Pretalx Vulnerabilities: How to get accepted at every conference

Pretalx Vulnerabilities: How to get accepted at every conference - Stefan Schiller, Sonar.

- Published: 2023-04-12
- Original: <https://www.sonarsource.com/blog/pretalx-vulnerabilities-how-to-get-accepted-at-every-conference/>
- Preserved from: https://www.sonarsource.com/blog/pretalx-vulnerabilities-how-to-get-accepted-at-every-conference/ (live) on 2026-08-09
- Licence: unknown

Content

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

[[image: sonar logo]](https://www.sonarsource.com/)

TL;DR overview

- Sonar's research found multiple vulnerabilities in pretalx—an open source conference management system—including cross-site scripting and SQL injection flaws accessible to unauthenticated or low-privileged users submitting talk proposals
- The vulnerabilities stem from insufficient input validation in submission and review features, where user-controlled data reaches database queries or template rendering without adequate sanitization.
- The irony: vulnerabilities in a conference submission system could theoretically allow an attacker to manipulate their submission record or access other speakers' information.
- Conference organizers running self-hosted pretalx instances should apply patches; the research demonstrates the value of scanning open source event infrastructure tools that handle personal data.

Pretalx is a web-based conference planning tool, which is used to manage call for papers (CfP) submissions, select talks, communicate with speakers, and publish conference schedules. Major IT security conferences like OffensiveCon, Hexacon, and TROOPERS are only a few of the numerous users of pretalx. Due to the call for papers functionality, a pretalx instance can contain data about yet undisclosed research, which makes it an interesting target for threat actors.

While submitting talks to some conferences, we wondered how secure the CfP platforms are and decided to audit the popular pretalx for security vulnerabilities. During this research, we identified

an arbitrary file read and a limited file write vulnerability. When determining the impact of these vulnerabilities, we found a **generic technique to turn a file write into code execution** by leveraging a specific feature of Python.

In this article, we outline the impact of the vulnerabilities and dive into the technical details. Furthermore, we introduce the generic technique to gain code execution via a file write vulnerability. In the end, we explain how the vulnerabilities can be mitigated by having a look at the applied patches.

Impact

We discovered the following vulnerabilities in pretalx, which affect versions `2.3.1` and prior:

- CVE-2023-28459: Arbitrary File Read
- CVE-2023-28458: Limited File Write

The first vulnerability allows a privileged user to **disclose any file** from the server's filesystem, which is accessible by the pretalx process.

The second vulnerability allows a user with access to a scheduled talk to write files on the server's filesystem. If the application is running in **debug mode**, the content of these files can be controlled, which leads to **remote code execution**.

Both vulnerabilities were fixed in pretalx version `2.3.2`, which was released in an incredible time of fewer than 3 hours after our notification. The SaaS platform pretalx.com was immediately patched. We strongly recommend updating any self-hosted instance with a version before this release.

Technical Details

In this section, we dive into the technical details of both vulnerabilities.

Arbitrary File Read (CVE-2023-28459)

Pretalx allows privileged users to create and download a static HTML export of a schedule. The creation of the exported HTML is also triggered automatically on a regular basis, usually via a cron job.

The function responsible for creating the export performs the following steps:

- Iterate over all URLs required for the schedule.
- Dump its content to a temporary folder which will later be archived in a zip file.
- Retrieve all URLs to additional assets.
- Dump all additional assets in a second iteration.

Since user-uploaded resources are also part of the schedule, attackers can make the application process arbitrary URLs in the second iteration by uploading an HTML file that references an asset using an `img` tag's `src` attribute.

URLs beginning with `STATIC_ROOT` or `MEDIA_ROOT` will first be read directly from disk:

Copy to clipboard

```
def get_mediastatic_content(url):
 if url.startswith(settings.STATIC_URL):
  local_path = settings.STATIC_ROOT / url[len(settings.STATIC_URL):]
 elif url.startswith(settings.MEDIA_URL):
  local_path = settings.MEDIA_ROOT / url[len(settings.MEDIA_URL):]
 else:
  raise FileNotFoundError()

 with open(local_path, "rb") as f:
  return f.read()
```

Since there is no check whether the final `local_path` is within the `STATIC_ROOT` or `MEDIA_ROOT` folder, arbitrary files can be referenced using the path traversal sequence `../`. This can also be achieved by using an absolute path. If the second part of the path begins with a slash (`/`), the first part of the path is ignored:

Copy to clipboard

```
MEDIA_ROOT = Path('/var/pretalx/data/media')
MEDIA_URL = '/media/'

url = '/media//etc/passwd'
local_path = MEDIA_ROOT / url[len(MEDIA_URL):]

print(local_path)
# '/etc/passwd'
```

You can read more about similar security pitfalls in Python in our blog post on 10 Unknown Security Pitfalls for Python.

Limited File Write (CVE-2023-28458)

The second vulnerability also resides within the HTML export feature. The function responsible for dumping the content retrieved from a URL is called `dump_content` and uses the URL (parameter `path`) to determine the destination path. Although leading slashes are removed from `path` before being added to the destination folder, it is not ensured that the final path is below the destination folder:

Copy to clipboard

```
def dump_content(destination, path, getter):
  # retrieve content (path is the URL)
  content = getter(path)

  # create folders if necessary
  path = Path(destination) / path.lstrip("/")
  path.parent.mkdir(parents=True, exist_ok=True)

  # write content to file
  with open(path, "wb") as f:
    f.write(content)
  return content
```

Attackers can again leverage the string sequence `../` to traverse out of the destination folder, resulting in arbitrary file write. This could be exploited by a self-registered user with access to a talk that has been added to a schedule. However, the content of the file cannot be controlled in most cases, because referencing an invalid URL returns the 404 error page. This is different if the application is running in `DEBUG` mode as we will demonstrate now.

In `DEBUG` mode, user-uploaded resources are served from the Django application itself instead of a reverse proxy. When the content of a URL is retrieved and cannot be read from disk via the `get_mediastatic_content` function, the Django test client is used to read the content from the application:

Copy to clipboard

```
def get(url):
  try:
    # Try getting the file from disk directly first, ...
    return get_mediastatic_content(url)
  except FileNotFoundError:
    # ... then fall back to asking the views.
    response = client.get(url, is_html_export=True, HTTP_ACCEPT="text/html")
    content = get_content(response)
    return content
```

In order to make the application read a user-uploaded resource via the Django test client, an attacker can simply URL-encode one of the first characters to prevent that the URL begins with `MEDIA_ROOT`. This is possible because the Django test client decodes the URL before accessing it.

The URL-decoding of the Django test client also introduces a significant difference from the filesystem path handling when the content is written to disk. The following URL is a valid reference to a user-uploaded resource for the Django test client:

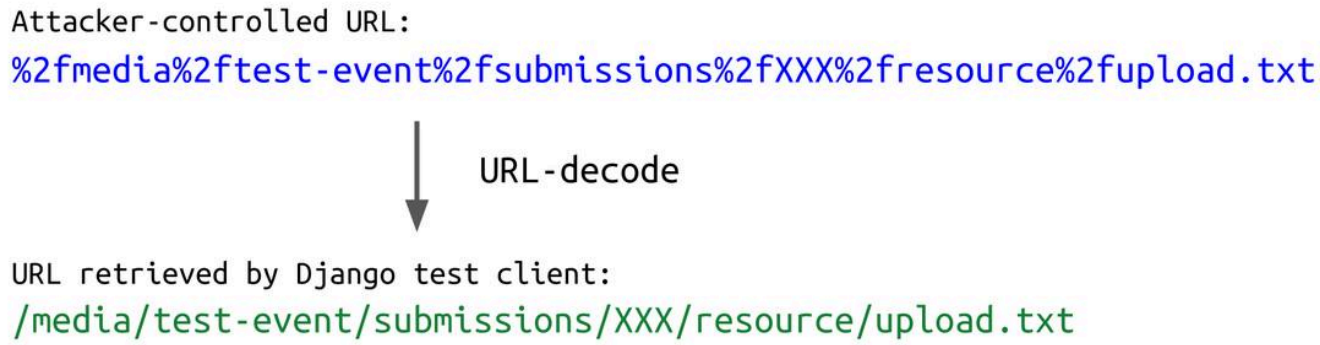

When the retrieved contents are written to disk, though, the path is **not** URL-decoded. This means that the following file is written:

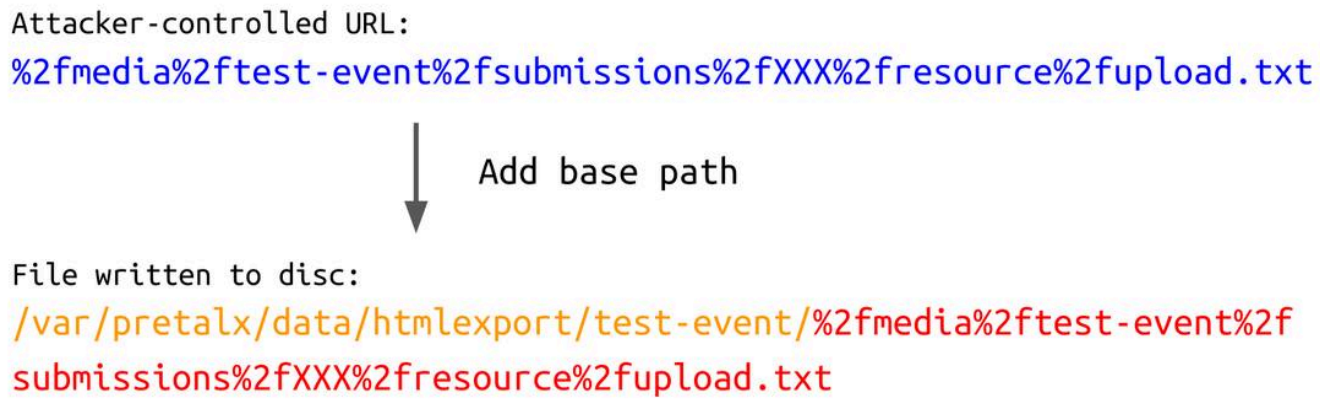

In conjunction with the path traversal, this can be leveraged to write the user-controlled resource to an arbitrary file. The following URL is still a valid reference to the user-uploaded resource for the Django test client:

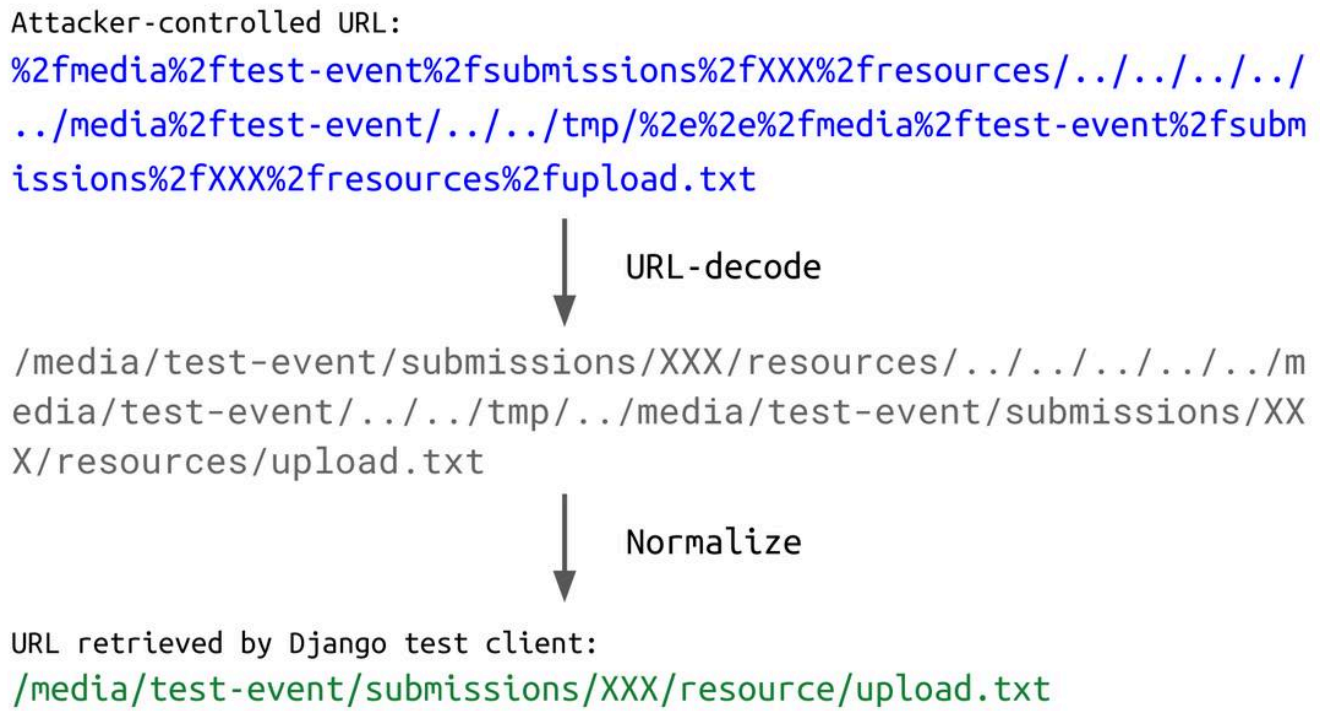

When the path is processed to write the file to disk, things look a little different:

Thus, the user-controlled content is written to `/tmp` in a file called `%2e%2e%2fmedia%2ftest-event%2fsubmissions%2fXXX%2fresources%2fupload.txt`.

When determining the impact of this arbitrary file write, we discovered a generic technique to gain code execution.

Code Execution via Site-Specific Configuration Hooks

The requirements for this technique are the following:

- Control over the file extension
- Control over the beginning of any line in the file
- Ability to write the file to `~/.local/lib/pythonX.Y/site-packages/`
- At some point, a new Python process is launched with the same identity

Python supports a feature called site-specific configuration hooks. Its main purpose is to add custom paths to the module search path. To do this, a `.pth` file with an arbitrary name can be put in the `.local/lib/pythonX.Y/site-packages/` folder in a user's home directory:

Copy to clipboard

```
user@host:~$ echo '/tmp' > ~/.local/lib/python3.10/site-packages/foo.pth
```

When a new Python process is spawned, the path `/tmp` is added to the module search path `sys.path`:

Copy to clipboard

```
user@host:~$ python3
>>> import sys
>>> sys.path
[ ... '/tmp', ... ]
```

Although there might be some cases where the ability to add a path to the module search path can be leveraged to gain code execution, bear with us. It even gets better.

Looking at the implementation of the site-specific configuration, the following part screams for attention:

Copy to clipboard

```
def addpackage(sitedir, name, known_paths):
 # ...
 for n, line in enumerate(f):
  # ...
  try:
   if line.startswith(("import ", "import\t")):
    exec(line)
    continue
```

If a line in a `.pth` file starts with `"import "` or `"import\t"`, it will be evaluated as Python code! This is clearly described in the docs:

[...] Lines starting with import (followed by space or tab) are executed. [...]

This video from anthonywritescode also mentions that this feature could be used to gain arbitrary code execution.

So, let's see this in action. We create a new `.pth` file, which pipes the output of the `whoami` command to `/tmp/x`. Once a new Python process is spawned, the command is executed:

Copy to clipboard

```
user@host:~$ echo 'import os;os.system("whoami>/tmp/x")' > .local/lib/python3.10/site-packages/arbitrary_name.pth
user@host:~$ cat /tmp/x
cat: /tmp/x: No such file or directory
user@host:~$ python3
>>> CTRL + D
user@host:~$ cat /tmp/x
user
```

The fact that this technique only requires a limited amount of file control makes it very appealing. Having control over the extension of a file and one single line in it is not that uncommon. The most

restrictive requirements are that the destination path must be controllable and that a new Python process is spawned in the context of the targeted user.

Regarding the destination path, it is worth mentioning that the base path, from which the `.pth` files are read, can be changed via the `PYTHONUSERBASE` environment variable. If a file write vulnerability does not allow to write to the user's home directory but it is possible to influence this environment variable of any spawned Python process, the `.pth` file can be stored in another directory (the subfolders `lib/pythonX.Y/site-packages/` are still required):

Copy to clipboard

```
user@host:~$ cat /tmp/x
cat: /tmp/x: No such file or directory
user@host:~$ mkdir -p /tmp/lib/python3.10/site-packages
user@host:~$ echo 'import os;os.system("whoami>/tmp/x")' > /tmp/lib/python3.10/site-packages/some_name.pth
user@host:~$ export PYTHONUSERBASE='/tmp'
user@host:~$ python3
>>> CTRL + D
user@host:~$ cat /tmp/x
user
```

Regarding pretalx, this technique can be used to turn the file write into code execution, if running in `DEBUG` mode. The payload is executed once a new Python process is spawned to perform periodic tasks.

Patch

The file read vulnerability in pretalx was fixed by first resolving the `local_path` and then ensuring that it is either within the `MEDIA_ROOT` or `STATIC_ROOT`:

Copy to clipboard

```
def get_mediastatic_content(url):
 # ...
 # Prevent directory traversal, make sure the path is inside the media or static root
 local_path = local_path.resolve(strict=True)
 if not any(
  path in local_path.parents
  for path in (settings.MEDIA_ROOT, settings.STATIC_ROOT)
 ):
  raise FileNotFoundError()
```

Similarly, the file write vulnerability was fixed by first resolving the destination path and then ensuring that it is below the destination folder:

Copy to clipboard

```
def dump_content(destination, path, getter):
  # ...
  path = (Path(destination) / path.lstrip("/")).resolve()
  if not Path(destination) in path.parents:
    raise CommandError("Path traversal detected, aborting.")
```

Timeline

Summary

In this article, we detailed a file read and file write vulnerability we discovered in the conference planning tool pretalx. Furthermore, we introduced a generic technique to turn a file write vulnerability into code execution by leveraging Python's `.pth` files. We also learned how to prevent these vulnerabilities by looking at the applied patches.

At last, we would like to thank the pretalx maintainer for acknowledging the issues and providing a patch in an astonishing time of fewer than 3 hours.

- 10 Unknown Security Pitfalls for Python
- Disclosing information with a side-channel in Django