2021-06-10 - Anuket: Hardware Acceleration Abstraction

Topic Leader(s)

- arkady kanevsky (no show)
- Petar Torre

Topic Overview

60m, arkady kanevsky

Anuket is promote HW independent RM, RAs and RC. For accelerators for RAN workload oRAN Alliance is defining Acceleration Abstraction Layer and APIs - https://oranalliance.atlassian.net/wiki/download/attachments/872841331/O-RAN.WG6.AAL-GAnP-v01.00.pdf?api=v2. However Kubernetes currently does not have an abstraction for accelerators and each accelerator is handled and need to be programmed to independently. This session will discuss current state of affairs and what can be done to remedy it.

Slides & Recording

YouTube

Please indicate your session type in the blank space below and then remove this Info field.

- Demo / Informational (non-interactive)
 - You may be asked to pre-record this session which will be made available on-demand.
- Live Interactive Session

LFN Staff may elect to publish some videos to YouTube. Please indicate here if you do not want your session to be published to YouTube.

Recording: Hardware Acceleration Abstraction.mp4

Agenda

Open discussion without presentation.

Minutes and next

- Current status in Anuket
 - RM Ch3 describes
 - Accelerators in an abstract way,
 - Motivations for accelerators (time precision, power or space envelopes, simple optimisation)
 - Categorization of accelerators
 - How to consume accelerators
 - ° RA2 Ch3 and Ch4
 - NFD and Device Plugins
- · What is missing
 - ° RM
- O-RAN AAL
 - The aim is to provide an abstract API to consume accelerators
 - RM mentions AAL
 - Anuket should wait till O-RAN finalizes AAL MVP and then refer to it. It will be radio-specific first, later more generalized HW acceleration model.
 - Stay aligned between Anuket and O-RAN: Anuket reps (Tomas Fredberg, Gergely Csatari, Karine Sevilla, Petar Torre
 ...) to sync with their companies' O-RAN reps to try to reduce overlaps and keep two communities in sync
 - Discussion about k8s networking models (so would be relevant to RA2) (Per Andersson)

° RA1

- Cyborg is planned to be introduced in Lakese. This was a strong motivation to change the OpenStack baseline version to Walaby. (led by Karine Sevilla and Pankaj Goyal)
- RA2Device Plugins:
 - Recommendation on what API Device Plugins should expose to be more portable between vendor implementations
 - Device plugin API-s are fragmented
 - We should aim for a more portable API design
 - At the moment there is no consensus on this, but maybe Anuket has enough vendors and operators to build a consensus and convince device plugin vendors to modify their API design
 - Gergely Csatari or Riccardo Gasparetto Stori to schedule discussion on some of RA2 calls

- ° RI1
 - Wait until RA1 work is finished
- ° RI2
- Support exist in the tools used to install RI2 (Intel BMRA), even some of these features are enabled (led by Michael Pedersen)
 Next need HW in lab

RC1RC2

Work towards basic test cases started (Michael Pedersen)