


The Need for Layer 2 in Modern Data Centers

Until recently, data centers have been designed with high availability as the main priority. Just like the organizations they serve, modern networks must now be agile and accommodate changes in a flexible way. The simple answer to this additional requirement would be to increase the size of the Layer 2 domain, because switching allows moving devices and modify the infrastructure in a way transparent to servers. However, existing switching technologies have inefficient forwarding schemes based on spanning trees and cannot be extended to the network as a whole. Therefore, current designs are a compromise between the flexibility provided by Layer 2 and the scaling offered by Layer 3.

Cisco FabricPath: Introducing Routing Concepts at Layer 2


Cisco® FabricPath is an innovation in Cisco NX-OS Software that brings the stability and scalability of routing to Layer 2. The switched domain does not have to be segmented anymore, providing data center-wide workload mobility. Because traffic is no longer forwarded along a spanning tree, the bisectional bandwidth of the network is not limited, and massive scalability is now possible. Figure 1 shows some of the advantages of a network that uses Cisco FabricPath.

Figure 1 Comparison Between Traditional Data Center Design and a Cisco FabricPath Design Using the Same Networking Equipment



Cisco FabricPath introduces an entirely new Layer 2 data plane by encapsulating the frames entering the fabric with a header that consists of routable source and destination addresses. These addresses are the address of the switch on which the frame was received and the address of the destination switch to which the frame is heading. From there, the frame is routed until it reaches the remote switch, where it is deencapsulated and delivered in its original Ethernet format. Figure 2 shows this simple process.

Figure 2 Frame Transported Across a Data Center Network Using Cisco FabricPath





The addresses of the switches are automatically assigned, and a routing table is computed for all unicast and multicast destinations. The forwarding process in the fabric never resorts to flooding: frames are always forwarded using a known destination address. The resulting solution still provides the simple and flexible behavior of Layer 2, while using the routing mechanisms that make IP reliable and scalable. While the Layer 2 and 3 boundary still exists in a data center based on Cisco FabricPath, there is now little difference in the way traffic is forwarded in those two distinct areas of the network.

Cisco FabricPath Benefits

- Simplified network, reducing operating expenses
 - FabricPath is extremely simple to configure. The only necessary configuration consists of distinguishing the core ports, which link the switches, from the edge ports, where end devices are attached. No parameters need to be tuned to achieve an optimal configuration, and switch addresses are assigned automatically.
 - A single control protocol is used for unicast forwarding, multicast forwarding, and VLAN pruning. This protocol requires less combined configuration than in an equivalent network based on Spanning Tree Protocol, further reducing the overall management needed for the solution.
 - Static network designs make some assumptions about traffic patterns and the locations of servers and services. If those assumptions are incorrect, which often becomes the case after a while, complex redesign may be necessary. A fabric switching system based on Cisco FabricPath can be modified as needed in a nondisruptive manner for the end stations without any constraints on the design.

- Switches that do not support Cisco FabricPath can still be attached to the Cisco FabricPath fabric in a redundant way without resorting to Spanning Tree Protocol.
- The capabilities of Cisco FabricPath troubleshooting tools surpass those of the tools currently available in the IP community. The ping and traceroute features now offered at Layer 2 can measure latency and test a particular path among the multiple equal-cost paths to a destination within the fabric.
- Reliability based on proven technology
 - Though Cisco FabricPath offers a plug-and-play user interface, its control protocol is built on top of the powerful Intermediate System-to-Intermediate System (IS-IS) routing protocol, an industry standard that provides fast convergence and that has been proven to scale up to the largest service provider environments.
 - Loop prevention and mitigation is available in the data plane, helping ensure safe forwarding that cannot be matched by any transparent bridging technology. Cisco FabricPath frames include a time-to-live (TTL) field similar to the one used in IP, and a reverse-path forwarding (RPF) check is also applied.
- Efficiency and high performance
 - Because equal-cost multipath (ECMP) can be used the data plane, the network can use all the links available between any two devices. The first-generation hardware supporting Cisco FabricPath can perform 16-way ECMP, which, when combined with 16-port 10-Gbps PortChannels, represents bandwidth of 2.56 terabits per second (Tbps) between switches.
 - Frames are forwarded along the shortest path to their destination, reducing the latency of the exchanges between end stations compared to a spanning-tree-based solution.

- Cisco FabricPath needs to learn at the edge of the fabric only a subset of the MAC addresses present in the network, allowing massive scalability of the switched domain.

Why Cisco FabricPath?

Cisco FabricPath is an innovation in Cisco NX-OS Software that uses routing principles to allow Layer 2 scaling in a way that was not possible before. This technology provides the flexibility and the performance required to make data centers a service for the applications that use them. Cisco and its partners can help customers deploy, optimize, and maintain a dependable design based on state-of-the-art technologies such as Cisco FabricPath.