Database Corruption Challenge #7 Coming Soon

Those who are subscribed to my newsletter will receive the exact time of the next corruption challenge. For those who have not subscribed, just keep checking back to find out when the next Database Corruption Challenge begins.

Here is the latest update on how the challenge has gone so far:

Week 1 was just a wild idea that I had to start this corruption challenge, and I quickly threw something together. Nobody knew it was coming, and it was about 8 hours before anyone even noticed the challenge. Week 1 was won by Brent Ozar, and the solution involved recovering the missing data from covering non-clustered indexes, removing the corruption, and replacing the data from what was stored from the NC indexes. See the solution from Brent Ozar. Also there is an alternate solution provided by Andre Kamman.

Week 2 built on week 1 utilizing a similar database, but with different problems. Week 2 was won by Rob Farley and involved pulling what was not corrupt out of the corrupt table, storing it off, then pulling the missing data from a backup, truncating the corrupt table, and then inserting the saved off data. The week 2 winning solution was provided just over 3 hours after the challenge was posted. An alternate solution was provided by Bogdan Sahlean.

Between week 2 and 3, I came up with the details on how scoring will work, decided on 3 criteria, the first correct solution each week gets 2 points, each other correct solution that is submitted before the deadline gets 1 point, and during the entire challenge, you can earn 1 point my linking to the Corruption Challenge from your or blog.

Week 3 was won by Randolph West just 37 minutes after the challenge was posted. This was the

first week that Randolph has competed in the contest. This solution involved using full backups, log backups, and pulling a tail of the log backup from a database where the mdf was missing and there was only a ldf file. This one was a lot of fun, however people were able to solve it quicker than I had expected.

Week 4 was the second win for Randolph West. Randolph won before the extra clue was provided, however many participants asked for an extra clue, and one was provided. Week 4 involved a database that had CDC or Change Data Capture enabled on the database. Using the WITH KEEP_CDC parameter on the restore allowed additional data to be included in the database which made it very straightforward to solve the corruption. An alternate one one of the more interesting solutions which completely skipped the CDC option and pulled the data from the corrupt pages to reconstruct the data using CTE's was posted by Patrick Flynn.

Week 5 was won by Patrick Flynn, and his week 5 solution is posted. Week 5 ended up being a bit too convoluted to configure. This involved having the exact right version of SQL Server, with a corrupt boot page. I ended up copying the boot page with a hex editor from a restored backup, I originally saw this done at a presentation at PASS summit by Paul Randal. Week 5 had a second corruption, once the database was restored and running, there was additional corruption that needed to be fixed. My solution based on the boot page copy was posted as the alternate for Week 5. Week 5 also involved an extra clue when it was requested by some participants.

Week 6, the winner was Raul Gonzalez. An extra clue was also provided. The tricky part with week 6 was there was corruption in both a clustered index, and an non-clustered index. Dropping the non clustered index prevented DBCC CheckTable from seeing the corruption in the clustered index. Week 6 also added an extra point for posting details to a Slow DBCC CheckDB blog post.

All of the winners listed in the post completed the corruption challenge without using the extra clue.

Updated Scoring

After Week 6, I have decided to modify the scoring going forward, additional details added in bold.

- Each week you can earn the following points:
    - 2 points for being the first to submit a correct solution, as long as you mention the contest on twitter, with a link.
    - 1 point for submitting a correct solution before the deadline.
    - **Bonus: 1 extra point to everyone who completes the challenge before any extra clues are provided. I will wait at least 24 hours before posting any extra clues.**
    - **There may be additional points described an a challenge description.**
- Once in the competition you can earn points for the following:
    - 1 point for creating a blog post **or LinkedIn post** that talks about the challenge, and links to http://SteveStedman.com/corruption (if you do this you should email a link to the post). **For the LinkedIn Post, status updates don't count, you need to click the "Publish a Post" link.**

The LinkedIn option was added when some people commented on the fact that they don't have a blog, and the "Publish a Post" feature on LinkedIn is similar to a blog post.

The view the current score leader board, visit the Current Scores page.

What About the T-Shirts

A while back I announced that there would be T-Shirts awarded as a prize, then number of shirts would depend on the number of participants. As we get closer I am finalizing the number of shirts that will be awarded. Right now I am thinking that I will award at least 10 shirts to the top 10 scoring participants. I will hand these out to the winners at PASS Summit in October. For those who won't be attending PASS Summit, I can ship the shirt to you, however there may be additional shipping charges depending on where you live.

So What About Week 7?

Week 7 will be more like what we have seen in weeks 1, 2, 3, and 6. I would say that it will be more diabolical, however I have learned that with all the participants, what one person finds easy, another finds diabolical and vice versa.

The Biggest Take Away Items

The biggest thing that I see as a take away after scoring many correct solutions, and some non-correct solutions.

- There is usually more than one way to fix the corruption.
- The winning solutions generally explore the data, or differences (from clustered to non-clustered, or from the current backup to an older backup). Those who understand what is different before attempting to fix any corruption usually get it right the first time.

Stay tuned for Database Corruption Challenge #7 coming soon.