

CLR records

The recovery manager is responsible for performing undos on aborted transactions, because it is the same action as done during the undo phase of recovery. When a transaction is aborted, the prevLSN record of the abort record is identified as the first record that needs to be processed.

Let's take the below transaction

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							

Record 995 is identified as the first LSN to be undone. The undo action is performed, and a CLR record is written. The undoNextLSN field is set to the prevLSN field being processed. In the above situation, P40 offset 20 for 4 bytes is set from 41 to 40.

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							
1003	1001	T4	CLR	P40	4	20	41	40	995	994

Record 994 is identified as the next LSN to be undone. The undo action is performed, and a CLR record is written. The undoNextLSN field is set to NULL.

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							
1003	1001	T4	CLR	P40	4	20	41	40	995	994
1006	1003	T4	CLR	P20	2	10	IJ	GH	994	0

And since prevLSN is set to NULL, this is the last record to undo, so we can write an end_trans record.

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							
1003	1001	T4	CLR	P40	4	20	41	40	995	994
1006	1003	T4	CLR	P20	2	10	IJ	GH	994	0
1008	1006	T4	End_trans							

What happens if we crash in the middle of undoing? We do not want to redo undone operations. Let's say we look like this at the time of the crash.

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							
1003	1001	T4	CLR	P40	4	20	41	40	995	994

The lastLSN of T4 is set to 1003. That is added to the toUndo list. When 1003's turn to be processed is reached, it is processed as a CLR record. The undoNextLSN record 994 is added to the toUndo list. 994 is processed as above:

									LSN	undo
		Trans		Page					being	Next
LSN	prevLSN	ID	Type	ID	Length	Offset	Before	After	undone	LSN
994	0	T4	Update	P20	2	10	GH	IJ		
995	994	T4	Update	P40	4	20	40	41		
1001	995	T4	Abort							
1003	1001	T4	CLR	P40	4	20	41	40	995	994
1004	1003	T4	CLR	P20	2	10	IJ	GH	994	0
1005	1004	T4	End_trans							

Those operations that had already been undone are not redone, those operations that still needed to be undone are.