Brussels, FOSDEM 2023

Open Source Switching: Upstreaming ONIE NVMEM and switch BSP drivers

Jakov Petrina Trnski

sartura

February 4th, 2023

About

- Delivering solutions based on GNU/Linux firmware
  - Focused on network edge, network switching, and CPEs
  - OpenWrt, Gentoo, Yocto, etc.
- Continuous participation in Open Source projects
  - Contributions for the Linux kernel, systemd, etc.

Open source switching

- Network switch
  - Multiport network device with packet switching ASIC
  - Forwards data at the data link layer (layer 2)
- `switchdev`
  - In-kernel Linux driver model for switch devices which offload the forwarding (data) plane from the kernel
- Open Compute Project *ONIE*
  - Install environment for bare-metal network switches

Challenges and solutions

- Limited mainline platform support
  - Marvell Prestera, NVIDIA Mellanox Spectrum, Microchip SparX-5
- Building a standardized open-source product
  - The Linux Foundation DENT project and community
- Upstreaming work lead by the DENT Upstream Working Group
  - Linux ONIE NVMEM driver by Bootlin
  - Linux Board Support Package (BSP) drivers by Sartura

Linux ONIE NVMEM support

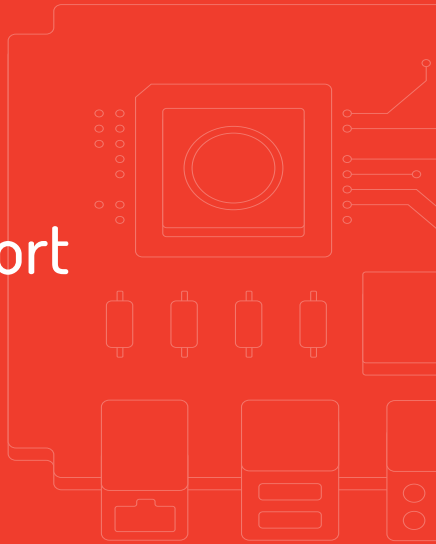

sartura

Linux ONIE NVMEM project

- Organized and funded by the DENT project, completed by Bootlin
- Hardware requirements from the ONIE project
  - "Hardware platform must provide a non-volatile storage (EEPROM, NOR, or NAND) which contains vital product data assigned by the manufacturer"
  - New `TlvInfo` EEPROM format
- Project goal to standardize access to the switch hardware platform data

Linux ONIE NVMEM project

- Project timeframe
  - Started on *August 18, 2022*
  - Status updates from September until December
  - Completed successfully on *January 8, 2023*
- End result
  - Support from Linux 6.3 for parsing information stored in an ONIE-compliant EEPROM or other flash memory

Linux ONIE NVMEM driver

- Pre-requisites
  - Necessary to define and implement concept of "NVMEM layouts"
  - *NVEM layouts* implementation
    - They can add NVMEM cells during runtime (!)
    - Post-processing for endian swapping, or ethernet offset handling, etc.
- Upstream switchdev drivers updated to use new API

Linux switch board-support

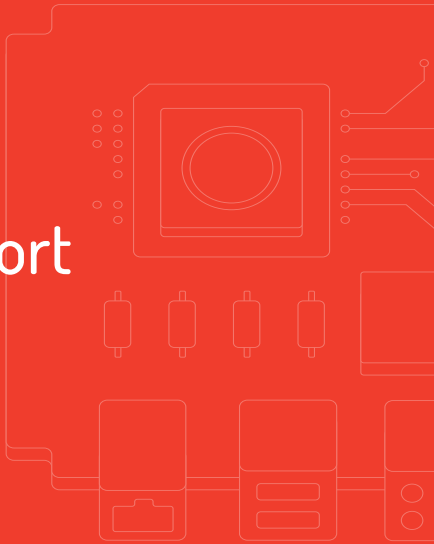

sartura

Linux switch BSP drivers

- Organized and done by Sartura for the Replica.one build system project
- Working on a variety of platforms and features
  - Delta DPS-920AB PSU driver
  - Microchip PD69200 PoE PSE driver
  - *Delta TN48M CPLD support* (upstreamed)
    - GPIO driver
    - CPLD reset controller

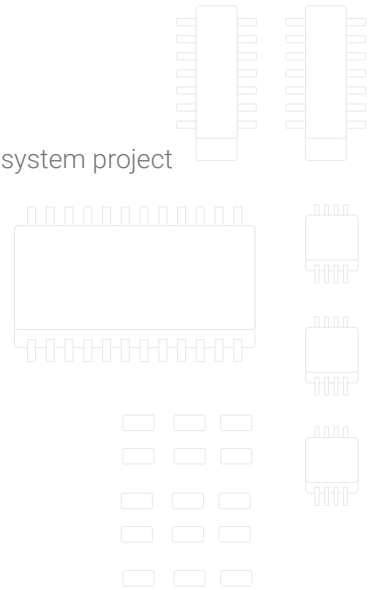

Linux switch BSP drivers

```
[PATCH v10 0/6] Add Delta TN48M CPLD support
[PATCH v10 1/6] mfd: simple-mfd-i2c: Add Delta TN48M CPLD support
[PATCH v10 2/6] gpio: Add Delta TN48M CPLD GPIO driver
[PATCH v10 3/6] dt-bindings: reset: Add Delta TN48M
[PATCH v10 4/6] reset: Add Delta TN48M CPLD reset controller
[PATCH v10 5/6] dt-bindings: mfd: Add Delta TN48M CPLD drivers bindings
[PATCH v10 6/6] MAINTAINERS: Add Delta Networks TN48M CPLD drivers
```

Planned activities

- Power over Ethernet (PoE) Linux kernel subsystem
  - standardize user-space API, PoE manager daemon
  - allow PoE controllers to be easily implemented
- User-space switch management tooling
  - ethtool, systemd-networkd, wired 802.1x support
- Questions? Ideas? Get in touch!

Open Source Switching: Upstreaming ONIE NVMEM and switch BSP drivers

jakov.petrina@sartura.hr

info@sartura.hr · www.sartura.hr