CAICT 中国信通院

TRUCS2019
TRUSTED CLOUD SUMMIT
可信云大会
中国·北京 2019.7.2-3

High-performance Heterogeneous Resource Management — Cyborg

Huawei
Luwei he

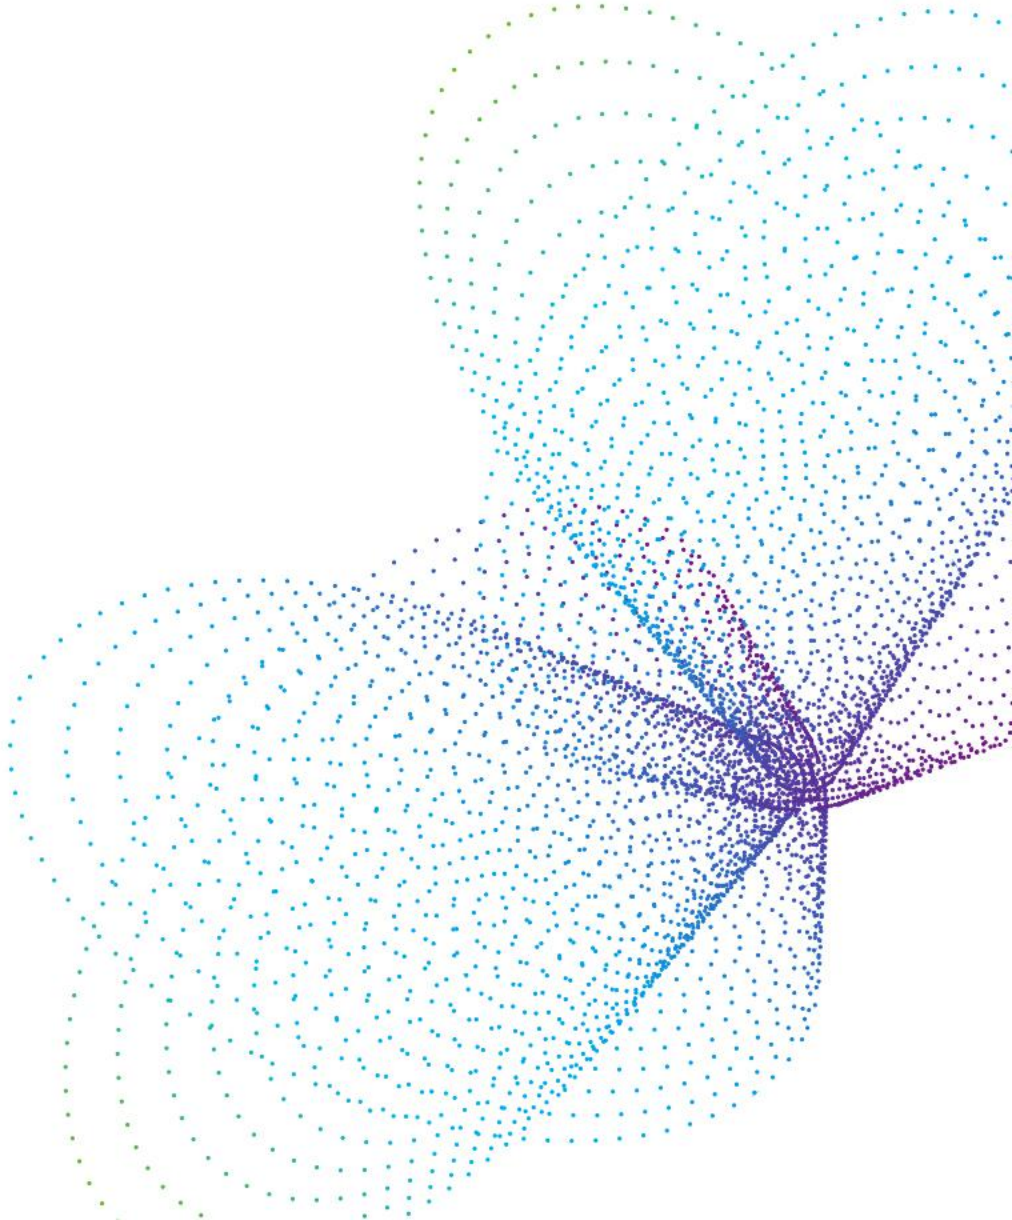

Agenda

1. Cyborg introdcution
2. Cyborg journey
3. Cyborg concept overview
4. Mainly work in Train release

New Era of Domain Specific Architecture

NPU

Neural network processors for machine learning

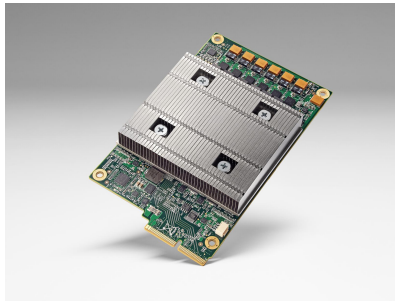

GPU

GPUs for graphics, virtual reality, ML

SmartNIC/FPGA

Programmable network switches and hardware

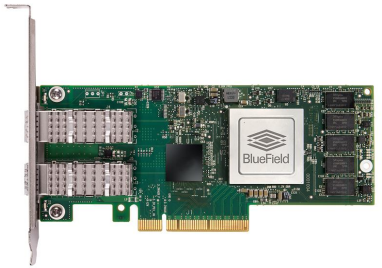

How To Manage These DSAs For Cloud ?

TRUSTED CLOUD SUMMIT
可 信 云 大 会

Cyborg is a general management framework for accelerators

Proud OpenStack Official Project since 2017.09
(https://github.com/openstack/cyborg)

An Unbelievable Journey

Cyborg's growth (popularity)

2016 .02- 2017.02

Cyborg's growth (popularity)

2016 .02- 2017.02

2017 .02- 2017.09

Cyborg's growth (popularity)

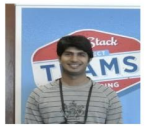
Rushil Chugh
Lenovo

Jinghua Gao
Lenovo


Yumeng Bao
ZTE


Shasha Guo
China Mobile


Justin Kilpatrick
Redhat


Xinran Wang
Intel


Yongfeng Du
Intel

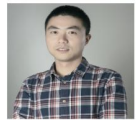
Rong Zhu
ZTE


Shaohe Feng
Intel

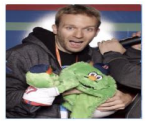
Jim Golden
NIST


Jinghan Sun
UIUC


Eric Fried
IBM

● ● ●

36
contributors
and growing


Xu Wang
Lenovo


Li Liu
Huawei, PTL


Li Zhu
Huawei


Pei Jia
Lenovo


Zhenghao Wang
Lenovo


Nadathur Sundar
Intel

2016 .02- 2017.02

2017 .02- 2017.09

2017 .09- 2019.06

Cyborg's growth (Diversity)

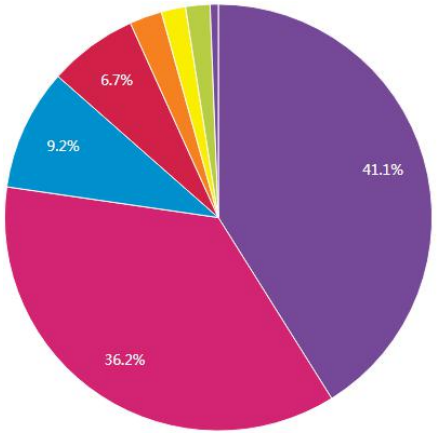

#	Company
1	Huawei
2	Lenovo
3	Fujitsu
4	Red Hat
5	Intel
6	SUSE
7	NEC
8	ZTE Corporation

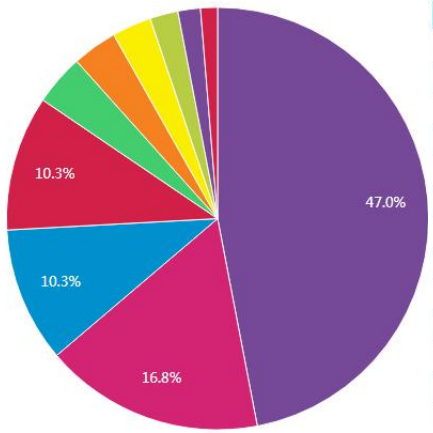

#	Company
1	Huawei
2	Intel
	*independent
3	Lenovo
4	99cloud
5	Red Hat
6	Fujitsu
7	NEC
8	IBM
9	Kylin Cloud
10	Yovole
11	Verizon Wireless
12	SUSE
13	Rackspace
14	Awcloud

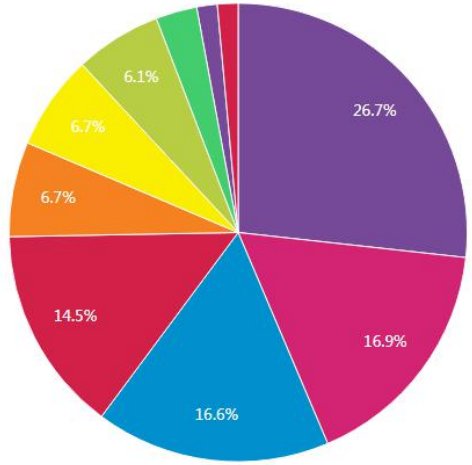

#	Company
1	Huawei
2	Intel
3	99cloud
4	Lenovo
5	Red Hat
	*independent
6	ZTE Corporation
7	SUSE
8	NEC
9	Kylin Cloud
10	Inspur
11	Ericsson
12	IBM

Q 2018.02

R 2018.08

S 2019.04

Cyborg's growth (Maturity)

Contribution Summary

Commits: **51**
LOCs: **13655**
Do not merge (-2): **2**
Patch needs further work (-1): **48**
Looks good (+1): **22**
Looks good for core (+2): **91**
Approve: **44**
Abandon: **0**
Change Requests: **50** (9 of them abandoned)
Patch Sets: **235**

Contribution Summary

Commits: **63**
LOCs: **19942**
Do not merge (-2): **15**
Patch needs further work (-1): **64**
Looks good (+1): **37**
Looks good for core (+2): **113**
Approve: **64**
Abandon: **3**
Change Requests: **103** (29 of them abandoned)
Patch Sets: **300**

Contribution Summary

Commits: **66**
LOCs: **14348**
Do not merge (-2): **13**
Patch needs further work (-1): **138**
Looks good (+1): **72**
Looks good for core (+2): **121**
Approve: **71**
Abandon: **0**
Change Requests: **128** (38 of them abandoned)
Patch Sets: **331**

Cyborg's growth (Arch Evolution)

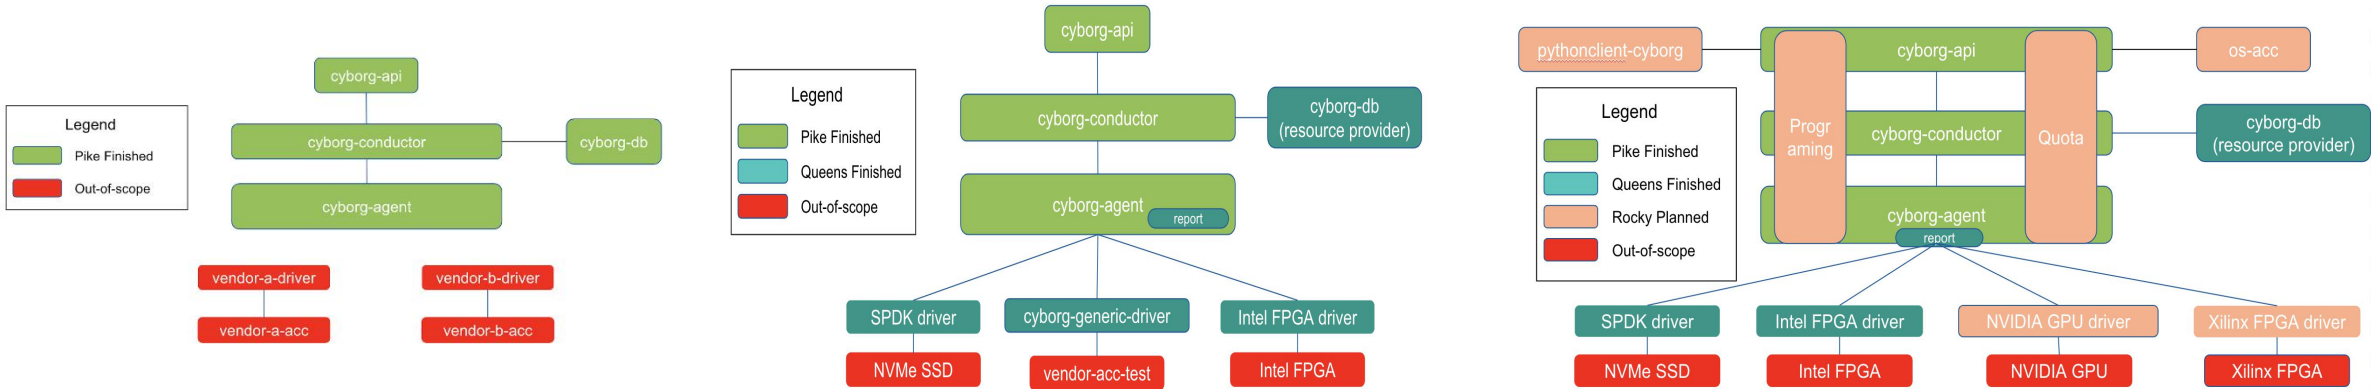

2017.09 2018.02 2019.04

Cyborg arch overview

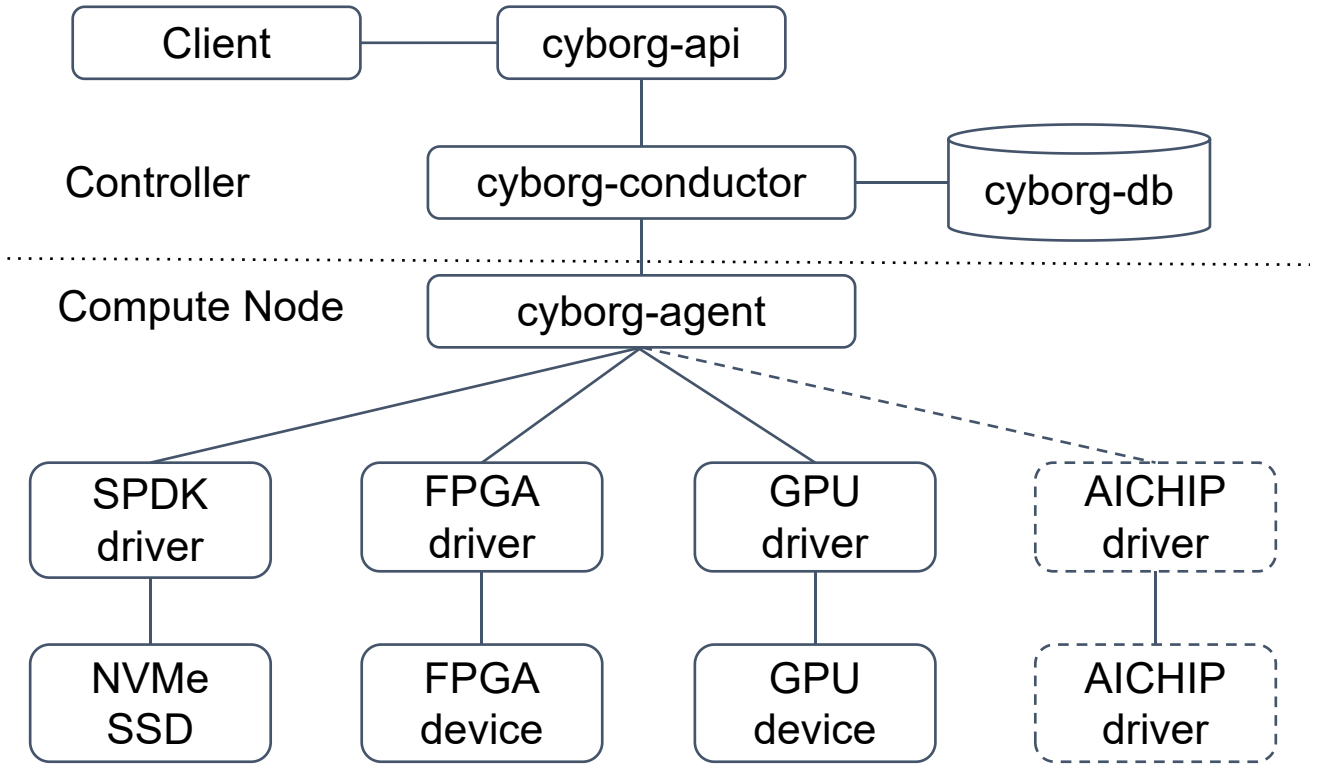

Abstract Device Model

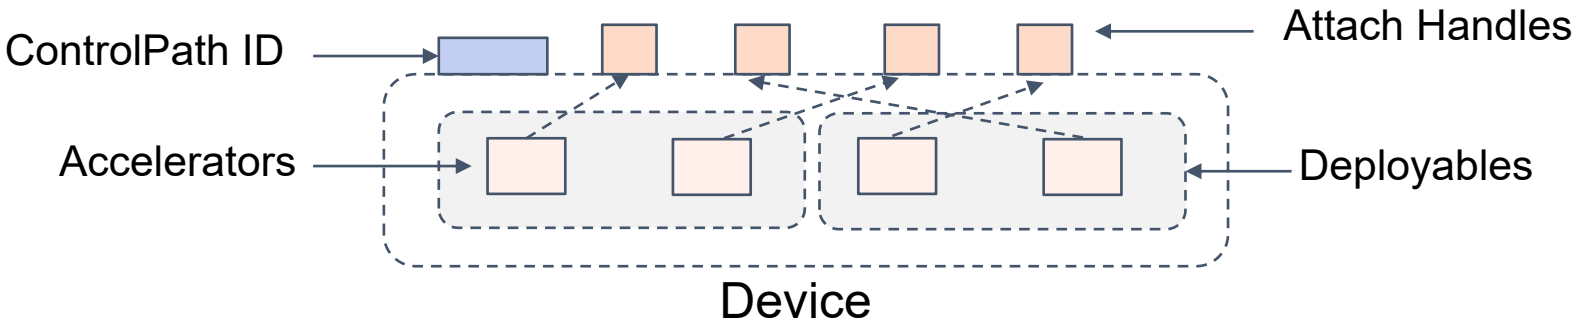

Term	Meaning	Placement Represent
Deployable	A logical structure in a device that provides a resource. A resource can be an accelerator, local memory, etc.	Resource Provider
Accelerator	A logical resource to offload computation, etc.	Resource Class Inventory
Device	Physical hardware. E.g. PCI card. Includes board (Flash/BMC).	--
ControlPath ID	Unique identifier to access the device. E.g. PCI PF.	--
Attach Handle	An ID of the handle to attach to an instance. E.g. PCI VF, mdev UUID.	--

Devices discover

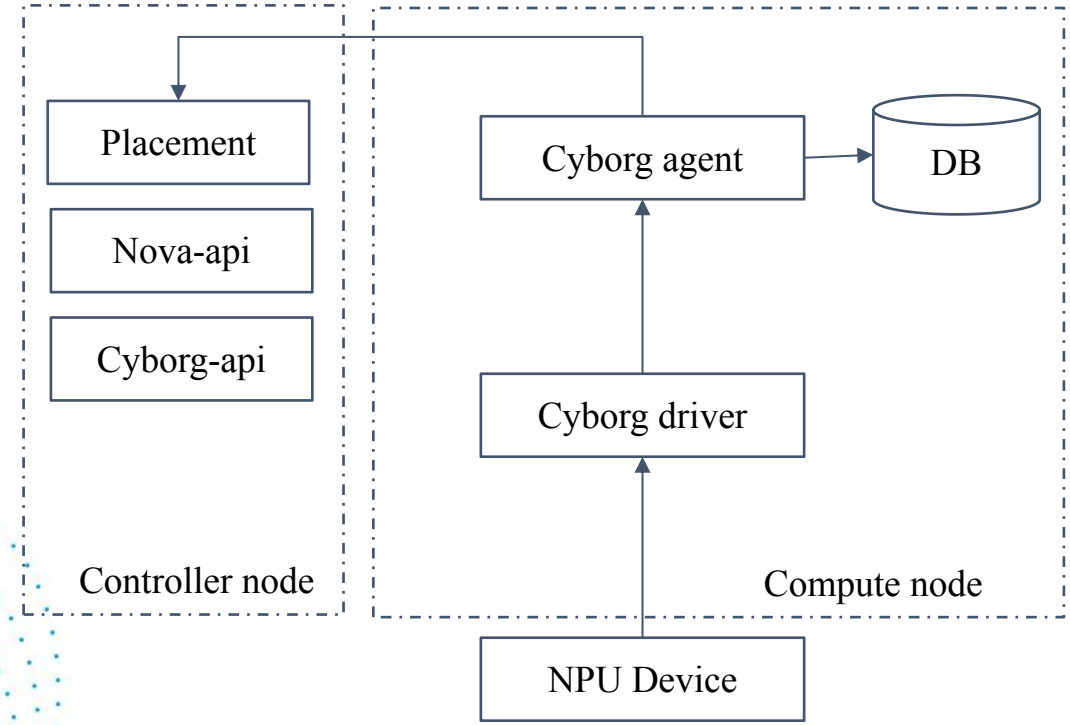

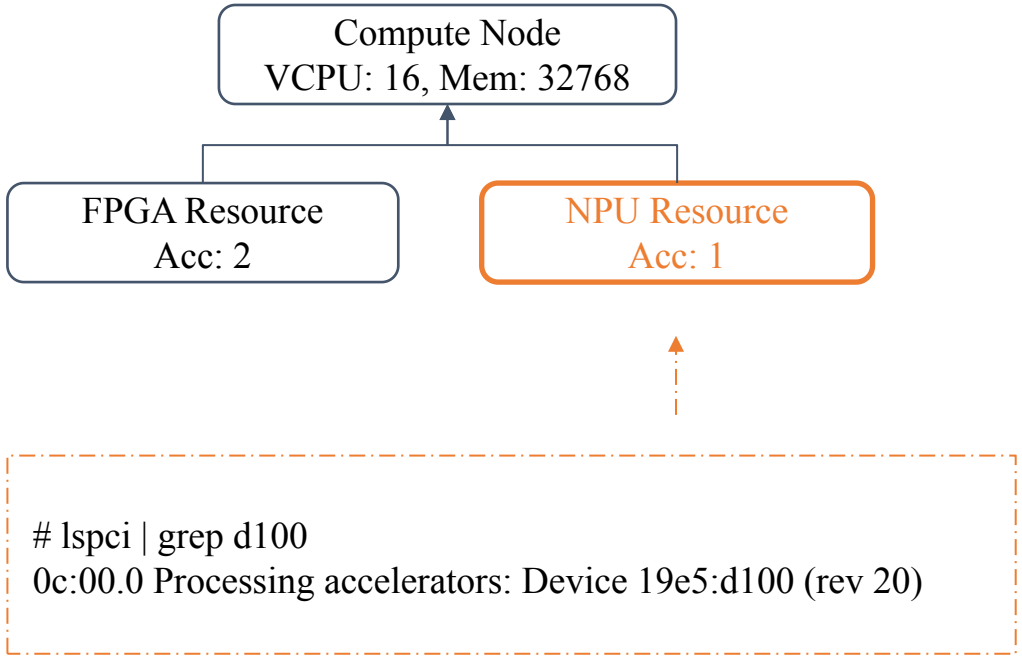

The mainly work in 2019 (Train Release)

- Nova-Cyborg Integration.
- Generic Driver Implementation.
- Driver improvement and support.
- Python 3 migration.
- Testing and validation.

https://etherpad.openstack.org/p/cyborg-train-goals

Nova-Cyborg Integration

Enable requesting an instance with one or more accelerators either preprogrammed or dynamically programmed.

This encompasses FPGAs managed by Cyborg as well as VGPUs (of multiple types) managed by Nova.

Note: This includes cross-project work with Cyborg. The Cyborg team's cycle priorities are aligned accordingly.

[1] https://specs.openstack.org/openstack/nova-specs/priorities/train-priorities.html

[2] Nova-Cyborg spec: https://review.opendev.org/#/c/603955/ Merged in 2019.06.20

Nova-Cyborg Integration

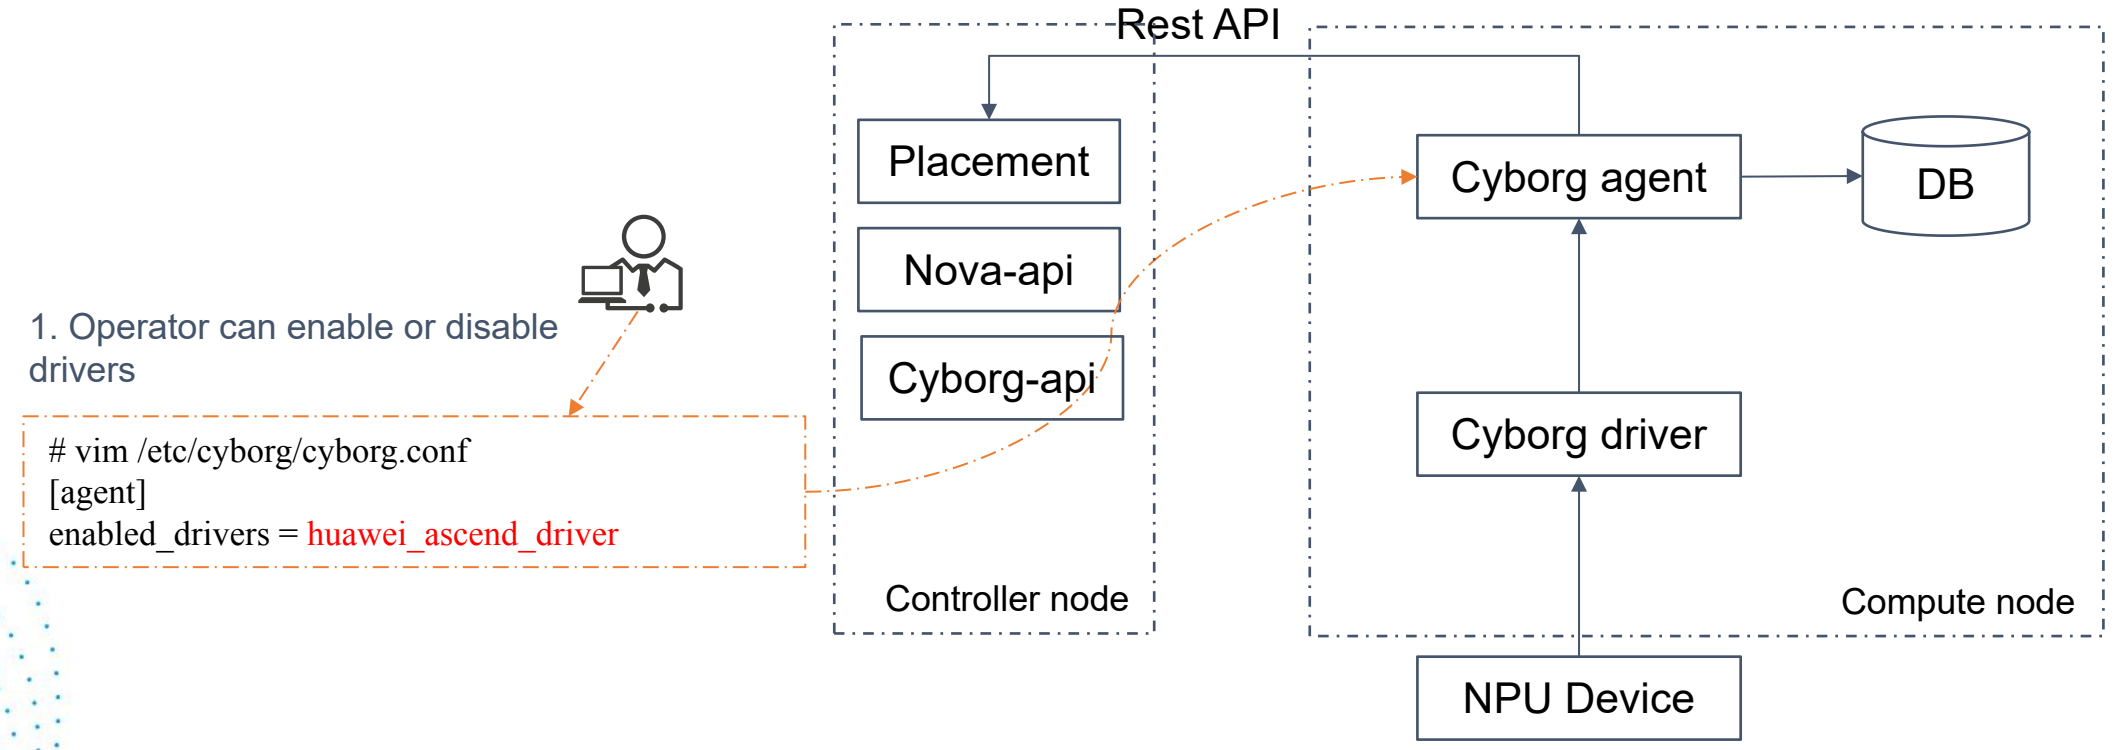

Nova-Cyborg Integration

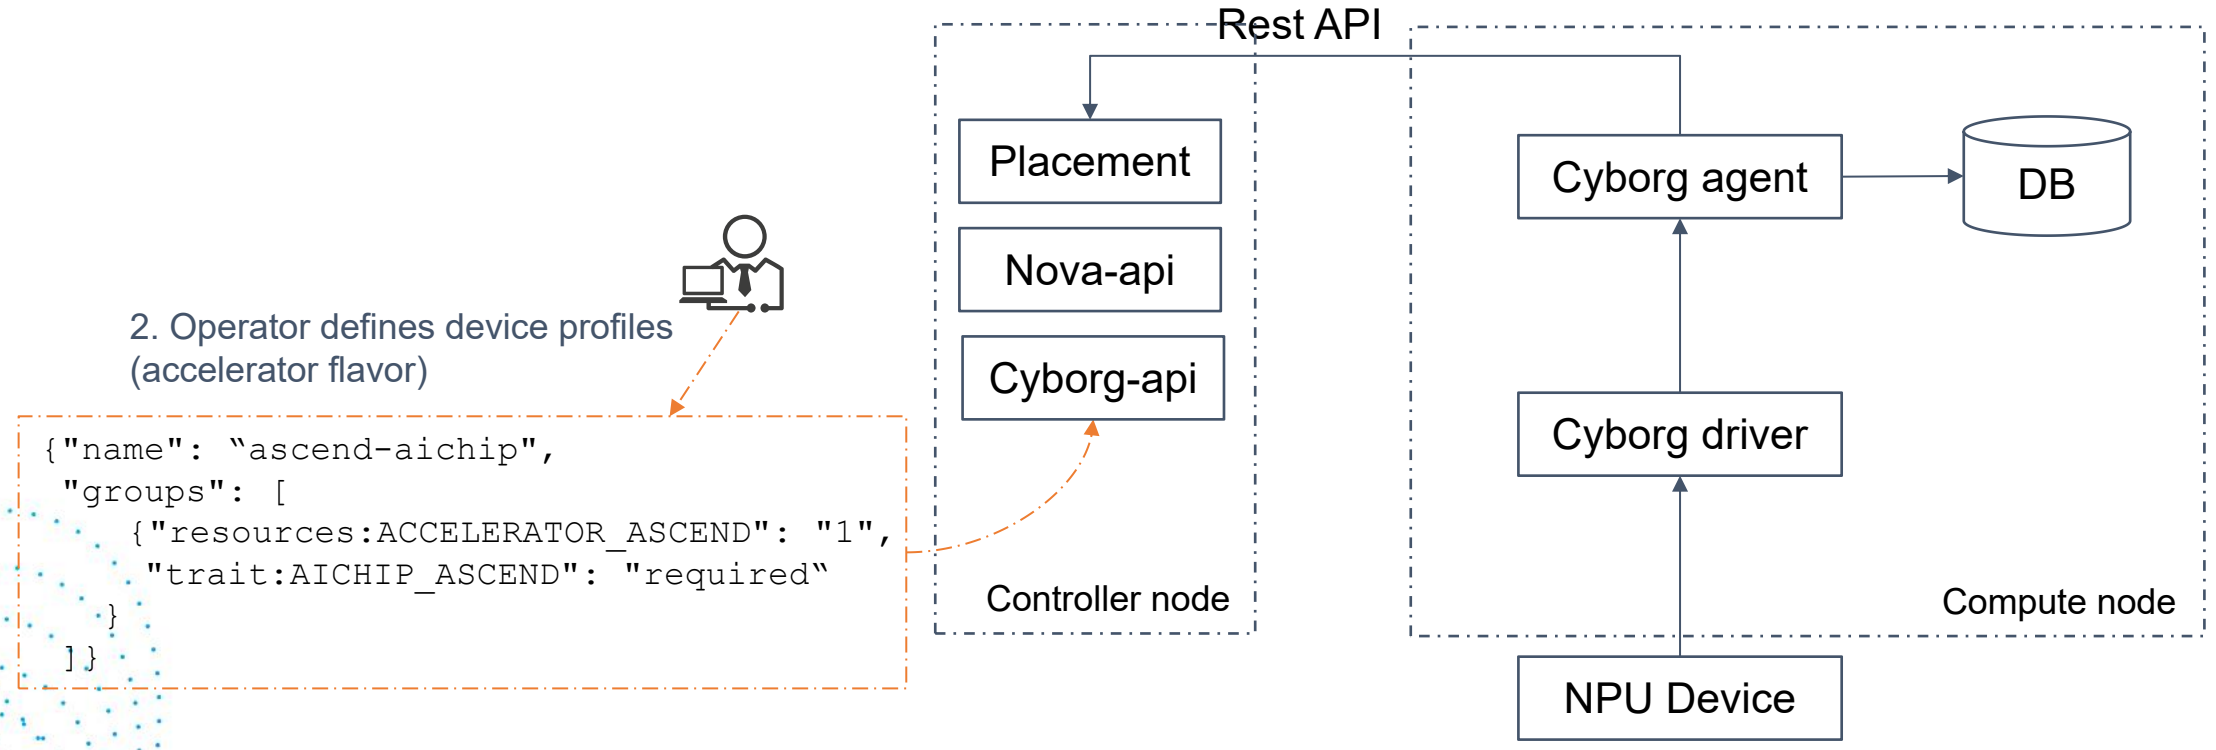

Nova-Cyborg Integration

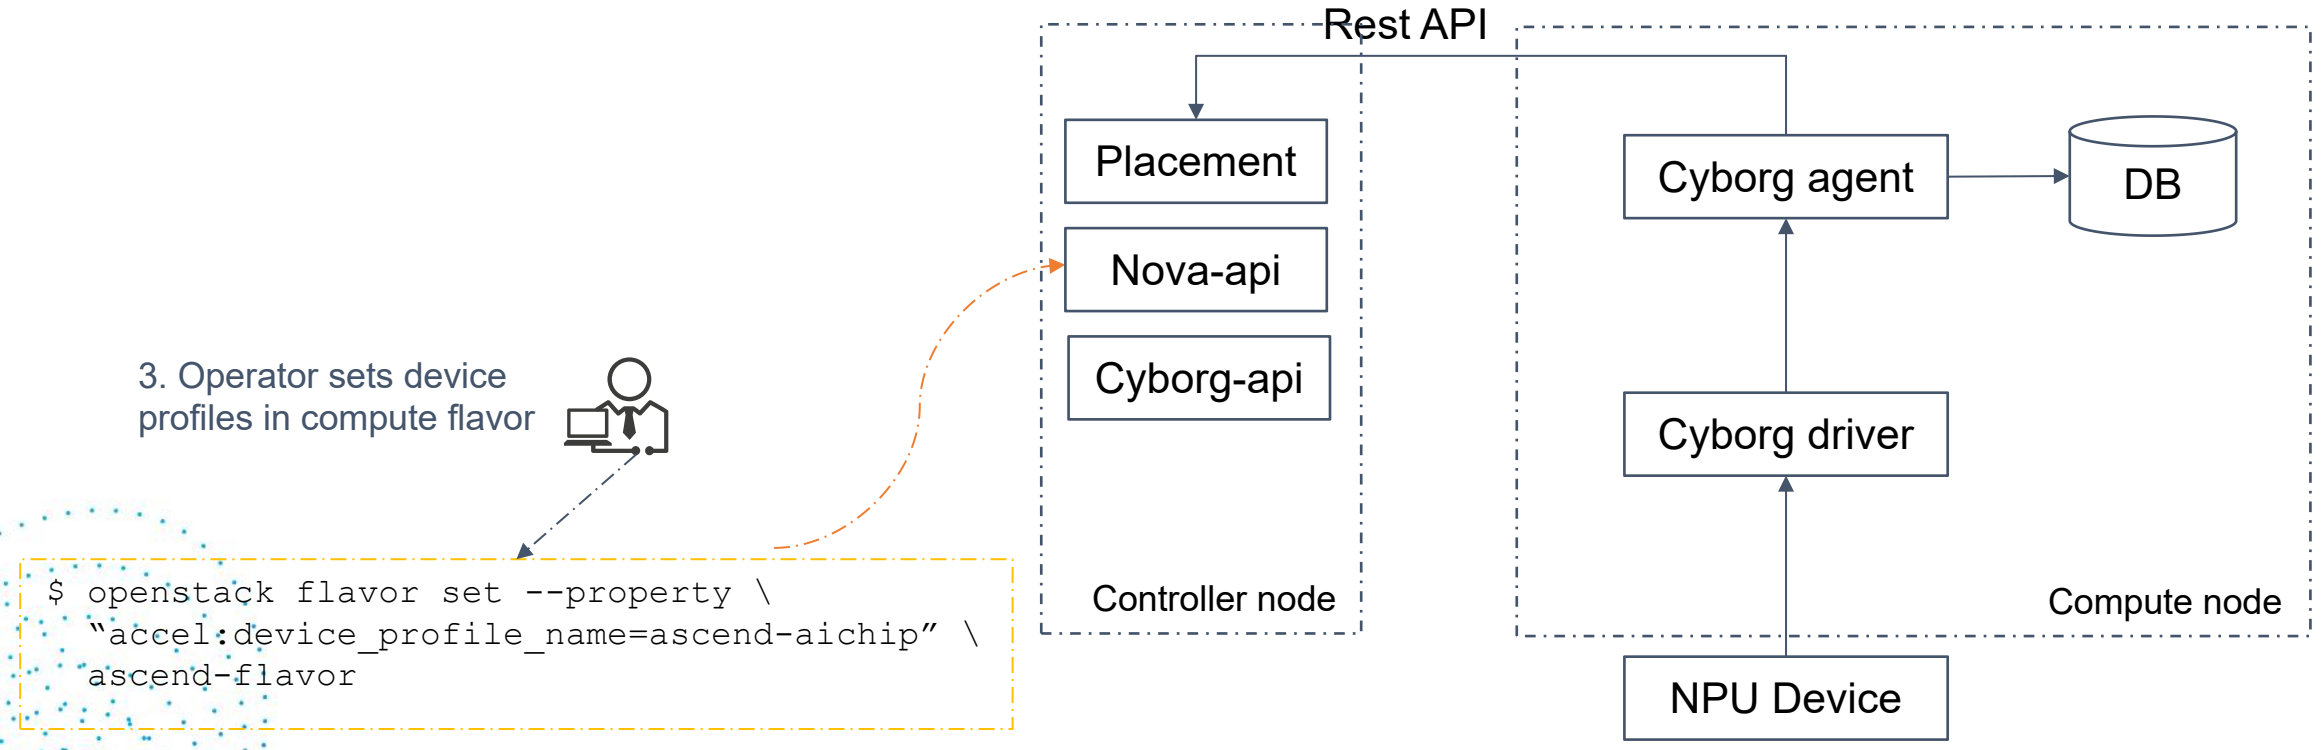

Nova-Cyborg Integration

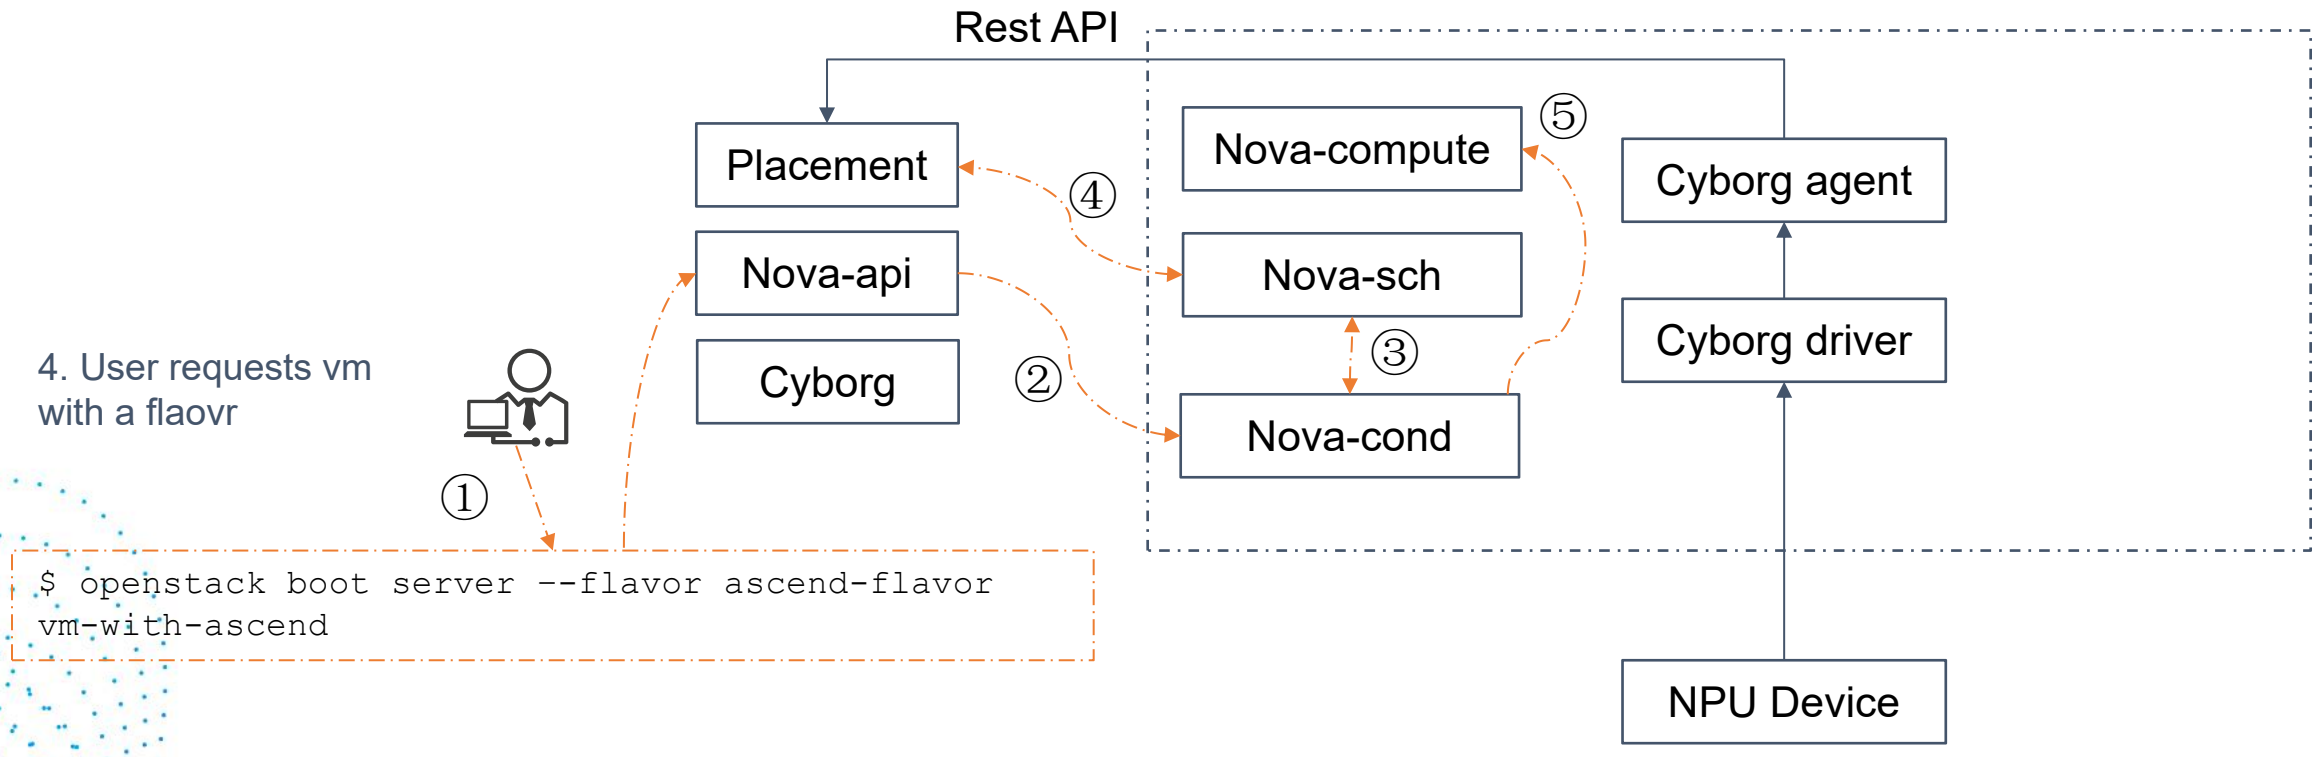

The key flow of nova-cyborg interaction

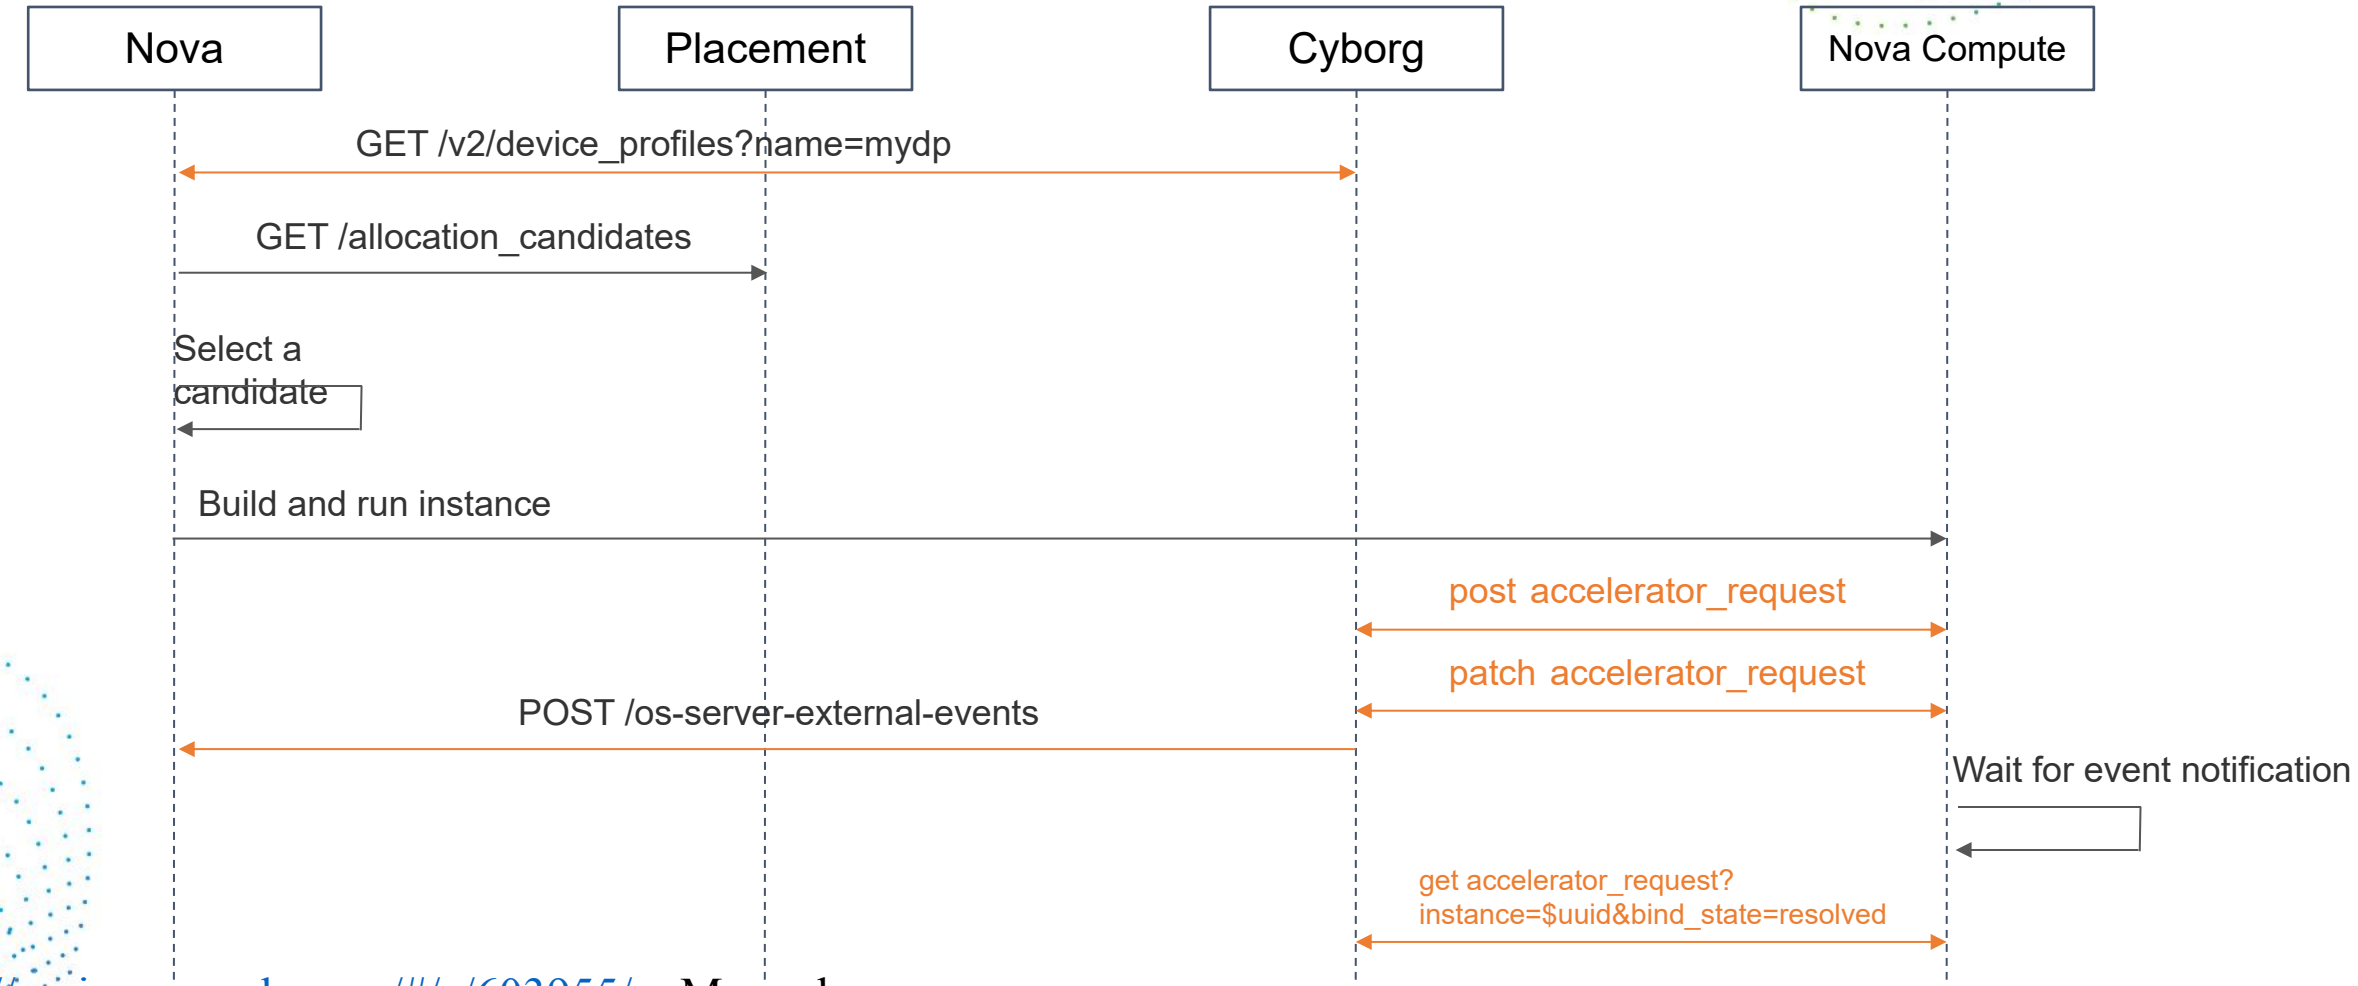

https://review.opendev.org/#/c/603955/ Merged

Generic driver

In order to easy to add the support for the common accelerator, we propose to improve the generic driver to manage these devices.

[1] https://github.com/openstack/cyborg/blob/master/cyborg/accelerator/drivers/generic_driver.py

```
class GenericDriver(object):

    # Discover a specified accelerator.
    def discover(self):

    # Update the device firmware with specific image.
    def update(self, control_path, image_path)

    # Collects device stats
    def get_stats(self):
```

What we want to be done

```
class FPGADriver(GenericDriver)

class GPUDriver(GenericDriver)

class NPUDRIVER(GenericDriver)

class FakeDRIVER(GenericDriver)
```

Driver improve and support

- FPGA Driver
- Nvidia GPU Driver
- （New）Huawei Ascend Driver
- （New）Intel Movidius Driver

Join us

IRC: #openstack-cyborg

Wechat: Cyborg中国开发讨论组

CAICT 中国信通院

TRUCS2019

TRUSTED CLOUD SUMMIT

可 信 云 大 会

中国·北京　2019.7.2-3

THANKS