

(March 27th, 2017)

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](#).



Quick Tips for our Insider friends!

Hey Insiders,

We heard great feedback about the new #TBT section (after the demo video section) – make sure to check it out! If you have any suggestions for topics to cover in any portion of the newsletter, [drop us a line](#) and we'll see what we can do.

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

SQLskills News

We've started a new initiative where we're all blogging about introductory topics, to help the burgeoning numbers of non-DBAs or junior DBAs who are responsible for SQL Server instances. We're calling it SQL101, and hopefully it'll even help refresh or reinforce topics for the more seasoned DBAs in the community. The blog posts will be automatically collected [here](#).

We're looking forward to seeing a bunch of you this year in our classes in Chicago in April/May and Bellevue in July/August. Note that several classes are sold out and all other classes are 2/3 full. **Also, please note that our May delivery of IEHADR will be the last.** Despite the importance of having a solid disaster recovery plan, many don't plan at all and very few put the resources behind having a well-designed and thoroughly-tested strategy. As a result, we've decided to convert this content into online training instead of in-person training. But, we're still looking at areas where there is more interest; we're particularly pleased to have a brand new [three-day class on using PowerShell](#) to administer SQL Server, taught by industry-expert, MCM, and MVP, Ben Miller. See [here](#) for the complete 2017 class schedule.

Our Spring SQLIntersection conference is also fast approaching at the end of May, and we have a phenomenal line up of workshops, sessions, and speakers (including Brent Ozar, Buck Woody, and Bob Ward!). Check out [this blog post](#) for all the details, and use the discount code 'sqlskills' when you register to save \$50.

Even though we're not teaching any Immersion Events in Europe this year, Kimberly and I will both be [presenting at SQLSaturday #620 in Dublin in June](#).

Finally, even if you can't join us in person, we're still taking requests for remote sessions for this year. We have 57 scheduled so far; if you'd like one of us to present for your user group, check out my blog post [here](#).

Book Review

The latest book I've read is Lawrence and Nancy Goldstone's [*Out of the Flames: The Remarkable Story of a Fearless Scholar, a Fatal Heresy, and One of the Rarest Books in the World*](#). This book describes the history of one of the rarest books in the world (only three copies remain): the *Christianismi Restitutio* by Michael Servetus, published in 1553, refuted orthodox Christianity and led to Servetus' persecution by John Calvin and eventual death by burning at the stake. The book covers the history of everyone related to the book's publication, suppression, and survival to the modern day, along with lots of interesting info on the Protestant Reformation. Highly recommended for all bibliophiles.

The Curious Case of...

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

Tim was working with a new client recently who had approached us for help with a migration to Azure SQL Database. The client presented us with their architectural design documents which indicated their primary user database would be running on an S3 database which would provide 100 DTUs at a cost of US\$150 per month.

When we asked how the client derived the 100 DTU value, they didn't have a good justification. We had them run the [DTU Calculator](#) and the results came back stating they needed a P6 database which is a premium tier 1000 DTU database at a cost of US\$3,720 per month – a slight disparity from their expectations!

The client was understandably very shocked at this discrepancy in cost.

In looking at the analysis of the DTU calculator, we could see which portion of their workload was driving the need for so much compute power. After discussing with the client, we engaged in a performance tuning project on their environment.

We found numerous stored procedures and processes that needed tuning or tweaking. Once we were done, we had them run the DTU Calculator again and the results came back recommending a P2 at 250 DTUs and a cost of \$930 per month. While this was still a much higher cost than the client originally thought they would be paying with an S3, they were very happy to have their workload running on a P2.

With a short performance tuning engagement, the client was able to reduce their Azure cost by US\$33k annually by having a more optimized database.

Bottom line: Make sure you run the DTU Calculator when planning a move to Azure SQL Database, and if you think the required performance seems too high, do some performance tuning to streamline your workload. A little bit of time invested in this can translate into a large amount of money saved!

Paul's Ponderings

My favorite user group session to present is *Performance Troubleshooting Using Wait Statistics* (I've done it three times already this year), and at the start of the deck I explain how using wait statistics for problem diagnosis is a lot easier if you have a performance baseline.

In fact, doing *any* kind of performance troubleshooting is far easier when you have a baseline so you can tell how the current performance is differing from the baseline – measurements can be taken and compared against the baseline, to identify areas that have changed. This allows you to target further investigations rather than doing what I call 'flailing' or 'knee-jerk' performance tuning, where you have no idea where to start or you pick a random symptom and focus on that, wasting precious time.

A baseline is a set of measurements of performance metrics when the overall system is performing normally (i.e. satisfactorily for users). Some examples of such metrics are:

- The I/O load during regular and peak operations, from the SQL Server and Windows perspectives
- Buffer pool health (e.g. Page Life Expectancy and Lazy Writer activity)
- Run-time characteristics of important queries
- Wait statistics

One thing we see quite commonly is that a new client has no way to accurately measure how long an operation takes. How many of you have heard things like 'the web site seems slower today' or 'my query seems to be taking longer sometimes'? These kind of subjective measures are really hard to work with.

Just like having a baseline, it's very important that operations in your environment can be measured accurately so you can determine where a slowdown is occurring. For instance, if you have a web site front-end, connecting to a mid-tier caching server, connecting to a back-end database server and things start to run slowly, how do you ascertain which of the components in the overall data flow is contributing to the slow down?

Only by having performance and timing metrics at each stage in the data flow can you do that, and only then by comparing against your baseline.

Call to action: Think about how you measure performance on your systems, and how you work out where performance problems are when they occur. Do you have a baseline to which you can compare current numbers? If not, I highly recommend that you work towards getting one as it'll make your life a lot easier. In this week's #TBT section, I'll point you at a bunch of resources around that.

Video Demo

In this demo video, Jonathan shows how to bypass trigger execution without disabling a trigger using the `CONTEXT_INFO` for a session in SQL Server. This is really useful to know if the trigger has to remain enabled to enforce business logic but you need to bypass it to manually fix some erroneous data.

The video is about 6 minutes long and you can get it:

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

And the demo code is [here](#).

Enjoy!

#TBT

(Turn Back Time...) This section of the newsletter highlights some older resources we've referred to recently that you may find useful, plus select blog posts we've published since the previous newsletter.

Continuing with the performance baseline theme from my editorial, here are some resources to help you produce a baseline:

- Erin's Pluralsight course on [SQL Server: Benchmarking and Baselineing](#) that has 3.5 hours covering what metrics to capture and how to capture them.
- If you're still using Trace/Profiler, Erin's Pluralsight course on [SQL Server: Replacing Profiler with Extended Events](#) will be really beneficial.
- And a bunch of blog posts:
 - Erin: [SQL Server Baselines Series on SQLServerCentral.com](#)
 - Erin (from our Accidental DBA series): [\(Day 18 of 30\): Baselines](#)
 - Erin: [\(Day 19 of 30\): Tools for On-Going Monitoring](#)
 - Jon: [\(Day 21 of 30\): Essential PerfMon counters](#)
 - Paul: [Capturing IO latencies for a period of time](#)
 - Paul: [Capturing wait statistics for a period of time](#)

Here are a few of the blog posts we've published since the last newsletter:

- TBD in Sorong

I hope you find these useful and interesting!

Upcoming Immersion Events

The classes for both Spring and Summer 2017 are available for registration!

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](#)
- [Community blog posts about our classes](#)
- [Immersion Event FAQ](#)

Chicago, IL, April/May 2017

- **IE0:** Immersion Event for Junior/Accidental DBAs
 - April 24-26 ****SOLD OUT****
- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - April 24-28 ****SOLD OUT****
- **IESSIS1:** Immersion Event on Learning SQL Server Integration Services
 - April 24-28 ****SOLD OUT****
- **IEBI:** Immersion Event on Business Intelligence
 - May 1-5
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - May 1-5
- **IEPS:** Immersion Event on Powershell
 - May 8-10 ****NEW** class**
- **IEPDS:** Immersion Event on Practical Data Science
 - May 8-12
- **IEHADR:** Immersion Event on High Availability and Disaster Recovery
 - May 8-12

Bellevue, WA, July/August 2017

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
 - July 31-August 4
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
 - August 7-11

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 and Kimberly@SQLskills.com