THE TIDELIFT MAINTAINER ADVANTAGE

How a Python project with 450 million monthly downloads improved its security practices

urllib3 is a HTTP client for Python. It sits at the underpinning of many components of the Python ecosystem:

- **requests,** a high-level HTTP library, is built on top of urllib3
- **boto,** the Amazon AWS SDK library, uses urllib3 to communicate with AWS
- **pip,** the Python package manager, uses urllib3

And that's just a fraction of its uses—urllib3 is used, directly or indirectly, by nearly 1.5 million direct dependent repositories. It's regularly among the top 3 downloaded Python projects. If you're building with Python, you're using urllib3.

For enterprise users, including many of Tidelift's largest customers, the security of urllib3 is particularly important because urllib3 handles web requests, TLS/SSL, certificate validation, and more. If any of that functionality has issues, it can leave users open to compromise from hostile websites, man-in-the-middle attacks, and more.

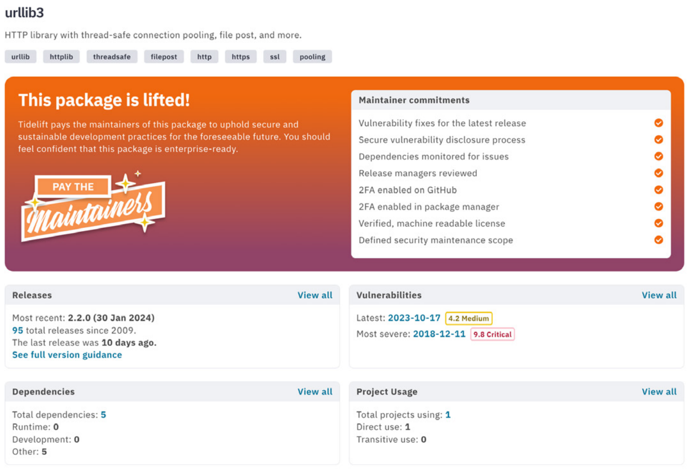

urllib3

HTTP library with thread-safe connection pooling, file post, and more.

urllib httplib threadsafe filepost http https ssl pooling

This package is lifted!

Tidelift pays the maintainers of this package to uphold secure and sustainable development practices for the foreseeable future. You should feel confident that this package is enterprise-ready.

Maintainer commitments

- Vulnerability fixes for the latest release ✓
- Secure vulnerability disclosure process ✓
- Dependencies monitored for issues ✓
- Release managers reviewed ✓
- 2FA enabled on GitHub ✓
- 2FA enabled in package manager ✓
- Verified, machine readable license ✓
- Defined security maintenance scope ✓

Releases View all

Most recent: 2.2.0 (30 Jan 2024)
95 total releases since 2009.
The last release was 10 days ago.
See full version guidance

Vulnerabilities View all

Latest: 2023-10-17 4.2 Medium
Most severe: 2018-12-11 9.8 Critical

Dependencies View all

Total dependencies: 5
Runtime: 0
Development: 0
Other: 5

Project Usage View all

Total projects using: 1
Direct use: 1
Transitive use: 0

Here is the urllib3 package page inside of the Tidelift application, showing the secure development practices the project is following, and also including additional useful information about releases, vulnerabilities, dependencies, and where it is being used within an example organization.

Also, the sheer scope of its usage—urllib3 is downloaded over 450 million times per month as part of usage, packaging, installation, and continuous integration workflows—means that it is especially important that the package stays well maintained and has a documented set of security practices and a long term security plan in place.

urllib3 is co-maintained by a team of maintainers, including lead maintainer Seth Michael Larson, Andrey Petrov (the original author of urllib3) and Quentin Pradet (who does much of the day to day development on the project).

In 2021, attackers were able to infiltrate Codecov, a service used for testing code coverage by over 20,000 projects and enterprises. By discovering a flaw in how Codecov built their images, they were able to modify a Codecov artifact that allowed the attackers to gain access to all environment variables that Codecov users used when accessing the service for a period of two months. The leaked environment variables could include passwords, tokens, and other secrets. Multiple organizations had their private source code and service access tokens leaked to the attackers.

When learning of this attack, the urllib3 team was able to investigate what variables were leaked, audit to confirm there wasn't any unauthorized use of the leaked API token, and then change the token going forwards. In addition, the team put in place a solution using short-lived tokens that are tightly bound to a specific workflow, another best practice to eliminate risk from attacks like the one that impacted Codecov. If they hadn't caught this issue quickly, it's possible that the API token could have been used to compromise accounts or release trojaned software. Many users of urllib3 may not even have been aware this breach happened, but the quick work of the maintainers ensured their supply chain was not compromised.

Many open source projects maintained by volunteer maintainers are not able to invest time in putting in place robust security practices to avoid situations like this. Thankfully, this is not the case with urllib3 in part due to income the project maintainers receive from Tidelift and its customers who rely on urllib3.

The urllib3 maintainers know how critical their software is to the Python ecosystem and the internet at large, so they've used the income from Tidelift to make significant investments in improving the secure development practices they have in place. A few examples:

- **Securing maintainer access:** They have performed a number of steps to harden their accounts, including implementing 2FA, separating permissions for their contributions between reviewers, and using API tokens with limited scope for all build, test, and release processes.
- **Backwards compatibility:** The team has worked to maintain backwards compatibility where each change passes a rigorous test suite to ensure there are no unintended breaks to functionality. This ensures that if they do need to issue an important change that users will be able to quickly pick up a release.
- **Automating and streamlining release processes:** The team built a checklist-based automated release process with a built-in ability to get approvals from maintainers. By automating the release process, they can ensure it is run the same way every time, embed any credentials that are needed, and scope them to only the permissions that are required. This ensures that the release process is safe, repeatable, and reliable, and that their users are protected from potential compromise and accidental mistakes.
- **Reproducible and verifiable builds:** By working to create reproducible builds with a chain of trust from source through the build system to the final artifact, the urllib3 team can avoid some particularly sneaky types of software supply chain attacks, and ensure that their build processes are not silently compromised.
- **Improving scores against industry standards:** The urllib3 team is constantly looking for ways to improve the security posture of their code, including against common industry standards like the OpenSSF scorecards. After a few weeks of investment, they became the first Python project to score at least 9.0 out of 10, and now score a 9.6 out of 10.

In his end of year report on urllib3 for 2023, Seth Michael Larson reported that Tidelift (and by proxy its customers that use urllib3) provided the primary income for this work, and helped enable them to bring on an additional co-maintainer, Illia Volochii, whose work has been critical to getting v 2.0 of urllib3 released.

Now, Tidelift customers can use urllib3—and other packages that rely on it—with confidence, knowing that a team of experienced maintainers are committed to ensuring the package follows a robust set of enterprise secure software development practices. What's more, Tidelift customers directly played a role in funding work on things like urllib3 v 2.0, ensuring it stays resilient and healthy into the future.