

Is Cisco Application Centric Infrastructure an SDN Technology?

Executive Summary

Software-defined networking (SDN) has garnered much attention in the networking industry over the past few years due to its promise of a more agile and programmable network infrastructure.

The initial OpenFlow 1.1 specification was introduced in 2011, the moment many industry analysts cited as the start of the modern SDN movement. Yet production deployments of SDN using OpenFlow technology (especially in the data center) as well as software-based network overlays, are still in their infancy¹. Many IT professionals view operations, scalability, and reliability as new challenges that SDN technologies must address.

In November 2013, Cisco announced the acquisition of Insieme Networks. In doing so, Cisco also announced its strategy to address the challenges that SDN is trying to solve, through the introduction of Cisco® Application Centric Infrastructure (ACI). Cisco ACI not only addresses the challenges of current networking technologies that OpenFlow and software-based overlay networks are trying to address, but it also presents solutions to the new challenges that SDN technologies are creating.

This document analyzes the benefits and limitations of current networking technologies that are fueling the movement to adopt SDN. It also looks at how Cisco ACI can meet the new challenges facing these SDN technologies. Finally, it addresses this fundamental question: Is Cisco ACI an SDN solution?

The Need for a New Network Architecture

The Open Networking Foundation (ONF) is a user-led organization dedicated to the promotion and adoption of SDN. The ONF has published a white paper titled “Software-Defined Networking: The New Norm for Networks”². This paper describes the traditional way of building hierarchical tree-structure, or tiered, networks, and explains why this approach is not suited to the dynamic nature of modern computing and storage needs. It also presents some of the computing trends shaping the need for new network architecture, including:

- **Changing traffic patterns:** The data center is shifting away from traditional client-server application architectures to models in which significantly more data is being transferred from machine to machine. The result is a shift from north-south traffic patterns to more east-west traffic in the data center. Content from the enterprise also needs to be accessible at any time, from anywhere. In addition, many corporate IT departments are showing great interest in moving to public, private, or hybrid cloud environments.
- **The consumerization of IT:** Users are demanding more bring-your-own-device (BYOD) flexibility, so that personal laptops, tablets, and smartphones can be used to access corporate information. A result of this trend is a need for greater emphasis on protection of corporate data with security policies and enforcement.

¹ [The State of SDN Adoption...](#), at blogs.gartner.com. Retrieved December 1, 2014.

² [Software-Defined Networking: The New Norm for Networks](#), at <http://www.opennetworking.org>. Retrieved December 1, 2014.

-
- **The rise of cloud services:** Public cloud services available from companies such as Amazon, Microsoft, and Google have given corporate IT departments a glimpse of self-service IT and demonstrate how agile applications and services can be. Organizations are now demanding the same service levels from their own IT departments. However, unlike public cloud environments, private cloud environments need to meet strict security and compliance requirements, which cannot be sacrificed for increased agility.
 - **Big data means more bandwidth:** Enterprises are investing in big data applications to facilitate better business decision making. However, these applications require massive parallel processing across hundreds or thousands of servers. The demand to handle huge data sets is placing greater stress and load on the network and driving the need for greater capacity.

Limitations of Current Networking Technology

The ONF paper also discusses significant limitations of current networking technologies that must be overcome to meet modern IT requirements. These challenges are presented in the context of traditional requirements: provide stable, resilient, yet static, connectivity. But the computing trends mentioned earlier require networks to support rapid deployment of applications. They also require the network to scale to accommodate increased workloads with greater agility, while also keeping costs at a minimum.

The traditional approach has substantial limitations:

- **Complexity that leads to stasis:** The abundance of networking protocols and features defined in isolation has greatly increased network complexity. The ONF paper states that each protocol is “solving a specific problem and without the benefit of any fundamental abstractions.” Additionally, old technologies were often recycled as quick fixes to address new business requirements. An example of this recycled approach is the loose use of VLANs in current networks: Initially, the purpose of VLANs was to create smaller broadcast domains. Today VLANs are being used as policy and security domains for isolation. This use has created complex dependencies that increase security risk and reduce agility, because a change in security policy requires a change in the broadcast and forwarding domain, while a change in VLANs may also impact security policy.
- **Inconsistent policies:** Security and quality-of-service (QoS) policies in current networks need to be manually configured or scripted across hundreds or thousands of network devices. This requirement makes policy changes extremely complicated for organizations to implement without significant investment in scripting language skills or tools that can automate configuration changes. Manual configuration is prone to error and can lead to many hours of troubleshooting to discover which line of a security or QoS application control list (ACL) was entered incorrectly for a given device.
- **Inability to scale:** As application workloads change and demand for network bandwidth increases, the IT department either needs to be satisfied with an oversubscribed static network or needs to grow with the demands of the organization. Unfortunately, the majority of traditional networks are statically provisioned in such a way that increasing the number of endpoints, services, or bandwidth requires substantial planning and redesign of the network. Server virtualization and private cloud deployments are challenging IT networking professionals to reevaluate their architecture. Some may choose to massively overprovision the network to accommodate the dynamic nature of virtual machines, which can be deployed on demand and instantiated anywhere on the network, but most will need to evaluate new ways to design the network.

-
- **Vendor dependence:** Although many early proponents of SDN have pointed to vendor lock-in and the high cost of networking equipment as the main reasons for moving to SDN, the industry is much more aware of the average selling price (ASP) of 10- and 40-Gbps ports in the data center, and incumbent vendors are also less inclined to offload high profit margins to the end customer. Hence, the discussion of vendor dependence and its challenges in the context of SDN focuses more on the capability of vertically integrated solutions from a given vendor to deliver “capabilities and services in rapid response to changing business needs or user demands.” The ONF paper argues that product cycles can take many years to respond to customer requirements, and that “lack of standard, open interfaces limits the ability of network operators to tailor the network to their individual environments.”

How SDN Can Help

Various surveys by networking publications³ plus informational RFCs, such as RFC 3535, list the requirements that end users have associated with SDN. These include:

- Capability to automate provisioning and management
- Capability to implement network-side policies
- Improved security
- Increased visibility into applications that are using the network
- Increased scalability
- Support for creation and dynamic movement of virtual machines
- Support for creation of a private or hybrid cloud
- Need for networks to be configured as a whole
- Need for text-based configuration for simplified revision control
- Need for network management tools more advanced than Simple Network Management Protocol (SNMP) and a command-line interface (CLI)

In March 2013, Gartner released a new research note titled “Ending the Confusion Around Software Defined Networking (SDN): A Taxonomy”⁴. The article states: “SDN is a new approach to designing, building and operating networks that supports business agility. SDN brings a similar degree of agility to networks that abstraction, virtualization and orchestration have brought to server infrastructure.” It goes on to describe three models for SDN deployment:

- **Switch based:** This model is well suited for greenfield (new) deployments in which the cost of physical infrastructure and multivendor options are important. Its drawback is that it does not use existing Layer 2 and 3 networking equipment.
- **Overlay:** This model is well suited for deployments over existing IP networks or those in which the server virtualization team manages the SDN environment. Here, the SDN endpoints reside in the hypervisor environment. The biggest drawbacks are that this model doesn’t address the overhead required to manage the underlying infrastructure, debugging problems in an overlay can be complicated, and the model does not support bare-metal hosts.

³ [2013 SDN Survey: Growing Pains](#), at [InformationWeek.com](#). Retrieved December 1, 2014.

⁴ [“Software Defined Networking Creates a New Approach to Delivering Business Agility,”](#) at <http://www.gartner.com>. Retrieved December 1, 2014.

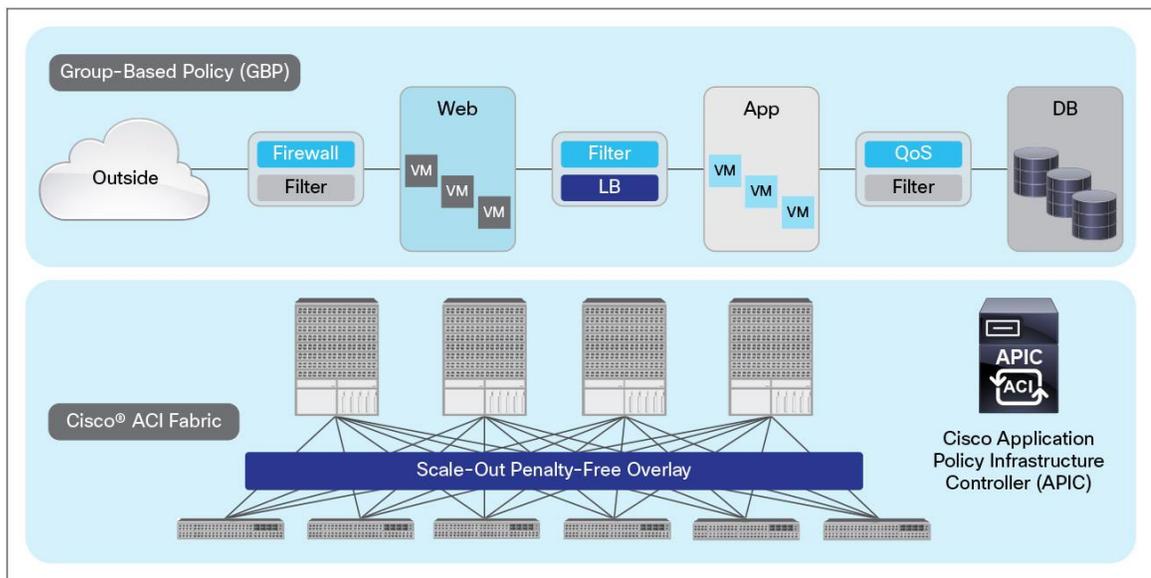
- **Hybrid:** This model is a combination of the other two approaches, with nondisruptive migration that can evolve to a switch-based model through time.

Overview of Cisco Application Centric Infrastructure

Cisco ACI is a new data center architecture designed to address the requirements of today's traditional networks, as well as to meet emerging demands that new computing trends and business factors are placing on the network.

Figure 1 provides an overview of Cisco ACI. A high-level summary of its main building blocks is presented here.

Figure 1. Cisco ACI



Application-Centric Policy Model Using Group-Based Policy

As mentioned earlier, one of the biggest challenges for current networking technologies is the tight coupling of networking protocols and features, forwarding, and policy. As a result of this coupling, a change in policy will likely adversely affect forwarding, and the converse. Furthermore, because the network protocols and features are designed for their own specific use cases, manipulation of these protocols and features requires a deep understanding of networking semantics.

To provide agility and simplicity in data center infrastructure, a new language describing the **abstracted** intent of connectivity is required so that the end user doesn't need significant networking knowledge to describe the requirements for connectivity. Additionally, this intent should be **decoupled** from network forwarding semantics so that the end user can describe the policy in such a way that a change in policy need not affect forwarding behavior, and the converse.

Because this abstracted, decoupled policy model did not exist prior to Cisco ACI, Cisco created such a model. It is called group-based policy (GBP) and is a working project in [OpenStack](#) and [OpenDaylight](#).

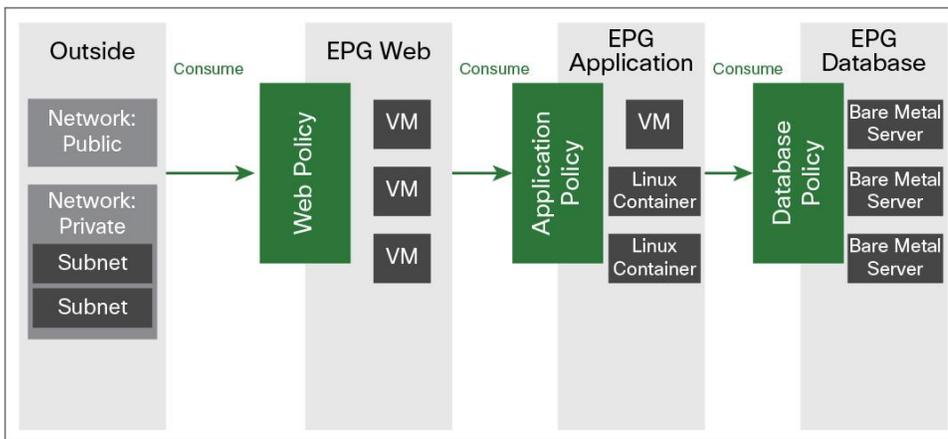
OpenDaylight describes group-based policy as “an application-centric policy model... that separates information about application connectivity requirements from information about the underlying details of the network infrastructure.”

This approach offers a number of advantages, including:

- **Easier, application-focused way of expressing policy:** By creating policies that mirror application semantics, this framework provides a simpler, self-documenting mechanism for capturing policy requirements without requiring detailed knowledge of networking.
- **Improved automation:** Grouping constructs allow higher-level automation tools to easily manipulate groups of network endpoints simultaneously.
- **Consistency:** By grouping endpoints and applying policy to groups, the framework offers a consistent and concise way to handle policy changes.
- **Extensible policy model:** Because the policy model is abstract and not tied to specific network implementations, it can easily capture connectivity, security, Layer 4 through 7, QoS, etc.

Cisco ACI makes extensive use of group-based policy in its application-centric policy model, in which connectivity is defined by consolidating endpoints (physical or virtual) into endpoint groups (EPGs). Connectivity is defined when the end user specifies a contractual relationship between one EPG and another. The end user does not need to understand the protocols or features that are employed to create this connectivity. Figure 2 provides an overview of this model.

Figure 2. Application-Centric Policy Model



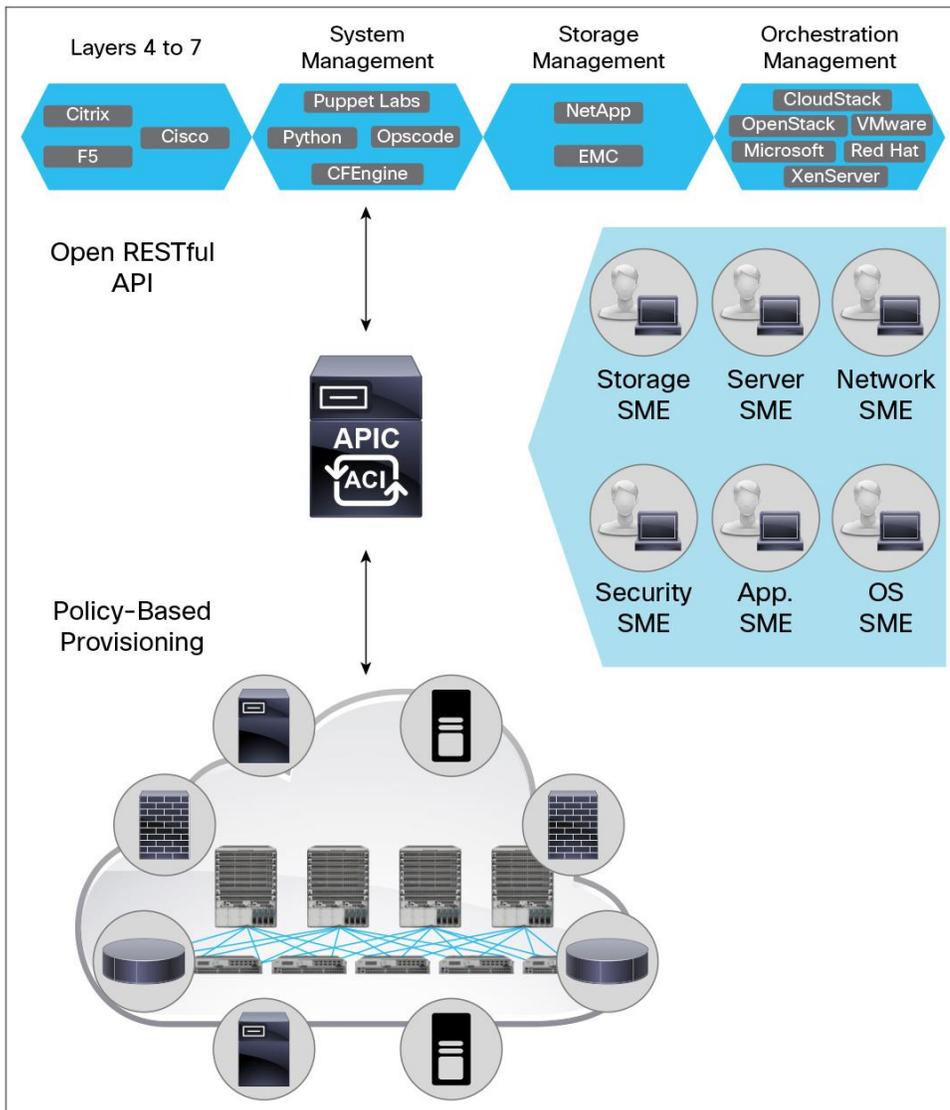
Cisco Application Policy Infrastructure Controller

In general, people want policies that are consistent across the entire network. However, one of the main challenges in managing policies in existing networks is the number of devices to which policies need to be applied coupled with the need to ensure consistency. The Cisco Application Policy Infrastructure Controller (APIC) addresses this issue.

Cisco APIC is a distributed system implemented as a cluster of controllers. It provides a single point of control, a central API, a central repository for global data, and a repository for group-based policy data for Cisco ACI.

Cisco APIC is a unified point for policy-based configuration expressed through group-based policy (Figure 3). The primary function of Cisco APIC is to provide policy authority and policy resolution mechanisms for the Cisco ACI fabric and devices attached to the fabric. Automation is provided as a direct result of policy resolution and renders its effects on the Cisco ACI fabric, so that end users no longer have to touch each network element and manually make sure that all policies are configured appropriately. Note that Cisco APIC is not involved in forwarding calculations or route provisioning, which provides additional scalability, stability, and performance.

Figure 3. The Role of Cisco APIC in the ACI Fabric



Cisco APIC communicates with the Cisco ACI fabric to distribute policies to the points of attachment and provide several critical administrative functions to the fabric. Cisco APIC is not directly involved in data-plane forwarding, so a complete failure or disconnection of all Cisco APIC elements in a cluster will not result in any loss of forwarding capabilities, increasing overall system reliability.

In general, policies are distributed to nodes as needed upon endpoint attachment or by an administrative static binding, allowing greater scalability across the entire fabric.

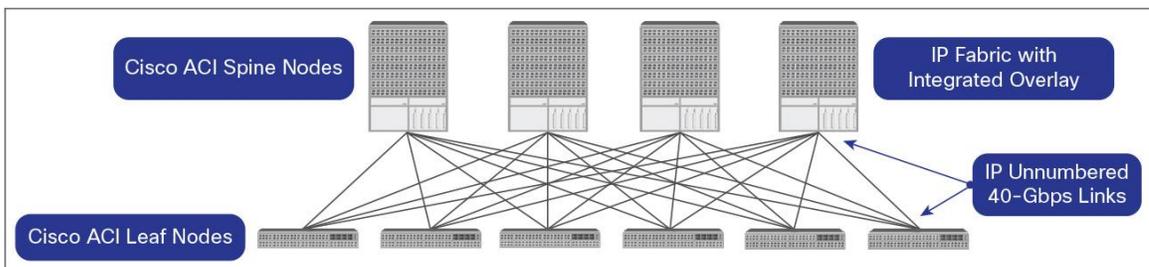
Cisco APIC also provides full native support for multitenancy so that multiple interested groups (internal or external to the organization) can share the Cisco ACI fabric securely, yet still be allowed access to shared resources if required. Cisco APIC also has full, detailed support for role-based access control (RBAC) down to each managed object in the system, so that privileges (read, write, or both) can be granted per role across the entire fabric.

Cisco APIC also has completely open APIs so that users can use Representational State Transfer (REST)-based calls (through XML or JavaScript Object Notation [JSON]) to provision, manage, monitor, or troubleshoot the system. Additionally, Cisco APIC includes a CLI and a GUI as central points of management for the entire Cisco ACI fabric.

Cisco ACI Fabric

As discussed previously, workloads continue to evolve, and traffic is becoming more east-west based. Networks need to respond faster to dynamic virtualized and cloud-based workloads and accommodate traffic growth as data sets become larger.

Figure 4. Cisco ACI Fabric



Scalability, extensibility, simplicity, flexibility, and efficiency are some of the main goals in the design of next-generation data center fabrics. When designing the Cisco ACI fabric, Cisco needed to consider all the new challenges facing the data center, but it also needed to understand and cater to existing challenges. The Cisco ACI fabric (Figure 4) is designed to address both today's requirements and tomorrow's requirements, with the following main goals:

- **Scalable fabric:** The Cisco ACI fabric is designed based on one of the most efficient and scalable network design models: a spine-and-leaf bipartite graph in which every leaf is connected to every spine, and the converse. To reduce the likelihood of hotspots of activity forming in the fabric, all devices (regardless of their functions) connect at the leaf nodes of the fabric. This approach allows the fabric to provide a simple way to scale the number of devices connected, by adding more leaf nodes. If the amount of cross-sectional bandwidth that is servicing the fabric needs to be increased, the administrator simply has to add spine nodes. This flexibility allows the fabric to start as a small environment, but gradually grow to a much larger environment if the need arises. The fabric is also built using standards-based IP routed interfaces, offering greater stability in larger scale-out deployments.
- **Extensibility:** The ACI fabric is highly extensible. The fabric administrator can integrate virtual networking (through integration of Microsoft System Center Virtual Machine Manager [SCVMM]) as well as Layer 4 through 7 services (firewalls, load balancers, etc.) today. This integration allows the end user to specify connectivity requirements using group-based policy on Cisco APIC, and the configuration for virtual networks and for Layer 4 through 7 services will automatically be rendered on the respective end systems, eliminating the need for the end user to coordinate connectivity and policies through those devices. Future software releases will also include WAN router integration.

-
- **Simplicity:** Although numerous protocols and features exist in the networking domain, the role of the fabric is very simple: to provide any connectivity anywhere. Rather than supporting numerous different protocols and features, the Cisco ACI fabric is designed with data center use cases in mind. The result is a simplified architecture without unnecessary complexity. A single Interior Gateway Protocol (IGP) has been chosen as the underlying fabric node discovery protocol: Intermediate System to Intermediate System (IS-IS). IS-IS is a link-state protocol that very efficiently detects link failures and recovers from such failures. Standards-based Virtual Extensible LAN (VXLAN) provides a simple overlay for tenant-facing traffic, supporting full Layer 2 bridging and Layer 3 routing across the entire fabric.
 - **Flexibility:** The Cisco ACI fabric supports the native capability to allow users to attach any host anywhere across the entire fabric. By using the integrated penalty-free VXLAN overlay, traffic can be flexibly bridged and routed across the entire fabric. Furthermore, the Cisco ACI fabric can provide normalization for multiple different encapsulation types arriving from hosts or their respective hypervisors, including VLAN, VXLAN, and Network Virtualization using Generic Routing Encapsulation (NVGRE). This feature allows physical, virtual, and container-based hosts to all co-exist on the same shared infrastructure. In addition, next-generation data center fabrics need to be backward-compatible with sunset-type applications, which may not be IP based or which may use network flooding semantics for discovery and communication. The Cisco ACI fabric can support both modern data center requirements and the requirements of traditional bare-metal and mainframe-based applications.
 - **Efficiency:** An inherent benefit to the spine-and-leaf bipartite graph architecture of the Cisco ACI fabric is that every host is exactly two physical hops away from every other host in the fabric. So for big data workloads that require a significant amount of east-west traffic between machines, the Cisco ACI fabric provides predictable low latency at scale. This approach delivers efficient support for traditional data center applications as well. The Cisco ACI fabric can exceed other traditional spine-and-leaf fabrics in fabric bandwidth efficiency, because it can take into account packet arrival time, end-to-end fabric congestion, and flowlet switching to make more intelligent load-balancing decisions. More information about these innovations is documented in the SIGCOMM paper "[CONGA: Distributed Congestion-Aware Load Balancing for Datacenters](#)".
 - **Investment Protection:** Customers may also want applications on their current IP networks to participate in the Cisco ACI fabric policy. The Cisco ACI Fabric allows for investment protection where the Cisco APIC manages policy for virtual or physical servers in the existing network. For virtual servers, they connect to an application-centric virtual switch (AVS) and enabled for Cisco ACI, in the existing Cisco Nexus network. AVS acts as a virtual leaf for the Cisco ACI spine-and-leaf fabric, and as an edge switch that is Cisco ACI aware, it can forward traffic according to Cisco ACI policy rules and apply Layer 4 through 7 services managed by Cisco ACI. For physical servers, a Cisco Nexus 9300 platform switch acts as an access-layer switch to the existing overlay network. Compared to other software-defined networking (SDN) overlay solutions, this solution provides a common infrastructure for physical and virtual workloads, along with a more advanced application-centric policy model. For more information, refer to the white paper: "[Transform Your Business and Protect Your Cisco Nexus Investment While Adopting Cisco Application Centric Infrastructure](#)".

Open APIs, Partner Ecosystem, and OpFlex

Cisco ACI supports an extensible partner ecosystem that includes Layer 4 through 7 services; hypervisors; and management, monitoring, and cloud orchestration platforms. All use Cisco ACI's open APIs and development kits, device packages, and plug-ins, as well as a new policy protocol, OpFlex, which is used to exchange group-based policy information.

- **Open APIs:** Cisco ACI supports API access through REST interfaces, GUIs, and the CLI as well as a number of software development kits (kits for Python and Ruby are available today). Cisco APIC supports a comprehensive suite of REST APIs over HTTP/HTTPS with XML and JSON encoding bindings. The API provides both class-level and tree-oriented data access. REST is a software architecture for distributed systems. It has emerged over the past few years as a leading web services design model and has increasingly displaced other design models such as Simple Object Access Protocol (SOAP) and Web Services Description Language (WSDL) because of its simpler approach.
- **Partner ecosystem and OpFlex:** OpFlex, the southbound API, is an open and extensible policy protocol used to transfer abstract policy in XML or JSON between a policy controller, such as Cisco APIC, and any device, including hypervisor switches, physical switches, and Layer 4 through 7 network services. Cisco and its partners, including Intel, Microsoft, Red Hat, Citrix, F5, Embrane, and Canonical, are working through the IETF and open source community to standardize OpFlex and provide a reference implementation.

OpFlex is a new mechanism for transferring abstract policy from a modern network controller to a set of smart devices capable of rendering policy. Whereas many existing protocols such as the Open vSwitch Database (OVSDB) management protocol focus on imperative control with fixed schemas, OpFlex is designed to work as part of a declarative control system, such as Cisco ACI, in which abstract policy can be shared on demand. One major benefit of this model is the capability to expose the complete feature set of an underlying device, allowing differentiation of hardware and software objects such as Layer 4 through 7 devices.

In addition to its implementations in the open source community, OpFlex is one of the primary mechanisms through which other devices can exchange and enforce policy with Cisco APIC. OpFlex defines that interaction. As a result, by integrating a number of devices from both Cisco and ecosystem partners using Cisco ACI, organizations can use it to gain investment protection.

How Cisco ACI Addresses Current Networking Limitations

Cisco ACI can address all the traditional networking limitations outlined in the ONF paper. Cisco ACI is built using a balanced approach that weighs the best software against the best hardware, custom silicon against merchant silicon, centralized models against distributed models, and the need to address old problems against the need to meet new challenges. It tackles business challenges rather than championing only one particular technology approach.

Cisco ACI addresses these specific limitations:

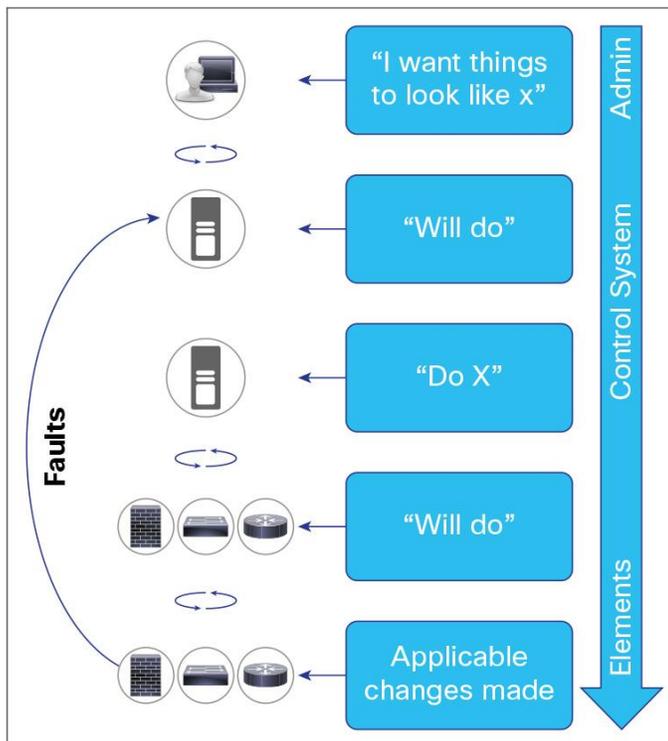
- **Complexity that leads to stasis:** Cisco ACI removes complexity from the network. Cisco ACI sets out to decouple policy from forwarding by allowing network routing and switching to be completely distributed across the entire fabric. Cisco ACI packet forwarding across the fabric uses a combination of merchant and custom silicon to deliver standards-based VXLAN bridging and routing with no performance penalty or negative impact on the user or application.

In addition, Cisco ACI can apply policy without the need to derive this information from network information (such as IP addresses.) It does this by populating each VXLAN frame with a 16-bit ID to uniquely identify the originating (source) group of the packet as specified in the group-based policy VXLAN IETF draft⁵. This approach provides outstanding flexibility for the end user, allowing users to modify network policies with little network knowledge and no negative impact on network forwarding.

Cisco ACI further simplifies policies by introducing an abstraction model - group-based policy - so that end users can define connectivity using higher-level abstracted constructs instead of concrete networking semantics. This model enables end users of Cisco ACI to define policy rules without the need for knowledge of networking, opening the way for application administrators and developers to directly interact with Cisco ACI policies to express their intent without the need to involve IT network administrators.

- **Inconsistent policies:** One of the biggest challenges in managing network policies across a large network is the requirement to touch a large number of devices and make sure that the policy configuration remains consistent. Cisco ACI addresses this challenge by offloading this task to Cisco APIC, which is the central policy authority and the central point of management for Cisco ACI and associated physical and virtual services. The end user simply needs to specify on Cisco APIC the desired intent of group-based policy, and Cisco APIC distributes the policy to all the nodes in the Cisco ACI fabric (Figure 5).

Figure 5. Using Cisco APIC to Implement Policy across the Fabric



Cisco APIC uses a variant of promise theory, with full formal separation between the abstract logical model and the concrete model, and with no configuration performed on concrete entities. Concrete entities are configured implicitly as a side effect of the logical model implementation. This implementation of promise theory provides policy consistency throughout the network at scale.

⁵ "VXLAN Group Policy Option" [[draft-smith-vxlan-group-policy-00](#)], M. Smith and L. Kreeger.

-
- **Inability to scale:** Cisco ACI is designed to scale transparently throughout its deployment, supporting changes in connectivity, bandwidth, tenants, and policies. The spine-and-leaf topology of the Cisco ACI fabric supports a scale-out design approach. If additional physical connectivity is required, leaf nodes can be added by connecting them to the spines. Similarly, if additional bandwidth or redundancy is required, additional spine nodes can be introduced. This scale-out deployment model also allows end users to start small and later scale to extremely large environments, thereby reducing the initial capital expenditure required to implement a scalable fabric. However, the addition of new devices does not mean an increased number of management points. After registering the new devices on the Cisco ACI fabric through Cisco APIC, the end user can administer the entire fabric, including the new devices, from the central Cisco APIC. Introduction of the new devices requires no intervention by the administrator.

Tenants and policies also use a scale-out approach. Policies are centrally stored on the fabric and are rendered to fabric nodes as required. The Cisco APIC policy repository itself is a scale-out clustered database. It can increase from 3 to more than 31 nodes in a single cluster, depending on the scale of tenants and policies required. Even with additional cluster nodes, all nodes are considered active, and policies can be managed on any cluster member.

- **Vendor dependence:** A complete Cisco ACI deployment will likely include Layer 4 through 7 services, virtual networking, computing and storage resources, WAN routers, and northbound orchestration services. A main strength of Cisco ACI is its openness, with its published APIs, Layer 4 through 7 device packages, and use of the open OpFlex protocol. With these open APIs, plug-ins, and protocols, end users can incrementally add functions to their solutions without the need to wait for a single vendor to introduce new capabilities.

How Cisco ACI Meets the Requirements of SDN

As mentioned previously, the ONF has prescribed that a new network model is needed to address the emerging trends in computing. A proposed solution is SDN, with SDN defined as “an approach to computer networking that allows network administrators to manage network services through abstraction of lower-level functionality. This is done by decoupling the system that makes decisions about where traffic is sent (the control plane) from the underlying systems that forward traffic to the selected destination (the data plane).”

As explained in this document, the Cisco ACI solution meets all the requirements of this definition, but to satisfy the requirements of SDN, it must also be able to accommodate the emerging trends in computing, including:

- **Changing traffic patterns:** As explained previously, east-west traffic in the data center is increasing, and next-generation data center architectures need to be optimized for such workloads. Cisco ACI uses a spine-and-leaf architecture that is well suited for east-west workloads, helping ensure consistent latency and performance from any source in the data center to any destination.

Another requirement for the modern data center is that content be accessible anytime, anywhere. Cisco ACI meets this requirement with its penalty-free overlay. It allows applications and services to be deployed anytime, anywhere across the entire fabric. It is also well suited as the platform of choice for any public, private, or hybrid cloud deployment. Because connectivity policies are abstracted away from network forwarding, higher-layer cloud orchestration platforms (such as OpenStack, CloudStack, and Microsoft AzurePack) can manage services instead of the underlying infrastructure.

-
- **The consumerization of IT:** Cisco ACI enables BYOD in the enterprise in a number of ways. First, security and compliance is inherently built in to the group-based policy model; the application administrator must explicitly grant access to outside groups for those groups to consume applications or services. Second, an audit of security policies is a much simpler task because policies are defined in an abstracted group-based policy representation rather than in concrete security access lists or firewall rules. Finally, group-based policies can be migrated to other areas of the network, such as campus and branch offices, where devices can be allocated to their own endpoint groups based on flexible classifications such as device type, access location, access method, time of day, and day of the week instead of network addresses.
 - **The rise of cloud services:** Cisco ACI enables enterprise IT departments to easily implement private clouds by abstracting applications, security policies, and IT services away from the underlying infrastructure. Organizations can spend more time creating connectivity policies through group-based policy rather than managing individual network devices. Additionally, many IT departments will prefer this self-service model and be more likely to adopt it because they can retain control over their sensitive content on their own premises, instead of giving control to a public cloud service. However, if an organization later wants to shift certain workloads to the public cloud, Cisco ACI group-based policy enables the organization to move network and security policies easily; a consistent group-based policy profile definition can exist in the on-premises private cloud as well as in the hosted cloud with very little modification.
 - **More bandwidth for big data:** Cisco ACI provides an optimal solution for hosting big data workloads. In addition to providing predictable low latency using 40-Gbps fabric links as mentioned earlier, Cisco ACI provides a cost-effective, scale-out fabric, allowing the end user to incrementally increase the cross-sectional bandwidth of the fabric simply by introducing an additional scale-out spine node. More computing nodes can also be incrementally added to a big data cluster to improve performance through the addition of leaf nodes. Also, compared to an equivalently built-out network fabric using the same number of devices and interface speeds, Cisco ACI can use its fabric bandwidth more efficiently, through the use of distributed congestion-aware load balancing (CONGA). This mechanism reduces flow completion time for network-intensive applications such as big data.

Cisco ACI also delivers other features that end users are seeking in SDN solutions, including:

- **Automated provisioning and management:** Cisco ACI provides complete automated provisioning and management of the network infrastructure. Fabric administrators can transparently add spine or leaf nodes to the Cisco ACI fabric without any configuration on the devices themselves. Provisioning of ecosystem partner solutions is also fully automated through the group-based policy model on Cisco APIC, so the end user doesn't need to manually and individually manage those devices and systems. Day-two management of these devices and systems is also enabled through Cisco APIC.
- **Capability to implement network-side policies:** Cisco ACI manages the entire system as a single entity, and any associated network-side policies (whether associated with physical, virtual, or Layer 4 through 7 services) can be defined, deployed, and managed through Cisco APIC.
- **Improved security:** Security is implicit in the application-centric policy model. Connectivity between groups is disallowed until the policy administrator defines, through group-based policy, which groups are allowed to talk, and which conduit each group can use to communicate. Cisco ACI is multitenant aware. Traffic, connectivity, and policies from different tenants can share the same infrastructure without leakage of information across tenants. All manageable entities (called objects) in Cisco ACI have unique privileges associated with them, and the administrator can assign highly specific security controls using RBAC definitions.

-
- **More visibility into applications that are using the network:** The definition of an application is provided directly to Cisco ACI in the form of a group-based policy profile (called an application network profile). From this profile, Cisco ACI can glean information about network dependencies and report on how well a given application is performing based on the underlying objects that are being monitored by Cisco ACI.
 - **Increased scalability:** The dynamic nature of the Cisco ACI fabric, the use of scale-out spine-and-leaf architecture and a clustered database, and hardware reachability for more than one million endpoints makes Cisco ACI the most scalable SDN solution on the market today. In addition, Cisco ACI can use endpoint location information to conservatively render policies in hardware only to enforcement points to which endpoints are attached. This feature allows greater scalability without the need to overprovision policy resources.
 - **Capability to create and move virtual machines dynamically:** The Cisco ACI fabric employs a penalty-free overlay, allowing outstanding flexibility for virtual machine deployment and unrestricted virtual machine mobility across the entire Cisco ACI fabric. Cisco ACI also includes innovations to make virtual machine mobility transparent.
 - **Capability to create a private or hybrid cloud:** Cisco ACI is an excellent solution for organizations wanting to deploy private, hybrid, and public clouds. Most Cisco ACI early adopters were public cloud service providers. In addition to its scalability, security, and flexibility, Cisco ACI comes out of the box with an abstracted representation of network connectivity through the group-based policy model. This feature is extremely attractive in cloud deployments because a main requirement for such deployments is the capability to abstract the complexity of the network infrastructure so that it is not visible to the consumer of the cloud service.
 - **Capability to configure the network as a whole:** Cisco ACI inherently allows the entire fabric to be viewed as a single Layer 2 and 3 virtual forwarder. Additional private networks can be created, but they always are configured as a single entity. This approach allows the end user to deploy new services (network services or applications) without the need to understand how the underlying network is provisioned or connected.
 - **Text-based configuration for simplified revision control:** All configuration settings are represented as managed objects through Cisco APIC. These managed objects can be accessed through APIs, which can be managed through structured REST commands in either XML or JSON format. Additionally, because these objects have a defined structure that is abstracted from the underlying physical and virtual networks, the same profile definitions can be used to implement the same connectivity policies in a completely different Cisco ACI fabric.
 - **Advanced network management tools beyond SNMP and CLI:** In addition to SNMP and the CLI, the Cisco ACI object model natively supports REST interfaces through XML and JSON, GUIs, and a number of software development kits. In addition to general network device management, Cisco ACI includes many innovative management capabilities that are not available in other SDN solutions. For example, Cisco ACI provides health scores, which capture dependencies across underlay, overlay, physical, virtual, and ecosystem devices, and contextualizes them into detailed health scores for specific objects. These scores can also be rolled into higher-level scores at the endpoint group level, application network profile level, or tenant level to provide an easy-to-see top-level view. The end user can then use this view to troubleshoot any anomalies throughout the Cisco ACI fabric. Atomic counters are another Cisco ACI innovation.

These provide end-to-end counts of packets entering and leaving the fabric, and they can also be scoped to individual leaf nodes. Atomic counters allow end users to easily see areas of the fabric with high, medium, and low use, through the creation of a traffic map, and to identify the location of any traffic drops in the Cisco ACI fabric - all from Cisco APIC.

Joe Skorupa, vice president and distinguished analyst at Gartner, defined the value of Cisco ACI succinctly when in March 2013 he described what he hoped SDN would deliver: “In a data center context, SDN is a component of the policy driven data center. It provides the programmable connectivity required to link the network to other components within the data center, delivering a more integrated, functional system. For example, a provisioning application could specify that an instance of the CRM application must have certain services delivered in a specific sequence and would ensure that the traffic flows through the appropriate devices in the correct sequence.”⁶

SDN Faces a New Set of Challenges

Although SDN, software-based virtual overlays, and OpenFlow present some interesting solutions for both traditional and emerging computing workloads, except in academic and research institutions, few data centers have adopted software overlays and OpenFlow. This lack of adoption can be attributed to a new set of challenges, including the following:

- **Software-based virtual overlays are difficult to manage and troubleshoot:** Joe Skorupa, vice president and distinguished analyst at Gartner, identified this limitation in a brief Q&A session in March 2013⁶: “The greatest limitations of this [software-based virtual overlay] approach are that it does not address the overhead of managing the underlying infrastructure, debugging problems in an overlay can be complex, and it does not support bare-metal hosts.”

The fundamental problem with software-based virtual overlays is that they have little or no relationship with the underlying physical infrastructure (which, of course, is always required), so any drops in the underlay network cannot easily be traced back to the service or application that was affected, and the converse. Furthermore, the software-based virtual overlay may be managed by a different team, with a different skillset, that isn't equipped to troubleshoot end-to-end network connectivity issues, leading to finger-pointing across IT operations departments, and possibly vendors. This drawback is potentially more severe in traditional networks, in which flexibility is limited, but at least such networks have predictability and deterministic points of failure for which a single IT operations group takes ownership.

- **OpenFlow protocols are too primitive and concrete for the data center:** OpenFlow in a network switch applies a match function on the incoming packet (typically based on an existing network field attribute such as a source MAC address or destination IP address) and then takes some form of action, which may be depend on the switch implementation but typically involves forwarding the packet out through a given physical or logical port.

For most data center use cases, such detailed control is not required, because bandwidth (and hence paths) in the data center is typically abundant. Therefore, detailed path decisions based on network header parameters may not be needed and may incur unnecessary overhead for data center network administrators to manage.

⁶ “Software Defined Networking Creates a New Approach to Delivering Business Agility” <http://www.gartner.com>. Retrieved December 1, 2014.

In addition, the OpenFlow protocol assumes a particular hardware architecture using a series of generic table look-ups that generate actions to apply to the packet. This specific view of the underlying hardware architecture limits scalability and portability. A higher level of abstraction allows different specific hardware architectures to provide specific benefits while still meeting the requirements of the APIs.

- **Merchant silicon is not optimized for OpenFlow today:** Most merchant silicon available in the market today has been optimized for general data center workloads. Typically, such deployments mandate allocation of a large amount of memory to forwarding tables to perform longest-prefix match and adjacency, Address Resolution Protocol (ARP) and MAC and IP address binding lookups; and a finite amount of memory to ACLs, typically using ternary content-addressable memory (TCAM) that is optimized for masking, packet matching, and action sets.

OpenFlow tables use the latter form of memory in switches. Unfortunately, forwarding memory is not easily changeable with ACL memory, so the total amount of memory available to install flow entries in today's modern merchant silicon is somewhat limited. As mentioned earlier, if any misses occur, packets may be unexpectedly dropped or forwarded to the SDN controller for further processing.

How Cisco ACI Addresses the New SDN Challenges

A next-generation SDN-based architecture must address all these challenges. Cisco ACI accomplishes this in context of the modern data center.

- **Software-based virtual overlays are difficult to manage and troubleshoot:** Cisco ACI does not rely solely on software-based virtual overlays for fabricwide connectivity. Although Cisco ACI does deploy an overlay, it is instantiated in hardware, and the management features of Cisco ACI have enough intelligence to provide full coordination and contextualization to show what services are affected by failures in the underlay or overlay network.

Note that Cisco ACI supports software-based overlays that either terminate on the Cisco ACI fabric or run over the top. These are completely optional modes of deployment, and the end user can decide how to deploy workloads. The point is that Cisco ACI does not rely on software-based overlays to achieve the flexibility and agility sought in SDN.

- **OpenFlow protocols are too primitive and concrete for the data center:** In addition to providing automation and flexibility, Cisco ACI introduces a new model to describe connectivity in the data center through group-based policy. Because the policy intent is abstracted, the end user can define how certain “things” connect to other “things” in the data center. This abstracted view of policies is a lot easier for cloud orchestration tools, applications, and security administrators to consume than is the case with OpenFlow protocols.
- **Merchant silicon is not optimized for OpenFlow today:** Cisco ACI uses a combination of merchant silicon and custom silicon developed by Cisco to provide the right level of capabilities at scale, while still delivering the solution at an attractive price. The custom silicon embedded in the Cisco Nexus® 9000 Series Switches includes memory table structures optimized so that certain functions can be offloaded to the on-board merchant silicon, with the custom silicon dedicated to functions such as policy enforcement, encapsulation normalization, and VXLAN routing.

Conclusion

The IT industry is going through a significant transformation, with BYOD, big data, cloud computing, IT as a service, and security now prominent concerns. At the same time, companies increasingly want to reduce overall IT spending (through reduction in both capital expenditures and operating expenses) and to provide much-improved levels of service to business units and functions by increasing overall IT agility. Many in the networking industry have cited SDN as the model to move the industry forward, but adoption of the prescribed technologies presents challenges.

Cisco ACI was developed not as a competitor to SDN, but as a catalyst to help promote the adoption of SDN throughout the IT industry: in essence, as an enabler of the SDN vision. Although some in the industry claim that Cisco ACI is not SDN, assessing Cisco ACI in the context of the definitions presented in the ONF SDN white paper shows that it does indeed meet the requirements of SDN. Cisco ACI also goes a step further, addressing areas that OpenFlow has not yet addressed.

Ultimately, the industry will be the final judge of whether Cisco ACI (and SDN) succeeds, but based on the current interest and adoption momentum, Cisco ACI is well positioned to become the new norm for data center networks.

For More Information

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)