

Cisco Application Centric Infrastructure Zero-Touch Fabric



What You Will Learn

Today's data center architects and operators have many choices to make in designing and implementing the data center LAN. After decisions to meet business requirements for application agility and network performance, the most critical decisions the data center architect makes pertain to the operation of the data center LAN. Operating costs often far outweigh the actual cost of the equipment that is purchased. A solution with implicit automation and self-provisioning can offer operation advantages over traditional switching infrastructure. It is also conceptually easier to manage the data center network as a whole entity rather than to manage multiple separate switches.

Cisco® Application Centric Infrastructure (ACI) combines traditional high-performance switching technologies with advanced management and automation capabilities. It will allow our customers to accelerate application deployment, simplify operations, and treat the network as a resource pool in a way similar to the way that servers and storage resources are treated today. This approach is known as a zero-touch fabric.


Several elements of the Cisco ACI fabric are designed to provide this zero-touch operation experience:

- A logically central but physically distributed controller for policy, bootstrap, and image management
- Easy startup with topology autodiscovery, automated configuration, and infrastructure addressing using industry-standard protocols Intermediate System-to-Intermediate System (IS-IS) Protocol, Link Layer Discovery Protocol (LLDP), and Dynamic Host Configuration Protocol (DHCP)
- A simple and automated policy-based upgrade process and automated image management
- Operation simplicity and minimal or zero-touch provisioning with loose wiring validation

Cisco ACI Fabric Overview

The Cisco ACI fabric (Figure 1) is a highly scalable, multipath, high-performance leaf-and-spine architecture (bipartite graph) which provides a Virtual Extensible LAN (VXLAN) overlay for the tenant space: the network used by the business applications, departments, and customers. The Cisco ACI fabric also implements the concept of infrastructure space, which is securely isolated in the fabric and is where all the topology discovery, fabric management, and infrastructure addressing is performed. Everything discussed in this document refers to functions in the infrastructure space of the Cisco ACI fabric.

Figure 1. Cisco ACI Fabric



The Cisco ACI fabric is composed of the Cisco Application Policy Infrastructure Controller (APIC) and the Cisco Nexus 9000 Series leaf and spine switches. As in any traditional Clos or bipartite graph design, the leaf top-of-rack (ToR) switches attach to the spines, and never each other. The spines attach only to the leaf switches, and possibly to a higher-level spine if the network design is hierarchical. The Cisco APIC (and all other devices in the data center) attach to the leaf switches only.

Cisco Application Policy Infrastructure Controller

The Cisco APIC is a physically distributed but logically centralized controller that provides DHCP, bootstrap configuration, and image management to the fabric for automated startup and upgrades. The Cisco Nexus[®] ACI fabric software is bundled as an ISO image, which can be installed on the Cisco APIC appliance server through the serial console. The Cisco Nexus[®] ACI Software ISO contains the Cisco APIC image, the firmware image for the leaf node, the firmware image for the spine node, default fabric infrastructure policies, and the protocols required for operation.

Easy Startup with Cisco ACI Fabric Discovery and Configuration

The Cisco ACI fabric bootstrap sequence begins when the Cisco ACI fabric is booted with factory-installed images on all the switches. Cisco Nexus 9000 switches running ACI firmware and Cisco APICs use a reserved overlay for the boot process. This infrastructure space is hard-coded on the switches. The Cisco APIC can connect to a leaf through the default overlay, or it can use a locally significant identifier. The Cisco ACI fabric is brought up in a cascading manner, starting with the leaf node nodes directly attached to the Cisco APIC. LLDP and control-plane IS-IS convergence occurs in parallel to this boot process.

The Cisco ACI fabric uses LLDP- and DHCP-based fabric discovery to automatically discover the fabric switch nodes, assign the infrastructure VXLAN tunnel endpoint (VTEP) addresses, and install the firmware on the switches. Prior to this automated process, a minimal bootstrap configuration must be performed on the Cisco APIC.

All Cisco ACI fabric management communication within the fabric takes place in the infrastructure space using internal private IP addresses. This addressing scheme allows Cisco APICs to communicate with fabric nodes and other Cisco APICs in the cluster. Cisco APICs discover the IP address and node information of other Cisco APICs in the cluster using the LLDP-based discovery process.

Easy Fabric Upgrade

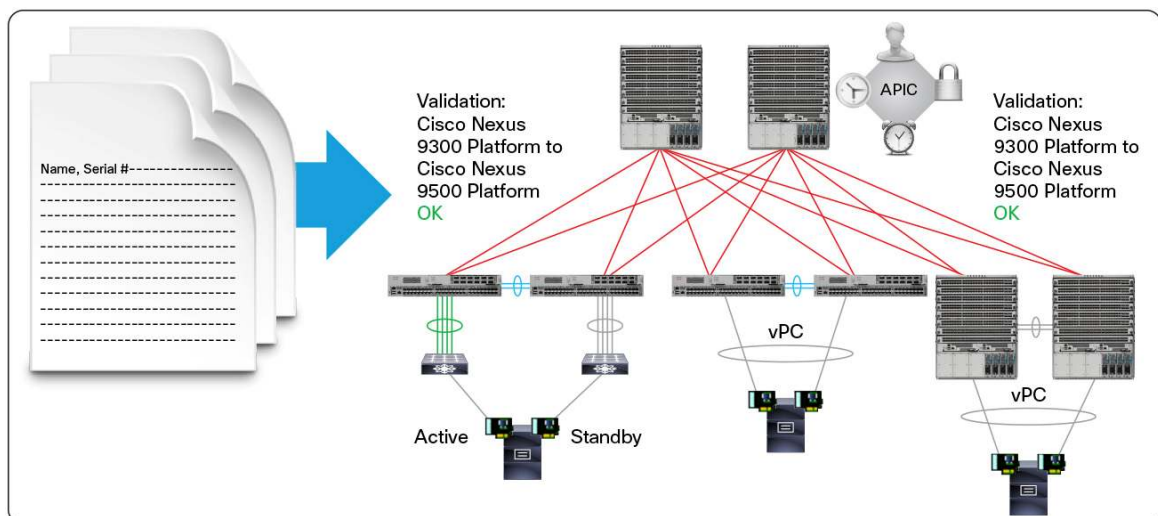
Firmware policies on the Cisco APIC define the firmware version needed on specific network nodes. Maintenance policies and schedules define when and how these firmware images are installed on the Cisco ACI fabric nodes. Maintenance policies define groups of nodes that can be upgraded together and assigns those maintenance groups to schedules. Default node groups include “All-Leaf nodes”, “All-Spine nodes,” and “All-APICs.” The maintenance schedule defines when and how many nodes in a maintenance group can be updated together. Upgrading the Cisco ACI fabric is as easy as selecting a policy and scheduling it to run.

Each firmware image has corresponding image metadata that identifies supported types and switch models. The Cisco APIC maintains a catalog of the firmware images and switch types and models allowed to use that firmware image. The Cisco APIC performs image management and has an image repository for both Cisco APIC and switch images. In the future, you will be able to download new images to the Cisco APIC from a Cisco download service.

Zero-Touch Provisioning

The Cisco APIC supports zero-touch provisioning: a method to automatically bring up the Cisco ACI fabric with the appropriate connections (Figure 2). After LLDP discovery learns all neighboring connections dynamically, these connections are validated against a loose specification rule such as “LEAF can connect to only SPINE-L1-*” or “SPINE-L1-* can connect to SPINE-L2-* or LEAF.” If a rule mismatch occurs, a fault occurs and the connection is blocked. In addition, an alarm is created indicating that the connection needs attention. The Cisco ACI fabric operator has the option of importing the names and serial numbers of all the fabric nodes from a simple text file into the Cisco APIC, or discovering the serial numbers automatically and assigning names from the Cisco APIC GUI, command-line interface (CLI), or API.

Figure 2. Zero-Touch Fabric Automation



Conclusion

Cisco ACI treats the network as a single entity rather than a collection of switches. It uses a central controller to implicitly automate common practices such as Cisco ACI fabric startup, upgrades, and individual element configuration. These innovations are designed to ease the operation of data center networks and provide a path to a fully automated fabric solution. The ultimate goal is to significantly reduce operating expenses and achieve a zero-touch fabric.

For More Information

Please visit <http://www.cisco.com/go/apic>.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)