

# Mellanox FlexFlow™



## Programmable Networking without Compromise

High-performance Ethernet switches make up the fabric of modern data centers, forwarding traffic at multi-terabit speeds. The changing world of virtualized and self-healing infrastructure requires switches that have dynamic, programmable capabilities in their DNA.

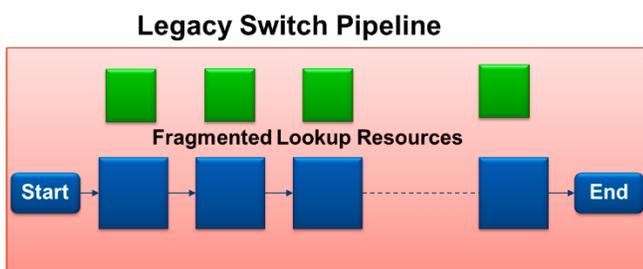
Mellanox FlexFlow™ is a set of advanced network functions that future-proof the switch, enabling custom applications, advanced telemetry, and new tunneling/overlay capabilities. FlexFlow combines a programmable, flexible and massively parallel packet processing pipeline with a fully shared & stateful forwarding database.

FlexFlow is implemented in the Spectrum® family of Ethernet switches and supports the following programmable capabilities:

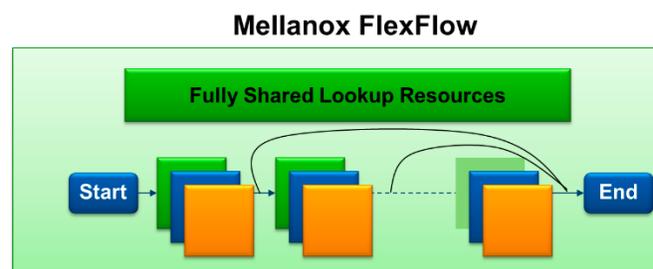
- Deep (512B) Packet Parsing
- Programmable Packet modification
- Flexible, Future-proof Tunneling
- Efficient and Fully Shared Lookup Tables
- Flow specific rules
- Comprehensive Telemetry
- Custom applications

## PROGRAMMABLE NETWORKING

- Ultra-Deep Packet Parsing – 512B
- Flexible Packet Modification
- Flow Specific Network Abstraction
- Future-Proof Tunneling
- Fully Shared Forwarding Database
- Customizable Telemetry
- Customer specific applications



- Rigid stages limit functionality
- Serialization restricts overall throughput
- High latency due to serialization
- Inefficient fragmented lookup resources
  - Lower scale
  - Higher power



- Programmable Parser, Editor & Stages
- Massively parallel to maximize throughput
- Lower Latency with per flow path optimization
- Efficient usage of unified lookup resources
  - Higher scale
  - Lower power

Figure 1: Legacy Switch Pipeline vs Mellanox FlexFlow

## Deep Parsing Packet

The parser is the first block in the switch packet- processing pipeline. This block is responsible for interpreting the incoming bytes into meaningful packet header fields that can be used for forwarding, policy enforcement, and QoS.

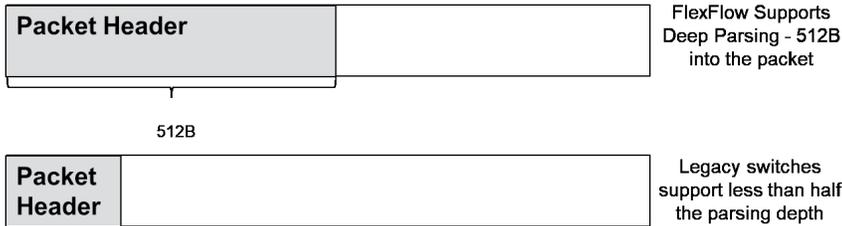


Figure 2: Deep Packet Parsing

Traditional switches expect packets to follow a pre-determined syntactically correct packet format. Packet parsing is halted when packets do not conform to the syntax. FlexFlow supports programmable packet grammar which allows custom modification of the packet grammar and addition of new packet formats. FlexFlow parses up to 512B deep into the packet 3-to-4 times deeper than other switches in the market. With greater parsing depth and programmable packet formats, Spectrum switches can support richer tunneling, advanced telemetry features while future proofing the infrastructure.

## Flexible and Flow-Specific Lookup

The FlexFlow programmable pipeline can support a flexible number of lookup tables. The keys for the lookup can be configured on a per flow basis. Table lookup actions include standard forwarding/policy actions as well as a pointer to the next look table(s). As a result, the number of tables, lookup sequence, matching keys and actions can be customized on a per-flow basis. Packet flows are also allowed to access the same tables multiple times. With flow level lookup sequence, FlexFlow can enable sophisticated network abstractions over a common underlying physical fabric.

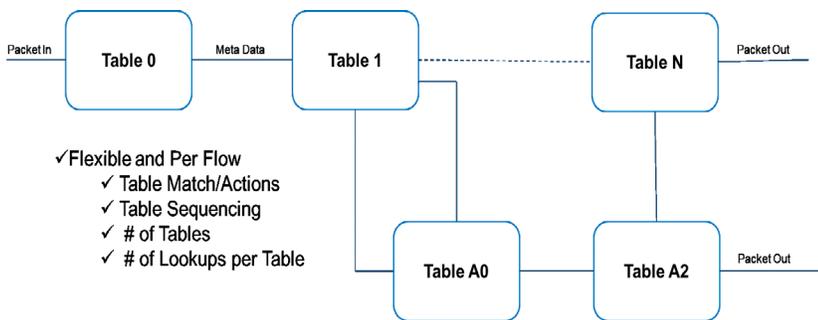


Figure 3: Flexible and Flow Specific Lookup

## Advanced Tunneling

Mellanox Spectrum flexible packet editing engine can programmatically add, modify, and delete multiple layers of packet headers. FlexFlow supports programmable encapsulation/decapsulation, IPv4 Options, IPv6 Extensions and a variety of overlay protocols including VXLAN, VXLAN-GPE, MPLS-over-GRE/UDP, NSH, NVGRE and MPLS/IPv6 based Segment routing. With FlexFlow, Spectrum Ethernet switches are ready for future tunneling and overlay technologies that are at the root Network Virtualization in the years to come.

## Fully Shared Lookup Tables

Legacy architectures use fragmented forwarding tables that are hardwired to specific pipeline stages. In such architectures, the table resources and lookup budget available for a stage is essentially hard-coded. Any table resources that are not consumed in a specific pipeline stage are wasted. If more lookup table resources are needed, more stages need to be used to accomplish the same task which reduces the number of features that can be simultaneously supported.

FlexFlow not only allows more efficient table sharing but also provides practically unlimited lookup budget. With this high lookup budget, Spectrum switches can support all the packet transformations that are needed for modern, emerging and future network protocols.

## Comprehensive Telemetry

In addition to packet transformations, Mellanox FlexFlow pipeline uses flexible and stateful pipeline stages to programmatically extract packet meta-data and provide real-time granular telemetry. As data center networks grow in complexity, Mellanox What Just Happened™ along with FlexFlow provide built-in and fully-integrated telemetry functionality to reduce time issue resolution, improve uptime and better optimize infrastructure utilization.

## Summary

Mellanox FlexFlow programmable and massively parallel packet processing pipeline performs at line rate with hardware-accelerated packet transformations to enable agile, software-defined data center networks. FlexFlow uses a programmable deep packet parser to effectively use a larger portion of the packet header to make forwarding and policy decisions. FlexFlow provides higher scale, feature concurrency, and flexibility with efficient sharing of on-chip lookup resources. FlexFlow futureproofs the infrastructure with support for programmable packet editing and advanced tunneling.

---

## About Mellanox

Mellanox Technologies (NASDAQ: MLNX) is a leading supplier of end-to-end Ethernet and InfiniBand intelligent interconnect solutions and services for servers, storage, and hyper-converged infrastructure. Mellanox intelligent interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance. Mellanox offers a choice of high performance solutions: network and multi-core processors, network adapters, switches, cables, software and silicon, that accelerate application runtime and maximize business results for a wide range of markets including high performance computing, enterprise data centers, Web 2.0, cloud, storage, network security, telecom and financial services. More information is available at: [www.mellanox.com](http://www.mellanox.com)



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085  
Tel: 408-970-3400 • Fax: 408-970-3403  
[www.mellanox.com](http://www.mellanox.com)