University of Groningen

A Low-Effort Analytics Platform for Visualizing Evolving Flask-Based Python Web Services

Vogel, Patrick; Klooster, Thijs; Andrikopoulos, Vasilios; Lungu, Micea-Filip

DOI:

[10.1109/VISSOFT.2017.13](https://doi.org/10.1109/VISSOFT.2017.13)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Vogel, P., Klooster, T., Andrikopoulos, V., & Lungu, M-F. (2017). *A Low-Effort Analytics Platform for Visualizing Evolving Flask-Based Python Web Services*. 109-113. Paper presented at 2017 IEEE Working Conference on Software Visualization, Shanghai, China. <https://doi.org/10.1109/VISSOFT.2017.13>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

A Low-Effort Analytics Platform for Visualizing Evolving Flask-Based Python Web Services

Patrick Vogel, Thijs Klooster, Vasilios Andrikopoulos, Mircea Lungu

Johann Bernoulli Institute for Mathematics and Computer Science
University of Groningen, Netherland

Email: {t.klooster,l.p.p.vogel}@student.rug.nl, {v.andrikopoulos,m.f.lungu}@rug.nl

Abstract—Tens of thousands of web applications are written in Flask, a Python-based web framework. Despite a rich ecosystem of extensions, there is none that supports the developer in gaining insight into the evolving performance of their service. In this paper, we introduce Flask Dashboard, a library that addresses this problem. We present the ease with which the library can be integrated in an already existing web application, discuss some of the visualization perspectives that the library provides and point to some future challenges for similar libraries.

I. INTRODUCTION

There is no getting around it: you are building a distributed system argues a recent article [1]. Indeed, even the simplest second-year student project is a web application implemented as two-tier architecture with a Javascript/HTML5 front-end a service backend, usually a REST API.

Python is one of the most popular programming language choices for implementing the back-end of web applications. GitHub contains more than 500K open source Python projects and the Tiobe Index¹ ranks Python as the 4th most popular programming language as of June 2016.

Within the Python community, Flask² is a very popular web framework³. It provides simplicity and flexibility by implementing a bare-minimum web server, and thus advertises as a micro-framework. The Flask tutorial shows how setting up a simple Flask “Hello World” web-service requires no more than 5 lines of Python code [2].

Despite their popularity, to the best of our knowledge, there is no simple solution for monitoring the evolving performance of Flask web applications. Thus, every one of the developers of these projects faces one of the following options when confronted with the need of gathering insight into the runtime behavior of their implemented services:

- 1) Use a commercial monitoring tool which treats the subject API as a black-box (e.g. Pingdom, Runscope).
- 2) Implement their own ad-hoc analytics solution, having to reinvent basic visualization and interaction strategies.
- 3) Live without analytics insight into their services.

¹TIOBE programming community index is a measure of popularity of programming languages, created and maintained by the TIOBE Company based in Eindhoven, the Netherlands

²<http://flask.pocoo.org/>

³More than 25K projects on GitHub (5% of all Python projects) are implemented with Flask (cf. a GitHub search for “language:Python Flask”)

For projects on a budget (e.g. research, startups) the first and the second options are often not available due to time and financial constraints. Even when using 3rd-party analytics solutions, a critical insight into the evolution of the exposed services of the web application, is missing because such solutions have no notion of versioning and no integration with the development life cycle. [3]

To avoid projects ending up in the third situation, that of living without analytics, in this paper we present Flask Dashboard — a low-effort service monitoring library for Flask-based Python web services that is easy to integrate and enables the *agile assessment* of service evolution. [4]

As a case study, on which we will illustrate our solution, we are going to use an open source API which was in the third situation presented above for more than one year.

II. CASE STUDY

Zeeguu⁴ is a platform and an ecosystem of applications for accelerating vocabulary acquisition in a foreign language [5]. The architecture of the ecosystem has at its core an API implemented with Flask and Python and a series of satellite applications that together offer three main intertwined features for the learner:

- 1) Reader applications that provide effortless translations for those texts which are too difficult for the readers.
- 2) Interactive exercises personally generated based on the preferences and past likes of the learner.
- 3) Article recommendations which are at the appropriate level of difficulty for the reader. The difficulty is estimated based on past exercise and reading activity.

The core API provides correspondingly three types of functionality: contextual translations, article recommendations, and personalized exercise suggestions. The core API of system is a research project, which sustains at the moment of writing this article the reading and practice of about two hundred active beta-tester accounts.

In the remainder of this paper, we will use the Zeeguu API as a case study. All the figures in this paper are captured from the actual deployment of Flask Dashboard in the context of the Zeeguu API⁵.

⁴<https://www.zeeguu.unibe.ch/>

⁵Within the Flask Dashboard the figures are interactive offering basic data exploration capabilities: filter, zoom, and details on demand[6]

III. THE FLASK DASHBOARD

In this paper we are introducing Flask Dashboard, a drop-in Python library that allows developers to monitor their Flask-based Python web applications with minimal effort.

The Flask Dashboard as well as the web application that is being monitored in the case study is written in Python using Flask. This makes binding to the web services of the application relatively easy, as well as adding additional routes to the service for interacting with the Flask Dashboard.

To start using our Python library for service visualization, and assuming Flask is already installed, one needs to install the Python package⁶ and simply add two lines of code to their Flask web service:

```
import dashboard
...
# flask_app is the Flask app object
dashboard.bind(flask_app)
...
```

After binding to the service, the Flask Dashboard becomes available at the /dashboard route of the Flask application. A custom route can also be defined by the programmer in a configuration file.

During binding, the Flask Dashboard will search for all endpoints defined in the target application. These will be presented to the user in the tool web interface, where the user can select the ones that should be monitored, see Fig. 1.

Rule	HTTP Method	Endpoint	Last accessed ▾	Monitor
/static/<path.filename>	OPTIONS, HEAD, GET	static	2017-06-23 23:41:11	☐
/user_words	OPTIONS, HEAD, GET	api.studied_words	2017-06-23 23:26:43	☒
/report_exercise_outcome<exerc...	OPTIONS, POST	api.report_exercise_outcome	2017-06-23 23:15:39	☒
/learned_language	OPTIONS, HEAD, GET	api.learned_language	2017-06-23 23:15:30	☒
/bookmarks_to_study<bookmark...	OPTIONS, HEAD, GET	api.bookmarks_to_study	2017-06-23 23:14:09	☒
/interesting_feeds/<language_id>	OPTIONS, HEAD, GET	api.get_interesting_feeds_for_lan...	2017-06-23 23:13:48	☐
/get_starred_articles	OPTIONS, HEAD, GET	api.get_starred_articles	2017-06-23 23:13:48	☐
/get_feeds_being_followed	OPTIONS, HEAD, GET	api.get_feeds_being_followed	2017-06-23 23:13:48	☐
/upload_user_activity_data	OPTIONS, POST	api.upload_user_activity_data	2017-06-23 23:13:43	☐
/get_possible_translations<from...	OPTIONS, POST	api.get_possible_translations	2017-06-23 23:13:37	☒
/native_language	OPTIONS, HEAD, GET	api.native_language	2017-06-23 23:10:52	☐

Fig. 1. All of the endpoints of the Zeeguu app are shown such that a selection can be made for monitoring them

In order to monitor an endpoint, the Flask Dashboard creates a function wrapper for the API function that corresponds to the endpoint. This way, the wrapper will be executed whenever that API call is made. The wrapper contains the code that takes care of monitoring an endpoint. Data collected by the wrappers are persisted in a local database.

There are three main categories of visual perspectives that are available using Flask Dashboard:

- 1) *Service Utilization* presents information about the usage of all the endpoints of interest,
- 2) *Endpoint Performance* presents response times of the various service endpoints,

⁶Section X shows how to install the package

- 3) *User Experience* presents information about the user-perceived performance of the service endpoints.

In the remainder of the paper we present several of these perspectives.⁷

IV. SERVICE UTILIZATION

The most fundamental insight that a service maintainer needs regards service utilization.

Figure 2 shows a first perspective on endpoint utilization that Flask Dashboard provides: a stacked bar chart of the number of hits to various endpoints grouped by day⁸. Figure 2 in particular shows that at its peak the API has about 12,000 hits per day. The way users interact with the platform can also be inferred since the endpoints are indicators of different activity types, e.g.:

- 1) *api.get_possible_translations* is an indicator of the amount of foreign language reading the users are doing, and
- 2) *api.report_exercise_outcome* is an indicator of the amount of foreign vocabulary practice the users are doing.

Fig. 2. The number of requests per endpoint per day view shows the overall utilization of the monitored application

Besides showing the overall utilization, this endpoint provides the maintainer with information relevant for decisions regarding endpoint deprecation — one of the most elementary ways of *understanding the needs of the downstream*[7]. In our case study, the maintainer realized that one endpoint which they thought was not being used (i.e. *words_to_study*), contrary to their expectations, was actually being used⁹.

⁷We recommend obtaining a color version of this paper for better readability

⁸Endpoint colors are the same in different views

⁹A complementary type of usage information can also be discovered in the view presented in Figure 1 where seeing that an endpoint is never accessed can increase the confidence of the maintainer that a given endpoint is not used, although it can never be used a proof.

A second type of utilization question that the Flask Dashboard can answer automatically regards *cyclic patterns of usage per hour of day* by means of a heatmap, as in Fig. 3.

Fig. 3. Usage patterns become easy to spot in the requests per hour heatmap

Figure 3 shows the API not being used during the early morning hours, with most of the activity focused around working hours and some light activity during the evening. This is consistent with the fact that the current users are all in the central European timezone. Also, the figure shows that the spike in utilization that was visible also in the previous graph happened in on afternoon/evening.

V. ENDPOINT PERFORMANCE

The Flask Dashboard also collects information regarding endpoint performance. The view in Fig. 4 summarizes the response times for various endpoints by using a box-and-whiskers plot.

Fig. 4. The response time (in ms) per monitored endpoint view allows for identifying performance variability and balancing issues

From this view it became clear to the maintainer that four of the endpoints had very large variation in performance. The most critical for the application and consequently the one optimized first was the `api.get_possible_translations`

endpoint which was part of an interactive loop in the reader applications that relied on the Zeeguu API. Moreover, cf. Fig. 2 this endpoint is one of the most used in the system.

However, with the current configuration of the tool, it would be impossible for the maintainer to see the improvements resulting from the optimization. One way to do this is to add an extra line of configuration to allow Flask Dashboard to find the `git`¹⁰ folder of the deployed service:

```
dashboard.config.git = 'path/to/.git'
```

With this extra configuration, the Flask Dashboard can now automatically detect the current version of the project, and group measurements by version¹¹. Fig. 5 is a zoomed-in version of such a view for `api.get_possible_translations` with versions increasing from top to bottom

Fig. 5. Visualizing The Performance Evolution of the `api.get_possible_translations` endpoint

This view confirms that the performance of the translation endpoint improved in the recent versions: the median of the last three versions is constantly moving towards the left, and progresses from 1.4 seconds (in the top-most box plot in Fig. 5) to 0.8 in the latest version (bottom-most box plot).

The Flask Dashboard collects **extra information about such outliers**: Python stack trace, CPU load, request parameters, etc. in order to allow the maintainer to investigate the causes of these exceptionally slow response times.

In order to address this, but without degrading overall performance, the Flask Dashboard tracks for every endpoint a running average value. When it detects that a given request is an outlier with respect to this past average running value, it triggers the *outlier data collection routine* which stores all the previously listed extra information about the current execution environment.

¹⁰<https://git-scm.com/>

¹¹Alternatively, the maintainer can add version identifiers manually for the web application through a configuration file if the system does not use git.

VI. USER EXPERIENCE

For service endpoints which run computations in real time, the maintainer of a system might want to understand the endpoint performance on a per-user basis, especially for situations where the system response time is a function of some individual user load¹².

To enable this, the Flask Dashboard must be configured to associate an API call with a given user. The simplest way is to take advantage of the architecture of Flask applications in which a global `flask.request` object can be used to retrieve the session which can in turn lead to user identification:

```
# app specific way of extracting the user
# from a flask request object
def get_user_id():
    sid = int(flask.request.args['session'])
    session = User.find_for_session(sid)
    return user_id

# attaching the get_user_id function
dashboard.config.get_group_by = get_user_id
```

In Zeeguu, `api.get_feed_items_with_metrics` retrieves a list of recommended articles for a given user. Cf. Fig. 4 it is the endpoint with the slowest response time and highest variability. The reason for this is that a user can be subscribed to anything from one to three dozen article sources and for each of the sources the system must compute the personalized difficulty of each article at every request.

Fig. 6. The `api.get_feed_items_with_metrics` shows a very high variability across users

A Per-User Performance perspective should show the different response times for different users. Figure 6 presents a subset of the corresponding view in the Flask Dashboard. The figure shows that the response times for this endpoint can vary considerably for different users with some extreme cases where a user has to wait a full minute until their recommended articles are shown¹³.

¹²E.g. in Gmail some users have two emails while other have twenty thousand and this induces different response times for different users

¹³After seeing this perspective, the maintainer refactored the architecture of the system to move part the difficulty computation out of the interactive loop

The limitation of the previous view is that it does not present the information also on a per version basis. To address this, a different visual perspective entitled `Multi-Version per-User Performance` can be defined. Figure 7 presents such a perspective by mapping the average execution time for a given user (lines) and given version (columns) on the area of the corresponding circle. The colors represent users. The figure shows average performance varying across users and versions with no clear trend: this is probably because varying user workload (i.e. number of sources to which the user is registered) is the reason for the variation in response times.

Fig. 7. This perspective shows that the evolution of response times for individual users (horizontal lines) across versions (the x-axis)

VII. DISCUSSION

A. Automatically Monitoring System Evolution

The main goal of the Flask Dashboard design was to allow analytics to be collected and insight to be gained by making the smallest possible changes to a running API. This technique assumes that the web application code which is the target of the monitoring is deployed using `git` in the following way:

- 1) The deployment engineer pulls the latest version of the code from the integration server; this will result in a new commit being pointed at by the HEAD pointer.
- 2) The deployment engineer restarts the new version of the service. At this point, the Flask Dashboard detects that a new HEAD is present in the local code base and consequently starts associating all the new data points with this new commit¹⁴.

The advantage of this approach is the need for minimal configuration effort, as discussed in the presentation of the tool. The disadvantage is that it will consider on equal ground the smallest of commits, even one that modifies a comment, and the shortest lived of commits, e.g. a commit which was active only for a half an hour before a new version with a bug fix was deployed, with major and minor releases of

¹⁴The Flask Dashboard detects the current version of the analyzed system the first time it is attached to the application object, and thus, assumes that the Flask application is restarted when a new version is deployed. This is in tune with the current version of Flask, but if the web server will support dynamic updates in the future, this might have to be taken into account

the software. A mechanism to control which versions are important for monitoring purposes is therefore required to be added. A further possible extension point here is supporting other version control systems (e.g. Mercurial). However, this is a straightforward extension.

B. User-Awareness

For the situations in which the user information is not available, the Flask Dashboard tracks by default information about different IPs and in some cases this might be a sufficiently good approximation of the user diversity and identity.

The visualizations for the user experience perspectives as presented in Section VI have been tested with several hundred users (of which about two hundred were active during the course of the study), but the scalability of the visualizations must be further investigated for web services with tens of thousands of users.

C. Other Possible Groupings

There are other groupings of service utilization and performance that could be important to the maintainer, that we did not explore in this paper. For example, if the service is using OAuth, then together with every request, in the header of the request there is information about the application which is sending a request. Grouping the information by application that sends the request could be important in such a context.

In general, providing a mechanism that would allow very easy specification of groupings (either as code annotations, as normal code, or as configuration options) is an open problem that Flask Dashboard and any other similar library will have to face.

VIII. RELATED WORK

There is a long tradition of using visualization for gaining insight into software performance. Tools like Jinsight [8] and Web Services Navigator [9] pioneered such an approach for Java and for Web Services that communicate with SOAP messages. Both have an “omniscient” view of the services / objects and their interactions. As opposed to them, in our work we present an analytics platform which focuses on monitoring a single Python web service from its own point of view.

From the perspective of service monitoring, our work falls within the server-side run-time monitoring of services [10]. While we don’t implement the more advanced features of related monitoring solutions like QoS policies driving the monitoring, it presents nevertheless an easy to use approach support improving the performance of web applications.

IX. CONCLUSION AND FUTURE WORK

In this paper we have shown that it is possible to create a monitoring solution which provides basic insight into web service utilization and performance with very little effort from the developer. The user group that we are aiming for with this work is application developers using Flask and Python to build web applications with limited or no budget for implementing their own monitoring solutions. The emphasis is in allowing

such users to gain insight into how the performance of the service evolves together with the application itself. We believe that the same architecture, and lessons can be applied to other frameworks and other languages.

In the future, we plan to perform case studies with other systems, with the goal of discover other needs and to wean out the less useful visualizations in the Flask Dashboard. We plan to also extend the tool towards supporting multiple deployments of the same applications across multiple nodes (e.g. for the situations where the application is deployed together with a load balancer). Finally, we plan to integrate Flask Dashboard with unit testing as a complementary source of information about performance evolution.

X. TOOL AND SOURCE CODE AVAILABILITY

The images in this paper are screenshots of the interactive visualizations from the deployment of Flask Dashboard in the context of the Zeeguu API. The deployment of the system can be consulted online by readers and reviewers of this paper¹⁵.

Flask Dashboard is implemented for Python 3.6 and is available on the Python Package Index repository¹⁶ from where it can be installed on any system that has Python installed by running `pip install flask_monitoring_dashboard` from the command line.

The source code of the Flask Dashboard is published under a permissive MIT license and is available on GitHub¹⁷.

REFERENCES

- [1] M. Cavage, “There is no getting around it: you are building a distributed system,” *Communications of the ACM*, vol. 56, no. 6, pp. 63–70, 2013.
- [2] A. Ronacher, “Quickstart,” <http://flask.pocoo.org/docs/0.12/quickstart/#quickstart>, 2010, [Online; accessed 25-June-2017].
- [3] M. Papazoglou, V. Andrikopoulos, and S. Benbernou, “Managing evolving services,” *Software*, vol. 28, no. 3, pp. 49–55, 2011.
- [4] O. Nierstrasz and M. Lungu, “Agile software assessment,” in *Proceedings of International Conference on Program Comprehension (ICPC 2012)*, 2012, pp. 3–10. [Online]. Available: <http://scg.unibe.ch/archive/papers/Nier12bASA.pdf>
- [5] M. F. Lungu, “Bootstrapping an ubiquitous monitoring ecosystem for accelerating vocabulary acquisition,” in *Proceedings of the 10th European Conference on Software Architecture Workshops*, ser. ECSAW 2016. New York, NY, USA: ACM, 2016, pp. 28:1–28:4. [Online]. Available: <http://doi.acm.org/10.1145/2993412.3003389>
- [6] B. Shneiderman, “The eyes have it: A task by data type taxonomy for information visualizations,” in *IEEE Visual Languages*, College Park, Maryland 20742, U.S.A., 1996, pp. 336–343.
- [7] N. Haenni, M. Lungu, N. Schwarz, and O. Nierstrasz, “A quantitative analysis of developer information needs in software ecosystems,” in *Proceedings of the 2nd Workshop on Ecosystem Architectures (WEA’14)*, 2014, pp. 1–6. [Online]. Available: <http://scg.unibe.ch/archive/papers/Haen14a-QuantitativeEcosystemInformationNeeds.pdf>
- [8] W. De Pauw, E. Jensen, N. Mitchell, G. Sevitsky, J. M. Vlissides, and J. Yang, “Visualizing the execution of java programs,” in *Revised Lectures on Software Visualization, International Seminar*. London, UK: Springer-Verlag, 2002, pp. 151–162.
- [9] W. D. Pauw, M. Lei, E. Pring, L. Villard, M. Arnold, and J. F. Morar, “Web services navigator: Visualizing the execution of web services,” *IBM Systems Journal*, vol. 44, no. 4, pp. 821–845, 2005.
- [10] C. Ghezzi and S. Guinea, “Run-time monitoring in service-oriented architectures,” in *Test and analysis of web services*. Springer, 2007, pp. 237–264.

¹⁵<https://zeeguu.unibe.ch/api/dashboard>; username: *guest*, password: *vissoft*

¹⁶<https://pypi.python.org/pypi/flask-monitoring-dashboard/1.8>

¹⁷<https://github.com/mircealungu/automatic-monitoring-dasboard>