

Commit

✓ Fix potential integer overflow in hash container create/resize

[Browse files](#)

The sized constructors, `reserve()`, and `rehash()` methods of `absl::{flat,node}_hash_{set,map}` did not impose an upper bound on their size argument. As a result, it was possible for a caller to pass a very large size that would cause an integer overflow when computing the size of the container's backing store. Subsequent accesses to the container might then access out-of-bounds memory.

The fix is in two parts:

- 1) Update `max_size()` to return the maximum number of items that can be stored in the container
- 2) Validate the size arguments to the constructors, `reserve()`, and `rehash()` methods, and abort the program when the argument is invalid

We've looked at uses of these containers in Google codebases like Chrome, and determined this vulnerability is likely to be difficult to exploit. This is primarily because container sizes are rarely attacker-controlled.

The bug was discovered by Dmitry Vyukov <dvyukov@google.com>.

PiperOrigin-RevId: 718841870

Change-Id: Ic09dc9de140a35dbb45ab9d90f58383cf2de8286

master

20250127.0 ... 20250127.rc1

 **derekmauro** authored and **copybara-github** committed last month 1 parent 5f8d605 commit 5a0e2cb

Showing 3 changed files with 38 additions and 1 deletion.

[Whitespa...](#)

[Ignore whitespace](#)

[Split](#)

[Unifi...](#)

Filter changed files

absl/container/internal

| 24 | 24 | #include "absl/base/config.h" | |
|----|----|---|--|
| 25 | 25 | #include "absl/base/dynamic_annotations.h" | |
| 26 | 26 | #include "absl/base/internal/endian.h" | |
| | 27 | + #include "absl/base/internal/raw_logging.h" | |
| 27 | 28 | #include "absl/base/optimization.h" | |
| 28 | 29 | #include | |
| | | "absl/container/internal/container_memory.h" | |

```

29      30      #include
          "absl/container/internal/hashtablez_sampler.h"
661     662      return target.offset;
662     663      }
663     664      }
664     665 + void HashTableSizeOverflow() {
665 +   ABSL_RAW_LOG(FATAL, "Hash table size overflow");
666 + }
668 +
669 } // namespace container_internal
670 ABSL_NAMESPACE_END
671 } // namespace absl

```

▼ ⏪ 26 [absl/container/internal/raw_hash_set.h](#)

```

1226    1226      // Given the capacity of a table, computes the
1227    1227      total size of the backing
1228    1228      // array.
1229    1229      size_t alloc_size(size_t slot_size) const {
1230    1230 +   ABSL_SWISSTABLE_ASSERT(
1231    1231 +     slot_size <=
1232    1232 +     ((std::numeric_limits<size_t>::max)() -
1233    1233         slot_offset_) / capacity_);
1234    1234         return slot_offset_ + capacity_ * slot_size;
1235    1235     }
1236    1236
1237    1237     return n ? ~size_t{} >> countl_zero(n) : 1;
1238    1238 }
1239    1239
1240    1240     + template <size_t kSlotSize>
1241    1241     + size_t MaxValidCapacity() {
1242    1242     +   return
1243    1243       NormalizeCapacity((std::numeric_limits<size_t>::ma-
1244    1244           x)() / 4 /
1245    1245           kSlotSize);
1246    1246     + }
1247    1247     +
1248    1248     + // Use a non-inlined function to avoid code bloat.
1249    1249     + [[noreturn]] void HashTableSizeOverflow();
1250    1250     +
1251    1251     // General notes on capacity/growth methods below:
1252    1252     // - We use 7/8th as maximum load factor. For 16-
1253    1253         wide groups, that gives an
1254    1254         // average of two empty slots per group.
1255    1255         :
1256    1256         settings_(CommonFields::CreateDefault<SooEnabled()
1257    1257             >(), hash, eq,
1258    1258                 alloc) {
1259    1259                 if (bucket_count > DefaultCapacity()) {
1260    1260 +                   if (ABSL_PREDICT_FALSE(bucket_count >
1261    1261 +                     MaxValidCapacity<sizeof(slot_type)>())) {
1262    1262 +                       HashTableSizeOverflow();

```

```

2663    +     }
2648    2664         resize(NormalizeCapacity(bucket_count));
2649    2665     }
2650    2666 }
2917    2933     ABSL_ASSUME(cap >= kDefaultCapacity);
2918    2934         return cap;
2919    2935     }
2920    - size_t max_size() const { return
2921    2936         (std::numeric_limits<size_t>::max()); }
2922    2937 + size_t max_size() const {
2923    2938 +     return
2924    2939     CapacityToGrowth(MaxValidCapacity<sizeof(slot_type
2925    2940     )>());
2926    2941 + }
2927    2942
2928    2943     ABSL_ATTRIBUTE_REINITIALIZES void clear() {
2929    2944         if (SwisstableGenerationsEnabled()) &&
2930    2945             auto m = NormalizeCapacity(n |
2931    2946                 GrowthToLowerboundCapacity(size()));
2932    2947             // n == 0 unconditionally rehashes as per the
2933    2948             // standard.
2934    2949             if (n == 0 || m > cap) {
2935    2950 +             if (ABSL_PREDICT_FALSE(m >
2936    2951                 MaxValidCapacity<sizeof(slot_type)>())) {
2937    2952 +                 HashTableSizeOverflow();
2938    2953 +             }
2939    2954         resize(m);
2940    2955
2941    2956         // This is after resize, to ensure that we
2942    2957         // have completed the allocation
2943    2958         const size_t max_size_before_growth =
2944    2959             is_soo() ? SooCapacity() : size() +
2945                 growth_left();
2946    2960             if (n > max_size_before_growth) {
2947    2961 +                 if (ABSL_PREDICT_FALSE(n > max_size())) {
2948    2962 +                     HashTableSizeOverflow();
2949    2963 +                 }
2950    2964             size_t m = GrowthToLowerboundCapacity(n);
2951    2965             resize(NormalizeCapacity(m));
2952    2966
2953    2967

```

▼ ⇕ 8  abs/containers/internal/raw_hash_set_test.cc 

```

3737    3737     }
3738    3738 }
3739    3739
3740    3740 + TEST(Table, MaxSizeOverflow) {
3741    3741 +     size_t overflow =
3742    3742         (std::numeric_limits<size_t>::max()) ;
3743    3743 +     EXPECT_DEATH_IF_SUPPORTED(IntTable t(overflow),
3744    3744         "Hash table size overflow");
3745    3745 +     IntTable t;

```

```
3744 + EXPECT_DEATH_IF_SUPPORTED(t.reserve(overflow),
  "Hash table size overflow");
3745 + EXPECT_DEATH_IF_SUPPORTED(t.rehash(overflow),
  "Hash table size overflow");
3746 + }
3747 +
3748 } // namespace
3749 } // namespace container_internal
3742 3750 ABSL_NAMESPACE_END
```

0 comments on commit 5a0e2cb

Please [sign in](#) to comment.