

(August 18th, 2014)

If you know someone who you think would benefit from being an Insider, feel free to forward this PDF to them so they can sign up [here](#).



Quick Tips for our Insider friends!

Hey Insiders!

This newsletter is coming to you from Redmond where we've been enjoying the excellent, better-than-average warm weather in the beautiful Pacific Northwest while working on new material and SQL Server 2014.

Note: There will be no newsletter on Monday, September 1st as we will be on vacation. The next newsletter will be Monday, September 15th.

Joe's latest Pluralsight course has just been published – *SQL Server: Common Query Tuning Problems and Solutions – Part 2* – check it out [here](#) and see below for a video excerpt.

Registrations are steadily coming in for the next of our popular SQLIntersection conferences, in Las Vegas this November – see [here](#) for details. Our show focuses on the troubleshooting and performance problems that we know you're facing today but will also highlight best practices in architecture, design, and SQL Server 2014. For a total of 6 full days – you can immerse yourself into great content with speakers that you know will deliver! Also, use the discount code “SQLskills” and you can save \$50 off registration.

The most recent book I've read is Kevin Mitnick's [Ghost in the Wires: My Adventures as the World's Most Wanted Hacker](#). I've been deeply interested in hacking since my earliest days with computers, and even almost got kicked out of college for hacking the CS department using some of the same tricks as Mitnick (read my post [We All Do Stupid Things When We're Young](#)). Mitnick's telling of his story is really well done, without coming across as annoyingly arrogant, and is full of great technical details. Although I was never into phone systems, I ran a student-run Unix cluster (called Tardis) and even crossed paths with a hacker while I was working on VMS at DEC who we assumed was Mitnick on the run so all his Unix and VMS info was great to read. Even more amazing are the social engineering skills he had and how easily people gave out privileged information. Highly recommended. If you're interested in this topic, also check out Stephen Levy's [Hackers](#), and Cliff Stoll's [The Cuckoo's Egg](#) among others.

Please [let us know](#) if you liked what you read/saw here and/or have any suggestions for future Quick Tips.

Note: you can get all the prior Insider newsletters [here](#).

Paul's Ponderings

Before I start, I had so much positive feedback on the last newsletter that I turned it into an expanded blog post – see [How to ask questions politely and correctly](#).

In this newsletter, I want to expound on a topic that came up on the Twitter #sqlhelp alias last week, where the questioner was concerned about how queries that use tempdb make use of multiple tempdb data files.

I'll paraphrase the question here:

If I have a tempdb with 8 data files, and I have a query that uses tempdb, will that query only use one of the tempdb data files?

The answer, of course, is It Depends! ☺ There are a number of scenarios to consider here, but before I begin, I'll explain simply how allocation works.

The first 8 pages that are allocated to an allocation unit (think of this simply as a portion of a table or index) are mixed pages from mixed extents, unless trace flag 1118 is enabled. After the magic 8-page threshold is reached, allocations are given to each allocation unit an extent at a time. If trace flag 1118 is enabled, allocations are immediately done in extent units, with no mixed pages being present.

See the following blog posts for more info:

- [Inside the Storage Engine: IAM pages, IAM chains, and allocation units](#)
- [Inside the Storage Engine: Anatomy of an extent](#)
- [Misconceptions around TF 1118](#)

Extents are allocated one at a time, from each file in the round-robin sequence – i.e. one from datafile1, one from datafile2, one from datafile3, one from datafile4, and then repeat – one from datafile1, etc. (Special cases apply to some bulk operations involving the –E startup parameter, or when the filegroup's proportional file is not uniform, but we'll ignore those as they rarely affect tempdb). The extent that is next allocated to an allocation unit depends on where the round-robin sequence is in the list of data files in the filegroup.

There's a lot more that goes on under the covers, and I've mentioned some of the complications, but that's enough to be able to answer the posed question.

Scenario 1: The query is single-threaded and nothing else is using tempdb space. In this case, if the query only allocates a few pages, it's likely that they will all be from a single data file, because they will be mixed pages all from the same mixed extent. It may be that there are only a few mixed pages available in the currently know mixed extent, so there may be some allocations from that extent, and some from another mixed extent, which may be in a second data file. If the query allocates more than 8 pages, the query will use space from multiple data files because of the round-robin process. A pathological case may occur where the query allocates less than 17

pages, with all 8 mixed pages being from the same extent in, say, datafile1, and the next file in the round-robin process is also datafile1, so the first full extent allocated for the query is from datafile1, in which case only a single file will be used.

Scenario 2: The query is single-threaded and other queries are using tempdb space. If other queries are also allocating tempdb space, then it depends how many allocations our query does:

- One data page: it will only use one data file.
- Less than 9 data pages: if it does all its allocations when no-one else is allocating then the mixed pages will be from one or two (at most) mixed extents, which will be in one or two data files. If the allocations are done while other threads are allocating too, each mixed page it allocates could be from a different mixed extent in a different data file, so the number of data files used will be less than or equal to the number of mixed pages used, in the worst case.
- More than 8 data pages: the more pages are allocated, the more extents are needed, and the more likely there will be one or more extents from each data file.

Scenario 3: The query is multi-threaded. In this case, the threads in the query are all calling into the allocation system for data pages (e.g. doing an index build using `SORT_IN_TEMPDB` in any version or doing a parallel `SELECT ... INTO` population of a large temp table in SQL Server 2014) and so it's almost guaranteed that all data files will be used.

Everything I've said above applies to all user databases as well.

With SQL Server 2014, there is a nice performance twist that pages that are used by a query and then deallocated by the same query should skip being written to disk under the covers (a process called eager writing). You can read about this on the CSS blog [here](#).

As you can see, a seemingly simple question has a complicated answer (and I didn't even get into proportional fill!) that really does justify beginning the answer with the classic It Depends!

Call to action: Not much of a call to action this time except to understand the relationship between queries, parallel threads, and multiple tempdb data files – cool stuff!

I'm curious to hear your thoughts on tempdb data file usage scenarios, so please feel free to [drop me a line](#), always treated confidentially, of course.

Video Demo

In this demo video excerpt from Joe's new Pluralsight course – [SQL Server: Common Query Tuning Problems and Solutions – Part 2](#) – Joe explains how forcing join order in a query using hints can have detrimental side effects.

The video is just over 5 minutes long and you can get it in MOV format [here](#).

You can get the demo code [here](#) and the Credit demo database [here](#).

Enjoy!

SQLskills Offerings

We've released all the classes for the remainder of 2014, including two in Australia in December. We expect to release the first half of our 2015 schedule around September.

Finally, to help your boss understand the importance of focused, technical training, we've added a few new items to help you justify spending your training dollars with us:

- [Letter to your boss explaining why SQLskills training is worthwhile](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

2014 Immersion Events

Chicago, IL

- October 6-8, 2014: **IE0**: Immersion Event for the Accidental/Junior DBA
- October 6-10, 2014: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)
- October 13-17, 2014: **IEPTO2**: Immersion Event on Performance Tuning and Optimization – Part 2 (formerly IE2)

Sydney, NSW, Australia

- December 8-12, 2014: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

Canberra, ACT, Australia

- December 15-19, 2014: **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1 (formerly IE1)

See [here](#) for the main Immersion Event Calendar page that allows you to drill through to each class for more details and registration links.

Fall SQLintersection

This year our Fall SQLintersection conference will be the week of November 10th in Las Vegas. See [here](#) for details. Don't forget to use the discount code "SQLskills" (without the quotes and it isn't case-sensitive) and you can save \$50 off registration. We hope to see you there!

Summary

We hope you've enjoyed this issue - we really enjoy putting these together.

If there is anything else you're interested in, we'd love to hear from you - [drop us a line](#).

Thanks,

Paul and Kimberly

Paul@SQLskills.com and Kimberly@SQLskills.com