


Application Centric Infrastructure Overview: Implement a Robust Transport Network for Dynamic Workloads



What You Will Learn

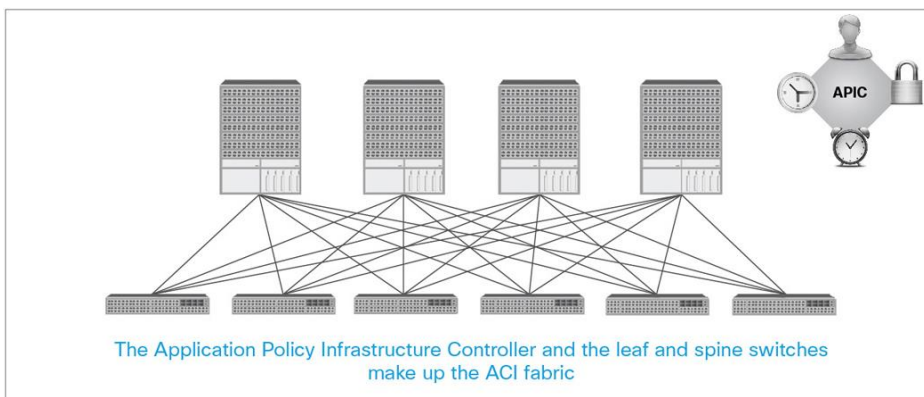
Application centric infrastructure (ACI) provides a robust transport network for today's dynamic workloads. ACI is built on a network fabric that combines time-tested protocols with new innovations to create a highly flexible, scalable, and resilient architecture of low-latency, high-bandwidth links. This fabric delivers a network that can support the most demanding and flexible data center environments.

Application Centric Infrastructure Fabric Overview

The ACI fabric is designed from the foundation to support emerging industry demands while maintaining a migration path for architecture already in place. The fabric is designed to support the industry move to management automation, programmatic policy, and dynamic "workload-anywhere" models. The ACI fabric accomplishes this with a combination of hardware, policy-based control systems, and software closely coupled to provide advantages not possible in other models.

The fabric consists of three major components: the Application Policy Infrastructure Controller, spine switches, and leaf switches. These three components handle both the application of network policy and the delivery of packets. Figure 1 shows these three components in the ACI fabric architecture.

Figure 1. ACI Fabric Architecture



In Figure 1, the fabric is designed in a leaf-and-spine architecture, with links connecting each leaf to each spine. This design enables linear scalability and robust multipathing within the fabric, optimized for the east-to-west traffic required by applications. No connections are created between leaf nodes or spine nodes because all nonlocal traffic flows from ingress leaf to egress leaf across a single spine switch. The only exceptions to this rule are certain failure scenarios.

In this architecture, the scalability of the fabric is limited only by the available ports on the spine; at least one port per leaf node is required. Bandwidth scales linearly with the addition of spine switches. Also, each spine switch added creates another network path, which is used to load-balance traffic on the fabric.

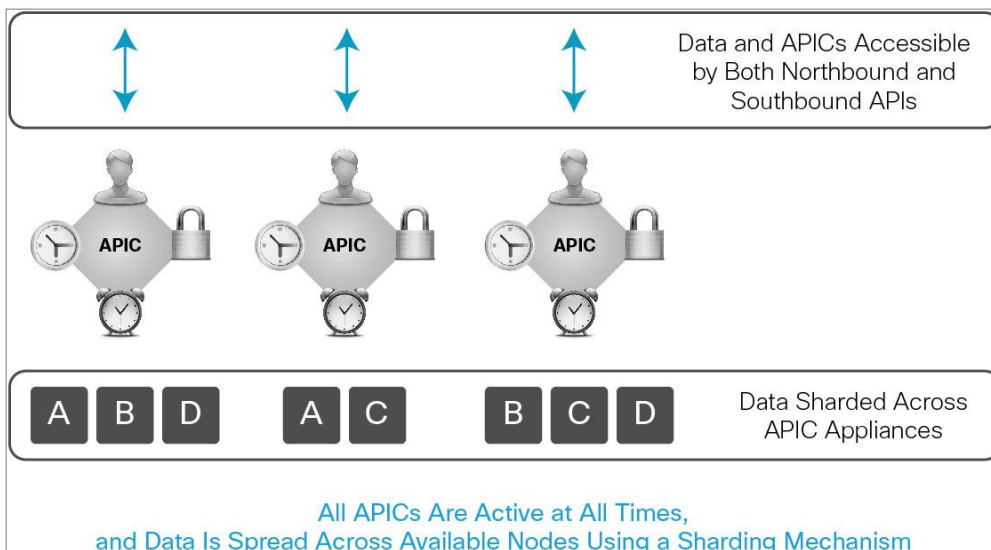
Fabric Management

The ACI fabric is designed from the foundation for programmability and simplified management. These capabilities are provided by the APIC, which is a clustered network control system. The APIC itself exposes a northbound API through XML and JavaScript Object Notation (JSON) and provides both a command-line interface (CLI) and GUI that use this API to manage the fabric. The system also provides an open source southbound API, which allows third-party network service vendors to implement policy control for supplied devices through the APIC.

The APIC is responsible for tasks from fabric activation and switch firmware management to network policy configuration and instantiation. While the APIC acts as the centralized policy and network management engine for the fabric, it is completely removed from the data path, including the forwarding topology. Therefore, the fabric can still forward traffic even when communication with the APIC is lost. The APIC itself is delivered as an appliance, and it typically is run as three or more appliances for performance and availability.

The design of the APIC is modeled on distributed computing to provide scalability and reliability that meets the needs of the data center now and in the future. Rather than using an active-standby configuration, each node is always active, processing data and accepting input. The fabric configuration data is sharded, or spread, across the appliances. Multiple copies are maintained for redundancy and performance. Figure 2 shows this clustering and sharding behavior.

Figure 2. Application Policy Infrastructure Controller Clustering and Sharding



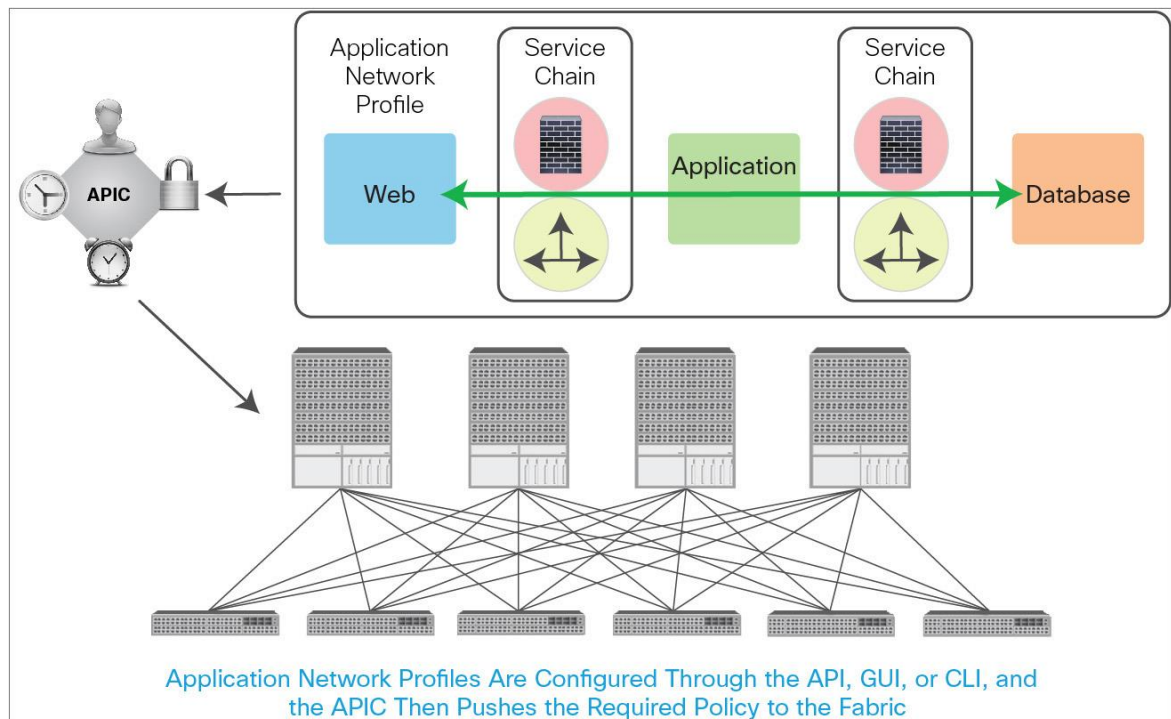
Applying Network Policy

The fabric is designed with application connectivity and policy at the core. This focus allows both traditional enterprise applications and internally developed applications to run side by side on a network infrastructure designed to support them in a dynamic and scalable way. The network configuration and logical topologies that traditionally have dictated application design are instead applied based on application needs. This approach is accomplished through the ACI object model.

Within the APIC, software applications are defined logically using constructs that are application centric, rather than network centric. For example, a group of physical and virtual web servers may be grouped in a single tier of a three-tier application. The communication between these tiers and the policies that define that communication make up the complete application. Within the APIC, this complete application definition is known as an Application Network Profile.

Application Network Profiles are defined based on the communication, security, and performance needs of the application. They are then used by the APIC to push the logical topology and policy definitions down to stateless network hardware in the fabric. This approach is the reverse of traditional architectures, in which VLANs, subnets, firewall rules, etc. dictate where and how an application can run. Figure 3 shows this behavior in the ACI fabric.

Figure 3. Application Deployment in ACI Fabric



Fabric Forwarding

The ACI fabric is designed for consistent low-latency forwarding across high-bandwidth links (40 Gbps, with 100-Gbps future capability). Traffic with the source and destination on the same leaf is handled locally, and all other traffic travels from the ingress leaf to the egress leaf through a single spine switch. Although this is a two-hop architecture from a physical perspective, it is a single Layer 3 hop because the fabric itself operates as a single Layer 3 switch. Figure 4 shows the basic forwarding within the fabric.

Figure 4. ACI Fabric Traffic Forwarding

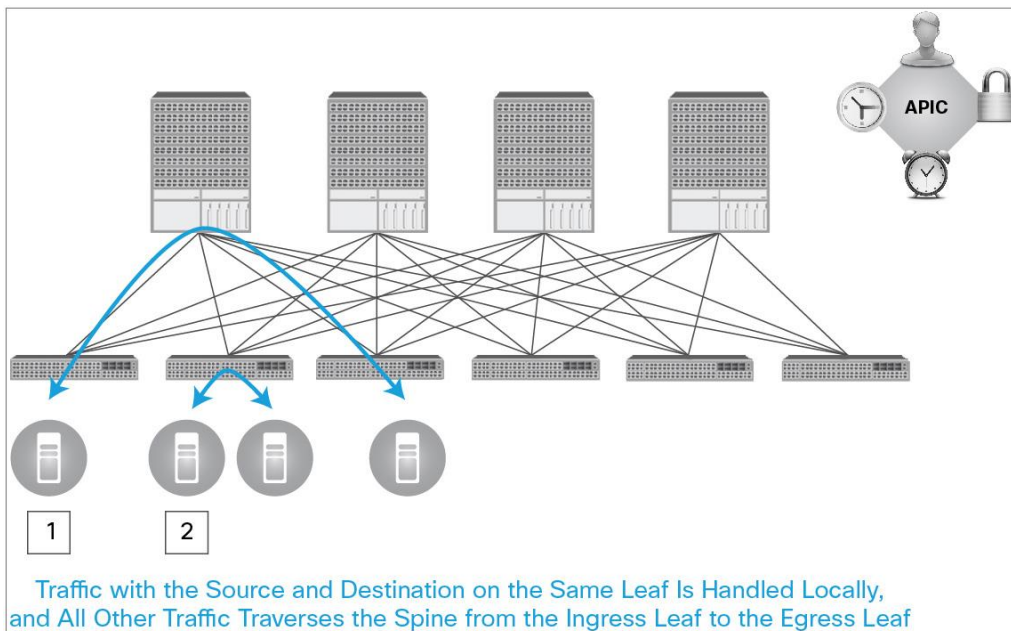


Figure 4 shows two basic forwarding behaviors. In the first example, the traffic destination is on a different leaf than the source. In this instance, a load-balancing algorithm chooses one of the spine switches to which to forward the packet. The spine then forwards the packet to the destination egress leaf. Any spine can be chosen with consistent latency across the fabric as a whole, enabling extremely efficient load balancing. The second example shows traffic in which the source and destination are on the same leaf. The traffic is forwarded locally without the need to traverse the fabric.

Conclusion

The ACI fabric uses a unique coupling of hardware and software to provide a robust set of networking features that are exceptional in the industry. Through the use of hardware aware overlays, policy-based connectivity, stateless network hardware, and an open ecosystem ACI is built for the needs of both today's workloads and tomorrow's changing demands.

For More Information

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




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)