

When Database Corruption Strikes

Presented by Steve Stedman

Founder/Owner of Stedman Solutions, LLC.

Partner at SQL Data Partners.

About Steve

- **SQL Data Partners – Weekly Podcast** Co-Host <http://SqlDataPartners/Podcast>
- Blog regularly at <http://SteveStedman.com>
- Founder of the **Database Corruption Challenge**
- Creator of **Database Health Monitor** <http://DatabaseHealth.com>
- 26 years of SQL Server experience
- Freelance Database Consultant (SQL Server) with Stedman Solutions, and SQL Data Partners.
- Bellingham SQL Server Users Group – Chapter Leader

Agenda

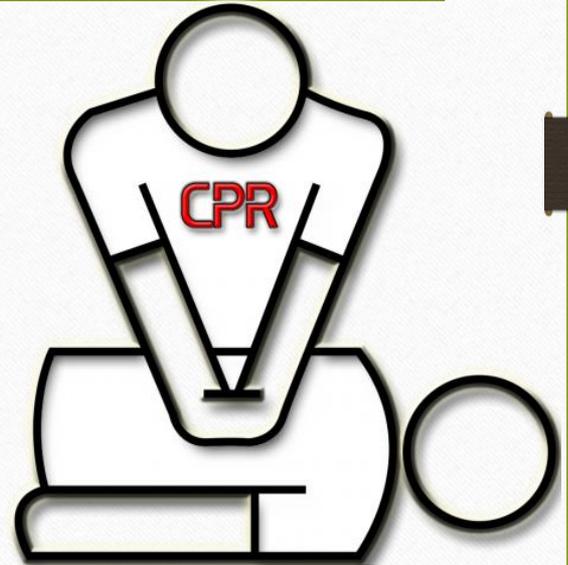
- Causes of Corruption
- Detecting Corruption
- Tracking Corruption (what has gone bad?)
- Before fixing or removing corruption
- Removing Corruption
- Examples (2)

Database Corruption Challenge

- Was an about weekly event that ran for 10 weeks.
- I created a corrupt database and confirm that it can be recovered.
- The corrupt database is posted.
- Participants send me their solutions.
- 10 Corrupt databases available for download from <http://SteveStedman.com/corruption>

Training for Database Corruption and Training for CPR

- Too late to learn when you need them.
- Outcomes can be very bad if not trained.
- Regular practice required to keep your skills sharp.



Confusion With Database Corruption

- Backup and restore of a corrupt database may help fix the corruption.
 - **FALSE.** When you back up a database, the corruption is backed up.
- Rebooting the SQL Server may help with the corruption.
 - **FALSE.** Once the file is corrupt a reboot will not help.
- If I just ignore the corruption it may go away or fix itself.
 - **UNLIKELY.** If your regular process truncates the table with the corruption, then it will go away... Otherwise, very unlikely.

Detecting Corruption

- DBCC CheckDB

```
DBCC CheckDB();
```

100 %

Messages

```
There are 0 rows in 0 pages for object "sys.sqlagent_jobs".
DBCC results for 'sys.sqlagent_jobsteps'.
There are 0 rows in 0 pages for object "sys.sqlagent_jobsteps".
DBCC results for 'sys.sqlagent_job_history'.
There are 0 rows in 0 pages for object "sys.sqlagent_job_history".
DBCC results for 'sys.sqlagent_jobsteps_logs'.
There are 0 rows in 0 pages for object "sys.sqlagent_jobsteps_logs".
Msg 8944, Level 16, State 13, Line 1
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc un
Msg 8944, Level 16, State 13, Line 1
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc un
Msg 8928, Level 16, State 1, Line 1
Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID 7205759
Msg 8976, Level 16, State 1, Line 1
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc un
DBCC results for 'Revenue'.
There are 27 rows in 1 pages for object "Revenue".
CHECKDB found 0 allocation errors and 4 consistency errors in table 'Revenue' (object
DBCC results for 'sys.queue_messages_1977058079'.
```

Detecting Corruption

- DBCC CheckDB
- DBCC CheckTable

```
DBCC CheckTable(Revenue);
```

100 % <

Messages

```
Msg 8944, Level 16, State 13, Line 8
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, a
Msg 8944, Level 16, State 13, Line 8
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, a
Msg 8928, Level 16, State 1, Line 8
Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID
Msg 8976, Level 16, State 1, Line 8
Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, a
DBCC results for 'Revenue'.
There are 27 rows in 1 pages for object "Revenue".
CHECKTABLE found 0 allocation errors and 4 consistency errors in table 'Revenue
repair_allow_data_loss is the minimum repair level for the errors found by DBCC
DBCC execution completed. If DBCC printed error messages, contact your system a
```

Detecting Corruption

- DBCC CheckDB
- DBCC CheckTable
- DBCC Check_____

 - Constraints, Catalog, Alloc, FileGroup, Ident

Detecting Corruption

- DBCC CheckDB
 - DBCC CheckTable
 - DBCC Check_____
 - Constraints, Catalog, Alloc, FileGroup, Ident
- msdb..suspect_pages

```
SELECT * FROM msdb..suspect_pages;
```

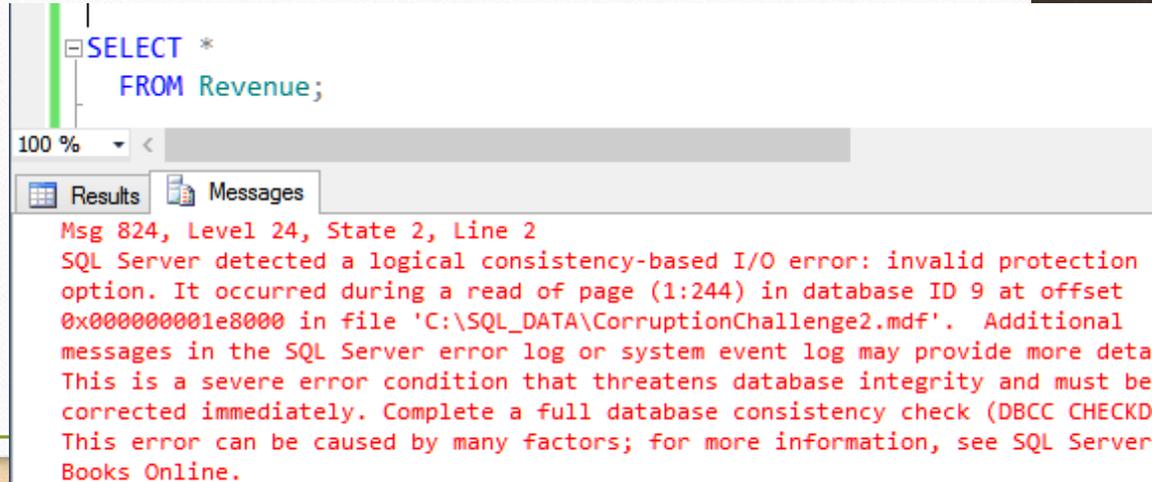
.00 % <

Results Messages

	database...	file_id	page...	event_ty...	error
1	11	1	244	4	8
2	11	1	244	1	4

Detecting Corruption

- DBCC CheckDB
- DBCC CheckTable
- DBCC Check_____
 - Constraints, Catalog, Alloc, FileGroup, Ident
- msdb..suspect_pages
- Just running a query may show corruption.



The screenshot shows a SQL Server query window with the following text:

```
SELECT *  
FROM Revenue;
```

Below the query, there are tabs for "Results" and "Messages". The "Messages" tab is active, displaying an error message in red text:

Msg 824, Level 24, State 2, Line 2
SQL Server detected a logical consistency-based I/O error: invalid protection option. It occurred during a read of page (1:244) in database ID 9 at offset 0x000000001e8000 in file 'C:\SQL_DATA\CorruptionChallenge2.mdf'. Additional messages in the SQL Server error log or system event log may provide more details. This is a severe error condition that threatens database integrity and must be corrected immediately. Complete a full database consistency check (DBCC CHECKDB). This error can be caused by many factors; for more information, see SQL Server Books Online.

Detecting Corruption

- DBCC CheckDB
- DBCC CheckTable
- DBCC Check_____
- Constraints, Catalog, Alloc, FileGroup, Ident
- msdb..suspect_pages
- Just running a query may show corruption.
- Recovery Pending

Object Explorer

Connect ▾



- [-] . (SQL Server 13.0.1601.5 - STEVEASUSLAPTOP\Steve)
- [-] Databases
 - [+] System Databases
 - [+] Database Snapshots
 - [+] CorruptionChallengeWeekA (Recovery Pending)
 - [+] Security

Tracking Corruption (what has gone bad?)

- Error messages and the error log

```
Msg 8944, Level 16, State 13, Line 1
```

```
Table error: Object ID 2105058535, index ID 1, partition ID 72057594038845440, alloc  
unit ID 72057594039762944 (type In-row data), page (1:158), row 3. Test  
(ColumnOffsets <= (nextRec - pRec)) failed. Values are 3139 and 288.
```

```
Msg 8944, Level 16, State 13, Line 1
```

```
Table error: Object ID 2105058535, index ID 1, partition ID 72057594038845440, alloc  
unit ID 72057594039762944 (type In-row data), page (1:158), row 3. Test  
(ColumnOffsets <= (nextRec - pRec)) failed. Values are 3139 and 288.
```

```
CHECKDB found 0 allocation errors and 4 consistency errors in table 'Revenue' (object ID
```

```
CHECKDB found 0 allocation errors and 4 consistency errors in database 'CorruptionChalle
```

```
repair_allow_data_loss is the minimum repair level for the errors found by DBCC CHECKDB
```

Tracking Corruption (what has gone bad?)

Selected row details:

Date	5/10/2015 4:20:36 PM
Log	SQL Server (Archive #1 - 5/10/2015 8:53:00 PM)
Source	spid52

Message

SQL Server detected a logical consistency-based I/O error: incorrect pageid (expected 1:9; actual 0:0). It occurred during a read of page (1:9) in database ID 8 at offset 0x000000000012000 in file 'C:\SQL_DATA\CorruptionChallenge5.mdf'. Additional messages in the SQL Server error log or system event log may provide more detail. This is a severe error condition that threatens database integrity and must be corrected immediately. Complete a full database consistency check (DBCC CHECKDB). This error can be caused by many factors; for more information, see SQL Server Books Online.

Tracking Corruption (what has gone bad?)

11/8/2015 2:44:01 ...	spid57	External dump process return code 0x20000001. External dump process returned no errors.
11/8/2015 2:43:58 ...	spid57	[INFO] Identity Begin End State Result Error Speculate Prepared LazyCommit ReadOnly
11/8/2015 2:43:58 ...	spid57	Stack Signature for the dump is 0x0000000000000074
11/8/2015 2:43:58 ...	spid57	* Short Stack Dump
11/8/2015 2:43:58 ...	spid57	* _____
11/8/2015 2:43:58 ...	spid57	* *****
11/8/2015 2:43:58 ...	spid57	*
11/8/2015 2:43:58 ...	spid57	* DBCC CheckDB(CorruptionChallenge1) WITH NO_INFOMSGS;
11/8/2015 2:43:58 ...	spid57	* Input Buffer 132 bytes -
11/8/2015 2:43:58 ...	spid57	*
11/8/2015 2:43:58 ...	spid57	* DBCC database corruption
11/8/2015 2:43:58 ...	spid57	*
11/8/2015 2:43:58 ...	spid57	* Private server build.
11/8/2015 2:43:58 ...	spid57	* 11/08/15 14:43:58 spid 57
11/8/2015 2:43:58 ...	spid57	* BEGIN STACK DUMP:
11/8/2015 2:43:58 ...	spid57	*
11/8/2015 2:43:58 ...	spid57	* *****
11/8/2015 2:43:58 ...	spid57	***Stack Dump being sent to D:\SQL2014\MSSQL12.SQL2014\MSSQL\LOG\SQLDump0022.txt

Tracking Corruption (what has gone bad?)

- Error messages and the error log
- See what you can query

-- lets see what we have in the corrupt table

```
SELECT *
```

```
FROM Revenue;
```

-- 54 rows

Tracking Corruption (what has gone bad?)

- Error Messages and the Error Log
- See what you can query
- Check your non-clustered indexes

```
-- pull from the non-clustered index without  
-- touching the clustered index
```

```
SELECT [id], [DepartmentID], [Revenue]  
FROM Revenue  
WITH (INDEX (ncDeptIdYear) );
```

Before Fixing or Removing Corruption

- Do you have a way to start over if something goes wrong?
- Do you have a backup of the current state?
- If your solution is going to cause data loss, can you save anything before causing that data loss?
- Do you have someone to review your ideas before proceeding?

Can I Get a “Do Over”?

What if you go through the whole process, but determine that part of your cleanup lost that could have been saved in the beginning?



Removing Corruption

- Restore from backup, prior to when the corruption was encountered.
 - Common solution. Lose data back to the point in time that corruption was encountered
 - Not always feasible.
 - Missing Backups.
 - Corruption has been there longer than your backup retention period.

Removing Corruption

- Drop/Recreate Index – if corruption is in a non-clustered index
- Truncate table – if you have a way to get the contents back
- `DBCC CheckTable(Revenue, REPAIR_REBUILD);`
- `DBCC CheckTable(Revenue, REPAIR_ALLOW_DATA_LOSS);`
- `DBCC CheckDB(database1, REPAIR_ALLOW_DATA_LOSS);`

Example 1

- You are given a .bak file with a corrupt database.
- That's it... That is all you have access to.
- The original database was already destroyed by a network admin attempting to solve the problem.
- Earlier non-corrupt backups were overwritten with the corrupt backup.

Example 1 – DBCC CheckDB

```
DBCC CheckDB(CorruptionChallenge1) WITH NO_INFOMSGS;
```

100 % <



Messages

Msg 8944, Level 16, State 13, Line 2

Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID 72057594045857792 (type In-row data), page (1:280), row 3. Test (ColumnOffsets <= (nextRec - pRec)) failed. Values are 3139 and 288.

Msg 8944, Level 16, State 13, Line 2

Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID 72057594045857792 (type In-row data), page (1:280), row 3. Test (ColumnOffsets <= (nextRec - pRec)) failed. Values are 3139 and 288.

Msg 8928, Level 16, State 1, Line 2

Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID 72057594045857792 (type In-row data): Page (1:280) could not be processed. See other errors for details.

Msg 8976, Level 16, State 1, Line 2

Table error: Object ID 245575913, index ID 1, partition ID 72057594040614912, alloc unit ID 72057594045857792 (type In-row data). Page (1:280) was not seen in the scan although its parent (1:281)

Stop - Think

- What would I do first?
- If my work makes something worse, how can I undo what I try?
- What can I do to determine what is corrupt?
- What can I do to recover missing data caused by the corruption?

Example 1

- Demo

Example 1

- Any Questions...

Example 2

- This was week 4 in the Database Corruption Challenge.
- Here is what we know. Everything was just fine, and all of a sudden users started reporting the following error when selecting from the database:

 Results

 Messages

```
Msg 605, Level 21, State 3, Line 1
Attempt to fetch logical page (3:22) in database 5 failed. It
belongs to allocation unit 1900544 not to 72057594039828480.
```

Example 2 – Errors from CheckDB

Msg 2534, Level 16, State 2, Line 1

Table error: page (3:22), whose header indicates that it is allocated to object ID 29, index ID 1, partition ID 281474978611200, alloc unit ID 1900544 (type In-row data), is allocated by another object.

Msg 2534, Level 16, State 2, Line 1

Table error: page (3:25), whose header indicates that it is allocated to object ID 29, index ID 1, partition ID 281474978611200, alloc unit ID 1900544 (type In-row data), is allocated by another object.

Msg 2511, Level 16, State 2, Line 1

... and many more

Example 2 – Additional Details

- The full backup file after the corruption occurred.
- Let's just say this database didn't have the best backup and restore schedule in place.
- There are no earlier backups to use.

Stop - Think

- What would I do first?
- If my work makes something worse, how can I undo what I try?
- What can I do to determine what is corrupt?
- What can I do to recover missing data caused by the corruption?

Example 2

- Demo

Example 2

- Any Questions?

Summary

- Detecting Corruption
- Tracking Corruption (what has gone bad?)
- Before fixing or removing corruption
- Removing Corruption
- Examples

Don't forget to ask for help...

- If you are stuck with a corrupt database don't be afraid to ask for help.

More Examples

- More Corruption Challenges available on my website.
- My website: <http://SteveStedman.com/Corruption>
- Connect
 - Twitter [@SqlEmt](#)
 - LinkedIn <http://linkedin.com/in/stevestedman>
 - Email Steve@StedmanSolutions.com
- Find me at PASS Summit 2016



Thank You

- Questions?



Example - On Your Own

- This was week 3 in the Database Corruption Challenge
- <http://stevestedman.com/fFQvC>
- Download the challenge and try it on your own.

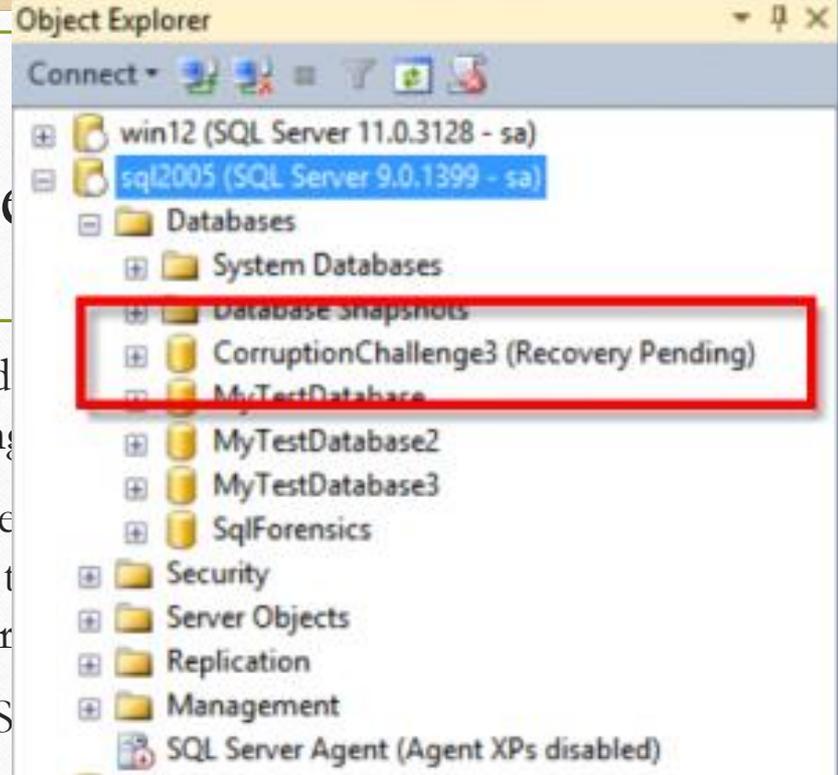
Example 4 – Timeline

- 2:53pm - full backup
- 2:54pm - transaction log backup
- 3:01pm - transaction log backup
- 3:12pm - transaction log backup
- 3:19 a catastrophic failure occurred causing the operating system running SQL Server to reboot.
 - Let's blame it on a power surge, or maybe it was just someone who accidentally bumped the power cord. I like the power surge idea, so that I don't get blamed for kicking the power cord.



Example

- After the server rebooted, several other databases were just fine, however the CorruptionChallenge3 database was in a Recovery Pending state.
- After examining the SQL Server data directory, the file for this database was gone. The only files left were the .ldf file and previous backups for the database.
- This was caused by some issue between SQL Server and the operating system.
- The database is in the Recovery Pending state.



Stop - Think

- What would I do first?
- If my work makes something worse, how can I undo what I try?
- What can I do to determine what is corrupt?
- What can I do to recover missing data caused by the corruption?