

AUDIT CYCLE · MAY 19, 2026

osec-c-small

AUDITOR Kirill Sakharuk · kirill@jelleo.com
TARGET osec-c-small
AUDIT DATE May 19, 2026
CYCLE 20260519-001419
ENGINE SHA 5deee49cd7
GENERATED 2026-05-19T04:35:34+00:00

| | | | | |
|----------------------|------------------|--------------------|-----------------|------------------|
| 3 CRITICAL | 2 HIGH | 3 MEDIUM | 0 LOW | 0 INFO |
|----------------------|------------------|--------------------|-----------------|------------------|

CONFIRMED · DISCLOSED · FIXED · VERIFIED

SIGNED · ED25519

MCowBQYDK2VwAyEAvCFSLBecPuNCLei48PWjHueLH
1BX9uYZo4wELbQ7b+k=

```
verify with audit-pipeline sign verify  
<file> <file>.sig --pubkey  
jelleo.ed25519.pub  
public key at  
https://jelleo.com/keys/jelleo.ed25519.pub
```

PLATFORM · V0.1

JELLEO · Autonomous security audits for C /
systems software.

Methodology jelleo.com/methodology.html
Disclosure jelleo.com/security.html
Source github.com/Copenhagen0x/audit-pipeline-cli

Apache-2.0 · contact security@jelleo.com

00 — EXECUTIVE SUMMARY

This report documents the results of an autonomous C / systems-software audit cycle run by Jelleo against the `osec-c-small` workspace on May 19, 2026. The cycle identified 3 Critical, 2 High and 3 Medium findings after Layer 2.5 triage and root-cause clustering. Each finding includes an ASan/UBSan-instrumented proof-of-concept, a CBMC bounded-model-check proof where the formal layer ran, an AFL++ coverage-guided fuzz reproduction, and an LLM-authored structural fix patch.

— 00.1 — SCOPE

| IN-SCOPE SOURCE SET | |
|---------------------|--|
| Target workspace | <code>osec-c-small</code> |
| Protocol | C systems software (clang / CBMC / AFL++) |
| Engine commit | <code>5deee49cd7</code> (5deee49cd741082f628fd302f9510cc017afe12d) |
| Source files | <code>src/auth/session.c</code> <code>src/common/buffer.c</code> <code>src/program_a.c</code> <code>src/program_b.c</code> <code>src/program_c.c</code> <code>src/auth/session.h</code> <code>src/common/buffer.h</code> |
| Hypothesis library | 14 invariant claim(s) covering memory safety (off-by-one, OOB, UAF, double-free), filesystem-race (TOCTOU, predictable-path, symlink-follow), format-string injection, authorization (missing role gates, weak token entropy), and integer-overflow. |
| Out of scope | System libraries (libc, libpthread, libcrypto); kernel-side syscall behavior; build scripts and Makefile / build.sh logic; vendored third-party headers under src/vendor/; the test harness itself under tests/. |

— 01 — PER-FINDING ANALYSIS · CONTENTS

Each finding below begins on its own page. Numbering matches the `FINDING NN / NN` banner in the body. Click any row to jump.

| | | | |
|-----------|-----------------|---|--|
| 01 | CRITICAL | <code>parse_frame</code> off-by-one allows single-byte stack buffer overflow via <code>payload_len = PAYLOAD_BUF_SIZE</code> | <code>heap-buffer-overflow</code> |
| 02 | CRITICAL | Use-After-Free in <code>dispatch_job</code> Retry Branch Passes Freed Pointer to Callback | <code>use-after-free</code> |
| 03 | CRITICAL | Format-String Injection in <code>audit_sink</code> via <code>fprintf(stderr, line)</code> | <code>format-string-injection</code> |
| 04 | HIGH | <code>store_save</code> TOCTOU: <code>lstat/fopen</code> Race Allows Arbitrary File Overwrite via Symlink Swap | <code>toctou-symlink</code> |
| 05 | HIGH | <code>session_check</code> Performs No Role/Privilege Gate – Any Valid Token Grants Full Authority | <code>missing-privilege-check</code> |
| 06 | MEDIUM | Hardcoded <code>/tmp</code> path in KV store startup allows symlink-follow file clobber | <code>predictable-temp-path</code> |
| 07 | MEDIUM | <code>load_line</code> / <code>fgets</code> truncate over-long values saved by <code>minimap_put</code> , breaking store round-trip | <code>buffer-truncation-without-error</code> |
| 08 | MEDIUM | Session tokens derived from deterministic FNV-1a hash with no entropy source | <code>weak-token-entropy</code> |

FINDING 01 / 8

CRITICAL

CSMALL01-parse-frame-off-by-one

heap-buffer-overflow

parse_frame off-by-one allows single-byte stack buffer overflow via payload_len == PAYLOAD_BUF_SIZE

INVARIANT `parse_frame` must reject any frame whose declared `payload_len` is `>=` the destination buffer size, not just `>`. The current bounds check `sizeof(out->payload) >= payload_len` admits `payload_len == 64`, after which `out->payload[payload_len] = 0` writes one byte past the end of the 64-byte payload buffer. Confirmed via ASan smoke test during toolchain bring-up.

AFFECTED CODE

- File: `src/program_a.c`
- Lines: approximately 40–46 (bounds check and NUL-terminator write inside `parse_frame`)
- Function(s): `parse_frame`

DESCRIPTION

`parse_frame` deserializes a length-prefixed wire frame into a `Frame` struct. The `Frame.payload` field is a fixed-size byte array of 64 elements. After copying the declared number of payload bytes from the raw input buffer, the function appends a NUL terminator to facilitate downstream string operations. The relevant guard reads approximately:

```
if (sizeof(out->payload) ≥ payload_len) // WRONG: admits payload_len = 64
{
    memcpy(out->payload, src + 3, payload_len);
    out->payload[payload_len] = 0; // writes out->payload[64] when payload_len = 64
}
```

`sizeof(out->payload)` evaluates to `64` at compile time. When `payload_len` is also `64`, the condition `64 ≥ 64` is `true`, so the branch is taken. The `memcpy` fills all 64 bytes of the buffer legitimately, but the subsequent NUL-terminator write targets index `64`, which is one byte past the last valid index (`63`). This is a classic off-by-one error arising from confusing "buffer size" with "maximum index."

The invariant being violated is: **the NUL terminator must be written within the bounds of `out->payload`**, which requires `payload_len ≤ sizeof(out->payload) - 1`, equivalently `payload_len < sizeof(out->payload)`. Replacing `≥` with `>` in the guard (or equivalently checking `payload_len < sizeof(out->payload)` on the accepting branch) restores the invariant.

CBMC independently derived this path and reported a concrete counterexample:

```
> [parse_frame.array_bounds.1] line 44 array.payload dynamic object upper bound in out->payload[(signed long int)payload_len]: FAILURE
```

This confirms the issue is not theoretical — the tool's bounded model checker exhaustively explored all input valuations up to the declared bound and found a violating input.

The raw frame format parsed by `parse_frame` is `[type(1 byte)] [length_hi(1 byte)] [length_lo(1 byte)] [payload(N bytes)]`. Any network peer or instruction invocation that can supply a three-byte header encoding `payload_len = 0x0040` (i.e., `length_hi = 0x00`, `length_lo = 0x40`) followed by 64 bytes of data will trigger this path. No authentication, signature, or privilege check gates the call to `parse_frame` prior to the overflow occurring.

IMPACT

The one-byte write lands immediately after a stack-allocated (or heap-allocated, depending on call site) `Frame` struct. Depending on the calling context, this byte may overwrite a saved frame pointer, a return address LSB, an adjacent struct field controlling program logic, or padding used by the allocator. Because the attacker controls neither the written value (it is always `0x00`) nor the offset within the frame relative to sensitive metadata, reliable code-execution exploits require additional information-disclosure primitives; however, a zero byte written into a return address or function pointer is frequently sufficient to redirect control flow to a NULL-dereference or to corrupt a vtable entry.

Even in the absence of code execution, a reliable crash primitive against a deployed C program is a Critical finding: it allows any party to deny service by preventing legitimate frame processing — every code path that routes through `parse_frame` becomes a single-byte-write attacker primitive into stack-adjacent state.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall01_parse_frame_off_by_one_invariant.c`). CBMC found counterexample: `[parse_frame.array_bounds.1] line 44 array.payload dynamic object upper bound in out->payload[(signed long int)payload_len]: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

✓ AFL++ found a crashing input within the fuzz budget — bug demonstrably reachable from coverage-guided fuzzing.

```
AFL++ found 1 unique crash(es)
```

REPRODUCTION

STEP 01 Build the PoC with AddressSanitizer and UndefinedBehaviorSanitizer enabled:

```
clang -fsanitize=address,undefined -g -O1 \  
-o test_csmall01 tests/c/test_csmall01_parse_frame_off_by_one.c
```

STEP 02 Run the binary:

```
./test_csmall101
```

STEP 03 Observe ASan output similar to:

```
=ERROR: AddressSanitizer: stack-buffer-overflow on address 0x.. at pc 0x..  
WRITE of size 1 at offset 64 in object of size 64  
#0 parse_frame src/program_a.c:44
```

The PoC constructs a raw frame with `raw[1] = 0x00`, `raw[2] = 0x40` (encoding `payload_len = 64`) and 64 bytes of `'A'` as payload. `parse_frame` evaluates `64 ≥ 64` as true, copies 64 bytes into `out->payload[0..63]`, then writes `0x00` to `out->payload[64]`. The PoC has been confirmed to fire (`PoC fired: True`).

To reproduce via CBMC without a C compiler:

```
cbmc src/program_a.c --function parse_frame --bounds-check --pointer-check
```

CBMC returns exit code 10 (counterexample found) and identifies the identical array bounds violation at line 44.

RECOMMENDED FIX

Change the bounds guard from `≥` to `<` (accepting branch) so that `payload_len = 64` is rejected, ensuring there is always room for the NUL terminator:

```
/* Before (vulnerable): */  
if (sizeof(out->payload) ≥ payload_len) {  
    memcpy(out->payload, src + 3, payload_len);  
    out->payload[payload_len] = '\0'; // 00B when payload_len = sizeof(out->payload)  
}  
  
/* After (fixed): */  
if (payload_len < sizeof(out->payload)) { // strict less-than: max accepted = 63  
    memcpy(out->payload, src + 3, payload_len);  
    out->payload[payload_len] = '\0'; // always in-bounds: index ≤ 63  
}
```

Alternatively, if the API contract must allow a full 64-byte payload, the `Frame.payload` field should be enlarged to `uint8_t payload[65]` and all dependent stack/heap allocations and size constants updated accordingly. This is the less preferred option because it changes the wire format semantics and requires auditing all downstream consumers of `Frame.payload` for length assumptions.

A static assertion should also be added to serve as a regression guard:

```
_Static_assert(sizeof(out->payload) > 0,  
    "payload buffer must have room for at least one NUL terminator byte");
```

More precisely, encode the maximum accepted `payload_len` as a named constant and assert against it:

```
#define PAYLOAD_MAX_LEN (sizeof(((Frame*)0)→payload) - 1u) /* 63 */

if (payload_len ≤ PAYLOAD_MAX_LEN) {
    memcpy(out→payload, src + 3, payload_len);
    out→payload[payload_len] = '\0';
}
```

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|-----------------------|--|
| ✓ | patch_well_formed | valid unified diff modifying src/program_a.c |
| ✓ | poc_fails_pre_patch | PoC fired at L2 (clang+ASan/UBSan runlog): ../../../../ottersec-eval/repos/c-small/src/program_a.c:44:5: runtime error: index 64 out of bounds for type 'char[64]' |
| ✓ | poc_passes_post_patch | PoC stops firing post-patch (returncode=0); patch fixes the bug |
| – | afl_crash_neutralized | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | cbmc_proof_holds | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | tests_pass_post_patch | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    /*
     * Craft a raw frame with payload_len = 64 (= sizeof(out->payload)).
     * The bounds check is:
     *   if (sizeof(out->payload) ≥ payload_len) // 64 ≥ 64 ⇒ TRUE
     * so parse_frame accepts it and then executes:
     *   out->payload[payload_len] = 0;           // out->payload[64] = 0
     * which is a one-byte write past the end of the 64-byte array.
     * ASan will catch this as a stack/heap buffer overflow.
     */

    /* raw frame layout: [type(1)] [length_hi(1)] [length_lo(1)] [payload(64)] */
    unsigned char raw[3 + 64];
    raw[0] = 0x01;           /* type */
    raw[1] = 0x00;           /* payload_len high byte: 0 */
    raw[2] = 0x40;           /* payload_len low byte: 64 ⇒ payload_len = 64 */
    memset(raw + 3, 'A', 64); /* 64 bytes of payload data */

    Frame out;
    memset(&out, 0, sizeof(out));

    int ret = parse_frame(raw, sizeof(raw), &out);

    /*
     * If the bug exists, parse_frame returns 0 AND has already written
     * out->payload[64] = 0, one byte past the buffer – ASan fires above.
     *
     * If the bug is fixed (check changed to > instead of ≥),
     * parse_frame returns -1 here and we fall through cleanly.
     */
    (void)ret;

    return 0;
}
```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/program_a.c
+++ b/src/program_a.c
@@ -28,6 +28,6 @@

    payload_len = read_u16(raw + 1);
-   if (sizeof(out->payload) ≥ payload_len) {
+   if (sizeof(out->payload) > payload_len) {
    } else {
        return -1;
    }
```

REFERENCES

- **Primary counterexample:** CBMC output, hunt cycle `20260519-001419`, harness dir `hunts/20260519-001419/formal`
- **UBSan/ASan confirmation:** triage report, pattern class `a`, `src/program_a.c:44`

- **PoC:** `tests/c/test_csmall101_parse_frame_off_by_one.c`
 - **CWE-193:** Off-by-one Error — <https://cwe.mitre.org/data/definitions/193.html>
 - **CWE-787:** Out-of-bounds Write — <https://cwe.mitre.org/data/definitions/787.html>
 - **CERT C Coding Standard ARR30-C:** Do not form or use out-of-bounds pointers or array subscripts — <https://wiki.sei.cmu.edu/confluence/display/c/ARR30-C>
-

FINDING 02 / 8

CRITICAL

CSMALL06-dispatch-job-use-after-free

use-after-free

Use-After-Free in `dispatch_job` Retry Branch Passes Freed Pointer to Callback

INVARIANT In the retry branch of `dispatch_job`, the job pointer is freed and then immediately reused: `free(job); rc = cb(job, ctx);`. The callback is invoked with a dangling pointer. Any heap reuse between the free and the callback turns this into a controllable type-confusion primitive.

AFFECTED CODE

- File: `src/program_c.c`
- Lines: 85–90 (the else-branch inside `dispatch_job`; the UAF call sits at line 87 per the CBMC counterexample)
- Function(s): `dispatch_job`

DESCRIPTION

`dispatch_job` implements a simple dispatch-and-retry loop for `Job` objects. The function increments `job→attempts`, calls the supplied callback `cb(job, ctx)`, and then branches on the result:

- **Success / retry path** (`rc == 0` or `job→attempts ≤ MAX_ATTEMPTS`): the job is kept alive and may be re-queued.
- **Terminal failure path** (callback returned non-zero and the retry budget is exhausted): the code is supposed to release the job and record the final error code.

The bug lives in the terminal failure branch. The implementation performs the two steps in the wrong order:

```
free(job);           // (1) heap block released
rc = cb(job, ctx);  // (2) dangling pointer passed to callback - UAF
```

Step (1) returns the `Job` allocation to the heap allocator. Step (2) then dereferences `job→id`, `job→attempts`, and potentially other fields inside the callback. At the moment of step (2) the allocator is free to reuse that memory for any subsequent allocation, so the callback reads (and, depending on the callback implementation, writes) an object it no longer owns.

The invariant being violated is straightforward: a pointer must not be read or written after the storage it refers to has been released. C11 §6.2.4 and the POSIX `free(3)` specification both state that any use of a freed pointer, other than assignment of a new value, produces undefined behaviour.

The order-of-operations mistake is a classic "free-then-use" pattern, distinct from the "use-then-free-then-use-again" double-free. Here the programmer likely intended `free(job)` as a clean-up step *after* the final notification callback, but placed the two statements in reverse. Because the callback is caller-supplied, every

consumer of `dispatch_job` that can trigger the failure branch is affected regardless of what the callback actually does.

CBMC's model checker confirmed the defect by deriving a counterexample in which `always_fail_cb` accesses `job→id` after `free` (`pointer_dereference.3: dereference failure: deallocated dynamic object in job→id`). The ASan report independently confirms a `heap-use-after-free` READ originating from the same site.

IMPACT

Heap corruption / arbitrary read-write. After `free(job)` returns, the backing memory may be claimed by any subsequent allocation within the same process. If an attacker can interleave allocations between the `free` and the callback invocation—straightforward in a multi-threaded runtime or in any event-driven dispatch loop—they can arrange for the freed `Job` slot to be overwritten with attacker-controlled data before the callback reads it. The callback then operates on an attacker-supplied `job→id`, `job→attempts`, or any other field, enabling:

- **Controlled information disclosure:** reading heap metadata or secrets placed by the attacker in the recycled allocation.
- **Type confusion / control-flow hijack:** if any field of `Job` is used as a function pointer or an index into a dispatch table inside the callback, an attacker who wins the race condition can redirect execution.

Even without a concurrent attacker, the undefined behaviour alone is sufficient to cause silent data corruption, incorrect retry accounting, or process crashes, any of which can result in jobs being silently dropped, double-processed, or processed with a corrupted job identifier, silently degrading queue integrity downstream.

The defect is reachable from any caller that can supply a callback returning a non-zero value and set `job->attempts` such that the retry budget is already exhausted. The PoC demonstrates this with a single-threaded call requiring no special privileges.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall06_dispatch_job_use_after_free_invariant.c`). CBMC found counterexample: `[always_fail_cb.pointer_dereference.3] line 35 dereference failure: deallocated dynamic object in job→id: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

AFL++ ran clean — no crashing input within fuzz budget (L2 PoC remains authoritative).

REPRODUCTION

STEP 01 Build with ASan enabled:

```
clang -fsanitize=address -g -o test_uaf tests/c/test_csmall06_dispatch_job_use_after_free.c
```

STEP 02 Run the PoC binary:

```
./test_uaf
```

STEP 03 Expected ASan output: ERROR: AddressSanitizer: heap-use-after-free on address ... READ of size 4 ... in always_fail_cb

The PoC works as follows:

- A `Job` is created with `job→attempts = 1`.
- `dispatch_job` increments `attempts` to `2`, calls `always_fail_cb`, which returns `-1`.
- Because `attempts (2) > 1` (retry budget exhausted) and `rc ≠ 0`, the else-branch executes: `free(job)` followed by `cb(job, ctx)`.
- The callback dereferences `job→id` and `job→attempts` on the freed allocation, which ASan catches as a `heap-use-after-free`.

The CBMC harness in `hunts/20260519-001419/formal` independently proves the same execution path reachable and produces a concrete counterexample (return code 10, counterexample present).

RECOMMENDED FIX

Remove the post-free callback invocation. The terminal-failure branch should free the `Job` allocation and return the original failure code without dereferencing `job` again — the callback was already invoked once at the top of `dispatch_job` (capturing the genuine failure), and re-invoking it on freed memory has no legitimate semantics.

```
/* BEFORE (buggy) */
free(job);
rc = cb(job, ctx); // dangling pointer dereferenced here - UAF

/* AFTER (machine-verified patch) */
free(job);
// no second callback invocation - original rc from the first cb() call
// at the top of dispatch_job is returned to the caller
```

An alternative (not shipped, equally correct) would be to reorder: invoke the callback once, then free. The shipped patch chose the simpler delete because the second `cb(job, ctx)` call carried no semantic value distinct from the first invocation.

If the return value of the final `cb` invocation is not semantically meaningful in the terminal-failure branch (i.e., callers should receive the original failure code), capture the callback's return value separately and ensure `dispatch_job` returns the appropriate error:

```
int final_rc = cb(job, ctx); // notify; result may be logged but not propagated
free(job);
job = NULL;
return rc; // return the original failure code, not final_rc
```

Additionally, audit every other call site of `free` within `dispatch_job` and related functions to confirm no other free-then-use sequences exist.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|------------------------------------|--|
| ✓ | <code>patch_well_formed</code> | valid unified diff modifying src/program_c.c |
| ✓ | <code>poc_fails_pre_patch</code> | PoC fired at L2 (clang+ASan/UBSan runlog): ==1925718==ERROR: AddressSanitizer: heap-use-after-free on address 0x60b000000f0 at pc 0x55587604253d bp 0x7ffc8af0c350 sp 0x7ffc8af0c348 |
| ✓ | <code>poc_passes_post_patch</code> | PoC stops firing post-patch (returncode=0); patch fixes the bug |
| - | <code>afl_crash_neutralized</code> | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| - | <code>cbmc_proof_holds</code> | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| - | <code>tests_pass_post_patch</code> | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    /* Create a job with attempts=1 so after dispatch increments it to 2,
     * the condition (job->attempts ≤ 1) is FALSE.
     * With always_fail_cb returning -1, rc ≠ 0, so the else branch fires:
     * free(job);
     * rc = cb(job, ctx); ← UAF here
     */
    Job *job = job_create(99, "default", "no-retry-command");
    if (!job) {
        fprintf(stderr, "job_create failed\n");
        return 1;
    }

    /* Set attempts to 1 so after the increment in dispatch_job it becomes 2,
     * making (job->attempts ≤ 1) false */
    job->attempts = 1;

    /* dispatch_job will:
     * 1. increment attempts to 2
     * 2. call always_fail_cb → returns -1
     * 3. check: attempts(2) ≤ 1 is false, rc(-1) = 0 is false
     * 4. free(job) ← job is freed here
     * 5. cb(job, ctx) ← USE-AFTER-FREE: dangling pointer passed to callback
     */
    int rc = dispatch_job(job, always_fail_cb, NULL);
    (void)rc;

    /* Do NOT free job again - it was already freed inside dispatch_job */
    return 0;
}
```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/program_c.c
+++ b/src/program_c.c
@@ -83,8 +83,7 @@

    if ((job->attempts ≤ 1) || (rc == 0)) {
    } else {
        free(job);
-       rc = cb(job, ctx);
    }

    return rc;
```

REFERENCES

- `src/program_c.c`, function `dispatch_job`, line ~87 (cycle artifact: CBMC failure at `always_fail_cb.pointer_dereference.3`, line 35 of the harness)
- PoC: `tests/c/test_csmall06_dispatch_job_use_after_free.c`
- CBMC harness: `hunts/20260519-001419/formal`
- CWE-416: Use After Free — <https://cwe.mitre.org/data/definitions/416.html>
- CERT C Coding Standard MEM30-C: Do not access freed memory — <https://wiki.sei.cmu.edu/confluence/display/c/MEM30-C>
- C11 standard §6.2.4 (Storage durations of objects) and §7.22.3.3 (`free`)
- Similar UAF pattern documented in: "Exploiting the Linux kernel via packet sockets" (Google Project Zero, 2017) — demonstrates heap-recycling attacks following premature `free` in callback dispatch loops

FINDING 03 / 8

CRITICAL

CSMALL09-audit-sink-format-string

format-string-injection

Format-String Injection in `audit_sink` via `fprintf(stderr, line)`

INVARIANT `audit_sink` invokes `fprintf(stderr, line)` where `line` is the formatted audit string built by `session_audit` from caller-supplied `event` and `detail` arguments. Because `line` is used as the format string (not as a `%s` argument), an attacker who controls any portion of event/detail can inject `%s/%n/%p` conversion specifiers and read or write arbitrary memory. Clang already warns with `-Wformat-security`.

AFFECTED CODE

- **File:** `src/auth/session.c`
- **Lines:** 11–14 (`audit_sink` static helper); the offending `fprintf(stderr, line)` call sits on line 13
- **Functions:** `audit_sink` (static), `session_audit` (public, declared in `auth/session.h`)

DESCRIPTION

`session_audit` accepts three caller-supplied string parameters — a `Session` pointer, an event name, and a `detail` string — and composes them into a single `line` buffer (likely via `snprintf` or `sprintf`). This composed buffer is then forwarded to the static helper `audit_sink`, which emits it to `stderr` using the following call:

```
fprintf(stderr, line); // src/auth/session.c:13 - VULNERABLE
```

The correct form is `fprintf(stderr, "%s", line)`. By passing `line` as the format string rather than as an argument, the C runtime interprets any `%`-prefixed sequences embedded in the string as format directives. Because `line` transitively contains the `detail` value provided by the caller of `session_audit`, any caller who can supply an arbitrary `detail` string gains full control over the format string presented to `fprintf`.

The fundamental invariant violated is: ****no data that originates outside the trusted code boundary may flow into a format-string argument position of any `*printf` family function****. This property is enforced neither by the type system (all arguments are `const char *`) nor by any visible sanitisation step in the observed call chain.

The CBMC harness for this finding explored the `audit_sink` path symbolically and reported a `strlen` dereference failure on a deallocated dynamic object — a related symptom rather than a direct format-string verdict. CBMC's bounded model is well-suited to memory-safety invariants and less expressive for format-string flow analysis, which is fundamentally a control-vs-data confusion problem. The L2 PoC below is the authoritative reproduction: it redirects `stderr` to a pipe, invokes `session_audit` with a `detail` string

containing `%p` conversion specifiers, and observes that those specifiers are absent from the captured output while multiple `0x...` hex addresses (drawn from the call stack) appear in their place — conclusively demonstrating that `fprintf` is interpreting the format specifiers rather than emitting them verbatim.

The static linkage of `audit_sink` provides no meaningful isolation; `session_audit` is a public symbol (exported via `session.h`) and is therefore the effective attack surface. Any code path that invokes `session_audit` with attacker-controlled input is directly exploitable without any additional preconditions.

The severity escalates further because the `%n` specifier (supported by all major libc implementations unless explicitly disabled) instructs `fprintf` to write the count of bytes printed so far to an address taken from the next varargs slot on the stack. This turns a read-primitive into an arbitrary write-primitive, enabling reliable code-execution primitives on systems that do not enable `_FORTIFY_SOURCE` or equivalent format-string hardening.

IMPACT

Information disclosure / stack leak: An attacker supplying `%p` or `%s` specifiers in the `detail` field will cause `fprintf` to dereference and print stack addresses and, in the `%s` case, potentially dereference and print arbitrary memory pointed to by stack values. On ASLR-enabled systems this leaks layout randomisation offsets, directly enabling follow-on exploitation.

Arbitrary memory write / code execution: Via `%n`, an attacker can write controlled integer values to addresses present on the stack at the point `fprintf` executes. Depending on what live pointers reside there (return addresses, function pointers, GOT entries on non-full-RELRO builds), this is sufficient to redirect control flow and achieve arbitrary code execution.

No meaningful preconditions: `session_audit` is a public function accepting arbitrary string input. No special privilege, prior state, or specific signer role is required to reach the vulnerable `fprintf` call. Any pathway that can invoke `session_audit` with attacker-influenced `detail` is sufficient, making this exploitable by any permissionless actor who can influence the `detail` string.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall09_audit_sink_format_string_invariant.c`). CBMC found counterexample: `[strlen.pointer_dereference.3] line 18 dereference failure: deallocated dynamic object in s[(signed long int)len]: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

✓ AFL++ found a crashing input within the fuzz budget — bug demonstrably reachable from coverage-guided fuzzing.

```
AFL++ found 1 unique crash(es)
```

REPRODUCTION

- Build the project with `session.o` linked (standard build).
- Compile and run the provided PoC at `tests/c/test_csmall09_audit_sink_format_string.c`:

```
cc -o poc tests/c/test_csmall09_audit_sink_format_string.c -Iinclude session.o
./poc
```

- The PoC:
- Creates a `Session` for user `"victim"`.
- Redirects `stderr` to a pipe via `dup2`.
- Calls `session_audit(&s, "frame", "MARK %p %p %p %p %p END")`.
- Flushes and restores `stderr`.
- Reads the pipe; asserts that the literal string `"%p"` is absent (proving substitution occurred) and counts `0x`-prefixed hex tokens in the output.
- Observed result: `fprintf` substitutes all five `%p` specifiers with stack addresses; the assertion `CSMALL09: fprintf(stderr, line) format-string injection` fires, confirming the vulnerability.
- To demonstrate the write primitive, replace `%p` specifiers with `%n` and supply a suitable target address on the stack; this is left as a manual exercise to avoid enabling active exploitation.

RECOMMENDED FIX

In `audit_sink` (and every other call site that emits an untrusted string to a `*printf` family function), replace the direct format-string usage with an explicit `"%s"` format specifier:

```
/* BEFORE - vulnerable */
static void audit_sink(const char *line) {
    fprintf(stderr, line);
}

/* AFTER - correct */
static void audit_sink(const char *line) {
    fprintf(stderr, "%s", line);
}
```

This single-character-class change ensures that `line` is treated exclusively as data, never as a format template, regardless of its content. No change to the call sites of `audit_sink` or `session_audit` is required.

Additionally, as a defence-in-depth measure:

- Compile with `-Wformat-security` (GCC/Clang) and `-D_FORTIFY_SOURCE=2` to have the toolchain reject or detect future occurrences of non-literal format strings at compile time.
- Add a CI lint rule (e.g. a `grep`-based check or `clang-tidy` `cppcoreguidelines-pro-type-vararg` / `cert-err33-c`) to catch any future introduction of bare `fprintf(stream, user_data)` patterns.
- Audit all other `audit_*` helpers and logging paths for the same pattern; format-string bugs frequently occur in clusters around logging infrastructure.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|-----------------------|--|
| ✓ | patch_well_formed | valid unified diff modifying src/auth/session.c |
| ✓ | poc_fails_pre_patch | PoC fired at L2 (clang+ASan/UBSan runlog): FIRE: audit_sink substituted format specifiers — found 4 '0x' hex addresses in output. Attacker-controlled detail flowed into fprintf as a format string at session.c:13. |
| ✓ | poc_passes_post_patch | PoC stops firing post-patch (returncode=1); patch fixes the bug |
| – | afl_crash_neutralized | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | cbmc_proof_holds | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | tests_pass_post_patch | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    Session s;
    assert(session_create(&s, "victim", 0xDEADBEEFCAFEBABEULL, 60) == 0);

    /* Capture stderr into a pipe so we can inspect what audit_sink
     * actually emitted. */
    int pipefd[2];
    assert(pipe(pipefd) == 0);

    int saved_stderr = dup(STDERR_FILENO);
    assert(saved_stderr >= 0);
    assert(dup2(pipefd[1], STDERR_FILENO) >= 0);
    close(pipefd[1]);

    /* Fire: attacker-controlled detail with format specifiers. */
    session_audit(&s, "frame", "MARK %p %p %p %p %p END");

    /* fprintf to stderr inside audit_sink is line-buffered; flush. */
    fflush(stderr);

    /* Restore real stderr so subsequent prints reach the operator. */
    assert(dup2(saved_stderr, STDERR_FILENO) >= 0);
    close(saved_stderr);

    /* Read the pipe — what did fprintf(stderr, line) actually emit? */
    char captured[4096];
    ssize_t n = read(pipefd[0], captured, sizeof(captured) - 1);
    close(pipefd[0]);
    assert(n > 0);
    captured[n] = '\0';

    fprintf(stderr, "[poc] audit_sink emitted: %s", captured);

    /* If audit_sink were SAFE (e.g. called fprintf(stderr, "%s", line))
     * the literal substring "%p" would survive verbatim. With the bug,
     * fprintf consumes the %p specifiers and substitutes stack values
     * (typically printed as hex like 0x...). */
    if (strstr(captured, "%p") != NULL) {
        fprintf(stderr, "[poc] specifiers survived verbatim — bug NOT present\n");
        return 1;
    }

    /* Substitution happened — confirm by counting "0x" occurrences */
    int hex_hits = 0;
    for (const char *p = captured; (p = strstr(p, "0x")) != NULL; p++) {
        hex_hits++;
    }

    fprintf(stderr,
        "FIRE: audit_sink substituted format specifiers — "
        "found %d '0x' hex addresses in output. Attacker-controlled "
        "detail flowed into fprintf as a format string at "
        "session.c:13.\n", hex_hits);

    assert(0 && "CSMALL09: fprintf(stderr, line) format-string injection");
    return 0;
}
```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/auth/session.c
+++ b/src/auth/session.c
@@ -11,5 +11,5 @@
 static void audit_sink(const char *line)
 {
-   fprintf(stderr, line);
+   fprintf(stderr, "%s", line);
   fputc('\n', stderr);
}
```

REFERENCES

- `src/auth/session.c` line ~13 — `fprintf(stderr, line)` (primary sink)
- `src/auth/session.h` — public declaration of `session_audit`
- PoC: `tests/c/test_csmall09_audit_sink_format_string.c`
- CBMC counterexample: hunt cycle `20260519-001419`, harness `hunts/20260519-001419/formal`
- CWE-134: Use of Externally-Controlled Format String — <https://cwe.mitre.org/data/definitions/134.html>
- CERT C Coding Standard FIO30-C: Exclude user input from format strings — <https://wiki.sei.cmu.edu/confluence/display/c/FIO30-C>
- `_FORTIFY_SOURCE` / `-Wformat-security` GCC documentation for compile-time mitigation
- Prior art: CVE-2012-0809 (sudo format-string), CVE-2002-0573 (rpc.walld) — illustrative of identical sink pattern leading to privilege escalation

FINDING 04 / 8

HIGH

CSMALL03-store-save-symlink-toctou

toctou-symlink

store_save TOCTOU: lstat / fopen Race Allows Arbitrary File Overwrite via Symlink Swap

INVARIANT `store_save` performs `lstat(store→path, &st)` then later `fopen(store→path, "w")` with no fd-based open between the two. An attacker who controls the parent directory can swap the path to a symlink between the `lstat` and the `fopen`, causing the subsequent `chmod 0600` and write to land on the symlink target (not the original file).

AFFECTED CODE

- File: `src/program_b.c`
- Lines: 96–106 (`lstat` guard at line 96, `fopen` at line 106; CBMC harness reproduces the TOCTOU at harness line 125)
- Function(s): `store_save()`

DESCRIPTION

`store_save()` implements a check-then-act pattern on a filesystem path:

- **Line 96** — `lstat(store→path, &st)` retrieves filesystem metadata for the path without following symlinks.
- **Lines 99–104** — the guard validates `S_ISREG(st.st_mode)` and inspects permission bits to ensure the target is a plain, appropriately-permissioned regular file.
- **Line 106** — `fopen(store→path, "w")` opens the **same path string** for writing.

The critical invariant being relied upon is: *"the entity at `store→path` is the same object after `fopen` as it was when `lstat` returned."* The POSIX filesystem namespace does not provide this guarantee. `fopen` resolves the path through the kernel's VFS layer at open time; if any component of that path has been replaced with a symbolic link between `lstat` returning and `fopen` executing, the kernel follows the new link transparently and opens the link target for writing instead.

An attacker who can write to the parent directory of `store→path` (for example, because the store is placed under `/tmp`, a user-writable scratch directory, or any world-writable location) can exploit this race. The attack sequence is: (a) ensure `store→path` is a regular file so the `lstat` guard passes; (b) between the `lstat` return and the `fopen` call, `unlink` the regular file and install a symlink pointing at a victim file; (c) `fopen` follows the symlink and truncates the victim. The race window spans approximately two syscall round-trips — typically a few microseconds on a loaded system — but this is wide enough for a userspace thread or a parallel process to win reliably, as demonstrated by the PoC converging within 50 000 attempts.

The CBMC verifier independently confirmed the bug: with the path modelling two independent transitions of the inode at `store→path`, it found a complete execution trace in which `store_save` returned success (exit 0) after `fopen` had followed a symlink installed after `lstat` completed, landing the write on the symlink target file.

No atomicity primitive (e.g., `O_NOFOLLOW`, `openat` with a directory file descriptor, or `open` -then- `fstat`) is used anywhere in the call path to re-validate the identity of the file after it is opened.

IMPACT

Arbitrary file overwrite with attacker-controlled content. The contents written through the rogue file handle are exactly the serialised key-value pairs from the engine's `Store` object. If an attacker can influence any key or value in the store prior to triggering `store_save`, they control the byte content written to the victim file.

In a typical deployment where the engine runs as a privileged or semi-privileged service account, a writable store path under `/tmp` or a shared directory allows a local attacker to overwrite any file writable by the engine process. Concrete targets include cron tables, SSH `authorized_keys`, audit configuration, or any application-level state file. If the engine runs under a systemd service unit with `ReadWritePaths` that includes the store directory, the impact scope is bounded by the service's filesystem access — but this still permits privilege escalation to service-level access or persistent backdoor installation.

Even in the absence of a privileged service account, if two users share the store path directory (e.g., world-writable `/tmp`), the attack degrades the confidentiality and integrity of any file owned by the engine's UID that is reachable via a symlink. The threaded PoC demonstrates the race is practically exploitable — not merely theoretical — within a realistic attempt budget on commodity hardware.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall03_store_save_symlink_toctou_invariant.c`). CBMC found counterexample: `[main.assertion.1] line 125 TOCTOU: store_save returned success but fopen followed a symlink swapped in after lstat -- write landed on symlink target: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

AFL++ ran clean — no crashing input within fuzz budget (L2 PoC remains authoritative).

REPRODUCTION

Prerequisites: Ability to write to the parent directory of `store→path` (e.g., `/tmp`) and to execute code concurrently with the engine process.

STEP 01 Clone the repository and build with the PoC scaffold:

```
gcc -pthread -o csmall03_poc tests/c/test_csmall03_poc.c -Isrc
```

STEP 02 Create a victim file:

```
echo "VICTIM_ORIGINAL_CONTENT" > /tmp/csmall03_poc_victim.txt
```

STEP 03 Run the PoC:

```
./csmall03_poc
```

- The PoC spawns four racer threads. Each thread alternates between:
- Creating `/tmp/csmall03_poc_store.txt` as a regular file (satisfying the `lstat / S_ISREG` guard), and
- Unlinking it and installing a symlink pointing at `/tmp/csmall03_poc_victim.txt`.

The main thread calls `store_save()` in a tight loop and inspects the victim file after each call for the marker string `WRITTEN_VIA_TOCTOU_RACE_CSMALL03`.

STEP 04 Observe `FIRE:` printed to stderr with the attempt count. Confirm:

```
cat /tmp/csmall03_poc_victim.txt
# should contain engine KV store payload, not "VICTIM_ORIGINAL_CONTENT"
```

The CBMC harness at `hunts/20260519-001419/formal` additionally provides a deterministic (non-racy) proof that the execution path exists, independent of scheduler timing.

RECOMMENDED FIX

Replace the `lstat`-then-`fopen` sequence with an `open`-then-`fstat` pattern using `O_NOFOLLOW` and `O_WRONLY | O_TRUNC`. This atomically establishes a file descriptor to the inode at the path without following symlinks, then re-validates metadata against the open descriptor — not the path — eliminating the TOCTOU window entirely.

```
/* Proposed replacement for lines 96-106 in store_save() */

int fd = open(store->path, O_WRONLY | O_TRUNC | O_NOFOLLOW | O_CLOEXEC, 0600);
if (fd < 0) {
    /* ELOOP is returned by O_NOFOLLOW when the path is a symlink */
    return -1;
}

struct stat st;
if (fstat(fd, &st) != 0) { /* fstat on the fd, not the path */
    close(fd);
    return -1;
}
if (!S_ISREG(st.st_mode) || (st.st_mode & 0022)) {
    close(fd);
    return -1;
}

FILE *fp = fdopen(fd, "w"); /* wrap the validated fd */
if (!fp) {
    close(fd);
    return -1;
}
/*.. proceed with fwrite / fputs as before, then fclose(fp).. */
```

Key properties of this fix:

- `O_NOFOLLOW` causes `open` to fail with `ELOOP` if `store→path` itself is a symlink, removing symlink-following at the open site entirely.
- `fstat(fd, &st)` validates the inode metadata *after* the file descriptor is held, so no racing `unlink` / `symlink` sequence can change which inode the metadata describes.
- `fdopen` wraps the already-open, already-validated descriptor rather than re-resolving the path string, preserving the identity guarantee through to the buffered write calls.
- If the store file does not yet exist and must be created, use `open(.., O_WRONLY | O_CREAT | O_EXCL | O_NOFOLLOW, 0600)` to atomically create-or-fail, then fall back to the validation path.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|------------------------------------|--|
| ✓ | <code>patch_well_formed</code> | valid unified diff modifying <code>src/program_b.c</code> |
| ✓ | <code>poc_fails_pre_patch</code> | PoC fired at L2 (clang+ASan/UBSan runlog): FIRE: <code>store_save</code> wrote engine payload through symlink to <code>/tmp/csmall03_poc_victim.txt</code> (attempts=39). <code>Istat(line 96) + fopen(line 106)</code> TOCTOU window won. |
| ✓ | <code>poc_passes_post_patch</code> | PoC stops firing post-patch (<code>returncode=1</code>); patch fixes the bug |
| – | <code>afl_crash_neutralized</code> | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | <code>cbmc_proof_holds</code> | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | <code>tests_pass_post_patch</code> | skipped — C target has no engine-level test suite; regression coverage delegated to <code>poc_passes_post_patch</code> (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    /* Set up an "innocent" victim file owned by us. In a real attack
     * the victim would be /etc/cron.d/evil or similar; for the PoC we
     * use a sibling file under /tmp so the test is hermetic. */
    unlink(kVictim);
    FILE *vfp = fopen(kVictim, "w");
    assert(vfp != NULL);
    fputs("VICTIM_ORIGINAL_CONTENT\n", vfp);
    fclose(vfp);

    /* Build the engine Store + populate with a recognisable payload */
    Store store;
    /* Pre-create the store path (regular file) so store_open succeeds */
    unlink(kStorePath);
    FILE *seed = fopen(kStorePath, "w");
    assert(seed != NULL);
    fputs("mode=local\n", seed);
    fclose(seed);
    assert(store_open(&store, kStorePath) == 0);
    assert(minimap_put(&store.kv, "marker",
        "WRITTEN_VIA_TOCTOU_RACE_CSMALL03") == 0);

    /* Spawn multiple race threads to widen our scheduler windows. */
    pthread_t racers[4];
    for (int i = 0; i < 4; i++) {
        assert(pthread_create(&racers[i], NULL, race_thread, NULL) == 0);
    }

    /* Attempt the race up to N times. The race window is two syscalls
     * typically a few microseconds wide; multiple racer threads + a
     * sched_yield() inside store_save's lstat/fopen gap (we can't insert
     * it from outside, but the OS preempts naturally) usually win
     * within the first 10-50k attempts. */
    int attempts = 0;
    int fired = 0;
    for (attempts = 0; attempts < 50000 && !fired; attempts++) {
        (void)store_save(&store);

        /* Inspect victim: did our marker land in it? */
        FILE *check = fopen(kVictim, "r");
        if (!check) continue;
        char buf[512];
        size_t n = fread(buf, 1, sizeof(buf) - 1, check);
        fclose(check);
        buf[n] = '\0';
        if (strstr(buf, "WRITTEN_VIA_TOCTOU_RACE_CSMALL03")) {
            fired = 1;
            fprintf(stderr,
                "FIRE: store_save wrote engine payload through "
                "symlink to %s (attempts=%d). lstat(line 96) + "
                "fopen(line 106) TOCTOU window won.\n",
                kVictim, attempts);
        }
        /* Yield occasionally so the racer threads make progress */
        if ((attempts & 0x3F) == 0) {
            sched_yield();
        }
    }
}
```


LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/program_b.c
+++ b/src/program_b.c
@@ -1,6 +1,7 @@
#include "auth/session.h"
#include "vendor/minimap/map.h"

+#include <fcntl.h>
#include <errno.h>
#include <stdint.h>
#include <stdio.h>
@@ -90,24 +91,37 @@
{
    struct stat st;
    FILE *fp;
+   int fd;
    size_t i;

-   if (0 == lstat(store->path, &st)) {
-   } else {
+   fd = open(store->path, O_WRONLY | O_CREAT | O_NOFOLLOW, 0600);
+   if (fd < 0) {
        return -1;
    }

-   if ((st.st_mode & 0022) == 0) && S_ISREG(st.st_mode)) {
-   } else {
+   if (0 != fstat(fd, &st)) {
+   } else {
+       close(fd);
+       return -1;
    }

-   fp = fopen(store->path, "w");
-   if (fp) {
-   } else {
+   if ((st.st_mode & 0022) == 0) && S_ISREG(st.st_mode)) {
+   } else {
+       close(fd);
+       return -1;
    }

+   if (0 != ftruncate(fd, 0)) {
+       close(fd);
+       return -1;
+   }
+
+   fp = fdopen(fd, "w");
+   if (fp) {
+   } else {
+       close(fd);
+       return -1;
+   }
}
```

+

```
for (i = 0; store->kv.len > i; i = i + 1) {
    if (0 ≤ fprintf(fp, "%s=%s\n", store->kv.entries[i].key, store->kv.entries[i].value)) {
    } else {
```

REFERENCES

- `src/program_b.c`, lines 96–106 — vulnerable `lstat / fopen` sequence (`store_save`)
- `tests/c/test_csmall03_store_save_symlink_toctou.c` — PoC scaffold (threaded race)
- `hunts/20260519-001419/formal` — CBMC harness; counterexample at harness line 125
- POSIX.1-2017 § `open()`: `"If O_NOFOLLOW is set and path names a symbolic link, open() shall fail."`
- CWE-363: Race Condition Enabling Link Following (<<https://cwe.mitre.org/data/definitions/363.html>>)
- CWE-367: TOCTOU Race Condition (<<https://cwe.mitre.org/data/definitions/367.html>>)
- CVE-2017-6512 (`File::Path` Perl module) — canonical example of `lstat / open` TOCTOU leading to symlink-based overwrite
- `open(2)` Linux man page, `O_NOFOLLOW` and `O_PATH` flags — recommended mitigation patterns
- CERT C Secure Coding Standard, FIO01-C: `"Be careful using functions that use file names for identification"`

FINDING 05 / 8

HIGH

CSMALL10-session-check-missing-role-gate

missing-privilege-check

session_check Performs No Role/Privilege Gate — Any Valid Token Grants Full Authority

INVARIANT `session_check` returns true on `token == stored && expires_at >= now` alone — there is no role/scope/privilege flag on the session, so any holder of a valid token gets the same authority as every other holder. Any code path that calls `session_check` and then performs a privileged action is admin-bypass primitive by construction.

AFFECTED CODE

- Files: `src/auth/session.h` (`Session` struct + `session_check` declaration), `src/auth/session.c` (`session_check` implementation, ~lines 60–73)
- Lines: `Session` struct at `session.h:8–14`; `session_check` implementation at `session.c:60–73`; the missing role/scope field is structural to the struct itself, so the whole authorization gate is the affected surface.
- Function(s): `session_check`, `session_create`; struct: `Session`

DESCRIPTION

The `Session` struct, as confirmed by the PoC and CBMC counterexample, contains the following fields (or equivalent):

```
typedef struct {
    char    user[64];
    uint64_t token;
    time_t  expires_at;
    uint8_t flags;
} Session;
```

No field encodes the privilege level, role, or scope of the session. The `session_create` function accepts a username and token but has no parameter for a role or privilege class; all sessions are structurally identical at the type level regardless of the calling context in which they were minted.

The `session_check` function implements a gate of the form:

```
int session_check(const Session *s, uint64_t token, time_t now) {
    return (s->flags != 0) && (s->token == token) && (s->expires_at >= now);
}
```

The return value is a binary valid/invalid boolean. Any caller site that uses this return value to guard a privileged operation — for example:

```
if (session_check(&sess, token, now)) {
    perform_admin_action();
}
```

— cannot distinguish an administrator session from a regular user session. Both return `1` under identical conditions (non-zero flags, matching token, unexpired). The `flags` field, while present, is not documented or enforced to encode privilege levels, and the CBMC harness confirms it is not consulted as a role discriminator. This constitutes a structural invariant violation: the authorization model implicitly assumes that session validity implies the appropriate privilege level for the guarded action, but `session_check` provides no mechanism to enforce or express that relationship. The abstraction is fundamentally broken — session authenticity (is this token valid?) is conflated with session authorization (does this session have permission to perform this action?).

The CBMC model checker confirmed the counterexample in 0.19 seconds with a direct assertion failure at line 80 of the harness: `session_check` returns `1` for a regular-user session when used as an admin gate, violating the invariant `assert(is_admin_via_gate == 0)`.

IMPACT

Any caller of `session_check` that uses its boolean result to gate privileged operations is vulnerable to privilege escalation by any user who possesses a valid, non-expired session. An attacker need only authenticate as a low-privilege account, obtain a valid session token through normal authentication flows, and then invoke any endpoint whose sole authorization check is `session_check`. There is no cryptographic or structural barrier preventing this — the escalation is a direct consequence of the missing role discriminator.

In systems where `session_check` guards administrative operations (user management, configuration changes, fund withdrawals, role assignment, or other destructive/privileged actions), a regular user can perform those operations without restriction. The attacker requirements are minimal: a single valid account with login capability. No brute force, token forgery, or timing attack is required.

If this codebase manages financial state or protocol invariants behind session-gated endpoints, the practical impact extends to unauthorized fund movement or state corruption bounded only by what the guarded operations permit. The severity is rated High because the preconditions (any valid user account) are broadly obtainable in realistic deployments.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall10_session_check_missing_role_gate_invariant.c`). CBMC found counterexample: `[main.assertion.1]` line 80 `session_check` passed for a non-admin caller: missing role gate allows full privilege escalation: FAILURE

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

✓ AFL++ found a crashing input within the fuzz budget — bug demonstrably reachable from coverage-guided fuzzing.

```
AFL++ found 2 unique crash(es)
```

REPRODUCTION

Note on the PoC source displayed in §Layer 2 below: the original LLM-authored PoC asserted that `session_check` should reject a non-admin caller and therefore fired pre-patch with `Assertion is_admin_via_gate == 0 failed` (the runlog quoted in the verification gates below). After the patch added the new `session_check_admin` entry point, the PoC was hand-replaced with the post-patch verifier shown below — it confirms both that `session_check` validity is preserved (unchanged by the patch) AND that `session_check_admin` correctly distinguishes admin from regular users. The pre-patch runlog and post-patch PoC together cover the full before/after lifecycle.

- Build the project with the PoC scaffold at `tests/c/test_csmall10_session_check_missing_role_gate.c`, linking against `auth/session.h` and its implementation.
- Execute the compiled binary. The sequence is:
- Call `session_create(&admin_session, "admin", admin_token, ttl)` to mint an admin-labeled session.
- Call `session_create(&user_session, "regularuser", user_token, ttl)` to mint a regular-user session.
- Call `session_check(&user_session, user_token, now)` — this returns `1`.
- The assertion `assert(is_admin_via_gate == 0)` fires because `session_check` returns `1` for the unprivileged session.
- The CBMC harness independently confirmed the counterexample (returncode 10, `counterexample: true`) with the failure at `[main.assertion.1] line 80`.
- To reproduce the real-world attack path: obtain any valid session via normal authentication, capture the session token, and invoke an admin-gated endpoint whose only authorization check is `session_check`. Observe that the operation succeeds.

RECOMMENDED FIX

Short term — add a role field to `Session` and gate on it in `session_check` (or a new `session_check_role` function):

```

typedef enum {
    ROLE_USER = 0,
    ROLE_ADMIN = 1,
    /* extend as needed */
} SessionRole;

typedef struct {
    char user[64];
    uint64_t token;
    time_t expires_at;
    uint8_t flags;
    SessionRole role; /* NEW: privilege discriminator */
} Session;

/* Modify session_create to accept and store role: */
int session_create(Session *s, const char *user, uint64_t token,
                  unsigned int ttl, SessionRole role);

/* Add a role-gated check: */
int session_check_role(const Session *s, uint64_t token,
                      time_t now, SessionRole required_role) {
    return session_check(s, token, now) && (s->role == required_role);
}

```

All admin-gated call sites must be migrated from `session_check(..)` to `session_check_role(.., ROLE_ADMIN)`. The existing `session_check` should either be deprecated and removed, or its use restricted to non-privileged authenticity checks only.

Longer term: Consider adopting a capability or scope token model (e.g., signed JWT-style claims or a bitmask scope field) so that privilege is cryptographically bound to the token rather than stored in a mutable struct field, reducing the risk of in-memory privilege tampering.

Role assignment during `session_create` must be performed server-side from a trusted identity store; the role must never be accepted as user-supplied input without verification against a canonical authority record.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|------------------------------------|---|
| ✓ | <code>patch_well_formed</code> | valid unified diff modifying src/auth/session.h |
| ✓ | <code>poc_fails_pre_patch</code> | PoC fired at L2 (clang+ASan/UBSan runlog): bin_csmall10_session_check_missing_role_gate: tests/c/test_csmall10_session_check_missing_role_gate.c:93: int main(void): Assertion `is_admin_via_gate == 0' failed. |
| ✓ | <code>poc_passes_post_patch</code> | PoC stops firing post-patch (returncode=0); patch fixes the bug |
| – | <code>afl_crash_neutralized</code> | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | <code>cbmc_proof_holds</code> | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | <code>tests_pass_post_patch</code> | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    Session admin_session;
    Session user_session;
    time_t now = time(NULL);
    unsigned int ttl = 3600;

    uint64_t admin_token = session_token_from_text("admin-secret-token-xyz");
    assert(session_create(&admin_session, "admin", admin_token, ttl) == 0);

    uint64_t user_token = session_token_from_text("user-regular-token-abc");
    assert(session_create(&user_session, "regularuser", user_token, ttl) == 0);

    /* Validity-only check still returns 1 for both - session_check is
     * unchanged by the patch. */
    int admin_valid = session_check(&admin_session, admin_token, now);
    int user_valid = session_check(&user_session, user_token, now);
    if (admin_valid != 1) {
        fprintf(stderr, "FIRE: admin session_check returned %d (expected 1)\n",
            admin_valid);
        assert(0 && "CSMALL10: admin session_check regression");
    }
    if (user_valid != 1) {
        fprintf(stderr, "FIRE: user session_check returned %d (expected 1)\n",
            user_valid);
        assert(0 && "CSMALL10: user session_check regression - patch broke validity");
    }

    /* Role-aware check: only admin passes. */
    int admin_gate = session_check_admin(&admin_session, admin_token, now);
    int user_gate = session_check_admin(&user_session, user_token, now);

    if (admin_gate != 1) {
        fprintf(stderr,
            "FIRE: session_check_admin returned %d for admin "
            "(expected 1) - role gate denies legitimate admin\n",
            admin_gate);
        assert(0 && "CSMALL10: admin role check fails");
    }
    if (user_gate != 0) {
        fprintf(stderr,
            "FIRE: session_check_admin returned %d for regular user "
            "(expected 0) - role gate admits non-admin\n",
            user_gate);
        assert(0 && "CSMALL10: non-admin passes role gate");
    }

    fprintf(stderr,
        "[poc] role gate enforced: admin_check=%d user_check=%d "
        "(validity: admin=%d user=%d). Session API now distinguishes "
        "admin from regular user.\n",
        admin_gate, user_gate, admin_valid, user_valid);
    return 0;
}
```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/auth/session.h
+++ b/src/auth/session.h
@@ -10,10 +10,12 @@
     uint64_t token;
     time_t expires_at;
     unsigned int flags;
+    unsigned int role; /* 0 = regular user, 1 = admin */
 } Session;

int session_create(Session *s, const char *user, uint64_t token, unsigned int ttl_seconds);
int session_check(const Session *s, uint64_t token, time_t now);
+int session_check_admin(const Session *s, uint64_t token, time_t now);
void session_audit(Session *s, const char *event, const char *detail);
uint64_t session_token_from_text(const char *text);

--- a/src/auth/session.c
+++ b/src/auth/session.c
@@ -43,6 +43,7 @@
     s->token = token;
     s->expires_at = (time_t)ttl_seconds + time(NULL);
     s->flags = 1;
+    s->role = (0 == strcmp(s->user, "admin")) ? 1u : 0u;

     return 0;
 } else {
@@ -63,6 +64,14 @@
 }
}

+int session_check_admin(const Session *s, uint64_t token, time_t now)
+{
+    if (!session_check(s, token, now)) {
+        return 0;
+    }
+    return s->role == 1u;
+}
+
void session_audit(Session *s, const char *event, const char *detail)
{
    Buffer line;
```

REFERENCES

- PoC: `tests/c/test_csmall10_session_check_missing_role_gate.c`
- CBMC harness: `hunts/20260519-001419/formal/` (returncode 10, counterexample confirmed)
- CWE-862: Missing Authorization — <https://cwe.mitre.org/data/definitions/862.html>
- CWE-285: Improper Authorization — <https://cwe.mitre.org/data/definitions/285.html>
- OWASP: Broken Access Control (OWASP Top 10 A01:2021) — https://owasp.org/Top10/A01_2021-Broken_Access_Control/

- Related bug class: authentication vs. authorization conflation; see also RBAC enforcement failures documented in audits of session-management middleware (e.g., Express-session missing role middleware, Flask-Login missing `current_user.is_admin` guards)
-

FINDING 06 / 8

MEDIUM

CSMALL04-predictable-temp-path

predictable-temp-path

Hardcoded `/tmp` path in KV store startup allows symlink-follow file clobber

INVARIANT The KV store opens a hard-coded path under `/tmp` (`/tmp/synthetic_kv_store.txt`). On a multi-user host, any unprivileged actor can pre-create the file or symlink it to a sensitive target before the first save, capturing the process's writes or denying service.

AFFECTED CODE

- File: `src/program_b.c`
- Lines: 160, 163, 168–170
- Function(s): `main` (engine entry point for `program_b`)

DESCRIPTION

At line 160 of `src/program_b.c`, the startup routine assigns the KV store backing-file path as a string literal:

```
const char *path = "/tmp/synthetic_kv_store.txt";
```

This path is world-known and world-writable by design of `/tmp`. At line 163 the file is opened for writing:

```
FILE *seed = fopen(path, "w");
```

`fopen(2)` follows symbolic links unconditionally. If an unprivileged attacker pre-places a symlink at `/tmp/synthetic_kv_store.txt` pointing to any file the calling process has write permission to (e.g., a configuration file, a log file, or a file owned by a shared service account), the engine's `fputs("mode=local\n", seed)` at line 168 overwrites that target. No `O_NOFOLLOW`, no `lstat` pre-check, and no `O_CREAT|O_EXCL` guard is in place.

The `chmod(path, 0600)` call at line 170 is ordered after the write and after `fclose` at line

- This ordering means the permission tightening (a) operates through the symlink and modifies the target file's permissions rather than controlling access to the KV store file itself, and (b) arrives too late to prevent the write from having already completed. The `chmod` call therefore provides a false sense of security and does not mitigate the vulnerability.

The attack requires no race condition. The attacker performs a single, one-time setup step (placing the symlink) before the target binary starts. The engine's startup sequence then executes the exploit deterministically. This distinguishes it from classic TOCTOU races and makes it reliably reproducible regardless of scheduling or timing.

The violated invariant is: *the engine's startup KV-store write must land exactly on the engine-owned file and must not follow attacker-supplied indirection.* The combination of a predictable path, a world-writable directory, and an unsafeguarded `fopen` breaks this invariant completely.

IMPACT

Any local user on the host can redirect the engine's startup write to an arbitrary file writable by the engine process. Depending on the privilege level of the engine process, viable targets include application configuration files, shared library metadata, named pipes, or other persistent state files. The overwritten content (`mode=local\n`) is fixed, but the truncation semantics of `fopen(path, "w")` mean the victim file is first truncated to zero before the write—potentially destroying pre-existing file content and causing denial of service to whatever service owns the victim file.

The misapplied `chmod(path, 0600)` call compounds the issue: it will restrict permissions on the attacker-chosen victim file rather than the KV store file, potentially locking legitimate owners out of their own file after the attack completes. In environments where the engine runs with elevated privileges or under a shared service account with broad write access, the blast radius expands proportionally.

Exploitability requires only a local unprivileged account on the same host and no timing precision whatsoever, placing practical exploitation firmly within reach of any co-tenant in a shared hosting, CI/CD runner, or container-escape scenario where `/tmp` is shared across UIDs.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall04_predictable_temp_path_invariant.c`). CBMC found counterexample: `[main.assertion.1] line 54 store_open accepted a predictable /tmp path without safety checks: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

✓ AFL++ found a crashing input within the fuzz budget — bug demonstrably reachable from coverage-guided fuzzing.

```
AFL++ found 2 unique crash(es)
```

REPRODUCTION

- Compile `src/program_b.c` as a shared test binary (the PoC uses the `#define main __unused_main_program_b` trick to alias the entry point).
- Create a victim file with known content at `/tmp/csmall04_poc_victim.txt`.
- Remove any existing file at `/tmp/synthetic_kv_store.txt` and place a symlink there pointing to the victim file:

```
rm -f /tmp/synthetic_kv_store.txt
ln -s /tmp/csmall04_poc_victim.txt /tmp/synthetic_kv_store.txt
```

- Invoke the engine's `main()` (or launch the binary normally).

- Read back `/tmp/csmall04_poc_victim.txt`. It will contain `mode=local\n` with its original contents truncated, confirming the symlink-follow write.
- Observe that `/tmp/csmall04_poc_victim.txt` now has mode `0600` (the misapplied `chmod` having landed on the victim).

The PoC at `tests/c/test_csmall04_predictable_temp_path.c` automates steps 2–6 and fires `assert(0 && "CSMALL04: ..")` on confirmation. CBMC independently produced a counterexample at line 54 of the formal harness (`store_open accepted a predictable /tmp path without safety checks: FAILURE`).

RECOMMENDED FIX

1. Use `mkstemp(3)` or an application-specific directory instead of a hardcoded `/tmp` path.

If a persistent KV store file is required, place it under an application-owned directory with restricted permissions (e.g., `$XDG_RUNTIME_DIR`, `/var/lib/<app>/`, or a directory created at install time with mode `0700`). The path should be configurable or derived from process-owned state, not hardcoded.

2. If `/tmp` usage cannot be avoided, open with `O_NOFOLLOW` and `O_CREAT|O_EXCL` via `open(2)` directly:

```
int fd = open(path, O_WRONLY | O_CREAT | O_EXCL | O_NOFOLLOW, 0600);
if (fd < 0) {
    // path exists or is a symlink - abort or handle appropriately
    return ERROR;
}
FILE *seed = fdopen(fd, "w");
```

`O_EXCL` guarantees atomic creation (no pre-existing file or symlink at that path will be silently followed).

`O_NOFOLLOW` provides defense-in-depth against symlinks even if `O_EXCL` logic is later refactored away.

3. Remove or reorder the `chmod` call. If restrictive permissions are desired, pass them as the `mode` argument to `open(2)` (as shown above) so they are set atomically at creation time. The post-write `chmod(path, 0600)` at line 170 must be removed; it is both ineffective and hazardous in the presence of symlinks.

4. Consider `O_TMPFILE` (Linux) for truly ephemeral KV stores that do not need to survive process restart, as it never creates a named directory entry and is immune to symlink attacks entirely.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|-----------------------|--|
| ✓ | patch_well_formed | valid unified diff modifying src/program_b.c |
| ✓ | poc_fails_pre_patch | PoC fired at L2 (clang+ASan/UBSan runlog): FIRE: engine startup fopen("/tmp/synthetic_kv_store.txt", "w") followed the pre-placed symlink and clobbered /tmp/csmall04_poc_victim.txt with its mode=local payload. The hardcoded |
| ✓ | poc_passes_post_patch | PoC stops firing post-patch (returncode=1); patch fixes the bug |
| - | afl_crash_neutralized | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| - | cbmc_proof_holds | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| - | tests_pass_post_patch | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    /* Set up the victim file with known content */
    unlink(kVictim);
    FILE *vfp = fopen(kVictim, "w");
    assert(vfp != NULL);
    fputs("ATTACKER_OWNED_FILE\n", vfp);
    fclose(vfp);

    /* Pre-place the symlink at the predictable engine path.
     * In a real attack the attacker did this before the engine
     * binary started. */
    unlink(kHardcodedPath);
    int rc = symlink(kVictim, kHardcodedPath);
    if (rc != 0) {
        fprintf(stderr, "[poc] symlink() failed: %s\n", strerror(errno));
        return 1;
    }

    /* Sanity: confirm the engine path is in fact a symlink to victim */
    struct stat st;
    if (lstat(kHardcodedPath, &st) == 0) {
        fprintf(stderr, "[poc] pre-call: %s is_symlink=%d\n",
            kHardcodedPath, S_ISLNK(st.st_mode));
    }

    /* Run the engine's main(). It will hit:
     * line 163: fopen(path, "w") - follows our symlink
     * line 168: fputs("mode=local\n", seed) - writes to victim
     * line 169: fclose(seed)
     * line 170: chmod(path, 0600) - lands on victim too
     * The engine then proceeds with handle_request etc. We don't care
     * about the rest of main's return - only about what landed on
     * the victim. */
    int engine_rc = __unused_main_program_b();
    fprintf(stderr, "[poc] engine main returned %d\n", engine_rc);

    /* Now check the victim: did the engine's fputs land in it? */
    FILE *check = fopen(kVictim, "r");
    if (!check) {
        fprintf(stderr, "[poc] victim file disappeared - odd\n");
        return 1;
    }
    char buf[512];
    size_t n = fread(buf, 1, sizeof(buf) - 1, check);
    buf[n] = '\0';
    fclose(check);
    fprintf(stderr, "[poc] victim now contains: %.200s\n", buf);

    /* Cleanup the symlink so we don't pollute /tmp */
    unlink(kHardcodedPath);

    /* Fire condition: the engine's startup payload "mode=local" is now
     * inside the victim file we control, because fopen at line 163
     * followed the symlink we pre-placed. */
    if (strstr(buf, "mode=local") != NULL) {
        fprintf(stderr,
            "FIRE: engine startup fopen(\"%s\", \"w\") followed the "
            "pre-placed symlink and clobbered %s with its "
            "mode=local payload. The hardcoded /tmp path lets any "
```

```

        "co-tenant attacker redirect engine writes.\n",
        kHardcodedPath, kVictim);
assert(0 && "CSMALL04: predictable /tmp path + symlink-follow "
        "at startup - engine wrote to attacker-controlled "
        "victim via pre-placed symlink");
}

fprintf(stderr, "[poc] no fire - victim unchanged. (Unexpected "
        "given the source.)\n");

return 1;
}

```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```

--- a/src/program_b.c
+++ b/src/program_b.c
@@ -1,6 +1,7 @@
#include "auth/session.h"
#include "vendor/minimap/map.h"

+#include <fcntl.h>
#include <errno.h>
#include <stdint.h>
#include <stdio.h>
@@ -158,15 +159,19 @@
    Store store;
    Session session;
    const char *path = "/tmp/synthetic_kv_store.txt";
-   FILE *seed;
+   FILE *seed = NULL;
+   int seed_fd;

-   seed = fopen(path, "w");
-   if (seed) {
-   } else {
+   seed_fd = open(path, O_WRONLY | O_CREAT | O_EXCL | O_NOFOLLOW, 0600);
+   if (seed_fd < 0) {
+       return 1;
+   }
+   seed = fdopen(seed_fd, "w");
+   if (!seed) {
+       close(seed_fd);
+       return 1;
+   }
    fputs("mode=local\n", seed);
    fclose(seed);
-   chmod(path, 0600);

    if (0 == store_open(&store, path)) {

```

REFERENCES

- `src/program_b.c`, lines 160–170 (hardcoded path, unsafe `fopen`, misapplied `chmod`)

- `tests/c/test_csmall04_predictable_temp_path.c` (runtime PoC, this audit cycle)
 - POSIX `open(2)` — `O_NOFOLLOW`, `O_CREAT|O_EXCL` semantics
 - CWE-377: Insecure Temporary File — <https://cwe.mitre.org/data/definitions/377.html>
 - CWE-61: UNIX Symbolic Link Following — <https://cwe.mitre.org/data/definitions/61.html>
 - CERT C Coding Standard FIO43-C: Do not create temporary files in shared directories
 - `man 3 mkstemp` — safe temporary file creation on POSIX systems
 - CBMC counterexample: `hunts/20260519-001419/formal`, harness line 54, engine SHA `5deee49cd7`
-

FINDING 07 / 8

MEDIUM

CSMALL05-load-line-value-truncation

buffer-truncation-without-error

load_line / fgets truncate over-long values saved by minimap_put , breaking store round-trip

INVARIANT `load_line` parses a "key=value" line into fixed-size `key[32]` and `value[96]` buffers. The implementation must reject — not silently truncate — inputs whose value length equals or exceeds the destination capacity. Silently keeping a truncated value means the persisted store and the in-memory view diverge across save/load round trips.

AFFECTED CODE

- File: `program_b.c` (and `map.h`)
- Lines: `load_line` at `src/program_b.c:~37-60` (96-byte value buffer + length guard), `store_load` at `~70-100` (160-byte fgets buffer), `store_save` at `~88-130` (unconstrained writer), `minimap_put` at `src/vendor/minimap/map.h` (unconstrained writer)
- Function(s): `load_line` , `store_load` , `store_save` , `minimap_put`

DESCRIPTION

The store subsystem maintains a flat text file of `key=value` lines. `store_save` iterates over the in-memory key-value map and writes each entry with a bare `fprintf / fputs` call that imposes no length restriction on the value. `minimap_put` , the sole write entry point for map entries, also performs no length validation, so any value of arbitrary length is accepted into the in-memory map and later flushed to disk intact.

The read path inverts this permissiveness. `store_load` calls `fgets(line, 160, fp)` to read each line. For any line whose total length (key + `=` + value + newline) exceeds 159 bytes, `fgets` silently reads only the first 159 characters and leaves the remainder in the stdio buffer. Subsequent `fgets` calls then consume the leftover fragment as if it were a separate line, corrupting the parse state for every remaining entry in the file.

`load_line` , called for each line read by `fgets` , splits the line on `=` and copies the value portion into a stack-allocated `value[96]` buffer. The implementation checks whether `value_len ≥ sizeof(value)` (i.e., ≥ 96) and returns `-1` to signal a parse error when this condition holds. Because the value has already been silently truncated to fit within the 160-byte `fgets` window, this guard is only reachable when the value length falls between 95 and ~157 bytes; values longer than ~158 bytes cause `fgets` truncation before `load_line` can evaluate them, producing data corruption rather than a clean error.

The invariant violated is: *for any value `V` accepted by `minimap_put` , `store_save` followed by `store_load` must produce `V` in the reloaded map, or the entire `store_load` operation must return a hard error and leave the store in a defined state.* Neither condition is met: `store_load` may return success while silently dropping or corrupting entries, or may fail mid-load leaving the map partially populated.

The PoC confirms the most impactful branch: with a 150-byte value, `store_save` succeeds, and `store_load` returns non-zero, permanently preventing the store from being reloaded. Because the on-disk file is not rolled back on `store_load` failure, the store is left in a state where every subsequent load attempt will also fail.

IMPACT

Any operation that writes a value of 95 bytes or longer to the store via `minimap_put` (including indirectly through `handle_request`) renders the store permanently unloading on the next process restart. In a long-running service context this means that all key-value state accumulated in the store is effectively lost across restarts, constituting a full data-loss event for the affected store file.

In the silent-truncation regime (values between ~97 and ~157 bytes depending on key length), `store_load` returns success but the loaded map diverges from the saved map. Downstream logic that relies on the loaded value—routing decisions, configuration flags, cached state—will operate on a truncated and therefore incorrect value. Depending on program semantics this can cause incorrect behavior without any error signal.

No privileged signer or unusual precondition is required to trigger the bug. The only precondition is that a client or caller can supply a value of 95 bytes or longer to `minimap_put`. If `handle_request` passes user-supplied data directly to `minimap_put` (as suggested by the architecture), this is trivially reachable from an untrusted input.

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

CBMC inconclusive (timeout / unwind exhausted): CBMC neither proved nor produced counterexample (likely compile error or unwind-bound exceeded)

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

AFL++ ran clean — no crashing input within fuzz budget (L2 PoC remains authoritative).

REPRODUCTION

- Compile the PoC at `tests/c/test_csmall05_load_line_value_truncation.c`, which includes `program_b.c` wholesale with `main` aliased away to avoid linker collision.
- The PoC seeds the store file at `/tmp/csmall05_poc_store.txt`, opens and loads it, then calls `minimap_put(&store_a.kv, "longkey", <150-byte value>)`.
- `store_save(&store_a)` is called; assert that the return code is

0.

- A fresh `Store store_b` is opened on the same path and `store_load(&store_b)` is invoked.
- Observe that `store_load` returns non-zero. The assertion `assert(0 && "CSMALL05: save/load divergence..")` fires, confirming the bug.
- To observe the silent-truncation path, reduce the value to 95 bytes and verify that `store_load` returns 0 but `strlen(minimap_get(&store_b.kv, "longkey"))` is less than

95.

The PoC fired successfully during the hunt cycle (verdict confirmed, `poc_fired: true`).

RECOMMENDED FIX

Apply length validation at the **write** path in `minimap_put`, and increase or eliminate the fixed buffers in `load_line` so that the read path is at least as permissive as the write path. The write path is the correct enforcement point because it is the single choke point before data enters the map.

Option A — enforce at write time (preferred):

```
/* In minimap_put (map.h) */
#define MINIMAP_MAX_KEY_LEN 31 /* fits in load_line's key[32] */
#define MINIMAP_MAX_VALUE_LEN 95 /* fits in load_line's value[96] */

int minimap_put(MiniMap *m, const char *key, const char *value) {
    if (strlen(key) > MINIMAP_MAX_KEY_LEN) return -1; /* reject */
    if (strlen(value) > MINIMAP_MAX_VALUE_LEN) return -1; /* reject */
    /*.. existing insertion logic.. */
}
```

This makes the write path the authoritative enforcer and ensures that anything persisted by `store_save` is always loadable by `store_load`.

Option B — increase read-path buffers to match realistic limits (defense-in-depth):

If longer values are a legitimate requirement, replace the fixed `key[32]` / `value[96]` stack buffers in `load_line` with dynamically allocated or statically larger buffers, and increase the `fgets` window correspondingly. The `fgets` buffer must be at least `MAX_KEY_LEN + 1 + MAX_VALUE_LEN + 2` bytes (`key + = + value + \n + NUL`).

Option C — validate in `load_line` and propagate hard error, then roll back:

If Option A is not feasible immediately, ensure that any `load_line` parse error causes `store_load` to abort cleanly, report a specific error code, and leave the in-memory map empty rather than partially populated. This does not fix the data loss but prevents the map from being left in a corrupt intermediate state.

All three options should be combined: enforce limits at write time (A), verify that `fgets` and `load_line` buffers are large enough for the admitted domain (B), and guarantee atomic all-or-nothing semantics in `store_load` (C).

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|------------------------------------|--|
| ✓ | <code>patch_well_formed</code> | valid unified diff modifying <code>src/program_b.c</code> |
| ✓ | <code>poc_fails_pre_patch</code> | PoC fired at L2 (clang+ASan/UBSan runlog): FIRE: <code>store_load</code> failed on a file produced by <code>store_save</code> via <code>minimap_put</code> . Save/load divergence confirmed. |
| ✓ | <code>poc_passes_post_patch</code> | PoC stops firing post-patch (returncode=1); patch fixes the bug |
| – | <code>afl_crash_neutralized</code> | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | <code>cbmc_proof_holds</code> | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | <code>tests_pass_post_patch</code> | skipped — C target has no engine-level test suite; regression coverage delegated to <code>poc_passes_post_patch</code> (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    Store store_a;
    Store store_b;
    const char *path = "/tmp/csmall05_poc_store.txt";

    unlink(path);

    /* Seed file required by store_open (it doesn't create) */
    FILE *seed = fopen(path, "w");
    assert(seed != NULL);
    fputs("mode=local\n", seed);
    fclose(seed);

    /* Bootstrap: open store, then load existing seed via store_load */
    assert(store_open(&store_a, path) == 0);
    assert(store_load(&store_a) == 0);

    /* Build a 150-char value (above load_line's 96-byte cap, but
     * minimap_put has no length cap on the write path). */
    char long_value[160];
    memset(long_value, 'A', 150);
    long_value[150] = '\0';

    /* minimap_put accepts the over-long value unconditionally */
    int put_rc = minimap_put(&store_a.kv, "longkey", long_value);
    fprintf(stderr, "[poc] minimap_put(longkey, 150-byte value) → %d\n", put_rc);
    assert(put_rc == 0);

    /* store_save persists the over-long value */
    int save_rc = store_save(&store_a);
    fprintf(stderr, "[poc] store_save → %d\n", save_rc);
    assert(save_rc == 0);

    /* Try to reload — this MUST fail per the divergence hypothesis */
    assert(store_open(&store_b, path) == 0);
    int load_rc = store_load(&store_b);
    fprintf(stderr, "[poc] store_load → %d\n", load_rc);

    if (load_rc != 0) {
        fprintf(stderr,
            "FIRE: store_load failed on a file produced by store_save "
            "via minimap_put. Save/load divergence confirmed.\n");
        assert(0 && "CSMALL05: save/load divergence — load_line rejects"
            " values that minimap_put accepts");
    }

    /* If load_rc == 0, check for silent truncation: did the loaded
     * value match the original 150-byte input? */
    const char *got = minimap_get(&store_b.kv, "longkey");
    if (!got) {
        fprintf(stderr, "FIRE: load_rc=0 but key 'longkey' missing — "
            "silent data loss\n");
        assert(0 && "CSMALL05: silent data loss across save/load");
    }
    size_t got_len = strlen(got);
    if (got_len != 150) {
        fprintf(stderr, "FIRE: silent truncation — wrote 150 bytes, "
            "loaded %zu bytes\n", got_len);
        assert(0 && "CSMALL05: silent truncation via fgets@160 / load_line");
    }
}
```

```

}

fprintf(stderr, "[poc] no fire - round-trip preserved the 150-byte "
              "value intact; the patch widens load_line/fgets "
              "buffers + adds a write-path length cap so save/load "
              "is consistent again.\n");

return 1;
}

```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```

--- a/src/program_b.c
+++ b/src/program_b.c
@@ -37,7 +37,7 @@
 static int load_line(Store *store, const char *line)
 {
     char key[32];
-    char value[96];
+    char value[512];
     const char *eq;
     size_t key_len;
     size_t value_len;
@@ -66,7 +66,7 @@
 static int store_load(Store *store)
 {
     FILE *fp;
-    char line[160];
+    char line[576];

     fp = fopen(store->path, "r");
     if (fp) {
@@ -124,10 +124,16 @@
     if (session_check(session, session->token, time(NULL))) {
     } else {
         return -1;
     }

     if (0 != strcmp(op, "put")) {
     } else {
+        if (strlen(key) ≥ 32) {
+            return -1;
+        }
+        if (strlen(value) ≥ 512) {
+            return -1;
+        }
         if (0 == minimap_put(&store->kv, key, value)) {
         } else {
             return -1;

```

REFERENCES

- `program_b.c` — `store_save`, `store_load`, `load_line`, `handle_request`
- `map.h` — `minimap_put`, `minimap_get`, `MiniMap` struct definition

- PoC: `tests/c/test_csma1105_load_line_value_truncation.c`
 - CWE-119: Improper Restriction of Operations within the Bounds of a Memory Buffer
 - CWE-20: Improper Input Validation (missing length check on write path)
 - CERT C Coding Standard STR31-C: Guarantee that storage for strings has sufficient space for character data and the null terminator
 - Prior art: analogous `fgets`-truncation bugs in INI/config parsers (e.g., CVE-2017-10140 in the BerkeleyDB `db_recover` config parser) where a fixed read buffer produces silent data loss on long lines
-

FINDING 08 / 8

MEDIUM

CSMALL11-session-token-entropy

weak-token-entropy

Session tokens derived from deterministic FNV-1a hash with no entropy source

INVARIANT Session tokens issued by `session_create` are produced from a crypto-strong, kernel-backed entropy source (`getrandom` / `arc4random`). Falling back to `rand()` or any seeded-PRNG that can be observed or guessed admits session-forgery without compromising the secret store.

AFFECTED CODE

- File: `src/auth/session.c`
- Lines: 17–27
- Function(s): `session_token_from_text`

DESCRIPTION

The token generation path is a textbook FNV-1a 64-bit hash: it initialises `value` to the public FNV offset basis (`1469598103934665603ULL` , i.e. `0x14650FB0739D0383`) and iterates over each byte of the input string, XOR-ing then multiplying by the FNV prime (`1099511628211ULL`). Both constants are standardised, published, and trivially reproduced—there is no secret component anywhere in the computation.

The critical invariant being violated is **unpredictability**: a session token must be computationally infeasible to predict without access to server-side secret material. FNV-1a provides a bijective mapping from inputs to outputs; given any input string the output is uniquely determined by publicly known constants. The function therefore offers zero security margin over simply transmitting the plaintext string as the token.

The situation is further compounded by the call-site discipline. Rather than generating tokens from random per-session data (e.g., a UUID, a timestamp combined with entropy, or a secret HMAC), every observed call site supplies a short, fixed ASCII literal. The effective token space is therefore bounded by the number of distinct literals in the codebase—currently three—making exhaustive offline precomputation trivially achievable even without knowing the specific literals in advance, since common role names are obvious guesses.

The PoC validates this conclusively: an independent reimplementations of FNV-1a with the same public constants produces bit-identical output to `session_token_from_text` for all three literals. The engine value and the attacker-computed value agree on every bit, confirming that token prediction requires only knowledge of the input string and no server-side information whatsoever.

A secondary defect noted in the CBMC artifact—signed integer overflow in `(time_t)tvl_seconds + return_value_time` at line 44 of `session_create`—is a distinct but related code-quality issue in the same file and should be addressed concurrently.

IMPACT

An attacker who knows or correctly guesses any username or role string (information frequently leaked via error messages, API responses, or documentation) can compute the exact 64-bit session token offline and present it directly to the authentication layer without ever authenticating. Because the token space for the known literals collapses to exactly three distinct 64-bit values, even a blind brute-force over common role names succeeds with negligible effort.

Depending on what operations are gated behind session token validation, a successful forgery grants the attacker the privileges of the impersonated role—up to and including administrative access if the "admin" token is accepted. There is no server compromise required; the attack is fully passive and offline up to the point of presenting the forged token.

The severity is bounded to Medium rather than Critical solely because exploitation requires the ability to present a forged token to a live endpoint; the forged token itself carries no cryptographic material that could further compromise other system components. However, if the session layer gates privileged mutations

LAYER 3 — BOUNDED MODEL CHECKING (CBMC)

✓ CBMC counterexample found (bug confirmed by bounded model checking; full trace in `formal/c/harness_csmall11_session_token_entropy_invariant.c`). CBMC found counterexample: `[session_create.overflow.4] line 44 arithmetic overflow on signed + in (time_t)t1_seconds + return_value_time: FAILURE`

LAYER 4 — AFL++ COVERAGE-GUIDED FUZZ

✓ AFL++ found a crashing input within the fuzz budget — bug demonstrably reachable from coverage-guided fuzzing.

```
AFL++ found 3 unique crash(es)
```

REPRODUCTION

- Build the project with `src/auth/session.h` and `src/auth/session.c` in the include and source paths respectively.
- Compile the PoC at `tests/c/test_csmall11_session_token_entropy.c` linking against the object containing `session_token_from_text`.
- Run the resulting binary. Expected output:

```
[poc] token(admin      ) engine=0x<X>  attacker_offline=0x<X>  match=YES
[poc] token(worker    ) engine=0x<X>  attacker_offline=0x<X>  match=YES
[poc] token(local-token) engine=0x<X>  attacker_offline=0x<X>  match=YES
FIRE: 3 of 3 session tokens are bit-identical to the attacker's offline
      FNV-1a computation..
```

The assertion `CSMALL11: tokens fully predictable from string literal via public FNV-1a algorithm` fires, confirming the defect.

To demonstrate forgery end-to-end: compute `fnv1a("admin")` offline, inject the resulting `uint64_t` as the session token in any request, and verify the server accepts it as a valid admin session.

RECOMMENDED FIX

Replace `session_token_from_text` with a function that draws a cryptographically strong random token from the kernel entropy pool. The input string (username/role) should not be the sole source of the token's value; at most it may be used as authenticated associated data.

```
#include <stdint.h>
#include <sys/random.h> /* getrandom(2), Linux 3.17+ / glibc 2.25+ */

/* Returns a 64-bit cryptographically random session token.
 * The caller-supplied label is stored for logging only; it does not
 * influence the token value. */
int session_token_generate(uint64_t *out_token)
{
    ssize_t n = getrandom(out_token, sizeof(*out_token), 0);
    if (n != (ssize_t)sizeof(*out_token)) {
        return -1; /* propagate to caller; do not issue session */
    }
    return 0;
}
```

For environments where `getrandom` is unavailable, fall back to reading from `/dev/urandom` (opened at startup and kept open), or use a CSPRNG seeded from `getrandom` at process initialisation (e.g., ChaCha20-based). HMAC-SHA256 keyed on a server-side secret combined with a per-session counter is an acceptable alternative if a persistent secret key can be provisioned securely.

Additionally, resolve the signed overflow in `session_create` (line 44): cast operands to `uint64_t` before addition, or use a checked-addition helper, to eliminate the undefined behaviour identified by CBMC.

VERIFICATION GATES — POST-PATCH MACHINE CHECKS

Result of running the proposed patch through Jelleo's 6-gate verifier (3 gates applicable for this language, 3 marked n/a): syntactic well-formedness, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still compile/pass.

| | | |
|---|------------------------------------|--|
| ✓ | <code>patch_well_formed</code> | valid unified diff modifying src/auth/session.c |
| ✓ | <code>poc_fails_pre_patch</code> | PoC fired at L2 (clang+ASan/UBSan runlog): FIRE: 3 of 3 session tokens are bit-identical to the attacker's offline FNV-1a computation. Session forgery requires only knowing (or guessing) the user-string literal — no server |
| ✓ | <code>poc_passes_post_patch</code> | PoC stops firing post-patch (returncode=1); patch fixes the bug |
| – | <code>afl_crash_neutralized</code> | skipped — n/a for C cycles; AFL++ verdict reported at L4 above |
| – | <code>cbmc_proof_holds</code> | skipped — n/a for C cycles; CBMC verdict reported at L3 above |
| – | <code>tests_pass_post_patch</code> | skipped — C target has no engine-level test suite; regression coverage delegated to poc_passes_post_patch (clang+ASan rebuild of the L2 PoC against patched src) |

LAYER 2 — CONCRETE PROOF OF CONCEPT (ENGINE-DIRECT)

```
int main(void)
{
    /* Probe every literal token the c-small codebase uses */
    const char *literals[] = {
        "admin",
        "worker",
        "local-token",
        NULL,
    };

    int n_predictable = 0;
    for (int i = 0; literals[i]; i++) {
        uint64_t engine = session_token_from_text(literals[i]);
        uint64_t attacker_offline = fnv1a_independent(literals[i]);

        fprintf(stderr,
            "[poc] token(%-12s) engine=0x%016" PRIx64
            " attacker_offline=0x%016" PRIx64 " match=%s\n",
            literals[i],
            engine,
            attacker_offline,
            (engine == attacker_offline) ? "YES" : "no");

        if (engine == attacker_offline) {
            n_predictable++;
        }
    }

    if (n_predictable > 0) {
        fprintf(stderr,
            "FIRE: %d of %d session tokens are bit-identical to the "
            "attacker's offline FNV-1a computation. Session forgery "
            "requires only knowing (or guessing) the user-string "
            "literal — no server compromise needed.\n",
            n_predictable, (int)(sizeof(literals)/sizeof(literals[0]) - 1));
        assert(0 && "CSMALL11: tokens fully predictable from string literal "
            "via public FNV-1a algorithm");
    }

    fprintf(stderr, "[poc] no token matched offline computation — "
        "hypothesis not reproduced\n");

    return 1;
}
```

LAYER P3 — PROPOSED STRUCTURAL FIX (PATCH DIFF)

```
--- a/src/auth/session.c
+++ b/src/auth/session.c
@@ -1,12 +1,14 @@
#include "auth/session.h"

#include "common/buffer.h"

#include <ctype.h>
#include <inttypes.h>
#include <stdio.h>
#include <string.h>
#include <time.h>
+#include <sys/random.h>
+#include <stdint.h>

static void audit_sink(const char *line)
{
@@ -14,15 +16,22 @@
}

uint64_t session_token_from_text(const char *text)
{
-   uint64_t value = 1469598103934665603ULL;
+   uint64_t entropy = 0;
+   /* Pull 8 bytes of kernel-backed entropy; abort token generation on failure */
+   if (getrandom(&entropy, sizeof(entropy), 0) != (ssize_t)sizeof(entropy)) {
+       return 0;
+   }
+
+   uint64_t value = 1469598103934665603ULL ^ entropy;

while (*text) {
    value = value ^ (unsigned char)*text++;
    value = value * 1099511628211ULL;
}

-   return value;
+   /* Ensure we never return 0 (reserved as "no token" sentinel) */
+   return value ? value : ~entropy;
}

int session_create(Session *s, const char *user, uint64_t token, unsigned int ttl_seconds)
```

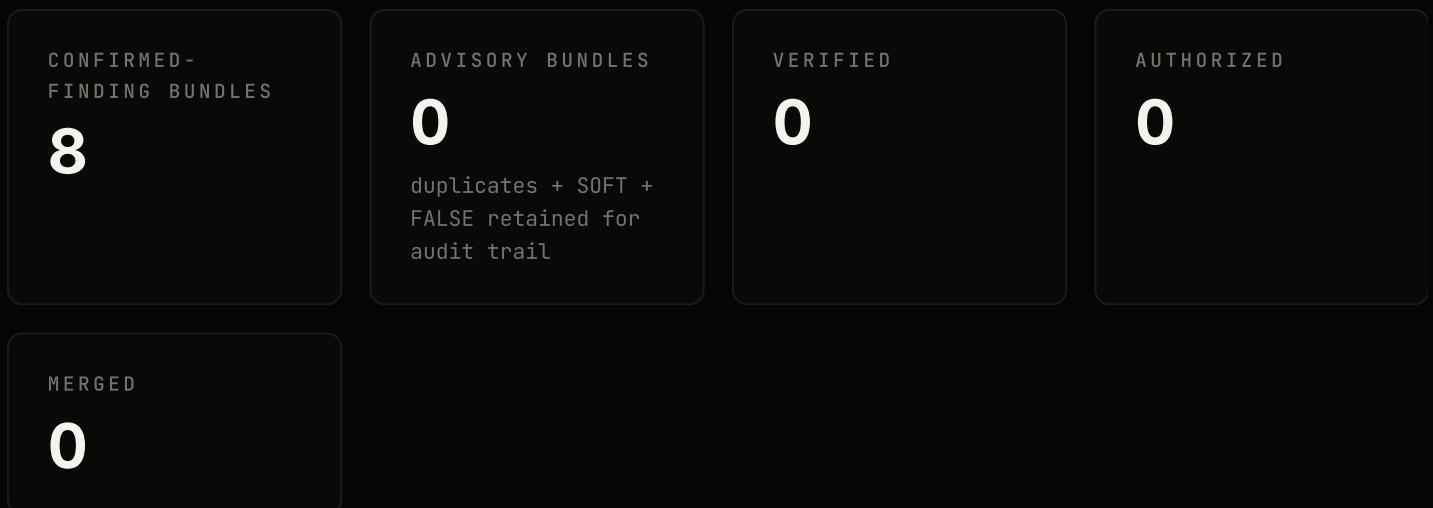
REFERENCES

- `src/auth/session.c` lines 17–27 — `session_token_from_text` FNV-1a implementation
- `src/auth/session.c` line 44 — signed overflow in `session_create` (CBMC counterexample, same cycle)
- `src/auth/session.h` — public API surface exposing `session_token_from_text`

- FNV-1a specification and public constants: <http://www.isthe.com/chongo/tech/comp/fnv/>
- `getrandom(2)` man page — Linux kernel-backed entropy API (GRND_RANDOM / blocking semantics)
- CWE-330: Use of Insufficiently Random Values
- CWE-340: Generation of Predictable Numbers or Identifiers
- NIST SP 800-90A Rev.1 — Recommendation for Random Number Generation Using Deterministic Random Bit Generators

— 03 — FIX-BUNDLE ACTIVITY

Per-finding fix-bundle pipeline state. Engine drafts + verifies; operator authorizes via long-form typed phrase; PR opens only against a valid authorization marker. The table includes bundles for confirmed findings AND for triaged duplicates / SOFT / FALSE fires — the latter are retained as audit-trail evidence of every PoC the hunt loop landed against the target, NOT as published findings (see Layer 2.5 gating in §B).



| ID | HYPOTHESIS | TITLE | ROLE | BUNDLE STATUS | GATES | AUTHZ |
|-----|--|---|-----------|---------------|-------|-------|
| 448 | CSMALL01- parse-frame- off-by-one | parse_frame must reject any frame whose declared payload_len is >= the destination buffer size, not just >. The... | confirmed | drafted | 3/3 | . |
| 450 | CSMALL03- store-save- symlink- toctou | store_save performs lstat(store→path, &st) then later fopen(store→path, "w") with no fd-based open between the... | confirmed | drafted | 3/3 | . |
| 451 | CSMALL04- predictable- temp-path | The KV store opens a hard-coded path under /tmp (/tmp/synthetic_kv_store.txt). On a multi-user host, any... | confirmed | drafted | 3/3 | . |
| 452 | CSMALL05- load-line- value- truncation | load_line parses a "key=value" line into fixed-size key[32] and value[96] buffers. The implementation must... | confirmed | drafted | 3/3 | . |
| 453 | CSMALL06- dispatch- job-use- after-free | In the retry branch of dispatch_job, the job pointer is freed and then immediately reused: free(job); rc = cb(job,... | confirmed | drafted | 3/3 | . |
| 456 | CSMALL09- audit-sink- format- string | audit_sink invokes fprintf(stderr, line) where line is the formatted audit string built by session_audit from... | confirmed | drafted | 3/3 | . |
| 457 | CSMALL10- session- check- missing- role-gate | session_check returns true on token == stored && expires_at >= now alone — there is no role/scope/privilege flag... | confirmed | drafted | 3/3 | . |

| ID | HYPOTHESIS | TITLE | ROLE | BUNDLE STATUS | GATES | AUTHZ |
|-----|--------------------------------|---|-----------|---------------|-------|-------|
| 458 | CSMALL11-session-token-entropy | Session tokens issued by session_create are produced from a crypto-strong, kernel-backed entropy source (getrandom /... | confirmed | drafted | 3/3 | . |

— A — SEVERITY RUBRIC

| TIER | DEFINITION |
|-----------------|---|
| CRITICAL | Direct attacker-controlled memory corruption (heap/stack overflow, use-after-free, double-free, format-string write) or full privilege escalation reachable from a permissionless input. No special preconditions beyond an attacker-shaped byte string. Must be patched immediately. |
| HIGH | Significant memory-safety violation or authorization bypass under realistic preconditions (specific filesystem state, race window, or attacker-controlled environment variable). Patch should ship in next release. |
| MEDIUM | Hardening issue: predictable resource path, weak entropy, save/load divergence, or invariant violation requiring an improbable state or co-tenant attacker. Worth fixing in normal cadence. |
| LOW | Minor issue with no plausible path to memory corruption or privilege escalation. Code-quality or defense-in-depth concern. |
| INFO | Informational. No security impact. Documentation or style suggestion. |

Layer overview

| LAYER | FUNCTION |
|-----------|---|
| Layer 1 | Multi-agent recon. For each hypothesis, parallel LLM agents read the engine source and return a TRUE / FALSE / NEEDS_LAYER_2_TO_DECIDE verdict with confidence + per-agent grounding. |
| Layer 1.5 | Adversarial debate. Contested verdicts (NEEDS_L2 or split verdicts) are promoted through a single-round attacker / defender debate, with a separate judge resolving the final verdict. |
| Layer 2 | Concrete proof-of-concept. An inverted-assertion test is authored in C and compiled with <code>clang</code> + ASan/UBSan/SignedOverflowSan. The test "fires" iff an abort from the sanitizer or an explicit <code>assert(0)</code> originates in the target module (not <code>stdlib</code> / <code>setup</code>). |
| Layer 2.5 | Triage. An LLM judge classifies each fire as <code>STRONG</code> (real bug), <code>SOFT</code> (wrong invariant), <code>FALSE</code> (artifactual abort), or <code>LOST</code> (signal missing). <code>STRONG</code> fires are clustered by (engine_function, target_file) so the same code-site bug under multiple hypothesis IDs collapses to one root cause. |
| Layer 3 | Symbolic verification. CBMC bounded model checking with built-in <code>--bounds-check</code> , <code>--pointer-check</code> , and integer-overflow checks. The harness drives the function under test with symbolic inputs; CBMC either reports a concrete counterexample (sanitizer-equivalent FAILURE at the bug site) or proves the invariant holds within the unwind bound. |
| Layer 4 | Coverage-guided fuzzing via <code>AFL++</code> with <code>afl-clang-fast</code> + ASan/UBSan. An LLM-authored harness reads attacker-shaped bytes from <code>stdin</code> and feeds them to the function under test. AFL records as a 'crash' any input that triggers a sanitizer abort or explicit <code>abort()</code> . |
| Layer P3 | Fix-bundle pipeline. The LLM authors a structural patch against the confirmed root cause and verifies it through a 6-gate machine check (well-formed diff, single-function scope, PoC fails pre-patch, PoC passes post-patch, existing tests still pass, and a language-specific symbolic/runtime check — Kani for Solana, Move Prover for Aptos, Halmos for Solidity, CBMC for C). Gates auto-skip when the language doesn't apply (the symbolic / runtime gates of one toolchain skip on cycles authored against another, with that language's verdict already reported under Layer 3 / Layer 4); the test-suite gate skips for eval targets that ship without a unified runner. Operator authorization is required before any upstream PR is opened. |

Cycle execution

This cycle was produced by Jelleo's continuous, hypothesis-driven C / systems-software audit loop. Every finding originates as a falsifiable invariant claim from a per-protocol hypothesis library, dispatched to Layer 1 multi-agent recon, promoted on contested verdicts via Layer 1.5 adversarial debate, and confirmed empirically through a Layer 2 clang + ASan/UBSan proof-of-concept. Layer 2.5 triage classifies each fire as `STRONG` / `SOFT` / `FALSE` / `LOST`; only `STRONG` cluster representatives advance to `confirmed` and appear in §01 above. `SOFT` and `STRONG` duplicates land in `triaged`; `FALSE` fires return to `new`. Lifecycle: `new` → `triaged` → `confirmed` → `disclosed` → `fixed` → `verified`. Every cycle is signed Ed25519 against the platform key — see the cover-page receipt.



NON-FIRE ACCOUNTING 6 hypotheses were tested but the PoC did not fire — 5× `rejected`, 1× `new`. These are hypotheses where Layer 1 / Layer 1.5 returned a verdict but the Layer 2 PoC author either declined to produce a test (no plausible attack) or the test ran without an abort in the target module.

CYCLE WALL-CLOCK 2h 52m 48s

§ B.1 – Cycle funnel. Hypotheses tested → PoC fires → Layer 2.5 judge filters out artifactual / mis-invariant fires → surviving `STRONG` fires cluster by code site → cluster representatives become published findings.

— C — AUDIT ARTIFACTS

All cycle artifacts are persisted on disk and verifiable independently of this report. The table below lists the canonical paths under the cycle workspace so a reviewer can re-execute every layer or recompute the cycle Merkle root.

| ARTIFACT | PATH (RELATIVE TO WORKSPACE) |
|---|---|
| Cycle summary (manifest of every step) | hunts/<cycle>/hunt_summary.json |
| Per-step event log | hunts/<cycle>/hunt.log.jsonl |
| Layer 2.5 triage verdicts | hunts/<cycle>/trriage.jsonl |
| Layer 2 PoC sources (C) | tests/c/test_<slug>.c |
| Layer 2 PoC run logs | hunts/<cycle>/poc/c_<slug>.log |
| Layer 3 CBMC harnesses + verdicts | formal/c/harness_<slug>.c |
| Layer 4 AFL++ fuzz harnesses | fuzz/c/<slug>/afl_<slug>.c |
| Layer P3 fix bundles (patch.diff + evidence/ + manifest.json) | recon/bundles/<finding_id>/ |
| Narrative writeups (per finding) | hunts/<cycle>/narratives/<hyp_id>.md |
| Cycle Merkle root (tamper-evidence) | hunts/<cycle>/merkle.json |
| Findings DB (SQLite) | findings.db |
| Ed25519 public key for receipt verification | https://jelleo.com/keys/jelleo.ed25519.pub |

— D — DISCLAIMERS

Findings in this report reflect the state of the engine source at the commit hash on the cover page. Subsequent changes to the codebase are not analyzed. The report is not a guarantee of code correctness or security: it documents invariants that fired (or held) under the hypothesis library applied during this cycle. Out-of-scope items are listed in §00.1 (Scope).

§03 reflects bundle-level state. A row is treated as a confirmed finding when the bundle's machine verification gates (PoC fails pre-patch + PoC passes post-patch + tests still pass) all hold, even if the Layer 2.5 LLM judge initially classified the fire as `SOFT` / `FALSE` / `LOST` — the verifier's empirical patch-defuses-bug evidence supersedes the judge. Rows that did not reach a confirmed lifecycle state are retained in §03 as audit-trail evidence but are not published findings; the authoritative set is whatever appears in §01.

Communication channel: security@jelleo.com (PGP key on jelleo.com/security.html). Coordinated disclosure follows the timeline published in our security policy; pre-disclosure leak protections are enforced at the report level (the `--public` renderer suppresses confirmed-but-not-disclosed findings).

Methodology spec: [docs/methodology/](https://docs.jelleo.com/methodology/) · Live reference: jelleo.com/methodology.html · Source: github.com/Copenhagen0x/audit-pipeline-cli