



Mellanox Linux Switch

Taking Open Ethernet to the next level, Mellanox Linux Switch provides users with the ability to deploy any standard Linux operating system on their switch devices

Mellanox Linux Switch enables users to natively install and use any standard Linux distribution as the switch operating system on the Open Ethernet Mellanox Spectrum® switch platforms. Mellanox Linux Switch is based on Switchdev, a Linux kernel driver model for Ethernet switches. Mellanox is the first switch vendor to embrace this model by implementing the switchdev driver for its Mellanox Spectrum switches. This revolutionary concept breaks the dependency of using vendor-specific software development kits (SDK). Instead of proprietary APIs, fully standard Linux kernel interfaces are used to control the switch chip. This allows users to natively install and use any standard Linux distribution as the switch operating system, while the switchdev driver ensures all network traffic is seamlessly offloaded to the switch hardware.

Openness and Freedom

The 100% free open-source Mellanox switchdev driver is developed and maintained as part of the Linux kernel and is promoted by the Linux community. There is no longer a need for closed code from switch vendors, no more binary blobs, and no need to sign a service level agreement (SLA).

Linux Switch provides users the flexibility to customize and optimize the switch to their exact needs without paying for expensive proprietary software that includes features they neither need nor will never use. It also reduces software maintenance expenses, resulting in lower TCO and improved ROI.

Uniformity

Installing network switches with a standard Linux distribution turns them into yet another server in the data center. This greatly reduces management efforts, as the same configuration and monitoring tools can be used for both servers and switches and be treated under the same support contract with your Linux distribution.

Abstraction

Mellanox's Linux Switch uses the Switchdev driver as an abstraction layer which provides open, standard Linux interfaces and ensures that any Linux application can run on top of it. These can be Open Ethernet protocol stacks, management tools, user-developed applications, and more.

This model also completely separates the switch hardware from the switch software. The same software can run on different hardware, such that the hardware can be replaced and upgraded without changing the software that runs on it. It also ensures full compatibility between all switches in the data center.

HIGHLIGHTS

- 100% free open source software
- L2/L3 switch system as standard Linux server
- Uniform Linux interfaces
- Fully supported by the Linux community
- Hardware independent abstraction layer
- User choice of Linux operating systems and network applications
- Industry's highest performance switch systems portfolio, including the high-density half-width Mellanox SN2100



L2 & L3 Support

Once the Mellanox Switchdev driver is loaded into the Linux Kernel, each of the switch's physical ports is registered as a net_device within the kernel. Using standard Linux tools (for example, bridge, tc, iproute), ports can be bridged, bonded, tunneled, divided into VLANs, configured for L3 routing and more. Linux switching and routing tables are reflected in the switch hardware. Network traffic is then handled directly by the switch.

Standard Linux networking applications can be natively deployed and run on switchdev. This may include open source routing protocol stacks, such as Quagga, Bird and XORP, OpenFlow applications, or user-specific implementations.

Chassis Monitoring & Management

The Mellanox switchdev driver also implements a Linux hardware monitoring (hwmon) driver. Chassis management can be performed by directly accessing 'sysfs' files or by using the lm_sensors free open source Linux tool for monitoring temperatures, voltage and fans.

Mellanox Spectrum Switches

Native Linux OS is supported on all Mellanox Spectrum SN2000 switch systems, as well as the Mellanox Spectrum switch ASIC. Mellanox Spectrum introduces the world's lowest latency for a 100GbE switching and routing element, and does so while having the lowest power consumption on the market. The SN2000 portfolio offers various combinations of number of ports and ports speeds, including 25/50GbE. For more information about Mellanox switches, please visit www.mellanox.com.

Installation

Mellanox SN2000 and SN3000 switches are shipped pre-installed with ONIE (Open Network Install Environment). Once the system boots, ONIE will connect to the management network, download the operating system from a pointed-to location, and then install it. This can be any standard Linux distribution (for example, RedHat, Ubuntu, Fedora, and so on). The following boot will load the newly installed operating system. If the switchdev driver is not in box in the operating system, a short procedure to update the kernel should be performed.

Performance

Mellanox Spectrum Linux Switch uses the Mellanox switchdev driver, which is natively implemented in the Linux kernel and interfaces directly with the switch silicon. This enables genuine maximal hardware performance when running network applications on top of it.

Features

Layer 2 Switching

- 10/25/40/50/56/100GbE
- Ports Bridging/Bonding
- VLANs
- STP
- LLDP
- FDB learning/aging
- IGMP
- SPAN

Layer 3 Routing

- Unicast IPv4/IPv6 router
- ECMP IPv4/IPv6
- Weighted ECMP
- VRRP
- Multicast IPv4/IPv6 router
- Linux protocol stack support (Quagga, BIRD, etc.)

Policy-based Routing

- ACLs using TC Flower
- Supports Linux Open Flow applications (OVS, etc.)

Virtualization

- GRE
- VXLAN

System Management

- Fans and temperature
- LEDs
- Power supplies
- Temperature gauge access
- Port Split
- Cable / Module details

QoS

- Buffer configuration
- Congestion control (RED,ECN)

Monitoring & Synchronization

- RSPAN/ERSPAN
- PTP 1588

*Standard Linux application which natively interfaces switchdev to configure the switch hardware. ** Readmap item.

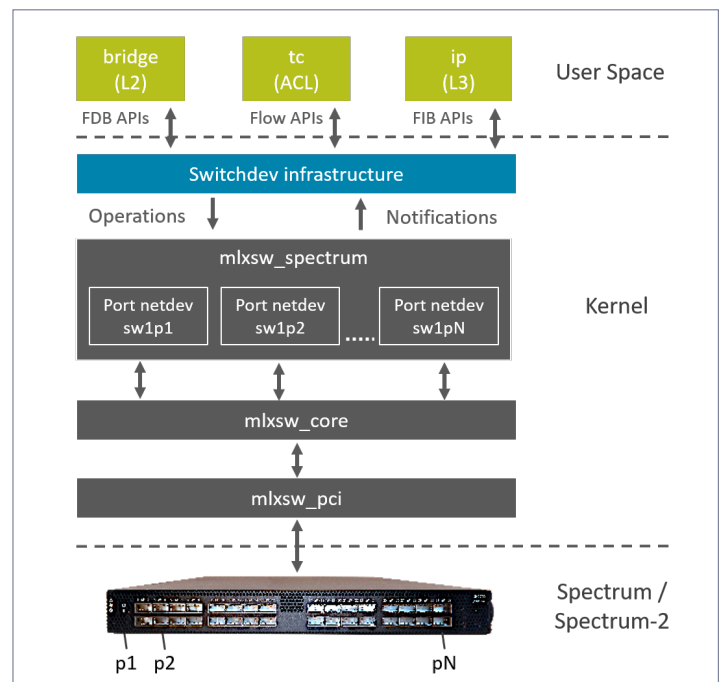


Figure 1. Linux Switch Architecture

Table 1 - Part Numbers and Descriptions

OPN	Description
SN2700	Spectrum-based 25GbE/100GbE 1U Open Ethernet switch with ONIE, 32 QSFP28 ports
SN2410	Spectrum-based 25GbE/100GbE 1U Open Ethernet switch with ONIE, 48 SFP28 and 4 QSFP28 ports
SN2100	Spectrum-based Half-width 25GbE/100GbE 1U Open Ethernet switch with ONIE, 16 QSFP28 ports
SN2010	Spectrum-based 25GbE/100GbE 1U Open Ethernet switch with ONIE, 18 SFP28 and 4 QSFP28 ports
SN3700	Spectrum-2 based 100GbE 1U Open Ethernet switch with ONIE, 32 QSFP28 ports
SN3700C	Spectrum-2 based 200GbE 1U Open Ethernet switch with ONIE, 32 QSFP56 ports
SN3800	Spectrum-2 based 100GbE 2U Open Ethernet switch with ONIE, 64 QSFP28 ports