

**Virtual  
Modular  
Switch**

# Mellanox Virtual Modular Switch®

## Simplified Network Operations for Leaf/Spine Networks

As new applications are constantly evolving, data centers must be flexible and future proof to meet the demand for higher throughput and higher scalability while protecting the original capital investment and preventing an increase in operating costs.

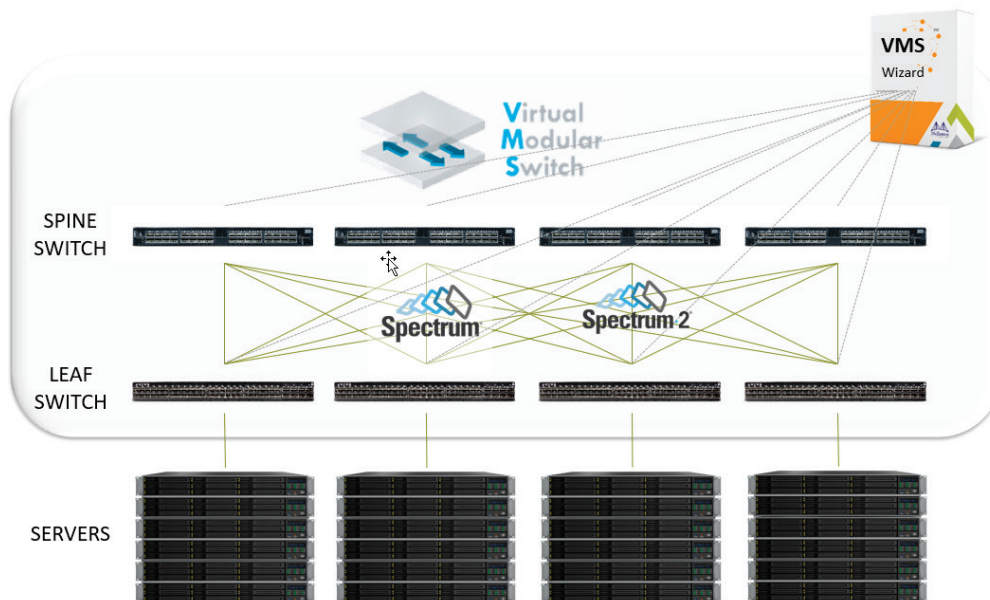
Traditionally, aggregation switching in the data centers of cloud providers, Web 2.0 providers, and large-scale enterprises has been based on modular switches. These switches are usually expensive and optimized for specific sizes of clusters and fail to provide the flexibility required in today's scale-out data centers.

To overcome this flexibility limitation in modular switches, users are shifting to high-density fixed switches. The Mellanox Virtual Modular Switch® (VMS) solution comprised of Mellanox fixed form-factor Spectrum® Ethernet switches provides an ideal, optimized approach for building data center networks. When compared to legacy modular systems, VMS excels with higher flexibility, better scalability, and energy efficiency, while future-proofing the investment and reducing expenses. VMS Wizard helps automate, orchestrate and validate the switch configurations necessary for leaf/spine data center networks based on Mellanox Spectrum Ethernet switches.

## HIGHLIGHTS

Advantages of using Mellanox VMS with Mellanox Spectrum fixed Ethernet switches:

- Control plane scalability
- Increased visibility and control
- Smaller fault domain
- Facilitates incremental growth
- Lower cost



**Figure 1:** VMS with Spectrum® Ethernet Switch-based Leaf/Spine Networks

## Leaf/Spine Networks with Fixed Ethernet Switches

Modern data center networks have standardized on fixed form-factor switch based leaf/spine networks.

The following are key design principles for a 2-tier leaf/spine design:

- All leaf switches are connected to all spine switches via uplinks. The maximum number of spines is determined by the number of uplink ports in the leaves.
- Spines are not connected to each other.
- Servers and endpoints can be single or dual homed but connect *only* to leaves.

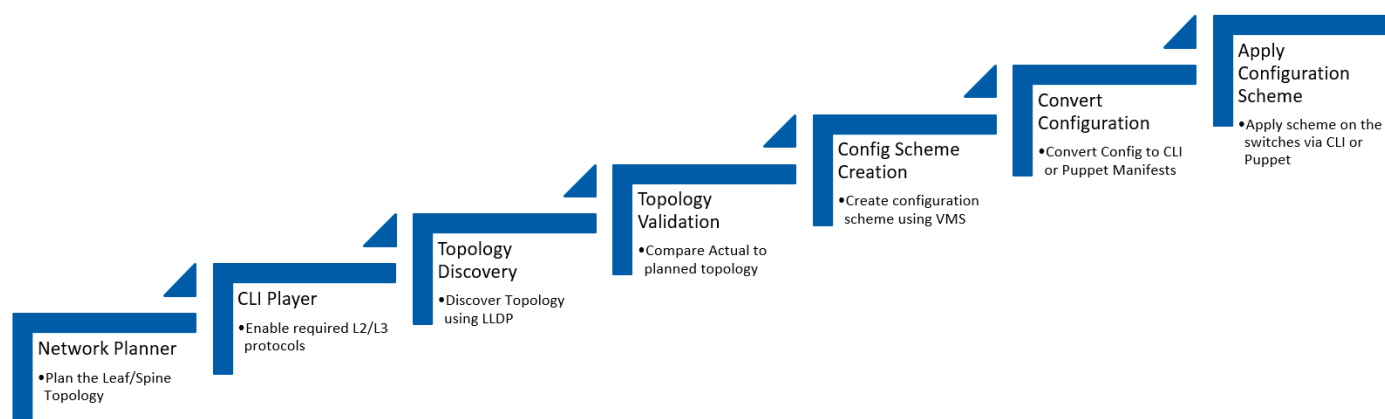
The leaf/spine topology can scale from a small infrastructure comprising of a few racks to a large infrastructure spanning several racks. Layer-3 based Equal Cost Multi-Paths (ECMP) between the leaf and spine switches is leveraged to provide high cross-sectional bandwidth. Mellanox has a strong portfolio of fixed form factor and high-performance Ethernet switches that can be used with VMS to build an automated, efficient and high-performance leaf/spine network.

Below is a summary of the advantages of using VMS with Mellanox Spectrum fixed Ethernet switches:

- Lower Cost:** Fixed switches typically cost 75% less than modular switches.
- Control Plane Scale:** Control plane scales linearly as there is no central performance bottleneck, like the supervisor card in a modular system. Each fixed-form factor switch has a local CPU and a control plane.
- Visibility & Control:** All traffic paths including the packets that flow between leaf and spine switches can be monitored. In contrast, modular systems are like a black box with no visibility into the packet flows between the line cards and the fabric cards.
- Smaller Fault Domain:** Fixed switch failures will affect only a small portion of the network. Big modular system failures can be catastrophic and can bring down a significant portion of the data center network.
- Incremental Growth:** Leaf/spine networks are easy to scale in small increments with switches added as additional capacity is needed.

## VMS Configuration and Management

VMS Wizard automation software generates a connectivity plan from simple user input through a 7-step process. After installation, it configures the fabric switches through Puppet™ manifests or CLI commands, and performs physical, layer 2, and layer 3 connectivity validations.




**Figure 2:** VMS Wizard 7-Step Process to Configure Leaf/Spine Networks

Mellanox VMS is practically a small-scale IP network and, therefore, is managed and configured using standard IP protocols. These protocols contribute to the ease-of-use of VMS and provide a robust, standard, and resilient solution. Each of the VMS switches runs dynamic routing protocols. The routing protocols gather link state information from the switches (routers) in the data center and constructs a topology map of the network to determine the IP routing tables. By building the routing tables, routing loops are prevented. It automatically detects topology changes such as link failures and updates the routing structure accordingly within seconds.

To complement the dynamic routing decisions, the switches also run Equal-Cost Multi-Path (ECMP) routing, which provides load balancing between routes of equal cost. By leveraging ECMP, congestion within the VMS and in the data center is avoided. Configuration of the VMS infrastructure is performed by the VMS Wizard automation software and applications such as Puppet (<http://www.puppetlabs.com/>). These applications automate the configuration and provisioning of the switches, including software upgrades, VLAN provisioning, port configuration, routing, and provide an easy way to scale the VMS. After configuration, the VMS Wizard and Puppet can report errors and mismatches.

## Summary

Legacy modular solutions are expensive and impose strict scale constraints. Leaf/spine topologies are becoming the de facto way of building modern data center networks. Mellanox VMS provides a simple, open and standards-based mechanism to automate leaf/spine networks that are based on the Mellanox Spectrum family of Ethernet switches. VMS can be used both for day-1 setup and fabric validation as well as to automate day-2 and ongoing operational configurations and monitoring.



## 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)