PRACTICAL SQL SERVER

SQL Server Database Corruption, Part VII: Backups

overachiever | Mar 16, 2012

In previous posts we looked at how to enable regular checks for corruption and how to set up alerts for IO subsystem problems when they occur. But, as covered previously, one of the key means for properly dealing with corruption is to detect it early – so that you have more options for correcting it at your disposal that you would if it's allowed to go unnoticed for long periods of time.

Consequently, in this post we'll take a peek at how you can use backups as an additional means of detecting corruption.

Bonus Points: Leveraging Backups in the War on Corruption

Simply stated: you need backups. While Mirroring and AlwaysOn do provide cool features that will allow you to recover from some forms of corruption auto-magically, the reality is that backups are still your best line of defense against certain forms of disaster. Consequently, since you need backups, there's no reason you couldn't, or shouldn't, leverage them in the war on corruption – as a way to more quickly discover physical corruption after it happens.

So, given that it's already assumed that you've got regular FULL/DIFFERENTIAL backups plus

regular transaction log backups on hand to meet your target RPOs and RTOs (and assuming that you're regularly testing those backups to make sure that they're viable), the fact that backups – by their very nature – are coming into direct contact with data that has been modified means that they're a great, additional, means of detecting corruption that might have occurred or gone unnoticed during write operations.

Enabling CHECKSUMs as part of your Backup Process

When it comes to using backups as a means of detecting corruption, the key thing to note about backups is that they're designed to copy pages (or blocks of data) from one file location to another – as a means of creating a duplicate copy (or backup) of SQL Server data. Accordingly, whenever a FULL backup kicks off, **IF** it could be directed to CHECK all pages and make sure they were valid, that would be a great, additional, way to ensure that corruption hadn't occurred at the page level when data had been written out. Likewise, as DIFFERENTIAL backups ONLY copy or backup data that has been modified since the last FULL backup, DIFFERENTIAL backups would, in turn, be a great way to check and see **IF** any data had been incorrectly saved in the process of being modified.

And, fortunately, SQL Server provides the ability to perform CHECKSUM operations against data as its being backed up. To do this, you just need to 'modify' any existing backups you're already taking to make use of the WITH CHECKSUM option that is provided by SQL Server.

It's also possible to use the CHECKSUM option when backing up transaction logs – just as a means of ensuring that data is actually written to transaction log backups correctly as well. And, by combining CHECKSUM operations with the RESTORE VERIFYONLY command, you can then explicitly tell SQL Server to both check each page as it's being backed up or copied to disk AND telling it to go back through when the operation is complete and make sure that all data was correctly written to disk – a powerful means of helping ensure backup quality.

It's also worth noting that if you're using 3rd party backup solutions, it's also **usually** possible to specify the WITH CHECKSUM option along with whatever options you've chosen in terms of compression or encryption. Though, in my experience too few 3rd party backup solutions provide the option to add the WITH CHECKSUM clause into the mix directly from the GUI. Instead, what I usually find is that you'll have to go in and modify whatever routines or scripts your 3rd party GUI is creating. (A bit of a bummer, obviously – but totally worth the effort in the vast majority of cases.)

Backups + CHECKSUM ARE Not a Replacement for Regular Corruption Checks

Finally, as great as the WITH CHECKSUM option is with your backups, CHECKSUMs + BACKUPS are NOT a substitute for regularly scheduled checks. Instead, using CHECKSUM'd backups merely serve as an additional, potential, option for early detection. Likewise, IF an error is found during execution of backups, you'll want to pay attention to HOW SQL Server addresses that by looking at the corresponding options for STOP_ON_ERROR or CONTINUE_AFTER_ERROR to see what makes best sense for you.

Part VIII: The Importance of Backups

TAGS: MICROSOFT SQL SERVER

2 COMMENTS

```
SELECT STRING_AGG(CHAR(n), '')
  WITHIN GROUP (ORDER BY CHAR(n) COLLATE Latin1_General_CS_
FROM dbo.GetNums(0, 255)
WHERE CHAR(n) COLLATE Latin1_General_CS_AS LIKE '[A-Z]';

SELECT STRING_AGG(CHAR(n), '')
  WITHIN GROUP (ORDER BY CHAR(n) COLLATE Latin1_General_BIN
FROM dbo.GetNums(0, 255)
WHERE CHAR(n) COLLATE Latin1_General_BIN LIKE '[A-Z]';
```

The Challenge Is on: Maintain a Custom Letter-Based Sequence

If you need to implement a custom sequence solution, and you want it to perform as well as possible, you need to ask yourself which guarantees the solution needs to provide versus which it doesn't.

Itzik Ben-Gan | Apr 01, 2018

Recently I delivered an advanced T-SQL class in Europe, and one of my students, Nicolas Barreau of Avanade France SAS, presented a nice little puzzle. The task wasn't very complex, but it was interesting and fun to work on, and I hope you will enjoy working on it, too.

The Challenge

The challenge Nicola presented was to maintain a custom, letter-based, cycling, gap-allowing, sequence, with values in the form LLL, where L is an alphabet letter in the range A through Z (26 symbols). Nicola's company uses the sequence value as a match flag for settlement, for instance, between payment and invoice. The first value produced should be AAA, and it should advance as if the letters were digits in base 26, like so:

AAA

AAB

…

AAY

AAZ

READ MORE

TAGS: SOFTWARE DEVELOPMENT

0 COMMENTS

RELATED

The Challenge Is on: Multi-Way Number Partitioning with T-SQL and SQL CLR
FEB 21, 2018

SQL Server: Avoiding a Sort with Descending Order
JAN 19, 2018

SQL Server: Avoiding a Sort with Merge Join Concatenation
DEC 18, 2017

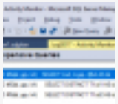

T-SQL Puzzle Challenge: Grouping Connected Items
NOV 27, 2017

Sign up for the ITPro Update newsletter.

Email address

United States

SIGN UP ▸

SECURITY > ENDPOINT SECURITY

Apple Is Ready for the Tech Backlash

The only real guarantee that companies won't go overboard in invading customers' privacy is that they have large revenue streams which make it an unnecessary risk.

Bloomberg | Mar 30, 2018

(Bloomberg View) --It's easy to understand Apple CEO Tim Cook's somewhat self-righteous comments on the Facebook-Cambridge Analytica scandal: His company has never emphasized customer data monetization, and privacy has been one of its persistent selling points. Apple still collects our data — and reserves the right to share it — but the company has a much better track record than newer tech titans.

"The truth is, we could make a ton of money if we monetized our customer — if our customer was our product," Cook said on Wednesday. "We've elected not to do that." Asked what he'd do if he were in Facebook CEO Mark Zuckerberg's shoes, he replied, "I wouldn't be in this situation."

In reality, Apple collects more information about its customers than Facebook because it offers more products and services. It knows everything from an iPhone's location to its owner's taste in

more products and services. It knows everything from an iPhone's location to its owner's taste in movies and music. It's all in the company's privacy statement, including the frank admission that Apple will share the data with "strategic partners that work with Apple to provide products and services, or that help Apple market to customers" and with law enforcement. Apple also lets applications that run on its platform (including Facebook) hoover up data. And it does have a small ad platform that allows targeted advertising.

Apple's closest equivalent in terms of data collection, sharing and "pass-through" to app developers is Google, which offers roughly the same array of services (plus search) and also owns mobile and desktop operating systems. The difference is that Apple has imposed some voluntary restrictions on itself: For example, it anonymizes location data (unless you're using the "Find My iPhone" feature) and generally employs "differential privacy" — a cryptography-based practice of obtaining usage and preference data without linking it to specific users. Google and Microsoft also use that to some degree, and a co-inventor of differential privacy, Cynthia Dwork, has even worked for Microsoft Research, but Apple was the first to market the technique as an advantage for its customers.

Cook clearly thinks of these self-restrictions as a good way of staving off external regulation. On Wednesday, he scolded Facebook for not being as far-sighted. "I think the best regulation is no regulation, is self-regulation," he said. "However, I think we're beyond that here."

As a user, however, I wouldn't put much trust in any company's ability to regulate itself. Cook has been wary of monetizing Apple's user base, but he won't be CEO forever. And people do change their mind, especially when it comes to picking up money that's there for the taking. Regulation — such as the EU's General Data Protection regulation, launching in May — is a more reliable way to ensure a level playing field. That's why Google, Apple and Microsoft all have good, convenient privacy controls which allow the user to turn off various kinds of data collection, erase data that companies already have and opt out of targeted advertising. Facebook, too, promises to put such controls in place soon.

But even external regulation isn't a reliable guarantee of privacy. Most people won't fiddle with the controls because they may not even realize what all the fuss is about. And by default, companies with advertising-based business models will want to collect as much information as they can and anonymize as little of it as they can. Neither Google nor Facebook, for example, will make the same commitment as Microsoft that no ads will be targeted based on your email and chat contents. Nor will they make it as easy as Apple and Microsoft make to shut off ad personalization.

No regulator will take care of that: It's not their goal to kill off the ad-based business model and