

Corruptions!

Michael Paquier – LT
PGCon 2018

Examples

- could not read block N of relation X/Y/Z: read only 0 of 8192 bytes
- catalog is missing N attribute(s) for relid M
- WARNING: page verification failed, calculated checksum %u but expected %u
ERROR: invalid page in block %u of relation %s

Corruption?

- Origin
 - Postgres => bugs or misconfiguration
 - File system
 - Kernel
 - Physical: CPU, memory, disks
 - Cosmic rays: flip bits.
- <https://wiki.postgresql.org/wiki/Corruption>

What to do?

- Breathe!
- Stop Postgres.
- Cold copy of data folder.
- Play with it and save as much as you can.
- Dump
 - zero_damaged_pages
 - ignore_checksum_failure

pg_dump

- Not a tool for corruption detection
- Still, sequential scan on all heap pages.

```
$ pg_dump > /dev/null 2>&1
```

pg_catcheck

- Check system catalog corruption
- https://github.com/EnterpriseDB/pg_catcheck

```
$ pg_catcheck
notice: pg_class row has invalid relnatts "1":
  attribute 1 does not exist in pg_attribute
row identity: oid="16385" relname="aa" relkind="r"
progress: done (1 inconsistencies, 0 warnings, 0 errors)
```

amcheck

- Present since v10
- Check consistency of relation structures
 - Support for btree
- heapallindexes since v11
- Direct ERROR on failure, wrap on top of plpgsql
- <https://www.postgresql.org/docs/devel/static/amcheck.html>

Page checksums

- `initdb -k`
- `SHOW data_checksums;`
- No way to enable checksums online. Yet.

pg_verify_checksums

- Introduced in v11
- Verify checksums if enabled.
- Server **cleanly** shut down.

```
$ pg_verify_checksums  
Checksum scan completed  
Data checksum version: 1  
Files scanned: 1225  
Blocks scanned: 3696  
Bad checksums: 0
```

pg_checksums

- Built on top of pg_verify_checksums
 - Verify
 - Enable
 - Disable
- Choose your version:
 - https://github.com/michaelpq/pg_plugins/tree/master/pg_checksums
 - https://github.com/credativ/pg_checksums

Questions for later.