Implementing Cisco Application Centric Infrastructure with Citrix NetScaler Application Delivery Controllers

Data Center Architecture Solutions

CITRIX®

Introduction

Traditional data center network architectures suffer from challenges when integrating Layer 4 through 7 services deployments with the network infrastructure and aligning them together with business application deployments. These challenges include the following:

- L4 to L7 service insertion
- Time to deploy
- Location rigidity
- Automation of services with network implementation
- No single central point of control
- Network and services alignment with specific application deployments
- Application-based service level agreements (SLAs) for the infrastructure
- No automated configuration cleanup

Cisco Application Centric Infrastructure (ACI) integrates with Citrix NetScaler Application Delivery Controller (ADC) appliances. This integration reduces deployment complexity and better aligns applications to infrastructure automation using centralized policy-based management. The combined solution provides a single point of management to define the network and L4 to L7 services requirements using policy-centric profiles, while elastically integrating them into the Cisco ACI network fabric. Cisco ACI and Citrix NetScaler ADC appliances combine to reduce deployment time to your physical and virtual application network services on a data center fabric. This combination offers freedom of deployment location, alignment with specific business applications, automated network configuration and services, plus ease of cleanup when the applications using services are decommissioned.

This paper describes the Cisco ACI and Citrix NetScaler ADC solutions. The paper includes details on how they integrate to solve traditional data center networking issues.

Audience

This document is intended for, but not limited to, the following profile of personnel.

- Data center IT systems architects and systems engineers
- Networks and services design and implementation engineers
- Application architects
- Application deployment engineers

Summary of Objectives

The primary goal of this document is to provide an overview of the features and benefits of Cisco ACI and Citrix NetScaler ADC appliances. The objective is to explain the key points of integration, benefits, and value these solutions bring to customers.

These improvements include the following:

- Reducing time to deploy
- Providing flexibility of services locations
- Automating deployment and cleanup of services with network implementation through a single central point of control
- Aligning network and services specific to application deployments
- Using advanced per-application-based telemetry to ensure infrastructure has visibility of SLAs relating to specific applications

The document starts with an overview of the features and benefits of Cisco ACI followed by an overview of Citrix NetScaler ADC appliances. The document then discusses how these benefits are better realized in a combined solution.

Cisco ACI Overview

What Is Cisco ACI?

Cisco ACI is an innovative new data center fabric. It is based on a combination of traditional protocols and new developments with centralized management, programmability, and automation. This centralization creates a highly flexible, scalable, resilient, high-speed network infrastructure with integrated overlay. Cisco ACI has been designed to support new industry demands in levels of programmability and flexibility, while maintaining integration with existing infrastructure.

Cisco ACI Fabric

The Cisco ACI fabric consists of three principal components, which are the Cisco Nexus 9000 spine and leaf switches, the Application Policy Infrastructure Controller (APIC), and a further optional Application Virtual Switch (AVS).

Figure 1: The Cisco ACI Fabric

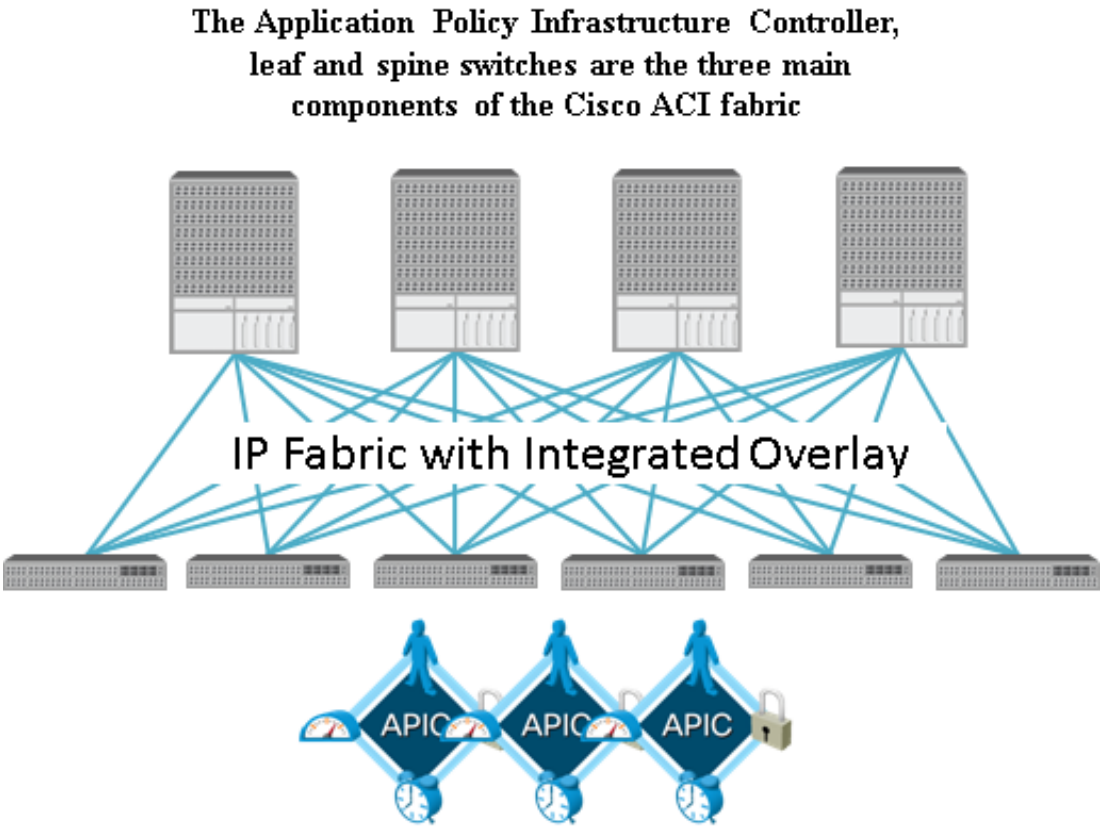

The Cisco ACI fabric provides a self-discovering IP fabric base with integrated overlay functionality. This feature allows any-to-any routing and switching of connected device flows, independent of the underlying IP structure of the network. Any device with any IP

address can exist on any leaf switch and communicate with any other device by configuring a policy-based application profile containing the communication rules.

Integrated Virtual and Physical Network

The Cisco ACI fabric integrates networking between physical and virtual switches to provide a single central point of management. This ability offers consolidated management for configuration, policy, monitoring, and troubleshooting. Cisco ACI integrates with hypervisors in what it calls a **virtual machine manager (VMM) domain**. Support for multiple hypervisors is built into Cisco ACI. Currently, support includes VMware ESXi with vCenter, with Microsoft Azure pack and Microsoft Hyper-V 2012 coming soon, and Linux-based hypervisors to follow.

The Cisco ACI fabric integrates with the virtual domain controller such as vCenter to instantiate a Cisco ACI-controlled virtual distributed switch and its port groups into the vCenter data center. This virtual distributed switch and its port groups become an extension of Cisco ACI. The virtual server administrator assigns the hosts and virtual machines to it in the usual fashion, meaning business as usual for the virtual server teams. All network provisioning and visibility for the network can be managed by the network team through Cisco ACI.

Figure 2: Cisco ACI VMM Domain with Cisco ACI Created and Managed Virtual Distributed Switches

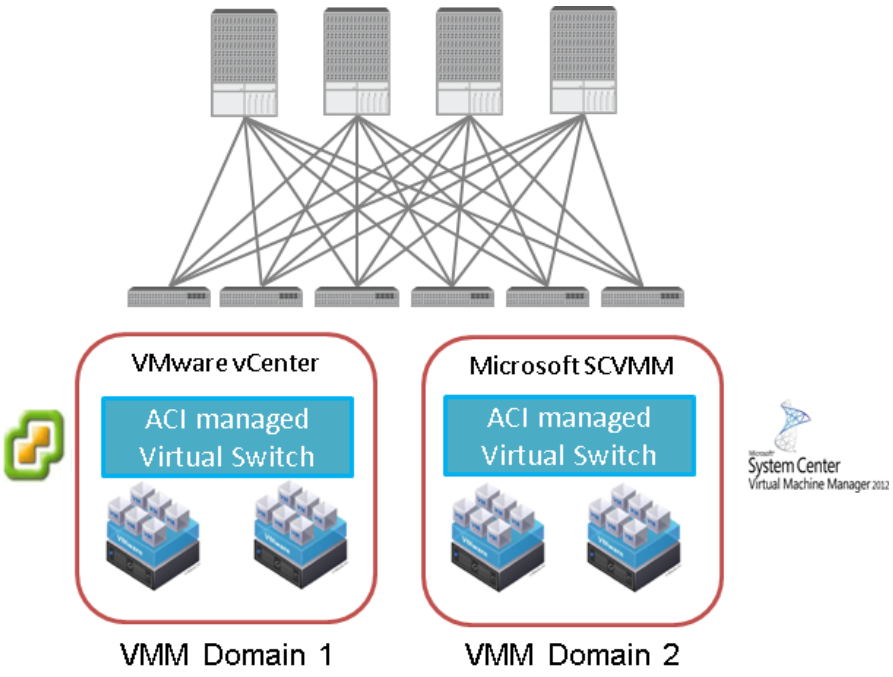

Cisco Nexus 9000 Series Switches

The Cisco ACI leaf and spine switches comprise the Cisco Nexus 9000 Series. This series provides high performance and density, low latency, and outstandingly power efficient switches, which come in a range of form factors.

The Cisco Nexus 9000 Series operates in both Cisco Nexus Operating System (NX-OS) mode and Cisco ACI mode. This ability allows you to benefit from investments in the hardware used in traditional networks by redeploying in the Cisco ACI fabric.

The Cisco Nexus 9000 Series Switches include the modular 4-slot, 8-slot, and 16-slot 9500 switches and the multiple fixed Cisco Nexus 9300 Series Switches, providing link speeds of 1/10/40Gbps.

Features of this series include the following:

- Industry-leading price per port for 10- and 40-Gbps
- The industry's most efficient chassis—no mid-plane—which allows optimum efficiency and airflow, up to 15 percent less power per port than competing solutions
- Industry-leading performance and port density with full line-rate non-blocking 40-Gbps connectivity in Cisco ACI fabric spine and up to 576 x 40-Gbps ports in the 16-slot chassis
- Highly compact leaf switches for top of rack (ToR) deployment such as the 1 rack unit (RU) Cisco Nexus 9372 with 48 x 1/10-Gbps edge-facing ports with 6 40-Gbps links to the Cisco ACI spine
- Cisco 40-Gbps (BiDi) optics allow standard multi-mode 10-Gbps fiber cables already placed in the data center to be redeployed for the 40-Gbps connections between the leaf and spine switches

Figure 3: Cisco Nexus 9000 Series Switches

Application Policy Infrastructure Controller

The Cisco ACI fabric is managed and programmed through the APIC, which is a clustered network control system, however, it is not in the network L2/L3 control plane itself. The network L2/L3 control plane remains in the leaf and spine switches. This feature means the network scales as more switches are added and the APICs do not become a central bottleneck or scalability constraint.

The APIC provides central, programmable, automated, and simplified management. The APIC exposes a northbound application programming interface (API) through XML or JavaScript Object Notation (JSON) and provides an integrated CLI and GUI to manage the fabric.

The APIC is a highly available cluster of at least three appliances that activates the fabric, manages switch firmware, and monitors and configures the fabric completely outside of the data plane. The fabric can still forward traffic without the APIC cluster.

APICs are designed based on distributed computing architecture for scalability and reliability. Each node in the cluster is always active, processing requests and data received. The fabric configuration is sharded and stored in multiple copies across each APIC in the cluster. Each shard is maintained by a primary APIC and copied at least twice to two other APICs. This allows for load to be spread evenly and multiple copies to provide high availability in the event of an APIC failure. Fabric configuration can also be partially or completely stored elsewhere in zip binary or XML format for recovery, if required.

Centralized Visibility

The Cisco APIC additionally provides centralized visibility with real-time application network health scores and telemetry. The APIC uses a health score system that is associated with a specific application using the Cisco ACI **application profile** as the container of all elements associated with the communication requirements for a particular business application. The health score looks at all aspects of the application profile. The score consists of the individual configuration elements within the Cisco ACI fabric, and the ongoing network performance metrics associated with the application flows and overall equipment health, including alarms and logs. Thus the health score can be immediately used to determine if the network is the cause of any poor behavior associated with a given application, without the need for additional probes and tools. These functions and tools are built into the Cisco ACI fabric and configurable through the APIC.

The analytics and telemetry can also include measuring latency for application flows across the fabric and recording any packets dropped for a given application. This information includes where the packets are being dropped such as at a given physical

or virtual switch port. Relating telemetry to a specific application is a big step forward in network innovation.

Figure 4: Cisco APIC Next-Generation Analytics Capabilities

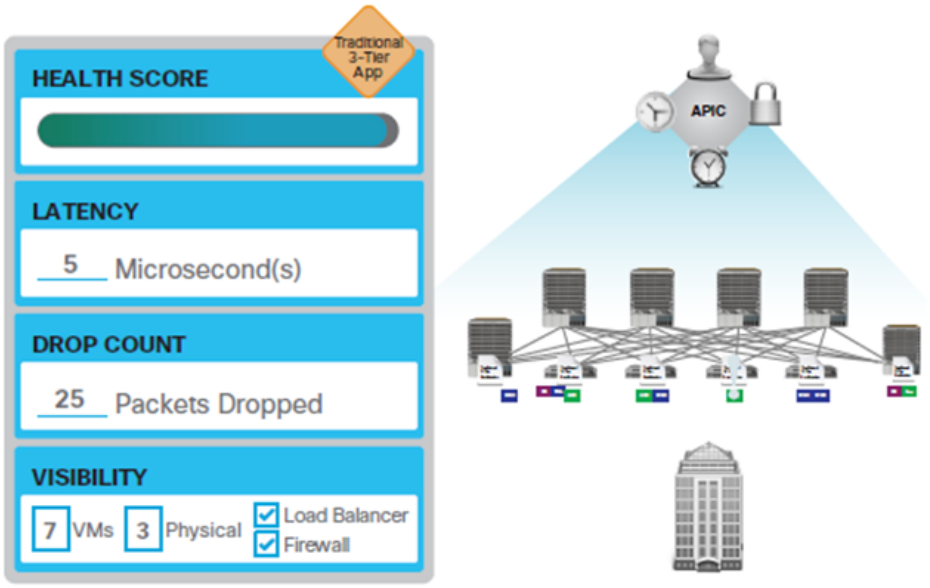

Application Virtual Switch

As an alternative to the native virtual switches of the hypervisor installed and configured by Cisco ACI, the APIC can also install and manage the Application Virtual Switch (AVS). The AVS is a virtual switch that has been designed by Cisco and lives in the hypervisor, similar to the Cisco Nexus 1000V Virtual Ethernet Module (VEM).

The AVS can be considered the implementation of the Cisco Nexus 1000V in Cisco ACI. The Cisco AVS uses the Cisco southbound API known as **OpFlex** to communicate with the APIC, which manages each AVS component within the hypervisor instead of using a Virtual Supervisor Module (VSM). One AVS element exists in each hypervisor within the VMM domain, behaving like a virtual leaf (vleaf) switch for the Cisco ACI fabric. The AVS license is bundled in at no additional cost and supports up to 32 hosts with vSphere 5.5. The APIC creates both the uplink profiles and the VM profiles for the AVS and sends them to vCenter. So once again, the operational model is business as usual with the virtual server administrator adding the hosts and VMs to the AVS for using the Cisco ACI fabric as the communication network.

Again, consolidating both virtual and physical networks under one central programmable management plane provides huge benefits for customers looking to streamline network operations in the data center, build towards private cloud infrastructures, and reduce costs and time to deploy IT for business applications. Cisco

ACI provides significant steps forward to achieve this goal of network unification while maintaining business as usual from the virtual server administration perspective to foster easier adoption.

Figure 5: The Cisco Application Virtual Switch

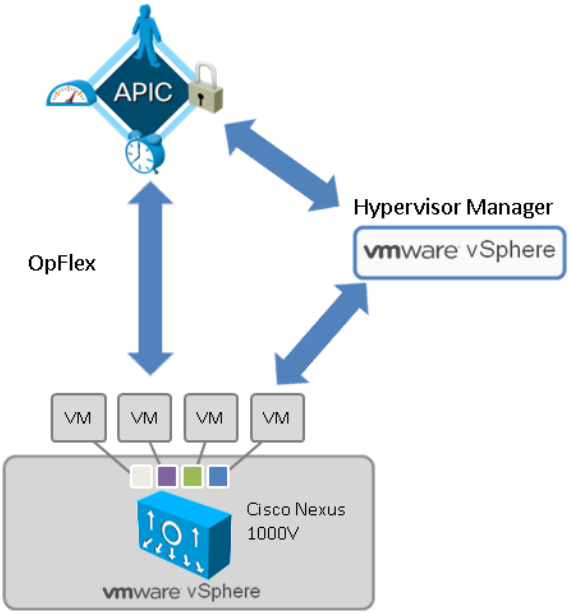

Application Network Profiles

The Cisco ACI fabric is designed around instantiating network connectivity by means of configuration profiles. These profiles, called **application profiles**, define a policy or set of policies that allow communications to pass over the fabric. Profile use allows communication rules to be distinctly defined and grouped together for ease of deployment and for adoption by automation and orchestration tools. The profiles are stored as managed objects in XML or JSON format. These application profiles are centered around the communication requirements between elements of a given business application. A given application profile, therefore, contains the desired communication paths for all the application elements to be deployed in the data center, using the Cisco ACI fabric as transport.

The APIC is used to define the application profile. What is key and different about this compared to all other network automation solutions is that the profile defines elements that are application centric rather than network centric. For example, a group of network physical or virtual servers may be grouped in a single tier of an n-tier application. Communication between the tiers and the policies that define that communication make up the application. The Cisco ACI application profile defines a set of policies that align with those communications. The network therefore is built around the application communication requirements, the health of the network is defined in terms of the

application communication requirements, and the network is programmed to only allow these specific communication paths when the application tier members are present on the fabric. This functionality means that removing an application profile also removes the communication elements in the fabric associated with it. This ability provides automatic cleanup, for example, when an application is retired. Additionally, this function also means that the network is very secure, since all allowed communication is specifically defined. If a disaster recovery site is being built, the same application profile can be deployed there to instantiate the required application connectivity in minutes, saving weeks or months of deployment staging and testing.

Application profiles drive the network overlays in Cisco ACI. Thus the Cisco ACI fabric is based around the underlying fabric—which you build once—and then overlay application profiles to provide application-specific communication by means of standards based virtual extensible LAN **(VXLAN) overlay technology**.

Therefore, only flows defined by the application profiles are allowed to pass across the fabric. All other flows are implicitly filtered. This feature means that there is no longer a need for large, difficult to manage access lists or stateless firewall rules.

Integrating Cisco and Cisco Partner firewalls and ADCs, such as Citrix NetScaler solutions, into Cisco ACI is also made easier for network simplification and rapid network service provisioning. This is the main focus of this paper and is discussed in detail later. Here we introduce the principal whereby the application profile in Cisco ACI is used additionally to define a service chain.

The service chain becomes an integral part of the policies defined for the inter-application tier communications. Here again, the configured functions in the service chain, such as firewalls and load balancing, are applied to the application tier flows on demand through association with the APIC application profile. For automated cleanup, these service function configurations are withdrawn with the application profile when it is decommissioned.

Figure 6: Cisco ACI Application Profile with Integrated Service Chain

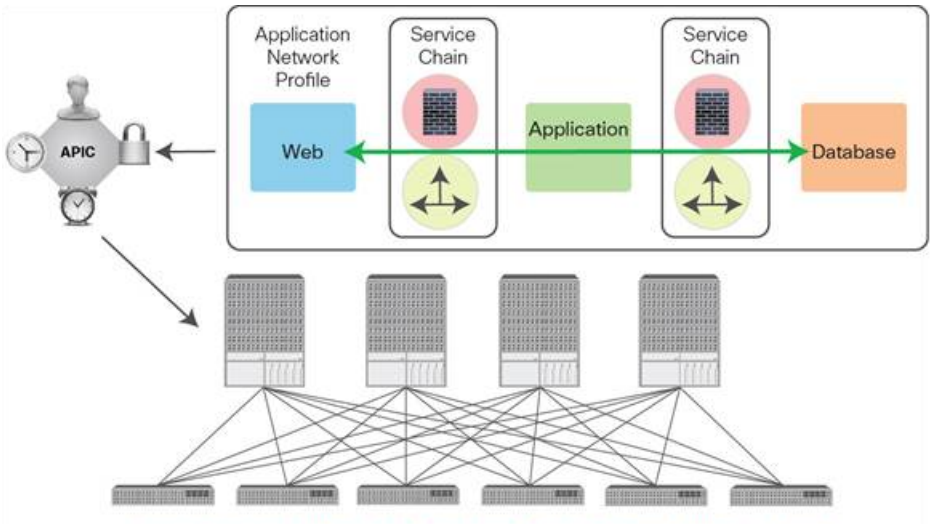

The application profiles are defined in terms of endpoint groups (EPGs) and the contracts between them. For example, in the illustration above, the application tiers such as Web, Application, and Database represent three EPGs, each having a number of servers—physical or virtual—associated with them. The servers create flows as part of the application tier interactions and these flows are allowed in the Cisco ACI fabric through contracts described in the Cisco ACI application profile. For example, the Application tier may initiate connections to the Database tier and so a contract is defined in the APIC describing what is allowed to pass between those two tiers, with filters and actions defined as part of the contract details.

Cisco OpFlex

As already mentioned, the Cisco APIC uses a new open southbound policy protocol for communicating with the Cisco ACI fabric elements. Those elements may include the AVS and the leaf switches to define the desired application profile communication policy.

OpFlex takes a different approach compared to previous southbound configuration and automation protocols used by network control systems, such as those used in other software defined network (SDN) and network virtualization solutions. Other solutions are all network centric. They define network configuration parameters and flow forwarding management centrally, and explicitly require detailed configuration to be designed and mapped out by the administrator. They also require huge tables of network flows and policies to be maintained in the central control plane.

In Cisco ACI the APIC is not, as previously explained, involved in flow forwarding. It does not have a routing table (RIB) or a forwarding table (FIB). The APIC instead

defines the application policies, as previously described, as simple profiles as an XML or JSON file. It then uses OpFlex to communicate with the leaf switches in the fabric the "Desired Communication Policy" between a set of application elements or tiers. This application profile does not tell the switches how to program the required communication. Rather, the application profile just expresses an end desired state and leaves it to the switches themselves to instantiate the configured elements. The switches have an **OpFlex Agent** built in to them that takes the desired state communicated by the APIC and turns this into a specific configuration on the switches. This feature is less burdensome on the administrator and far more scalable in terms of resource allocation in the controller function.

This model is known as a **declarative model**, as compared to the imperative model used by other solutions. The declarative model has many advantages including abstracting the complexity of the network configuration from the actual implementation so that, for example, an application developer could deploy network services for applications without knowing the specifics of the network infrastructure. The declarative model allows the application profile to be deployed anywhere in the fabric, wherever the tiers of the application using the profile happen to be connected. This ability makes for a flexible, on-demand anywhere-to-anywhere, policy-based network utility. A programming layer used for automation no longer needs to be aware of where application tiers are located. The network provisions the required profile policy based on the detected location of the application tiers and not based on a predefined explicit location set up by the network administrator.

Open Framework

The Cisco APIC framework enables a broad ecosystem and industry interoperability for Cisco ACI with management, Layer 4 through 7 services, orchestration, and virtualization across a wide range of vendors.

The APIC provides centralized access to Cisco ACI through an object-oriented representational state transfer (REST)ful API framework with XML and JSON. This framework allows Dev-Ops teams to innovate internal business application alignment with IT infrastructure. They can also accelerate ecosystem partner orchestration and management framework integration including open source solutions such as OpenStack.

Southbound APIs provide extension to Cisco ACI policies through existing virtualization platforms, Layer 4 through 7 services and networking components, and supporting compute and storage components in the future.

Cisco has submitted to the IETF the OpFlex protocol. The goal of the protocol is to maintain control of intelligence within the network for better scalability and abstraction rather than centralize it into another controller for imperative control as other SDN

solutions have done. Cisco is actively working with other vendors including Citrix to develop the declarative model further to extend it to be used by ADCs for partnering L4 to L7 services into Cisco ACI in the future. Cisco, together with its partners, plans to make the OpFlex draft an open source version available to further encourage broader adoption.

Scalable Performance and Multi-Tenancy in Hardware

As discussed, the OpFlex declarative model scales better than centralized imperative solutions. The declarative model scales out as more network and service elements are added to the fabric. As you add more elements, you add more control plane in a distributed fashion, instead of adding to the burden of the imperative, centralized SDN controller model or network virtualization model.

Cisco ACI is designed to scale to one million endpoints, that is, one million routes to one million application elements. The scale is determined by the tables held in the physical network elements and so performed in hardware rather than being held in the software-based controllers of alternative solutions. Cisco believes this is a more robust way to distribute load and ultimately provides a more stable, consistent performing fabric infrastructure as the network in the data center grows. Learning Layer 2/3 paths is performed in the data path with no additional latency for learning and lookup. Again, this represents a step forward over existing SDN and network virtualization solutions.

Cisco ACI is designed also to be multi-tenant in use. Infrastructure elements, application profiles, external routes, and VMM domains are assigned to tenants in the fabric and administered within the tenant for deployment and operation.

Each tenant maintains its own network containers called **contexts**. These contexts have their own address space, VLANs, and subnets. Layer 4 through 7 services are integrated into the Cisco ACI fabric and associated with application profiles. Services, VMM domains, and physical connectivity can be allocated to one tenant or shared by multiple tenants. But even when shared, logical elements such as subnets and policies can exist side by side without overlapping, interfering, or having knowledge of each other from an administration perspective.

The application profiles align to the elements within the tenant structure. For example, application profile EPGs exist in a container called a **bridge domain**, which has one or more subnets associated with it. The bridge domain lives in a given context under the main tenant hierarchy. This is shown in the following figure.

Figure 7: Tenant-Based Hierarchical Object Structure in Cisco ACI

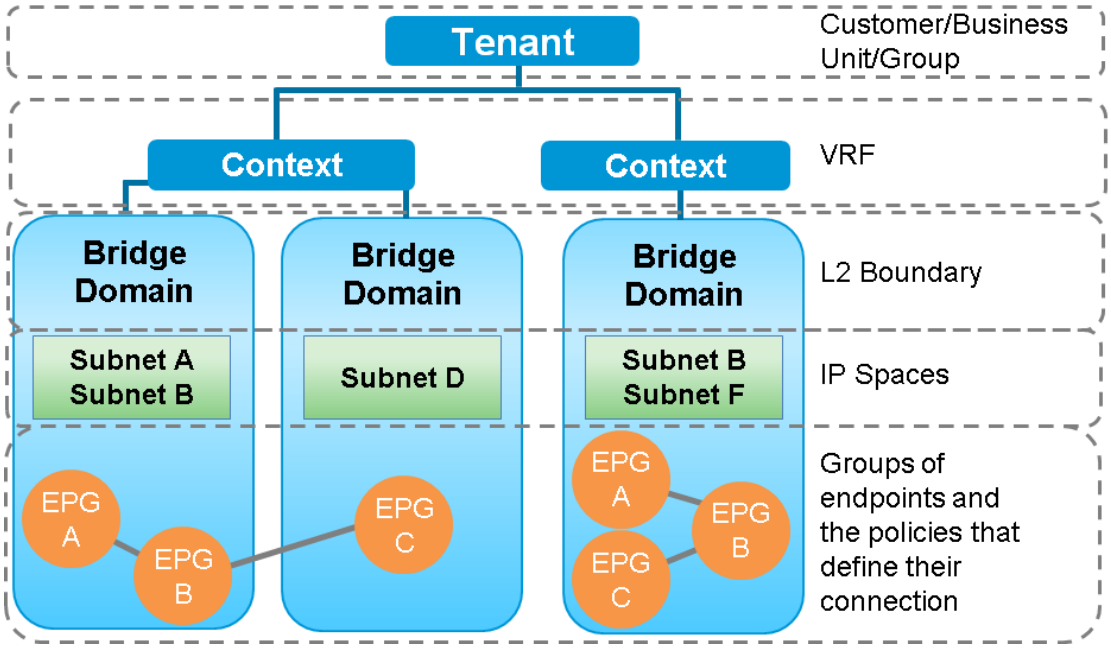

The tenant object is the top container as far as user partition goes. A tenant can be an organization, a department, a customer, or a given application. The way the partitioning is used is down to the best operational model for the customer. The context, as described, acts as a private network space within a tenant. The intention of this setup is that objects and constructs within a context do not communicate with another context unless some specific route leaking is configured. Within a context, there are multiple bridge domains. These bridge domains act as containers for network forwarding behavior for one or more subnets. The application profile elements and EPGs are assigned to bridge domains and thus their networking characteristics are defined therein. EPGs can be reused, as shown in the figure, in more than one context if required.

This approach allows for complete flexibility as well as control. This keeps services defined in terms of their behavior and networking characteristics, and uses traditional networking constructs such as VLANs and virtual routing and forwarding instances (VRFs), but maintains them under a tenancy hierarchy to enable a truly multi-user fabric.

Cisco ACI Advantages

Cisco ACI provides the following advantages:

Centralized policy-driven management:

- Integration of virtual and physical networking
- Reduction in time to deliver application service
- Automatic fabric deployment and configuration with single point of management
- Application-based networking solution that provides flexible, on-demand any-to-any provisioning with automation and programmability for agile IT
- Automation of repetitive tasks to reduce errors, fabric-wide object-based provisioning for integration with orchestration and duplicate deployments such as disaster recovery sites

Automated provisioning of networks and network services with tight integration:

- On-demand scale out and tear down network services for dynamic alignment with business requirements
- Network and services tightly bound to specific application deployments
- Promotion of consistent policies and services
- Real-time network visibility, telemetry, and application health scores
- Centralized real-time health monitoring for physical and virtual networks and their services
- Always-on visibility into application performance coupled with placement correlation
- Faster integrated troubleshooting with application-centric focus

Application agility:

- Application deployment independent of underlying infrastructure for ease of deployment, management, and migration
- Application focused network policy for closer integration between application and IT teams
- Management of application infrastructure lifecycle

Open and comprehensive end-to-end security:

- Open APIs and standards to enable software flexibility for Dev-Ops teams, firewalls, and ADC ecosystem partner integration
- Automatic capture of all configuration changes integrated with existing audit and compliance tracking solutions
- Detailed role-based access control (RBAC) with fine grained tenant and network context-based separation
- White list based forwarding rules, meaning only policies explicitly stated in the application profile allow flows to be forwarded

Citrix NetScaler ADC Overview

Data Center Challenges That Are Addressed by an ADC: Scalability, Security, and Resiliency

The modern data center is a challenging environment in which to deliver applications. “Consumerization” of IT means that there are many different device types, with different screen resolutions, and security capabilities. The various needs require features such as **content switching** optimization of mobile phones, tablets, and laptops, database optimization, or static and dynamic caching. All have the capability to access applications globally via the Internet or cloud and require load balancing of application servers to ensure resilience and responsiveness.

Businesses rely heavily on e-commerce and e-business. Some businesses run 24/7 entirely within the virtual world, requiring distributed, highly available, on-demand services. As engineers, we face the challenge of delivering these applications with a good end-user experience in a secure, reliable, scalable, and performance-oriented manner.

Application agility, mobility, and rapid deployment, mandated by hypervisor requirements, are key drivers within most data center designs. Businesses require data center infrastructure to dynamically respond to application needs as a result of changing requirements. NetScaler’s integration with the Cisco ACI policy-based, centrally managed architecture, facilitates this application-centric delivery. Cisco ACI currently supports the NetScaler 10.1 feature set.

From a scalability perspective, Citrix NetScaler ADC uniquely has the capability to scale-up, scale-in, or scale-out.

NetScaler is able to **scale-up** by using licensing. This ability allows the administrator to unlock, when necessary, additional performance, such as bandwidth or Secure Sockets Layer (SSL) offload up to the limit of the underlying device.

The **SDX platform** delivers **scale-in** capability with up to 80 virtual, hardware-based instances on larger NetScaler devices. This ability provides dedicated, configurable CPU, memory, and SSL offload per instance, fully supporting multi-tenant solutions and especially suitable in service-chaining with Cisco ACI.

NetScaler’s **clustering** ability makes it possible to **scale-out** to 32 individual NetScaler appliances (physical or virtual). These appliances work in concert to deliver one or more distributed applications, using a single system image across the cluster, with a single management and configuration point. An implementation can start small yet scale

beyond 3 Tbps in total capacity. Should one cluster member die, the other members can dynamically intervene to absorb its workload, bringing inherent resiliency.

NetScaler can also be configured as an active/standby **high-availability** pair, which brings design flexibility between data centers, as the devices can be deployed either in the same or different subnets.

In modern application delivery, **security** is paramount. Most countries have a legal requirement for data protection and processing. NetScaler mitigates multiple types of common denial of service (DoS) attacks, such as Internet Control Message Protocol (ICMP) floods (including Smurf attacks), TCP SYN floods, and HTTP Request attacks and performs a full TCP proxy of incoming traffic, further protecting back-end servers. With platinum edition software, there is also the opportunity to leverage AppFirewall. This intelligent, integrated L7 firewall has a rich set of security features, including SQL injection, TLS man-in-the-middle attacks, cross-site scripting, day-zero and buffer overflow attack mitigation and facilitating secure Software as a Service (SaaS). AppFirewall uses a hybrid security model, combining negative and positive models. The negative model uses signatures to detect known exploits, which are automatically updated. The positive model, uses patterns of good behavior to detect zero-day exploits and adaptive learning to observe application behavior and then dynamically create regular expression (regex) rules. The rule base can then be viewed within the GUI with the Visualizer.

NetScaler also supports multiple types of access lists, rate limiting, surge protection, and traffic domains, which allow creation of separate routing domains.

Design Mode Flexibility

Citrix NetScaler can be deployed in multiple modes, even within the same device. The most widely-adopted is routed-mode (L3), whereby NetScaler is inline and all traffic passes through the ADC in a two-armed topology, from a client-side subnet to server-side subnets. In this deployment, the real IP addresses of the servers are never visible to the client, thus ensuring security. NetScaler also performs a full TCP proxy by default, ensuring that real servers are protected from security attacks, such as SYN floods.

NetScaler can also be deployed in bridge mode (L2). In this scenario, even though there are still two distinct client and server side VLANs and NetScaler remains inline; all devices reside within the same subnet. This design type can be used where clients need to talk directly to servers with their real IP addresses, as well as to the Virtual IP address (VIP).

Figure 8: Citrix NetScaler Inline Design

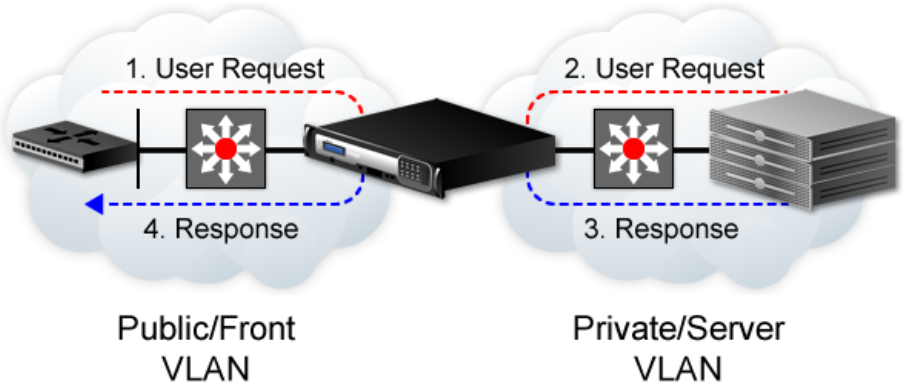

One-arm mode is particularly suitable when it is necessary to partition traffic into load-balanced and non load-balanced traffic. In this design, NetScaler is not inline, again facilitating direct client-to-server communications but without utilizing ADC resources. Load-balanced traffic alone passes through the ADC, typically via a Link Aggregation Control Protocol (LACP) port-channeled interface, which carries both client and server VLANs.

Figure 9: Citrix NetScaler One-Arm Mode Design

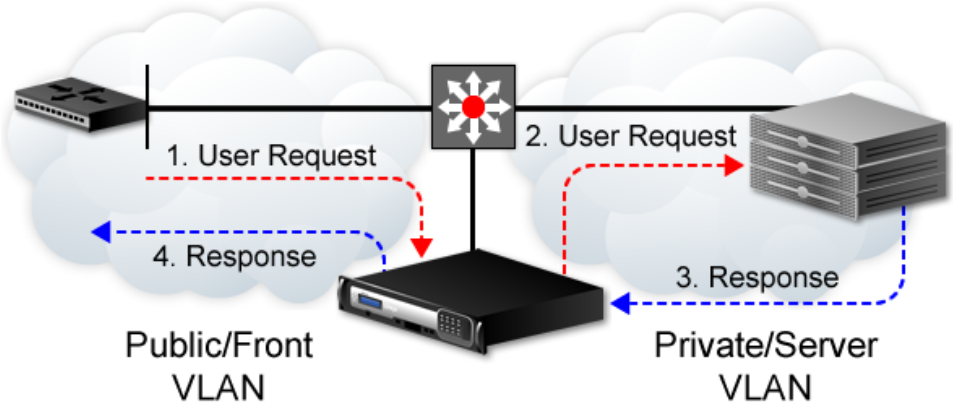

Direct server return is topologically similar to one-arm mode, however, traffic from the server in the return path does not traverse the ADC. The client would expect the return traffic to have the source IP address of the VIP (its original ingress destination). However, without intervention, the server would respond with its own real address, resulting in the client sending a TCP reset. To mitigate this issue, a loopback is configured to have the same IP address as the VIP on each server and the reply is sourced from this address.

There are some inherent constraints with this design, predominantly that the ADC will not see any return path errors issued by the server, such as *404 Page Unavailable* messages. This may result in a lack of awareness of server problems by the ADC. However, this design works well in certain application designs where very high, return path traffic ensues. An example would be content delivery networking or high-bandwidth applications like video.

Citrix NetScaler Product Portfolio

NetScaler is divided into three platforms. The first is **MPX**, which is a standalone, non-virtualized ADC but may be upgraded to an SDX. **SDX** is an ADC appliance but is divided into multiple, hardware-based, virtual instances. **VPX** is a virtual ADC and can be run on multiple hypervisors.

Figure 10: Citrix NetScaler Product Portfolio

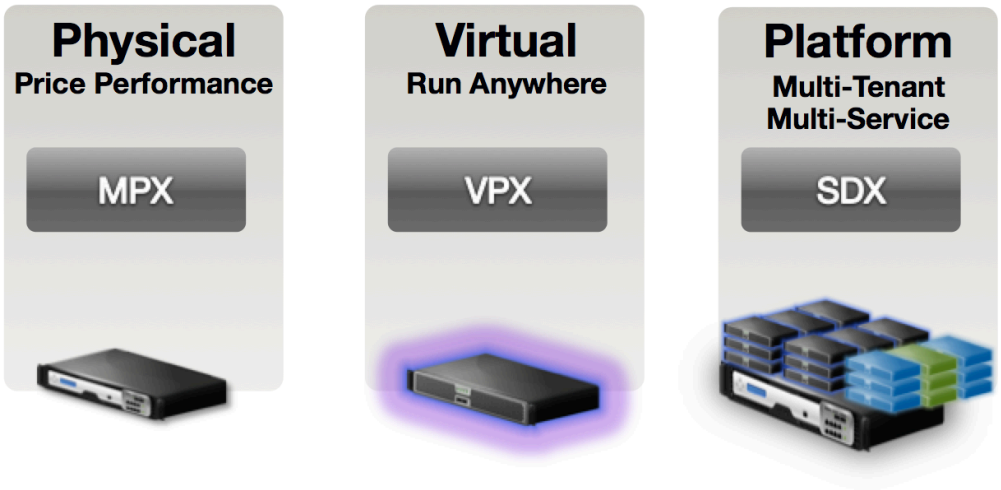

MPX has a wide choice of offerings with throughput capabilities ranging from 500 Mbps to 120 Gbps and between 17,000 to 560,000 SSL TPS.

SDX is designed to support multi-tenancy and enables fully isolated and independent NetScaler instances to run on a single appliance. Subject to the device's underlying hardware, SDX can scale from 5 to 80 instances on the largest appliances. Memory, CPU cycles, and SSL cards become resources, which can be divided and definitively assigned to different NetScaler instances. SDX instances can also be used to support third-party solutions including PaloAlto Firewall, Bluecat, CSE, and Websense.

VPX is a virtual ADC, which can run on a variety of hypervisors, including Xen, Hyper-V, VMware, or in the cloud in Amazon Web Services.

VPX can also be used as a NetScaler 1100V Virtual Service Blade (VSB) in **Cisco Nexus 1100 Series Cloud Services Platform** hardware appliance.

Figure 11: Cisco Nexus 1100 Series Cloud Services Platform

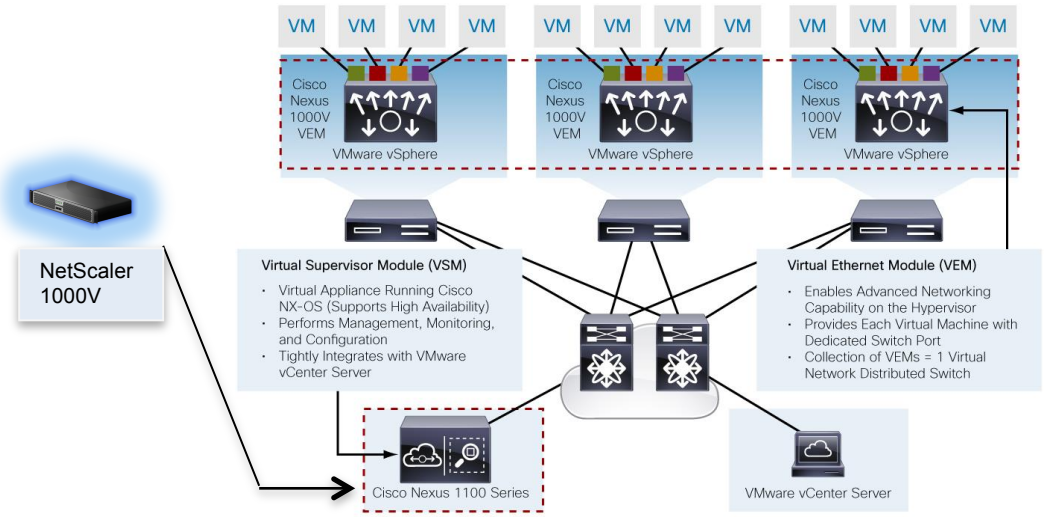

The cohesive solution integrated with the Cisco ACIs policy-driven, application-aware architecture enables design flexibility. This is because the Cisco Nexus provides a hypervisor agnostic solution, and a rich subset of VSBs such as the **Virtual Service Gateway** (VSG), which provides a firewall for east-west traffic and the **Cisco Prime Network Analysis Module** (NAM), which provides in-depth analytics.

All NetScaler platforms are able to integrate with Cisco ACI, whether they are physical or virtual, VSB or standalone.

Load Balancing and Content Switching for Mobile Devices

Citrix NetScaler has an impressive portfolio of load-balancing methods, allowing granular, transparent, fault-tolerant application support. Load balancing allows an application to be represented by a VIP, which could, for example, be the fully-qualified domain name (FQDN) of a web service. Utilizing a VIP means that the client still thinks that it is requesting content directly from the real server, but remains unaware of any changes or failures in the server-side architecture. A load-balancing algorithm is used to determine which server will respond, based on administrator-defined criteria. Citrix NetScaler has multiple algorithms available to support diverse application profiles and careful consideration should be given to decide which would be the most appropriate for each application.

Least Connections is the default method and is especially suitable for web-based applications, which have short requests but long responses. There are a variety of others, such as Least Bandwidth, for long-lived flows, like FTP, Least Response Time

for HTTP or HTTPS traffic, Token or Hash for caching or Database Server Load Balancing, and so on. Each focuses on optimizing application delivery.

To ensure that traffic is not sent to a server, which is either out of service or incapable of response, administrators configure monitors to verify server health state. Monitors can be as simplistic as a TCP or ICMP probe, or more complex, able to check header fields within HTTP or verify application payload.

Content switching allows intelligent application handling of HTTP and HTTPS traffic, beyond merely inspecting an IP address or port. Traffic can be selectively manipulated based on many diverse criteria, including device type, client browser type, ISO language code, geographic location, or character set. This ability is particularly useful when dealing with mobile devices such as smartphones or tablets. With those devices, it may be desirable to respond with content tailored to the device capabilities or when performing caching of static or dynamic content to maximize response times and improve bandwidth utilization.

Figure 12: Typical Content Switching Design for Static and Dynamic Content

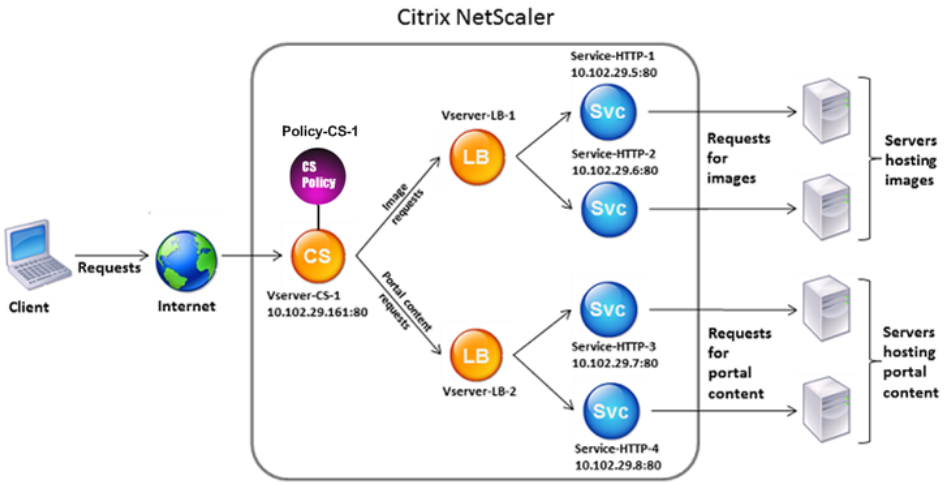

Database Optimization

DataStream also uses content switching for optimization of MySQL and MSSQL databases. (Oracle will be supported in the future.) DataStream allows for deep inspection of database queries and design consolidation.

NetScaler can perform the following DataStream optimizations:

- SQL multiplexing
- SQL connection offload
- SQL caching
- SQL aware policies
- Intelligent monitoring and failover

DataStream scale-up is achieved by converting large numbers of short-lived client connections into fewer persistent server connections, achieving TCP connection scale-up. Since NetScaler acts as a SQL proxy, it offloads the connection management from the server. This relieves database server CPU and memory for more critically needed database activity. Scaling up with NetScaler achieves server consolidation with tangible reduction in database server infrastructure.

Figure 13: Master/Slave Databases Being Intelligently Content Switched with NetScaler

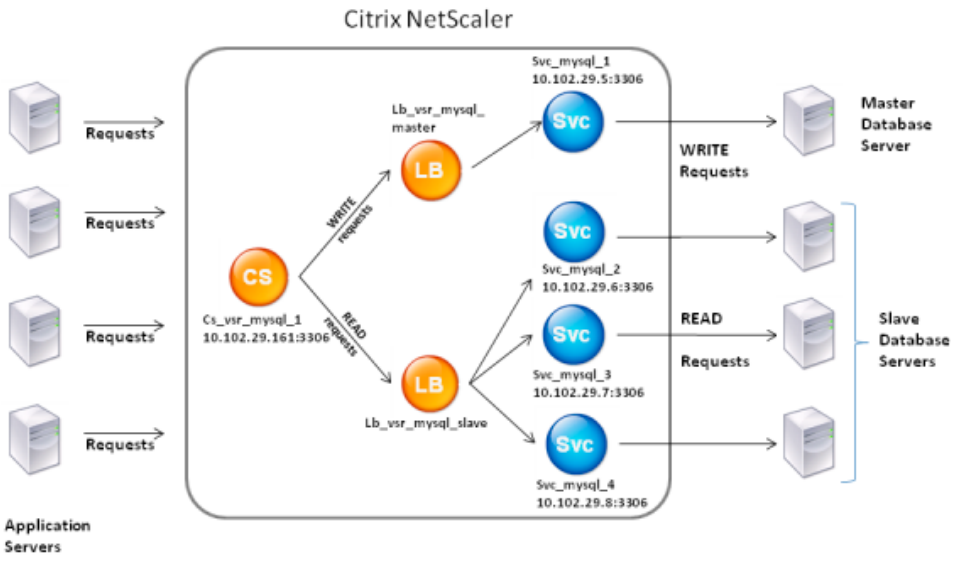

Intelligent SQL load balancing allows scale-out to optimally route SQL requests to the most available servers based on extensive health monitoring capabilities. NetScaler's SQL-aware content switching policy engine is able to perform the read/write split between the read and write server pools. NetScaler's SQL connection offload with SQL load balancing enables high performance, allowing the addition of more servers and expansion of website capacity, with no increase in latency.

A critical requirement is the ability to **monitor** the replication state of the database servers as to how far behind they are with respect to the master. Due to the nature of the unbalanced queries, where some queries take minutes to complete versus others executed within milliseconds, it is essential to have a monitoring solution that is aware of this condition. NetScaler, in addition to its VIP-based automated IP failover, provides for a replication-aware monitor that can estimate the amount of time the slave server is behind the master. This replication-aware monitoring, in conjunction with the SQL request switching, enables NetScaler to redirect SQL queries to other active servers in the pool, while allowing existing queries to the slow server to run to completion without aborting application connections.

SSL Offload

SSL offload is a commonly deployed feature of ADCs. As the Internet of Everything (IoE) becomes more commonplace, SSL plays an increasingly greater part in protecting businesses by encrypting secure communications, such as credit cards and even access to internal corporate architecture. Whenever you use encryption, there is, of course, a significant impact on the devices involved. This impact includes additional design complexity to ensure that traffic can be inspected "in the clear" by a firewall or intrusion detection system (IDS) or intrusion prevention system (IPS) device, at some point. NetScaler MPX and SDX have impressively high, hardware-based SSL offload capabilities. In SDX, it is possible to configure dedicated CPU, memory, and SSL offload per instance, ensuring true isolation and guaranteed resources.

Figure 14: SSL Traffic Being Decrypted to HTTP

Integration of NetScaler with **AppFirewall** brings additional benefits. In this scenario, it is possible to perform SSL offload in hardware with either MPX or SDX and then, while still in the NetScaler backplane, to direct the decrypted traffic for inspection by the integrated AppFirewall. Once this has been performed, the traffic can then be encrypted again for SSL termination on the back-end server. The entire client-server traffic path is encrypted but with the benefit of firewall inspection.

Integrating with the Cisco Nexus 1100 Series Cloud Services Platform

Because VPX is deployed as a virtual machine (VM), it cannot deliver hardware-based SSL offload, but it is capable of terminating lower levels of SSL in software.

When VPX is integrated into a Cisco Nexus 1100 Series Cloud Services Platform hardware appliance as a VSB, it is then possible to perform hardware-based SSL offload with VPX by adding an SSL crypto card to NetScaler 1000V.

Integrating Citrix NetScaler 1000V and Cisco ACI brings new levels of control of Layer 2 to Layer 7 services. NetScaler 1000V uses the APIC to programmatically automate network provisioning and control, based on the application requirements and policies. Cisco ACI defines a policy-based service insertion mechanism for both physical and virtual NetScaler 1000V appliances.

Simplicity of Configuration with AppExpert Engine

The NetScaler system uses policies to evaluate specified conditions and to define actions to be taken if conditions are met. A policy is always the following:

Rule: Rules are logical expressions that are used to evaluate traffic or another object and implement if-then-else logic.

Action: Actions depend on individual features. Multiple actions can apply to the same traffic flow.

Binding: Bindings activate the policy on NetScaler entities, such as virtual servers.

Policies provide the foundation for the behavior of most NetScaler features, enabling the features to interpret the data, such as SSL requests, HTTP requests, TCP requests, and HTTP responses that pass through it. Typically, if a policy matches the data being evaluated, the NetScaler system takes one or more actions that are associated with the policy. For example, if an HTTP request matches a rewrite policy, the NetScaler can take an action to change (rewrite) the information in the request.

Most NetScaler features have built-in policies. The number and type of user-defined policies that a feature requires differs for each feature based on implementation requirements.

A NetScaler feature can use one of two policy types:

> **Classic policies:** Classic policies evaluate basic characteristics of traffic and other data. For example, classic policies can identify if an HTTP request or response contains a particular type of header or URL.
>
> **Default policies:** Default policies can perform the same type of evaluations as classic policies. In addition, default policies enable you to analyze more data (for example, the body of an HTTP request) and to configure more operations in the policy rule (for example, transforming data in the body of a request into an HTTP header).

Given the number of features that a traffic flow may be subject to and that multiple traffic flows will require different features, a framework is needed to define how traffic is matched and how actions are applied.

NetScaler has a built-in **expression builder** and both **expression and action evaluators**. These features simplify the creation of policies and allow the administrator to check the outcome of either the policy or action before they are associated. There are also additional features within **AppExpert** that minimize CPU and memory utilization, such as **pattern sets** and **string maps**.

When integrated into Cisco ACI, instead of the policy being defined locally on each device, the policy will be defined within the APIC and then pushed to the NetScaler estate via a **device package**. The device package interprets the configuration commands for the NetScaler operating system. In future, device packages will not be necessary as vendors move to support OpFlex.

Integrated Caching, Cache Redirection and Compression

Caching: NetScaler can perform both integrated caching or/and cache redirection. It has the largest integrated cache in the industry, with up to 48GB of configurable memory. The cache can be used to serve web content to users without requiring a round-trip to an origin server.

Figure 15: Benefits of NetScaler Caching

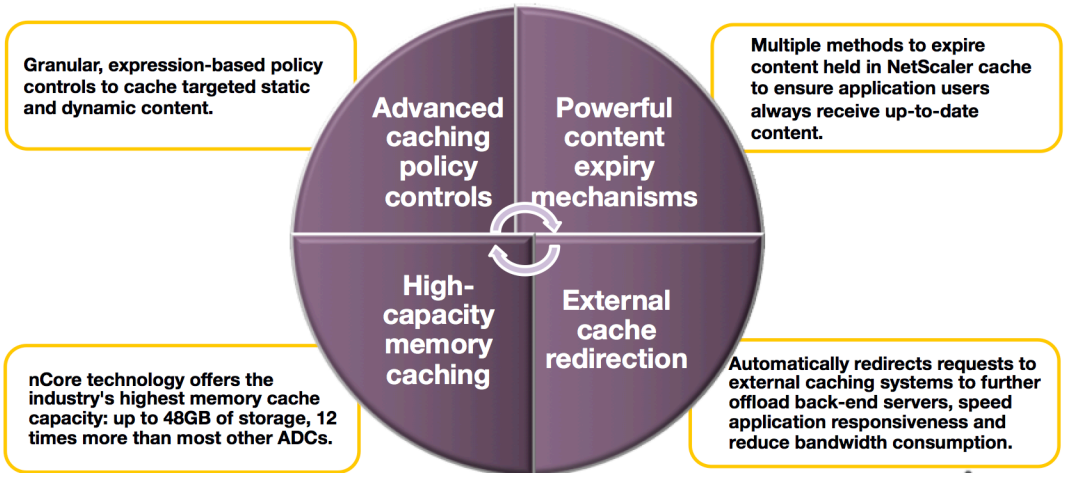

For **static content**, the integrated cache requires only that the feature is enabled and that basic setup, such as optionally specifying the amount of memory, is performed. The integrated cache uses **built-in policies** to store and serve specific types of static content, including simple web pages and image files. You can also configure the integrated cache to store and serve **dynamic content** that is usually marked as non-cacheable by web and application servers, such as database records and stock quotes.

You can tune the performance of the integrated cache, using methods such as preloading cached objects before they are scheduled to expire. To manage cached data once it leaves the NetScaler appliance, you can configure caching-related headers that are inserted into responses. The integrated cache can also act as a forward proxy for other cache servers.

NetScaler can be configured with a rate limit to specify how much of the integrated cache is used. Once that value has been reached, it can then redirect content requests to external caching servers.

Compression: Typically in web traffic, 80 percent of data comes in the response from the server. Compression can be invoked at server level but this is CPU intensive. NetScaler's compression feature compresses the size of HTTP responses from the server to a compatible browser (this is determined via handshake), improving performance by reducing the download time, saving bandwidth, and encrypting the data.

Once compression is enabled, NetScaler invokes built-in compression policies and compresses all traffic that matches these policies. You can create bespoke

compression actions and policies and bind the policies globally or to particular virtual servers.

The greatest benefit with compression comes from fewer bytes on the wire on slow-speed links, as the client has to wait less time for the complete response.

High Availability and Clustering Solutions

High availability: In high availability mode, NetScalers are deployed as an active-standby pair in either a different or the same subnet. The primary node services connections and performs load balancing. The secondary node monitors the primary by sending periodic heartbeats to confirm that it is able to accept connections, but does not pass application traffic. If a heartbeat fails, the secondary node retries the connection for a configurable period, after which it determines that the primary node is not functioning normally and failover occurs. All clients must reestablish their connections to the back-end servers, but session persistence rules ensure that the clients remain connected as before.

Failover can be forced if the primary node is able to determine that the status of the secondary node is up, the health of the secondary node is good, and is not configured to stay secondary.

Configuration is automatically synchronized but can also be forced at any time.

Figure 16: Device Failover

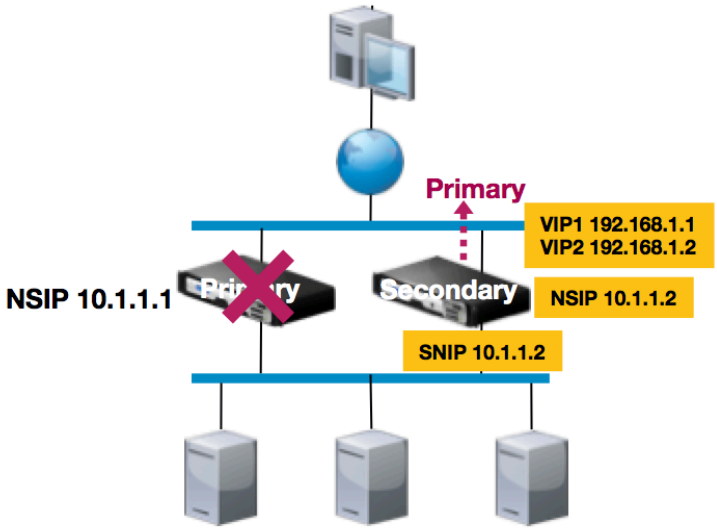

Clustering: In cluster mode, from 2 to 32 nCore, (including SDX instances) or VPX devices can be deployed to work as an all-active group with up to 3-Tbps throughput

with inherent fault tolerance. Devices can be added seamlessly to the cluster to give dynamic scalability. All devices must be of the exact same hardware and software build.

Within each cluster, one device is elected to be the Cluster Coordinator (CCO), which is the single configuration point and then the entire cluster is managed via one IP address, the Cluster IP (CLIP). All devices are connected physically via a backplane and each has a unique cluster ID. VIPs, mapped IPs (MIPs), and subnet IPs (SNIPs) can be common across the cluster (striped) or local to only one cluster node (spotted).

Load balancing of traffic to the cluster can be achieved in several ways:

1. Equal cost multipath (ECMP), using a routing protocol (L3), which allows the use of a routing protocol to distribute traffic
2. Cluster Link Aggregation Protocol (CLAP) (L2), which allows multiple interfaces to be bound together inside a bundle
3. Linksets, which allow non-directly connected devices to use the interfaces specified in the linkset to communicate with the external network through the cluster backplane

Once a device has been selected to receive the flow (Flow Receiver), it will then pass the traffic to the cluster backplane and the internal cluster logic will choose a node to process the flow (Flow Processor).

Figure 17: Flow Distribution within a Cluster

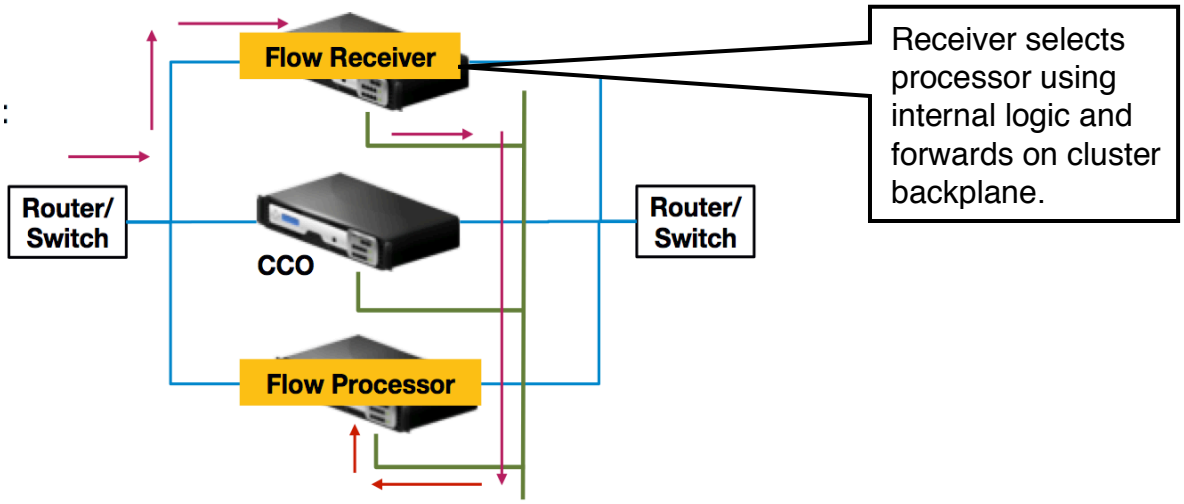

Fault tolerance within the cluster is a function of aggregate throughput per device. For example, if you have a three-node cluster and each node is working at 66 percent of its capabilities and subsequently, one node fails, then there would be redundancy because

the remaining two nodes would each be capable of processing the 66 percent that was being processed originally by the failed node, 33 percent each. The same would not be true if they had been working at 80 percent of their capabilities. Cluster fault tolerance requires consideration within the design phase of a project.

Citrix NetScaler Advantages

Citrix NetScaler provides the following features and advantages:

Scalability, security, and resiliency:

- Multi-dimensional scalability
- Integrated security with AppFirewall
- Resiliency with fault tolerance and clustering

Design mode flexibility:

- Flexible design at either Layer 3 or Layer 2 designs
- Inline or not inline deployment support
- Support for applications, which require high bandwidth

Broad, price-focused product portfolio:

- High performance
- Physical, virtual, and hardware-based virtualized platforms, suitable for any topology
- Feature parity across all platforms
- Active-standby, fault tolerance, and clustering of any NetScaler platform
- Cisco ACI integration across all platforms

Rich feature set:

- Comprehensive load balancing and content switching capabilities
- Intuitive database optimization
- High performance SSL offload
- Integration with Cisco Nexus 1000V in the NetScaler 1100 Cloud Services Platform
- Intelligent, GUI-based configuration and management tool with AppExpert
- Largest integrated cache within the industry and also cache redirection support

Integrating Cisco ACI with Citrix NetScaler Platforms

Cisco ACI integrates with NetScaler Series MPX, SDX, VPX, and NetScaler 1000V.

There are two key fundamental aspects to the integration of Citrix NetScaler with Cisco ACI:

1. The integration of the network provisioning for Citrix NetScaler
2. The automated provisioning of Citrix NetScaler function configuration with the Cisco ACI application profile coupled with integrated application infrastructure telemetry

Cisco ACI Integrated Network Provisioning for Citrix NetScaler

Traditionally when services are inserted into data center networks there is a high degree of manual effort and complicated stitching together of network elements such as VLANs (Layer 2) or VRF (Layer 3) to make the services accessible to the data flows. This makes location of services both rigid and inflexible plus there are complications associated with changes that add risk and delays in provisioning cycles and headaches for the infrastructure teams.

Services have become localized as a result, positioned close to where they are required to avoid the complexity of running VLANs and VRFs between racks and rows across the data center. This, of course, can mean lower consumption of resources due to low sharing ratios.

Additions, moves, and changes become difficult to deploy without risk, and policies tend to remain configured in services appliances long after the applications themselves are decommissioned. This results from the fear of the consequences of removing the wrong policy configurations.

A solution that offers the flexibility to use localized services or centralized shared services with automated network Layer 2 and Layer 3 provisioning based on consumption provides a better overall solution. The advantages of this solution include the following:

- Lower operational costs and time to deploy, no need to manually plan and configure individual switches with service insertion
- Flexible, anywhere location of services for any application flows automatically stitched into any part of the fabric based on need
- Automatic, dynamic provisioning of networking elements for multiple services built into the application profile, prepared in advance, to be deployed on demand through manual or programmatic orchestration tools

- Automatic removal of networking elements and service elements when the application profile is decommissioned
- Integrated with hypervisor operations for high availability and mobility

Cisco ACI integrates the deployment of networking constructs such as interfaces, VLANs, and VRF routing along with the deployment of L4 through L7 services in the fabric. Interfaces on the appliances and leaf switches are automatically configured with the correct VLAN constructs when the service elements are instantiated through the policy-based deployment. There is no need to manually configure the network. All the elements are defined as part of the services deployment for a given application. They are automated into the provisioning to save time, prevent errors, and allow programmatic deployments to be built. VLAN pools are assigned from which the service-specific VLANs are instantiated.

Figure 18: Automated Provisioning of Network for Services Insertion

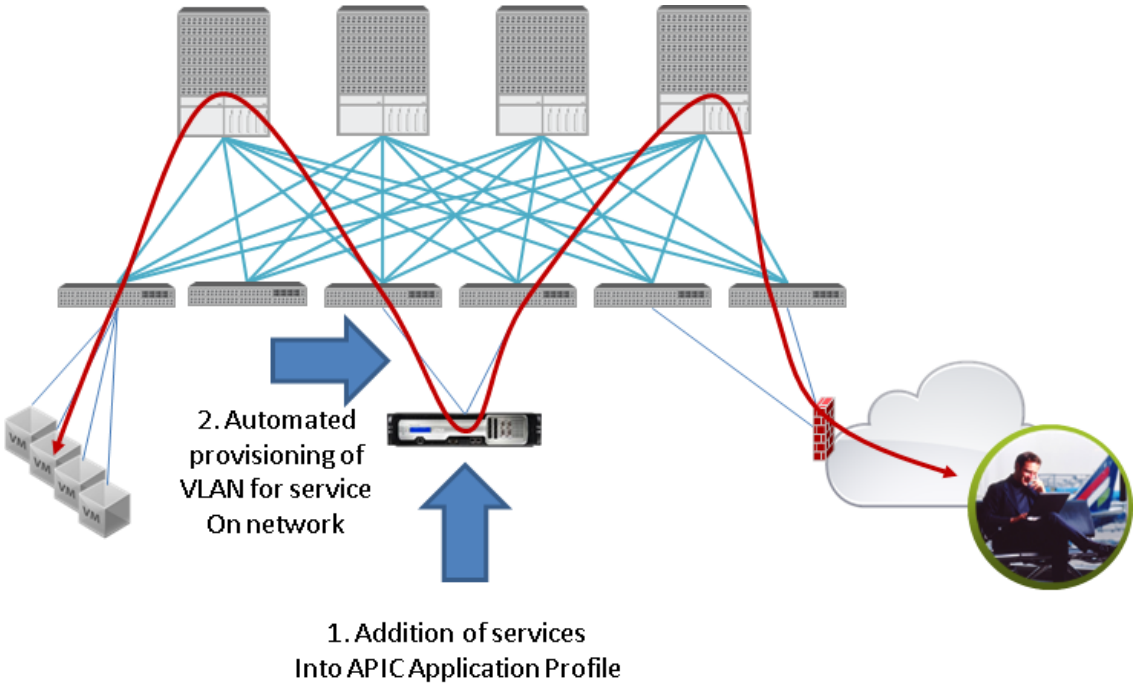

Automated Provisioning of Citrix NetScaler with Cisco ACI Application Profiles

Within the Cisco APIC, the Citrix NetScaler device is integrated with the Cisco ACI fabric. Citrix has developed a device package for Cisco ACI that is uploaded into the APIC. The package provides all the required scripts for provisioning, configuring, and

editing the NetScaler for server load balancing, SSL acceleration, remote access, content switching, web compression and other services.

There is one device package for Citrix NetScaler MPX, SDX, and VPX Series and another device package for the NetScaler 1000V.

With the Citrix NetScaler device package, the APIC can be a single point of configuration and provisioning for the network and service elements of an application profile. This tight integration is a key strategy for reducing time to administer and coordinate provisioning through automation services on the Cisco ACI data center fabric.

Figure 19: Citrix NetScaler Device Package for Cisco ACI

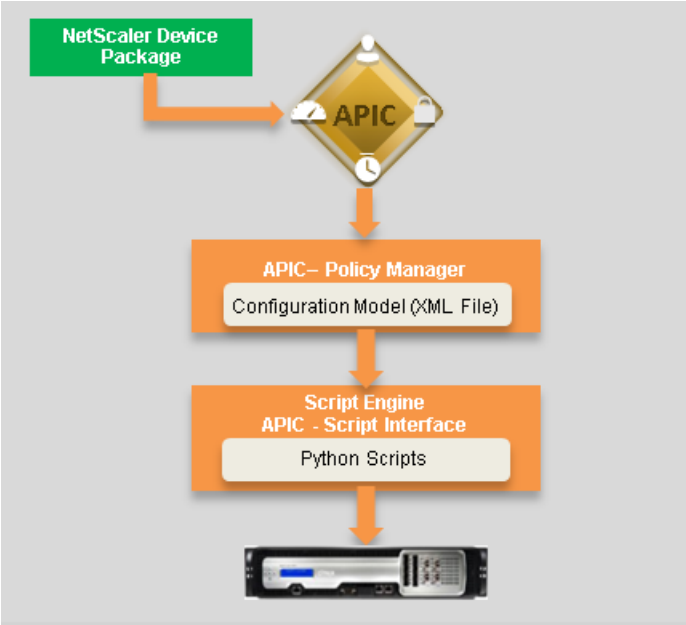

Once the Citrix NetScaler devices are initially configured and have a management address that the APIC can communicate with, the device package can be installed. Then the APIC can begin to create to logical device instances used to define services within the NetScaler as part of a service chain. The following steps outline the simple process of integration and configuration:

Step 1: Install the NetScaler devices.

Step 2: Configure management interfaces and credentials.

Step 3: Install the NetScaler **device package** in the APIC.

Step 4: Create a **device cluster** in the APIC for the APIC to manage the NetScaler.

Step 5: Within the APIC, define **logical interfaces** associated with the device cluster and define a pool of VLANs that the logical interfaces use for connectivity.

Step 6: Define the **service graph**, including the functions to be configured on the NetScaler for a given service deployment.

Step 7: Associate the service graph and device cluster with a **logical device context** with interfaces from Step 5 and add to the **application profile contract**, to make the NetScaler part of the Cisco ACI application profile when it is provisioned.

As previously discussed, the APIC provisions network communications by means of application profiles. The application profiles are made up of EPGs and contracts. Essentially, the EPGs represent tiers of servers in an application. The contracts are rules regarding how servers in EPGs can communicate with servers in other EPGs. In the example shown in the figure that follows, the Web servers want to communicate with the App servers. They can do so according to the contract being provided by the App EPG.

When adding services into the path between EPGs, a service graph is created as described and associated to the provided contract in the providing EPG. In the example shown in the figure, the service graph containing the logical device context, using the load-balancing function configured by the APIC on the NetScaler, is mapped to the App EPG provided contract. This causes the configuration within the service graph to be applied to the traffic flows from the Web servers in the Web EPG communicating with the App servers in the App EPG. These flows are load balanced according to the rules in the service graph configuration, along with the contract rules configured in the application profile.

In this way, Citrix NetScaler services can be configured with and integrated into Cisco ACI application profiles in a tightly bound fashion. These profiles can be applied to specific applications, edited and managed during the life-cycle of the application, and maintained as part of the telemetry and health of the application reporting through the single management point of the APIC interface.

Figure 20: Service Integration with Application Profile

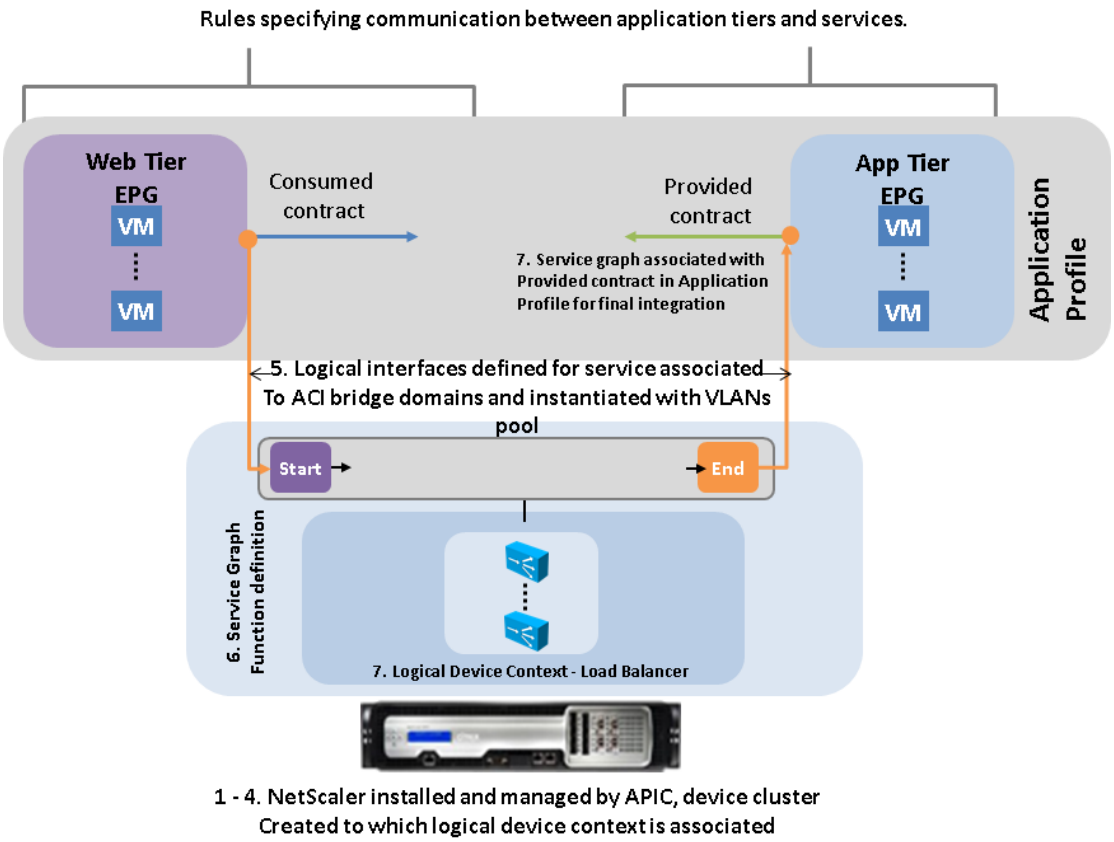

Citrix NetScaler ADC and Cisco Application Profile Integration Summary

- Citrix NetScaler ADC is explicitly configured and provided network connectivity on demand by Cisco ACI APIC using device packages written by Citrix and imported into the APIC.
- There is a single point of management and deployment to foster programmatic orchestration.
- Cisco ACI defines a configuration policy within the Citrix NetScaler ADC and applies it to a service within the Cisco ACI application profile using a service graph tool in the APIC.
- L4 to L7 policy and network provisioning are coupled into the Cisco ACI application profile for dynamic instantiation into Cisco ACI fabric to provide a tightly coupled complete network service for a given application.
- There is no need to manually and separately configure network and L4 to L7 services, and no need to stitch network interfaces together using VLANs. The Cisco ACI application profile contains all the requirements.

- The Cisco ACI application profile and NetScaler-configured service graph allow for flexible deployment anywhere in the fabric, with cleanup of all configuration when and application is decommissioned and the Cisco ACI application profile is removed.
- Integrated telemetry and application health association include the service elements in health scores associated with a given application profile.

Future Developments

In the future, the device package will be replaced or complimented with the use of the OpFlex protocol in the Citrix NetScaler. With OpFlex, the APIC will be able to send declaratively a desired end state to the NetScaler device. The device will then be able, together with an integrated OpFlex agent, to interpret the end state request and instantiate the functional configuration to be used in the service graph. This ability will allow for a simpler more abstract policy to be configured in the APIC, similar to the EPG and contract model used today but leaning towards the service function. This abstract policy leaves it to the NetScaler device to interpret the requirements independently, allowing, for example, new features to be commissioned that are not explicitly understood by the APIC. This ability allows these new features to become much more quickly adopted for automation than traditional imperative models of orchestration that always lag by several months or even years.

Better Together: Cisco ACI with Citrix NetScaler ADC

Cisco ACI integrates tightly with Citrix NetScaler ADC physical and virtual appliances to do the following:

- Reduce deployment complexity.
- Improve the way that applications are deployed with existing and next-generation data center infrastructure.
- Protect investment in existing infrastructure including of NetScaler including NetScaler SDX, NetScaler MPX, and NetScaler VPX, and existing networking gear NetScaler 1000V.
- Operational cost savings by reducing time to deploy production infrastructure for new application implementations.
- Consistent automation of Citrix NetScaler services on any device type and anywhere in the customer’s data center networking environment without causing network disruption and interruption to existing services.
- Help customers move to a more agile, automated, application-centric data center infrastructure model where services and network integration are implemented independent of location and encompass both the hardware and software infrastructure elements.

Corporate Headquarters
Fort Lauderdale, FL, USA

Silicon Valley Headquarters
Santa Clara, CA, USA

EMEA Headquarters
Schaffhausen, Switzerland

India Development Center
Bangalore, India

Online Division Headquarters
Santa Barbara, CA, USA

Pacific Headquarters
Hong Kong, China

Latin America Headquarters
Coral Gables, FL, USA

UK Development Center
Chalfont, United Kingdom

About Citrix
Citrix (NASDAQ:CTXS) is a leader in mobile workspaces, providing virtualization, mobility management, networking and cloud services to enable new ways to work better. Citrix solutions power business mobility through secure, personal workspaces that provide people with instant access to apps, desktops, data and communications on any device, over any network and cloud. This year Citrix is celebrating 25 years of innovation, making IT simpler and people more productive. With annual revenue in 2013 of $2.9 billion, Citrix solutions are in use at more than 330,000 organizations and by over 100 million users globally. Learn more at www.citrix.com.