

Chad Miller


Aaron Nelson @sqlvariant


Kendal Van Dyke @sgldba

Mike Fa1 @mike_fa1

SQL Server / Powershell gurus


VLFs in the transaction log –
discussion of how the log wraps
around and there are not empty VLFs
in the middle of the log






T \ DB \ - log - ldf


Transaction log growth – the log will grow if it is not truncated (e.g. by a log backup) and consume all space on the drive.



Full (F) and Log (L) backups – if there is a missing or bad full backup, you can still restore to a point in time using all log backups (top). If you are missing a log backup (bottom), then the last log backup is the point to which you can recover.



If full backups are also done via Veem, make sure they are not configured to truncate the transaction log. This also breaks the backup chain, and you can then only recover to the log backup taken right before the Veem full.



Example backup schedule used before an upgrade – full backup at 10 PM, stop services at 1AM, then run a differential, then two transaction log backups. The differential is not needed, just use a final log backup after everyone is out of the system to capture all changes.

Heap



A heap is a table that does not have a clustered index, so data is unordered. As data is inserted into the table, it is just added to pages already allocated that have free space. If no pages have free space, new pages are allocated and data is written to them.

CLUSTERED IDX

Person

ID

FN

LN

MI

SSN

DOB

CL key = ID

A clustered index has a clustering key (e.g. ID) by which the data is ordered. A clustered index has all the data in it (ID, first name, last name, middle initial, SSN, date of birth, shoe size, etc.).


NCI

LN

include (FN, DOB)

A nonclustered index has one or more columns from the table (whether a heap or clustered index) that are ordered by the index key (e.g. last name)

LN



Index structure (b-tree index, such as a clustered or nonclustered index in SQL Server). The top page in the index is also known as the root page. The bottom level in the index is known as the leaf level. Any levels in between (if they exist) are referred to as intermediate levels. The rows in the leaf level of the nonclustered index include the clustering key or row ID (RID) in case it is needed to look up additional columns from the clustered index or heap, respectively.

} leaf

CI



Clustered indexes, nonclustered indexes, and heaps are all comprised of pages that point to each other. The difference is whether the data is ordered (CI, NCI) and what data it contains (CI and heap have all the data).

NCI
(LN) inc. FN DUB



heap

