
Python support for the Linux perf profiler

Release 3.12.2

Guido van Rossum and the Python development team

February 06, 2024

Python Software Foundation
Email: docs@python.org

Contents

1	How to enable <code>perf</code> profiling support	4
2	How to obtain the best results	5
	Index	6

author

Pablo Galindo

The `Linux perf profiler` is a very powerful tool that allows you to profile and obtain information about the performance of your application. `perf` also has a very vibrant ecosystem of tools that aid with the analysis of the data that it produces.

The main problem with using the `perf` profiler with Python applications is that `perf` only gets information about native symbols, that is, the names of functions and procedures written in C. This means that the names and file names of Python functions in your code will not appear in the output of `perf`.

Since Python 3.12, the interpreter can run in a special mode that allows Python functions to appear in the output of the `perf` profiler. When this mode is enabled, the interpreter will interpose a small piece of code compiled on the fly before the execution of every Python function and it will teach `perf` the relationship between this piece of code and the associated Python function using `perf map` files.

Note: Support for the `perf` profiler is currently only available for Linux on select architectures. Check the output of the `configure` build step or check the output of `python -m sysconfig | grep HAVE_PERF_TRAMPOLINE` to see if your system is supported.

For example, consider the following script:

```
def foo(n):  
    result = 0  
    for _ in range(n):
```

(continues on next page)

(continued from previous page)

```
        result += 1
    return result

def bar(n):
    foo(n)

def baz(n):
    bar(n)

if __name__ == "__main__":
    baz(1000000)
```

We can run perf to sample CPU stack traces at 9999 hertz:

```
$ perf record -F 9999 -g -o perf.data python my_script.py
```

Then we can use perf report to analyze the data:

```
$ perf report --stdio -n -g
```

#	Children	Self	Samples	Command	Shared Object	Symbol
#
↪
#	91.08%	0.00%	0	python.exe	python.exe	[.] _start
		---				__start
		--90.71%				__libc_start_main
		Py_BytesMain				
						--56.88%--pymain_run_python.constprop.0
						--56.13%--_PyRun_AnyFileObject
						_PyRun_SimpleFileObject
						--55.02%--run_mod
						--54.65%--PyEval_EvalCode
						PyEval
↪EvalFrameDefault						PyObject_
↪Vectorcall						_PyEval_Vector
						PyEval
↪EvalFrameDefault						PyObject_
↪Vectorcall						_PyEval_Vector
						PyEval
↪EvalFrameDefault						PyObject_
↪Vectorcall						_PyEval_Vector
						--51.67%--_
↪PyEval_EvalFrameDefault						

(continues on next page)

2 How to obtain the best results

For best results, Python should be compiled with `CFLAGS="-fno-omit-frame-pointer -mno-omit-leaf-frame-pointer"` as this allows profilers to unwind using only the frame pointer and not on DWARF debug information. This is because as the code that is interposed to allow `perf` support is dynamically generated it doesn't have any DWARF debugging information available.

You can check if your system has been compiled with this flag by running:

```
$ python -m sysconfig | grep 'no-omit-frame-pointer'
```

If you don't see any output it means that your interpreter has not been compiled with frame pointers and therefore it may not be able to show Python functions in the output of `perf`.

Index

E

environment variable
 PYTHONPERFSUPPORT, 4

P

PYTHONPERFSUPPORT, 4