

(April 8<sup>th</sup>, 2019)

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!

Hey Insiders,

This newsletter is coming to you from the Revillagigedo Archipelago of Mexico, where we're taking a short dive trip to the islands there including Socorro Island. These islands lie about 400 kilometers off the southwest coast of the Baja peninsula where there are lots of sharks (we might see as many as 10 species!), mantas (these are the pelagic *manta birostris* that reach a "wing span" size of 20+ feet), and lots of other pelagic sea life. And, it's humpback whale season! It's definitely going to be a packed week of big critters!

**Note: our April/May classes are almost or actually sold out – don't delay and be disappointed!**

## SQLskills News

You can find all the details of our upcoming classes (pricing, logistics, registration, curricula) by drilling down from the class schedule page [here](#), and there's a list of classes at the bottom of the newsletter.

If you're looking for a great place to learn best practices in migration, troubleshooting, and performance, join us in Florida in June for our [SQLintersection](#) conference where we're going to focus on this incredibly important process. SQLintersection will not only help you migrate to a newer version of SQL Server but it will teach you the best new features to leverage once you're there! We've got an incredible line-up this Spring and ALL of our sessions have been posted; check out our website to view our spectacular line-up including Bob Ward, Pedro Lopes, Pam Lahoud, Brent Ozar, and most of our SQLskills team! We hope you can join us in June! Check out my blog post [here](#).

**Live, IN-PERSON classes:** we're back in Chicago in April/May 2019 and our classes are open for registration, including a new class on Power BI taught by Tim. This is the only time during the year where we offer so many options and combinations for in-depth SQL immersion – all in a central location (Chicago) and a great hotel (every room is a suite with a fridge, microwave, and sitting area separate from the bedroom). We are not planning any events in Europe this year due to other commitments, conferences, and online training. **NOTE: Book ASAP as many classes are close to selling out or have already sold out!**

**Live, ONLINE classes:** we've finished our first set of live, online classes and we'll announce the Fall line-up at the start of the summer.

**Finally, even if you can't join us in person,** I've put out a call for **2019 remote user group sessions** and we've got 32 scheduled this year already! If you'd like one of us to present for your user group, check out my blog post [here](#).

### **Book Review**

The most recent book I've read is Michael Dirda's [Browsings: A Year of Reading, Collecting, and Living with Books](#). I'm a bibliophile, as you all know, and so is Dirda, who's the book columnist for the Washington Post and an author in his own right. This book is a collection of his weekly essays for a year for a literary column in The American Scholar journal. He covers a wider range of literary topics in an erudite and light-hearted way, including favorite books (some great annotated lists), writing, forgotten authors, conventions, pets, and lots more. Great stuff and 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.*

Jonathan was working with a client last week who had a CLR assembly failure after an AG failover. He figured it out but it's a really interesting problem...

Read more about it in my blog post [here](#).

### **Ponderings...**

*(This time from Rafal Lukawiecki, our fantastic Machine Learning instructor...)*

ML (Machine Learning) Services in SQL Server 2017, and recently added to Azure SQL Database, are underappreciated by data scientists, who are not aware of just how fast the new T-SQL *PREDICT* statement is, also known as *native scoring*. It is near real-time (<20ms), written as a native C++ extension to SQL. If you want even more, there is a *real-time* scoring variant, a.k.a. *sp\_rxPredict*, a CLR extension (<10ms), but unlike *PREDICT* it is not yet in Azure SQL Database.

Above, < really means less—I have seen 2-5ms. No need to move data out of the database either! Nothing compares to *PREDICT* in plain R or Python, and you really need to roll-out the mighty (and in my opinion overly complex) set of Spark/Hadoop clusters to get that performance for the majority of business-focused predictions (scoring). Sure, there are exceptions, notably anything to do with image analysis, when the big cluster guns of deep learning-focused frameworks work faster, especially on dedicated hardware—but this is not what the majority of my customers, and

to the best of my knowledge, majority of today's users of ML require. Add to that the simplicity of deployment, not having to worry about package/dependency management, and you can quickly grasp why this SQL Server/Azure SQL Database secret deserves more light.

How do you get to use this fabulous *PREDICT* statement? To start with, you need to build a predictive model using Microsoft RevoScaleR (in R) or revoscalepy (in Python) library. To do that, you do not want to be using *sp\_execute\_external\_script* despite many Microsoft demos and conference sessions suggesting that—tedious, hard to debug, just wrong workflow. Instead, please set-up a nice development environment using a standalone ML Server (like in the VM you will get if you attend the IEPML12 class in Chicago in May. The installer comes with the SQL Server ISO but it is better to download the newer version, 9.3. You can install it on its own, but it is usually better to have it on the same machine as a (dev) SQL Server. Then you need to select an IDE, such as RStudio (free, easy, and recommended) or R Tools for Visual Studio 2017 (deprecated in VS 2019), or any Python environment you like. Ideally, use an RMarkdown notebook for easy reproducibility of your workflow.

Once you understand the business need, and you have expressed it as a testable *hypothesis* that can be turned into a predictive machine learning model, and which—very important!—you have tested to be accurate (easy to test) and reliable (much harder), you deploy the model, in a serialized form, by storing it in something like a varbinary(max). You won't need the full 2GB this allows. Most of my models are merely 10–50kB in size. It is that model that you later retrieve from your table and that you feed into the *PREDICT* statement in either an *INSERT* or a *SELECT*, at prediction time.

Simple, but the key is to understand how to *build* and *test* the model, and, above all, how to prepare your data, because your data matters *much* more than your choice of the modelling algorithm or the ML framework. Needless to say, this is all the subject of my IEPML12 super-intensive training. It takes someone who knows how to work with data and brings them up to the level of being able to build machine learning models, fulfilling the role of a practical data scientist or an AI engineer.

On that last subject...AI and AI engineering are poised to become the next most-in-demand job titles (move over data scientists, you are so '10s). More seriously, I think that everyone who works with data, DBAs and data-focused developers, are uniquely qualified for this role, because it is quite easy to pick up the ML skills but much harder to learn how to deal with complexities of organizational data. I often need more than one data expert to each data scientist I see on projects, and with AI this is only going to get busier.

Oh, while you are at it, please, do not build Artificial Stupidity. Sadly, too many organizations are doing just that today (like a highly desirable autonomous car that otherwise mysteriously drives itself into a parked fire truck on the side of the road). Interested? Have a look at my two,

free videos, that introduce the concept of AI and Artificial Stupidity from a very practical perspective – hope you find it useful, or at least amusing – see [here](#).

**Call to action:** Consider if your organization is asking for predictive functionality at all, like predicting customer churn, finding VIP customers, calculating profit potential, product failures, capacity or sales planning, recommendation engines, outlier or fraud detection, next best action, and so on. If the answer is yes, make sure they consider the value of SQL Server/Azure SQL Database for this need. Also, make sure they understand where this journey is headed, especially if autonomous decision-making, aka AI, is part of the strategy, to make sure that you avoid Artificial Stupidity—which, I am afraid, is cheaper and easier to deliver than good AI. See what I have to say online about AI and the future [here](#) and join me at IEPML12 in Chicago May 6-10.

### **Glenn's Tech Insights**

*Recent news and views from the hardware and Windows worlds that we think will be interesting to SQL Server community members.*

In Glenn's regular roundup of tech news this time he talks about the following:

- Intel's Data Centric Innovation Day
- Computex 2019 keynote speech by AMD's CEO

Check it all out on his blog [here](#).

### **#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.*

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

- Paul: [SQLintersection Spring 2019](#)
- Glenn: [Intel Cascade Lake-SP Processor Analysis for SQL Server](#)

Instead of a TBT theme this time, I want to tell you about the new locations of the various SQL Server team blogs that moved recently.

SQL Server: <https://techcommunity.microsoft.com/t5/SQL-Server/bg-p/SQLServer>

Azure SQL DB: <https://techcommunity.microsoft.com/t5/Azure-SQL-Database/bg-p/Azure-SQL-Database>

The SQL Server blog is now a combined blog that has labels for the different channels. If someone wants to follow releases for example, they can use this URL: <https://techcommunity.microsoft.com/t5/SQL-Server/bg-p/SQLServer/label-name/SQLReleases>

I hope you find these useful and interesting!

### **Video Demo**

In this week's Insider Video Glenn discusses the new 2nd Generation Intel Xeon Scalable Processor family and how it compares to the previous Intel Xeon processor family from a SQL Server perspective.

The video is just over 9 minutes long and you can get it [here](#).

A PDF of Glenn's deck is [here](#).

Enjoy!

### **Upcoming SQLskills Events**

Our set of 2019 live, in-person events for Chicago in April/May are coming up quickly – many have already filled – and, we've just wrapped up our Q1 live, online classes.

You still 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 wide-ranging topics at SQLintersection (see [here](#) for details of the June show) where you can learn more effectively how to keep moving forward in both your environment and your career! 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, IN-PERSON Immersion Events:**

Chicago, IL, April/May 2019

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1

- April 29-May 3 **\*\* SOLD OUT \*\***
- **IECAG:** Immersion Event on Clustering and Availability Groups
  - April 29-30 **\*\* 4 seats remaining \*\***
- **IEPowerBI:** Immersion Event on Power BI, Power BI Report Server, and SSRS
  - April 29-30 **\*\* NEW \*\***
- **IE0:** Immersion Event for the Junior/Accidental DBA
  - May 1-3 **\*\* 3 seats remaining \*\***
- **IEUpgrade:** Immersion Event on Upgrading SQL Server
  - May 1-3
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - May 6-10 **\*\* 3 seats remaining \*\***
- **IEPML:** Immersion Event on Practical Machine Learning
  - May 6-10 **\*\* 4 seats remaining \*\***
- **IEAzure:** Immersion Event on Azure SQL Database, Azure VMs, and Azure Managed Instance
  - May 6-9 **\*\* SOLD OUT \*\***

#### **LIVE, ONLINE Immersion Events:**

- The Fall class lineup will be announced at the start of the summer

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)