



VM layout

This slide and the next two slides were from the discussion about storage configuration for VM servers. If the disk layout physically for a SQL Server VM is a single shared VMFS volume, then there is no point in having separate VMDK's on the same storage for separating transaction logs, data files, tempdb, and backups since the underlying storage is just a single shared pool. The goal in splitting things up is to have isolation of the underlying physical disk configurations to separate out the sequential writes for the logs from the random reads of the data files, and to isolate tempdb on its own for optimum performance benefit.





KB 816517

This is the KB article that discusses the Optimize Performance for Background Services option in Windows Server, which was asked about during class.

Page Life Expectancy

The established value for Page Life Expectancy being ≥ 300 is not a good measure for current servers. Consider a server with 144GB RAM installed in it. If we set aside 24GB for the OS and memory requirements outside of the buffer pool for SQL Server, we would have a buffer cache that is 120GB in size. At 300 seconds of page life this equates to:

$$122880\text{MB}/300\text{sec} = 409.6 \text{ MB/sec}$$

This is the sustained transfer for Random I/O that the disk I/O subsystem would have to be able to do to keep up with the page reads occurring on the system. While this number is not impossible it is far short of the typical disk I/O configuration for most SQL Server environments. However, if you know your I/O subsystem throughput limits, you can inverse the formula to figure out what a low PLE would be.

$$122880\text{MB}/280\text{MB/sec} = 438.85 \sim 439\text{sec PLE}$$

This would be your disk saturation point so you would want to ideally shoot for something higher than this as a problem point, possibly 600 which gives you just over 1/3rd additional for planning or damage control.