

#sqlhelp

Your friend on twitter...

MORE RESTRICTIVE



M1 Discussing how to limit MAXDOP.

Instance MAXDOP is least restrictive as it can be overridden by a query hint, but only up to the limit imposed by Resource Governor.


RG WORKLOAD GRP
MAXDOP


QUERY HINT MAXDOP

INSTANCE MAXDOP




M1S7 Explaining that dividing a large database into filegroups allows targeted restores during disaster recovery, rather than having to wait for the entire single file database to restore.





M1S7 Explaining a different partitioning scenario – auto parts sales to WA/CA/OR/UT.

If things go down, it makes sense to restore the portion of the database that will bring in the most sales quickly – CA, as it has the most population.




SA

DBCC PAGE

M1S14 Explaining why instant file initialization is off by default. Sharing a volume between SQL Server and 'secure files' can lead to information leakage if the secure files are deleted, then SQL grows a data file, and an SA can use DBCC PAGE to get a hex dump of the bits/bytes that comprised the deleted secure file. If you don't have this situation, turn instant file initialization on.

M1S33 Explaining that backup striping makes restores and backups faster. One reader thread per distinct drive letter/mount point, and one writer thread similarly.


Bwin in 2005 could backup a 2TB database in 36 minutes, backing up to 12 files each on a separate drive.




A hand-drawn red diagram. On the left, a red arrow points upwards and then curves to the right, pointing towards the number '1117'. The number '1117' is drawn in red, with three vertical lines for the '1's and a curved line for the '7'. Below the number, a long red horizontal line with an arrow at its right end points to the right.

M1 Discussing trace flag 1117. It's a way to work around the tempdb bug where only the largest data file will grow. Beware – it applies to all databases and whenever a file needs to grow, all files in that filegroup will grow at the same time.

PFS



TF 1118 - 328551




M1S34 Explaining how adding more tempdb data files leads to round-robin allocation and a spread of the contention to multiple files, leading to lower overall contention.

PFS





PFS



M5S29 Explaining the
ACCESS_METHODS_HOBT_VIRTUAL_ROOT
latch. It protects access to the Storage
Engine metadata that holds the page ID of
an index's root page.

M5S29 Explaining PAGELATCH waits and a hotspot in an index with an INT identity column as the key. Many concurrent threads, each inserting 80 byte records, will lead to 100 inserts per page. There will be contention among the threads waiting to insert onto the same page.





SPINLOCK

LOCKS

ACQUIRE

LATCHES

INTERLOCKED

TEST-AND-SET-IF-CLEAR

4ms




M5 Explaining how spinlocks work. Code does a test-and-set-if-clear on the memory that implements the spinlock. If it can't get it, it spins (tries again) up to 1000 times and then backs off. SOS_SCHEDULER_YIELD is the wait type when a backoff occurs.

SPINLOCKS




can't



M6 Explaining why actual query plans are so expensive to create. The query plan has an extra operator for each 'real' operator, whose job is just to count the rows coming from each other operator.

100



M6 Explaining what a ring buffer is. Think of it as an array of data elements which is filled in a circular fashion. E.g. with 100 elements, filling the array and then writing the 101st element will cause element 1 to be overwritten.