

Why SSD developers need pynvme? and why pynvme needs SPDK?

<https://github.com/cranechu/pynvme>

SSD Testing

SAMSUNG	Kingston	intel	WD	Lenovo	TOSHIBA Leading Innovation >>>	HIKVISION 海康威视	PLEXTOR
hp	SanDisk 闪迪	GLOWAY	COLORFUL 七彩虹	威刚 (ADATA)	Asgard	tigo 金泰克	crucial by Micron
STECLAST	KingBank 金百达	影驰	acer	Netac 朗科	Pioneer	Lexar 雷克沙	<u>KingSpec</u>
建兴 LITEON	EAGET 忆捷	ThinkPad	GIGABYTE®	紫光	美商海盗船	TJ	亿储 (Billion reserv...

SSD Testing

[back](#) Data Center SSDs X

Intel® SSD D3-S4610 Series	Intel® SSD DC S3100 Series	Intel® SSD D5-P4326 Series
Intel® SSD DC S4600 Series	Intel® Optane™ SSD DC P4801X Series	Intel® SSD D5-P4320 Series
Intel® SSD D3-S4510 Series	Intel® Optane™ SSD DC P4800X Series	Intel® SSD DC P4101 Series
Intel® SSD DC S4500 Series	Intel® SSD DC P4618 Series	Intel® SSD DC P3700 Series
Intel® SSD DC S3710 Series	Intel® SSD DC P4610 Series	Intel® SSD DC P3608 Series
Intel® SSD DC S3700 Series	Intel® SSD DC P4608 Series	Intel® SSD DC P3600 Series
Intel® SSD DC S3610 Series	Intel® SSD DC P4600 Series	Intel® SSD DC P3520 Series
Intel® SSD DC S3520 Series	Intel® SSD DC P4511 Series	Intel® SSD DC P3500 Series
Intel® SSD DC S3510 Series	Intel® SSD DC P4510 Series	Intel® SSD DC P3100 Series
Intel® SSD DC S3500 Series	Intel® SSD DC P4501 Series	Intel® Optane™ SSD DC D4800X Series
Intel® SSD DC S3320 Series	Intel® SSD DC P4500 Series	Intel® SSD DC D3700 Series
Intel® SSD DC S3110 Series	Intel® SSD D5-P4420 Series	Intel® SSD DC D3600 Series

SSD Testing



INNOGRIT



F A D U

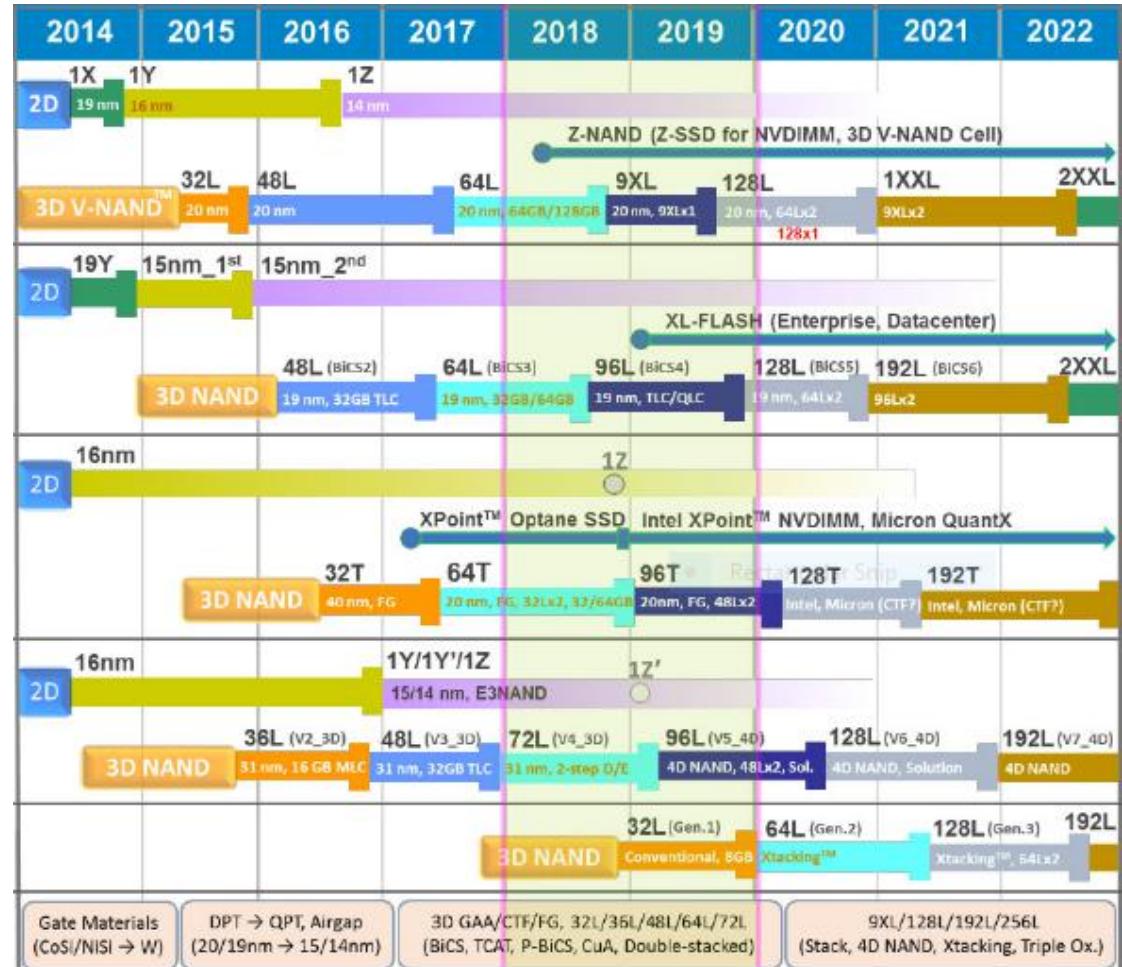


CNEXLABS

SSD Testing

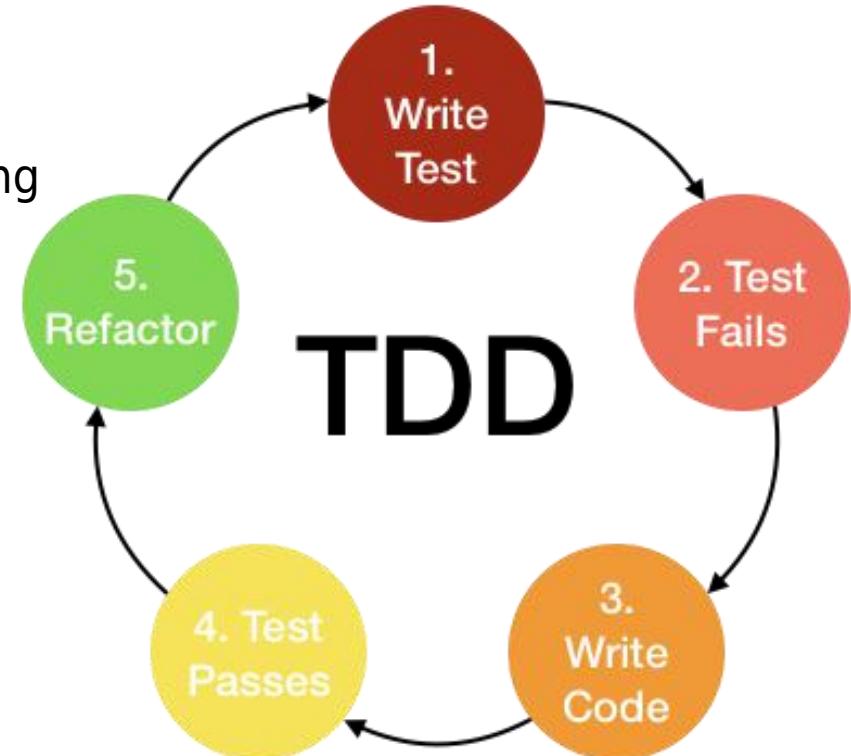
NVMe spec:

- 1.0e (Jan. 2013)
- 1.1b (July 2014)
- 1.2 (Nov. 2014)
 - 1.2a (Oct. 2015)
 - 1.2b (June 2016)
 - 1.2.1 (June 2016)
- 1.3 (May 2017)
 - 1.3a (Oct. 2017)
 - 1.3b (May 2018)
 - 1.3c (May 2018)
 - 1.3d (March 2019)
- 1.4 (June 2019)



Agile and TDD

- Challenge
 - NAND is changing
 - Applications and specifications are changing
 - Diversity on NAND and controllers
- SSD development should be ...
 - fast iteration
 - customer-oriented
 - open to change
- Waterfall v.s. Agile
 - Test-driven development (TDD)



Test Driver

- Embedded devices provide very limited resources
- We need A test-dedicated NVMe driver in host platforms:
 - exports device's features to host
 - exports device's flaws to host
 - exports device's performance to host
 - friendly to test script development
 - friendly to firmware debug
 - friendly to CI

Existed Tools

	tnvme	DM	fio	*Marks
feature	√	√	✗	✗
performance	✗	√	√	√
scripts	√	✗	√	✗
debug	✗	√	✗	✗
CI	√	✗	√	✗
driver	dnvme	OFA	Linux	Windows

pynvme

The pynvme is a python extension module.

Users can operate NVMe SSD **intuitively** in Python scripts. It is designed for NVMe SSD testing with **performance** considered.

Integrated with third-party tools, vscode and pytest, pynvme provides a **convenient** and professional solution to test NVMe devices.



pynvme

Why SPDK?

✓ user space:

- easy for debugging
- maintainness

✓ well modularized:

- jsonrpc: for the vscode plugin
- memzone: share memory between processers
- crc32

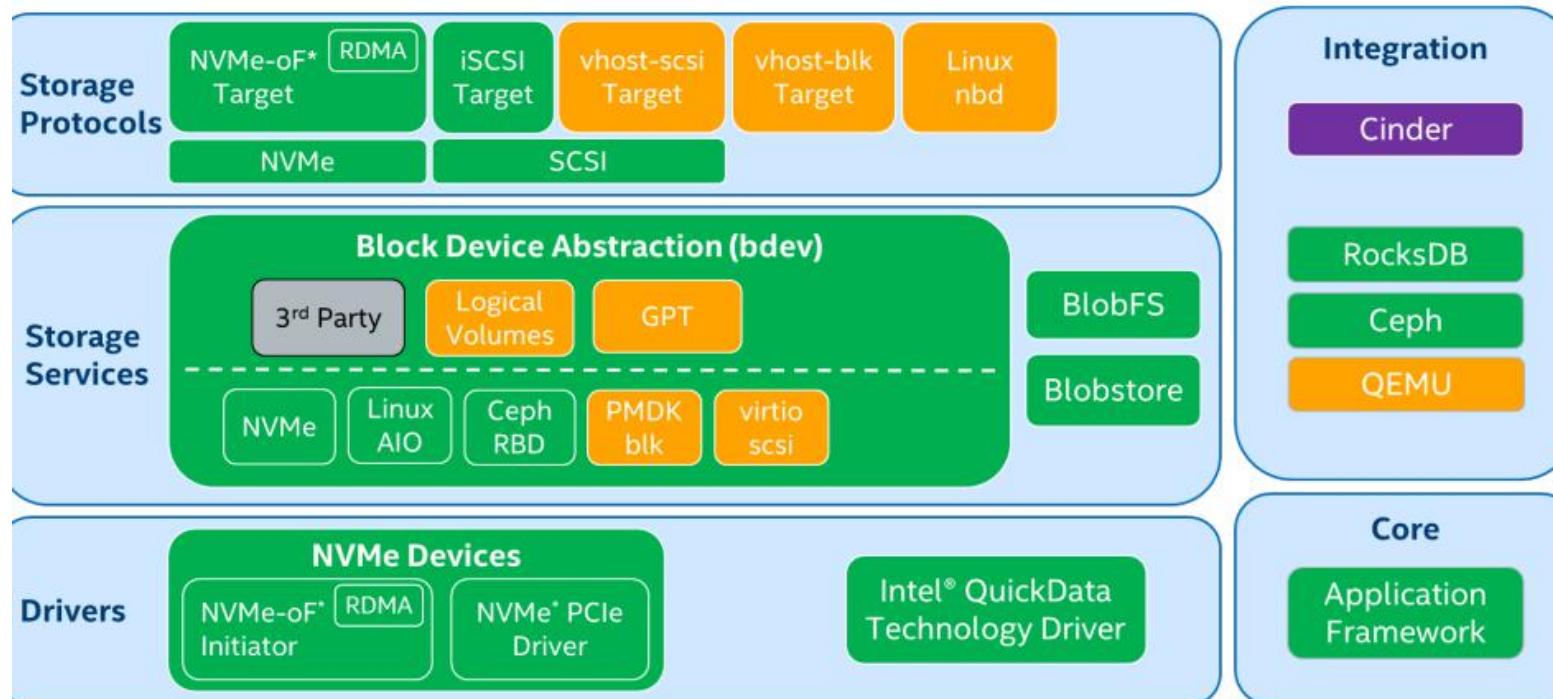
✓ best performance:

- test efficiency
- stress test

✓ open and active

- SSD, NVMe, NAND are all keeping changing!

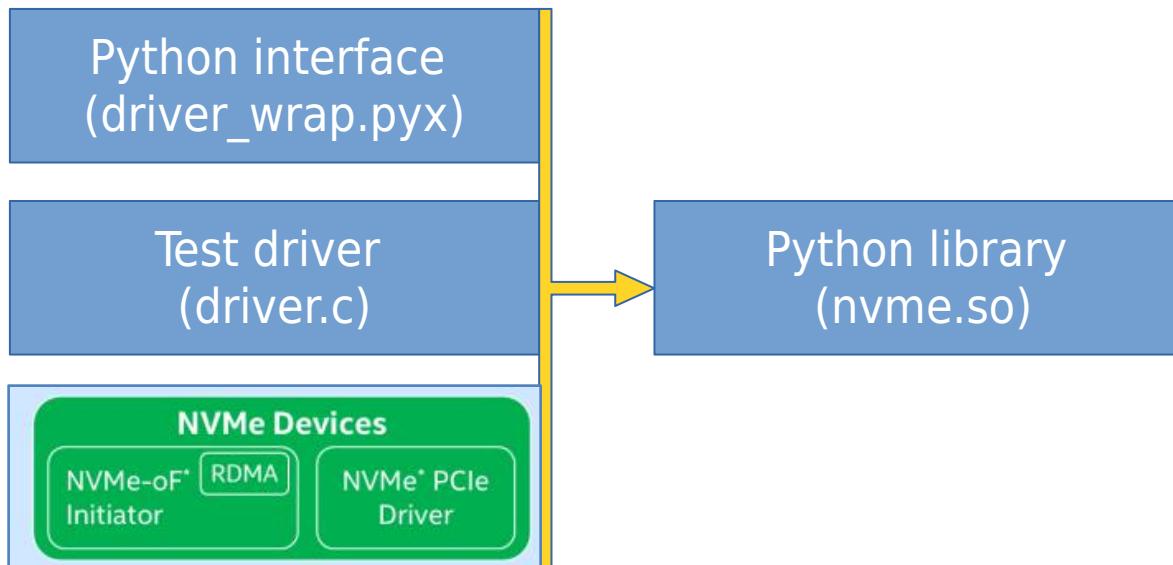
Pynvme is based on SPDK/DPDK



Pynvme Architecture

Build python library with [Cython](#):

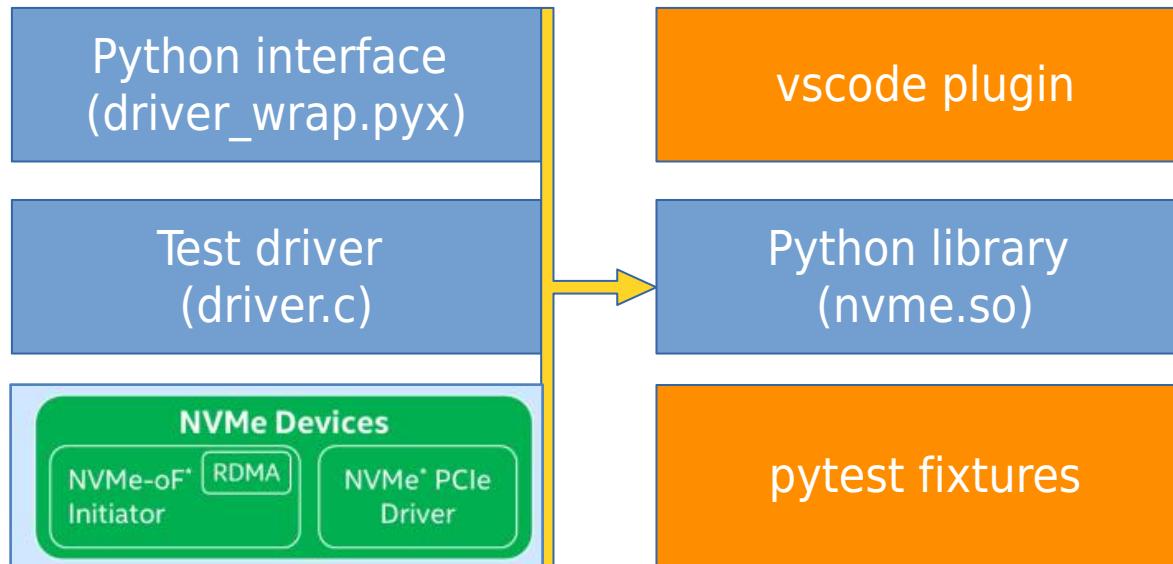
- setup.py
- driver.c
- driver.h
- cdriver.pxd
- driver_wrap.pyx
- Makefile



Pynvme Architecture

Organize test cases in `pytest`:

- `mvme.so`
- `pytest.ini`
- `conftest.py`
- `driver_test.py`



Pynvme Architecture

Python interface (driver_wrap.pyx)

Controller

Namespace

Qpair

Buffer

Test driver (driver.c)

cmdlog

checksum

MSIx

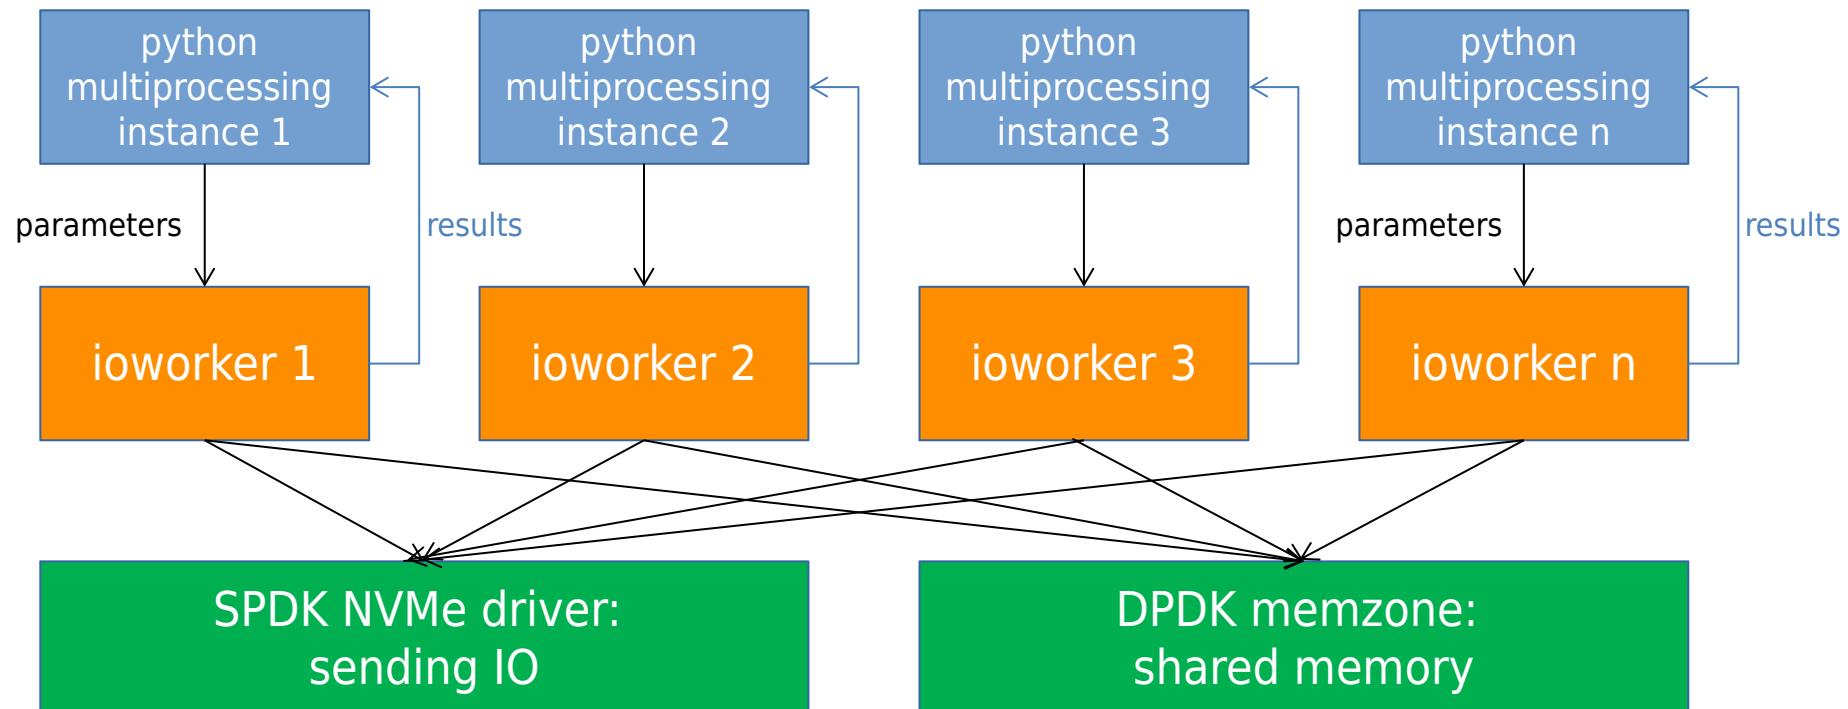
ioworker

SPDK NVMe Driver

PCIe

TCP

IOWorker



Why Python?



- ✓ Many beautiful and mature libraries
 - ✓ Cython
 - ✓ pytest
 - ✓ logging
 - ✓ multiprocessing
 - ✓ pydoc, os, io, time, pytemperature, statistics, yaml, json, struct, matplotlib, ...
- ✓ Friendly to test script development
 - VSCode, Emacs, Pycharm, ...
- ✓ Friendly to CI: develop firmware softly
 - Introducing software methodologies, processes and tools to firmware.

pipeline passed

Pytest Execution



pytest

- "The pytest framework makes it easy to write small tests, yet scales to support complex functional testing for applications and libraries."
- "pytest fixtures offer dramatic improvements over the classic xUnit style of setup/teardown functions"
- use “make test” to start pytest sessions
 - make test
 - make test TESTS=scripts
 - make test TESTS=scripts/demo_test.py
 - make test TESTS=scripts/utility_test.py::test_download_firmware
- find test logs in test.log

Fixtures of pynvme

- create/delete test objects. in conftest.py:
 - nvme0
 - nvme0n1
 - pcie
 - ...
- parametrize of tests
 - `@pytest.mark.parametrize("qcount", [1, 2, 4, 8, 15])`
 - `@pytest.mark.parametrize("repeat", range(10))`
- test control
 - `@pytest.mark.skip("nvme over tcp")`
- doc: <https://docs.pytest.org/en/latest/fixture.html>

Visual Studio Code



- VSCode is [the most popular IDE](#).
 - root user is not recommended by vscode, so users need to run sudo without a password: *sudo visudo*
- Pynvme also providers an extension to monitor device status and cmdlog of every qpair, via jsonrpc. To install the extension:
 - *code --install-extension pynvme-console-1.x.x.vsix*
- Add DUT pci address to .vscode/settings.json
 - get the BDF address with *lspci*
- *make setup; code . # launch the vscode*

DEBUG

Variables

Locals

cdw0: 0

nvme0n1: <nvme.Namespace object at 0x7ffffe02e1bd8>

qpair: <qpair: 2>

read_buf: <buffer name: buffer>

status1: 1

WATCH

write_buf[10:21]: NameError("name 'write_buf' is not defined")

read_buf[10:21]: b'\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00'

CALL STACK

PAUSED ON BREAKPOINT

demo_test.py 11:1

demo_test.py 13:1

testlauncher.py 50:1

testlauncher.py 62:1

demo_test.py 11:1

testlauncher.py 50:1

testlauncher.py 62:1

BREAKPOINTS

Raised Exceptions

Uncaught Exceptions

demo_test.py scripts

demo_test.py

import pytest

import nvme as d

Run Test | Debug Test

def test_hello_world(nvme0, nvme0n1: d.Namespace):

 write_buf = d.Buffer(512)

 write_buf[10:21] = b'hello world'

 qpair = d.Qpair(nvme0, 16)

 read_buf = d.Buffer(512)

 def write_cb(cdw0, status1):

 nvme0n1.read(qpair, read_buf, 0, 1)

 nvme0n1.write(qpair, write_buf, 0, 1, cb=write_cb)

 qpair.waitdone(2)

 assert read_buf[10:21] == b'hello world'

PROBLEMS

OUTPUT

DEBUG CONSOLE

TERMINAL

2: Python Debug Console

sudo /usr/bin/python3 /home/cranechu/.vscode/extensions/ms-python.python-2019.4.1
1987/pythonFiles/ptvsd_launcher.py --default --client --host localhost --port 3878
1 /home/cranechu/.vscode/extensions/ms-python.python-2019.4.1987/pythonFiles/test
launcher.py /home/cranechu/pynvme pytest --junitxml=/tmp/tmp-12479aQxuEn0mpwap.xml
--pciaddr=01:00.0 ./scripts/demo_test.py::test_hello_world
===== test session starts =====
platform linux -- Python 3.7.3, pytest-4.3.1, py-1.8.0, pluggy-0.9.0 -- /usr/bin/p
yton3
cachedir: .pytest_cache
rootdir: /home/cranechu/pynvme, ini file: pytest.ini
plugins: cov-2.6.1
collected 1 item

scripts/demo_test.py::test_hello_world
----- live log setup -----
[2019-04-29 22:20:36.206] INFO pciaddr(19): running tests on DUT 01:00.0

LN 11, COL 1 SPACES: 4 UTF-8 LF Python

This screenshot shows a Python test session in Visual Studio Code (VS Code) for a project named 'pynvme'. The interface includes a sidebar with icons for file operations, search, and configuration, and a bottom status bar showing the current file, Python version, and other details.

The main area displays the code for `test_trim_basic.py`. A yellow callout labeled "test items" points to the code, which contains several logging statements and assertions. A yellow callout labeled "qpairs" points to the "PYNVMQ QPAIRS" section in the sidebar, which lists three entries: QPair 00, QPair 01, and QPair 02.

The right side of the screen shows the "OUTPUT" tab displaying the test logs. The logs indicate a test session starting on a Linux platform with Python 3.7.3, pytest 4.4.0, and pluggy 0.9.0. It shows the execution of the test, including live log setup, log calls, and teardown, along with XML generation and duration information. A yellow callout labeled "test log" points to the top portion of the logs, while another yellow callout labeled "test scripts" points to the bottom portion.

```
TEST
PYTHON
  driver_test.py
scripts
  test_trim_basic.py
    test_trim_basic

PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL
Python Test Log
platform linux -- Python 3.7.3, pytest-4.4.0, py-1.8.0, pluggy-0.9.0 -- /usr/bin/python3
cachedir: .pytest_cache
rootdir: /home/cranechu/pynvme, infile: pytest.ini
plugins: cov-2.6.1
collecting ... collected 1 item

scripts/test_trim_basic.py::test_trim_basic
----- live log setup -----
[2019-04-29 12:37:06.237] INFO pciaddr(19): running tests on DUT 01:00.0
----- live log call -----
[2019-04-29 12:37:10.777] INFO test_trim_basic(14): model number: SM961 NVMe SAMSUNG 1024GB
[2019-04-29 12:37:10.778] INFO test_trim_basic(15): firmware revision: CXA75D0Q
[2019-04-29 12:37:10.779] INFO test_trim_basic(18): write data in 10G ~ 20G
[2019-04-29 12:37:16.761] INFO test_trim_basic(42): trim the 10G data from LBA 0x1400000
[2019-04-29 12:37:16.763] INFO test_trim_basic(47): trim bandwidth: 16219.27GB/s
[PASSED]
----- live log teardown -----
[2019-04-29 12:37:16.766] INFO script(33): test duration: 5.992 sec

----- generated xml file: /tmp/tmp-5551D69216jfy7tm.xml -----
===== 1 passed in 10.65 seconds =====
platform linux -- Python 3.7.3, pytest-4.4.0, py-1.8.0, pluggy-0.9.0 -- /usr/bin/python3
cachedir: .pytest_cache
rootdir: /home/cranechu/pynvme, infile: pytest.ini
plugins: cov-2.6.1
collecting ... collected 1 item

scripts/test_trim_basic.py::test_trim_basic
----- live log setup -----
[2019-04-29 12:37:45.654] INFO pciaddr(19): running tests on DUT 01:00.0
----- live log call -----
```

PYNVMQ QPAIRS

QPair 00:
QPair 01:
QPair 02: [REDACTED]

test items

test scripts

trim_script* Python 3.7.3 64-bit 0 ▲ 0 Running Tests / Ln 50, Col 63 Spaces: 4 UTF-8 LF Python

demo_test.py x

```
1 import time
2 import pytest
3 import logging
4 import nvme as d
5 from pytemperature import k2c
6
7 Run Test | Debug Test
8
9 def test_ioworker_with_temperature(nvme0, nvme0n1):
10     smart_log = d.Buffer(512, "smart log")
11     with nvme0n1.ioworker(io_size=8, lba_align=16, lba_random=True,
12                           qdepth=16, read_percentage=0, time=30):
13         for i in range(40):
14             nvme0.getlogpage(0x02, smart_log, 512).waitdone()
15             ktemp = smart_log.data(2, 1)
16             ctemp = k2c(ktemp)
17             logging.info("temperature: %0.2f degreeC" % ctemp)
18             time.sleep(1)
```

≡ CMDLOG Q0 x

```
1 1556546728.807225: [cmd: Get Log Page]
2 0x005f0002, 0xffffffff, 0x00000000, 0x00000000
3 0x00000000, 0x00000000, 0x7439a000, 0x00000001
4 0x00000000, 0x00000000, 0x007f0002, 0x00000000
5 0x00000000, 0x00000000, 0x00000000, 0x00000000
6 1556546728.809293: [cpl: SUCCESS]
7 0x00000000, 0x00000000, 0x00000023, 0x0001005f
8
9 1556546727.804234: [cmd: Get Log Page]
10 0x005f0002, 0xffffffff, 0x00000000, 0x00000000
11 0x00000000, 0x00000000, 0x7439a000, 0x00000001
12 0x00000000, 0x00000000, 0x007f0002, 0x00000000
13 0x00000000, 0x00000000, 0x00000000, 0x00000000
14 1556546727.806340: [cpl: SUCCESS]
15 0x00000000, 0x00000000, 0x00000022, 0x0001005f
16
17 1556546726.801221: [cmd: Get Log Page]
```

test scripts

PROBLEMS OUTPUT ...

Python Test Log

```
[2019-04-29 22:05:11.721] INFO test_ioworker_with_temperature(15):
temperature: 44.85 degreeC (15):
[2019-04-29 22:05:15.756] INFO temperature: 44.85 degreeC (15):
[2019-04-29 22:05:16.760] INFO temperature: 45.85 degreeC (15):
[2019-04-29 22:05:18.775] INFO temperature: 45.85 degreeC (15):
[2019-04-29 22:05:19.779] INFO temperature: 45.85 degreeC (15):
[2019-04-29 22:05:20.782] INFO temperature: 45.85 degreeC (15):
[2019-04-29 22:05:21.786] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:22.789] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:23.792] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:24.795] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:25.800] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:26.803] INFO temperature: 46.85 degreeC (15):
[2019-04-29 22:05:27.806] INFO temperature: 47.85 degreeC (15):
[2019-04-29 22:05:28.809] INFO temperature: 47.85 degreeC (15):
```

define IO patterns in ioworker's parameter list, and run in a separated process

monitor temperature at the same time

test log

cmdlog



master*

Python 3.7.3 64-bit

x

0

▲ 0

Running Tests \

Ln 1, Col 1 Spaces: 4 Plain Text



pynvme goes to 1.x

cranechu / pynvme

Watch 8 Unstar 32 Fork 15

Code Issues 11 Pull requests 0 Projects 2 Wiki Security Insights Settings

Branch: master pynvme / README.md Find file Copy path

cranechu change homepage video 6648db0 on Jun 3

1 contributor

1228 lines (902 sloc) | 39.1 KB Raw Blame History

↪ pynvme

test NVMe devices in Python. [<https://github.com/cranechu/pynvme>]

pipeline passed license BSD-3-Clause release v1.0



pynvme

Thanks!



pynvme

Q & A

Live Demo