


Continuous Performance Analysis

for Python with



whoami

Arthur Pastel
Software Engineer



twitter.com/art049



github.com/art049



linkedin.com/in/arthurpastel

Weapons of choice:



Open Source Developer:

Built ODMantic 

Building

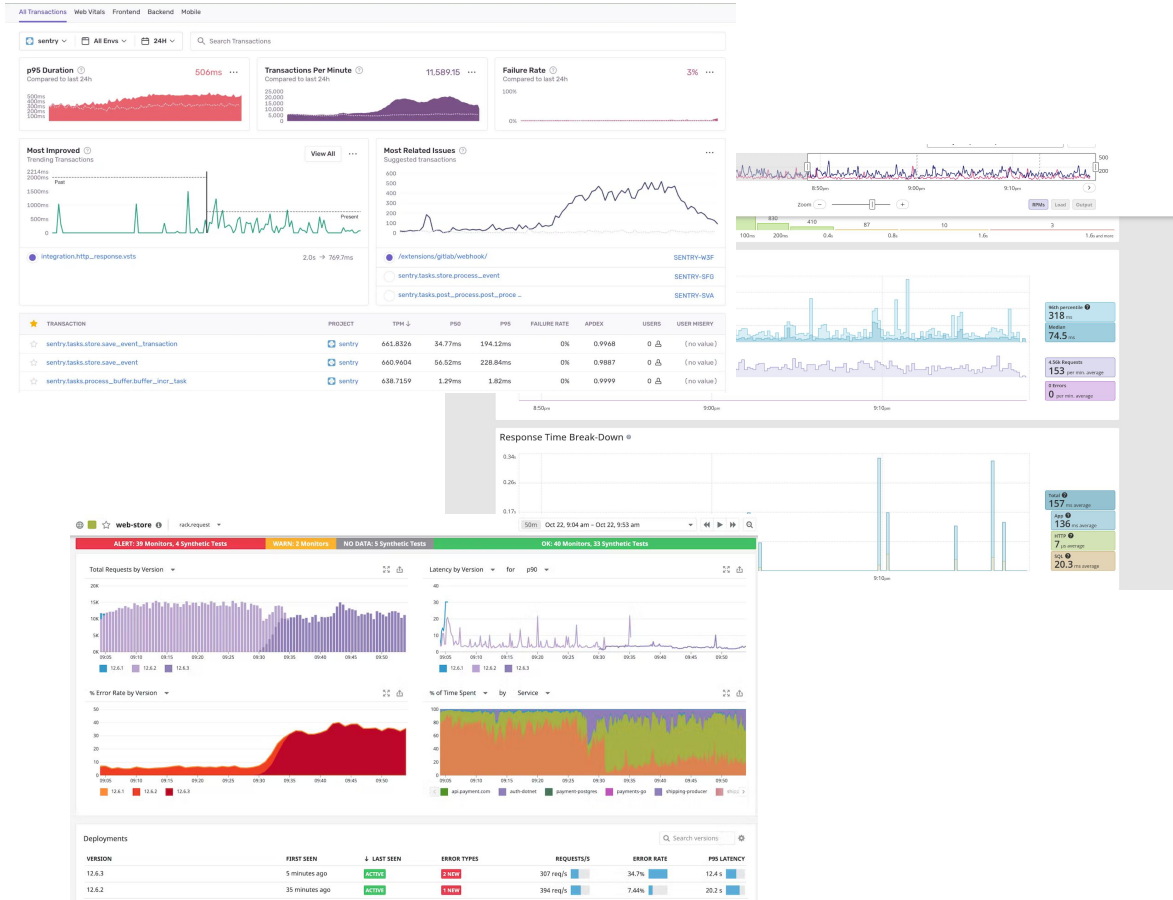
{ **codspeed**  }

Software Performance?

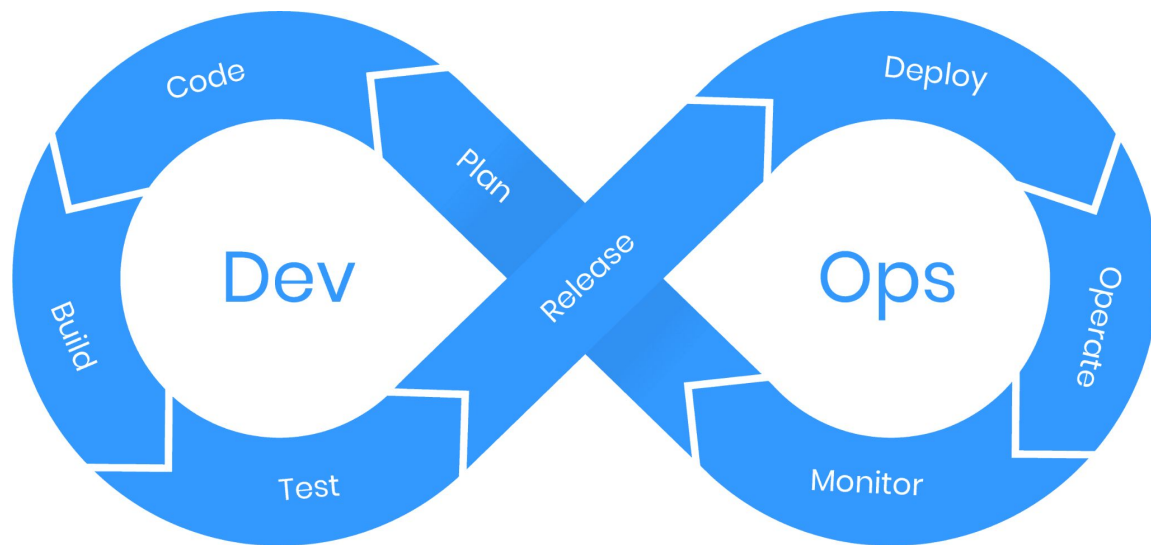
- Execution speed
- Throughput

APMs


- Sentry
- Datadog
- Blackfire
- Cloud Providers



The Software Development Life Cycle(aka SDLC)




Performance feedback in the SDLC



Problems

- Performance issues identified in production
- Experiments in production environments
- Dev != Ops
- Costs

Shifting left



The ideal workflow

- Performance checks as a testing flavour
- Block merging upon regression
- Reliable performance history
- Compare performance from anywhere

Requirements for the performance metric

- Consistent
- Repeatable
- Hardware agnostic

Measuring performance

The toy algorithm



```
def fibonacci(n: int) → int:  
    if n ≤ 1:  
        return n  
    return fibonacci(n - 1) + fibonacci(n - 2)
```

Basic approach



```
import time

start = time.time()
fibonacci(10)
end = time.time()

elapsed_time = round((end - start) * 10**6, 3)
print(f"Elapsed time: {elapsed_time} μs")
```



```
$ python main.py
Elapsed time: 15.84 μs

$ python main.py
Elapsed time: 23.869 μs

$ python main.py
Elapsed time: 18.054 μs

$ python main.py
Elapsed time: 17.041 μs
```

`time.time`: system clock

`time.perf_counter`: high resolution timer

Stats to the rescue

```
import time

samples = []

for i in range(100):
    start = time.perf_counter()
    fibonacci(10)
    end = time.perf_counter()
    samples.append(end - start)

mean = sum(samples) / len(samples)
mean_micros = round(mean * 10**6, 3)
print(f"Mean: {mean_micros} μs")
```

```
$ python main.py
Mean: 12.46 μs

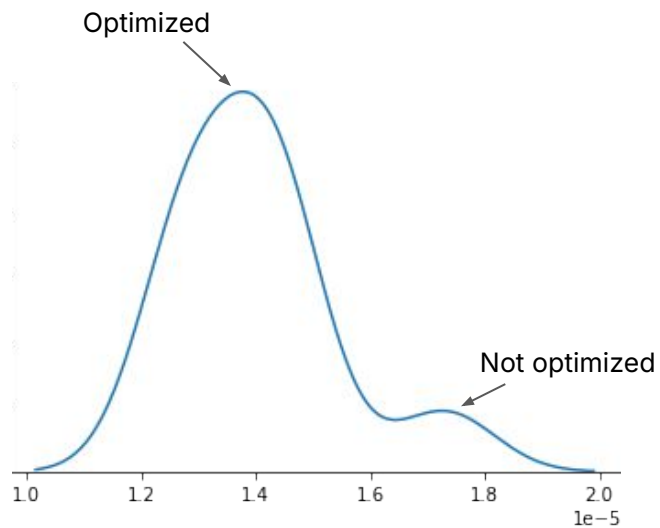

$ python main.py
Mean: 12.503 μs

$ python main.py
Mean: 13.343 μs

$ python main.py
Mean: 12.218 μs
```

~30% faster

Under the hood



Sample distribution

Solutions:

- Warmup rounds
- Disable garbage collection

Stats to the rescue(with warmup)

```
import time, gc

gc.disable() # Disable garbage collection

# Warm up
for i in range(1000):
    fibonacci(10)

# Measure
samples = []
for i in range(100):
    start = time.perf_counter()
    fibonacci(10)
    end = time.perf_counter()
    samples.append(end - start)

mean = sum(samples) / len(samples)
mean_micros = round(mean * 10**6, 3)
print(f"Mean: {mean_micros} μs")
```

```
$ python main.py
Mean: 12.46 μs

$ python main.py
Mean: 12.408 μs

$ python main.py
Mean: 12.243 μs

$ python main.py
Mean: 12.218 μs
```

Improvements: more samples, conditional warmup, outlier removal

Further improvements

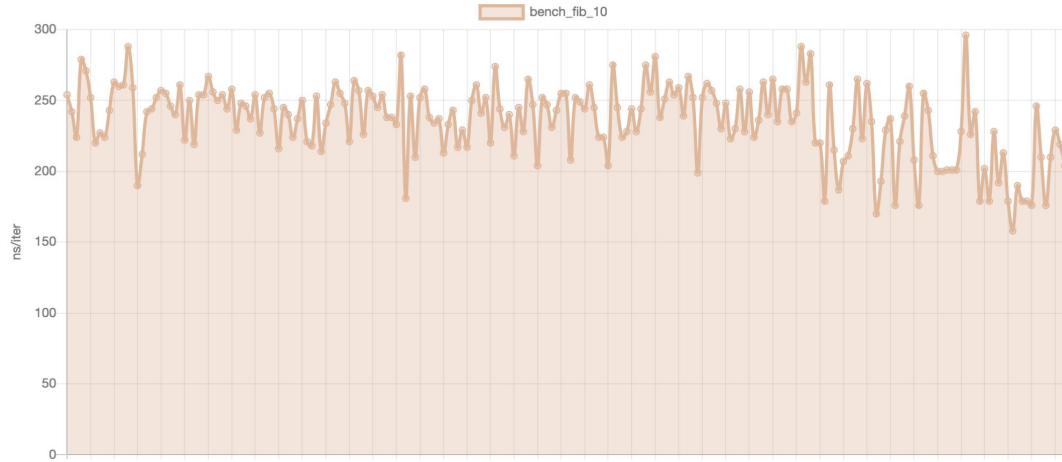
- Conditional warmup
- Remove outliers
- More samples
- Disable the Garbage Collector

Frameworks

- **pytest-benchmark**
- airspeed velocity
- pyperformance


```
def test_fibo_5(benchmark):  
    @benchmark  
    def _():  
        iterative_fibonacci(5)  
  
def test_fibo_10(benchmark):  
    @benchmark  
    def _():  
        iterative_fibonacci(10)  
  
def test_fibo_15(benchmark):  
    @benchmark  
    def _():  
        iterative_fibonacci(15)
```

In a CI environment

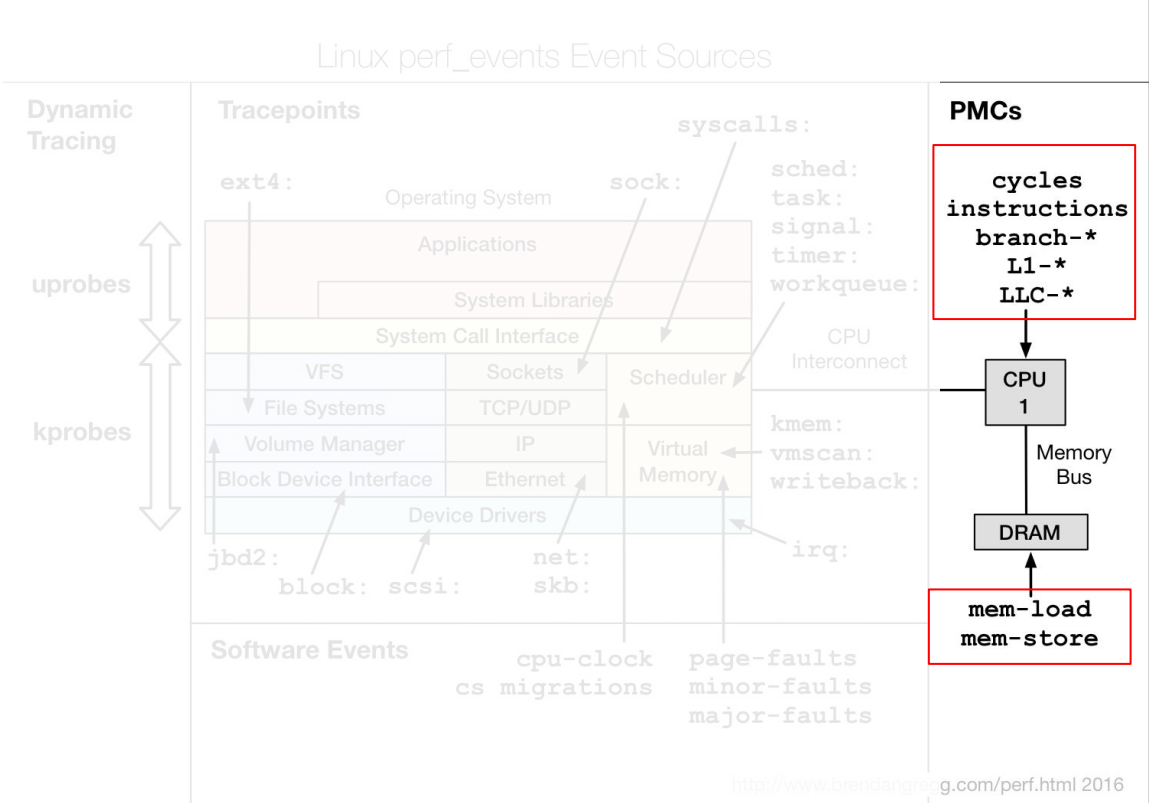


Time measurement for a Fibonacci sequence computation(runs from GitHub Action)

What else?




Performance monitoring counters

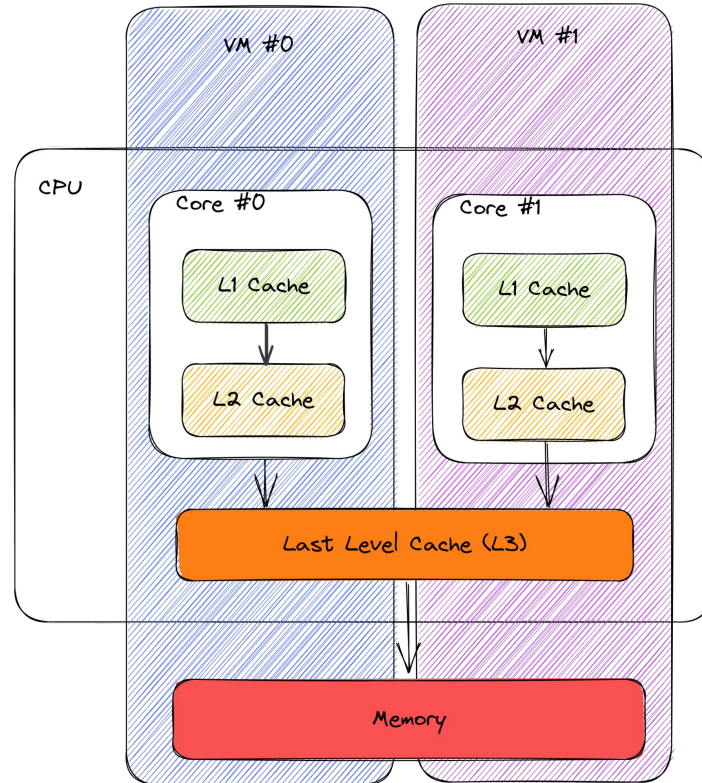


CPU Cache Architecture

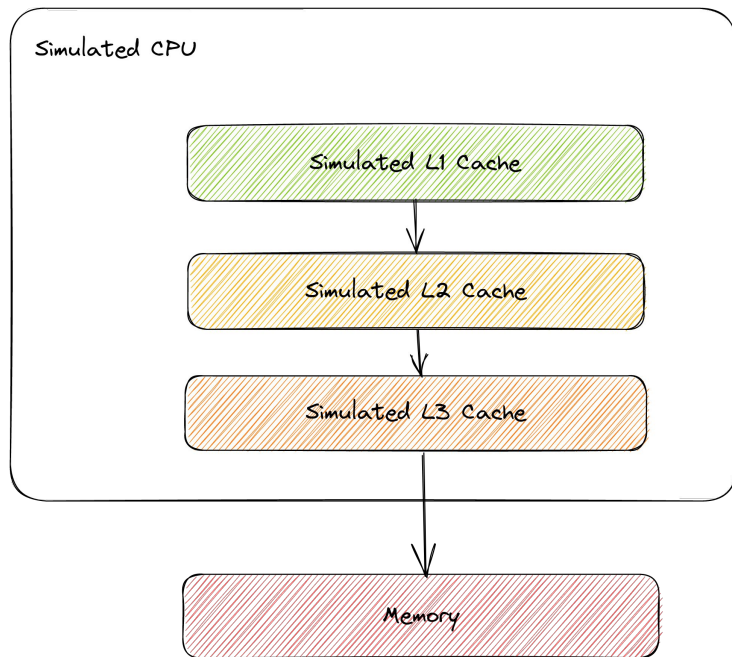
Size	Latency
64KB	1ns
256KB	2.5ns
16MB	11ns
	45ns



Virtualization issue



Simulated caches



Results with CodSpeed



CodSpeed Measurement for a Fibonacci sequence computation(runs from GitHub Action)

Pydantic

Data validation using Python type hints

 FastAPI

 RAY



Strawberry GraphQL

ORMs, ODMs, ...

V2 rewritten in Rust 


```
from datetime import datetime
from pydantic import BaseModel

class User(BaseModel):
    id: int
    name = 'John Doe'
    signup_ts: datetime | None = None
    friends: list[int] = []

external_data = {
    'id': '123',
    'signup_ts': '2019-06-01 12:22',
    'friends': [1, 2, '3'],
}
user = User(**external_data)
print(user.id)
#> 123
print(repr(user.signup_ts))
#> datetime.datetime(2019, 6, 1, 12, 22)
print(user.friends)
#> [1, 2, 3]
print(user.dict())
"""
{
  'id': 123,
  'signup_ts': datetime.datetime(2019, 6, 1, 12, 22),
  'friends': [1, 2, 3],
  'name': 'John Doe',
}
"""
```

Demo

Recap



Performance reporting during development

Improve iterative fibonacci speed for big values (#4)
Merging `improve-iterative-fibo-speed` into `master`

Commits

- No performance changes normally
[BASE] →5c1c151 a month ago by art049
- 54%
Improve iterative fibonacci speed for...
→7821d64 a month ago by art049
- 0%
No performance changes on PR...
→95d2c85 a month ago by art049
- 0%
No performance changes on PR...
→8607178 a month ago by art049
- +65%
[Additional commit]

Benchmarks

- `iterative_fibo_small`: -1%
- `iterative_fibo_big`: +11%
- `iterative_fibo_huge`: +11.2%

CodSpeed Performance Report

Merging #4 `improve-iterative-fibo-speed` (487a918) will improve performances by 11.2%

Summary

- 🔥 1 improvements
- ❌ 0 regressions
- ✅ 7 untouched benchmarks
- 📄 0 new benchmarks
- ❓ 0 dropped benchmarks

Visibility for future deliveries



Supported languages



Thank you!