Alpine Linux - Bug #9488

Crash while running a headless libreoffice spreadsheet conversion

09/27/2018 02:53 PM - Alexander Stepanov

Status:	Closed	**Start date:**	09/27/2018
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	100%
Category:		**Estimated time:**	0.00 hour
Target version:	3.8.2		
Affected versions:	3.6.0, 3.7.0, 3.8.0	**Security IDs:**	

Description

When Libreoffice is used in the headless mode to convert office documents (texts, spreadsheets, presentations, etc.) inside an Alpine-based Docker container via UNO interface, it crashes with a SIGSEGV while converting some documents (see the attached files).

Steps to reproduce:

1. Run a docker container:
docker run -it --rm alpine:3.8

2. Install libreoffice, python and unoconv in the container:
apk add --update wget libreoffice python3 pip3 install unoconv

3. Get the test documents and try to convert them:
unoconv -f html test.ods # has formulas, libreoffice crash - SIGSEGV unoconv -f html test2.ods # no formulas, no crash

A typical stack trace from debug build shows more details:

```
Thread 7 "cppu_threadpool" received signal SIGSEGV, Segmentation fault.
[Switching to LWP 113]
0x00007fdb8b425d70 in formula::FormulaCompiler::GetToken() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
(gdb) bt
#0  0x00007fdb8b425d70 in formula::FormulaCompiler::GetToken() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#1  0x00007fdb8b4277b4 in formula::FormulaCompiler::NextToken() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#2  0x00007fdb8b426540 in formula::FormulaCompiler::Factor() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#3  0x00007fdb8b427fef in formula::FormulaCompiler::RangeLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#4  0x00007fdb8b428214 in formula::FormulaCompiler::IntersectionLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#5  0x00007fdb8b4284bd in formula::FormulaCompiler::UnionLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#6  0x00007fdb8b4285f2 in formula::FormulaCompiler::UnaryLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#7  0x00007fdb8b4286af in formula::FormulaCompiler::PowLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#8  0x00007fdb8b4287ed in formula::FormulaCompiler::MulDivLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#9  0x00007fdb8b4288dd in formula::FormulaCompiler::AddSubLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#10 0x00007fdb8b4289cd in formula::FormulaCompiler::ConcatLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#11 0x00007fdb8b428a9d in formula::FormulaCompiler::CompareLine() ()
   from /opt/libreoffice6.1/lib/libreoffice/program/../program/libforlo.so
#12 0x00007fdb8b427ed8 in formula::FormulaCompiler::Expression() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#13 0x00007fdb8b426cb7 in formula::FormulaCompiler::Factor() () from /opt/libreoffice6.1/lib/libre
office/program/../program/libforlo.so
```

```
#14 0x00007fdb8b427fef in formula::FormulaCompiler::RangeLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#15 0x00007fdb8b428214 in formula::FormulaCompiler::IntersectionLine() () from /opt/libreoffice6.1
/lib/libreoffice/program/../program/libforlo.so
#16 0x00007fdb8b4284bd in formula::FormulaCompiler::UnionLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#17 0x00007fdb8b4285f2 in formula::FormulaCompiler::UnaryLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#18 0x00007fdb8b4286af in formula::FormulaCompiler::PowLine() () from /opt/libreoffice6.1/lib/libr
eoffice/program/../program/libforlo.so
#19 0x00007fdb8b4287ed in formula::FormulaCompiler::MulDivLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#20 0x00007fdb8b4288dd in formula::FormulaCompiler::AddSubLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#21 0x00007fdb8b4289cd in formula::FormulaCompiler::ConcatLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#22 0x00007fdb8b428a9d in formula::FormulaCompiler::CompareLine() () from /opt/libreoffice6.1/lib/
libreoffice/program/../program/libforlo.so
#23 0x00007fdb8b427ed8 in formula::FormulaCompiler::Expression() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#24 0x00007fdb8b42741b in formula::FormulaCompiler::Factor() () from /opt/libreoffice6.1/lib/libre
office/program/../program/libforlo.so
#25 0x00007fdb8b427fef in formula::FormulaCompiler::RangeLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#26 0x00007fdb8b428214 in formula::FormulaCompiler::IntersectionLine() () from /opt/libreoffice6.1
/lib/libreoffice/program/../program/libforlo.so
#27 0x00007fdb8b4284bd in formula::FormulaCompiler::UnionLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#28 0x00007fdb8b4285f2 in formula::FormulaCompiler::UnaryLine() () from /opt/libreoffice6.1/lib/li
breoffice/program/../program/libforlo.so
#29 0x00007fdb8b4286af in formula::FormulaCompiler::PowLine() () from /opt/libreoffice6.1/lib/libr
eoffice/program/../program/libforlo.so
#30 0x00007fdb8b4287ed in formula::FormulaCompiler::MulDivLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#31 0x00007fdb8b4288dd in formula::FormulaCompiler::AddSubLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#32 0x00007fdb8b4289cd in formula::FormulaCompiler::ConcatLine() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#33 0x00007fdb8b428a9d in formula::FormulaCompiler::CompareLine() () from /opt/libreoffice6.1/lib/
libreoffice/program/../program/libforlo.so
#34 0x00007fdb8b427ed8 in formula::FormulaCompiler::Expression() () from /opt/libreoffice6.1/lib/l
ibreoffice/program/../program/libforlo.so
#35 0x00007fdb8b428c62 in formula::FormulaCompiler::CompileTokenArray() () from /opt/libreoffice6.
1/lib/libreoffice/program/../program/libforlo.so
#36 0x00007fdb8bcab136 in ScFormulaCell::CompileXML(sc::CompileFormulaContext&, ScProgress&) () fr
om /opt/libreoffice6.1/lib/libreoffice/program/../program/libsclo.so
#37 0x00007fdb8bb788d4 in ScColumn::CompileXML(sc::CompileFormulaContext&, ScProgress&) () from /o
pt/libreoffice6.1/lib/libreoffice/program/../program/libsclo.so
#38 0x00007fdb8bcf0453 in ScTable::CompileXML(sc::CompileFormulaContext&, ScProgress&) () from /opt
/libreoffice6.1/lib/libreoffice/program/../program/libsclo.so
#39 0x00007fdb8bc1ee04 in ScDocument::CompileXML() () from /opt/libreoffice6.1/lib/libreoffice/pro
gram/../program/libsclo.so
#40 0x00007fdb8bfabdb8 in ScXMLImport::endDocument() () from /opt/libreoffice6.1/lib/libreoffice/p
rogram/../program/libsclo.so
```

Associated revisions

Revision 18821e8e - 09/28/2018 04:55 PM - Natanael Copa

community/libreoffice: upgrade to 6.1.0.3

and fix tread stack size issue
ref #9488

Revision 96e1e57f - 10/01/2018 12:00 PM - Natanael Copa

community/libreoffice: fix tread stack size issue

fixes #9488

History

#1 - 09/27/2018 03:10 PM - Alexander Stepanov

Correct steps:

2. Install libreoffice, python and unoconv in the container:

```
apk add --update wget libreoffice python3
pip3 install unoconv
```

3. Get the test documents and try to convert them:

```
wget https://bugs.alpinelinux.org/attachments/download/520/test.ods
wget https://bugs.alpinelinux.org/attachments/521/test2.ods

unoconv -f html test2.ods # no formulas, no crash, created test2.html

unoconv -f html test.ods # has formulas, libreoffice crash - SIGSEGV
...
unoconv: RuntimeException during import phase:
Office probably died. Binary URP bridge disposed during call
...
```

#2 - 09/27/2018 03:15 PM - Alexander Stepanov

The test documents were created with the latest LO on my desktop, and LO on Alpine Linux can open them correctly (not in the UNO mode). E.g. it is possible as a workaround to convert them via soffice CLI options:

```
soffice --version
LibreOffice 5.4.5.1 8316fbc25979b56ad2d85a05feea815f66d5e128

soffice --convert-to html test.ods
javaldx: Could not find a Java Runtime Environment!
Warning: failed to read path from javaldx
convert /test.ods -> //test.html using filter : HTML (StarCalc) # test.html is created!
```

#3 - 09/27/2018 03:38 PM - Alexander Stepanov

I have tried the recent LO (6.1.0.3) building it from source (almost the same way the libreoffice package is built in Alpine APK), but the problem persisted. The debug build has given me just more info in the stack trace (see the very first post in this issue).

My investigation revealed the following difference in MUSL vs. GLibc:

> The default stack size for new threads on glibc is determined based on the resource limit governing the main thread's stack (RLIMIT_STACK). It generally ends up being 2-10 MB.
> musl provides a default stack size of 80k. This does not include the guard page, nor does it include the space used for TLS unless total TLS size is very small. So the actual map size may appear closer to 90k, with around 80k usable by the application. This size was determined empirically with the goals of not gratuitously breaking applications but also not causing large amounts of memory and virtual address space to be committed in programs with large numbers of threads. Programs needing larger stacks, or which explicitly want a smaller stack, should make this explicit with pthread_attr_setstacksize.

See also https://wiki.musl-libc.org/functional-differences-from-glibc.html

So I've added the following patch to the LO vanilla source to fix it:

```
diff -Naur libreoffice-6.1.0.3/sal/osl/unx/thread.cxx libreoffice-6.1.0.3-patched/sal/osl/unx/thread.cxx
--- libreoffice-6.1.0.3/sal/osl/unx/thread.cxx    2018-08-02 22:54:54.000000000 +0300
+++ libreoffice-6.1.0.3-patched/sal/osl/unx/thread.cxx    2018-09-05 18:21:38.552838233 +0300
@@ -249,7 +249,7 @@
     short             nFlags)
 {
     Thread_Impl* pImpl;
-#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS)
+#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS) || (defined LINUX
 && !defined __GLIBC__)
     pthread_attr_t attr;
     size_t stacksize;
 #endif
@@ -265,7 +265,7 @@
```

```
     pthread_mutex_lock (&(pImpl->m_Lock));

-#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS)
+#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS) || (defined LINUX
 && !defined __GLIBC__)
     if (pthread_attr_init(&attr) != 0)
         return nullptr;

@@ -282,7 +282,7 @@

     if ((nRet = pthread_create (
         &(pImpl->m_hThread),
-#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS)
+#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS) || (defined LINUX
 && !defined __GLIBC__)
         &attr,
 #else
         PTHREAD_ATTR_DEFAULT,
@@ -301,7 +301,7 @@
         return nullptr;
     }

-#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS)
+#if defined OPENBSD || ((defined MACOSX || defined LINUX) && !ENABLE_RUNTIME_OPTIMIZATIONS) || (defined LINUX
 && !defined __GLIBC__)
     pthread_attr_destroy(&attr);
 #endif
```

This enabled a larger thread stack by default and fixed the problem.

Is it possible to include something like this fix in the libreoffice APK build patches (and also to migrate to LO 6.1.x) in the Alpine Edge?

#4 - 09/28/2018 12:56 PM - Natanael Copa

- Target version set to 3.8.2

Wow, good job!

Yes, default thread stack size is the usual suspect for this kind of segfaults. The patch looks correct too.

Have you reported this upstream?

#5 - 10/01/2018 12:06 PM - Natanael Copa

- Status changed from New to Resolved

- % Done changed from 0 to 100

Applied in changeset 96e1e57fed146a3449b7070143861b7e22ba57f4.

#6 - 10/01/2018 02:17 PM - Alexander Stepanov

Natanael Copa wrote:

> Wow, good job!
>
> Yes, default thread stack size is the usual suspect for this kind of segfaults. The patch looks correct too.
>
> Have you reported this upstream?

Thanks, Natanael.
I definitely want to report this upstream and I'll post an update here when it is done.
My concern is that LO supports different Linux flavors and I'm not sure the patch will work correctly for other combinations of Linux and other libc, e.g. for Linux+uLibc (and because the MUSL team refused to create a specific define for it like "__GLIBC__", as stated in https://www.openwall.com/lists/musl/2013/02/08/9).

By the way, do you plan (or have already tried) to report other MUSL-specific LibreOffice patches there?
I can see the patch for -lint in your commit (https://git.alpinelinux.org/cgit/aports/commit/?id=18821e8e6f9714bf14cac60e1c23d35e3554b332), this is something I've stumbled upon during my LO rebuild too (but I forgot to mention it in the issue, sorry). I also suggest you simplify my patch and replace libreoffice-6.1.0.3/libreoffice-6.1.0.3-patched with a/b because I haven't used the git diff (diff with --git option).

I've also successfully rebuilt LO with clang during my experiments with Alpine. It also required a small additional patch for configure.ac, and I can share it if somebody is interested in LO clang builds here.

#7 - 12/20/2018 03:47 PM - Natanael Copa

- Status changed from Resolved to Closed

Alexander Stepanov wrote:

> By the way, do you plan (or have already tried) to report other MUSL-specific LibreOffice patches there?

I dont know.

> I can see the patch for -lint in your commit (https://git.alpinelinux.org/cgit/aports/commit/?id=18821e8e6f9714bf14cac60e1c23d35e3554b332), this is something I've stumbled upon during my LO rebuild too (but I forgot to mention it in the issue, sorry).

I don't remember, but it is quite possible it was a quick fix. I don't think I have reported it upstream. Would be great if you can help me with that.

> I also suggest you simplify my patch and replace libreoffice-6.1.0.3/libreoffice-6.1.0.3-patched with a/b because I haven't used the git diff (diff with --git option).

That does not matter.

Files

test.ods	8 KB	09/27/2018	Alexander Stepanov
test2.ods	7.9 KB	09/27/2018	Alexander Stepanov