

# Exposing NIC Information to Userland

Netconf 2019

Zorik Machulsky

AWS

zorik@amazon.com

# What are we trying to solve?

- Running cloud infrastructure poses multiple challenges, among them simple and intuitive information exposure to the users.
- We are looking for the way to communicate certain network infrastructure related information directly to the user through well known OS mechanisms/tools
- Examples:
  - Identify OS network interface serving as Primary ENI (Elastic Network Interface) – *Info Message*
  - Identify ENI associated with EFA (Elastic Fabric Adapter) – *Info message*
    - EFA is a different device (HPC interconnect) that shares the same ethernet infrastructure with ENA
  - Recommend to upgrade an outdated/buggy driver - *Warning message*
  - More to come in the future

# What are we trying to solve? (cont.)

- This information should be easily extendable, ready to reflect future enhancements and evolving infrastructure.
- Driver is rarely updated by the customers, and hence usually is not up-to-date with the recent infrastructure changes.
- In other words, we need to be able to establish a direct NIC-Userland channel for communicating info/warning messages to the user, where driver would handle those messages as a “black box” and forward them up to the user space tool.

# What did we try?

- We hoped that ethtool private flags could be a legitimate mechanism for exposing device private info to userland
- This was criticized and rejected by the community due to the following reasons:
  - In our patch the flags and strings are retrieved from the NIC and the driver handles them as opaque objects and just forward them to the user space.
  - Private flags are not expected to serve as labels, tags or info/warning messages of any kind
- Sysfs option was also discussed and found to be confusing and not serving all use cases.

# What we propose and want to discuss

- `ethtool -S` today exposes driver and NIC arbitrary stats
- In a similar way we might consider introducing a new `ethtool` function that would be used to query info/warning messages from the NIC

# What's the outcome of the discussion?

- During netconf 2019 discussions we've been advised to evaluate the devlink tool.
- It sounds like devlink provides functionality we need
- In a long run devlink is expected to supersede the ethtool
- We are going to study devlink, and, if found suitable for our needs, implement its APIs in the ENA driver.