U++ - Bug #2077
String0::IsEqual causing warnings with newer GCC

09/25/2020 11:03 PM - Zbigniew Rebacz

Status:	Approved	**Start date:**	09/25/2020
Priority:	Normal	**Due date:**	
Assignee:	Zbigniew Rebacz	**% Done:**	0%
Category:	Core	**Estimated time:**	0.00 hour
Target version:	Release 2020.2	**Spent time:**	0.00 hour

Description

Here is the warnings generated by gcc on Linux:

```
/home/klugier/upp/uppsrc/Core/AString.hpp: In member function ‘void Ide::ResolveUvsConflict()’:
/home/klugier/upp/uppsrc/Core/AString.hpp:269:36: warning: ‘int __builtin_memcmp_eq(const void*, const void*, long unsigned int)’ reading 17 bytes from a region of size 16 [-Wstringop-overflow=]
  269 |  return len == GetCount() && memcmp(begin(), s, len) == 0; // compiler is happy to optimize memcmp out...
      |                       ~~~~~~^~~~~~~~~~~~~~~~~
/home/klugier/upp/uppsrc/Core/AString.hpp:269:36: warning: ‘int __builtin_memcmp_eq(const void*, const void*, long unsigned int)’ reading 17 bytes from a region of size 16 [-Wstringop-overflow=]
  269 |  return len == GetCount() && memcmp(begin(), s, len) == 0; // compiler is happy to optimize memcmp out...
      |                       ~~~~~~^~~~~~~~~~~~~~~~~
/home/klugier/upp/uppsrc/Core/AString.hpp:269:36: warning: ‘int __builtin_memcmp_eq(const void*, const void*, long unsigned int)’ reading 21 bytes from a region of size 16 [-Wstringop-overflow=]
  269 |  return len == GetCount() && memcmp(begin(), s, len) == 0; // compiler is happy to optimize memcmp out...
      |
```

Solution that fix the warning (not sure about performance):

```
inline
bool String0::IsEqual(const char *s) const
{ // This optimized for comparison with string literals...
   size_t len = strlen(s);
   return len == GetCount() && strncmp(begin(), s, len) == 0;
}
```

Tested on GCC 10.2.

History

#1 - 09/26/2020 08:49 AM - Miroslav Fidler

- Status changed from New to Ready for QA

Posting GCC version would be really helpful.

Anyway, while perhaps negligible, the performance impact would be there. I believe this is basically GCC bug - he does not see the logic in begin()...

So we should find a way how to silence that... Maybe just replace with memeq8?

#2 - 09/26/2020 08:51 AM - Miroslav Fidler

For record, no warnings with gcc 7.5.0

#3 - 09/26/2020 09:34 AM - Zbigniew Rebacz

This is gcc 10.2

#4 - 09/26/2020 09:34 AM - Zbigniew Rebacz

- Description updated

#5 - 09/26/2020 09:43 AM - Zbigniew Rebacz

No warning with following line:

```
return len == GetCount() && memeq8(begin(), s, len) == 0; // compiler is happy to optimize memcmp out...
```

#6 - 09/26/2020 10:27 AM - Zbigniew Rebacz

With memeq8 strange things happen inside TheIDE - files are not loaded correctly, so it is definitely no function we would like to use...

#7 - 09/26/2020 02:12 PM - Miroslav Fidler

Without == 0!!!!

memeq returns true if equal...

#8 - 09/26/2020 04:03 PM - Zbigniew Rebacz

Works fine without comparison and doesn't produce warning.

#9 - 09/29/2020 10:45 AM - Miroslav Fidler

Unfortunately, the performance impact is still ugly. With memcmp, compiler can optimize constant comparison (like s == "test") into basically simple integer cmp...

What distro are you testing with? I will install in VirtualBox and find a way to silence that...

Mirek

#10 - 09/29/2020 11:31 AM - Zbigniew Rebacz

I use Manjaro KDE edition. This is rolling distro, so you all the time should have decent compilers.

#11 - 09/30/2020 02:17 PM - Miroslav Fidler

OK, I have installed manjaro.

gcc --version

10.2.0

With -Wall it gave warning in IsEqual0 about comparing size_t with int (ok, that is worth fixing), but not the warning you have posted.

Maybe you should update your gcc?

Mirek

#12 - 09/30/2020 02:17 PM - Miroslav Fidler

- Status changed from Ready for QA to In Progress

- Assignee changed from Miroslav Fidler to Zbigniew Rebacz

#13 - 09/30/2020 02:21 PM - Miroslav Fidler

Ah, ok, it is in release only...

#14 - 09/30/2020 05:52 PM - Zbigniew Rebacz

Yes, I compile TheIDE in release mode. GCC already updated :)

#15 - 10/01/2020 10:54 AM - Miroslav Fidler

- Status changed from In Progress to Ready for QA

I took me two days to figure all warnings out, but it should now compile with -Wall (except that && parenthesis issue).

The problem with String0::IsEqual is strongly GCC bug, but in the end I have figured out a fix that actually might improve the performance a bit...

#16 - 10/01/2020 11:33 PM - Zbigniew Rebacz

- Status changed from Ready for QA to Approved

No warnings anymore - good job :)

#17 - 10/04/2020 07:34 PM - Zbigniew Rebacz

- Status changed from Approved to New

- Assignee changed from Zbigniew Rebacz to Miroslav Fidler

Seems like the problem returned (tutorial/Skylark10 package), however in different place:

```
                 from /home/klugier/upp/uppsrc/Skylark/Dispatch.cpp:1:
/home/klugier/upp/uppsrc/Core/AString.hpp: In member function ‘void Upp::Http::Dispatch(Upp::TcpSocket&)’:
/home/klugier/upp/uppsrc/Core/AString.hpp:171:15: warning: ‘int __builtin_memcmp_eq(const void*, const void*, long unsigned int)’ reading 33 bytes from a region of size 16 [-Wstringop-overflow=]
  171 |  return memcmp(s, B::Begin(), len * sizeof(tchar)) == 0;
      |
```

Potentially all places with memcpy are affected in AString.hpp not onlu StartsWith, but EndsWith, ReversFind etc..

#18 - 10/08/2020 12:40 PM - Zbigniew Rebacz

- Status changed from New to Approved

- Assignee changed from Miroslav Fidler to Zbigniew Rebacz

This is not obvious warning and probably bug in gcc. We could close it for now.