

(November 17<sup>th</sup>, 2020)

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



Note: As an Insider, you can read all prior Insider newsletters [here](#).

### **Quick Tips for our Insider friends!**

This newsletter is coming to you from Redmond, where we're working on our schedule for early 2021 and preparing to release some exciting new online offerings in December!

**Stay safe out there!**

### **SQLskills News**

**Live, ONLINE classes in 2021:** our classes next year will all be live-streamed online, as it's not feasible to run in-person classes for the foreseeable future. Just like that last two months, we'll be running them as a series of half-day sessions so it's less onerous than the usual 9-hour days in our in-person classes. Attendees will also get streaming access to a recording of the entire class – something we'd never done before this year!

We're putting together the schedule at the moment and we'll publicize it soon.

**Kandio job candidate assessments:** we've teamed up with Kandio to produce technical assessments to help companies screen candidates for job recruitment. If you want to make sure someone really knows what they say they know, check out the assessments [here](#).

### **Book Review**

From the 2013 archives, I present Allie Brosh's [\*Hyperbole and a Half: Unfortunate Situations, Flawed Coping Mechanisms, Mayhem, and Other Things That Happened\*](#). If you're in need of a laugh this year, get this book – it's just wonderful. She also just released her second book, which will arrive for me on Wednesday. If you need more laughs, also check out [The Oatmeal](#) and [XKCD](#) (and don't blame me when you realize you spent a whole afternoon clicking through their websites!). Highly recommended!

### **The Curious Case of...**

*This section of the newsletter explains recent problems we've helped with on client systems; they might be something you're experiencing too.*

Last week I had a discussion with someone about whether they should cancel a long-running transaction or not... You can read through my explanation [here](#)...

## **Ponderings...**

*(From Erin this time – enjoy!)*

Hi friends, this week I want to talk about data, and its importance when troubleshooting an issue in SQL Server. With each release, we have more and more information available to us within SQL Server, which is great, but it can also be overwhelming! I've talked before about how we at SQLskills always start with wait statistics when troubleshooting, but today I want to share a story from a recent customer experience where wait statistics weren't something we could use.

Our customer reached out after multiple SQL Server instances failed over unexpectedly, within a series of hours, in the same evening. They were obviously concerned it could occur again, so they wanted to understand why it happened to then address the root problem. When I joined the call there were numerous individuals from multiple teams on the call. Sometimes, these critical calls with lots of people can be tricky. But with this group it was very organized and methodical – each team was talking through what data they had reviewed thus far. There was no finger-pointing, no arguing, just a calm and collected approach to working through the problem; I so appreciated it.

Sidebar: I recognize that when something “bad” happens or there is a performance issue, there's a ton of pressure, from multiple directions, to fix it...NOW. *I get it.* But for the hundreds of calls that I've been on with that type of situation, the steady and composed method of troubleshooting consistently works very well.

Because there was a failover, my first request was to review the error log files, the cluster logs, and all .xel files (system\_health, AlwaysOn, \*SQLDiag\*). In looking in the error log, it was obvious the shutdown was not controlled, and there was nothing in any of the .xel files that pointed to a culprit. The cluster logs showed a lease failure. At that point we went back to the error logs, as the DBA team had originally mentioned they noticed long I/O requests. Indeed, there were many occurrences of the message:

*SQL Server has encountered X occurrence(s) of I/O requests taking longer than 15 seconds to complete on file Y*

where X ranged from one to several hundred, and they occurred across multiple data files on different drives. Going back further in the logs, those messages did not exist with near the same frequency, if at all. We had very little data that demonstrated any problem with SQL Server. This happens. Sometimes, the problem is not **WITH** SQL Server; SQL Server just ends up being the victim. I said as much to the team, and the different groups continued working through

the other data they had – storage, network – and eventually it was determined that the problem related to a bad switch.

It's easy to get inundated with information, and to feel overwhelmed by everything you might want to try and review. Understanding what data is important, as it relates to a specific problem, is something you learn over time when working with SQL Server. You also have to learn what that data is trying to tell you. In this case, the I/O requests told me that SQL Server wasn't getting data returned to it quickly. But I had nothing else that I could pair with it to say what *part* of the storage fabric was having an issue (remember that when we talk about storage, especially SANs, it's not just a direct path from the SQL Server to the disks...there are NICs, HBAs, switches and more that data has to go through). I needed people from other teams to look at *their data* to pin down the root issue.

Therefore, my takeaway for you is to collect as much relevant data as you can, take time in your analysis, and be clear about what findings you can, and cannot, make based on that information. Happy sleuthing!

## **#TBT**

*(Turn Back Time...) Blog posts we've published since the previous newsletter plus some older resources we've referred to recently that you may find useful.*

The theme for the TBT this time is database mirroring. Although database mirroring has been deprecated, it's still in wide use and it's still a very easy way to perform a rolling upgrade or move data around. Here are some useful links:

- Pluralsight course: [SQL Server: Understanding, Configuring and Troubleshooting Database Mirroring](#)
- Whitepapers:
  - [Database Mirroring and Log Shipping Part I – Prescriptive Guidance](#)
  - [Prior Clustering and Mirroring Deployments Part I – Prescriptive Guidance](#)
  - [Proven SQL Server Architectures for High Availability and Disaster Recovery](#)  
(written by me)
  - [SQL Server Replication: Providing High-Availability Using Database Mirroring](#)  
(written by me)
  - [Database Mirroring and Log Shipping Working Together](#)
  - [Database Mirroring: Alerting on Database Mirroring Events](#)
  - [Database Mirroring: Best Practices and Performance Considerations](#)
- Doc article [Prerequisites, Restrictions, and Recommendations for Database Mirroring](#)
- My blog posts:
  - [Importance of network latency when using database mirroring](#)
  - [Importance of monitoring a database mirroring session](#)

- [Database mirroring blog category](#)

I hope you find these useful and interesting!

### **Video Demo**

Query Store captures a large number of queries, so how do you find one that's of interest to you without scrolling through rows and rows of output? In this week's Insider video, Erin will show you how to find a query based on the query's text – whether it's an ad hoc query or part of a stored procedure. Tune in to learn how!

The video is about 7 minutes long and you can get it in WMV format [here](#).

The demo code is [here](#).

Enjoy!

### **Upcoming SQLskills Events**

Our 2021 classes will be entirely online in the Spring and we'll announce them soon!

You have multiple learning opportunities as every event has a different focus as well as different benefits – from deep-technical training in our Immersion Events to a more broad set of topics at conferences! And, of course, one benefit you'll always get from in-person events is networking; we hope to meet/see you at an event soon!

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

- [Letter to your boss explaining why SQLskills training is worthwhile](#)
- [So why do you want to come to our training? And the winners are...](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

### **LIVE, Online Immersion Events:**

Spring 2021

- We'll be announcing our Spring 2021 online class schedule soon!

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

### **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](mailto:Paul@SQLskills.com) and [Kimberly@SQLskills.com](mailto:Kimberly@SQLskills.com)