Automated Leaf-Spine Network

Solution Brief

Introduction

With the advent of big data applications and virtualization of workloads and storage, traffic patterns within data centers are shifting from predominantly north-south (outside the data center) to east-west (machine-to-machine within the data center). Network architects now face the challenge of taking data centers that were built with a three-tier design optimized for north-south traffic and retrofitting them with leaf-spine designs that provide more robust support for still-growing volumes of east-west traffic. One of the posterchildren for this trend of more and more east-west traffic is Facebook. In 2014 it was estimated that their east-west traffic levels were 1000 times higher than their north-south traffic levels.

Leaf-spine architecture is a useful approach to the problem, but like all solutions it has its own pain points. The dense interconnections that provide robust support for east-west traffic increase the complexity of adding capacity or moving equipment from one part of the structure to another. Glass Core solutions are designed to make physical network structures more dynamic through software control and automation.

Leaf-Spine 101

The concept of a leaf-spine network originated from a telecommunication network design called Clos, named after the scientist Charles Clos. Instead of aggregating devices into hierarchical tiers with a single path out of base-level devices, a leaf-spine architecture provides multiple paths to the same point, thus improving the throughput and latency of the network. A leaf-spine network typically looks like this:

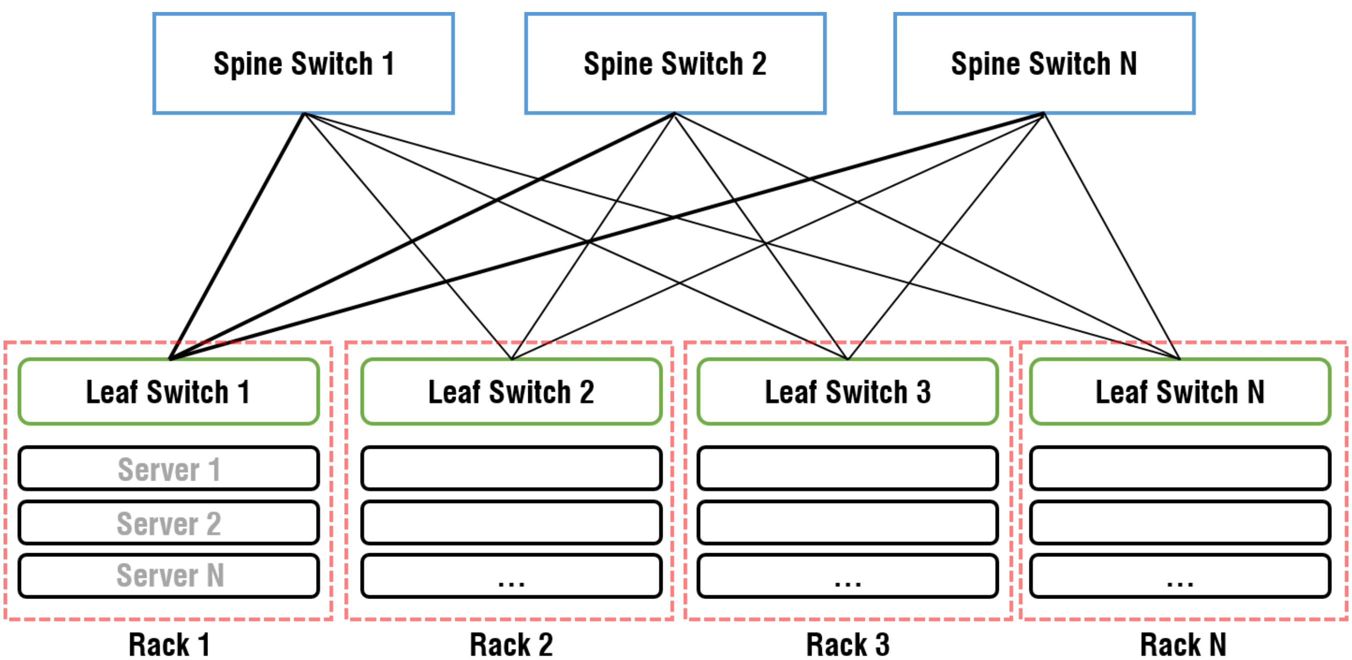

As the previous diagram shows, the physically connectivity in a leaf-spine network is based on two key elements:

1. All servers in a rack connect to the top of rack (TOR) leaf switch.
2. All leaf switches connect to all spine switches.

With this connection pattern, any east-west traffic can be facilitated by one spine hop, and there are multiple paths between any two points.

Pain Points - Cabling

One of the major pain points of the leaf-spine architecture is the need to change physical cabling when capacity is added or removed from the architecture. Adding and removing cables in a leaf-spine architecture can be prohibitively time consuming or sometimes physically impossible, depending on equipment locations, physical cabling already in place, remaining cable tray or conduit capacity and how accurately the physical network structure has been documented. Here are two examples:

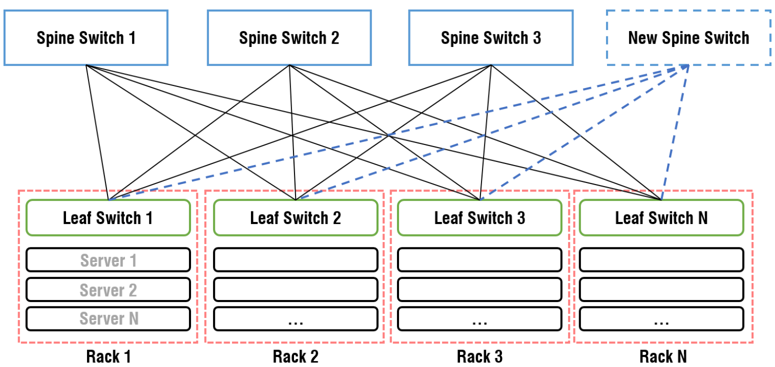

When a new Spine Switch is added, it requires the addition of a new physical connection to each existing leaf switch.

Likewise, when a new Leaf Switch is added, it requires the addition of a new physical connection to each existing spine switch.

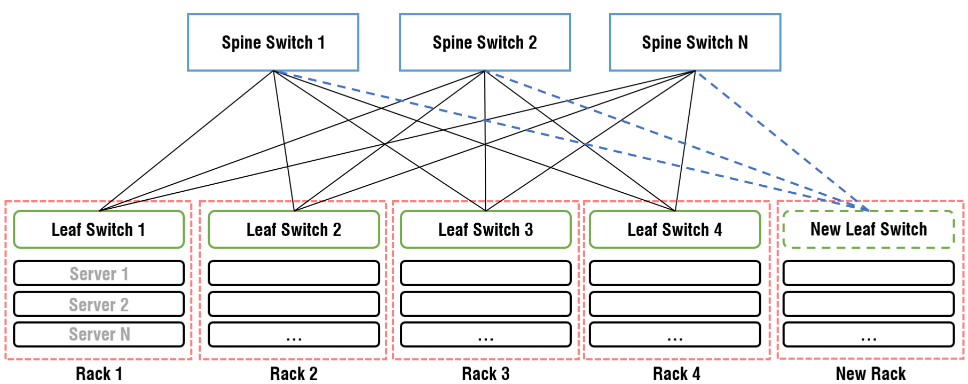

Automated Leaf-Spine Network

Solution Brief

FIBER MOUNTAIN™
The Glass Core™ Company

Fiber optic cables take up less space than copper cables for 10, 25 and 100 Gbps network connections, but the extra connection density they enable makes the sheer volume of connections unmanageable with manual systems. Automation and software management of physical connections is an essential component in the race to support ongoing data center growth.

Pain Points – Architecture Change

Leaf-spine architectures come in many flavors. For example, a large network may connect spine switches to super-spine switches like this:

These super spine switches can be used to connect multiple spine-leaf pods and extend the network structure beyond current capacity for a full connection fabric.

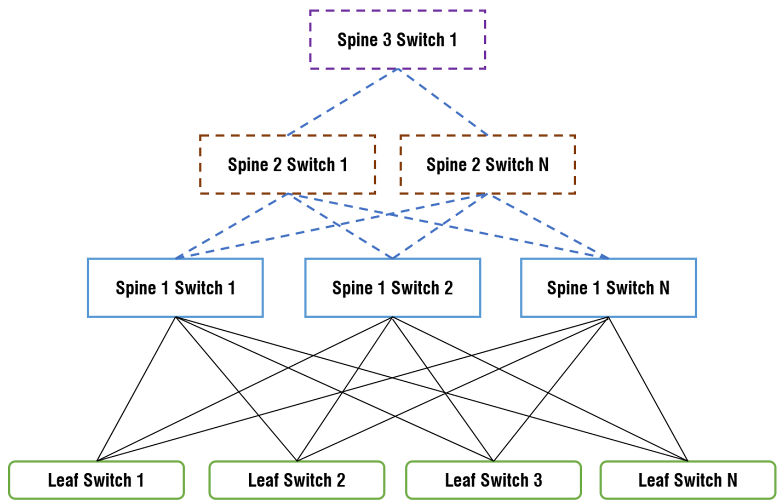

The Facebook design published November 14, 2014 achieved this goal by mapping leaves and spines across a third plane, in a system of consistently structured pods. (https://code.facebook.com/posts/360346274145943/introducing-data-center-fabric-the-next-generation-facebook-data-center-network/)

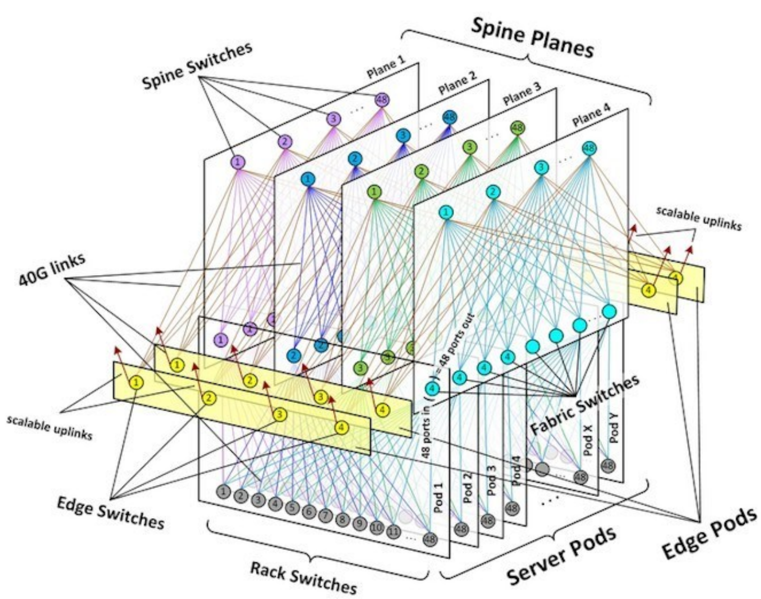

They managed the complexity of cabling such systems through uniformity of design from pod to pod, but how does an enterprise with varied needs and smaller technology teams manage this kind of complexity, especially when it comes to cabling?

Enterprise network operators need a scalable and automated way to design and implement a dynamic layer 1 for their leaf-spine network.

Automated Leaf-Spine Network

Solution Brief

Automating Layer 1

Fiber Mountain's Glass Core simplifies the process of changing network structures. Using software to automate management of layer one connections, Glass Core solutions allow the dynamic reconfiguration of connection paths, and therefore network configurations, whether leaf-spine, super leaf-spine, or a custom hybrid architecture that changes according to business and network needs.

In the network shown here, the physical connections from each spine switch, leaf switch, server and storage device are connected via cable to the FPAs (fiber port aggregators) and OPXs (Optical Path Exchanges) which make up the physical side of the Glass Core.

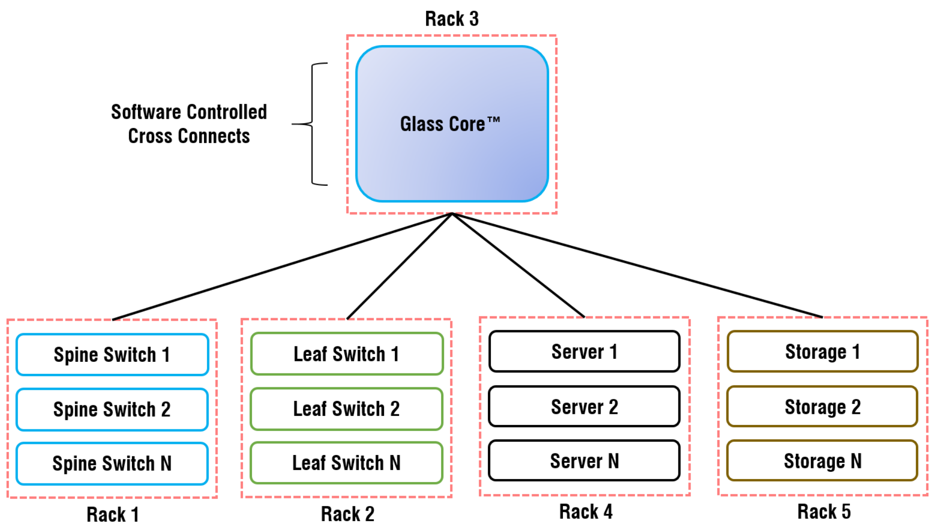

All of those connections can then be controlled via Fiber Mountain's AllPath Director orchestration software. The network architect can then configure any connection profile that is desired, and switch between different network structures on demand, without the need to manually re-run cable. Because the cross connections are software defined, the physical network can now be updated in seconds.

In addition, the Glass Core also provides the following:

1. **Security** – All physical changes are auditable
2. **Documentation** – All physical connections are auto discovered and documented
3. **Future Proof** – Can be used to configure leaf-spine, core-distribution-access or any future physical network design

Conclusion

Everyone who manages a network needs to plan for future changes in demands and available technology. Fiber Mountain's Glass Core provides much-needed flexibility, with the power to update your network architecture via software, skipping the months of cable wrangling necessary in today's networks. The Glass Core solution not only automates connectivity to enable implementation and upgrade of leaf-spine networks, it also provides physical layer security and documentation to offer network administrators unprecedented visibility and control.