

Wt - Bug #9595

Consider improvements to Wt Dbo floating-point round-trip/precision (Postgres backend)

01/10/2022 12:41 PM - Bruce Toll

Status:	InProgress	Start date:	01/10/2022
Priority:	Normal	Due date:	
Assignee:	Roel Standaert	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	4.9.0		

Description

As of Wt 4.6.1, it appears that some floating-point values are changed slightly when round-tripped (saved and loaded) through Wt Dbo. This can be observed with the attached patch that provides some extra round-trip floating-point tests intended to work with any of the Wt Dbo backends. To date, I've only tested them with test.sqlite3, which passes all tests, and test.postgres, which experiences some failures using either the 4.5.1 or 4.6.1 versions of Postgres.C.

A second file, attached for your review, provides a patch intended to improve the round-tripping of floating-point for the Postgres backend. The patch is fairly minimal, employing the C++ 17 `std::to_chars()` and `std::from_chars()` functions, if available. It falls back to the existing 4.5.1 code, if not, and relies on Boost Spirit to handle subnormals if they are not handled by the underlying implementation. With this patch, the round-trip floating-point tests pass on my system (when floating-point `to_chars/from_chars` are available), although there are warnings about two subnormals that do not match exactly (due to the continued use of Boost Spirit for subnormals).

Here are some additional notes, along with details on files and testing:

1. The extra tests in 0001-Enhance-round-trip-testing-of-floats-in-Wt-Dbo.patch are meant to apply on top of wt-4.6.1. The following is a reasonable starting point for testing:
`./test.postgres -l warning -t */dbo_test46, dbo_test47' -e stdout 2>/dev/null`
NOTE: Running the tests on either 4.5.1 or 4.6.1, there are comparison failures, as well as two Postgres exceptions: 'ERROR: "1.797693134862316e308" is out of range for type double precision'. This error does not occur with the proposed `to_chars` implementation.
2. The attached 0002-Improve-floating-point-round-trip-for-Postgres.patch can be applied on top of the previous patch. It should enable the tests to pass with two warnings on subnormals, if built with a C++17 compiler that fully supports `std::to_chars` (indicated by the preprocessor definition, `__cpp_lib_to_chars` in the `charconv` header file). NOTE: If `__cpp_lib_to_chars` is not defined, the code will fall-back to the 4.5.1 `stof/stod` behavior, augmented with the use of Boost Spirit for subnormals.
3. I did not do any performance testing of `to_char/from_char` vs. the existing approach. I also did not test the overhead associated with putting `stof/stof` in a try block when `to_chars` support is not available.
4. The test environment used gcc 11.2.0 and wt was built with `-DCMAKE_CXX_STANDARD=17` and boost 1.74.0. This version of gcc supports `to_chars/from_chars` with floating point; it provides the `charconv` header and defines `__cpp_lib_to_chars`.
5. The version of Postgres used in testing was 12.9. This could be relevant since the tests round-trip through Postgres' parsing (and generation) of floating-point text.
6. When I describe running tests against Wt Dbo 4.5.1, I'm only using the 4.5.1 version of Postgres.C. The rest of the build is from Wt 4.6.1.
7. In the process of working on the tests, I noticed that the behavior of 4.5.1 and 4.6.1 were different, although both had values that were changed while round-tripping. I wasn't initially sure which version of Wt was more correct, but I now believe that 4.6.1 may have a regression. I will file a separate issue for this.

If I can provide any additional information, please let me know. If you choose to distribute the test code (which incorporates values and comments from Postgres), please review it for any license/notice concerns.

I realize this is a lot of material to review and test. Hopefully, some of it will be useful. Thanks for all of the hard work on Wt 4.6 and I wish you all the best in the New Year!

History

#1 - 01/10/2022 06:59 PM - Roel Standaert

- Target version set to 4.6.2

#2 - 02/07/2022 05:06 PM - Roel Standaert

- Status changed from New to InProgress

#3 - 02/07/2022 05:06 PM - Roel Standaert

- Assignee set to Roel Standaert

#4 - 02/14/2022 12:47 PM - Roel Standaert

- Target version changed from 4.6.2 to 4.7.0

#5 - 03/10/2022 03:24 PM - Roel Standaert

- Target version changed from 4.7.0 to 4.8.0

#6 - 07/06/2022 11:49 AM - Roel Standaert

- Target version changed from 4.8.0 to 4.9.0

Files

0001-Enhance-round-trip-testing-of-floats-in-Wt-Dbo.patch	18.1 KB	01/10/2022	Bruce Toll
0002-Improve-floating-point-round-trip-for-Postgres.patch	5.87 KB	01/10/2022	Bruce Toll