



Bug 1835736 (CVE-2018-20225) - CVE-2018-20225 python-pip: when --extra-index-url option is used and package does not already exist in the public index, the installation of malicious package with arbitrary version number is possible.

Keywords: Security

Reported: 2020-05-14 12:10 UTC by Michael Kaplan

Status: CLOSED WONTFIX

Modified: 2025-02-10 12:04 UTC ([History](#))

Alias: CVE-2018-20225

CC List: 22 users ([show](#))

Product: Security Response

Fixed In Version:

Component: vulnerability

Clone Of:

Environment:

Version: unspecified

Last Closed: 2020-05-25 21:15:19 UTC

Embargoed:

Hardware: All

OS: Linux

Priority: low

Severity: low

Target Milestone: ---

Assignee: Red Hat Product Security

QA Contact:

Docs Contact:

URL:

Whiteboard:

Depends On: [1835737](#) [1835742](#)

Blocks: [1835743](#)

TreeView+ [depends on](#) / [blocked](#)

Attachments ([Terms of Use](#))

Michael Kaplan 2020-05-14 12:10:14 UTC

[Description](#)

A vulnerability was discovered in pip (all versions), because it installs the version with the highest version number, even if the user had intended to obtain a private package from a private index. This only affects use of the --extra-index-url option, and exploitation requires that the package does not already exist in the public index (and thus the attacker can put the package there with an arbitrary version number).

Michael Kaplan 2020-05-14 12:10:37 UTC

[Comment 1](#)

Created python-pip tracking bugs for this issue:

Affects: fedora-all [[bug-1835737](#)]

Michael Kaplan 2020-05-14 12:12:38 UTC

[Comment 2](#)

External References:

<https://cowlicks.website/posts/arbitrary-code-execution-from-pips-extra-index-url.html>

Michael Kaplan 2020-05-14 12:17:12 UTC

[Comment 3](#)

Created python-pip-epel tracking bugs for this issue:

Affects: epel-7 [[bug-1835742](#)]

Petr Viktorin (pviktori) 2020-05-14 12:40:45 UTC

[Comment 4](#)

I don't think we can do anything on the pip side.

The --extra-index-url means an *extra* source of packages; it's right there in the name. If you want to use a private server, pass --index-url

If you need some packages in PyPI, add them to the private server. Use a more advanced server like devpi if you need more advanced mirroring.

Anyone can upload packages to PyPI; there's no vetting of the uploads, so there can definitely be malicious packages.

Miro Hrončok 2020-05-14 13:59:19 UTC

[Comment 5](#)

I agree with Petr. This is like reporting a CVE in curl: "Malicious content found on the internet"

Lumír Balhar 2020-05-15 04:30:41 UTC

[Comment 6](#)

As the author of the CVE states in the post mentioned in the [comment#2](#): "I disclosed this to the security list. Unfortunately they said there is currently no path to fix this. They recommended using "version-pinning and hash-pinning for deployments" to avoid this."

Christian Heimes 2020-05-15 09:18:25 UTC

[Comment 7](#)

By the way dnf, yum, and apt have basically the same kind of "vulnerability". Distro package managers also select the highest version of a package across all enabled repositories.

Miro Hrončok 2020-05-15 09:39:37 UTC

[Comment 8](#)

Michael, please let us know if you don't agree. If you agree, I'll close the Fedora/EPEL bugzillas as NOTABUG.

Michael Kaplan 2020-05-15 10:03:57 UTC

[Comment 9](#)

In reply to [comment #8](#):
> Michael, please let us know if you don't agree. If you agree, I'll close the
> Fedora/EPEL bugzillas as NOTABUG.

I agree with this, Please do. Thanks

Jason Shepherd 2020-05-18 19:49:40 UTC

[Comment 10](#)

While the version of pip used to build Red Hat Quay is affected by this, the --extra-index-url option is not used. Therefore we will not fix this issue in Red Hat Quay.

Todd Cullum 2020-05-22 20:23:00 UTC

[Comment 11](#)

According to the pip documentation[1]:

```
====BEGIN QUOTATION====  
-i, --index-url <url>
```

Base URL of the Python Package Index (default <https://pypi.org/simple>). This should point to a repository compliant with PEP 503 (the simple repository API) or a local

directory laid out in the same format.

```
--extra-index-url <url>
```

Extra URLs of package indexes to use in addition to --index-url. Should follow the same rules as --index-url.
====END QUOTATION====

This means that this "flaw" is actually intended behavior as per the docs. The default index-url when --index-url is omitted is <https://pypi.org/simple>, thus when --extra-index-url is used, pip will try to grab the latest package from there, hence the "extra" in --extra-index-url. To protect from any unintended behavior, simply use --index-url and do not use --extra-index-url OR explicitly set --index-url and use --extra-index-url. Blake Griffith, the reporter of this issue, states[2] that some CERN packages use --extra-index-url without --index-url specified[3], thus pip will by default, grab the latest version of the package from PyPI before using the specified PyPI server, which is apparently what they intended. As per the pip documentation, this is an incorrect understanding/usage of --extra-index-url. For this reason, this issue is WONTFIX for pip as shipped with Red Hat Enterprise Linux and Red Hat Software Collections.

1. https://pip.pypa.io/en/stable/reference/pip_install/#install-index-url
2. <https://cowlicks.website/posts/arbitrary-code-execution-from-pips-extra-index-url.html>
3. <https://twiki.cern.ch/twiki/bin/view/LHCb/PythonPyPIServer>

Todd Cullum 2020-05-22 20:25:26 UTC

[Comment 12](#)

Mitigation:

To protect from any unintended behavior, use --index-url and do not use --extra-index-url OR explicitly set --index-url and use --extra-index-url.

Todd Cullum 2020-05-22 20:34:28 UTC

[Comment 15](#)

Also note that the CERN example above uses --trusted-host which ignores HTTPS requirement and could lead to a MITM[1], which is also intended functionality in this case.

1. <https://pip.pypa.io/en/stable/reference/pip/#trusted-host>

Marco Benatto 2020-05-25 15:22:34 UTC

[Comment 16](#)

Statement:

Although this issue affects versions of pip shipped with Red Hat Enterprise Linux, Red Hat Software Collections and Red Hat CodeReady Workspaces, Red Hat Product Security does not consider this to be a security vulnerability because per the pip documentation, this is intended behavior of pip when using the --extra-index-url flag. Therefore, this issue has been marked WONTFIX.

Product Security DevOps Team 2020-05-25 21:15:19 UTC

[Comment 17](#)

This bug is now closed. Further updates for individual products will be reflected on the CVE page(s):

<https://access.redhat.com/security/cve/cve-2018-20225>

cowlicks 2020-05-26 21:20:50 UTC

[Comment 18](#)

(In reply to Christian Heimes from [comment #7](#))
> By the way dnf, yum, and apt have basically the same kind of
> "vulnerability". Distro package managers also select the
> highest version of
> a package across all enabled repositories.

Distro package managers (dnf, yum, apt, and others) aren't normally vulnerable to this since only vetted packages are allowed in the default repositories.

cowlicks 2020-05-26 22:18:44 UTC

[Comment 19](#)

Hi, I'm Blake Griffith

> This means that this "flaw" is actually intended behavior as per the docs.

I think this is true in some sense. However, I don't think the current situation was foreseen when --extra-index-url was added.

> To protect from any unintended behavior, simply use
> --index-url and do not use --extra-index-url OR explicitly
> set --index-url
> and use --extra-index-url.

As stated, this is insufficient.

> Blake Griffith, the reporter of this issue,
> states[2] that some CERN packages use --extra-index-url
> without --index-url
> specified[3], thus pip will by default, grab the latest

```
version of the
> package from PyPI before using the specified PyPI server,
which is
> apparently what they intended.
```

This is not what CERN intended. If they had, they would have claimed the names of their packages on PyPI.

```
> As per the pip documentation, this is an
> incorrect understanding/usage of --extra-index-url.
```

I think you are saying CERN (not me?) has an incorrect understanding. I agree. However, it is a very common misunderstanding.

To demonstrate, I found another vulnerable package while writing this. Here is what I did:

I made this google query:

```
https://www.google.com/search?
ei=i43NXvXCGMrA0PEPzK03kAg&q=%22pip+install+--extra-index-
url%22&oq=%22pip+install+--extra-index-
url%22&gs_lcp=CgZwc3ktYWIQAzoECAAQ1DADFihKMDcLWgAcAF4AIABRogB
iwGSAQEymAEAoAEBqgEHZ3dzLXdpeg&sclient=psy-
ab&ved=0ahUKewi1tLHov9LpAhVKIDQIHczRDYIQ4dUDCAw&uact=5
```

The fourth result (for me) was this:

```
https://pypi.org/project/cngi-prototype/0.0.22/
```

There they say "pip install -extra-index-url <https://casa-pip.nrao.edu/repository/pypi-group/simple> casatools"

I checked if the name "casatools" was taken on PyPI. It wasn't.

So I used a script I wrote to squat that name. It now exists at <https://pypi.org/project/casatools/99999.9.9/>

So now, when anyone follows the instructions given by the 'cngi-prototype' package, they will install my package.

This was just to demonstrate how easy this is. There are **many** more examples you can find with some googling.

My point is the "incorrect understand" is so wide-spread pip itself should probably do something to fix it.

Thanks,

Petr Viktorin (pviktori) 2020-05-27 06:34:09 UTC

[Comment 20](#)

What do you suggest? Removing the --extra-index-url option? It can be useful when used correctly.

cowlicks 2020-05-28 23:46:17 UTC

[Comment 21](#)

(In reply to Petr Viktorin from [comment #20](#))
> What do you suggest?

Are you asking what I suggest Red Hat should recommend for mitigation? I think the security folks can give a better answer than me.

In the thread for this issue there glyph said:
"it's a best practice to always pin all of your dependencies to exact versions, and ideally, more than that, to hashes"

So maybe something like that

cowlicks 2021-02-12 17:59:18 UTC

[Comment 22](#)

Hello again,

Someone recently rediscovered this bug, and got a lot of press for it.

<https://medium.com/@alex.birsan/dependency-confusion-4a5d60fec610>

Just thought y'all might want to know, in case you want to revisit this.

Cheers,
Blake

~~Petr Viktorin (pviktori)~~ 2021-02-17 13:45:48 UTC[Comment 23](#)

There are also a lot of other kinds of unsafe instructions, like ``curl ... | sudo bash``, or ``sudo pip install ...``. And a lot of projects with pinned but unreviewed dependencies.

> In the thread for this issue there glyph said:
> "it's a best practice to always pin all of your dependencies to exact versions, and ideally, more than that, to hashes"

Nope, best practice is to review all code you're using. Which is what we provide RPM packages. (The archive format is a small detail compared to the reviews; we use RPM so the can be the same across language ecosystems.)

The downside of that is that reviews are costly, so people take shortcuts, follow dubious instructions online, and use PyPI without any verification. I don't think that's a reason for Red Hat to patch pip and deviate from its upstream behavior.

(Disclaimer: I'm not the authority here and I'm not on the

security team.)

Note

You need to [log in](#) before you can comment on or make changes to this bug.

