

Aries example - Redo phase

									LSN			
									being	undo		
LSN	prevLSN	TransID	Type	PageID	Length	Offset	Before	After	undone	NextLSN		
990	0	T3	Update	P30	3	20	ABC	DEF				
991	989	T2	Update	P10	3	1	123	456				
992	991	T2	Commit									
993	992	T2	End_trans									
994	0	T4	Update	P20	2	10	GH	IJ				
995	994	T4	Update	P40	4	20	40	41				
996	0	0	Beg_ckpt									
997	0		End_ckpt (Trans table, {(T3, 990), (T4, 995)}, Dirty Page Table {(P30, 990), (P40, 995), (P20, 994)})									
998	990	T3	Update	P30	4	10	10	11				
999	0	T5	Update	P50	4	20	12	13				
1000	0	T6	Update	P60	4	19	15	16				
1001	995	T4	Abort									
1002	998	T3	Update	P30	4	10	11	12				
1003	1001	T4	CLR	P40	4	20	41	40	995	994		
1004	999	T5	Update	P50	2	2	Mm	Nn				
1005	1004	T5	Update	P50	2	2	Nn	Oo				
1006	1003	T4	CLR	P20	2	10	IJ	GH	994	0		
1007	1000	T6	Update	P60	3	8	ABC	XYZ				
1008	1006	T4	End_trans									
1009	0	T7	Update	P70	1	0	K	M				
1010	1005	T5	Commit									
CRASH												

PageID	PageLSN
P10	991
P20	980
P30	981
P40	1003
P50	983
P60	984
P70	985

Disk images at the time of the crash

Trans table			Dirty Page Table		
Trans ID	LastLSN	Status		PageID	reLSN
T3	1002	U		P30	990
T5	1010	C		P40	995
T6	1007	U		P20	994
T7	1009	U		P50	999
				P60	1000
				P70	1009

Transaction and dirty page tables

Aries example - Redo phase

Redo phase:

First, find the smallest recLSN in the Dirty Page table - that is where we start redoing: 990

			Log LSN		
	Is this a	Page	greater	Log LSN	
Log	redoable	in dirty	than or	greater	
Seq.	log	page	equal to	than	
Num.	record?	table?	recLSN	PageLSN	Results
990	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P30 to 990
991	Y	N	N/A	N	Result - do nothing (not in DPT)
992	N				Result - do nothing (not a redoable log record)
993	N				Result - do nothing (not a redoable log record)
994	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P20 to 994
995	Y	Y	Y	N	Result - do nothing (PageLSN >= Log entry)
996	N				Result - do nothing (not a redoable log record)
997	N				Result - do nothing (not a redoable log record)
998	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P30 to 998
999	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P50 to 999
1000	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P60 to 1000
1001	N				Result - do nothing (not a redoable log record)
1002	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P30 to 1002
1003	Y	Y	Y	N	Result - do nothing (PageLSN >= Log entry)
1004	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P50 to 1004
1005	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P50 to 1005
1006	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P20 to 1006
1007	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P60 to 1007
1008	N				Result - do nothing (not a redoable log record)
1009	Y	Y	Y	Y	Result - redo changes, set the pageLSN of P70 to 1009
1010	N				Result - do nothing (not a redoable log record)

PageID	PageLSN
P10	991
P20	1006
P30	1002
P40	1003
P50	1005
P60	1007
P70	1009

Pages in buffers after redo

Finally, write an end_trans for each committed trans and remove from trans table