

P A S S

PROFESSIONAL ASSOCIATION
FOR SQL SERVER

PASS Deutschland e.V.


Wie lege ich eine Datenbank richtig an

Original Title
CREATE DATABASE...

=tg= Thomas Grohser

select * from =tg=

@@Version	Remark
SQL 4.21	First SQL Server ever used (1994)
SQL 6.0	First Log Shipping with failover
SQL 6.5	First SQL Server Cluster (NT4.0 + Wolfpack)
SQL 7.0	2+ billion rows / month in a single Table
SQL 2000	938 days with 100% availability
SQL 2000 IA64	First SQL Server on Itanium IA64
SQL 2005 IA64	First OLTP long distance database mirroring
SQL 2008 IA64	First Replication into mirrored databases
SQL 2008R2 IA64	First 256 CPUs & >500.000 STMT/sec
SQL 11 (Denali)	Can't wait to push the limits even further



Thomas Grohser, SQL Server MVP, bwin Interactive Entertainment AG

<http://www.grohser.com/>

Focus on SQL Server Infrastructure Architecture and Implementation

Close Relationship with Microsoft

SQLCAT (SQL Server Customer Advisory Team)

SCAN (SQL Server Customer Advisory Network)

TAP (Technology Adoption Program SQL2008R2 and SQL11)

Active PASS member and PASS Summit Speaker





World's biggest publicly listed online gaming platform

World's **leading provider** of online Sports Betting

One of the largest **Poker networks**

Comprehensive range of **Payment Service Providing**

Integrated gaming portal - **22 languages, 25 core markets**

Gross gaming revenues 2008 (GGR):
EUR **421 million**

More than **20 million registered customers**

1,500 employees

bwin builds on the strengths of the web in order to **tie up responsibility and gaming**

15 million page views and up to **980,000 users a day**



CREATE DATABASE - Agenda

- **Elements of a database**
- **File groups**
- **Data files**
- **Log file(s)**
- **Options**
- **Planning**

Elements of a Database

- **File groups**
 - Used to group data logical
 - Separate partitions
 - Allow piecemeal restore
- **Data files**
 - Holding the actual data, indexes, metadata
- **Log file(s)**
 - Record all transactions
- **Filestream files**
 - Large binary or text data stored as one file per record

File groups

- **Each database has at least one file group**
- **When to use more than one?**
 - Every time I want to separate data
- **You can specify the file group for**
 - a table
 - an index
 - a partition

Data files

- **Each file group must have at least one data file**
- **When do I want to have more than one file?**
 - Manageability: you don't want a file to be too large
 - Performance: you can spread the workload
 - Careful: 2 or more files on the same physical volume have no effect on read / write performance during normal operation
 - Exception one physical disk, multiple partitions, different drive letters for each partition and fast file initialization option is not active (windows user that runs SQL server does not have the "Perform volume maintenance task" user right), then restore database or create database works faster

Data files

- **Optimize for fast read**
 - Random for OLTP
 - Sequential for DWH
 - Always 100% write cache
- **Size recommendation**
 - I like to limit the size to be smaller than the smallest physical drive I use
 - right now 146GB → keep a single file below 100 GB
 - !!! Never run a 100GB database on a single disk in production.

Data files

- **Initial size**

- No effect on internal structure, make a good guess

- **Resizing**

- Grow No problem
 - Exception multiple files on a disk → physical fragmentation but no effect on internal structure
- Truncate No problem
- Shrink Don't, never ever

Data files

- **Multiple files for performance**
 - Make sure they all have the same size
 - If you add files later try to keep the free space in all files the same

Log file(s)

- **Most important file(s) of a database**
- **Usually one per database**
 - Reasons for more than one file
 - Disk space too small (often during maintenance)
 - Faster restores
 - Strict performance requirements
(> 400.000 TSQL per sec)

Log file

- **Initial size matters**
 - Log file is internally organized in virtual log files (VLF)
 - Number of VLF's depends on size/size change of log file
 - between 4 and 16

Log file

- **Growing**

- New VLF's are created based on the size of the growth

- Always change the default from 10% auto grow to something useful

Log file

- **Optimize Drive for**
 - Low Latency
 - Sequential Write
 - RAID 1
 - RAID 1/0 Careful to add enough spindles
 - Solid State Device
 - Always 100% write cache

Log file

- **Recommendation**
- **Create in steps to keep VLF size where it should be:**
 - Max size 25% of the size of your log drive controller cache
 - 512MB cache

OPTIONS

- **Recovery Model**
- **Collation**
- **Auto Shrink – DON'T EVER USE**
- **Thorn Page Detection and Checksum**
- **Statistics**

Recovery Model

- **You care about your data**
 - FULL RECOVERY

- **You don't**
 - Simple or Bulk

Collation

- **Don't mess with it**
- **If you need specific collations specify them on the table/column level**

Auto shrink

**MESSAGE TO THE MICROSOFT SQL
DEV TEAM:**

**THIS FEATURE SHOULD BE REMOVED
FROM THE PRODUCT !!!**

!!! DO NOT USE EVER !!!

Thorn Page & Checksum

- **Activate, use, turn on, ...**
- **Believe me if your database has a problem you want to know...**

Statistics

- **Auto create statistics**
- **Auto update statistics**
 - Yes you want them
- **Auto update statistics async**
 - Depends
 - Your query pays the bill
 - Next guy pays the bill

Planning

- **On a less than \$ 5000 Server bought in the last year**
- **OLTP**
 - Database size < 10GB hard to screw it up
 - Database size < 50GB start taking care
 - Database size < 100GB take care
 - Database size > 500GB start extensive planning
- **DWH**
 - Database size < 300GB hard to screw it up
 - Database size < 1 TB start taking care
 - Database size < 10 TB take care
 - Database size > 50 TB start extensive planning

Questions?

=tg= Thomas Grohser

tg@grohser.com



<http://www.grohser.com/>