25      23           |    |--Clustered Index Seek(OBJECT:([pubs].[dbo].[titleauthor].[..
25      25           |--Clustered Index Seek(OBJECT:([pubs].[dbo].[titles].[UPKCL_titl..
```


Index Concepts

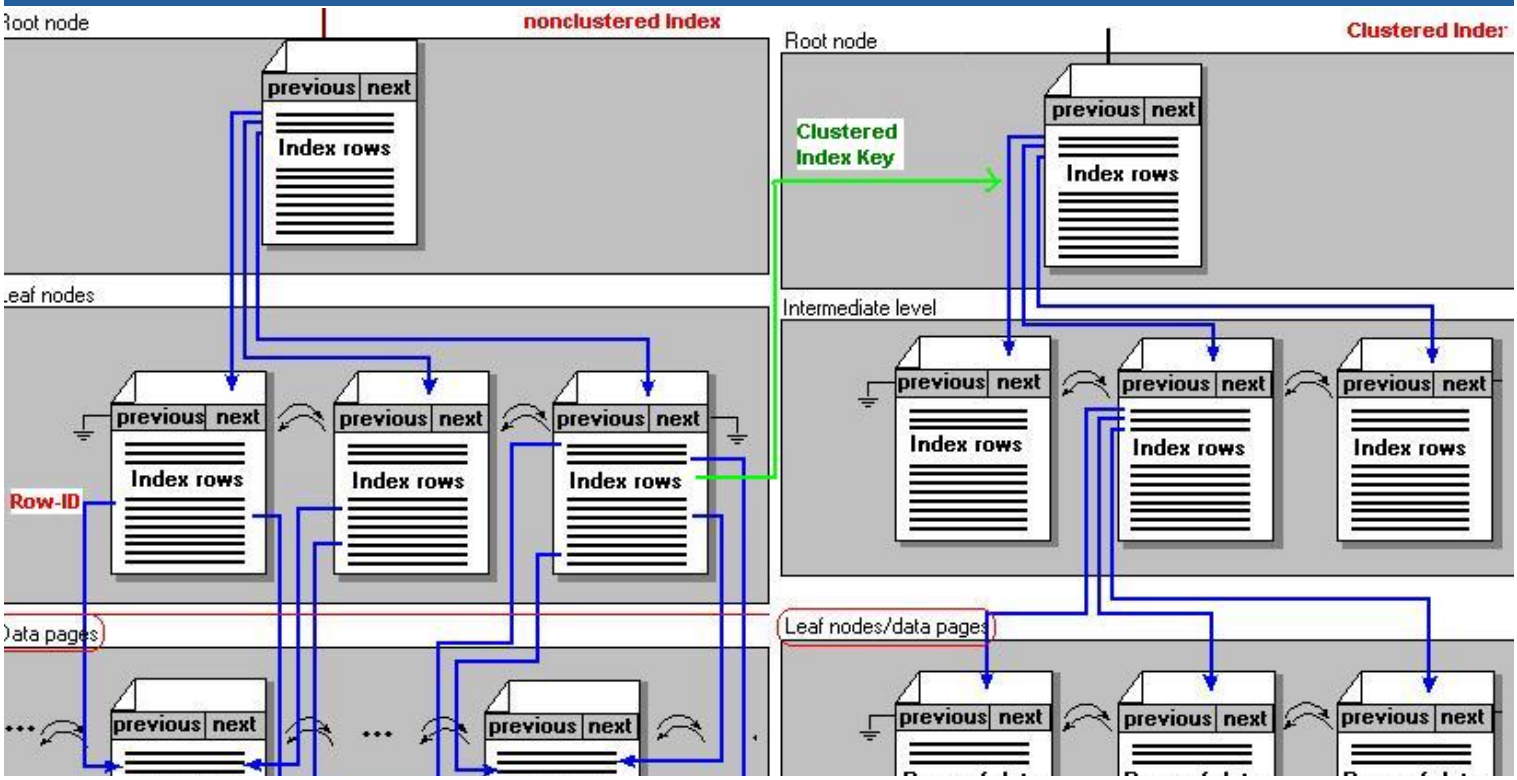

- Index Seek - vertical scan - root-leaf-data
- Index Scan - horizontal scan - root-leaf-leaf-leaf ...

- Cluster I.Seek - I. contains data ! look for specific value
- Cluster I.Scan - scan data pages (sorted)
- Table Scan - scan data pages (unsorted)

- Note: Files read in parallel.

- **Possibility to set Execution Time Limits**
  - **query governor cost limit (ms)**

Indexes

```
set statistics io on

SELECT au_lName, au_fName, city FROM authors ...

...with (index (0)) WHERE au_lName LIKE '[M-Z]%'
  -- read via cluster index on au_id
  -> clust. index scan: Scan count 1, logical reads 2, physical
  reads 0

...with (index (iLName)) WHERE au_lName LIKE '[M-Z]%
   -- read via lastname index
  -> index seek + bookmarklookup  Scan count 1, logical reads
  23, phys0

created covering index containing all columns
...WHERE au_lName LIKE '[M-Z]%'
```

```
Query 1: Query cost (relative to the batch): 20,00%
Query text: SELECT au_lName, au_fName, city FROM authors with (index (0)) WHEI
```

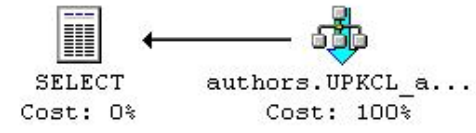


```
Query 2: Query cost (relative to the batch): 40,00%
Query text: SELECT au_lName, au_fName, city FROM authors with (index (iLName))
```


```
Query 3: Query cost (relative to the batch): 40,00%
Query text: SELECT au_lName, au_fName, city FROM authors  WHERE au_lName LIKE
```

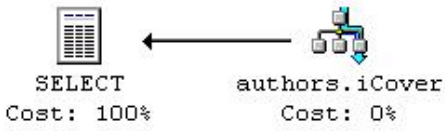

Indexes ..

- SET STATISTICS IO ON
- DBCC DROPCLEANBUFFERS
- SELECT COUNT(*) FROM Orders
  - Table 'Orders'. Scan count 1, logical reads 3, physical reads 1, read-ahead reads 2.
- SELECT COUNT(CustomerId) FROM Orders
  - Table 'Orders'. Scan count 1, logical reads 21, physical reads 1, read-ahead reads 20.
- SELECT * FROM Orders
  - Table 'Orders'. Scan count 1, logical reads 21, physical

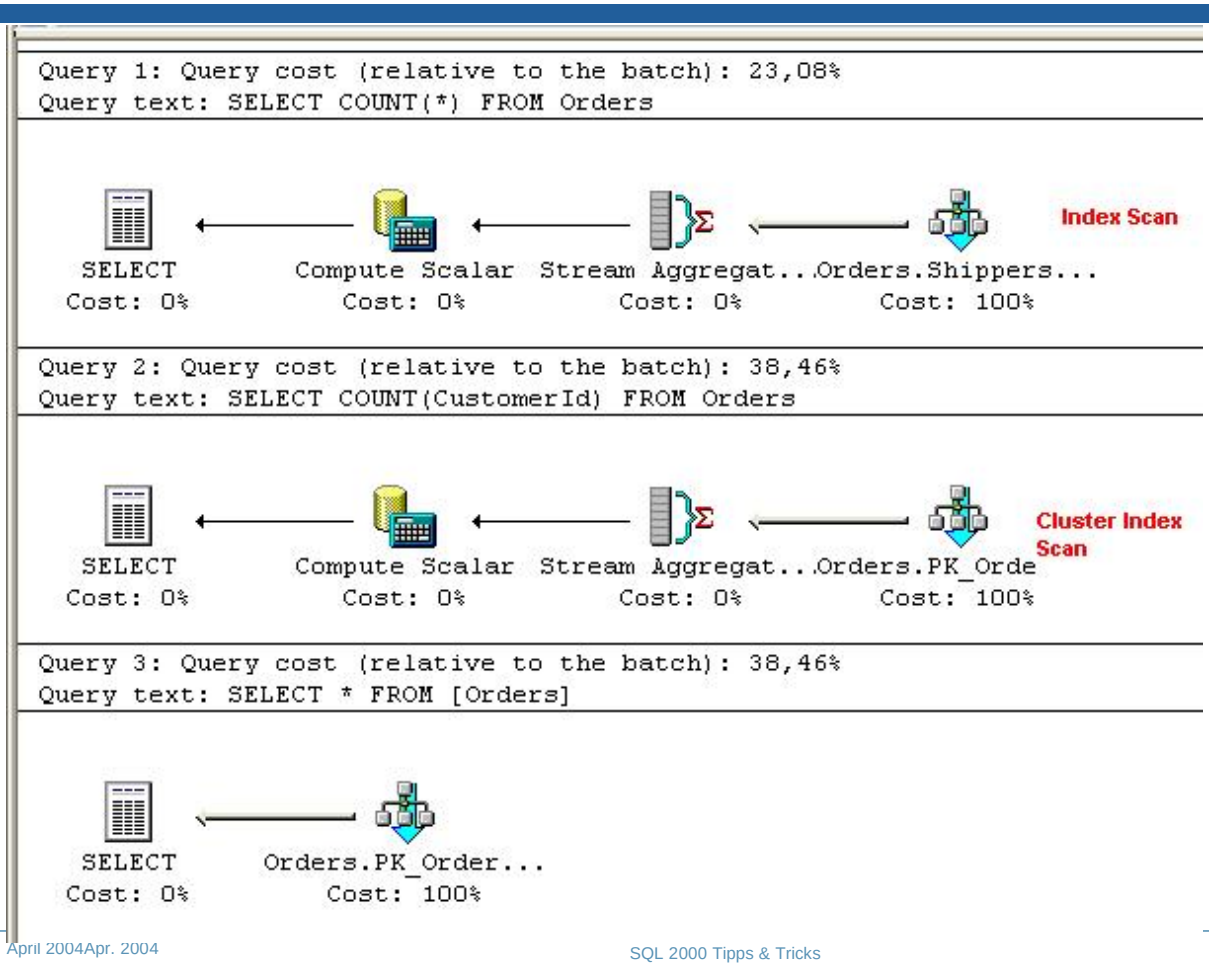

DBCC Showcontig('tablename') [with all_indexes]

DBCC SHOWCONTIG scanning ‚mytest' table...
Table: ‚mytest' (933578364); index ID: 1, database ID: 5
TABLE level scan performed.
- Pages Scanned.....: 4167
- Extents Scanned.....: 521
- Extent Switches.....: 520
- Avg. Pages per Extent.....: 8.0
- Scan Density [Best Count:Actual Count].....: 100.00% [521:521]
- Logical Scan Fragmentation .....: 11.21% (%out of order)
- Extent Scan Fragmentation .....: 0.96%
- Avg. Bytes Free per Page.....: 198.6 (low)
- Avg. Page Density (full).....: 97.55% (high)

DBCC execution completed. If DBCC printed error messages, contact your system administrator.

- **DBCC DBREINDEX**

ait Type	Requests	Wait Time	Signal Wait Time
ISCELLANEOUS	0.0	0.0	0.0
CK_M_S	2.0	881.0	0.0
LEEP	5760.0	5767433.0	5767433.0
O_COMPLETION	242.0	11525.0	40.0
SYNC_IO_COMPLETION	4.0	1942.0	0.0
ESOURCE_SEMAPHORE	0.0	0.0	0.0
TC	0.0	0.0	0.0
LEDB	0.0	0.0	0.0
AILPOINT	0.0	0.0	0.0
ESOURCE_QUEUE	1210.0	5879965.0	5763297.0
SYNC_DISKPOOL_LOCK	33.0	0.0	0.0
MS_THREAD	0.0	0.0	0.0
.....			
RITELOG	28.0	921.0	0.0
EMPOBJ	0.0	0.0	0.0
OGMGR	0.0	0.0	0.0
URSOR	0.0	0.0	0.0
XECSYNC	0.0	0.0	0.0
ATCH DT	0.0	0.0	0.0
AGELATCH_DT	0.0	0.0	0.0
AGEIOLATCH_DT	0.0	0.0	0.0
RAN_MARK_DT	0.0	0.0	0.0
ETWORKIO	8.0	10.0	0.0
otal	7581.0	1.1691357E+7	1.1550869E+7

Sqlmag.com

- Sysprocesses is not cumulative. Only last wait info
- Per Process , not systemwide

- Provides two stored procedures
  - Track_waitstats
    - Creates *waitstats* table
    - Populates table with output from *dbcc sqlperf (waitstats)*
  - Get_waitstats
    - Displays accumulated summary !
    - Sorted output !

- Similiar from MS

```
wait type                              wait time              percentage
-------------------------------------- ----------------       -------------
***total***                             71.0                     100.0
WRITELOG                                60.0                     84.5
NETWORKIO                               6.5                      9.1
LCK_M_X                                 4.5                      6.4
MISCELLANEOUS                           .0                       .0
IO_COMPLETION                           .0                       .0
PSS_CHILD                               .0                       .0
EXCHANGE                                .0                       .0
...
...
...
...
```


I/O

- Waittypes
  - IO_COMPLETION -WRITELOG
  - ASYNC_IO_COMPLETION -LOGMGR
- Possible Reasons
  - Disk Throughput
- Check
  - fn_virtualfilestats iostallms
  - Perfmon:

    Physical Disk: Disk sec/read Disk sec/write Disk queues Q-length

    SQL Access: Full Scans/sec Index seeks/sec pagesplits/sec

    SQL Buffer: Page Life Expectancy Checkpoint pages/sec

- Waittypes
  - LATCH_xxx / PAGELATCH_xxxx
    short term light weight synchronization objects to protect internal structures.
  - possible Reasons
    - Memory shortage
    - TempDB Stress
    (http://support.microsoft.com/default.aspx?scid=kb;en-us;328551)
  - check
    - Perfmon: SQL Latch waits, Page Faults/sec, pages/sec
  - PAGEIOLATCH_xxxx
    used for disk to memory transfers.
    - Check diskio

Locks

- Waittypes
  - LCK_xxxx
- Possible Reasons
  - Application design
  - Nested / Long running transactions
  - Memory problems causing high I/O
- check
  - `Select * from syslockinfo`
  - Sp_who2
  - Sp_locks

- Waittypes
  - UMS_THREAD
- Possible Reasons
  - CPU
  - Shortage on workerthreads
- check
  - **Dbcc Sqlperf (umsstats)**
  - **Increase „max worker threads“**

Network

- Waittypes
  - NETWORKIO
    waiting for an acknowledgement
- Possible Reasons
  - Network contentions / collisions / Throughput
- Check
  - **Test bandwidth**
  - **Perfmon: Packets/sec Bytes/Sec**
  - **optimize „network packet size“**

- Track Recompiles
  - Causes locks on the „shared compile plan“
- Profiler
  - SP:recompile, SP:Starting, SP:StmtStarting
    (Pay attention to EventSubClass. 2=threshhold, 3=not found, 4=set)
  - Object: autostats
- Reasons
  - Row threshholds changed (different for temp/perm tables)
  - Usage of SET ANSI_xxx, Forceplan, Quoted_identifer ....
    (=> use SET at connectionstart instead)
- What you might do
  - Use Select .... KeepPlan on Temp Tables to delay recompile (>6)
  - Use Select .... Keeplan fixed to turn off statistical recompile
  - Use sp_executesql (one compile, parameters)
  - Use sp_autostats to disable autostats on certain tables

http://support.microsoft.com/default.aspx?scid=kb;en-us;Q243586
http://msdn.microsoft.com/library/default.asp?url=/library/en-us/trblsql/tr_reslsyserr_2_54ky.asp

General Hints

- **Consider EE vs Standard Edition**
  - Parallelism, indexed views, >4 proc.
- **Check**
  - Consider Fibers (light weight pooling) if cpu almost 100% and high ctxtswitch
  - OS-Setting Max throuput for network apps.
  - OS-Setting Optimize performance for background apps (6*quantum)
- **Physical Optimizations**
  - fit Rows into Pages
  - Veritical partitioning (split long tables)
  - Horizontal partitioning (consider archiv tables. Sales1999)
  - Distribute Indexes according to their usage
- **Try Best Practice Analyzer**
  - Rulebased check of your DB (mem settings, torn page, implicit Trans ....)

- Books Online
  - Database,Application Design
  - Query Tuning, Tempdb Tuning
  - Optimizing Tools , Optimizing Servers

- www.sqlservercentral.com
- www.sql-server-performance.com
- www.sqlmag.com
- www.umachandar.com/resources.htm

- SELECT * FROM :: fn_virtualfilestats(*dbid, fileid*)
- In master do -1

```
SELECT sysdatabases.name AS Database_Name,
       sysdatabases.dbid AS Database_ID,
       sysaltfiles.fileid AS File_ID
FROM   sysaltfiles INNER JOIN sysdatabases
       ON sysaltfiles.dbid = sysdatabases.dbid
ORDER BY sysdatabases.name, sysaltfiles.fileid
```

DbId	FileId	TimeStamp	NumberReads	NumberWrites	BytesRead	BytesWritten	IoStallMS
7	2	75041640	25	160	2380288	4946944	0

Calculate IoStallMS /(NumberReads+NumberWrites) >15-20msec ?