Steve Stedman
Freelance SQL Server Consultant
http://stevestedman.com

Database Corruption Challenge #9 - How I Corrupted the Database.

I was asked how I caused the corruption in Database Corruption Challenge #9. Here is how I did it.

To cause the corruption, I used the undocumented DBCC WritePage, however the same thing could have been accomplished by detaching the database, opening it with a hex editor, and then re-attaching the db.

DBCC WritePage

WARNING: DBCC WritePage is a dangerous command, that should never be used on any production database. It may invalidate your ability to get support from Microsoft on issues that arise with that database going forward. It is not my intention to encourage anyone to use DBCC WritePage ever. This is just what I used to create a corrupt database, and since creating corrupt databases is not part of the role of most DBAs, you should not use DBCC WritePage. Consider yourself warned.

First I used DBCC IND to find the pages that were part of the index, then I used those page (189, 226, 227, and 230) with DBCC WritePage to stomp on some memory to simulate corruption.

Then something similar for [Record1000] clustered index.

And again something similar for [Record1005].

This was all done on SQL Server 2005 so that people using the corruption challenge could restore on SQL Server 2005 or newer. When you restore to a newer version of SQL Server, the database gets upgraded, and some of the page numbers may be different than what I used to create it.

Once the corruption was introduced, I then created a full database backup and tested it on other (newer) version of SQL Server to confirm that it would work for everyone.

-Steve Stedman