Donald Stufft

Staff Software Engineer

Staff Software Engineer with 17+ years of experience. I enjoy tackling large scale projects that impact tens of millions of users; my skills are at the core of PyPI's success as the package index for the world's most popular programming language.

Contact Info

Philadelphia, PA
linkedin.com/in/dstufft
github.com/dstufft
donald@stufft.io
484-235-5050

Programming Languages

Python
Rust
Go
JavaScript

Skills

API Design
Build Systems
Cloud Native Design
Developer Experience
Distributed Systems
Large Systems Design
Open Source
Packaging
Security
Supply Chain Security
Technical Leadership
Threat Modeling

Work Experience

Datadog

Staff Software Engineer | Jan 2024 - Apr 2025

- Designed the system for managing the lifecycle and security of VCS repositories, working with stakeholders from over 5 teams.
- Designed and led the implementation of the system for automated dependency upgrades, ensuring critical vulnerabilities were resolved.
- Designed and led the implementation of a system for proactively tracking and mirroring upstream packages and related metadata, which was able to ingest new upstream releases in under 5 minutes.

Senior Software Engineer | Jan 2020 - Jan 2024

- Refactored deployment of internal package mirror, reducing p95 latency to sub ms as well as limiting the blast radius of a bad deploy with automatic error detection.
- Migrated internal package mirror data storage from OCI to S3.
- Built system for reactive mirroring of upstream packages for Python, Go, Java, and Bazel.
- Developed plan for and started initial migration of a large Python 2 code base to Python 3, which included migrating to a completely different web framework as part of the migration.
- Designed and built a system for automatically scanning images published to an internal container registry for CVEs, working with 2 other teams to ensure it satisfied compliance requirements.

Amazon Web Services

Software Development Engineer II | Dec 2016 - Oct 2019

- Provided support for AWS customers using the Python SDK for AWS.
- Worked with internal product teams to provide customization to how their API would be exposed via Python.
- Refactored the packaging for aws-cli and the Python SDK to use modern standards.

Hewlett Packard Enterprise

Systems/Software Engineer V | May 2015 - Oct 2016

• Implemented a new PyPI code base from scratch designed to be maintainable and modernized.

• Built a statistics gathering pipeline that took 500 million incoming download events a month, processed them, and stored them durably for later querying.

• Worked with downstream distributors of pip to ensure that they were able to distribute the latest versions of pip in a consistent and stable way.

• Founding member of the Python Packaging Working Group.

• Lead a growing team of PyPI and Warehouse contributors, including developers, a designer, and a project manager.

Rackspace

Software Developer V | Dec 2013 - May 2015

• Built a library for interfacing with a Certificate Authority vendor for issuing x509 certificates.

• Designed a new versioning standard for Python projects (PEP 440) and implemented support within PyPI and pip.

• Removed insecure methods of downloading and installing files from PyPI.

Nebula

Software Engineer | Oct 2012 - Dec 2013

• Implemented remote support capabilities which injected one time use SSH keys across an OpenStack cluster.

• Implemented an automatic security testing framework for a Linux OS cluster to ensure a consistent secured state on all hosts in the product.

• Implemented a secure web based password management tool for customer passwords.

Eldarion

Developer | Jun 2011 - Sep 2012

- Implemented custom web applications for clients in Django.
- Contributed to Gondor; a PaaS deployment engine for Django.

Xenofox

Owner / Developer | Jan 2008 - Jun 2011

- Implemented custom web applications for clients primarily using Django.
- Ported existing PHP and flat file websites into custom written CMS.

Open Source Contributions

PyPI

2013 - Current

- Led the design and implementation of the rewrite of PyPI from the legacy code base to the new, modern "Warehouse" code base.
- Designed and implemented the CDN caching strategy, allowing PyPI to scale to serving billions of requests and petabytes of traffic per month with a handful of servers and mostly volunteers.
- Designed and implemented the standard JSON API for package installation (PEP 691).
- Designed and implemented support for "Yanking" files (PEP 592).
- Designed, implemented, and managed the removal of "external file hosting" on PyPI (PEP 470).
- Optimized the legacy JSON API to require 50% less database queries and reduce the average response time by over 80%.
- Designed and implemented the Macaroon based API token authentication and authorization mechanism for PyPI uploads.
- Prevented 100% of compromised passwords from being used on PyPI by checking passwords against known data breaches at authentication time.

Python Packaging

- Designed and implemented the versioning standard for Python packages (PEP 440).
- Designed and implemented the bootstrapping of pip with Python and within the stdlib venv module (PEP 453).
- Helped design the mechanism for defining pluggable "build backends" for Python packages, introducing the pyproject.toml file which provided a foundation for nearly all standards since (PEP 518).
- Helped design the mechanism for interoperability between OS provided Python installations and Python packaging, working with stakeholders across multiple Linux distributions as well as the greater Python community (PEP 668).