cesnet

Precise Time with SDN Whiteboxes

Michal Hažlinský, Jan Kundrát

CESNET, a.l.e.

TNC25

2025-06-10

CESNET

- E-infrastructure service provider
  - Czechia, EU
  - Science, research & education
  - Services
    - Network
    - Compute
    - Storage
    - Collaborative environment

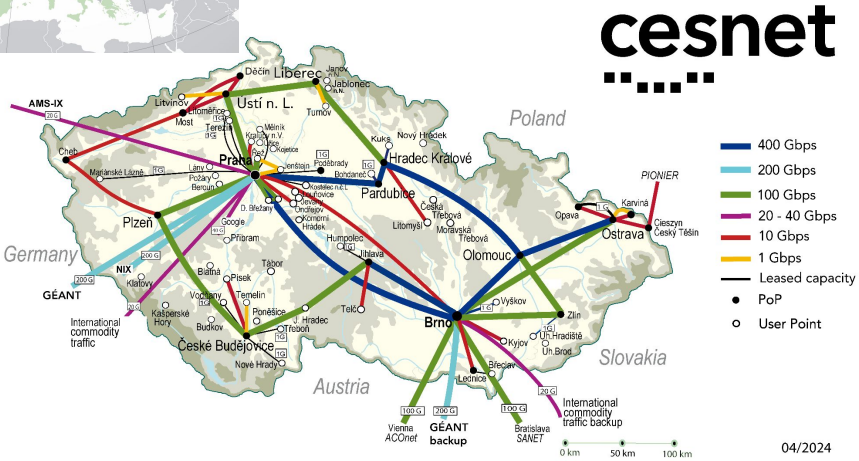

Time Distribution Service

- Nationwide
- Long Haul White Rabbit with bidirectional amplification
- Over 20 PoPs and 40 amplifiers
- Shared fibers with production 400G network for Research and Education
- Integration into management and monitoring tools

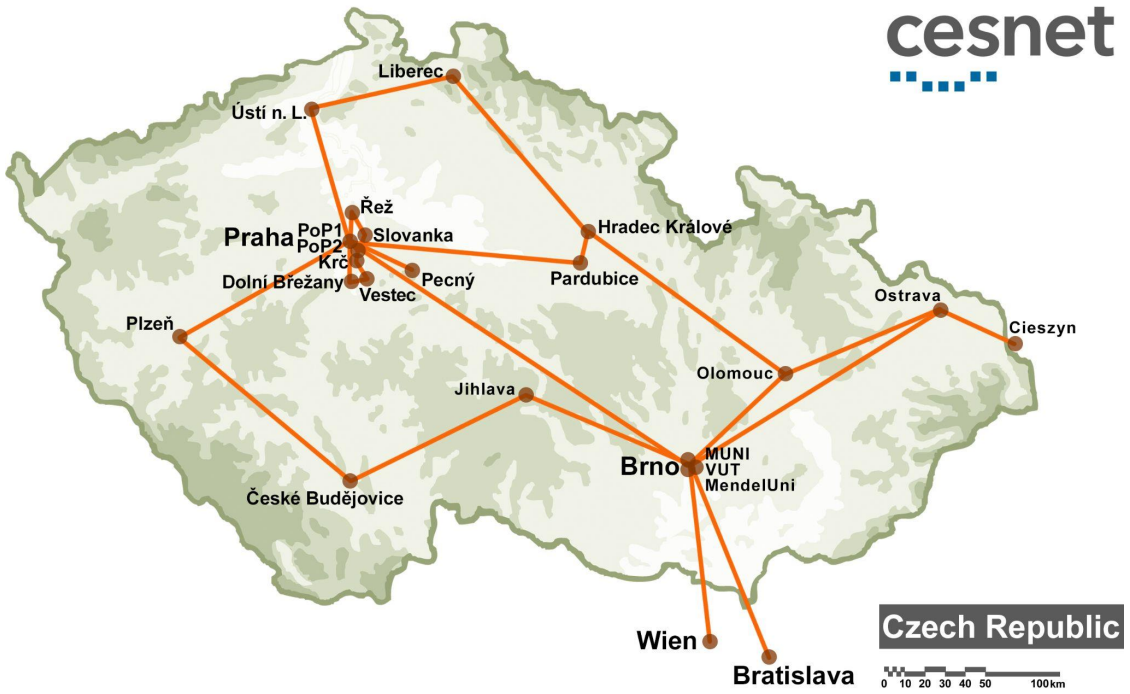

cesnet

Challenges

- Bidirectional transmission over existing fibre lines equipped by commodity transmission systems
  - Need to add/drop BiDi channels at each inline node for amplification
  - Filters have already been part of the CESNET3 design (future proof design work)
- Long Haul White Rabbit with L-Band DWDM SFPs
  - There are not many vendors supplying 1GE DWDM transceivers for L-band
  - Compatibility issues
- Integrate the management and monitoring of new equipment into the existing ecosystem
  - SDN driven solution using open source tools and CzechLight amplifiers

Full Demo Setup

- Fiber line (3 spans)
- Coherent 400G signal via fiber pair
- White Rabbit protocol for precise time transfer over one fiber of the pair
- CzechLight OLS
  - 2x ROADM
  - 2x InLine EDFA
  - 2x BiDi EDFA
  - 6x DWDM filters for BiDi channels
- Rubidium timebase and time interval counter for time comparison

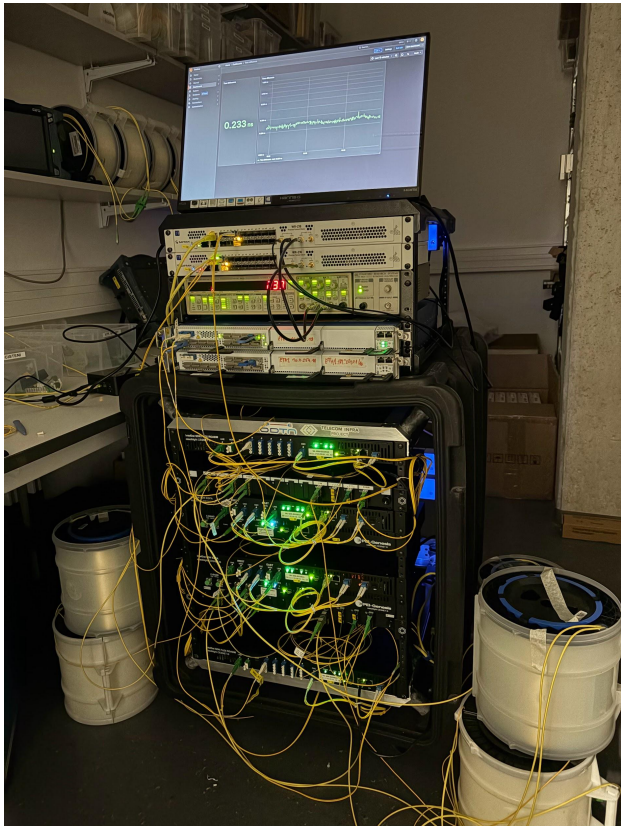

Staging Setup Topology

- 300 km fiber line in lab
- 1GE LR DWDM SFPs for L-Band
  - 2 channels 1570 nm and 1571 nm
  - Filters on Tx and Rx side
- 2 Bidirectional amplifiers for time transfer channels
- DWDM filters

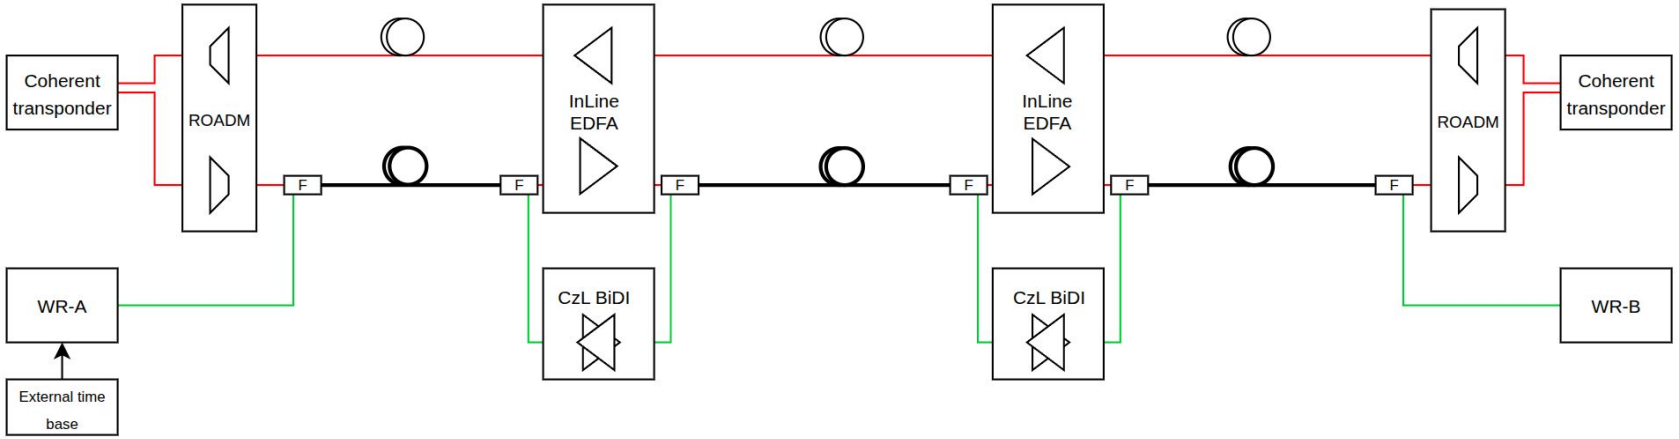

BiDi Channels over Shared Fiber

- Passive filters “everywhere”
  - L-band channels 6–9 for WR
  - C-band channels 39–46 for stable frequency (fixed frequency requirement)
  - Rest of the C-band for data channels

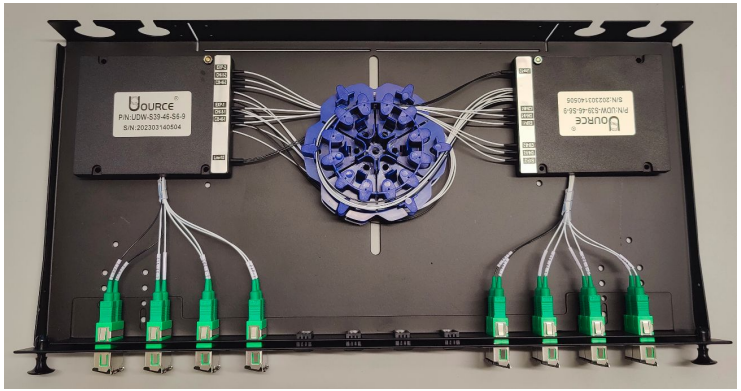

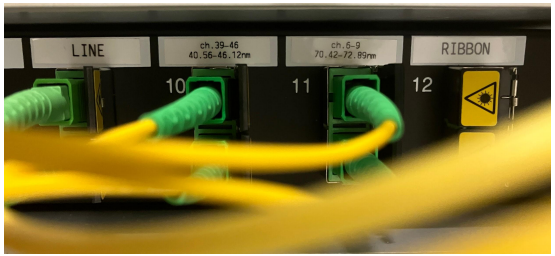

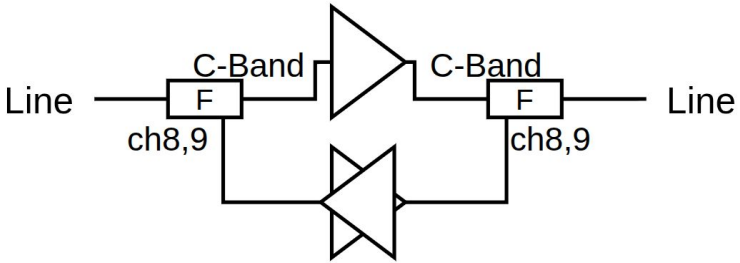

BiDi & Multiband is “Different”

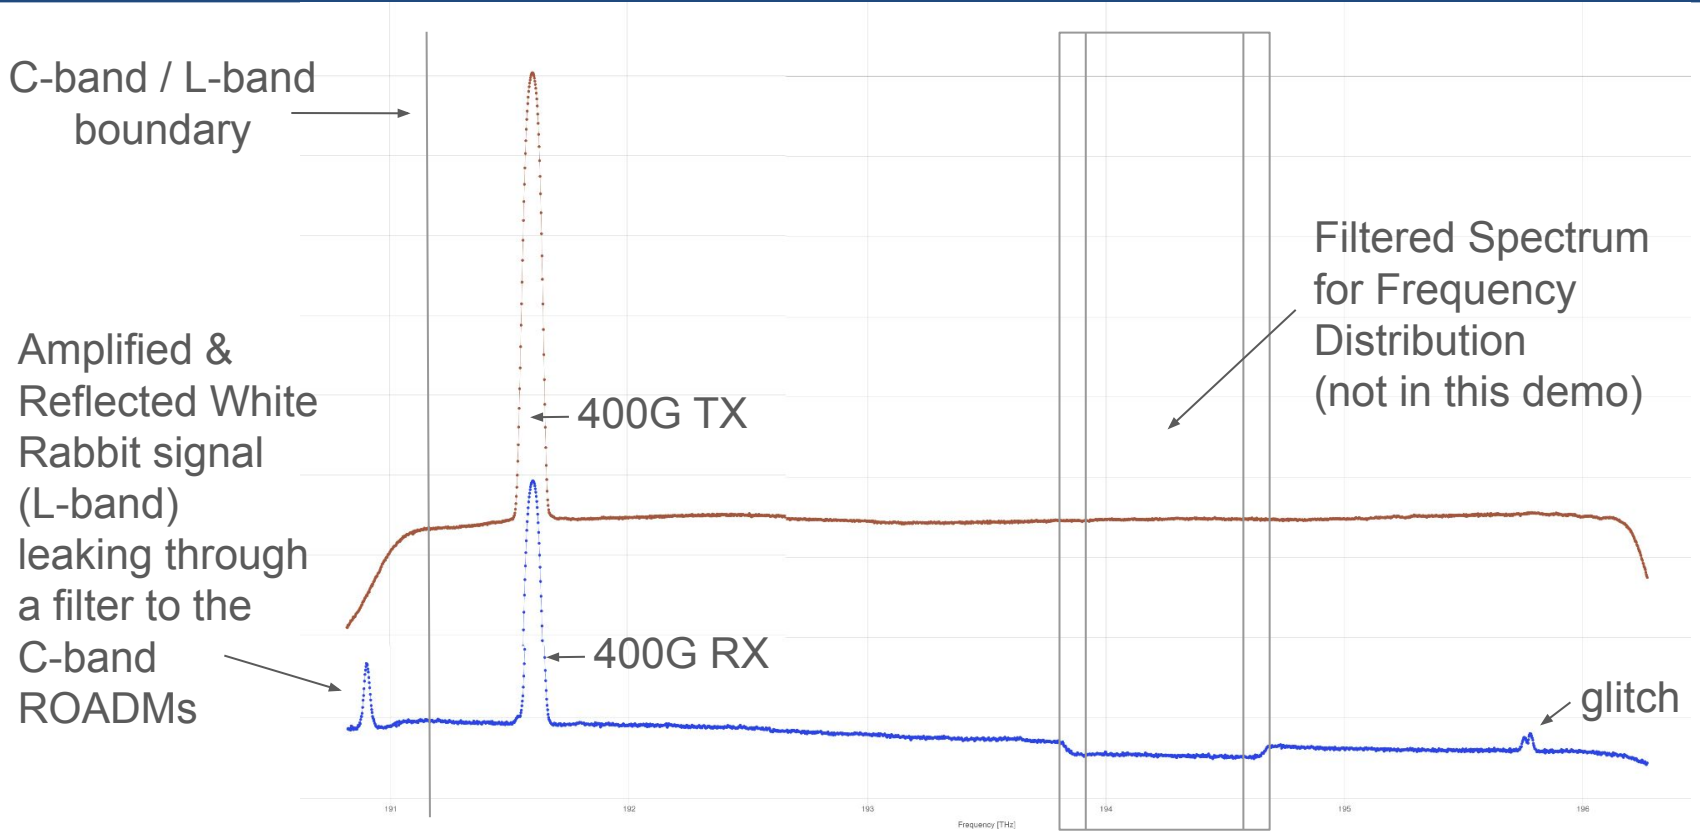

Challenges: Transceivers compatibility

- One manufacturer's mistake meets another manufacturer's mistake during testing of White Rabbit switches.
  - When the transceiver manufacturer saves one resistor and connects two inputs together (the transceiver works normally elsewhere),
  - and the White Rabbit CERN open hardware design has an error in the SFP transceiver connection – a 0 Ohm resistor directly to the +2.5V power supply,
- = a power short circuit in the WR device when the SFP transceiver is inserted.

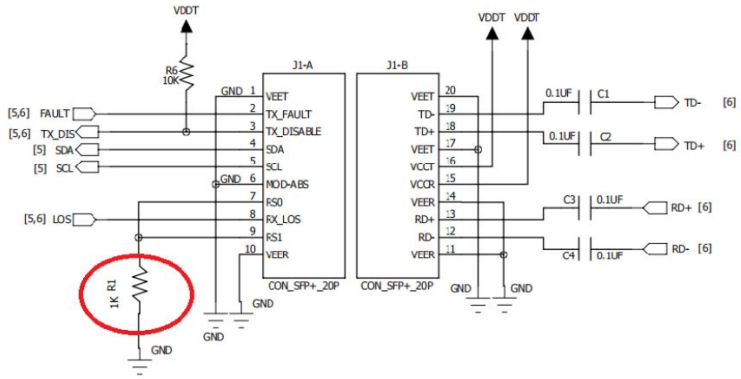

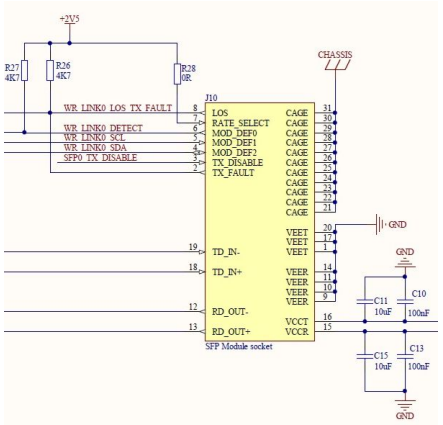

CzechLight SDN BiDi EDFA

- Bidirectional optical amplifiers for precise time distribution in CESNET3
- 47 devices delivered by PEI-GENESIS
- Key features
  - Multiband: C-band and L-band
  - Dual power supply AC 230V/DC 48V
  - Hot swap fans
  - Open Source OS - images released by CESNET→full control over the device (security)
  - NETCONF, RESTCONF, local CLI
  - Streaming telemetry

Source of truth in NETBOX

- Already in use in CESNET3 network
- All topology data
- IPAM
- Serve as the Inventory for ansible
- DWDM layer modeling has its limitations with Netbox

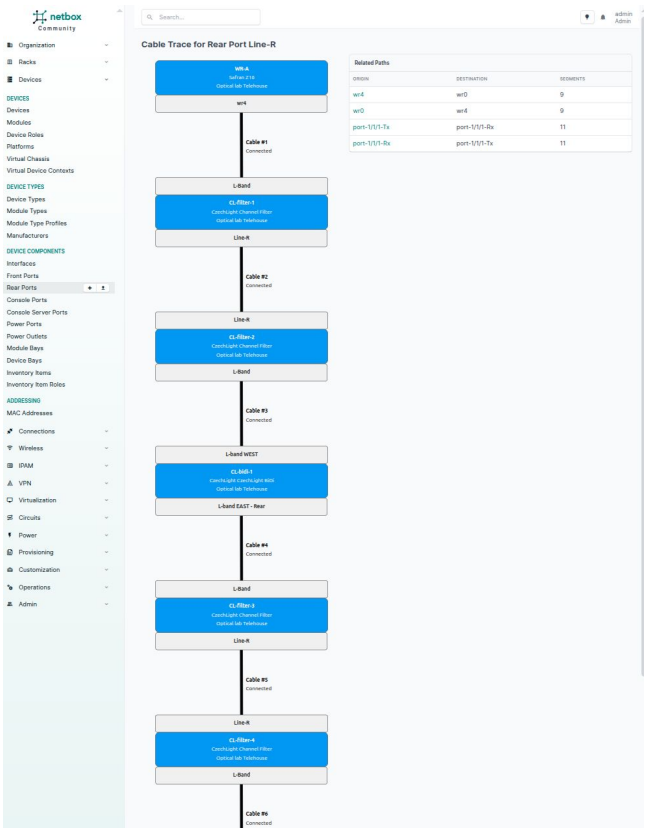

NETBOX Challenges: WDM, BiDi vs Duplex

- Some impedance mismatch
  - NETBOX started as a datacenter tool
  - WDM was not a native concept
  - “Link”: fiber pair vs. one fiber

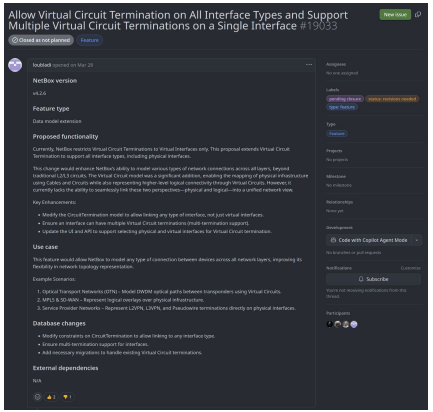

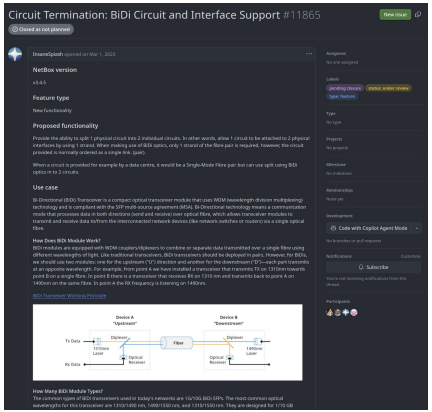

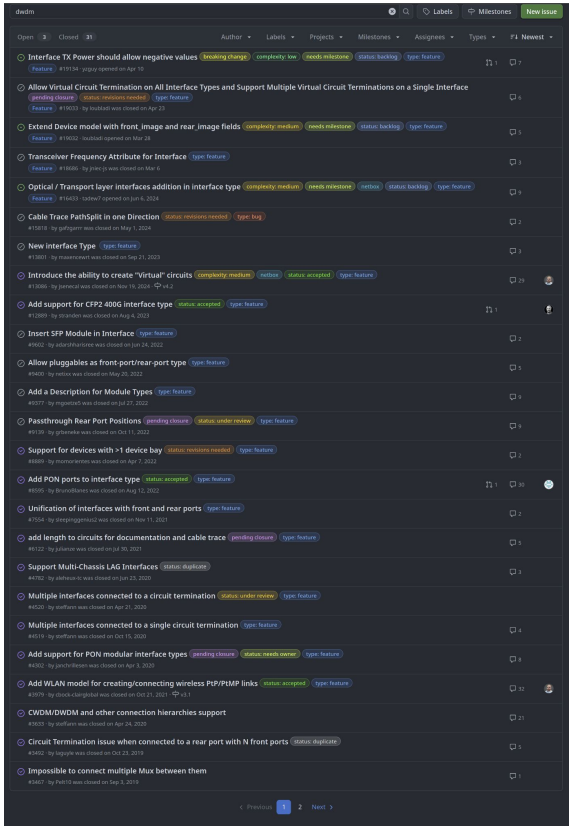

Configuration via Ansible

Ansible project structure for BiDi amplifiers

```
/
├── bidi-configure.yaml
├── inventory/
│   └── netbox-inventory.yml
├── host_vars/
│   ├── CzechLight-BiDi-A.yml
│   └── CzechLight-BiDi-B.yml
├── roles/
│   └── cl-bidi/
│       ├── tasks/
│       │   └── main.yml
│       ├── templates/
│       │   └── edfa_pump.j2
│       └── defaults/
│           └── main.yml
└── vault-password.txt
```

```
channels:
  narrow-1572:
    edfa_mode: "acc"
    edfa_setting: "84"
  c-band:
    edfa_mode: "off"
    edfa_setting: "0"
persistent: true
```

Configuration play

```
---
- name: "Configure CL BiDi"
  hosts:
  - "CzechLight-BiDi-A"
  - "CzechLight-BiDi-B"
  gather_facts: no
  tasks:
    - name: "Configure CL BiDi using restconf"
      include_role:
        name: "cl-bidi"

      vars:
        channel: "{{ item.key }}"
        edfa_mode: "{{ item.value.edfa_mode }}"
        edfa_setting: "{{ item.value.edfa_setting }}"
      loop: "{{ channels | dict2items }}"
```

Netbox names

ANSIBLE

git

ansible.netcommon

- Official module for configuring Network Elements
  - NETCONF & RESTCONF
- Sufficient for config distribution
  - Rather coarse features compared to native Linux modules
  - No model-driven knowledge

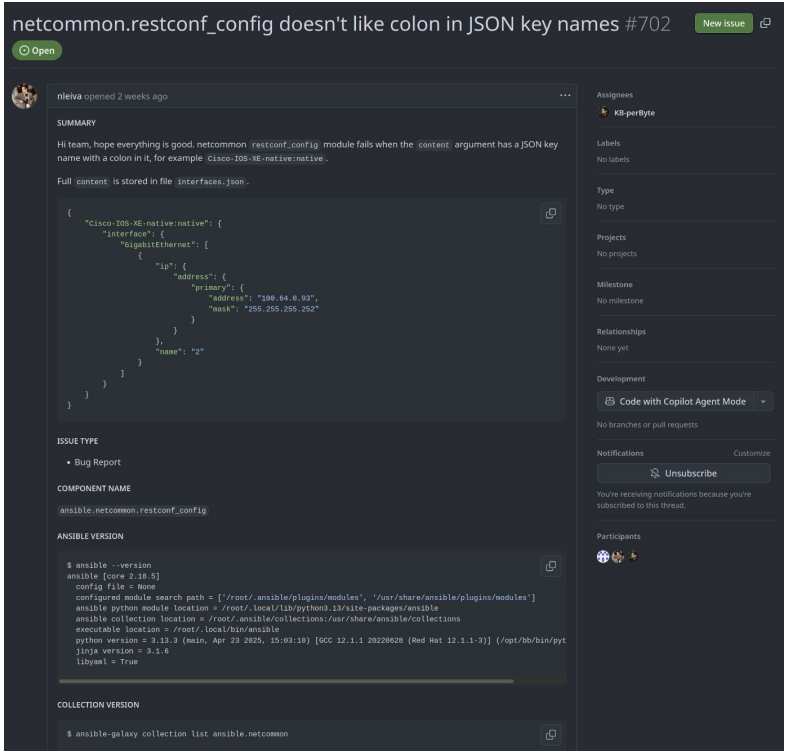

cesnet

Monitoring & Telemetry

- Streaming Telemetry
  - All the metrics
  - Sub-second latencies
- I/O Formats
  - IETF YANG-push
  - OpenMetrics (Prometheus)
  - Grafana

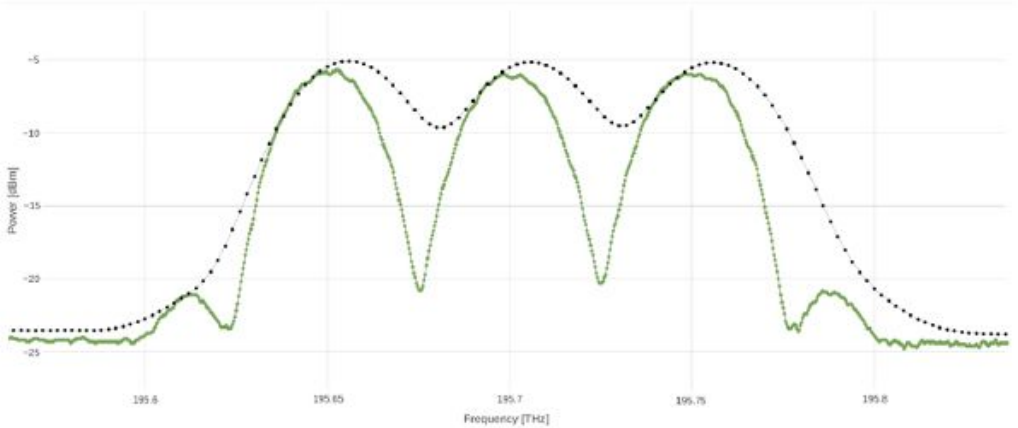

VICTORIA METRICS

Grafana

https://youtu.be/9vJex2p0JOA

Last 15 minutes | Refresh 5s

Time difference

Line view

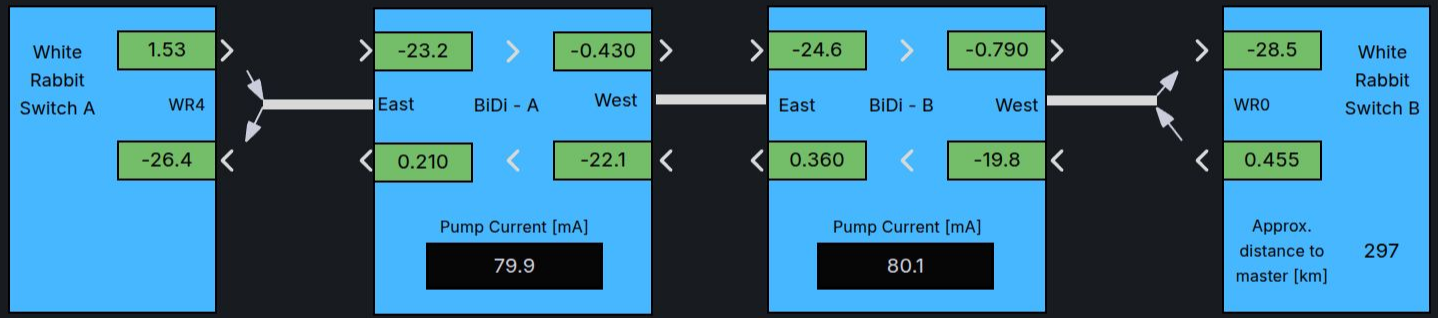

Service Quality - MDEV for last 6h

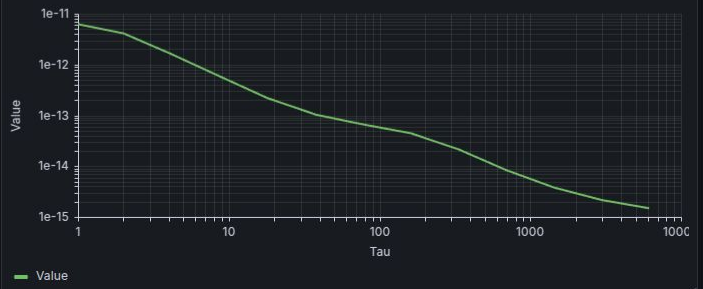

Time diference

Grafana

VICTORIA METRICS

cesnet

YANG

- Data model
  - Tree structure
  - Configuration vs. State
  - “Leaf” data types
- Defines Data Semantics
- Machine Validation

- Bjorklund, Martin. "YANG-a data modeling language for the network configuration protocol (NETCONF)." RFC 6020, 2010.
- Bjorklund, Martin. "The YANG 1.1 data modeling language." RFC 7950, 2016.


```
module: czechlight-bidi-amp
  +--rw c-band {dualband-c-plus-1572}?
  |  +--rw pump
  |  |  +--rw (control)?
  |  |  |  +--:(disabled)
  |  |  |  |  +--rw disabled?   empty
  |  |  |  +--:(agc)
  |  |  |  |  +--rw agc    gain-dB
  |  |  |  +--:(manual-current)
  |  |  |     +--rw manual-current    laser-current-mA
  |  |  +--ro measured-current?   laser-current-mA
  |  +--ro east-to-west
  |  |  +--ro input-power          odt:optical-power-dBm
  |  |  +--ro output-power         odt:optical-power-dBm
  |  |  +--ro output-spectrum (c-band-ocm}?
  |  |     +--ro lowest-frequency frequency-ghz
  |  |     +--ro step             frequency-ghz
  |  |     +--ro p?               anyxml
  |  +--ro west-to-east
  |  |  +--ro input-power          odt:optical-power-dBm
  |  |  +--ro output-power         odt:optical-power-dBm
  |  |  +--ro output-spectrum (c-band-ocm}?
  |  |     +--ro lowest-frequency frequency-ghz
  |  |     +--ro step             frequency-ghz
  |  |     +--ro p?               anyxml
  |  +--rw ocm-period?            cla-common:spectral-scan-frequency (c-band-ocm}?
  |
[...]
```

cesnet

- YANG
- NETCONF
- RESTCONF
- CLI

SW Architecture with Sysrepo

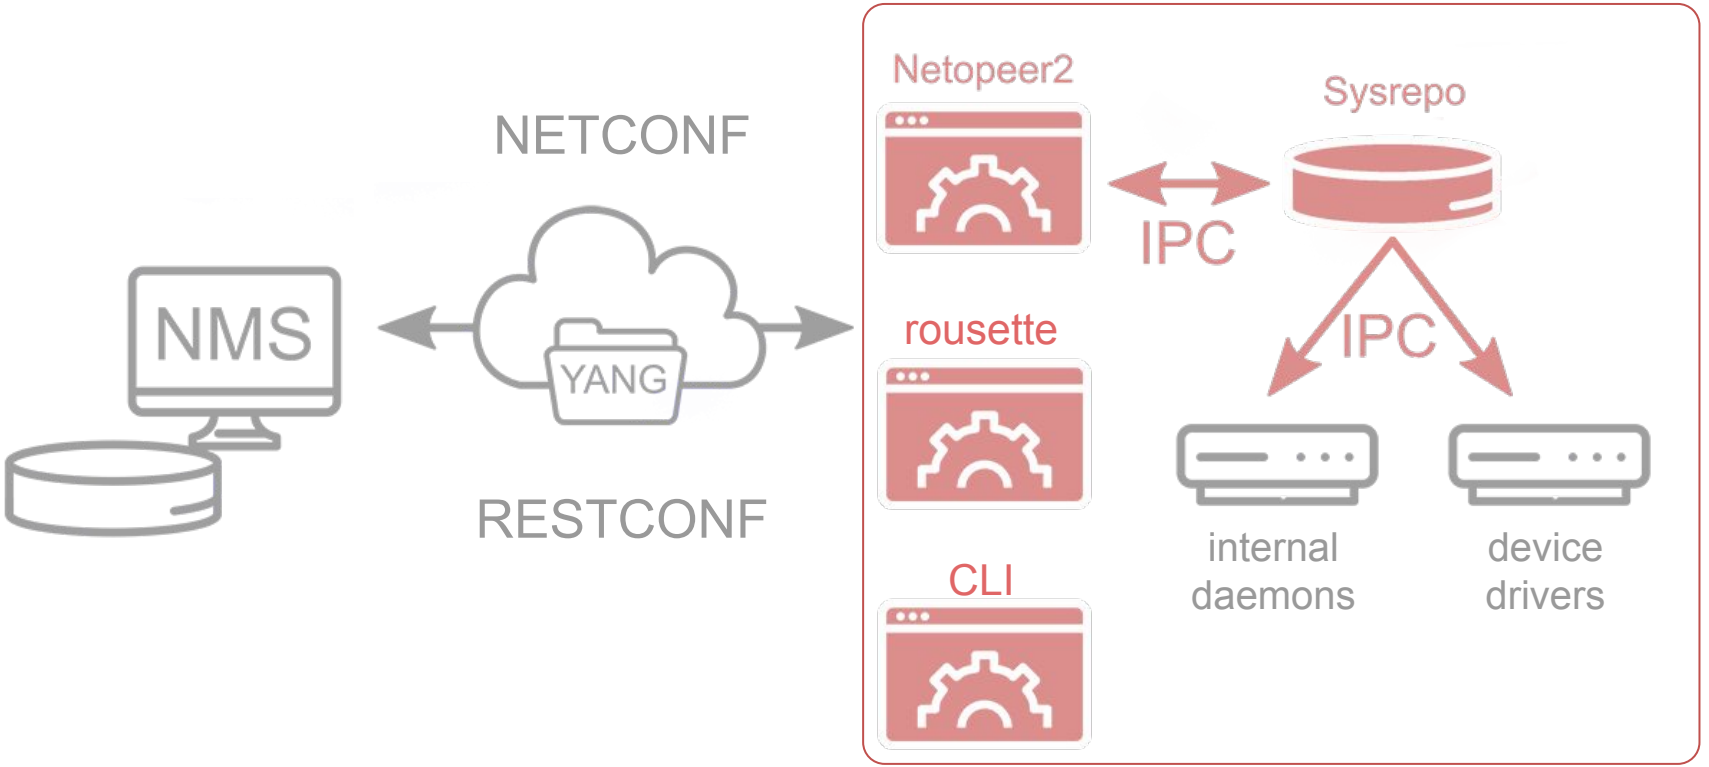

cesnet

YANG models

- czechlight-firewall
- czechlight-inline-amp
- czechlight-lldp
- czechlight-netconf-server
- czechlight-network
- czechlight-roadm-common
- czechlight-system
- iana-afn-safi
- iana-crypt-hash
- iana-hardware
- iana-if-type
- iana-ssh-encryption-algs
- iana-ssh-key-exchange-algs
- iana-ssh-mac-algs
- iana-ssh-public-key-algs
- iana-tls-cipher-suite-algs
- ietf-access-control-list
- ietf-alarms
- ietf-crypto-types
- ietf-datastores
- ietf-ethertypes
- ietf-factory-default
- ietf-hardware
- ietf-interfaces
- ietf-ip
- ietf-ipv4-unicast-routing
- ietf-ipv6-unicast-routing
- ietf-keystore
- ietf-netconf-acm
- ietf-netconf-monitoring
- ietf-netconf-nmda
- ietf-netconf-notifications
- ietf-netconf-server
- ietf-netconf-with-defaults
- ietf-netconf
- ietf-network-instance
- ietf-notification-capabilities
- ietf-origin
- ietf-packet-fields
- ietf-restconf-monitoring
- ietf-restconf-subscribed-notifications
- ietf-restconf
- ietf-routing
- ietf-ssh-common
- ietf-ssh-server
- ietf-subscribed-notifications
- ietf-system-capabilities
- ietf-system
- ietf-tcp-client
- ietf-tcp-common
- ietf-tcp-server
- ietf-tls-common
- ietf-tls-server
- ietf-truststore
- ietf-x509-cert-to-name
- ietf-yang-library
- ietf-yang-patch
- ietf-yang-push
- ietf-yang-schema-mount
- libnetconf2-netconf-server
- nc-notifications
- netopeer-notifications
- notifications
- opendevice-types
- sysrepo-factory-default
- sysrepo-ietf-alarms
- sysrepo-monitoring
- sysrepo-notifications
- sysrepo-plugind
- sysrepo
- velia-alarms
- yang

cesnet

Operating System

- Read-only Linux `rootfs` via **`buildroot`**
  - Stateless system
  - R/W config with YANG data
- Userland based on **`systemd`**
- A/B software slots via **`rauc`**
  - Integrated with HW watchdog
  - Atomic system updates

Open Source

[● ◀] systemd

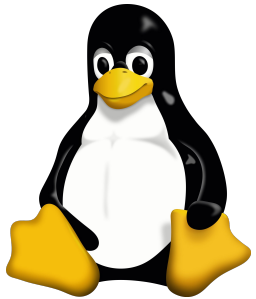

BuildRoot
Making Embedded Linux Easy

RAUC

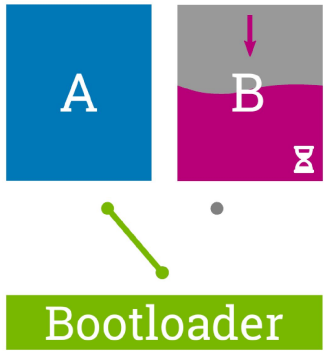

Image source: © Enrico Jörns, Pengutronix.

Source Code + CI/CD

- Open Source SW
  - One closed part due to NDAs
- github.com/CESNET
- CI/CD: gerrit.cesnet.cz

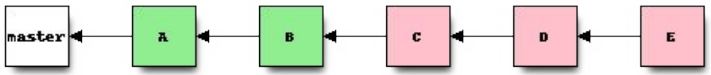

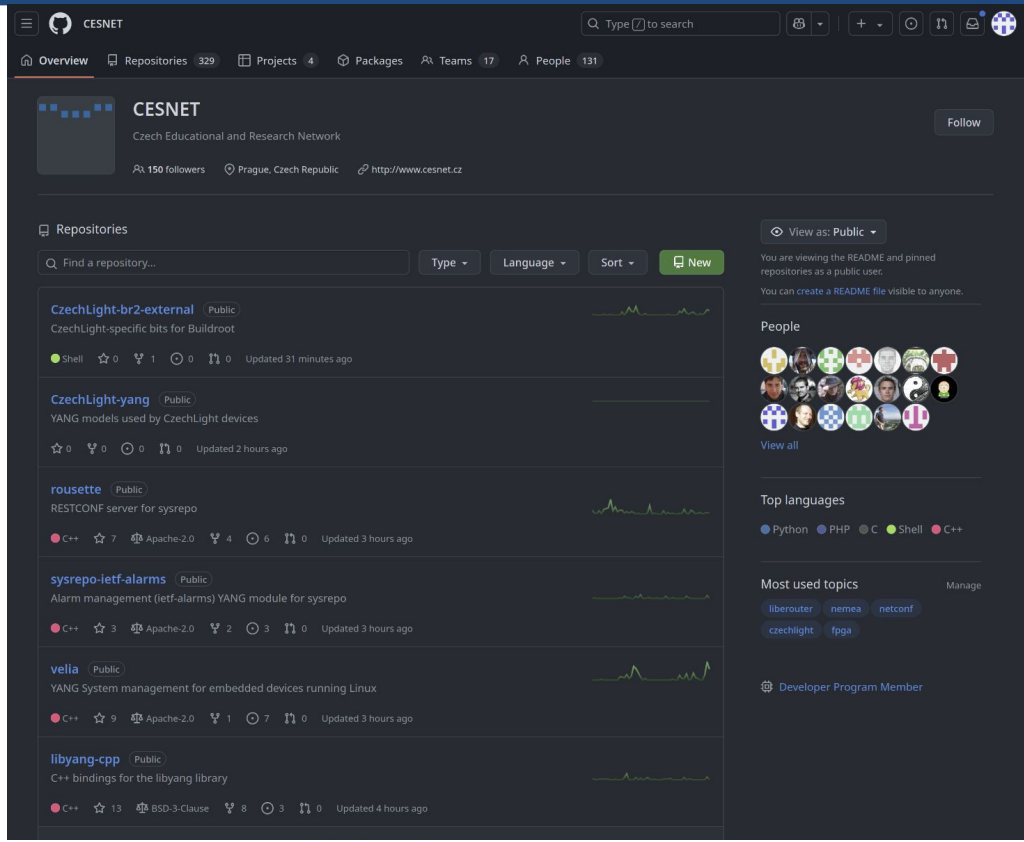

cesnet

Open Hardware Platform

Hotswap
AC & DC PSUs

Redundant fans

HW health
telemetry & alarms

cesnet

Czech Light ™ SDN OLS

Scalable,
Modular,
Open

ROADMs,
EDFAs,
BiDi EDFAs

SDN Control
& Monitoring

cesnet

Commercial Availability

IntelDat

OPTOKON®
FIBER OPTIC TECHNOLOGY
CALIBRATION LABORATORY

CZECHOS
Czech Optical Solutions

PEI-Genesis™

Photonic
Technologies

cesnet

TL;DR

- CESNET's upcoming production service
  - Bringing precise time to existing users
- Reusing Existing Fiber Footprint
  - Multiband & BiDi is possible
  - …but some aspects are "tricky"
- DevOps Approach
  - Treating Network Elements as "servers with an API"
  - High-quality telemetry as a key for successful deployment
- Open Hardware
  - With open source software and commercial support
  - Not just the BiDi amplifiers

cesnet

Come see the demonstration
at our stand

cesnet

https://ces.net/tnc25

cesnet

Q&A

michal.hazlinsky@cesnet.cz

cesnet

Slides for Offline Reading

Czech Light ™ SDN ROADM

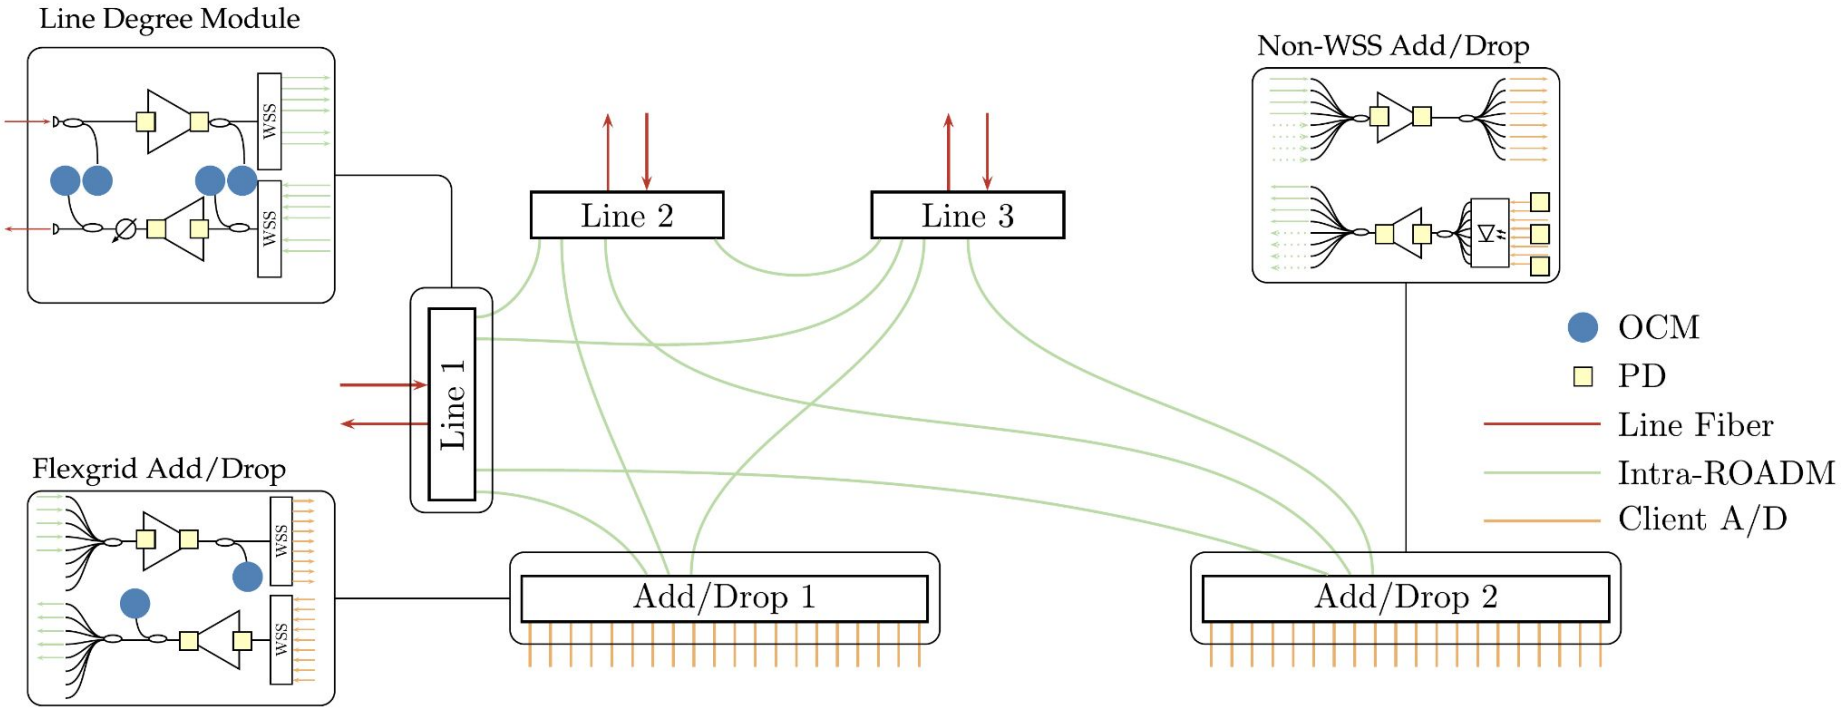

Line Degree

- Long-haul connection
  - Internal amplification, up to 25 dB spans (VOA)
  - 0 dBm/ch typ. TX power
  - Integrated OSC
  - OTDR option via SFP
- Flexgrid
  - 12.5 GHz granularity
  - min. 50 GHz MCs
- 9 Express ports
  - -12 dBm/ch TX
  - -15..+5 dBm/ch RX

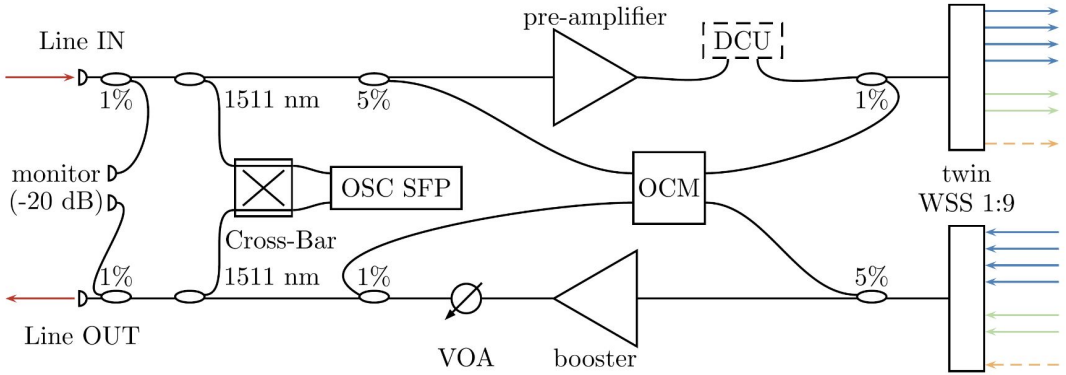

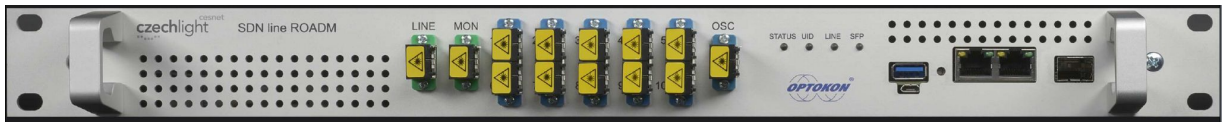

Add/Drop Options

- Passive
  - Direct-connection to Express ports
  - 2-deg only
  - 8 client ports
  - Spectrum control via Line Degrees
  - -12 dBm/ch

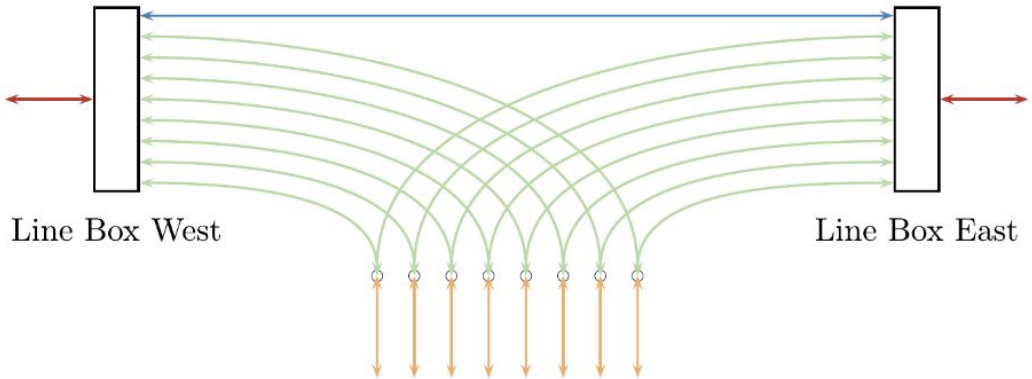

Add/Drop Options

- Coherent Signals
  - 8ch/1U
  - Per-port power monitoring
  - Power equalization via Line Degree ROADM

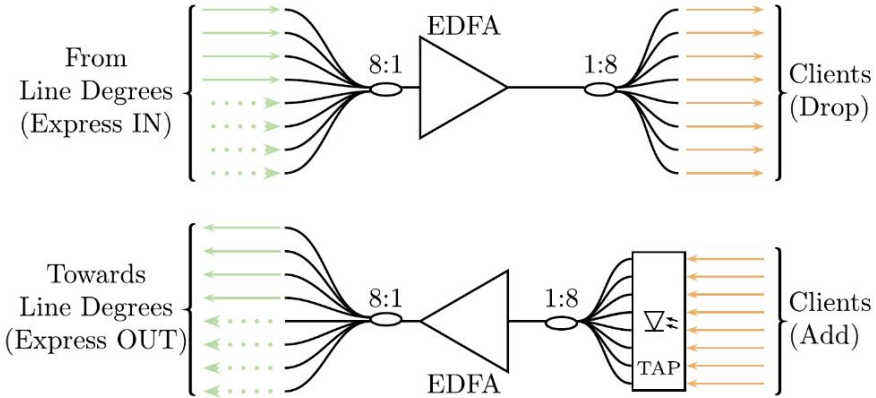

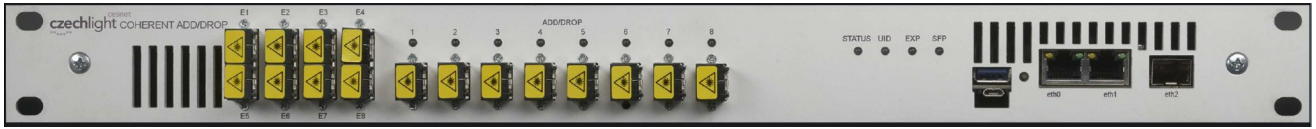

Add/Drop Options

- WSS-based
  - 20ch/1U
  - 12.5 GHz routing (min. 37.5 GHz per MC)
  - 12.5/6.25 GHz monitoring
  - Indirect client-port monitoring

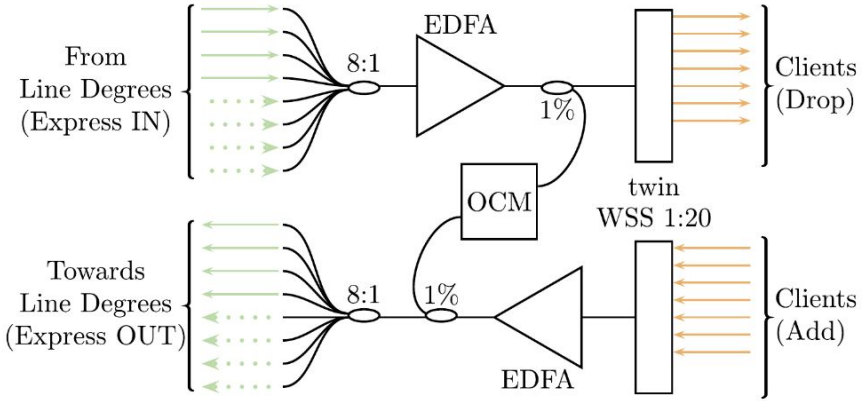

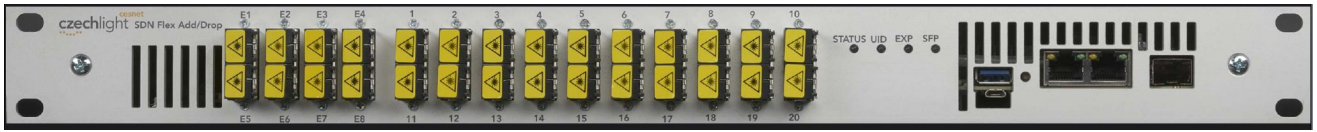

Add/Drop Options

- Alien Wavelength
  - 20ch/1U
  - 3.25 GHz routing
  - 0.3125 GHz monitoring (2.5 pm)
  - Direct client-port spectrum monitoring

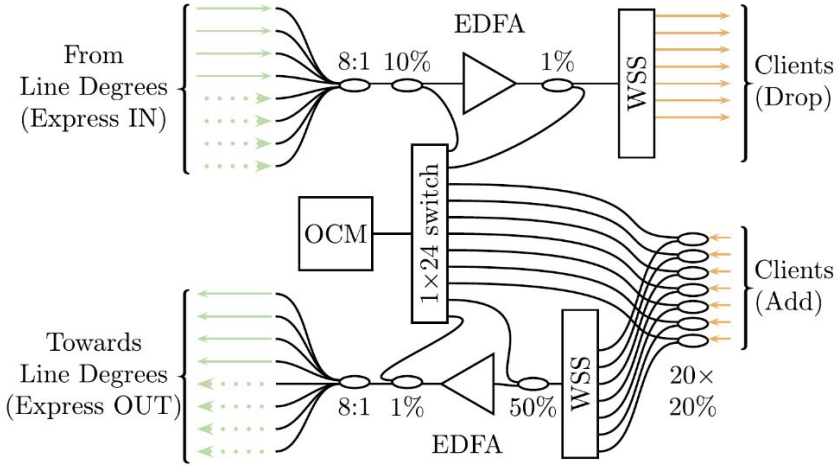

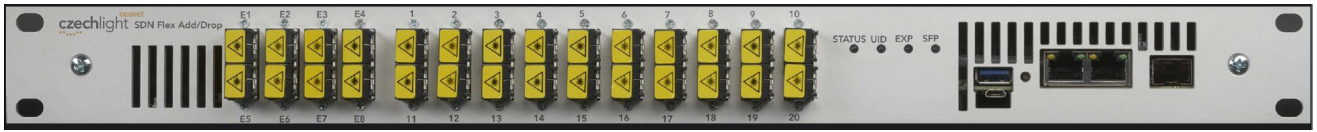

In-Line Amplifiers

- Dual EDFA
  - Flat gain, 27 dB
  - NF ≤ 5.5
  - No tilt control
  - Up to 25 dB spans,
    output VOA in both directions
  - Class-1M product

- OTS
  - Power monitoring
  - Integrated OSC
  - OTDR option via SFP

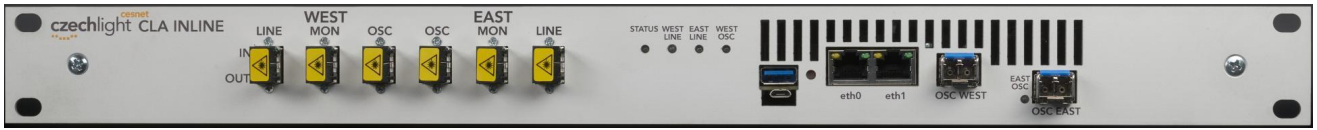

BiDi Amplifiers

- EDFA
  - Dual-band by default
    - C-band: 1540 – 1547 nm
    - L-band: 1572 nm
    - Other options possible on request
  - Optional built-in OCM
  - Auto Gain Control, 23 dB
  - Auto Current Control, Manual pump setting