

(March 18th, 2013)

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 Kimberly and I are decompressing from our wildlife tour in Japan and preparing for a plethora of classes, conferences, and crazy travel.

First up is our [SQLintersection Conference](#) in Las Vegas the week of April 8! There's a fabulous line-up of top speakers and sessions. Check out the blog post that Kimberly did on this performance-centric event [here](#), and one that Brent just posted [here](#).

Next up, for our Immersion Events in Chicago, we've added another IE1 (which is more than 50% full already) and we added another IEBI event after our first one sold-out. See the final section of this email for links/details. Most other classes are at 75% or higher capacity with IEDev, IE4 and IE2 about to sell out.

The most recent book I've read is Dorothy Dunnett's [King Hereafter](#). The book is a brilliant fictional account of King Macbeth's life as Earl of Orkney and King of Alba in the 11th Century. This was particularly interesting to me as it covers the time immediately prior to William the Conqueror's invasion and conquest of England in 1066. Dunnett has a more involved writing style than my other favorite, Sharon Kay Penman, so the books take more reading time. This is a good thing for me as I like spending longer with the characters and story line. I liked this so much that I got a bunch of Dunnett's other books to try out.

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

One of the questions I was asked in email this week concerned running multiple DBCC SHRINKFILE operations in parallel for a single database. The sender knew my stance on data file shrinking and was wondering about the possibility of causing corruption and/or long-running operations because of the shrinks.

There's never a possibility of causing corruption when running any of the shrink options – I haven't heard of a bug in that code since SQL Server 2000 (a very interesting scenario where shrinking to an exact GAM interval boundary caused an IAM page to stay erroneously allocated sometimes).

Running multiple shrink operations in parallel will work, but will take longer because of the parallel use of CPU, I/O, and logging. I was particularly concerned about the logging, as everything that shrink does is fully logged and having multiple shrink operations running at the same time would potentially generate a lot of transaction log very quickly.

Apart from slowing down operations on the log, by causing more log flushes, the sheer amount of transaction log being generated could cause problems with log backup size, log shipping, database mirroring, replication, and availability groups (all of which have to transport the log in some way to a remote system).

Apart from parallel shrink operations, there are some other algorithmic reasons why a shrink operation may run very slowly:

1. Moving any data file page in shrink requires acquiring an exclusive lock on the page. Shrink will wait forever for that lock to be acquired, which can obviously lead to a long-running operation.
2. Moving a text page requires a table scan. This is because there is nothing in a text page which indicates which data/index/text-tree records point to the text records on the text page being moved, so a scan is required to find the records that need to be updated with the new location of the moved text page. This has always been the case.
3. Moving a heap data page requires the query processor to update the nonclustered index records that link to the data records on the page being moved. This used to be done under the covers by the Storage Engine before SQL Server 2005, and was much more efficient – now that's not possible any more.

Issue #1 is easy to diagnose though – look in the output of `sys.dm_exec_requests` for the shrink task (the `command` column will say `DbccFileCompact`) and if it's `SUSPENDED` and the `wait_type` column is `LCK_M_X` then you can see who's blocking shrink in the `blocking_session_id` column.

For issues #2 and #3 you need to know what's in your database to be able to tell if these may be affecting the run time of a shrink operation. You can do this by running a simple query like the following:

```
SELECT
    [f].[name] AS N'Filegroup',
    [o].[name] AS N'Object',
    [p].[index_id],
    [a].[type_desc] AS N'Type',
    [a].[total_pages] AS N'Pages'
FROM sys.system_internals_allocation_units [a]
JOIN sys.partitions [p]
    ON [p].[partition_id] = [a].[container_id]
JOIN sys.filegroups [f]
    ON [a].[filegroup_id] = [f].[data_space_id]
JOIN sys.objects [o]
```

```
        ON [o].[object_id] = [p].[object_id]
WHERE
    ([p].[index_id] = 0
    OR [a].[type_desc] = N'LOB_DATA'
    OR [a].[type_desc] = N'ROW_OVERFLOW_DATA')
AND [a].[total_pages] > 0;
GO
```

This will show you how many heap data pages and text pages you have in your database, and in which filegroups.

Call to action: Apart from the obvious consideration of whether you can perform the effect of a shrink operation in a more efficient way (see [this blog post](#)), if you're going to have to run a shrink operation then be prepared for it to take a long time. Using the methods I've described above, you might be able to limit the impact of the shrink operation on regular database operations by working out in advance that it's going to take a long time and scheduling the operation in a period of light work.

I'm curious to hear your thoughts about shrink performance (we all know that shrink should be avoided if possible), so please feel free to [drop me a line](#), treated confidentially of course.

Video Demo

From Erin:

Profiler has helped DBAs find problem queries since SQL Server 7. But with its deprecation in the SQL Server 2012 release, DBAs will soon be required to use a new method to trace database activity – and that method is Extended Events. Available since SQL Server 2008, Extended Events is a lighter, better, and more powerful way for DBAs to troubleshoot and diagnose problems. However, making the transition from the tried and true world of Profiler and SQL Trace requires a bit of re-thinking. The terminology may be different, but there are a lot of similarities, as we'll see in this demo.

First we'll take a pre-existing trace in a SQL 2012 instance and generate the syntax to create an Extended Events session using the script from Jonathan's post, [Converting SQL Trace to Extended Events in SQL Server 2012](#). Then we'll run a SQL Trace and Extended Events session side-by-side while running Jonathan's [AdventureWorks BOL Workload script](#) to generate a workload. Finally, we'll end with a quick review of the output from the Extended Events session, showing that what you used to capture in Profiler is still there in Extended Events. Once you understand how similar the two are, you can then review Jonathan's [Introduction to Extended Events](#) course on Pluralsight, or if you'll be in Chicago for SQLSaturday #211 on April 13, check out my session, [Making the Leap From Profiler to Extended Events](#). It's time to start exploring and understanding Extended Events. Change is never easy and sometimes changes are uncomfortable, but trust me, it's worth it.

The video is about 13 minutes long and you can get it:

- In WMV format [here](#)
- In MOV format [here](#)

You can get a zip of the demo code [here](#).

Enjoy!

SQLskills Offerings

All of our 2013 public classes are filling up fast! Based on requests from people, attendee ratings of the hotels we used last year, and the ease of using hotels we know, we're using the same locations again. This means we cover both sides of the US, central US, and Europe.

Please know that these classes are final as the hotel contracts are signed, and the classes will not be cancelled or moved for any reason, nor will the dates change. We are not planning any other locations for 2013 at this time.

- April 29-May 3, 2013: Internals and Performance (**IE1**) in Chicago, IL **SOLD OUT!**
- April 29-May 3, 2013: Immersion Event for BI (**IEBI**) in Chicago, IL **SOLD OUT!**
- May 6-10, 2013: Performance Tuning (**IE2**) in Chicago, IL – USA
- **NEW** May 6-10, 2013: Immersion Event for BI (**IEBI**) in Chicago, IL – USA (co-located but in a different training room. Attendance is for one event or the other; these cannot be combined for one attendee where they move back/forth.)
- May 13-17, 2013: High Availability & Disaster Recovery (**IE3**) in Chicago, IL – USA
- May 13-17, 2013: Immersion Event for Developers (**IEDev**) in Chicago, IL – USA (co-located but in a different training room. Attendance is for one event or the other; these cannot be combined for one attendee where they move back/forth.)
- **NEW** May 20-24, 2013: Internals and Performance (**IE1**) in Chicago, IL – USA
- May 20-24, 2013: Development Support (**IE4**) in Chicago, IL – USA (co-located but in a different training room. Attendance is for one event or the other; these cannot be combined for one attendee where they move back/forth.)
- June 3-7, 2013: Internals and Performance (**IE1**) in London – UK
- June 10-14, 2013: Performance Tuning (**IE2**) in London – UK
- June 17-21, 2013: High Availability & Disaster Recovery (**IE3**) in London – UK
- June 24-28, 2013: Development Support (**IE4**) in London – UK
- September 16-20, 2013: Internals and Performance (**IE1**) in Bellevue, WA – USA
- September 23-27, 2013: Performance Tuning (**IE2**) in Bellevue, WA – USA

One thing to note is that the course prices have increased slightly for 2013, reflecting increasing food, logistics, travel, and accommodation costs. We kept our prices the same for the last three years but now we have to raise them a little.

For US classes, the new early-bird price is US\$3,295 and the full-price is US\$3,795. However, for all past attendees in the 12 months prior to registration, we will only charge the 2012 early bird price of US\$2,995. As an alumnus, send us an email prior to registering and we'll give you a code to use to access this special rate. **And be sure to get your registrations in early!**

For UK classes, the new early-bird price is US\$3,795 and the full-price is US\$4,295. There is a similar past-attendee price equal to the 2012 UK early bird price of US\$3,495. **Again, be sure to get your registrations in early!**

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

So, that's it for now. We hope to see you soon!

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