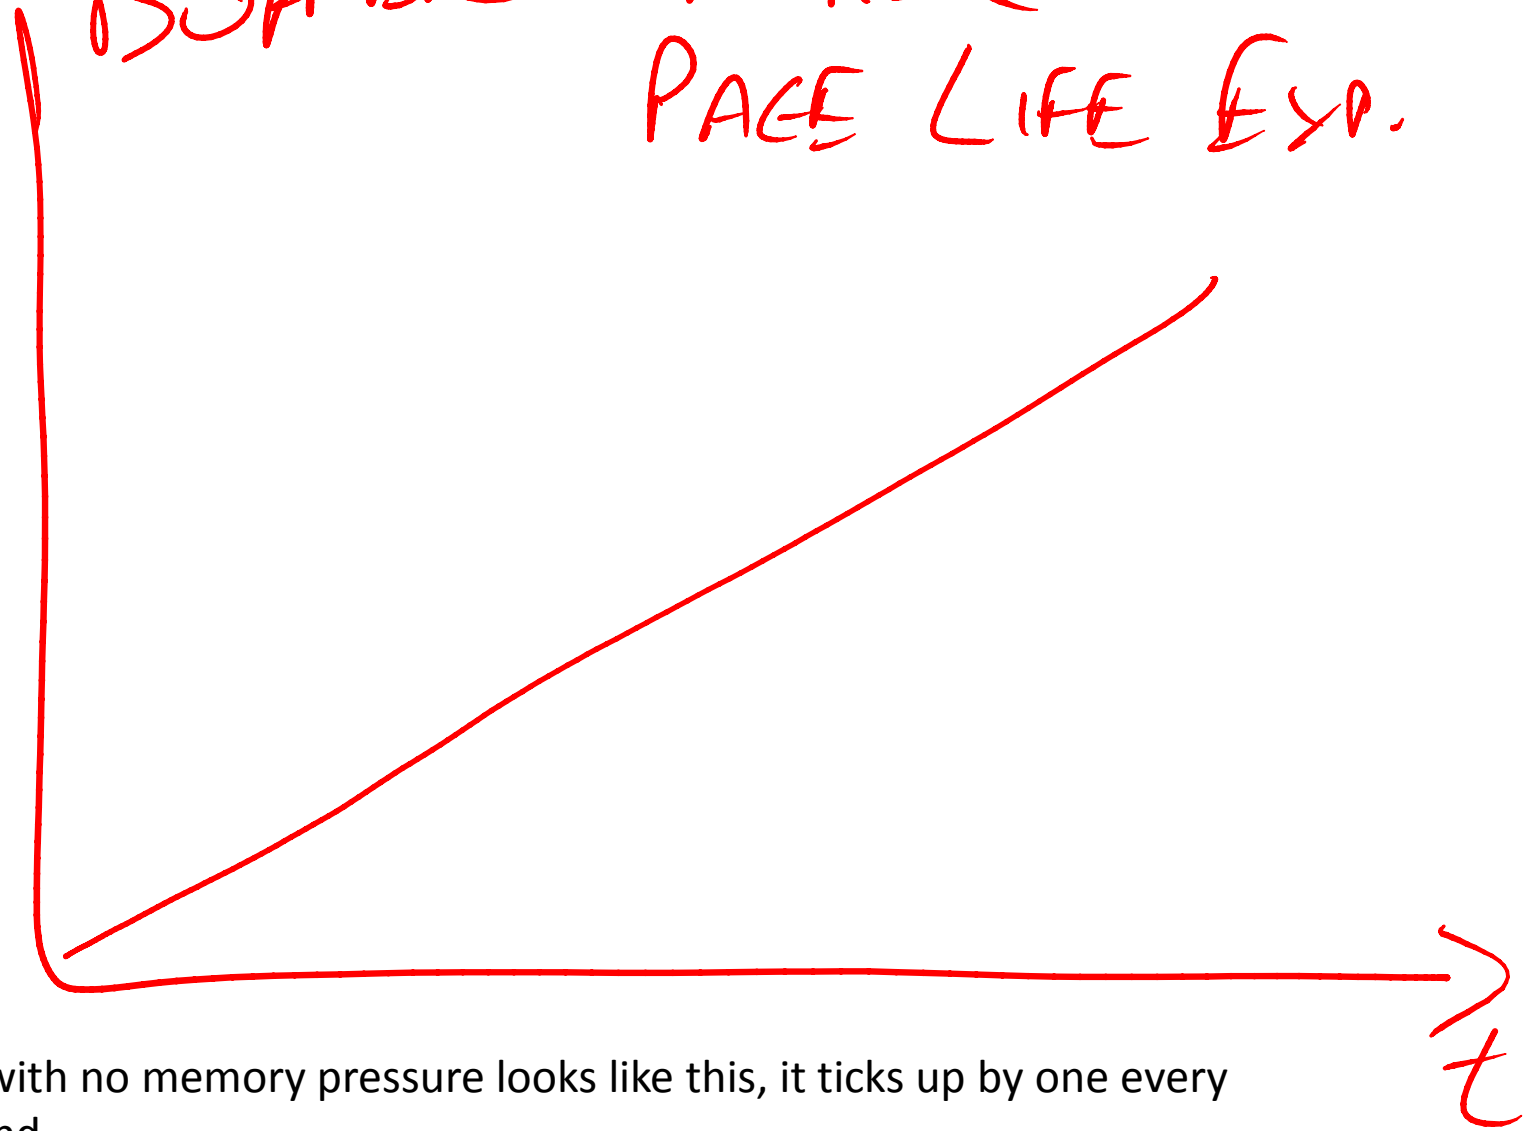



BUFFER MANAGER PAGE LIFE EXP.





PLE with no memory pressure looks like this, it ticks up by one every second





A graph like this is normal, as some things cause memory pressure, and then PLE comes back up again.

If you hear anyone saying a threshold for caring about PLE is 300, run the other way. Very old advice. That means your entire buffer pool is being flushed and reread every 300 seconds. The 'right' PLE value is whatever is normal on your server. If it drops below that and stays low, you have a problem.







NUMA = Non-Uniform Memory Access. Each node has some processor cores and node-local memory. Buffer pool is split between NUMA nodes.




On NUMA systems, you need to monitor each Buffer Node PLE, not the Buffer Manager PLE (as it's a mean). See <http://www.sqlskills.com/blogs/paul/page-life-expectancy-isnt-what-you-think/>




Explaining how automatic page repair may not be instant if the redo queue on the mirror hasn't caught up to the LSN on the principal when the bad page was discovered.







Explaining how full logging of an index rebuild logs complete page images rather than individual inserts, as it's more efficient that way.



PROBING


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.



PFS
(2:1:1)

SGAM
(2:1:3)

PAGE LATCH_UP
_EX



PFS

SGAM

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