



High Performance Storage : A Cloud Story


Luwei He
Standard Engineer, Huawei



Outline

- ▣ **Background**
- ▣ Cloud Storage Core Technologies
- ▣ Proposed Design
- ▣ Conclusion

Background



More and more applications rely on high performance storage

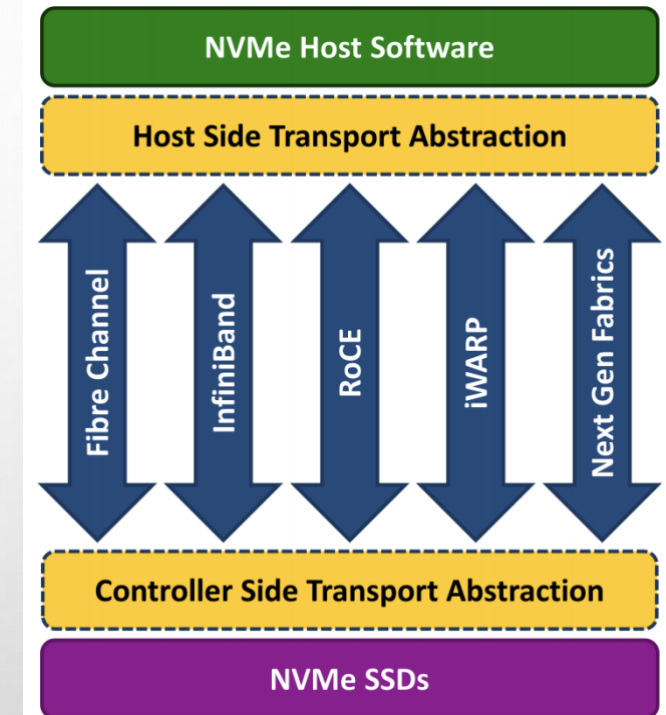
Outline

- Background
- Cloud Storage Core Technologies
 - NVMf
 - SPDK
 - Cyborg
 - OpenSDS
- Proposed Design
- Conclusion

Cloud Storage Core Technologies——NVMe

NVMe over Fabrics

- a technology specification
- transfer data between a host computer and a target SSD over a network
- Fibre Channel, Infiniband, RoCE, ...




Cloud Storage Core Technologies——NVMf

Benefits

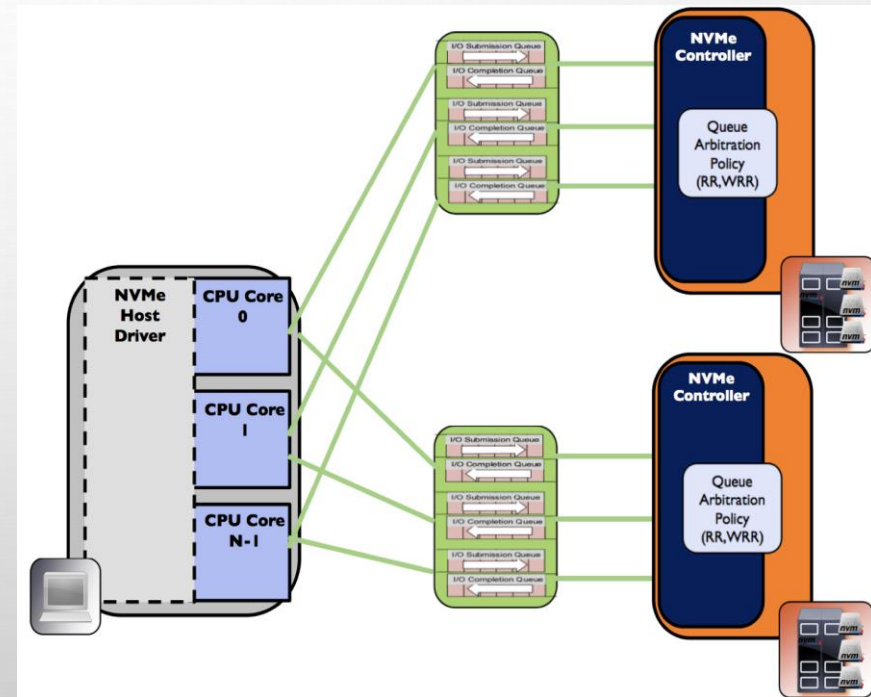
- low latency
- high IOPS



IO path is shorter

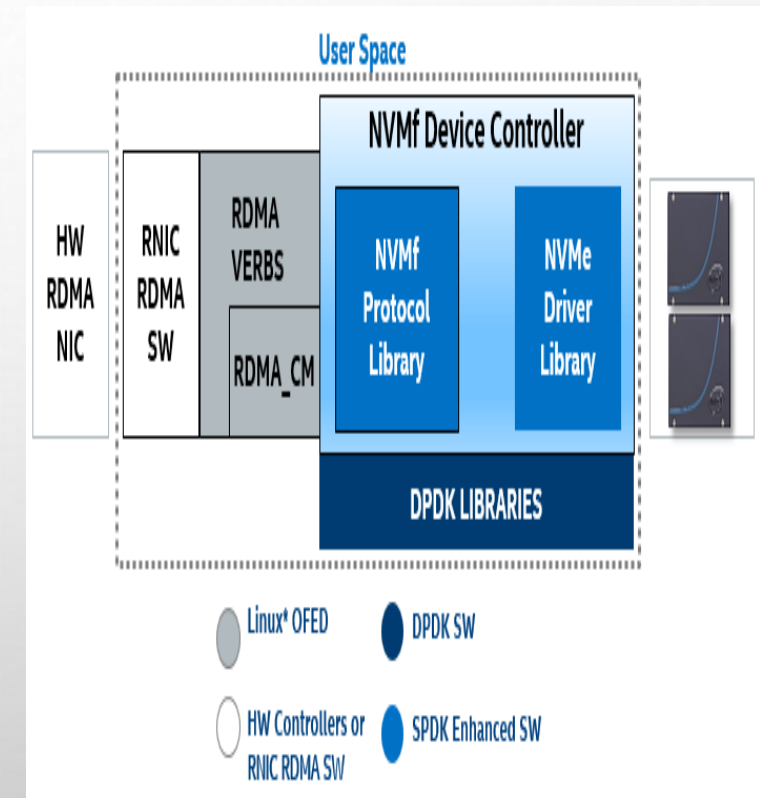
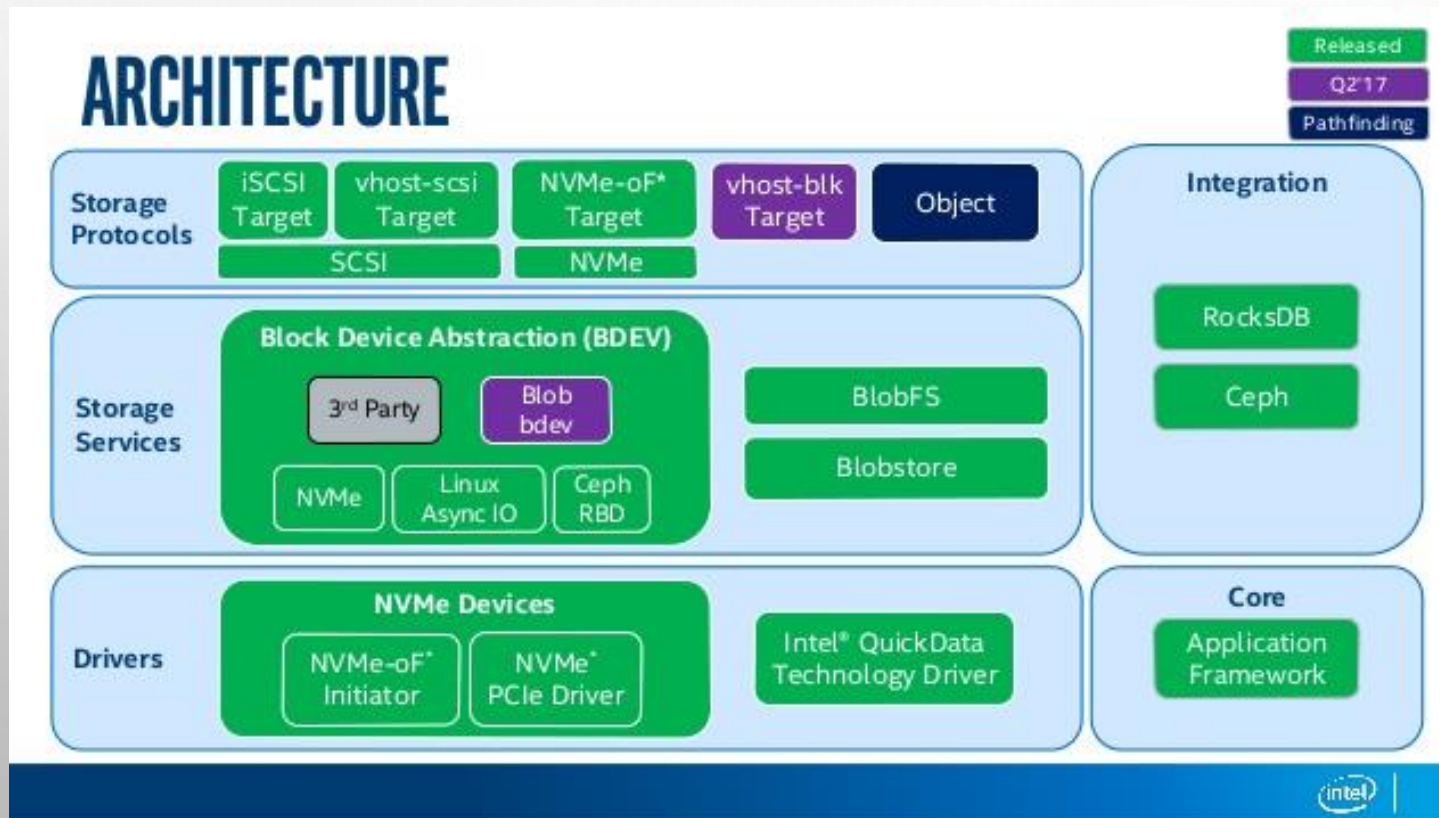


Multi-queue, Lockless



Cloud Storage Core Technologies——SPDK


Storage Performance Development Kit: User Space Driver and Poll-mode



Cloud Storage Core Technologies——SPDK

NVMf

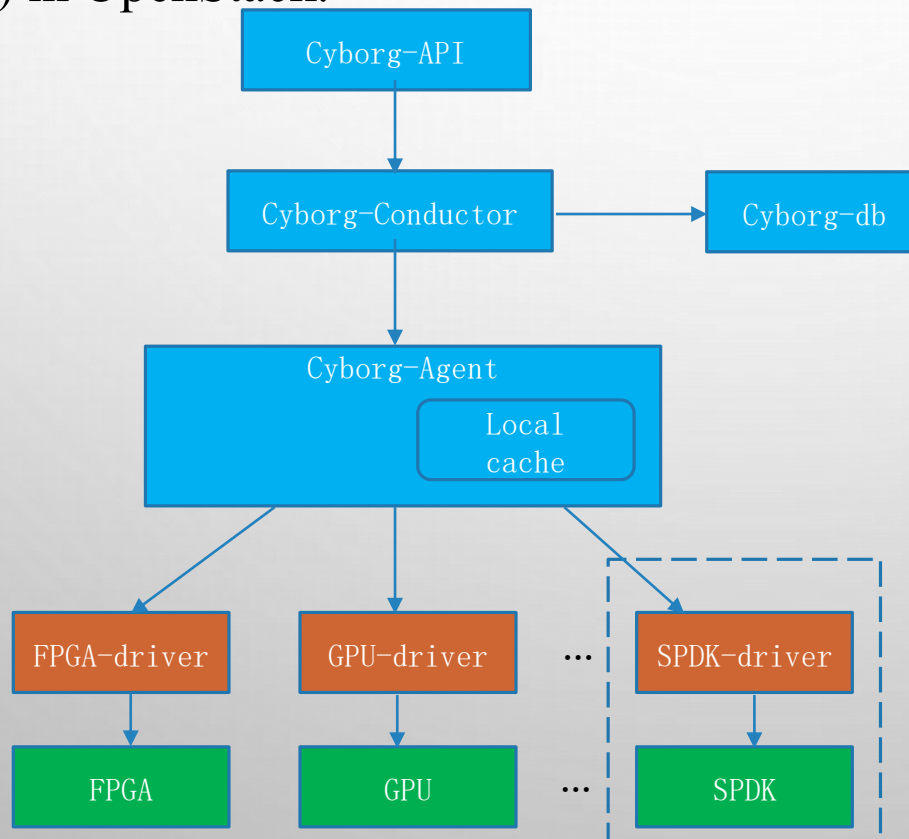
- Moving the NVMe drivers into user space, which avoids syscalls and enables zero-copy access from the application.
- Polling hardware for completions instead of relying on interrupts, which lowers latency.
- Avoiding all locks in the I/O path, instead relying on message passing.



Cloud Storage Core Technologies — OpenStack

Cyborg

Cyborg aims to provide a management framework for dedicated devices (e.g. FPGA, GPU) as well as various accelerators (e.g. DPDK, SPDK) in OpenStack.

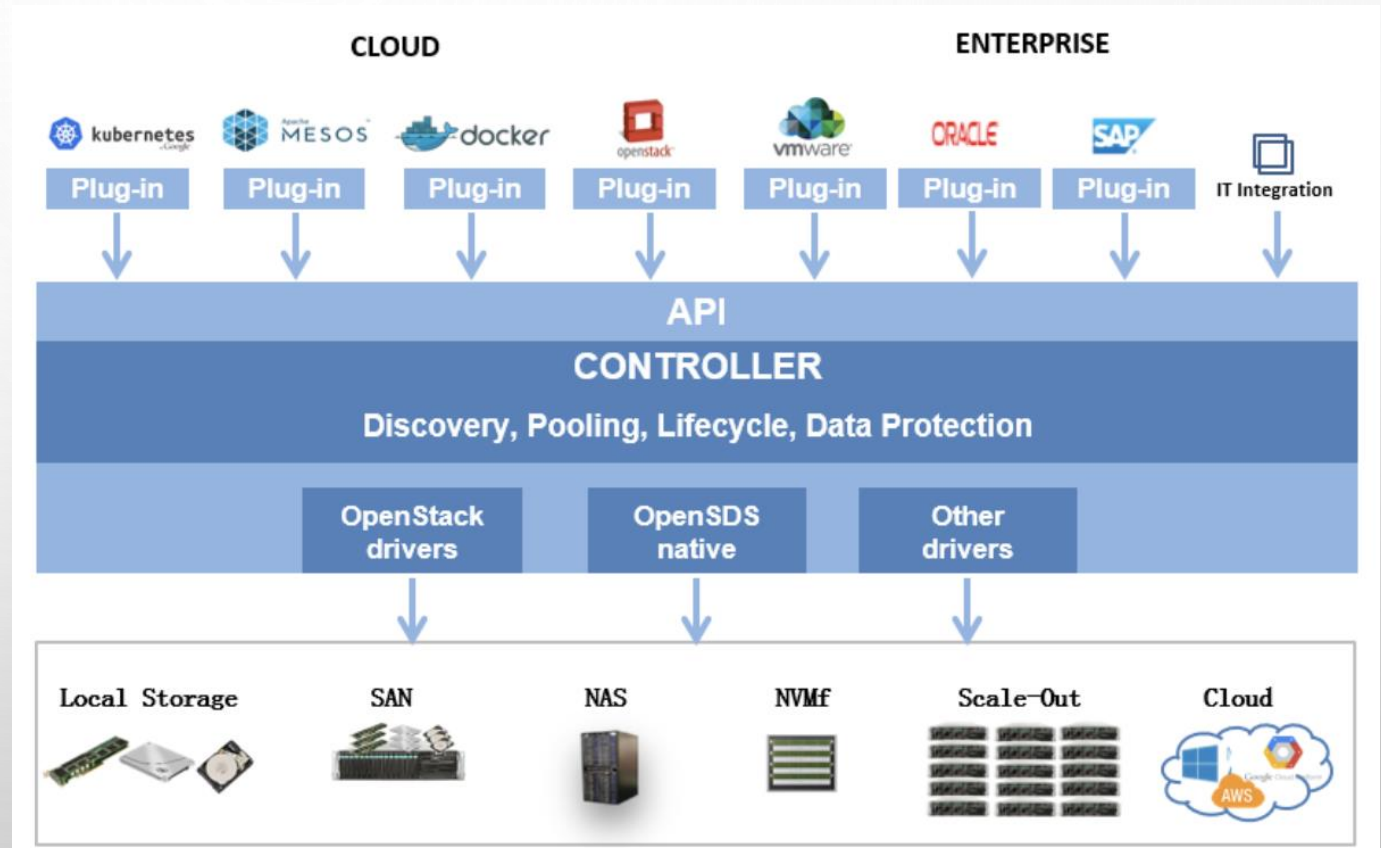


```
'server': 'nvmf_tgt',
'bdevs': [{"num_blocks": 131072,
           "name": "nvme0n1",
           "block_size": 512
          }],
'subsystems': [{"nqn": "nqn.2016-06.io.spdk:cnode1",
                 "subtype": "NVMe",
                 "listen_addresses": [
                 {
                 "trtype": "RDMA",
                 "adrfam": "IPv4",
                 "traddr": "192.168.0.123",
                 "trsvcid": "4420"
                 }
                 ],
                 "hosts": [
                 {"nqn": "nqn.2016-06.io.spdk:host1"}
                 ],
                 "allow_any_host": false,
                 "serial_number": "abcdef",
                 "namespaces": [
                 {"nsid": 1, "name": "Malloc2"},
                 {"nsid": 2, "name": "Nvme0n1"}
                 ]
                }
                ]
...

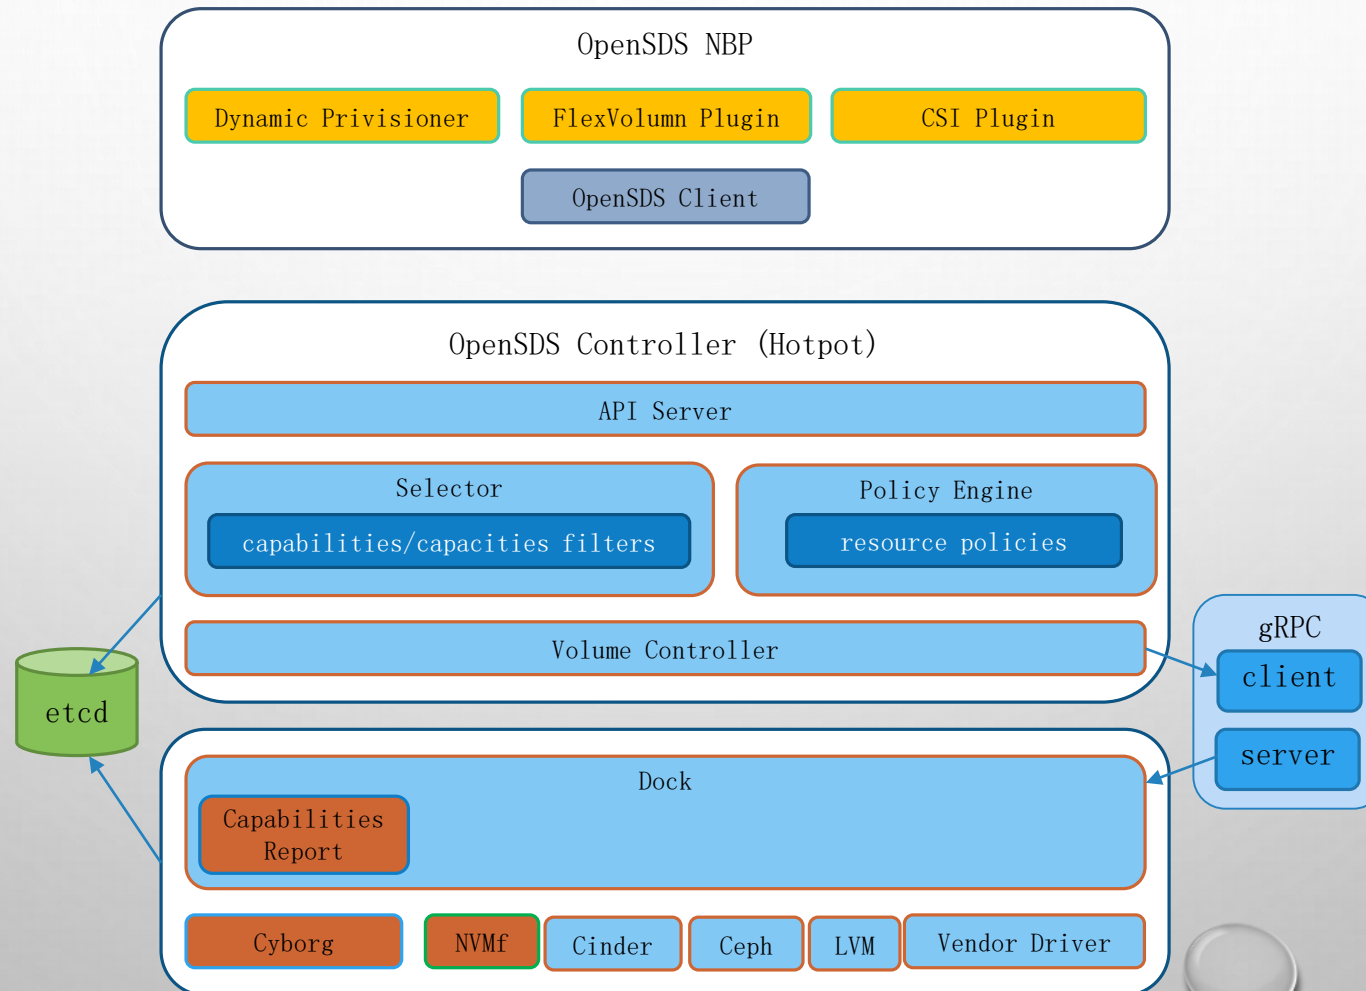
```

Cloud Storage Core Technologies——OpenSDS

- **Unified Plug-in Interface:** for different frameworks
- **Policy-Based Control:** Pooling of storage resources with policy-based control
- **Wide Storage Support:** support for wide range of storage drivers
- **Enterprise Class:** supports various hardware platforms (x86, ARM, ARM64, PowerPC,..)



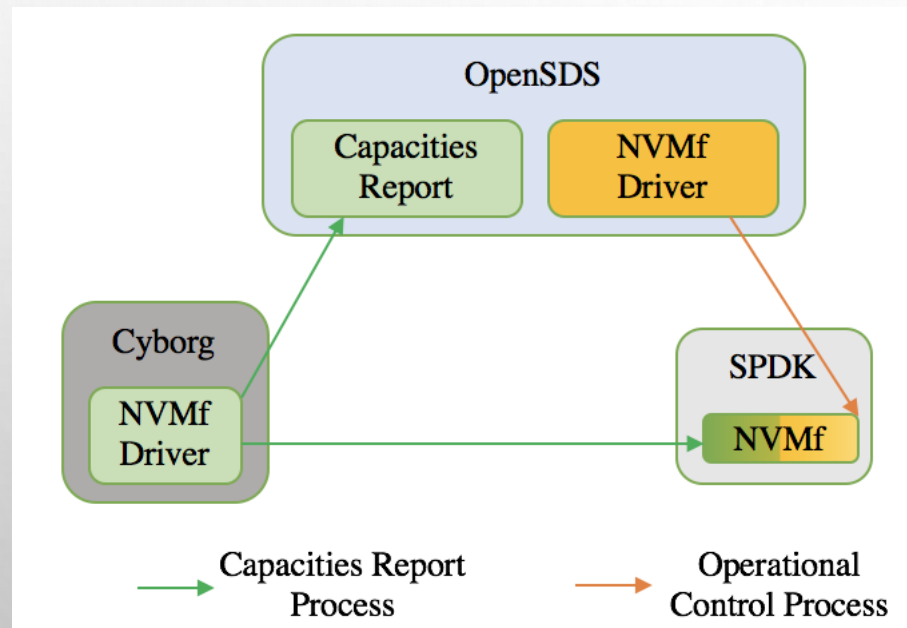
Cloud Storage Core Technologies——OpenSDS Enabled Storage Management



Outline

- ❑ Background
- ❑ Cloud Storage Core Technologies
- ❑ **Proposed Design**
- ❑ Conclusion

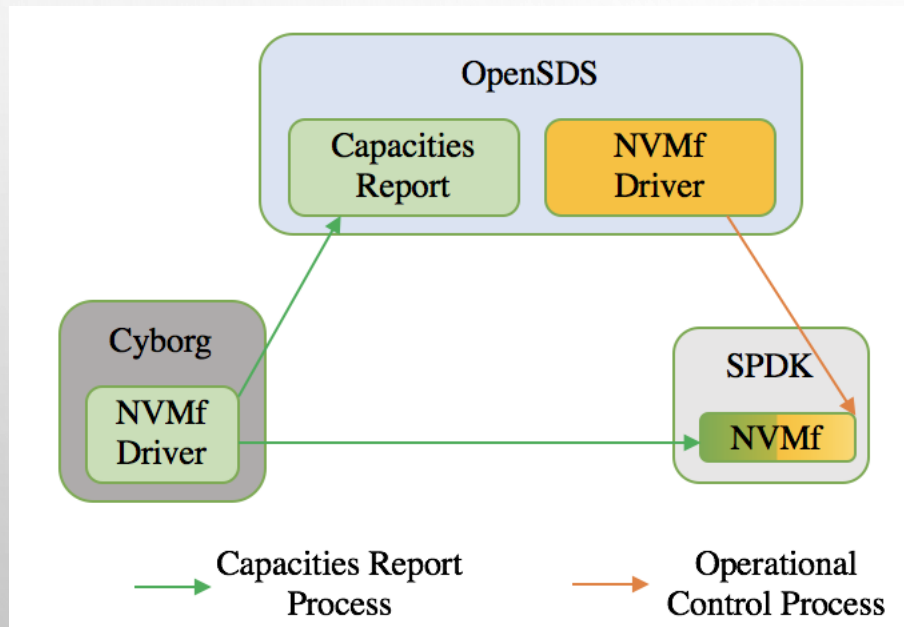
Proposed Design: Fine Grained High Performance Storage Management (1)



Capacities Report Process

1. OpenSDS start
2. OpenSDS Dock informs Cyborg to report the abilities of SPDK NVMf_target through 'Capacities Report'
3. Cyborg SPDK driver initialize the huge pages and start NVMf_target server, then get the abilities of accelerator (like subsystem, namespace, port, bdev, ...)
4. Cyborg report these abilities to Dock through 'Capacities Report'

Proposed Design: Fine Grained High Performance Storage Management (2)



Operational Control Process

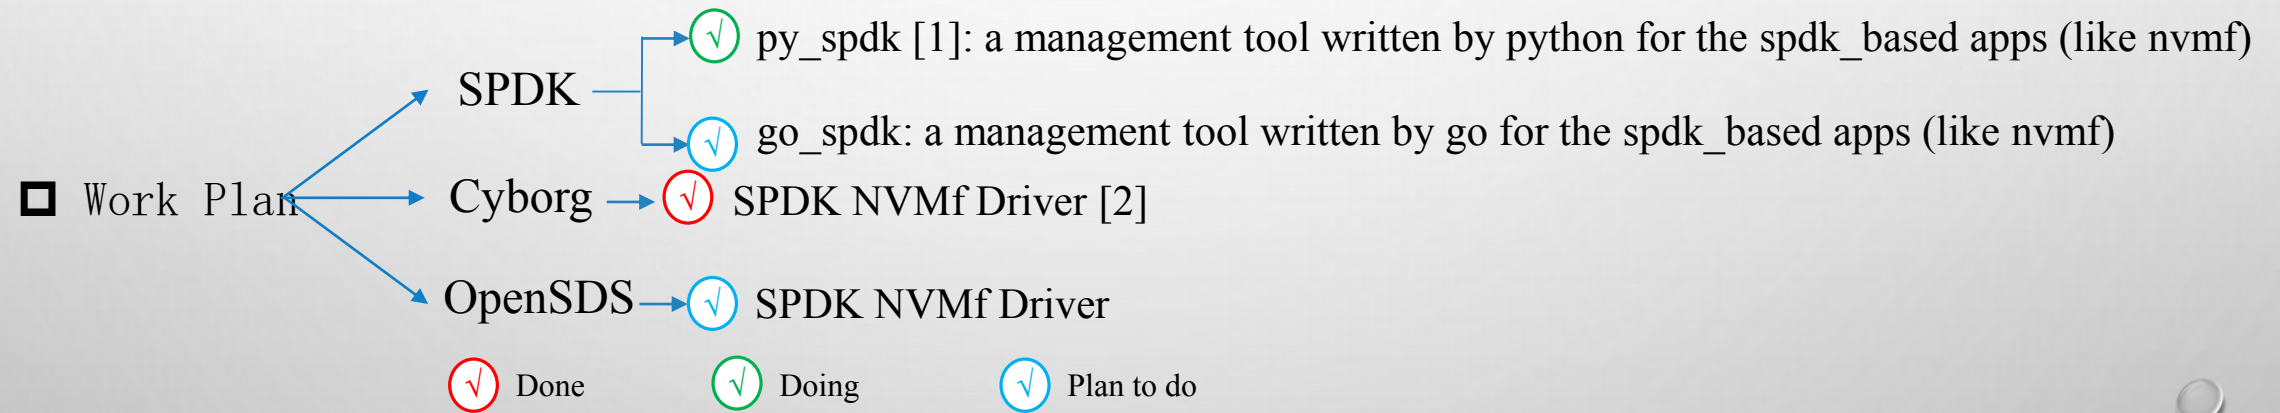
1. According to the returned abilities, OpenSDS starts creating the pool which is the basic operating unit of OpenSDS .
2. OpenSDS Dock starts executing specific actions (like add/delete namespace) through nvmf driver.

Outline

- ❑ Background
- ❑ Cloud Storage Core Technologies
- ❑ Proposed Design
- ❑ **Conclusion**

Conclusion

- Provide NVMe high performance storage backend for cloud computing:
 - SPDK implements the user space NVMe
 - Cyborg reports the abilities of NVMe to OpenSDS
 - OpenSDS manages and schedules the NVMe



【1】 <https://review.gerrithub.io/#/c/403506/>

【2】 <https://github.com/openstack/cyborg/tree/master/cyborg/accelerator/drivers/spdk/nvme>

The background features a light gray gradient with several realistic water droplets of various sizes scattered in the corners. The droplets have highlights and shadows, giving them a three-dimensional appearance.

Thank You

Q&A