**(August 29th, 2016)**


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,

Welcome back! We hope you've had a great summer and have had a chance to do some fun things with your family and friends. We focused our summer vacation on Norway, and it did not disappoint! After driving the coast of Norway, we headed up to Svalbard to cruise around the archipelago and get some time up in the Arctic pack ice. Our Svalbard adventures were joined by some techie friends and their families; we had an absolute blast! Check out our photography Facebook page for a bunch of Kimberly's fantastic photos of polar bears and scenery. You can also follow Kimberly on Instagram here.

Now we're in Redmond working on content and getting ready for our Bellevue classes in a couple of weeks (still a few spots in those and our Chicago Fall classes – all the schedule details are here).

We've published two new Pluralsight courses since the last newsletter:

- Tim's first course *SQL Server: Consolidation Tactics and Best Practices* – check it out here.
- Glenn's latest course *SQL Server: Improving Storage Subsystem Performance* – check it out here.

SQLintersection session and workshop details have been posted; we have an exciting line-up scheduled for our show! If your team needs architectural advice and sessions from speakers who not only know their technology but know how to convey it – this is the place to be! Be sure to use the 'SQLskills' discount code to save $50 on registration. Check it out here.

We've decided to make our special offer **single-instance health check for only US$2,500 the permanent price for the first health check for new clients**, which is more than 1/3 off the original price! Details about why our health checks are so cost-effective are here.

The latest book I've read is Peter Frankopan's *The Silk Roads: A New History of the World*. This is a very interesting book tracing the varied history of the countries along the Silk Road, including empires, explorers, and religions that affected the various routes. A lot of what's in the book I already knew, but having it all presented in one volume in a chronological sequence was excellent. The last 100 pages or so detailed the quite despicable British and American machinations around the countries in the Middle East for their own (mostly oil-related) gains, to

the huge detriment of the native populations, which I felt quite ashamed to read about, being of both nationalities. Highly recommended for history fans!

Note: you can get all the prior Insider newsletters here.

## 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.*

A quick one this time: Glenn was working with one of our "remote DBA" clients recently who encountered weird errors while trying to create a view using the SSMS view designer – they basically couldn't create the view they wanted.

Glenn noticed schema-modification (SCH-M) locks on some of the tables referenced in the view that was trying to be created. After further investigation, he found that the root cause was another person at the client was editing an existing view in the view designer that referenced the same tables – hence the schema-modification locks.

If someone is going to use a tool that holds long-term blocking locks (like the SSMS view designer), it's essential to coordinate with other team members to prevent this kind of problem occurring.

## Paul's Ponderings

One topic that comes up in class sometimes is using an int identity as a clustering key and the possibility of running out of integers. Depending on your insert volume, this might be quite likely, as an int can only store 2^32 (^ = 'to the power') or about 4 billion values, between -2^31 and 2^31-1.

Imagine that you have a theoretical system that can create a thousand data rows per second. Using an int identity value increasing by 1 and starting at 1, you'll run out of values when the value hits 2^31-1 and tries to insert the next value. Let's simplify the math by just saying that 2^31 is the limit. With a thousand values per second, that would mean 2^31 / 1,000 seconds = 2.15 million seconds or approximately 25 days. While many of you don't sustain 1,000 rows per second, this is still a very problematic limitation.

One solution is to use a bigint identity as the key. This can store 2^64 or about 18.5 quintillion (18.5 billion billion) values, between -2^63 and 2^63-1.

Every so often someone asks whether it's possible to run out of bigint values. My answer is no. Well, technically yes, there is a limit, but in practical terms the answer is no.

Now imagine that you have a theoretical system that can create a million data rows per second, with a bigint identity value increasing by 1 and starting at 1. You'll run out of values when the

value hits 2^63-1 and tries to insert the next value. With a million values per second, that would mean 2^63 / 10^6 seconds = 9.2 trillion seconds or approximately 292.5 thousand years. And by then it's someone else's problem ☺.

Now what about the storage for those values? Doing a quick test of a heap with a single bigint identity column shows me that I can get 453 rows per 8KB data file page (don't forget the record overhead, slot array overhead, and that the heap pages won't be filled completely because of the way the free space caching and searching works). A terabyte of data would store roughly 61 billion rows.

So with 1 million rows per second, you'll be generating 1 million x 3,600 x 24 = 86.4 billion rows per day, so you'll need about 1.4 terabytes of new storage per day. If the values never get removed, so there's always new space required, you'll need almost exactly 0.5 petabytes of new storage every year.

At that rate, actually running out of bigints AND storing them would take roughly 150 thousand petabytes. This is clearly impractical – especially when you consider that simply storing a bigint is pretty pointless – you'd be storing a bigint and some other data too – probably doubling the storage necessary, at least.

Why is this interesting as a newsletter topic? We've had several clients over the years that didn't consider their data volume and designed a schema using int keys instead of bigint keys. When the inevitable happened and they ran out of ints, the process of changing to a bigint key was quite painful – as there's no really easy, space and log efficient way to do it once you have the 2 billion rows, and especially if constraints are involved, and application changes need to be made to allow the 8-byte instead of 4-byte value in result sets.

A common solution is to just reset the seed to -1 and then decrement by -1. As a short-term solution this is reasonable but it's not ideal as a long term solution as you'll only run out again. Ultimately, you'll need to make the int to bigint change.

**Call to action:** make sure that when you're designing a new schema, you think through the maximum values required and pick appropriate data types then and there. Changing data types can be very painful once the system has been in production for a while and there's a lot of data in the schema.

## Video Demo

The demo video this time is from another course we published in July, Erin's *SQL Server: Replacing Profiler with Extended Events*. In the demo she shows how to use the histogram target to track database file growth.

The video is just over 7 minutes long and you can get it:

- In MOV format here.

The demo code is available here.

Enjoy!

## SQLskills Offerings

Our 2016 classes are all open for registration (listed below), including three new classes in Chicago in November added because of popularity. We hope to see you at a class or a conference this year!

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

## Upcoming Immersion Events

Bellevue, WA

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - September 12-16
- **IEPTO2:** Immersion Event on Performance Tuning and Optimization – Part 2
  - September 19-23
- **IESSIS2:** Immersion Event on Advanced SQL Server Integration Services
  - September 19-22 **New course!!**

Dublin, Ireland (*returning to Europe in 2018* ***NOT 2017***)

- **IEPTO1:** Immersion Event on Performance Tuning and Optimization – Part 1
  - October 3-7

Chicago, IL

- **IE0**: Immersion Event for Junior/Accidental DBAs
  - November 7-9 **Just added!!**
- **IEPTO1**: Immersion Event on Performance Tuning and Optimization – Part 1
  - November 7-11 **Just added!!**
- **IEPDS**: Immersion Event on Practical Data Science
  - November 7-11 **Just added!!**

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