

# Implementing next generation web server in Galaxy

Wai Yi Leung<sup>1</sup>, Leon Mei<sup>1,2</sup>

<sup>1</sup> Leiden University Medical Centre, Sequence Analysis Support Core, Leiden, The Netherlands

<sup>2</sup> Netherlands Bioinformatics Centre, Nijmegen, The Netherlands

[w.y.leung@lumc.nl](mailto:w.y.leung@lumc.nl) / [github.com/wyleung](https://github.com/wyleung)

## Introduction

Our interest is to see whether we can push Galaxy to a new limit on serving more request per second. The reason for this is simple: web request are (relatively) not cpu intensive. Web request are mostly database-connection and/or file-i/o bound because of the web-templates.

We expect performance gains with modern WSGI-servers compared to the default setup with the Paste-server for Pylons.

A secondary goal of this project is to provide a package for sysadmins to easily change the WSGI server in their setup.

## Material and methods


Two virtual machines with 2 virtual cores on an Intel i5-3570K, 8GB ram and 50GB SSD was setup for this benchmark. Both VMs were connected to a GBit-network.

Both machines were installed with Ubuntu 12.04 LTS server edition with the base tools. The webserver had additional packages installed as required by a clean installation of Galaxy (bitbucket revision 8870:09c81e81952d). Server settings were tuned to accept high number of connections and open sockets.

A similar setup was applied on the client-machine where we conducted the benchmark. We choose to use AB v2.3 (Apache Bench) for the benchmarking.


Each WSGI server was installed with the default installation instructions from the package. The configuration was tuned for production usage, e.g. logging and debugging options were turned off.

### WSGI Benchmark with Galaxy



Picture 1: Benchmarking the connectivity of WSGI-servers in a Galaxy setup. Numbers for the Paste are not reliable for  $n > 1$ .

### Memory usage



Picture 2: Measurement of system resource usage. Numbers taken from top/htop while tests were performed.

## Discussion

In the initial setup we started testing with 1 concurrent connection. We didn't see improvements by using Tornado or Gunicorn. Only Fapws3 could significantly outperform Paste by 230%.

To simulate real world applications we open multiple connections to the webserver (e.g. pictures, ajax-snippets etc.). We tried to profile concurrent connections by opening more connections to the server per client.

We observed not only gain in the number of connections, but also a reduction of memory and swap usage compared to the Paste profile. In multi-concurrency mode, we also noticed errors in serving pages using Paste. Paste was not able to open multiple sessions at the same time. About 280/10.000 connections were dropped because of a failure.

As these results were measured in a closed lab-environment, we cannot really tell whether real-world applications would benefit from changing the WSGI server. Ideally this test should be tested in a course where more participants can attend

We created a python-package called **galaxyservers**, which is installable from the PyPi repository and available on GitHub for those managing high traffic Galaxy servers.

Command to install galaxyservers:

```
$> pip install galaxyservers
```

or checkout the source at:

<https://github.com/wyleung/galaxyservers>

```
[server:main]
```

```
use = egg:Paste#http
```

```
# The replacements:
```

```
#use = egg:galaxyservers#tornado
```

```
#use = egg:galaxyservers#gunicorn
```

```
#use = egg:galaxyservers#fapws3
```

Code listing 1: Change the WSGI server by just changing 1 line of code in your configuration. Installable using "pip install galaxyservers"