

# Devicetree Overlay use at Juniper Networks

Guenter Roeck  
groeck@juniper.net

# System Overview

- PTX5000 Packet Transport Router
  - Routing Engine
    - Routing protocols, administrative tasks
    - Interfaces to other cards in the system
  - 8 x FPC (Flexible PIC Concentrator)
    - 2 x PIC per FPC
  - Control Board
    - 9 x SIB (Switch Interface Board) per CB
  - All cards identified using I2C EEPROMs
  - Card connectors use multiple interface types
    - I2C, GPIO, PCIe, SERDES, ...
  - Various CPU types
    - P2020, P5020, P5040, x86

# Devicetree overlay use

- All OIR capable cards managed with devicetree overlays
  - RE
    - FPCs, Fan tray, power supply, ...
  - FPC
    - PICs
  - Control Board
    - SIBs
- Each card represented as 'connector' node in devicetree data

# 'connector' nodes

```
pic0 {
    compatible = "jnx,pic-connector", "simple-bus";
    slot = <0>;
    auto-enable;
    ovname = "jnx_pic0", "jnx_pic0_pwr";
    presence-detect-gpios = <&gpio20 148 0x1>; /* active low */
    attention-button-gpios = <&gpio20 150 0x1>; /* active low */
    power-enable-gpios = <&gpio20 154 0x0>; /* active high */
    power-status-gpios = <&gpio20 151 0x0>; /* active high */
    reset-gpios = <&gpio20 153 0x1>; /* active low */
    power-enable-timeout = <2000>; /* in ms */
    attention-button-holdtime = <3000>; /* in ms */
    activation-timeout = <5000>; /* in ms */
    debounce-interval = <1>;
    led-green = <&pic0_green>;
    led-red = <&pic0_red>;

    i2c-bus {
        #address-cells = <1>;
        #size-cells = <0>;

        i2c-parent = <&pic0i2c>;

        eeprom@54 {
            compatible = "atmel,24c02";
            reg = <0x54>;
            ideeprom;
        };
    };
};
```

# Connector driver

- **Functionality**
  - Manages card insertion and removal
  - Responsible for loading and removing devicetree overlays
  - State machine with 10 states and 12 events
- **Status**
  - Reliably loads and removes overlays
  - Some limitations and concerns

# Limitations

- Power management
  - After enabling power, chips may be immediately visible on bus
    - PCIe: hotplug driver attempts to load driver before overlay is loaded
  - Kind of solved by using layered overlays
    - First overlay inserted after card identified, prior to enabling power
    - Second overlay inserted after power enabled and stable

# Limitations

- Indirect target support
  - Currently requires information within overlay for each slot
  - Problematic if card is re-used in a different chassis
  - Limited scalability
  - Proposal: Simplify API by providing reference(s) from calling code
    - `of_overlay_indirect()` gets reference(s) instead of slot number as parameter

# Limitations

- No DT / DT Overlay support on x86
  - Mandatory for us
  - Other solutions either not feasible or not scalable
    - ACPI
      - Not supported on all architectures
      - No overlays
    - Platform data is clumsy
      - Requires new driver / code for each new card
    - Card management from user space does not work
      - Yes, we tried ...
  - Implemented and working with small patch set on top of upstream kernel