Compression of big WAL records
(path to wholesale WAL compression)

Andrey Borodin
amborodin@acm.org
Telegram: @x4mmm

Below is a diagram of the **current structure** of a WAL record (XLogRecord). Each Full Page Image (FPI) is compressed individually.

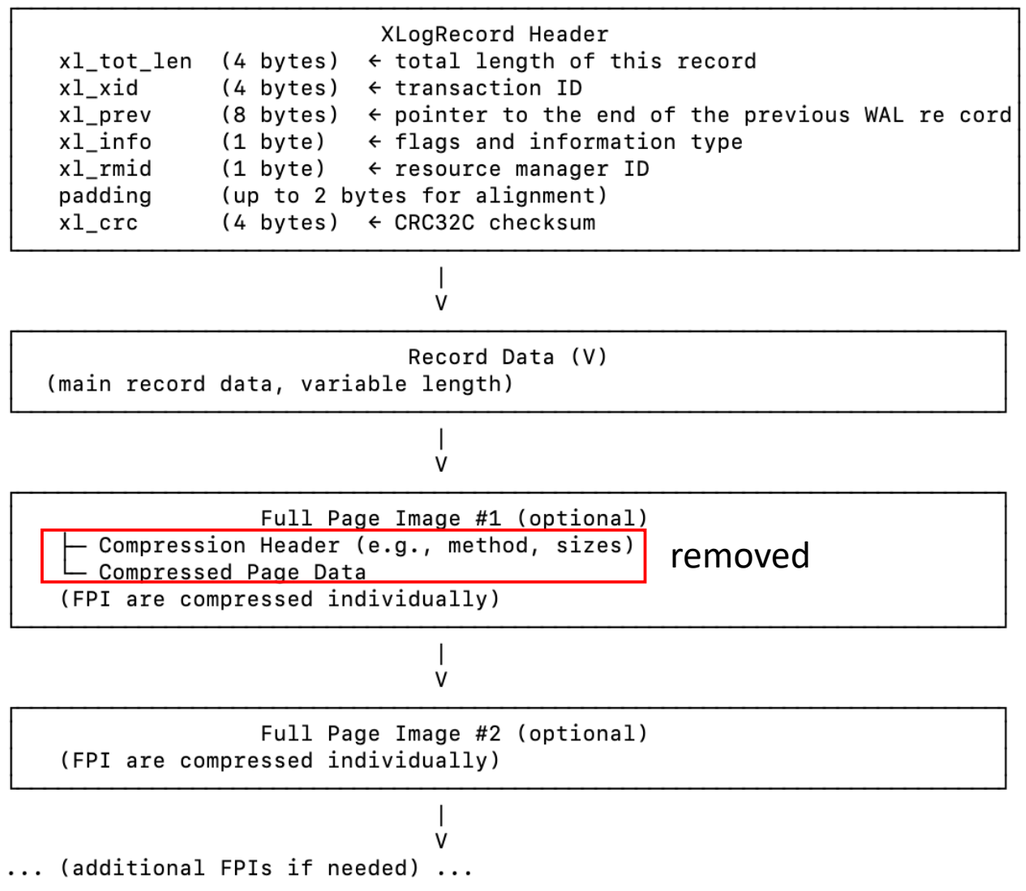

Patch idea: compress body of big records, if it's bigger than some threshold (e.g. 860 bytes).

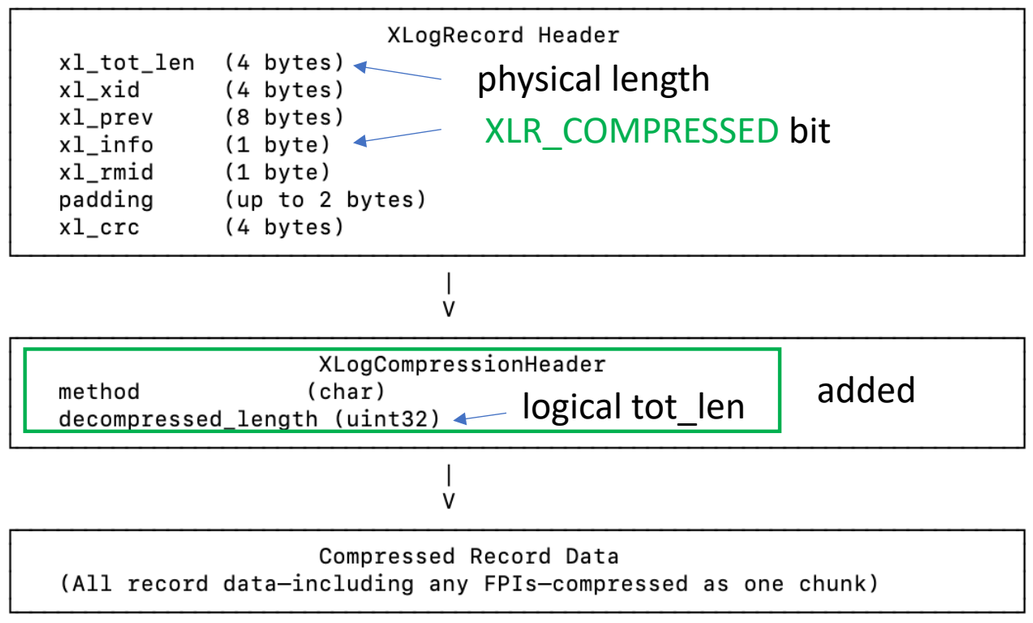

Proposed patch converts individual FPI compression headers to one compression header per record. Minimal threshold for compression is controlled via GUC **wal_compression_threshold.**

Benchmarks of current patch version indicate substantial WAL reduction in cases, when single WAL records contains many FPIs. For example, B-tree index creation:

```
create table a as
    select random() from generate_series(1,1e7);
create index on a(random );
```

WAL footprint of creating index:

method	HEAD	patched
pglz	193 MB	193 MB
lz4	160 MB	132 MB
zstd	125 MB	97 MB

Future version can reduce **wal_compression_threshold** by retaining compression context between records. But this limits ability to to read WAL at random positions. Contexts cannot cross CHECKPOINT boundaries.

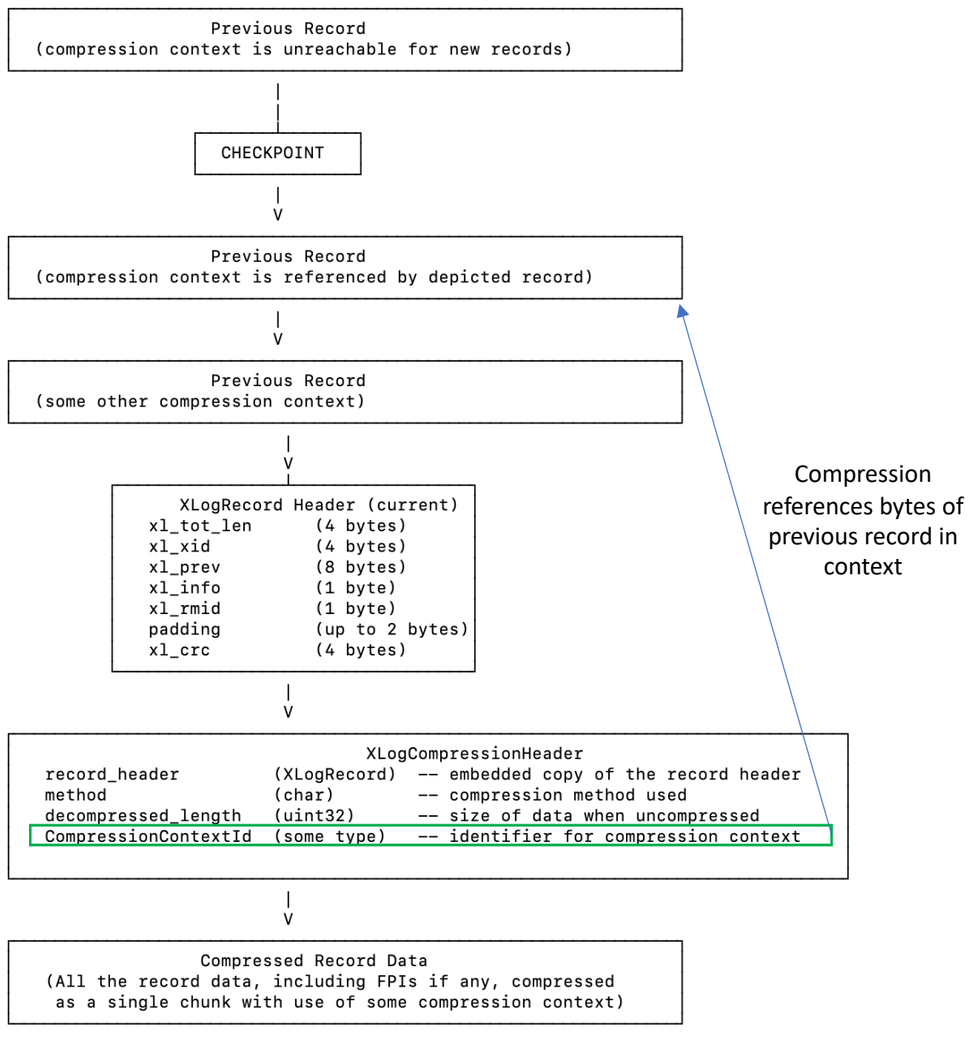

Open questions:

1. Memory allocations for compression\decompression buffers, that are hard to predict. Some are known only in critical sections.
2. Logical replication and WAL dump want to start at random LSN, not from CHECKPOINT.
3. Codecs do not play well with compression continuation: need a copy of previously compressed data.
4. Unknown unknowns.

Collaboration wanted:

1. Overall design review needed.
2. The design of compression context needs refinement.
3. Benchmarks and testing.

Commitfest item:

PGConf.dev 2025