

Changes and New Features History

Upstream Kernel 5.18				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
2883524	Multiple VF/SF Interfaces in the Shared Block in TC Offload	[ConnectX-5 and above] Added support for creation of TC flows on shared block of VF representors.	N/A	N/A
2429749	sFlow Sampling Rules Offload	[ConnectX-5 and above] Added support for sFlow sampling rules offloads. sFlow is an industry standard technology for monitoring high speed switched networks. Open vSwitch integrated sFlow to extend the visibility into virtual servers, ensuring data center visibility and control.	Not supported	xx.29.1016
2204161	VLAN Push on RX and Pop on TX	[ConnectX-6 Dx and above] Added support in SwitchDev mode to push VLAN on RX (rules on uplink representor) and pop on TX (rules on VF representor). This feature is only supported in Software Steering mode.	N/A	N/A
Core Features				
2626743	Control SF Class	[All HCAs] Added support for Control SF Class. Each PCI, PF, VF, SF function, by default, has netdevice, RDMA, and vdpa-net devices always enabled. This feature enables the user to control which device functionality to enable/disable.	N/A	N/A
RDMA Features				
2634958	mlx5 VFIO CQ Interrupt Mode	[ConnectX-5 and above] Added support for mlx5 VFIO CQ interrupt mode. This feature allows applications to listen on and capture CQ completion events via the Event Queue mechanism with mlx5 VFIO driver.	rdma-core v41	N/A
2600493	mlx5 VFIO Asynchronous Event	[ConnectX-5 and above] Added support for mlx5 VFIO asynchronous event. This feature supports applications to listen on and capture device asynchronous events via the Event Queue mechanism with mlx5 VFIO driver.	rdma-core v41	N/A
Software Steering Features				
2895043	Matching Granularity Change	[ConnectX-5 and above] Added support for matching granularity change. Currently when creating FDB flow with destination of VPORT, a src_port matching must be added. This feature removes that limitation so a FDB flow can match all VPorts and go to a VPORT destination. The new behavior is the same as that on firmware steering.	rdma-core	N/A
2855804	Full Tunnel Header Matching	[ConnectX-6 Dx and above] Added support for full tunnel header matching. Currently, the full-tunnel-header matching cannot be used along with many other criteria within one matcher. This feature is to add the support for that by using the new definer index, defined in the firmware, to build matcher so that the full tunnel header matching can be used along with all other criteria.	rdma-core	xx.32.1010

Upstream Kernel 5.17				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
2668112	Setting VF Group Rate Limit	[ConnectX-5 and above] Added support for setting VF group rate limit using Devlink command.	iproute2 (devlink) 5.16	N/A
Core Features				
2658197	Steering of IPoIB Packets via Ethtool	[All HCAs] Enabled steering of IPoIB packets via Ethtool, in the same way it is done today for Ethernet packets.	N/A	N/A
2633768	Memory Improvements	[ConnectX-4 and above] Added support for providing knobs which enable users to minimize memory consumption of mlx5 functions (PF/VF/SF).	N/A	xx.32.1010

2339680	512 vDPA Interfaces	[ConnectX-6 Dx and BlueField-2] Added support in multiple flow steering priorities for FDB rules.	N/A	N/A
NetDev Features				
2801496	Expose FEC Counters via Standard Ethtool Interface	[ConnectX-4 and above] Exposed Ethtool counters which report FEC (forward error correction) statistics via the 'ethtool -l --show-fec ethX' command. <ul style="list-style-type: none"> fc_fec_corrected_blocks_laneX rs_fec_corrected_blocks fc_fec_uncorrectable_blocks_laneX rs_fec_uncorrectable_blocks phy_corrected_bitz 	ethtool v5.13	N/A
2800578	Ethtool CQE Mode Control	[ConnectX-4 and above] Replaced the vendor-specific Ethtool API (priv-flag) with a standard Ethtool API (replaced 'ethtool --set-priv-flags ethX rx_cqe_moder on/off tx_cqe_moder on/off' with 'ethtool -C ethX cqe-mode-rx on/off cqe-mode-tx on/off'). This decreases the amount of vendor-specific configurations and aligns mlx5 driver with the upstream Ethtool API.	ethtool v5.15	N/A
2482544	Reduce Memory Consumption	[ConnectX-4 and above] Reduced memory consumption. Previously, several hardware/memory resources were pre-allocated per NetDev to serve a prefixed maximum allowed value. Now, the actual maximum is used instead of the prefixed theoretical maximum and resources are allocated on demand, improving memory consumption.	N/A	N/A
Software Steering Features				
2338344	Dumping Software Steering Information	[ConnectX-5 and above] Extended mlx5 debugfs support to present Software Steering resources: dr_domain including its tables, matchers, and rules. The interface is read-only. While dump is being presented, new steering rules cannot be inserted /deleted. This allows to see the existing steering configuration and to be able to debug steering issues.	N/A	N/A

Upstream Kernel 5.16				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
2754419	Offload Forwarding to Multiple Destinations	[ConnectX-5 and above] Added support for offloading packet replication to up to 32 destination through the use of TC rule.	N/A	N/A
2666887	Bridge Spoof Check	[All HCAs] Added support for spoof check with TC flower rules on representors attached to bridge to mirror spoof check SR-IOV functionality.	N/A	N/A
2293534	Tunnel Rules Offload	[ConnectX-6 Dx and above] Added support for offloading tunnel rules when the source interface is OpenvSwitch bridge (internal port).	N/A	N/A
Core Features				
2713656	Single IRQ for PCI Function	[ConnectX-4 and above] Added support for single IRQ for PCI function. To use a high number of VFs, a large amount of IRQs is required which the device cannot always support. This feature enables VFs to function with a minimum of a single IRQ instead of two. This is done via dynamic MSIX feature. In the case that dynamic MSIX feature is not supported (old kernels), the following configuration will probe all VFs with single IRQ: \$ mlxconfig -d s NUM_VF_MSIX=0 STRICT_VF_MSIX_NUM=1.	N/A	xx.31.1014
2655918	Steering-Based LAG	[ConnectX-6 Dx and BlueField-2] Added support to have physical port selection based on the hash function defined by the bond so that different packets of the same flow will be egress from the same physical port.	N/A	xx.31.2006
NetDev Features				
2816603	Firmware Health Asserts Enhancement	[ConnectX-4 and above] Health buffer now contains more debug information like the epoch time in sec of the error and the error's severity. The print to dmesg is done with the debug level corresponding to the error's severity. This allows the user to use dmesg attribute: dmesg --level to focus on different severity levels of firmware errors.	N/A	xx.32.1010
2300278	SHAMPO (GRO Offload)	[HCA?] Added support for SHAMPO (Split Header And Merge Payload Offload), a new functionality that enables stitching of packets data to unified skb representation in the network stack to reduce per-skb overhead.	N/A	N/A

RDMA Features				
2438531	Lossy RoCE	[ConnectX-5 and above] Extended the RDMA statistic tool to support setting vendor-specific optional counters dynamically using netlink. Added to mlx5_ib the following optional counters: cc_rx_ce_pkts, cc_rx_cnp_pkts, cc_tx_cnp_pkts. The cc_rx_cnp_pkts, cc_tx_cnp_pkts Example: \$ rdma statistic mode supported link rocep8s0f0/1 link rocep8s0f0/1 supported optional-counters cc_rx_ce_pkts, cc_rx_cnp_pkts, cc_tx_cnp_pkts \$ sudo rdma statistic set link rocep8s0f0/1 optional-counters cc_rx_ce_pkts, cc_rx_cnp_pkts \$ rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts, cc_rx_cnp_pkts \$ sudo rdma statistic set link rocep8s0f0/1 optional-counters cc_rx_ce_pkts \$ rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts \$ sudo rdma statistic unset link rocep8s0f0/1 optional-counters Note: Need ECN/PFC counter statistics based on VF	iprout2 v5.15.0	xx.31.1014
Steering UserSpace Features				
2680924	Flow Direct to Port	[ConnectX-6 Dx] When using Socket-Direct NICs, each interface may have different PCIe IDs and even different domains where the NIC resides on multiple NUMA nodes. The current logic is that interfaces which have the same "system_image_guid" reside in the same physical NICs. "native_port_num" represents interface index. When using firmware older than xx.32.1xxx, the original PCIe BDF approach is used and can be queried at sysfs "phys_switch_id" and "phys_port_name".	mlx_dpdk_20.11	xx.32.1010
2441633	MLX5DR SF	[ConnectX-5 and above] Added support for up to 512 SFs with the mlx5dv_dr API.	rdma-core v38	N/A
Upstream Kernel 5.15				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP ² Features				
2250409	Kernel VDPA Network Devices	[ConnectX-6 Dx] Added support for virtio network drivers utilizing VDPA interfaces. This requires recent libvirt and qemu.	qemu and libvirt >= 6.0	N/A
2193092	VF Group Throttling	[ConnectX-5 and above] Added support to allow user to apply QoS configuration of single function (VF, SF) or group of functions using iproute2 devlink command.	iproute2 v5.13.0 korg kernel v5.15-rc1	N/A
NetDev Features				
2192365	Application Device Queues (ADQ)	[ConnectX-4 and above] Added driver-level support for Application Device Queues. It allows partition defining over the RX/TX queues into groups and isolates traffic of different applications. This mainly improves predictability and tail latency.	N/A	N/A
RDMA Features				
2565548	DV API for DMA GGA memcpy	[BlueField-2 and above] DMA memcpy is one of several Memory-to-Memory Offloads (MMO) available from BlueField-2 onwards. It utilizes the GGA modules on the DPU to perform DMA memcpy, thus improving performance. The memcpy can be done locally, on the same host, or between the host and the Arm. To use this feature, expose DV API.	rdma-core only	N/A
2452045	mlx5 over VFIO	[ConnectX-4 and above] Added support for mlx5 user space driver over VFIO. This feature enables an application to take full ownership on the opened device and run any firmware command (e.g., port up/down) safely. The application look and feel is like regular RDMA application over DEVX. It uses verbs API to open /close a device and then mostly uses DEVX APIs to interact with the device. New mlx5 DV APIs were added to get ibv_device for a given mlx5 PCI name and to manage device-specific events. For description of the relevant APIs and expected usage of those APIs, look up the following: mlx5dv_get_vfio_device_list() mlx5dv_vfio_get_events_fd() mlx5dv_vfio_process_events()	rdma-core v37	N/A
2297914	DV API for AES-XTS	[ConnectX-6 and above] Added DV API that allows configuration of MKey with AES-XTS crypto offloads. The MKey can be configured for both crypto and signature offloads.	rdma-core only	xx.31.1014
1711066	Single FDB	[All HCAs] Added support for bonding the uplink representors under the same bond device so that a single InfiniBand device is created which can be used to manage VPorts which are on both E-Switch managers.	N/A	N/A
Steering UserSpace Features				
2669516	Extended Flex Parser Ability	[ConnectX-5 and ConnectX-6 Dx] Added userspace ability to expose flex parsers 4-7 provided by misc4 to extend matching ability of flex parsers. Now all flex parsers can be matched at the same time.	rdma-core v38	xx.32.1xxx

2669419	Setting DR Matcher Layout	[ConnectX-6 Dx] Added support for a new RDMA CORE DR API to set the DR matcher layout by calling <code>mlx5dv_dr_matcher_set_layout</code> . Setting the matcher layout allows presetting the matcher size and increasing matcher rule capacity as well as other performance improvements in case matcher size is known.	rdma-core v37	N/A
2616897	Software Steering Stability	[ConnectX-6 Dx] Improved Software Steering rule creation stability so that all rule's insertion occur in a defined time using defined (export) size of <code>Htbl</code> and decreased use of dynamic allocation.	rdma-core v37	N/A
2436999	Software Steering Rule Optimization	[ConnectX-6 Dx] Added Software Steering Rule Optimization which improved rate of updating steering rules, insertion, and deletion. The feature includes definers, multi-qp approach, and better memory usage.	rdma-core v37	N/A
2350115	Software Steering Insertion Rate Optimizations	[ConnectX-6 Dx] Added support for better insertion rate in software steering. This includes multi-QP which skips areas in the code that may be for debug only.	rdma-core v37	N/A
2248827	mlx5dv_dr Support of Scalable Function	[All HCAs] Added <code>mlx5dv_dr</code> SF support. This feature allows directing traffic using <code>mlx5dv_dr_dest_ib_port</code> action to a SF VPort.	rdma core v37	N/A
2239853	mlx5dv_dr Support for VF LAG	[All HCAs] Feature Description: This feature allows directing traffic using <code>mlx5dv_dr_dest_ib_port</code> action to any port on the e-switch, including the other PF.	rdma core v37	N/A

Upstream Kernel 5.14

Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
2198071	Enlarge Switchdev Tables	[ConnectX-5 and above] Added support for allowing OVS kernel to support up to 128 matches (groups) per table and 16M entries per group.	N/A	N/A
2047159	Bridge Offloads with VLAN Support	[ConnectX-6 Dx and above] Added support for bridge offloads with VLAN support that works on top of <code>mlx5</code> representors in switchdev mode.	N/A	N/A
ASAP² and Steering Features				
2200978	sFlow Sampling Rules Offload	[ConnectX-5 and above] Added support for offloading sFlow sampling rules.	N/A	N/A
RDMA Features				
2431025	Raw WQE	[ConnectX-5 and above] Added support for Raw WQE (<code>mlx5dv_wr_raw_wqe</code>). This feature allows applications to build a new custom work request (WQE) that is not supported by the verbs or driver and post it on normal QP. It is an extension for IBV work request (<code>ibv_wr_*</code>) with <code>mlx5</code> -specific features for sending a work request.	rdma-core v36	N/A
2356957	ah_to_qp Mapping	[ConnectX-6 Dx] Added support for mapping a QP to AH over DEVX API, which enables DC/UD QPs to use multiple CC algorithms in the same data center.	rdma-core v36	xx.31.1014
2352179	Relaxed Ordering for Kernel ULPs	[ConnectX-4 and above] Added support for enabling Relaxed Ordering for Kernel ULPs. Using relaxed ordering can improve performance in some setups. Since kernel ULPs are expected to support RO, it is enabled for them by default so they can benefit from it.	N/A	N/A
2339814	Timestamp Format in QP Database	When the device works in real-time mode, and the CQ was created with <code>IBV_WC_EX_WITH_COMPLETION_TIMESTAMP_WALLCLOCK</code> , the CQEs timestamps will be in a real-time format.	N/A	N/A
2163186	ibv_query_qp_data_in_order() verb	[All HCAs] Added support for <code>ibv_query_qp_data_in_order()</code> API. This API enables an application to check if the given QP data is guaranteed to be in order, enabling poll for data instead of poll for completion.	rdma-core v36	xx.31.1014
1748169	DV "Signature API"	[ConnectX-5 and above] Added support for "Signature API" which allows application-level data-integrity checks via a signature handover mechanism. Various signature types, including CRC32 and T10-DIF, can be automatically calculated and checked, stripped, or appended during the transfer at full wire speed.	rdma-core v36	N/A

Steering UserSpace Features				
2446723	Duplicate Rules Insertion	[ConnectX-5 and above] Added support for ability to allow or prevent insertion of duplicate rules. This enables to choose between the following behaviors: 1. Prevent duplicate rules, so that already-existing rule and fail can be detected. 2. Allow duplicate rules, to enable updating the rule's action (this will only take effect once the previous rule is deleted). By default, duplicate rules are allowed.	rdma-core v36	N/A
2368372	MLX5DR Packet OK and Checksum Checks	[ConnectX-6 Dx] Added support for new matching fields ipv4_checksum_ok and I4_checksum_ok. I3_ok, I4_ok.	rdma-core v36	xx.30.1004
2350116	MLX5DR Match Definer	[ConnectX-6 Dx] Added support for match definers which are used internally in the mlx5dv_dr API. Definers allow filtering on more packet fields, improving the packet rate and accelerates mlx5dv_dr API.	rdma-core v36	N/A
2334773	Matching on RAW Tunnel Headers	[ConnectX-5 and above] Added DR support for matching on RAW tunnel headers using the misc5 parameters. This feature allows matching on each bit of the header, including reserved fields.	rdma-core v36	N/A

Upstream Kernel 5.13				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP ² Features				
2576790	Offloading Extended ct_state Flags	[ConnectX-5 and above] Added support to offload ct_state flags rpl, inv, and rel. For rpl, support was added for both set and not set matching offload (i.e., +rpl and -rpl). For inv and rel, support was added only for the not set option (i.e., -rel and -inv).	ovs 2.14.1	N/A
2392316	vDPA Instances on Scalable Functions	[ConnectX-5 and above] Allow to create vDPA instances on Scalable Functions (subfunctions) using devlink commands.	iproute 5.13	N/A
2055404	Uplink Represent or Modes	[ConnectX-5, ConnectX-6, ConnectX-6 Dx and above] Support for new_netdev mode in SwitchDev mode. The new default behaviour is to always keep the NIC netdev.	N/A	N/A
ASAP ² and Steering Features				
1886181	Offloading sFlow Sampling Rules	[ConnectX-6 Dx] Added support for offloading sFlow sampling rules.	N/A	N/A
2330570	Simultaneous Activation of CQE Compression and PTP	[ConnectX-4 and above] Added support for the activation of PTP and CQE compression simultaneously. Since CQE compression might harm the accuracy of the PTP, the feature enables PTP packets to be moved to a dedicated queue where they are not subjected to compression. However, this configuration conflicts with setting aRFS. Turning off CQE compression, causes a hiccup in traffic which may cause a loss of synchronization. To overcome this, restart the synchronization. Note: this combination is supported only for Ethernet driver. Other driver profiles, like IPoIB and representors, do not support this combination.	N/A	N/A
NetDev Features				
1821120	TLS RX Hardware Offload	[ConnectX-6 Dx and above] Added GA-level support for hardware offload decryption of TLS traffic over crypto-enabled ConnectX-6 Dx NICs and above.	N/A	xx.28.2006
2257993	TLS RX Resynchronization Resiliency	[ConnectX-6 Dx and above] Added support to have driver resiliency against high load of RX resync operations.	N/A	N/A
2225882	ethtool Extended Link State in mlx5e	[ConnectX-4 and above] Added ethtool Extended Link State to mlx5e. ethtool can be used to get more information to help the user troubleshoot the state. For example, if there is no link due to a missing cable (\$ ethtool eth1 ... Link detected: no (No cable)). Besides the general extended state, drivers can pass additional information about the link state using the sub-state field (for example: \$ ethtool eth1 ... Link detected: no (Autoneg, No partner detected)). The extended state is available only for some cases of no link; in other cases, ethtool will print only "Link detected: no", as before.	N/A	N/A
Core Features				

2041295	Scalable Functions (Subfunctions)	<p>[ConnectX-5 and above] Added support for scalable functions (also called subfunctions). The feature enables the user to create, configure, and deploy a scalable functions (e.g., RDMA and networking applications) and to assign them to a container when a container is started via mlxdevm tool.</p> <p>A scalable function can also be deployed in an untrusted guest/host system from the NIC/DPU. This enables full configuration of the function and its representors from the NIC/DPU before giving the function for a container to run in a host system.</p> <p>For more information, see https://github.com/Mellanox/scalablefunctions/wiki/MLNX_OFED-step-by-step-guide.</p>	devm	xx.30.1004 (ConnectX-5, 6, 6-Dx) xx.31.1014 (Bluefield)
2401360	Reading DSFP Module Information	<p>[ConnectX-4 and above] Added support for reading DSFP module information. The change included adding new options to ethtool netlink EEPROM module read API, to read specific page and bank.</p>	iproute2 5.13	xx.30.1004

RDMA Features

2404639	Export Object IDs to Users	<p>[ConnectX-4 and above] Extended support for the "rdma res show" command to SRQ and context resources.</p>	iproute2 v5.14	N/A
2360568	Huge Page Support for devx umems	<p>[ConnectX-4 and above] Added support to allow DEVX UMEM to be created with larger page sizes than 4K. This is a must for some device objects (e.g. RegEx) and, in any case, page size larger than 4K may need less MTTs which may improve performance.</p> <p>A new API <code>mlx5dv_devx_umem_reg_ex()</code> was added which requests a specific page sizes. It enables better application control on the required UMEM page size.</p>	rdma-core V35	N/A
2335586	Dynamically Assign MSI-X Vectors Count	<p>[ConnectX-5, ConnectX-6, ConnectX-6 Dx and above] The number of MSI-X vectors is PCI property visible through lspci, that field is read-only and configured by the device. The static assignment of an amount of MSI-X vectors does not allow utilize the newly created VF because it is not known to the device the future load and configuration where that VF will be used. The VFs are created on the hypervisor and forwarded to the VMs that have different properties (for example number of CPUs). To overcome the inefficiency in the spread of such MSI-X vectors, we allow the kernel to instruct the device with the needed number of such vectors, before VF is initialized and bounded to the driver.</p>	N/A	N/A

Steering Features

2239829	Match on Dynamic Configured Flex Parser	<p>[ConnectX-5 and ConnectX-6 Dx] Added support to allow to match on dynamic configured flex parser. The PRM allows for configuring a flex parser by creating a PARSE_GRAPH_NODE general object. This version contains the support to match the parsed value using this configured flex parser using misc4 PRM struct containing an index and value for the configured flex parser.</p>	N/A	N/A
---------	---	--	-----	-----

Steering UserSpace Features

2380582	Pop VLAN on VF/SF Tx Direction	<p>[ConnectX-6 Dx and above] Added support to pop VLAN on VF/SF Tx direction.</p>	rdma-core v35	N/A
2363360	Dump Single Flow	<p>[ConnectX-5, ConnectX-6 Dx and above] Added support to dump single flow/rule with flow-id.</p>	rdma-core v35	N/A

Upstream Kernel 5.12

Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
---------------	---------	-------------	-------------------------------------	-----------------------------------

ASAP² Features

2402429	Connection Tracking Replay State	<p>[BlueField-2 with CentOS kernels] Added support for matching on CT state replay.</p>	iproute2 V5.12, OVS 2.14.1	N/A
2198083	Connection Tracking Mirroring Offload	<p>[ConnectX-5 and above] Added support for using Mirroring Offload with Connection Tracking.</p>	N/A	N/A
2103546	Tunnel Rules Offload	<p>[ConnectX-6 Dx] Added support for offloading tunnel rules when the source interface is VF (in addition to uplink) in the Hypervisor.</p>	N/A	N/A
2089370	Kernel Software Steering ConnectX-6 Dx Support	<p>[ConnectX-6 Dx] Added support for kernel software steering on ConnectX-6 Dx adapter cards.</p>	N/A	N/A

NetDev Features				
2252726	KTLS RX and TX Hardware Offload Support for BOND	[ConnectX-6 Dx and above] Added support in bond driver to benefit from kTLS offload on capable lower devices.	N/A	N/A
1871320	Traffic Engineering—Hierarchical QoS	[ConnectX-5 and above] This feature allows the user to offload the HTB qdisc to the NIC, allowing it to scale better by eliminating a single locking point. The configuration is done with the:) similar tc commands.	iproute2 V5.12	xx.29.2002
1869050	WHAT-JUST-HAPPENED-support for NICs	[ConnectX-4 and above] Added support for WJH on NICs. WJH allows visibility of dropped packets and their contents to increase ability for analysis and debugging. WJH is a service in devlink context and it is already implemented in the switch. Note: processing dropped packets (even for visibility purposes) may cause a degradation in performance and leaves the driver vulnerable for malicious attacks. The feature is disabled by default. <ul style="list-style-type: none"> VLAN mismatch – existing generic trap DEVLINK_TRAP_GENERIC_ID_DMACH_MISMATCH—Traps received packets with wrong VLAN tag DMAC mismatch – new generic trap DEVLINK_TRAP_GENERIC_ID_DMACH_MISMATCH—Traps received packets with wrong destination MAC Devlink provides an infrastructure called devlink trap which allow a device to register / unregister and to enable / disable traps. Devlink traps also provide traps grouping and policing. The trapped packets are monitored and then forward to the drop monitor. Drop monitor is used to send notifications to user space about dropped packets. Note: For this release, WJH support for NICs will not implement Policy.	Devlink infrastructure	N/A
1839182	Synchronize Hardware Clock to PTP Signal on ConnectX-6 DX	[ConnectX-6 Dx] Added support for the hardware clock device to be adjusted and provide timestamps which are translated into real-time nanoseconds. This can be used by the driver for PTP protocol.	rdma-core v34	xx.28.1002
Core Features				
2479696	Flex IO	[ConnectX-7 and above] Added support for auxiliary processor on NICs called Flex IO. It is possible to open a dedicated CQ for this processor in order to trigger it. This requires support for opening a CQ without a valid EQ number when the following conditions are met: (1) HCA_CAP.apu = 1, (2) apu_thread_cq = 1 in CQC. In such a case, EQ number should be set to 0xFFs on CQ creation.	N/A	N/A
RDMA Features				
2392058	ODP Locking Optimization	[ConnectX-4 and above] Added support for cleanup of the synchronize_srcu() from the ODP flow because it was a time consuming part of dereg_mr. Note: This only affects the driver and not the firmware.	N/A	N/A
2359452	NDR Speed	[ConnectX-7] Added support for NDR Speed (400Gb/s InfiniBand rate) in mlx5 driver.	N/A	N/A
2320662	DCT Support for Connection Establishment with RDMA_CM	[ConnectX-5 and above] Added support for the dv APIs to allocate/deallocate a unique QP number that can be used as DC QPN in RDMA_CM connection establishment.	rdma-core v34	xx.29.2002
Steering UserSpace Features				
2330345	Connection Tracking Window Validation	[ConnectX-6 Dx] Added support for ASO connection tracking of action creation and modification. This action allows performing TCP connection tracking using hardware offloads. Using this offload, the validity of the connection state of the incoming or outgoing packets on this TCP connection can be examined. Also added the ability for an ASO CT action created on one GVMI to be used on different GVMI.	rdma-core v34	xx30.1004
2275829	mlx5dv_dr API Support for GTP-U Extension Header	[ConnectX-5 and ConnectX-6 Dx] Added mlx5dv_dr API support for matching on a new field "gtpu_first_ext_dw_0". This field enables packet filtering based on the GTP-U first extension header (first dword only). To enable parsing of tunnel GTP-U extension header, run the following command. ./cloud_fw_reset.py FLEX_PARSER_PROFILE_ENABLE=3.	rdma-core v34	xx.29.1016
2264190	Modifying GTPU TEID	[ConnectX-6 Dx and BlueField-2] Added support to modify GTPU TEID. This support is experimental. Modifying TEID requires configuring flex parser.	rdma-core v34	N/A
Upstream Kernel 5.11				
Internal Ref.	Feature	Description	Support Added in User Space Version	Support Added in Firmware Version
Core Features				

2329342	Auxiliary Bus in mlx5 Driver	[ConnectX-4 and above] Updated mlx5 driver to use auxiliary bus in order to integrate different driver components into driver core and optimize module load/unload sequences.	N/A	N/A
NetDev Features				
2263652	TX CQE Error Resiliency	[ConnectX-3 and ConnectX-3 Pro] Added TX CQE error resiliency on mlx4 driver. When error CQE was found while polling TX CQ, the QP is in error state and all posted WQEs generated error CQEs without any data transmitted. Added resiliency flow to handle it by reopening the channels.	N/A	N/A
2143621	TX Port Timestamping for PTP Packets	[ConnectX-6 Dx] Added support for TX Port Timestamping for PTP Packets. Transmitted packet timestamping accuracy can be improved when using timestamp generated at the port level instead of timestamp generated when CQE is generated. TX port timestamping better reflects the actual time of a packet's transmit and has less jitter. As part of the feature, the driver is to open a special SQ per TC (PTP SQ). The driver is to steer the following streams to these special SQs: 1. SKBTX_HW_TSTAMP was set at tx_flag (SO_TIMESTAMPING was set via setsockopt() or similar) 2. Packet type is (a) non-IP, with EtherType of PTP over IEEE 802.3 (0x88f7) or (b) UDP over IPv4/IPv6. By default, the feature is disabled (to avoid the extra memory allocated for these SQs). The feature can be enabled/disabled by the following command: ethtool --set-priv-flags <ifs-name> tx_port_ts on / off.	N/A	xx.29.1016
RDMA Features				
2256043	Multi-Application QoS	[ConnectX-5 and above] Added support for configuring QoS on a single QP or on a group of QPs.	N/A	N/A
RDMA UserSpace (RDMA-Core) Features				
2315048	mlx5dv API	[ConnectX-4 and above] Added support for mlx5dv API to modify the configured UDP source port for RoCE packets of a given RC/UC QP when QP is in RTS state.	rdma-core v33	xx.29.1016
Steering UserSpace Features				
1884522	Kernel Software Managed Flow Steering (SMFS) Performance	[ConnectX-5, ConnectX-6 Dx & BlueField, BlueField-2] Improved the performance of Kernel software steering by reducing its memory consumption.	rdma-core v33	N/A
2254030	mlx5dv_dr API Matching on Geneve Tunnel	[ConnectX-5, ConnectX-6 Dx & BlueField, BlueField-2] Added support for the option to match mlx5dv_dr API on Geneve tunnel using a dynamic flex parser. The option header consists of class, type, length, and data. The parser should be configured using the devx command, after which a rule can be created to match on parser ID and data.	rdma-core v33	N/A
2239811	Pop/Push VLAN	[ConnectX-6 Dx] Enabled pop VLAN action in TX and push VLAN action in RX. This will allow the host to send and receive packets with a VLAN, but, in steering, this VLAN can be removed or replaced.	rdma-core v33	N/A
2238570	mlx5dv_dr API ASO First Hit	[ConnectX-6 Dx & BlueField-2] Added support for ASO first hit using the mlx5dv_dr API which allows for tracking rule hits by packets. When a packet hits a rule with the ASO first-hit action, a flag is set indicating this event and the original value of the flag is copied to one of the C registers.	rdma-core v33	xx.29.1016
2228100	mlx5dv_dr API ASO Flow Meter	[ConnectX-6 Dx & BlueField-2] Added support for ASO flow meter using the mlx5dv_dr API which allows for monitoring the packet rate for specific flows. When a packet hits a flow that is connected to a flow meter, the rate of packets through this meter is evaluated and the packet is marked with a color copied into one of the C registers.. It is marked according to the current rate compared to the reference rate.	rdma-core v33	xx.31.10xx
2154380	Software Steering Parallel Rules Insertion	[ConnectX-6 Dx & BlueField-2] Added support for a locking mechanism to enable parallel insertion of rules into the software steering using the mlx5dv_dr API. The parallel insertion improves the insertion rate and takes place when adding RX and TX rules via the FDB domain.	rdma-core v33	N/A
Upstream Kernel 5.10				
Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP Features				
2282002	IPv6 TOS Rewrite	[ConnectX-5 and above] Added support for offloading TOS value rewrite for IPv6 packets.	N/A	N/A
2275891	Connection Tracking with Hairpin	[ConnectX-5 and above] Added support for adding connection-tracking rules on virtual functions (VF) in order to forward traffic from one VF to the other.	N/A	xx.28.4000
1890766	Connection Tracking with Hairpin	[ConnectX-5 and above] Added support for adding connection-tracking rules on uplinks in order to forward traffic from one uplink to the other.	N/A	N/A

2234654	TC Trap	[ConnectX-5 and above] Added support for offloading TC trap actions, allowing for redirection of packets to slow path (software).	N/A	N/A
1750259	SR-IOV Live Migration	[ConnectX-5 and above] Added support for performing a live migration for a VM with an SR-IOV NIC VF attached to it and with minimal to no traffic disruption. This feature is supported in SwitchDev mode; enabling users to fully leverage VF TC/OVS offloads, where the failover inbox driver is in the Guest VM, and the bonding driver is in the Hypervisor. Note that you must use the latest QEMU and libvirt from the Upstream github.com sources.	N/A	N/A
RDMA Core Features				
1995519	Command Interface Resiliency From A Lost Interrupt	[ConnectX-4 and above] Added support for command interface resiliency from a lost interrupt, which manually polls the command EQ in case of a command timeout. If the resiliency mechanism finds non-handled EQE (due to a lost interrupt), it will consume it and return the actual status to the caller. Without this feature, when the driver detects a command interface command timeout, it warns the user and returns a timeout error to the caller and the entry of the command is not evacuated (because only real event interrupt can clear a command interface entry). If the hardware event interrupt never arrives, this entry is left unused forever. Because command interface entries are limited, the ability to post a new command can eventually be lost. In addition, if the driver does not consume the EQE of the lost interrupt and rearm of EQ, no new interrupts will arrive for other completed commands.	N/A	N/A
1959838	Firmware Live Patching	[ConnectX-4 and above] Added support for firmware live patching, which allows to update firmware without reset and applies to minor changes in firmware. The firmware checks if the firmware diff can be applied by live patching.	iproute2 v5.10	xx.28.2006
1889839	Firmware Reset on Devlink	[ConnectX-4 and above] Added support for initiating NIC firmware reset for a firmware upgrade through the devlink tool.	iproute2 v5.10	xx.28.2006
1869155	Multi-Host Reset	[ConnectX-4 and above] Added support for multi-host firmware reset for firmware upgrades. When a new firmware was burnt on flash and is pending activation, the multi-host reset loads the new firmware. If no new firmware is pending, the current firmware image will be reloaded from flash.	iproute2 v5.10	xx.28.2006
NetDev Features				
2176729	CQE Compressi on Multi-Strides	[ConnectX-5 and above] Added support for RX CQE compression for multi-strides packets.	N/A	N/A
2145391	IPsec Crypto Offload	[ConnectX-6 Lx and above] Added support for IPsec crypto offload. IPsec offload is a set of features aimed at reducing the CPU overhead of using IPsec. It targets the heavy crypto (encryption/decryption) operation, offloading them to the hardware. Connectx6-Lx is the first ASIC to support IPsec crypto offload and handles all encryption, decryption, and authentication, leaving the rest (replay, encapsulation, and decapsulation) to the software IPsec implementation. This approach benefits from increased flexibility in software, at the cost of some performance overhead.	N/A	xx.28.1002
1885835	Enhanced TX MPWQE	[ConnectX-5 and above] Added support for Enhanced TX MPWQE which allows for acceleration of TX datapath by saving PCI bandwidth and CPU utilization. The savings are achieved by aggregating multiple packets into a single WQE. The feature is driven by xmit_more for certain traffic types, such as UDP.	N/A	N/A
RDMA Features				
2149631	Memory Consumpti on of Memory Registration	[ALL] Added support for optimization of memory consumption of memory registration. Without this change, a huge pages system with 2MB page size would have a memory consumption of x512).	N/A	N/A
2048075	Query GID API	[ALL] Added support for a new query GID API which works over ioctl and allows querying a single GID entry by its port and GID index or querying all the GID tables of a specific device. This is much faster than querying it over sysfs.	rdma-core v32	N/A
1974041	ODP On Demand Synchroniz ation	[ConnectX-4 and above] Added support to expose an option to prefetch ODP MR without faulting. This enables updating the device page table with the presenting CPU pages and reducing page faults in the system.	rdma-core v32	N/A
Steering UserSpace Features				
2228217	Remote Mirroring	[ConnectX-5 and above] Added support for remote mirroring, which allows for duplication of packets to multiple different destinations as well as performing a variety of different actions per destination.	rdma-core v32	N/A
2068596	Software Steering	[ConnectX-6 Dx and above] Added support for software steering on ConnectX-6 Dx adapter cards in the user-space RDMA-Core library through the mlx5dv_dr API.	rdma-core v32	N/A
2066943	Sampling Action on Direct Rule API	[ConnectX-5 and above] Added support for Direct Rule (DR) API, enabling sampling action to allow duplication and sampling of a ratio of the traffic. The API also allows to perform a set of different DR actions on the sampled traffic.	rdma-core v33	xx.28.1002
Upstream Kernel 5.9				

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
NetDev Features				
2067103	Relaxed Ordering	[ConnectX-4 and above] Relaxed ordering is a PCIe feature which allows flexibility in the transaction order over the PCIe. This reduces the number of retransmissions on the lane, and increases performance up to 4 times. By default, mlx5e buffers are created with Relaxed Ordering support when firmware capabilities are supported. <code>pci_ro_write [on/off]</code> . Note: Some CPUs which are not listed in the kernel's blacklist may suffer from buggy implementation of relaxed ordering, in which case the user may experience a degradation in performance and even unexpected behavior. To turn off relaxed ordering: "RlxOrd-" : <code>setpci -s82:00.0 CAP_EXP+8.w=294e</code>	N/A	N/A
1888097	Devlink Port Health Reporters	[ConnectX-4 and above] Added per-port reporters to devlink health to manage per-port health activities. Users can now access the devlink port reporters by specifying the port index in addition to the device devlink name through the devlink health commands API. This update was first introduced in iproute2 v5.8. As part of this feature, mlx5e Tx and Rx reporters are now redefined as devlink port reporters. For examples, please see devlink-health manpage.	iproute2 v5.9	N/A
1821120	TLS RX Hardware Offload	[ConnectX-6 Dx and above] Added beta-level support for hardware offload decryption of TLS traffic over crypto-enabled ConnectX-6 Dx NICs and above.	N/A	xx.28.2006
RDMA Features				
2067098	CQ and QP Context Exposure	[ConnectX-4 and above] Exposed QP, CQ and key context in raw format via RDMA tool.	N/A	xx.28.0214
2056081	IPoIB QP Number Creation	[ConnectX-4 and above] Enabled setting the QP number of an IPoIB PKey interface in Enhanced mode. This is done using the standard <code>ip link add</code> command while padding the hardware address of the newly created interface. The QP number is the 2nd-4th bytes. To enable the feature, the <code>MKEY_BY_NAME</code> configuration should firstly be enabled in NVCONFIG.	N/A	N/A
2047060	RoCE Selective Repeat	[ConnectX-6 and above] This feature introduces a new QP retransmission mode in RoCE in which dropped packet recovery is done by re-sending the packet instead of re-sending the PSN window only (Go-Back-N protocol). This feature is enabled by default when RDMA-CM is being used and both connection nodes support it.	rdma-core v31	xx.27.1016
1955266	Shared PD	[ConnectX-4 and above] Added the ability to share a PD between processes for RDMA usage. This feature is particularly critical for business logic solutions that are based on multi-process designs, such as: <ul style="list-style-type: none"> • NGINX, with TCP load balancing, sharing the RSS indirection table with RQ per process. • HPC frameworks with multi-rank (process) solution on single hosts. UCX can share IB resources using the shared PD and can help dispatch data to multiple processes/MR's in a single RDMA operation. • Primary processes registering a large shared memory range, where each worker process spawned can create a private QP on the shared PD, and use the shared MR to save the registration time per-process. 	rdma-core v31	N/A
Upstream Kernel 5.8				
Internal Reference Number	Feature/Change	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
1752536	MPLS-over-UDP Hardware Offload Support	[ConnectX-5 and above] Added support for encap/decap hardware offload of IPv4 traffic over MPLS-over-UDP. This can be used in networks with MPLS routers to achieve more efficient routing.	iproute2 v5.8	N/A
NVMe Features				
1915224	Metadata/T10-PI Offload Support for NVMe over Fabrics (RDMA) Host and Target	[ConnectX-4 and above] Added support for providing offloaded end-to-end data integrity and validation according to T10-PI standard, also known as Metadata in NVMe specification.	N/A	N/A
Steering Features				
2077839	Software Steering Support for Memory Reclaiming	[ConnectX-5 and above] Added support for reclaiming device memory to the system when it is not in use. This feature is disabled by default and can be enabled using the command <code>mlx5dv_dr_domain_set_reclaim_device_memory()</code> .	rdma-core v30	N/A

1899086	mlx5dv_dr Memory Allocation	[ConnectX-5 and above] Reduced memory allocation time when using the mlx5dv_dr API. This is particularly significant for the first inserted rules on which memory is allocated.	rdma-core v30	N/A
1894989	Default Miss Action Support	[ConnectX-5 and above] Added support for a new API for creating a Default Miss Action to achieve better performance. When a packet is matched with the default "miss action", the packet will be steered to the default "miss" of the steering domain.	rdma-core v30	N/A
RDMA Features				
2053467	mlx5dv Helper APIs for Tx Affinity Port Selection	[ConnectX-4 and above] Added support for the following mlx5dv helper APIs which enable the user application to query or set a QP's Tx affinity port number in a LAG configuration. <ul style="list-style-type: none"> mlx5dv_query_qp_lag_port mlx5dv_modify_qp_lag_port 	rdma-core v31	xx.28.0214
2051494	RDMA Tx Steering	[ConnectX-4 and above] Enabled RDMA Tx steering flow table. Rules in this flow table will allow for steering transmitted RDMA traffic.	rdma-core v30	xx.28.xxxx
1919657	RDMA-CM Enhanced Connection Establishment (ECE)	[ConnectX-6 and above] Added support for allowing automatic enabling/disabling of vendor specific features during connection establishment between network nodes, which is performed over RDMA-CM messaging interface.	rdma-core v31	xx.28.xxxx
1888194	RDMA-CM Path Alignment	[ConnectX-4 and above] Added support for RoCE network path alignment between RDMA-CM message and QP data. The drivers and network components in RoCE calculate the same hash results for egress port selection both on the NICs and the switches.	N/A	xx.28.0214
1888159	RDMA-CM Disassociate Support	[All] Added support for connecting kernel and RDMA-CM in a reliable way based on device index.	rdma-core v30	N/A
1779659	RoCEv2 Flow Label and UDP Source Port Definition	[ConnectX-4 and above] This feature provides flow label and UDP source port definition in RoCE v2. Those fields are used to create entropy for network routes (ECMP), load balancers and 802.3ad link aggregation switching that are not aware of RoCE headers.	rdma-core v30	xx.24.0222

Upstream Kernel 5.7

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
ASAP² Features				
2074669	TC Query Stats and Flow Insertion Dependency Removal	[ConnectX-5 and above] Optimized the kernel's TC subsystem to use fine-grained per-action locks instead of global RTNL lock when constructing flow-action intermediate representation for the driver. This change allows for rules insertion to run in parallel with rule dump, which significantly improves performance of live Open vSwitch system.	N/A	N/A
1818081	Hairpin Support in Switch Mode	[ConnectX-5 and above] Added support for receiving tunneled traffic from the uplink port, after which it can be decapsulated and sent them back to the uplink port, possibly encapsulated with new tunnel information.	N/A	xx.28.1000
1765236	Connection Tracking Offload	[ConnectX-5 and above] Added support for offloading TC filters containing connection tracking matches and actions.	OVS v2.13 + iproute2 v5.4.0	xx.26.1040
Core Features				
1928787	Firmware Reactivation	[ConnectX-4 and above] Added support for safely inserting consecutive firmware images without the need to reset the NIC in between.	N/A	xx.28.1000
Steering Features				
1997168	GPRS Tunneling Protocol (GTP) Header	[ConnectX-5 and above] Added support for matching (filtering) GTP header-based packets using mlx5dv_dr API over user-space RDMA-Core library.	rdma-core v29	N/A
RDMA Features				
2067127	Custom Parent-Domain Allocators for CQ	[ConnectX-4 and above] Enabled specific custom allocations for CQs.	rdma-core v29	N/A
1890393	RDMA-CM DevX Support	[ConnectX-4 and above] Added support for DevX in RDMA-CM applications.	rdma-core v29	N/A

1830364	Packet Pacing DevX Support	[ConnectX-5 and above] Enabled RiverMax to work over DevX with packet pacing functionality by exposing a few DV APIs from rdma-core to enable allocating/destroying a packet pacing index. For further details on usage, see man page for: <code>mlx5dv_pp_alloc()</code> and <code>mlx5dv_pp_free()</code> .	rdma-core v29	xx.27.1016
NetDev Features				
1888096	Devlink Port Support in Non-represorator Mode	[ConnectX-4 and above] Added support for viewing the mlx5e physical devlink ports using the 'devlink port' command. This also may affect network interface names, if predictable naming scheme is configured. Suffix indicating a port number will be added to interface name.	N/A	N/A
1888020	Get FEC Status on PAM4 /50G	[ConnectX-6 & ConnectX-6-Dx] Allowed configuration of Reed Solomon and Low Latency Reed Solomon over PAM4 link modes.	ethtool v5.5	N/A
1702814/1702706	Devlink Health Dump	[ConnectX-4 and above] Added support for the <code>devlink health dump show</code> command in the Rx and Tx reporters. The dump contains the content given by the firmware for <code>wqump/rxdump/sxdump</code> in binary format. The user can retrieve the binary data via <code>devlink (v5.3 and on)</code> and use a NVIDIA parsing tool to analyze the dump.	iproute2 v5.3	xx.27.1016

Upstream Kernel 5.6				
----------------------------	--	--	--	--

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
---------------------------	-----------------	-------------	-------------------------------------	-----------------------------------

RDMA Features				
----------------------	--	--	--	--

1978779	ODP Huge Pages Support	[ConnectX-4 and above] Enabled ODP MR to work with huge pages by exposing <code>IBV_ACCESS_HUGETLB</code> access flag to indicate MR range mapping by huge pages. This flag is applicable only with conjunction of <code>IBV_ACCESS_ON_DEMAND</code> .	rdma-core v28	N/A
1973016	RoCE Accelerator Counters	[ConnectX-4 and above] Added the following RoCE accelerator counters: <ul style="list-style-type: none"> <code>roce_adp_retrans</code> - counts the number of adaptive retransmissions for RoCE traffic <code>roce_adp_retrans_to</code> - counts the number of times RoCE traffic reached timeout due to adaptive retransmission <code>roce_slow_restart</code> - counts the number of times RoCE slow restart was used <code>roce_slow_restart_cnps</code> - counts the number of times RoCE slow restart generated CNP packets <code>roce_slow_restart_trans</code> - counts the number of times RoCE slow restart changed state to slow restart 	N/A	N/A
1885435	vhost Data Path Acceleration (vDPA)	[ConnectX-5 and above] Added support to enable mapping the Virtio access region (VAR) to be used for doorbells by vDPA applications. Specifically, the following DV APIs were introduced (see man page for more details): <code>mlx5dv_alloc_var()</code> <code>mlx5dv_free_var()</code> .	rdma-core v28	xx.27.1016
1882973	Memory Region	[All] Added support for the user to register memory regions with a relaxed ordering access flag. This can enhance performance, depending on architecture and scenario.	rdma-core v28	xx.27.1016
1769019	Kernel Recommendation on User Memory	[ConnectX-4 and above] Added support for the Kernel to advise on the user memory region that is registered in the HCA. In order to be able to use this functionality, virtual address and memory key must be known to the caller.	N/A	N/A
1711108	User Memory Kernel Registration	[ConnectX-4 and above] Added support for the Kernel to register user virtual memory in the HCA. This requires the user to pass the address and size of the memory to the kernel in some way and the kernel must register the memory from the contexts of the user process.	N/A	N/A

Steering Features				
--------------------------	--	--	--	--

1961704	Kernel Software Steering Remote Mirroring	[ConnectX-5] Added kernel support for OVS remote mirroring to allow hardware traffic mirroring to multiple ports.	N/A	N/A
1915174	Kernel Software Steering for Connection Tracking (CT)	[ConnectX-5] Added required kernel software steering support for CT.	N/A	N/A
1840759	Userspace Software Steering Debugging API	[ConnectX-5] Added support for software steering to dump flows for debugging purposes in the userspace RDMA-Core library through the <code>mlx5dv_dr</code> API.	rdma-core v28	N/A
1754441	Discards Counters ^[1]	[ConnectX-4 and above] Exposed <code>rx_prio[p]_discards</code> discard counters per priority that count the number of received packets dropped due to lack of buffers on the physical port.	N/A	N/A

NetDev Features				
------------------------	--	--	--	--

2042124	Devlink Health State Notifications	[ConnectX-4 and above] Added support for receiving notifications on devlink health state changes when an error is reported or recovered by one of the reporters. These notifications can be seen using the userspace 'devlink monitor' command.	iproute2 v5.6	N/A
1996062	mlx5e Max Combined Channels	[ConnectX-4 and above] Increased the driver's maximal combined channels value from 64 to 128 (however, note that OOB value will not cross 64). 128 is the upper bound. Lower maximal value can be seen on the host, depending on the number of cores and MSIX's configured by the firmware.	N/A	N/A

Upstream Kernel 5.5

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1895928	Resource Allocation on External Memory	[ConnectX-5 and above] Added support to enable overriding mlx5 internal allocations in order to let applications allocate some resources on external memory, such as that of the GPU. The above is achieved by extending the parent domain object with custom allocation callbacks. Currently supported verbs objects are: QP, DBR, RWQ, SRQ.	rdma-core v27	N/A
1756530	ODP Diagnostic counters	[ConnectX-4 and above] Added ODP diagnostics counters for the following items per MR (Memory Region) within IB/mlx5 driver: <ol style="list-style-type: none"> 1. Page faults: Total number of faulted pages. 2. Page invalidations: Total number of pages invalidated by the OS during all invalidation events. The translations can no longer be valid due to either non-present pages or mapping changes. 3. Prefetched pages: When prefetching a page, a page fault is generated in order to bring the page to the main memory. 	iproute v5.4	N/A
1747149	Performance Improvements	[ConnectX-4 and above] <ul style="list-style-type: none"> • Updated Blueflame capability reporting to prevent redundant use of Blueflame when Write-combining is not supported. • Added Blueflame capabilities over VFs. 	N/A	xx.26.1040
1686052	RoCE Disablement	[ConnectX-4 and above] Added the option to disable RoCE traffic handling. This enables forwarding of traffic over UDP port 4791 that is handled as RoCE traffic when RoCE is enabled. When RoCE is disabled, there is no GID table, only Raw Ethernet QP type is supported and RoCE traffic is handled as regular Ethernet traffic.	N/A	N/A

rdma-core Only Features^[1]

1871545	Hardware Clock Exposure	[ConnectX-5 and above] Added support for querying the adapter clock via mlx5dv_query_device.	rdma-core v27	N/A
---------	-------------------------	--	---------------	-----

Steering Features

1915338	Direct Verbs Support for Batch Counters on Root Table	[ConnectX-5] Added support for mlx5dv_dr API to set batch counters for root tables.	rdma-core v27	N/A
1907530	Modify Header	[ConnectX-5 & BlueField] Added support for mlx5dv_dr_actions to support up to 32 modify actions.	rdma-core v27	N/A
1757878	GENEVE Encap/Decap Rules Offload	[ConnectX-5] Added support for GENEVE encapsulation/decapsulation rules offload using software steering.	rdma-core v27	N/A

ASAP² Features

1920181	OVS-Kernel ToS Rewrite	[ConnectX-5 and above] Added support for Type of Service (ToS) rewrite in the OVS-Kernel.	N/A	N/A
---------	------------------------	---	-----	-----

Upstream Kernel 5.4

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1717501	RDMA_Rx RoCE Flow	[ConnectX-4 and above] Extended support for mlx5 with RDMA_RX RoCE flow steering to DevX objects.	rdma-core v26	xx.26.1000

1747205	Steering Support	[ConnectX-4 and above] Extended support for mlx5 with RDMA_RX RoCE flow steering to QP objects.	rdma-core v 26	xx.26.1000
1821108	ODP Support over DC Transport	[ConnectX-4 and above] Added support for On-Demand-Paging over DC transport.	rdma-core v 26	N/A
Steering Features				
1751025	Push/Pop VLAN over Software Steering	[ConnectX-5 and above] Added support to pop and push VLAN actions in software steering in SwitchDev mode.	N/A	N/A
1762181	Flow and Recycle Counters	[ConnectX-4 and above] Added support for a pool of flow counters, based on flow counter bulks, with the option to acquire and release flow counters from and to the pool. This shortens the time needed to allocate a flow counter using a FW command. The effective memory footprint of a flow counter is also increased by 10% thanks to this feature.	N/A	16.25.1000
1703811	Kernel Software Steering for eSwitch	[ConnectX-5] Added software steering capabilities to the SR-IOV eSwitch in SwitchDev mode. Software steering enables better rules update rate compared to the current firmware-based solution. This is achieved by performing calculations on the main CPU which allows for higher update rates.	N/A	N/A
NetDev Features				
1700763	PF to Show VF Counters	[ConnectX-4 and above] Added support for the VF to send real-time rings statistics (Rx/Tx packets/bytes) to the PF over HyperV PCI configuration channel. These statistics can be monitored via perfmon tool.	N/A	N/A
1821239	Ethtool Counters	[ConnectX-5 and above] Added counter dev_out_of_buffer to ethtool statistics to count the number of times there was no enough buffer allocated to the device internal queue.	N/A	xx.26.1040
1821554	IP-in-IP Tunnel Offload for Checksum and TSO	[ConnectX-5 and above] Added support for the driver to offload checksum and TSO in IP-in-IP tunnels.	N/A	xx.26.4012
1858107	AF_XDP need_wakeup Flag Support	[ConnectX-4 and above] Added driver support for the flag "need_wakeup" in AF_XDP Tx and FILL rings. When both the AF_XDP application and the driver support this flag, the driver is able to signal to the application when it needs to invoke a syscall to wake the driver explicitly, and when it needs not. This significantly improves single-core performance by reducing resource contention and avoiding wasting resources on busy polling and context switches.	N/A	N/A
1821229	VF LAG Load Balancing	[ConnectX-4 and above] When VF LAG is in use, round-robin the Tx affinity of channels among the different ports, if supported by the firmware, enables all SQs of a channel to share the same port affinity. This allows the distribution of traffic sent from a VF between two ports, as well as round-robin the starting port among VFs to distribute traffic originating from single-core VMs.	N/A	xx.26.1040
1760559	Virtual Output Queuing (VoQ) Counters	[ConnectX-6 Dx] Exposed rx_prio[p]_buf_discard, rx_prio[p]_wred_discard and rx_prio[p]_marked firmware counters that count the number of packets that were dropped due to insufficient resources.	N/A	See note below ¹
1821559	IP-in-IP RSS Offload	[ConnectX-4 and above] Added support for receive side scaling (RSS) offload in IP-in-IP (IPv4 and IPv6).	N/A	N/A
1782820	Devlink Health Reporters	[ConnectX-4 and above] Added support for monitoring and recovering from errors that occur on the Rx queue, such as CQE errors and timeout.	N/A	N/A
¹ ConnectX-6 Dx FW version supporting this feature is not available yet.				
ASAP² Features				
1682148	Parallel TC Rules Insertion	[ConnectX-5 and above] Removed global RTNL lock dependency in the kernel and in mlx5 TC layers to allow updating parallel TC rules.	N/A	N/A
1767933	Offloading Ingress Rate Control Offload on E-Switch Ports	[ConnectX-5 and above] Added support for configuring the ingress rate of an E-Switch port. This is achieved by the ability to control the transmit rate of a virtual function whose data appears at the ingress of the Switch port. Thus, controlling the VF Tx rate effectively controls the ingress rate of an E-Switch Port.	N/A	N/A
Upstream Kernel 5.3				

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1755657	Device Emulation Infrastructure	[BlueField] Added support for Device Emulation in BlueField. This mechanism allows function-A to perform operations "on behalf" of function-B. The emulation manager creates a channel (named VHCA_TUNNEL general object) that acts as the direct command interface between the emulated function host and the HCA hardware. The emulation software creates this tunnel for every managed function and issues commands via the DevX general command interface.	N/A	xx.26.1040
1711102	DevX Async Events Support	[ConnectX-4 and above] Added support for reading device events over DevX API.	rdma-core v25	xx.25.1020
1686017	QP Counters and Firmware Errors per PID	[ConnectX-5 and above] QP counters and flow counters are now set per Process ID (PID) to allow better visibility of RDMA error states. Users will be able to manually tune the Q counter to monitor specific QPs, or automatically monitor QPs according to predefined criteria, such as the QP type.	iproute2 v5.3	xx.25.1020
1679990	Dynamic Interrupt Moderation for RDMA Kernel	[ConnectX-5 and above] Added dynamic interrupt moderation (DIM) library for RDMA to improve bandwidth and reduce CPU overhead.	iproute2 v5.3	N/A
1160889	Signature Offload Kernel Verbs Enhancements	[ConnectX-4 and above] Added a new API which enables posting a single WR that completes the Protection Information (PI) operation internally. This reduces CPU utilization for posting and processing multiple WRs and improves performance by choosing the optimal mkey for the hardware according to the buffer memory layout.	N/A	N/A
1704114	Message Signaled Interrupts-X (MSI-X) Vectors	[ConnectX-4 and above] Added support for using a single MSI-X vector for all control event queues instead of one MSI-X vector per queue in a virtual function driver. This frees extra MSI-X vectors to be used for completion event queue, allowing for additional traffic channels in the network device.	N/A	xx.26.1040
NetDev Features				
1750279	Firmware Flash Update via Devlink	[ConnectX-4 and above] Added the option to update the firmware image in the flash using the devlink tool. Usage: devlink dev flash <dev> file <file_name>.mfa2	iproute2 v5.3	xx.25.1020
1747672	Firmware Versions Query via Devlink	[ConnectX-4 and above] Added the option to query for running and stored firmware versions using devlink tool.	iproute2 v5.3	N/A
1706472	Generic Segmentation Offload (GSO) Optimization	[ConnectX-4 and above] Improved GSO workload performance by decreasing doorbells usage to the minimum required.	N/A	N/A
1704398	Devlink Firmware Fatal Errors Handling	[ConnectX-4 and above] Added support for the driver to detect firmware fatal errors and initiate CR-Dump and recovery flow.	iproute2 v5.3	N/A
1700834	AF_XDP Socket Zero Copy	[ConnectX-4 and above] Added zero-copy support for AF_XDP sockets to improve performance.	N/A	N/A
1700781	Devlink Health CR-Space Dump	[ConnectX-4 and above] Added the option to dump configuration space via the devlink tool to improve debuggability.	iproute2 v5.3	N/A
1610373	TLS TX Hardware Offload	[ConnectX-6 Dx] Added support for hardware offload encryption of kTLS traffic to improve performance. Note: This is an infrastructure feature only, to enable future support for ConnectX-6 Dx adapter cards.	N/A	xx.26.1040
1186273	Devlink Health Firmware Core Dump	[ConnectX-4 and above] Added support for dumping firmware core data via the devlink tool to improve debuggability. Note that to allow firmware core dumping, FW Tracer should be enabled.	iproute2 v5.3	N/A

ASAP ² Features				
1767550	Dual-Port RoCE Support	[ConnectX-5 and above] Enabled simultaneous operation of dual-port RoCE and Ethernet in SwitchDev mode.	N/A	xx.25.6000
1765208	Geneve Encap /Decap Rules Offload	[ConnectX-5 and above] Added support for Geneve encapsulation/decapsulation rules offload.	N/A	xx.25.1020
1693782	VLAN PUSH /POP Offload (VGT)	[ConnectX-5 and above] Allowed offload of VLAN push/pop operations on both transmitted and received packets.	N/A	xx.25.1020

Upstream Kernel 5.2				
Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1688852	DevX Representors	[ConnectX-4 and above] Changed InfiniBand representors into a single InfiniBand device with multiple ports, where each port represents a different virtual function (port 1 is for the uplink), to allow DevX control by exposing a further uniformed software model.	N/A	xx.25.1020
1706795	Userspace Software Steering for eSwitch	[ConnectX-5] Added software steering capabilities to the SR-IOV eSwitch. Software steering enables better rules insertion rate compared to the current firmware-based solution. This is achieved by performing calculations on the main CPU which allows for higher insertion rates.	rdma-core v24	xx.25.1020
1686065	Userspace Software Steering for NICs	[ConnectX-5] Added software steering capabilities to NIC Rx/Tx. Software steering enables better rules insertion rate compared to the current firmware-based solution. This is achieved by performing calculations on the main CPU which allows for higher insertion rates. This solution was designed to work with Virtio DPDK. Note: Support will be enabled by default once the support for GID change is added.	rdma-core v24	xx.25.1020
rdma-core Only Features ^[1]				
1677617	Dynamic Connection (DC) Data Path	[ConnectX-4 and above] Added DC QP data-path support using new Send APIs introduced in Direct Verbs (DV).	rdma-core v24	N/A
1686034	Send APIs	[ConnectX-4 and above] Introduced a new set of QP Send operations (APIs) that allow extensibility for new Send opcodes.	rdma-core v24	N/A
1705623	User-Mode Memory Registration (UMR)	[ConnectX-4 and above] Enabled registration of memory patterns that can be used for future RDMA operations.	rdma-core v24	xx.25.1020

[1] These features require rdma-core update only - no kernel update needed.

NetDev Features				
1705845	XDP Inline Transmission of Small Packets	[ConnectX-5 and above] Added support for when forwarding packets with XDP, a packet smaller than 256 bytes would be sent inline within its WQE Tx descriptor for better performance. The number of packets that are transmitted inline depends on CPUs load, where lower load leads to a higher number of inline transmission.	N/A	N/A
1706862	EEPROM Device Thresholds via Ethtool	[ConnectX-4 and above] Added support to read additional EEPROM information from high pages of modules such as SFF-8436 and SFF-8636. Such information can be: 1. Application Select table 2. User writable EEPROM 3. Thresholds and alarms - Ethtool dump works on active cables only (e.g. optic), but thresholds and alarms can be read with "offset" and "length" parameters in any cable by running: <code>ethtool -m <DEVNAME> offset X length Y</code>	N/A	N/A

1699196	PCIe Power State	<p>[ConnectX-6] Added support for the following PCIe power state indications to be printed to dmesg:</p> <ol style="list-style-type: none"> 1. Info message #1: PCIe slot power capability was not advertised. 2. Warning message: Detected insufficient power on the PCIe slot (xxxW). 3. Info message #2: PCIe slot advertised sufficient power (xxxW). <p>When indication #1 or #2 appear in dmesg, user should make sure to use a PCIe slot that is capable of supplying the required power.</p>	N/A	xx.25.1020
1535442	Multiple Outstanding UMR Posts	<p>[ConnectX-4 Lx and above] Added support for multiple outstanding UMR posts in order to allow faster gap closure between consuming Multi-Packet WQEs (MPWQEs) and re-posting them back into the WQ. This improves Rx packet rate on Striding RQ as it solves scaling and stability issues.</p>	N/A	N/A
1703927	GENEVE Tunnel Stateless Offload	<p>[ConnectX-4 and above] Added support for Generic Network Virtualization Encapsulation (GENEVE) tunneled hardware offload of TSO, CSUM and RSS.</p>	N/A	N/A
ASAP² Features				
1699405	VLAN Rewrite	<p>[ConnectX-5 and above] Added support for offloading VLAN ID modify operation, allowing the user to replace the VLAN tag of the incoming frame with a user-specified VLAN tag value.</p>	N/A	xx.25.1020
1699042	VLAN Support over Uplink	<p>[ConnectX-5 and above] Added support for offloading TC rules from VF representor and forwarding them to a VLAN device over the physical function (PF), and vice-versa.</p>	N/A	N/A
Upstream Kernel 5.1				
Internal Reference Number	Feature/Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1678107	SRQ and XRC Support on On Demand Paging (ODP) Memory Region (MR)	<p>[ConnectX-4 and above] Added support for using ODP MR with SRQ WQEs and XRC transport.</p>	rdma-core v23	xx.25.1000
1698300	PCI Atomic Operations	<p>[ConnectX-5 and above] Added the ability to run atomic operations on local memory without involving verbs API or compromising the operation's atomicity.</p>	rdma-core v23	N/A
1685107	Indirect Mkey ODP	<p>[ConnectX-4 and above] Added the ability to create indirect Mkeys with ODP support over DevX interface.</p>	N/A	xx.25.1000
1698230	DevX Asynchronous Query Commands	<p>[ConnectX-4 and above] Added support for running QUERY commands over the DevX interface in an asynchronous mode. This enables applications to issue many commands in parallel while firmware processes the commands.</p>	rdma-core v23	xx.25.1000
1686056	RDMA-CM QP Timeout Control	<p>[All] Added a new option to rdma_set_option that allows applications to override the RDMA-CM's QP ACK timeout value.</p>	rdma-core v23	N/A
1647527	Object IDs Exportation	<p>[All] Added a unique ID for each verbs object to allow direct query over rdma-tool and rdma-netlink for enhanced debuggability.</p>	iproute v5.1	N/A
rdma-core Only Features^[1]				
1698278	Implicit ODP	<p>[ConnectX-4 and above] Added support for reporting implicit ODP support to user applications in order to allow better granularity over ODP creation.</p>	rdma-core v23	N/A
1698290	RDMA-CM Application Managed QP	<p>[All] Added support for the RDMA application to manage its own QPs and use RDMA-CM only for exchanging Address information.</p>	rdma-core v23	N/A
1680025	mlx4 Internal Resources Sharing	<p>[ConnectX-3/ConnectX-3 Pro] Added support for a new API for sharing internal resources, such as UAR and memory allocation, which will enable higher layers, such as DPDK, to implement a secondary process.</p>	rdma-core v23	N/A
NetDev Features				

1700770	Devlink Health Utility	<p>[ConnectX-4 and above] Added support for real-time alerting of functionality issues that may be found in a system component (reporter). This utility helps detect and recover from a problem with a PCI device. It provides a centralized status of drivers' health activities in the generic Devlink instance and inter alia, supports the following:</p> <ul style="list-style-type: none"> • Storing real-time error dumps • Performing automatic (configurable) real-time reporter recovery • Performing real-time reporter diagnosis • Indicating real-time reporter's health status • Providing admins with the ability to dump, diagnose and recover a reporter • Providing admins with the ability to configure a reporter 	iproute v5.1	N/A
1698287	Ethtool 200Gbps	<p>[ConnectX-6 and above] ConnectX-6 hardware introduces support for 200Gbps and 50Gbps-per-lane link mode. The driver supports full backward compatibility with previous configurations. Note that in order to advertise newly added link-modes, the full bitmap related to the link modes must be advertised from ethtool man page.</p> <p>NOTE: This feature is firmware-dependent. Currently, ConnectX-6 Ethernet firmware supports up to 100Gbps only. Thus, this capability may not function properly using the current driver and firmware versions.</p>	ethtool v5.1 (tentative)	xx.25.1000
1650372	Virtual Ethernet Port Aggregator (VEPA)	<p>[ConnectX-4 and above] Added support for activating/deactivating Virtual Ethernet Port Aggregator (VEPA) mode on a single virtual function (VF). To turn on VEPA on the second VF, run:</p> <pre>bridge link dev <netdev> hwmode vepa</pre>	N/A	N/A
1656472	ConnectX Device IDs	<p>Added support for the following new device IDs:</p> <ul style="list-style-type: none"> • ConnectX-6 Dx (PF) • ConnectX Family mlx5Gen Virtual Function (VF) <p>Note that every new device (adapter) VF will be identified with this device ID. Different VF models will be distinguished by their revision ID.</p>	N/A	N/A

ASAP² Features

1700775	Equal Cost Multi-Path (ECMP)	[ConnectX-4 Lx and above] Added support for offloading ECMP rules by tracking software multipath route and related next-hops, and reflecting this as port affinity to the hardware.	N/A	N/A
1600548	GRE Entropy Control	[ConnectX-4 Lx and above] Added support for controlling port tunnel entropy calculation in order to set the correct GRE key value.	N/A	xx.25.1000

BlueField Features

1351321	Respsenators for Host Physical Functions	[BlueField] When ECPF is the E-switch manager, ECPF driver will create a representor for the host PF in Offload mode in order to ensure secure control of the host PF traffic in a bare-metal environment.	N/A	18.24.1000
1351323	ICM Pages Allocation on Host PFs in Non Secure Mode	<p>[BlueField] Added the ability to use host memory in order to increase scalability in virtual cloud scenarios where the host PF is trusted. ICM page allocation is now managed in firmware as follows:</p> <ol style="list-style-type: none"> 1. Function provides pages for itself 2. PF provides pages for its VF 3. ECPF provides pages to itself and for another function 	N/A	18.24.1000

Upstream Kernel 5.0

Internal Reference Number	Feature /Change	Description	Support Added in User Space Version	Support Added in Firmware Version
RDMA Features				
1474608	ODP Pre-fetch	[ConnectX-4 and above] Added support for pre-fetching a range of an on-demand paging (ODP) memory region (MR), this way reducing latency by making pages present with RO/RW permissions before the actual IO is conducted.	rdma-core v22	N/A
1470521	HDR Link Speed Exposure	[ConnectX-6 and above] Added support for HDR link speed in CapabilityMask2 field in port attributes.	rdma-core v22	xx.25.1000
1470528	QP Packet Based Credit Mode	[ConnectX-6 and above] Added support for an alternative end-to-end credit mode for QP creation. Credits transported from the responder to the requester are now issued per packet. This is particularly useful for sending large RDMA messages from HCA to switches that are short in memory.	rdma-core v22	xx.24.1000

1463974	Fragmented QPs Buffer	[ConnectX-4 and above] Added the ability to allocate a fragmented buffer to in-kernel QP creation requests, in cases of large QP size requests that used to fail due to low memory resources on the host.	N/A	N/A
1498575	Flow Counters Batch Query	[ConnectX-4 and above] Allowed flow counters created with the DevX interface to be attached to flows created with the raw flow creation API.	rdma-core v22	N/A
1550997	DevX Privilege Enforcement	[ConnectX-4 and above] Enforced DevX privilege by firmware. This enables future device functionality without the need to make driver changes unless a new privilege type is introduced.	N/A	xx.25.1000
1471289	DevX Interoperability APIs	[ConnectX-4 and above] Added support for modifying and/or querying for a verb object (including CQ, QP, SRQ, WQ, and IND_TBL APIs) via the DevX interface. This enables interoperability between verbs and DevX.	rdma-core v22	xx.25.1000
1519601	Restored Verbs Fork Support	Re-added support for libverbs users to open ibv_context using ibv_open_device(), call fork and continue to use that context from the fork.	rdma-core v22	N/A
1551627	rdma-core Python Bindings (Pyverbs)	Added support for Python API over IB verbs. Pyverbs package shortens development time and provides a comfortable API to work with.	rdma-core v22	N/A
NetDev Features				
1558262	CQE Padding	[ConnectX-5 and above] Added support for padding 64B CQEs to 128B cache lines to improve performance on 128B cache line systems, such as PPC.	N/A	xx.21.1000
1464598	Counters Monitoring	[ConnectX-4 and above] Added support for monitoring selected counters and generating a notification event (Monitor_Counter_Change event) upon changes made to these counters. The counters to be monitored are selected using the SET_MONITOR_COUNTER command.	N/A	xx.24.1000
1437218	Rx Hash Fields Configuration	[ConnectX-4 and above] Added the ability to configure Rx hash fields used for traffic spreading into Rx queues using ETHTOOL_SRXFH and ETHTOOL_GRXFH ethtool commands. Built-in Receive Side Scaling (RSS) profiles can now be changed on the following traffic types: UDP4, UDP6, TCP4 and TCP6. This configuration affects both outer and inner headers.	ethtool v2.6.33	N/A
1190960	XDP Multi-Packet Tx Work Queue Element (WQE)	[ConnectX-5 and above] Added support for Multi-Packet Tx WQEs in XDP transmit flows to work with a new and improved WQE layout that describes several packets. This saves PCI bandwidth and transactions, and improves transmit packet rate.	N/A	N/A
ASAP² Features				
1082182	VF LAG	[ConnectX-4 Lx and above] Added support for High Availability and load balancing for Virtual Functions of different physical ports in SwitchDev SR-IOV mode.	N/A	N/A
1441639	Remote Mirroring	[ConnectX-5 and above] Added support for additional mirroring output in SwitchDev mode. The mirroring port may either be a local or a remote VF, using VxLAN or GRE encapsulations.	OVS v2.10.90	xx.24.1000
	GRE HW Offloading	[ConnectX-4 Lx and above] Added the ability to offload TC filters set on GRE interfaces to NICs that support ASAP ² .	- OVS v2.10.90 - iproute v5.0	xx.24.1000
1592719	VxLAN/GRE Tunneling over VLAN	[ConnectX-4 Lx and above] Added support for VxLAN and GRE tunnel encap/decap offload over Ethernet tagged packets.	N/A	N/A
1306741	Uplink Representors	[ConnectX-4 Lx and above] Exposed PF (uplink) representors in SwitchDev mode, similarly to VF representors, as an infrastructure improvement for SmartNICs.	N/A	N/A