

TF LACS


610
— 11
12

2528
79
52
89


M1 Explaining how trace flags work. The first 1-2 digits are the area (6 is Access Methods, 25,52,79,89 are DBCC).

CACHE LINE INVALIDATION

CACHE LINES
64b / 128b



M1S06 Discussing
why the NULL
bitmap is a perf
optimization by
avoiding having to
read record parts
into the CPU's L3
cache



ADD COLUMN

RE sys.syscolpars - syscolumns


SE sys.sysrscols } sysindexes
sys.sysiscols }

M1 Explaining the names of the hidden system tables that store column metadata.

DAC

M1 Explaining how you can only see the hidden system tables using the Dedicated Admin Connection. See [http://technet.microsoft.com/en-us/library/ms189595\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms189595(v=sql.105).aspx)

M1 1-many relationships for a table. The opposite does not apply. A table can only have 1 heap or 1 clustered index.



c1, c2... c10



HEAP
CL

SEL c1, c6, 8

W. c2 = x

PTR

PHYSICAL RID

LOGICAL



c7

M1S09 Explaining NC index back-links to the data records as part of describing why forwarding/forwarded records exist. Without them, moving a heap row would entail changing all NC index records with a back-link to that row.

(FILE : PAGE : SLOT)
(CL KEY(S))

UCL =


c2	c4
----	----

NUCL =


U	c2	c4
---	----	----

UNIQUEIFIER

M1S8 Explaining how for a nonunique clustered index, SQL Server adds a hidden uniquifier (tie-breaker) column to the cluster key(s) to guarantee uniqueness.




IN-ROW < 8000





M1S11 Text structure. Google Marc Rasmussen for his blog. He's blogged text tree structures.

100	1900 byte
-----	-----------




4 M1S12 Having a large LOB value in row that's not used much makes for large rows and low data density. Choose how to store LOB data wisely.






M1S14 When an update occurs to a record and snapshot isolation is enabled, the pre-change record is copied into the version store in tempdb. The changed record has a 14-byte tag on the end with the timestamp and pointer to previous record version. A chain of previous versions is possible.



M1S15 When a record is deleted, it's marked as a ghost record. The ghost cleanup task does not physically delete the records – it removes the slot array entry that points to the record on the page. I.e. the record is unmarked as being a record and the area on the page is now free space – that just happens to have bits and bytes that used to be a valid record.




0 — 7 8 — 15 16 — 23





M1S24 Explaining the difference between mixed and dedicated extents. Mixed extent is one in which the 8 pages could be allocated to 8 different objects. Dedicated extent is one where all 8 pages are reserved for exclusive use of a single object.

Later in the deck I refine 'object' to really be 'allocation unit'.



M1S36 Example of IAM chains.
Allocations from each 4GB
chunk of a data file require an
IAM page to track the
allocation. Multiple IAM pages
make an IAM chain.





0 - FH

1 - PFS

2 - GAM

3 - SGAM



4 - UNUSED

5 - UNUSED


6 - DIFF

7 - ML

M1 The page types in the first extent in a data file




M2S7 Splitting a large database up into multiple filegroups can enable you to do small, targeted restores or even a partial restore from scratch – in the event that the database is damaged or destroyed.



M2S7 Explaining some ways to split up a database to allow targeted restore. On the left, auto parts sales system – makes sense to bring parts and NSW online first.


STRIPING



R0


http://en.wikipedia.org/wiki/Standard_RAID_levels

MIRRORING



http://en.wikipedia.org/wiki/Standard_RAID_levels

NOTATING PARITY




http://en.wikipedia.org/wiki/Standard_RAID_levels

$$R10 = R1 + G$$

= MIRRORING

+ STRIPING



http://en.wikipedia.org/wiki/Standard_RAID_levels

sys.dm_io_virtual_file_stats
dm_io_pending_io_requests


M2 Tracking I/O latencies: see <http://www.sqlskills.com/blogs/paul/how-to-examine-io-subsystem-latencies-from-within-sql-server/>

wmic

diskpart


M2S8 The two commands you can use to check partition alignment. Don't use diskpart though as it rounds and can give false results.

RR




PROP FILL
WEIGHTING

~~8100~~ ←
SO (8099)

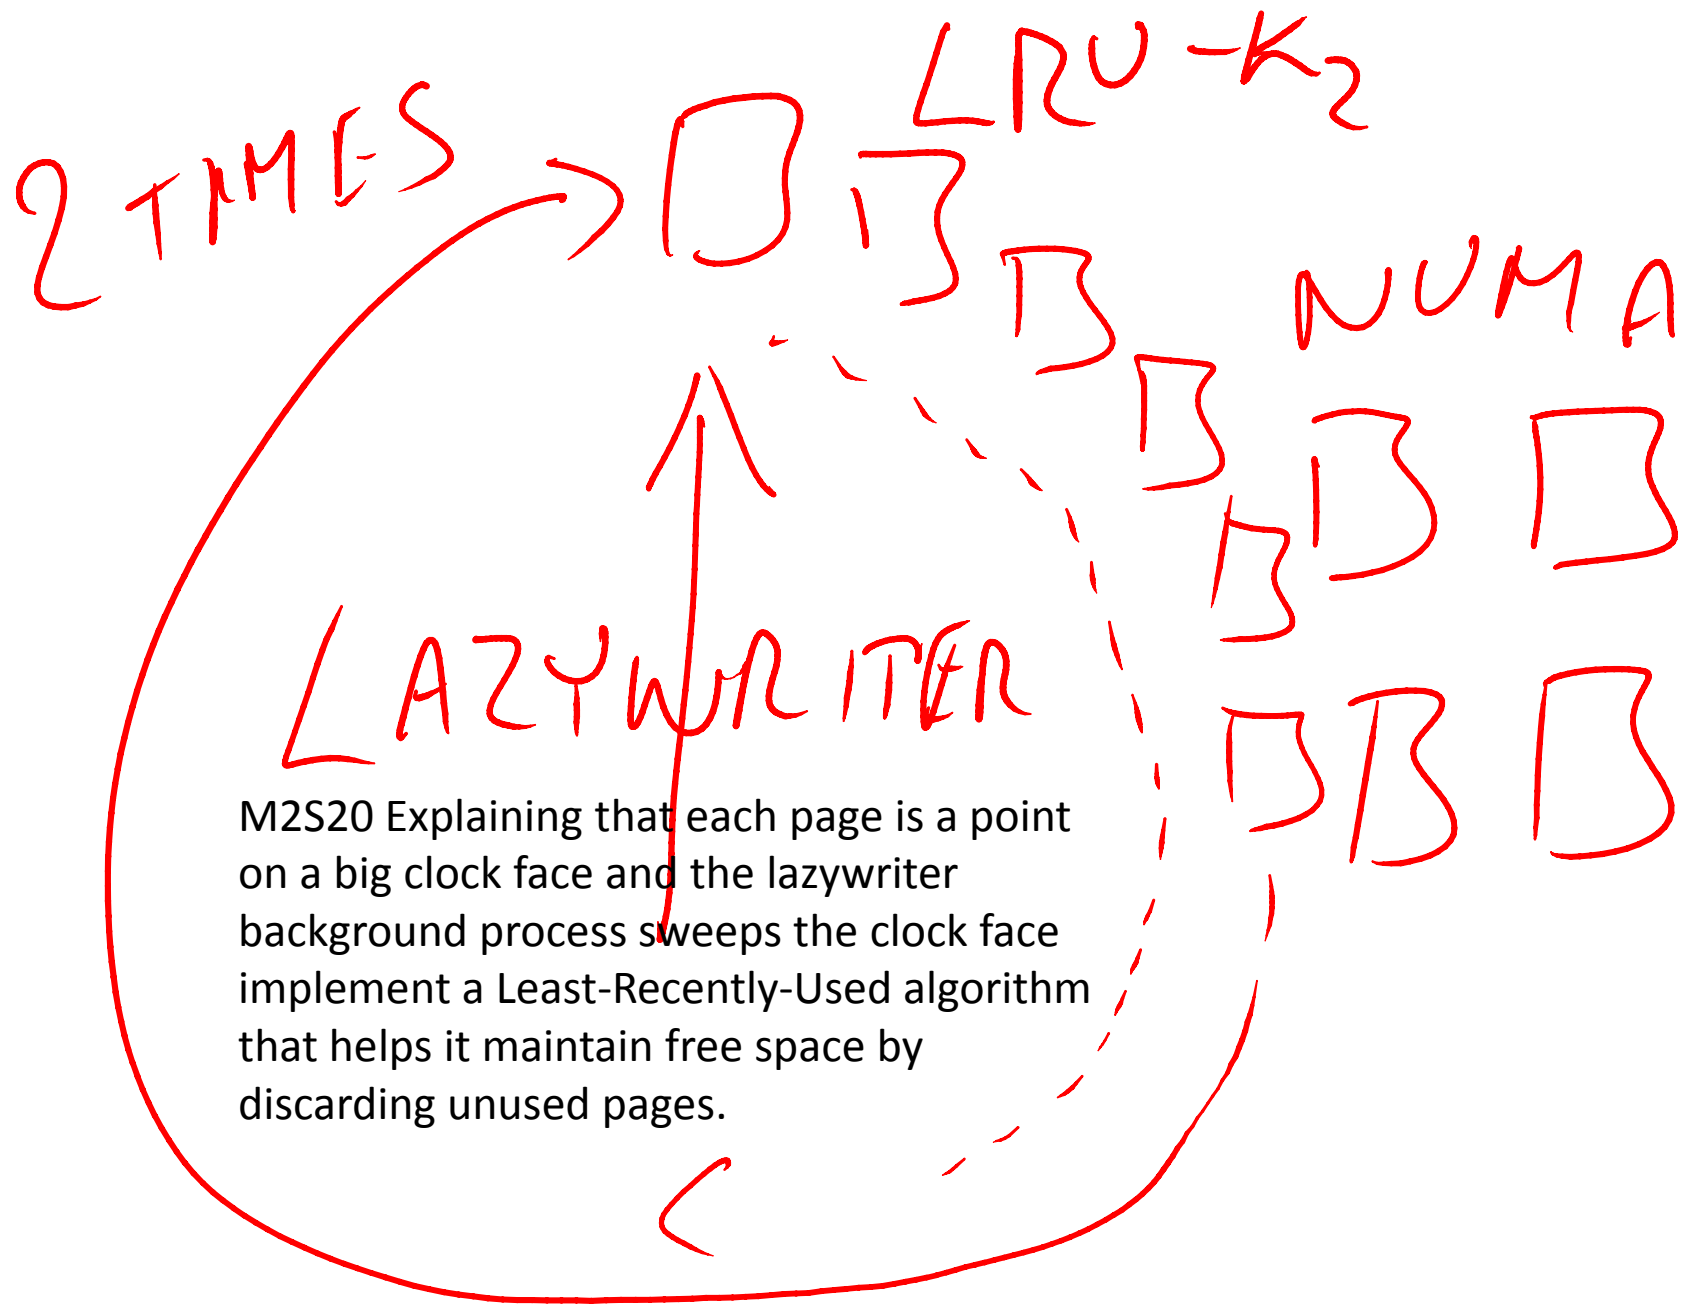


~~8100~~ ←
SO (8099)




1 ← ✓
8100

M2S19 Explaining round-robin
allocation and proportional fill



See <http://www.sqlskills.com/BLOGS/PAUL/post/Page-Life-Expectancy-isnt-what-you-think.aspx>



PLE
4000

PLE
4000

PLE
4000

PLE
4000

$$BM/PLE = (1+2+3+4) / 4 \downarrow$$

= 4000

1000


This should be Buffer
Node instead of Buffer
Partition - my mistake.

3625
BUFFER MANAGER $\Rightarrow$ BUFFER PARTITION


NUMA

NON-UNIFORM

MEM ACCESS





M2 Explaining NUMA, how each node has local memory that is faster to access than non-local (foreign) memory.




M2S22 Explaining zero initialization and how NTFS does not trust newly allocated space until the highwatermark of the file has increased, either through consecutive writes, or setting the HWM through a called to SetFileValidData.

M1S14 & M2S35 Explaining how there's a new portion of the version store created every minute.






M2S40 Explaining PFS and SGAM contention in tempdb and how adding multiple files spreads the contention over multiple files, thus reducing it.



M5S46 Explaining log flushing



ASYNC

< 2012

32 logblocks

48-16

2012 >


128-16

112 logblocks


sys.dm-io-pending-io-requests

M5S46 Explaining log flushing. There is a limit of 32 outstanding async log block I/Os pre-2012. The limit is 112 in 2012.


chkpt pages/sec



M5S35 Explaining how regular checkpoints look in perfmon (using Buffer Manager/Checkpoint pages/sec counter) and how they may exceed I/O capacity (the dotted line). Doing more frequent manual checkpoints, or using the new indirect checkpoints in SQL 2012, can result in more constant, lower IO load.



M5S42 It's possible for concurrent log backups to occur with a full backup from 2005 onward. Log clearing cannot happen though. It's deferred until the full backup completes, giving the impression that a full backup clears the log.



M5S48 Explaining the circular log demo. We discussed log space reservation that happens to ensure the active transactions can roll back without having to grow the log – this accounts for extra log growths. Space reserved in the log does not make a VLF become active – as there are no log records written out, just empty space reserved.

100,000


10

M5S60 Explaining that full recovery doesn't mean there's always one log record per operation. An index rebuild in full recovery will log full LOP_FORMAT_PAGE log records instead of multiple LOP_INSERT_ROWS log records. It's more efficient, but is still fully logged.


~~INS
INS
INS~~

F - PAGE
F - PAGE


10



M5S60 If a log backup contains a minimally-logged operation, it cannot be used to do a point-in-time restore using STOPAT for any point in time between the beginning and end of the backup. It can be used as part of a restore sequence as normal though.



M5S60 Explaining that if a minimally-logged operation is present in a log backup, you cannot do a point-in-time restore to any time covered by that backup (except the start or end of that time period). You also cannot do a tail-of-the-log backup after a crash that destroys the data files if there has been a minimally-logged operation since the last log backup. Well in 2008 R2/2012, you can, but the log backup is invalid.




PROBE



M5 Explaining part of the deferred-drop functionality. Before an extent can be deallocated, an X lock is acquired on it and all 8 page locks are probed (instant acquire X lock and release) to make sure nothing else has them locked.



M5S25 VLF fragmentation is bad because of having to search through the VLFs for a particular VLF sequence number. See the KB article link for more info.





M8S5 Explaining how a record is found on a page using binary search rather than brute force search.



M8S13 Explaining how fixed-size large rows leads to low page density and wasted space.

M8S15 Showing all
the operations
required when a
page split occurs.





LOP_DELETE_SPLIT

M8S18 Append-only insert pattern fills pages on right-hand side of index. New page allocations in this case are called page splits too – making all methods of tracking page splits largely useless (until 2012 with Xevents). You can look for LOP_DELETE_SPLIT log records to see splits occurring.

	NO SPLIT	SPLIT	X
1000	1256	7572	x5.5
100	356	6764	x18
10	268	11164	x47


M8S22 The numbers from the page split logging cost demo.

COMMENTS MEMBER ID

PAUL	PETER	SIMON	
------	-------	-------	--

DBCC DBREINDEX KIMBERLY

M8S23 The MySpace story. The 'K' is the tiny portion of comments on Kimberly's page because she wasn't very popular back then... ☺



$$FF = 70$$

M8S25 To leave 30% free space, set the fillfactor to 70.

ITB



M8S37 Explaining staggered index maintenance with database mirroring. See <http://bit.ly/IS04W5> for more explanation