

bluestore - Bug #52230

crash in pread64 during collection_list()

08/11/2021 06:47 PM - Telemetry Bot

Status: Triaged	% Done: 0%
Priority: Normal	
Assignee:	
Category:	
Target version:	
Source: Telemetry	Affected Versions: v16.2.4, v16.2.5, v16.2.7
Tags:	ceph-qa-suite:
Backport:	Pull request ID:
Regression: No	Crash signature (v1): 0ae3ccdd27d47a526bdd327d6213e2260 b3ecd9d3c28a499adbe082f1758e51c 247261499ff94d0dda7a61ceb0ccb61c80 d3ec4bc3794df90489e8740f7b9b8f 62de4b1a9a0b46a54387992523bd04d7 94016cae2a0b1338a4bc7a9681b15368 7a39cf2b5e28efa7ddd88d1a08bbfd781f 283949f5a5cac03d62dfbe09eb70c3 476de4d52c41f5c3f03850502369f4a580 0706c20593d0d9e9aba9f685315f11
Severity: 3 - minor	Crash signature (v2): "b137bc0d6f50e94d9f5effc9ab6677955f b2c7853ba7c9acad4e907ea9d8cc80":htt p://telemetry.front.sepia.ceph.com:4000/ d/jByk5HaMz/crash-spec-x-ray?orgId=1& var-sig_v2=b137bc0d6f50e94d9f5effc9a b6677955fb2c7853ba7c9acad4e907ea9 d8cc80 "b1f94a1ad99c5aa802e6daea93f28515c 12a58c5cacccbf8d612b4d4f600afc8":htt p://telemetry.front.sepia.ceph.com:4000/ d/jByk5HaMz/crash-spec-x-ray?orgId=1& var-sig_v2=b1f94a1ad99c5aa802e6daea 93f28515c12a58c5cacccbf8d612b4d4f6 00afc8
Reviewed:	

Description

http://telemetry.front.sepia.ceph.com:4000/d/jByk5HaMz/crash-spec-x-ray?orgId=1&var-sig_v2=62de4b1a9a0b46a54387992523bd04d794016cae2a0b1338a4bc7a9681b15368

Sanitized backtrace:

```
/lib/x86_64-linux-gnu/libpthread.so.0(  
  pread64()  
  KernelDevice::read(unsigned long, unsigned long, ceph::buffer::v15_2_0::list*, IOContext*, boo  
1)  
  BlueFS::_read(BlueFS::FileReader*, unsigned long, unsigned long, ceph::buffer::v15_2_0::list*,  
  char*)  
  BlueRocksRandomAccessFile::Prefetch(unsigned long, unsigned long)  
  rocksdb::LegacyRandomAccessFileWrapper::Prefetch(unsigned long, unsigned long, rocksdb::IOOpti  
ons const&, rocksdb::IODebugContext*)  
  rocksdb::RandomAccessFileReader::Prefetch(unsigned long, unsigned long) const  
  rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::InitDataBlock()  
  rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::FindBlockForward()  
  rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::FindKeyForward()  
  rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::Next()
```

```

    rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::NextAndGetResult(rocksdb::IterateResult*)
    rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()
    /usr/bin/ceph-osd(
    /usr/bin/ceph-osd(
    rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()
    rocksdb::MergingIterator::Next()
    rocksdb::MergingIterator::NextAndGetResult(rocksdb::IterateResult*)
    rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()
    rocksdb::DBIter::FindNextUserEntryInternal(bool, rocksdb::Slice const*)
    rocksdb::DBIter::FindNextUserEntry(bool, rocksdb::Slice const*)
    rocksdb::DBIter::Next()
    rocksdb::ArenaWrappedDBIter::Next()
    RocksDBStore::RocksDBWholeSpaceIteratorImpl::next()
    /usr/bin/ceph-osd(
    BlueStore::_collection_list(BlueStore::Collection*, gobject_t const&, gobject_t const&, int, bool, std::vector<gobject_t, std::allocator<gobject_t> >*, gobject_t*)
    BlueStore::collection_list(boost::intrusive_ptr<ObjectStore::CollectionImpl>&, gobject_t const&, gobject_t const&, int, std::vector<gobject_t, std::allocator<gobject_t> >*, gobject_t*)
    PGBackend::objects_list_partial(hobject_t const&, int, int, std::vector<hobject_t, std::allocator<hobject_t> >*, hobject_t*)
    PrimaryLogPG::scan_range(int, int, BackfillInterval*, ThreadPool::TPHandle&)
    PrimaryLogPG::do_scan(boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)
    PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)
    OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)
    ceph::osd::scheduler::PGOpItem::run(OSD*, OSDShard*, boost::intrusive_ptr<PG>&, ThreadPool::TPHandle&)
    OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)
    ShardedThreadPool::shardedthreadpool_worker(unsigned int)
    ShardedThreadPool::WorkThreadSharded::entry()
    /lib/x86_64-linux-gnu/libpthread.so.0(
    clone()

```

Crash dump sample:

```

{
  "archived": "2021-07-28 09:18:38.737870",
  "backtrace": [
    "/lib/x86_64-linux-gnu/libpthread.so.0(+0x14140) [0x7f213e8ad140]",
    "pread64()",
    "(KernelDevice::read(unsigned long, unsigned long, ceph::buffer::v15_2_0::list*, IOContext*, bool)+0xf4) [0x559a9ff9b9c4]",
    "(BlueFS::_read(BlueFS::FileReader*, unsigned long, unsigned long, ceph::buffer::v15_2_0::list*, char*)+0x5bd) [0x559a9fb3391d]",
    "(BlueRocksRandomAccessFile::Prefetch(unsigned long, unsigned long)+0x1b) [0x559a9fb569bb]",
    "(rocksdb::LegacyRandomAccessFileWrapper::Prefetch(unsigned long, unsigned long, rocksdb::IOOptions const&, rocksdb::IODebugContext*)+0x4d) [0x559aa00179bd]",
    "(rocksdb::RandomAccessFileReader::Prefetch(unsigned long, unsigned long) const+0x5d) [0x559aa026b825]",
    "(rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::InitDataBlock()+0x24b) [0x559aa0289d8b]",
    "(rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::FindBlockForward()+0x2c2) [0x559aa028bb00]",
    "(rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::FindKeyForward()+0x90) [0x559aa028a57e]",
    "(rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::Next()+0x8a) [0x559aa0287ea4]",
    "(rocksdb::BlockBasedTableIterator<rocksdb::DataBlockIter, rocksdb::Slice>::NextAndGetResult(rocksdb::IterateResult*)+0x1d) [0x559aa0287ed3]",
    "(rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()+0x62) [0x559aa00d9cde]",
    "/usr/bin/ceph-osd(+0x17e1cd9) [0x559aa0170cd9]",
    "/usr/bin/ceph-osd(+0x17e26e1) [0x559aa01716e1]",
    "(rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()+0x62) [0x559aa00d9cde]",

```

```

    "(rocksdb::MergingIterator::Next()+0xaf) [0x559aa02aa1ab]",
    "(rocksdb::MergingIterator::NextAndGetResult(rocksdb::IterateResult*)+0x28) [0x559aa02aa2e6]",
    "(rocksdb::IteratorWrapperBase<rocksdb::Slice>::Next()+0x62) [0x559aa00d9cde]",
    "(rocksdb::DBIter::FindNextUserEntryInternal(bool, rocksdb::Slice const*)+0xd7b) [0x559aa00d2787]",
    "(rocksdb::DBIter::FindNextUserEntry(bool, rocksdb::Slice const*)+0x70) [0x559aa00d19d8]",
    "(rocksdb::DBIter::Next()+0x2f3) [0x559aa00d1863]",
    "(rocksdb::ArenaWrappedDBIter::Next()+0x1c) [0x559aa033c914]",
    "(RocksDBStore::RocksDBWholeSpaceIteratorImpl::next()+0x4a) [0x559a9ffb21ba]",
    "/usr/bin/ceph-osd(+0x107354d) [0x559a9fa0254d]",
    "(BlueStore::_collection_list(BlueStore::Collection*, ghobject_t const&, ghobject_t const&, int, bool, std::vector<ghobject_t, std::allocator<ghobject_t> >*, ghobject_t*)+0xe00) [0x559a9fa2d630]",
    "(BlueStore::collection_list(boost::intrusive_ptr<ObjectStore::CollectionImpl>&, ghobject_t const&, ghobject_t const&, int, std::vector<ghobject_t, std::allocator<ghobject_t> >*, ghobject_t*)+0xc7) [0x559a9fa2e6e7]",
    "(PGBackend::objects_list_partial(hobject_t const&, int, int, std::vector<hobject_t, std::allocator<hobject_t> >*, hobject_t*)+0xc4d) [0x559a9f6f92fd]",
    "(PrimaryLogPG::scan_range(int, int, BackfillInterval*, ThreadPool::TPHandle&)+0xfc) [0x559a9f69443c]",
    "(PrimaryLogPG::do_scan(boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)+0x9a4) [0x559a9f695b34]",
    "(PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x989) [0x559a9f696ae9]",
    "(OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)+0x1bc) [0x559a9f502edc]",
    "(ceph::osd::scheduler::PGOpItem::run(OSD*, OSDShard*, boost::intrusive_ptr<PG>&, ThreadPool::TPHandle&)+0x65) [0x559a9f780005]",
    "(OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x9f8) [0x559a9f52e778]",
    "(ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x41a) [0x559a9fbc90ba]",
    "(ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x559a9fbc420]",
    "/lib/x86_64-linux-gnu/libpthread.so.0(+0x8ea7) [0x7f213e8a1ea7]",
    "clone()"
],
"ceph_version": "16.2.5",
"crash_id": "2021-07-26T01:16:26.865479Z_9034324f-d8c4-4bec-a35e-d890df3f2a3e",
"entity_name": "osd.69aa0d3c7e2c0923efff3fcee8fa629e59041fd",
"os_id": "11",
"os_name": "Debian GNU/Linux 11 (bullseye)",
"os_version": "11 (bullseye)",
"os_version_id": "11",
"process_name": "ceph-osd",
"stack_sig": "0ae3ccdd27d47a526bdd327d6213e2260b3ecd9d3c28a499adbe082f1758e51c",
"timestamp": "2021-07-26T01:16:26.865479Z",
"utsname_machine": "x86_64",
"utsname_release": "5.11.22-2-pve",
"utsname_sysname": "Linux",
"utsname_version": "#1 SMP PVE 5.11.22-4 (Tue, 20 Jul 2021 21:40:02 +0200)"
}

```

Related issues:

Duplicated by bluestore - Bug #52228: crash: /lib/x86_64-linux-gnu/libpthread...

Duplicate

History

#1 - 08/11/2021 06:47 PM - Telemetry Bot

- Crash signature (v1) updated
- Crash signature (v2) updated
- Affected Versions v16.2.4, v16.2.5 added

#2 - 08/12/2021 04:15 PM - Sage Weil

- Subject changed from crash: /lib/x86_64-linux-gnu/libpthread.so.0(to crash: /lib/x86_64-linux-gnu/libpthread.so.0(rocksdb pread)

#3 - 08/12/2021 04:15 PM - Sage Weil

- Status changed from New to Triaged

#4 - 08/12/2021 04:16 PM - Sage Weil

- Duplicated by Bug #52228: crash: /lib/x86_64-linux-gnu/libpthread.so.0(added

#5 - 10/07/2021 02:32 PM - Neha Ojha

- Subject changed from crash: /lib/x86_64-linux-gnu/libpthread.so.0(rocksdb pread) to crash in pread64 during collection_list()

#6 - 03/17/2022 12:17 AM - Telemetry Bot

- Crash signature (v1) updated

- Crash signature (v2) updated

#7 - 03/17/2022 03:55 PM - Telemetry Bot

- Crash signature (v1) updated

- Crash signature (v2) updated

#8 - 03/19/2022 01:24 AM - Telemetry Bot

- Crash signature (v1) updated

- Affected Versions v16.2.7 added