

Integrating Application Delivery Solutions into the Data Center Infrastructure

What You Will Learn

IT organizations are being challenged to provide employees with a growing number of business-critical applications, often with strict requirements for high availability, low latency, and scalability. Many companies are currently working on data center consolidation and virtualization projects with the goal of meeting these requirements while reducing costs and adding the flexibility needed to manage the continuous evolution and migration of applications to meet their business needs.

Application delivery devices are crucial components of the data center infrastructure, significantly improving application performance and resiliency by providing load balancing, server offloading, application optimization, and security features. Although standalone application delivery appliances are excellent for many environments such as small to medium-sized server farms or hosting environments with dedicated hardware for individual customers, more demanding data centers can greatly benefit from the deployment of application delivery devices that natively integrate with the switching and routing data center infrastructure.

This document examines the value of using integrated application delivery devices to meet the challenges of delivering hundreds of dynamic business-critical applications while containing costs and simplifying infrastructure. In particular, it describes the Cisco® solution. The Cisco ACE Application Control Engine Module, which provides the highest levels of performance and scalability, secure virtualization, and role-based administration capabilities, is natively integrated with the Cisco Catalyst® 6500 Series Switches and Cisco 7600 Series Routers. This document describes some of the solution's main integration features from a technical standpoint and its benefits in comparison to those of traditional appliances: cabling, power, and cooling savings; greater availability; higher performance; and enhanced security.

Challenges of Today's Application Infrastructure

Enterprises are challenged to adapt their data centers to efficiently deliver hundreds (and in some cases, thousands) of business-critical applications to their employees, regardless of their physical location, providing high availability, scalability, and security.

As business processes continue to be automated and streamlined, applications continue to grow in number and evolve, incorporating new functions or adapting to more efficient technologies such as the Web and Extensible Markup Language (XML). At the same time, many applications continue to be developed and improved using proprietary protocols or industry-specific standards (for example, Financial Information Exchange [FIX] for financial applications and Picture Archiving and Communication Systems [PACS] for the healthcare industry), creating a mix of technologies and requirements that need to coexist and transparently work together to help ensure the success of the business. In addition, in many environments, applications are still deployed and managed in separate silos across the network, with performance and security often being secondary concerns.

All this adds to the complexity and dynamic nature of application requirements for Layer 2 and 3 infrastructure (for example, routing and switching) and Layer 4 to 7 services (server load balancing, application delivery and acceleration, security, etc.). To address this complexity, IT organizations want new solutions that allow them to achieve these goals:

- Easily deploy and migrate new applications
- Scale the application environment whenever needed
- Handle a very distributed workflow across different IT teams
- Consolidate functions and devices
- Increase application performance as seen by the end users, be they employees, partners, or customers

To improve efficiency and reduce costs, many data centers are being migrated and consolidated, taking advantage of more powerful devices as well as virtualization technologies in many layers of the data center:

- At the networking level, with VLANs and virtual routing and forwarding (VRF) instances
- At the security level, with virtualized firewalls, such as the Cisco Catalyst 6500 Series Firewall Services Module
- At the storage level, with virtual storage area networks (VSANs)
- At the server level, with virtual operating systems running on a single physical machine

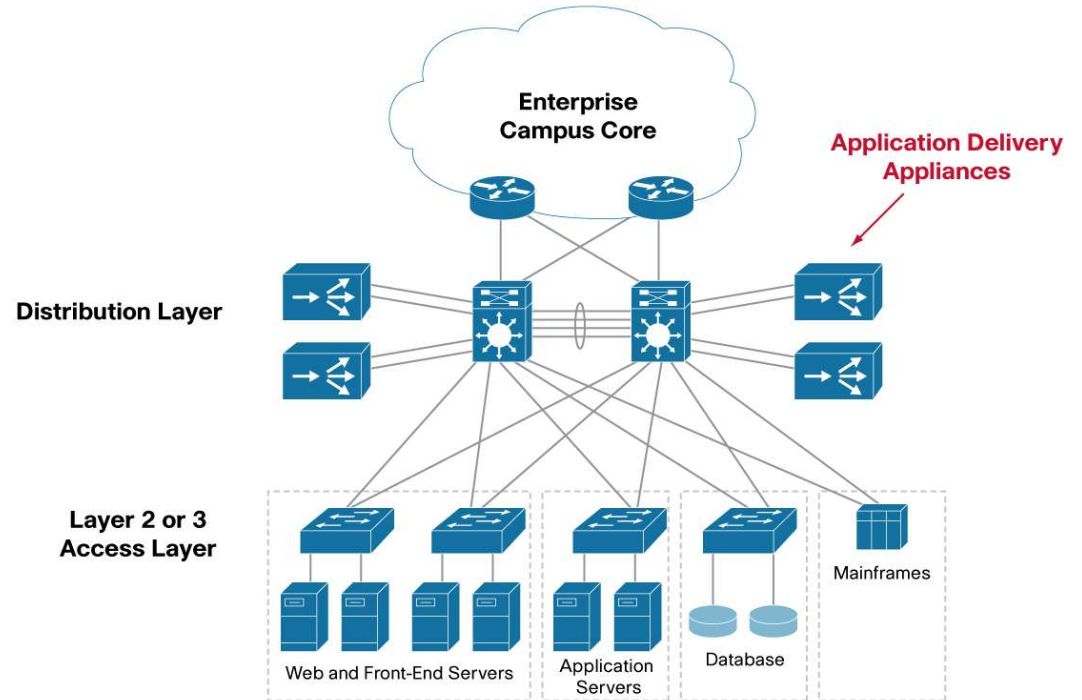
Application delivery devices are crucial components of this infrastructure, especially given their influence on application performance and their close interaction with both servers and applications. In data center environments, scalability and virtualization are not the only requirements that an application delivery device needs to meet: Transparent network integration is also a primary requirement, to improve application rollout and migration, as are high availability and scalability.

Traditional Solution

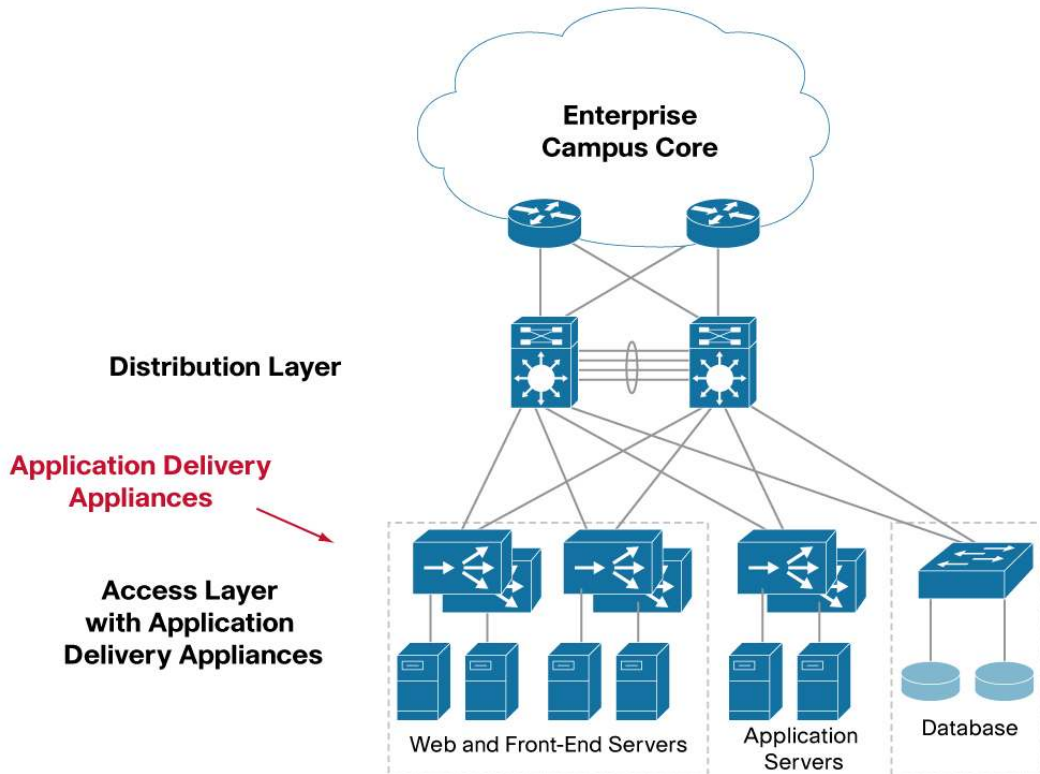
The traditional solution has been to rely on standalone application delivery appliances, typically connected to Layer 2 to 3 switches in the distribution layer. Multiple physical links are required to protect against link failure, and dedicated links to redundant Layer 2 to 3 switches are often a requirement to account for all possible link and device failures. When additional performance is required, or when specific applications or organizations require dedicated devices, the solution includes additional physical devices connected to the same Layer 2 to 3 switches.

Figure 1 shows a data center with two pairs of redundant load balancers connected to the distribution layer and serving requests for one or more tiers of servers.

Figure 1. Data Center with Application Delivery Appliances Deployed in the Distribution Layer



In a variation on this approach, servers are connected directly to the appliances. Dedicated appliances are used in redundant pairs and associated with individual server farms, connecting upstream to redundant switches in the distribution layer, as shown in Figure 2.

Figure 2. Data Center with Application Delivery Appliances Deployed in the Access Layer

Although some vendors do sometimes recommend this type of design, and although it may at first appear to be simple and fast to deploy, it poses many critical challenges:

The application delivery appliance needs to have enough physical interfaces to support a growing numbers of servers within the same server farms.

- Even when the appliance is modular, it likely will not offer the same port density and choice of media and speed as a dedicated Layer 2 or 3 switch, thus forcing trade-offs in terms of connectivity and requiring more devices.
- Application delivery appliances rarely implement all the Layer 2 features required in today's demanding access layers, thus often limiting security or high-availability features or introducing compatibility concerns.
- The application delivery appliance must handle all the traffic to and from the servers, not just the flows that require advanced Layer 4 to 7 services.
- A server that needs to be allocated to a different application or server farm may need to be physically reconnected to a different device.
- As data centers consolidate and applications evolve, transitioning applications and servers becomes very costly because physical changes and network redesigns are often required.

Overall, the number of application delivery devices required to support a given application infrastructure is significantly higher and the design less adaptable to changes than with an integrated system, thus making this solution less flexible and more costly in terms of capital and operational expenditures.

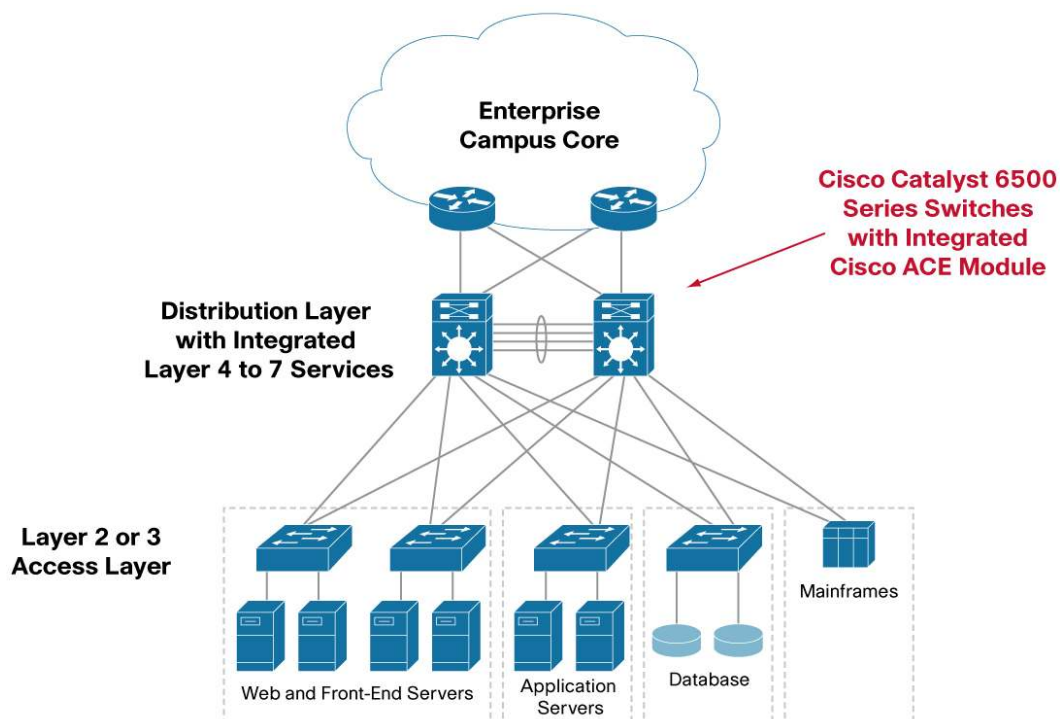
A Different Approach: Network Integration

A different approach that addresses many of today's data center application infrastructure challenges fully integrates Layer 4 to 7 application delivery services within the routing and switching infrastructure, especially within modular switches. Whether this is accomplished in software alone or with dedicated hardware, the resulting design is much simpler, reducing the need for extra cabling and dedicated power supplies, and much more flexible, allowing rapid changes without the need for any physical modification of the infrastructure. The user retains full control over which flows have to be processed for load balancing or application delivery functions, while all other flows are simply Layer 2 or Layer 3 switched without any affect on performance. In addition, specific integration features can enhance even further the collaboration between the Layer 2 and 3 functions and the application delivery engines and deliver additional security and resiliency.

Cisco delivers such a solution with its unique integration of the Cisco ACE Module into the Cisco Catalyst 6500 Series Switches.

Cisco has pioneered the network integration approach with the services modules for the Cisco Catalyst 6500 Series modular switches, introducing the Cisco Content Switching Module (CSM) in 2000 and recently adding the Cisco ACE Module, which enhances integration benefits by including full support for virtualization and role-based network management. Figure 3 shows the topology when the redundant Cisco Catalyst 6500 Series Switches in the distribution layer of a data center are equipped with Cisco ACE Modules.

Figure 3. Data Center with Cisco ACE Module Integrated into the Distribution Layer



Some of the benefits of this approach are immediate and directly related to the design:

- **Cabling:** Less cable needs to be connected and managed, with no need to recable to accommodate redesigns. The Cisco ACE Module has a 16-Gbps, full-duplex connection to the switch fabric of the Cisco Catalyst 6500 Series Switches, which is configured as an

IEEE 802.1q trunk and can carry any of the 4096 VLANs supported by the switch.

Communication with upstream routers and downstream switches and server farms happens through VLANs, which can be remotely reconfigured and mapped to different networks in minutes, without any change in cabling.

- **Space:** Requirements for rack units (which can be very costly in data center environments) are significantly lower compared to requirements for dedicated appliances. The Cisco ACE Module delivers up to 16-Gbps of Layer 4 to 7 performance while occupying just one slot in the switch. For a Cisco Catalyst 6513 Switch, which has 13 available slots and is 19 rack units (19RUs) high, that is equivalent to approximately 1.46RUs, significantly less than for any appliance with the same range of performance. In data center environments with available Cisco Catalyst 6500 Series slots, or where low-density interface line cards have been replaced with high-density ones, freeing up slots, essentially no additional RU space is needed to add a Cisco ACE Module.
- **Port density:** Modular Layer 2 to 3 switches offer a much higher density of physical interfaces than Layer 4 to 7 appliances, providing more flexible design options, even when the distribution and access layers are collapsed and hundreds of servers need to be directly connected. The Cisco Catalyst 6500 Series Switches support 48 to 576 10/100/1000 Ethernet and 48 to 1152 10/100 Ethernet port configurations.
- **Interface media types:** Greater flexibility is provided in interface media type and speed. The Cisco Catalyst 6500 Series Switches offer multiple media and density choices for 10/100 and 10/100/100 Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, and WAN interfaces.
- **Power and cooling:** The same (often redundant) power supplies required for the routing and switching infrastructure are used for application delivery, and the power requirements are lower: 220 watts (W) for the Cisco ACE Module compared to a 363W for a comparable (and lower-performance) competitive product, about a 40 percent savings in power consumption. Adjusting for performance (traffic throughput) parity multiplies that advantage by a factor of 3. In addition, in many environments, virtualization can be used; virtual devices, rather than physical appliances, are allocated to distinct departments or applications, thus reducing power consumption even further.
- **High availability:** Fewer failure scenarios need to be considered in the design because it has fewer physical links, thus reducing the chance of failure. Critical hardware components such as power supplies, fans, and clocks, which are already designed to be redundant to support the switching and routing infrastructure, are employed automatically.
- **Layer 2 and 3 features:** A broad range of Layer 2 and 3 features are natively supported and do not need to be implemented in an appliance, nor do they require additional interoperability tests. Spanning tree for loop detection and fast reconvergence, Cisco EtherChannel® for aggregation of multiple physical links, quality of service (QoS), and Remote Switch Port Analyzer (RSPAN) to generate and forward a copy of the traffic for analysis purposes are some of the many leading Layer 2 and 3 features that the Cisco Catalyst 6500 Series platform provides and that Cisco ACE Module deployments can immediately use.

Cisco ACE Solution: Enhanced Network Integration and Virtualization

Additional benefits of the network integration approach are specific to the implementation and to the level of integration achieved with the hosting platform. The Cisco ACE Module and Cisco

Catalyst 6500 Series solution has three additional features that significantly enhance the network integration model:

- The Cisco ACE Module integrates management functions into the Cisco Catalyst 6500 Series through a dedicated, secure connection, helping ensure that data traffic, regardless of load, does not affect the management of the device. This approach allows the exchange of information about configuration changes and other events, which can trigger specific actions in the Cisco Catalyst 6500 Supervisor Engine or the Cisco ACE Module.
- The Cisco ACE Module supports full virtualization and role-based management. The Cisco Catalyst 6500 Series supports virtualization with VLANs and VRF instances. The Cisco ACE Module represents the logical extension of the Cisco Catalyst 6500 Series virtualization capabilities. The Cisco ACE Module supports virtualization of all its functions, allowing the administrator to create multiple virtual devices operating independently on a single physical device. Each virtual device has a separate configuration file, user and network management rules, performance limits, application delivery and security rules, and routing tables. In a virtualized data center, Cisco ACE virtual devices can be mapped to VLANs and VRF instances, using their secure and segmented approach to connectivity and increasing the flexibility of the deployment. The Cisco ACE Module also supports role-based management, allowing the administrator to granularly define the network management actions that each user can perform on each of the Cisco ACE virtual devices. This behavior, together with the separation of network management functions on the Cisco Catalyst Supervisor Engine and the Cisco ACE Module, provides additional flexibility for defining which groups within the organization control which functions and for configuring the Cisco ACE network integration.
- The switch fabric architecture scales linearly when multiple modules are used in parallel. The Cisco Catalyst 6500 Series switch fabric and distributed forwarding architecture support the addition of new modules without inherent packet-forwarding bottlenecks.

With these features, the integrated Cisco ACE Module and Cisco Catalyst 6500 Series solution offers some additional advantages.

Scalability

Four Cisco ACE Modules can be deployed within the same Cisco Catalyst 6500 Series chassis and operate in parallel, scaling linearly to more than 60 Gbps of measured throughput for load-balanced traffic. The Cisco Catalyst 6500 Series routing engine can be used to distribute the traffic at wire speed between multiple Cisco ACE Modules, either through equal-cost multipath routing or policy-based routing (PBR) rules.

Manageability

The Cisco ACE module has its own dedicated console and configuration file (or files, when multiple virtual devices are configured), thus providing a simple and clean separation of the Layer 2 to 3 configuration on the Cisco Catalyst 6500 Supervisor Engine and the Layer 4 to 7 configuration on the module and its virtual devices. At the same time, however, access to the Cisco ACE Module command-line interface (CLI) can be allowed through the supervisor console, providing additional flexibility in managing the module.

While each Cisco ACE virtual device can send its syslog file to a dedicated server, the user can specify which syslogs are relevant enough to be tunneled to the Cisco Catalyst 6500 Supervisor

Engine, so that critical events can also be captured by the management systems configured for the routing and switching infrastructure. Additional mechanisms for communication between the Cisco ACE Module and the Cisco Catalyst 6500 Supervisor Engine provide information about misconfigured VLAN mappings between the two, reducing debugging time.

The role-based management support on the Cisco ACE Module allows organizations to customize the management of the Cisco ACE Module functions: for example, the Cisco Catalyst 6500 Series administrator could be authorized to manage the Cisco ACE on-board VLAN and IP configuration, but to only monitor the server and virtual server statistics, without being allowed to modify any of them.

Easy Upgrades

The Cisco ACE Module retains the same Cisco IOS® Software CLI and look and feel as the Cisco Catalyst 6500 Supervisor Engine, but it runs its own independent operating system. This approach allows the user to transparently upgrade the Cisco ACE Module and benefit from the latest Cisco ACE application delivery features without any requirement to upgrade the Cisco IOS Software for the supervisor.

Improved Configuration Management

The Cisco ACE Module on-board Cisco IOS CLI has been improved to better handle application delivery Layer 4 to 7 configurations, which tend to be more dynamic and to include more user-defined object names (servers, policies, virtual servers, server health probes, etc.) compared to the Layer 2 to 3 configurations typically found on the Cisco Catalyst 6500 Supervisor Engine.

The Cisco ACE Module also supports configuration rollback, which allows users to save up to 10 snapshots of different configurations for each Cisco ACE virtual device and to revert to any of those previously saved configurations within seconds and without the need to reboot the module. This feature simplifies configuration management tasks and reduces potential downtime by allowing the user to keep a backup configuration ready to use at any time should newly deployed changes negatively affect some of the application environments.

Advanced Cisco IOS Routing

The Cisco ACE Module can take advantage of the full range of routing features supported by Cisco IOS Software on the Cisco Catalyst 6500 Supervisor Engine. The Route Health Injection (RHI) feature allows the Cisco ACE Module to inject or remove routes for its virtual servers, based on the health of servers and applications, directly managing the routes in the Cisco Catalyst 6500 Supervisor Engine routing tables. The routes to the virtual servers can then be propagated to the rest of the network using the Cisco IOS routing protocol of choice, supported by the supervisor.

The RHI feature can be used to distribute load for the same virtual server across multiple Cisco ACE Modules, or it can be used as a disaster-recovery mechanism where the route to a specific virtual server is propagated with a different cost for different Cisco ACE Modules either within the same data center or across data centers.

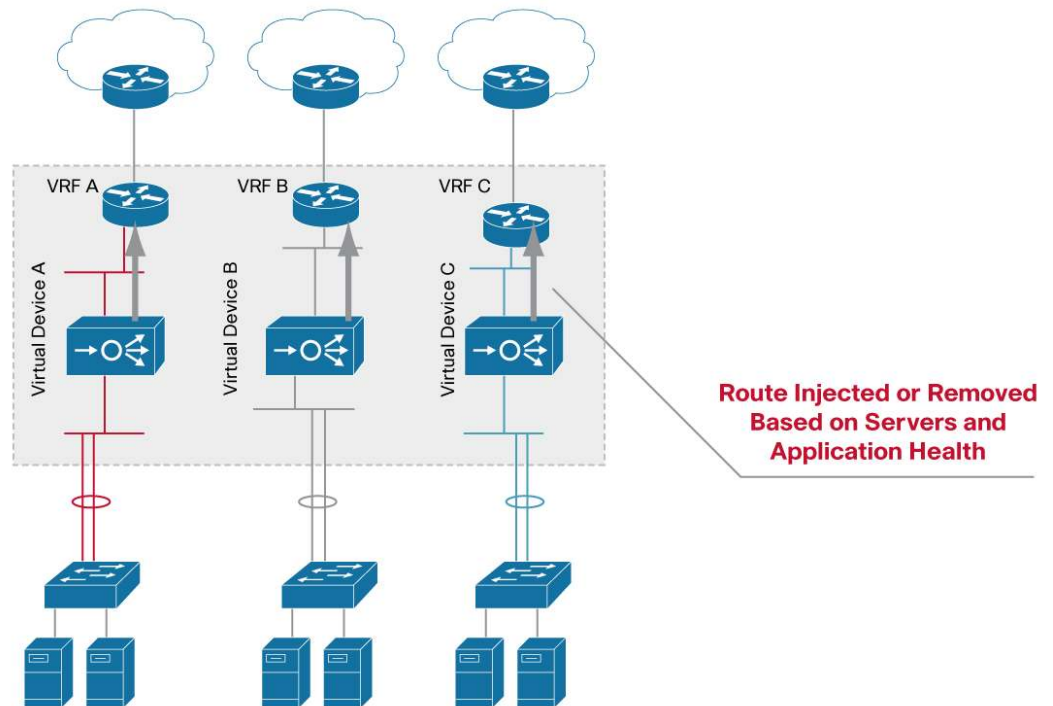
Virtualization and Network Segmentation

The Cisco ACE Module supports virtualization through the extension of the logic to the application delivery space of the Layer 2 and 3 VLANs and VRF instances that the Cisco Catalyst 6500 Series natively supports. It is simple to map Cisco ACE virtual devices to VLANs and VRF instances, thus

associating a separate network instance on the Cisco Catalyst 6500 Supervisor Engine with a completely independent application delivery instance.

The RHI feature is aware of VRF instances, so Cisco ACE virtual devices can inject and remove routes directly from VRF routing tables in the Cisco Catalyst 6500 Supervisor Engine. Figure 4 illustrates this concept.

Figure 4. Interaction Between Cisco ACE Virtual Devices and VRFs



Each virtual device can be dedicated to a set of applications, to an organization within the enterprise, or to a customer in a hosted environment. Overlapping IP addresses are supported, and each virtual device benefits from independent network management and policies and also from a dedicated virtual routing instance with full Cisco IOS routing protocol support.

Flexible Security

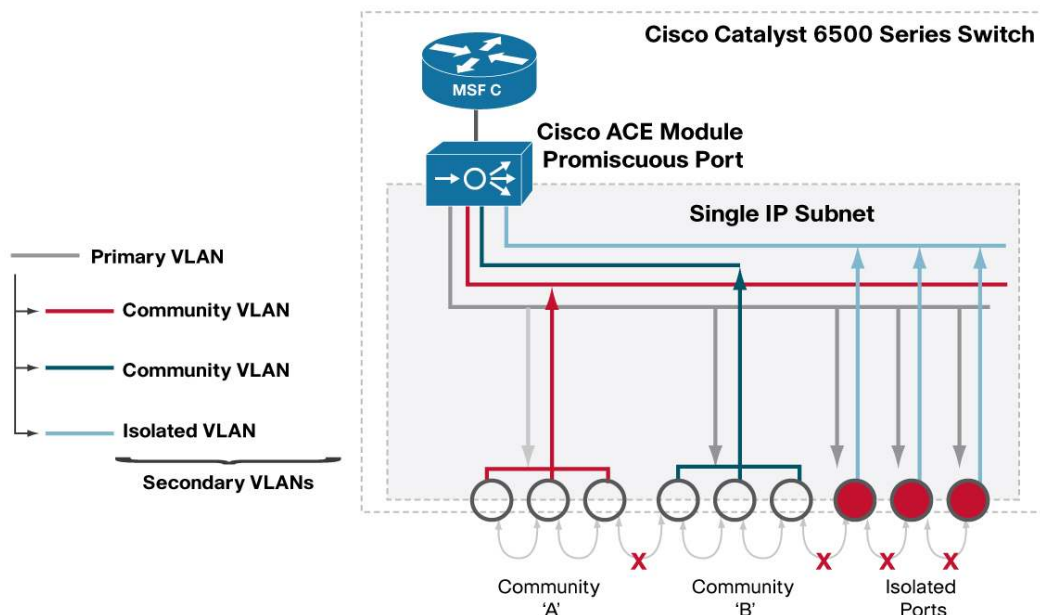
The Cisco ACE Module itself supports a wide set of security features (TCP and IP compliance and option checks, distributed-denial-of-service [DDoS] SYN attack protection, HTTP, Domain Name System [DNS], FTP, Real Time Streaming Protocol [RTSP], voice and video protocol compliance, access control lists [ACLs], etc.), but the integration with the Cisco Catalyst 6500 Series platform adds a layer of security features that can be used to protect servers and applications and the resources of the Cisco ACE Module against specific attacks. All traffic that reaches or leaves the Cisco ACE Module is automatically processed by the Cisco Catalyst 6500 Series hardware, so features such as user-based or per-flow rate limiting can be directly used to limit the amount of traffic that a single user can send, for example, to application servers.

One security feature in particular, private VLANs, has been enhanced to automatically integrate with the Cisco ACE Module without requiring any specific configuration on the module itself.

Private VLANs allow administrators to restrict traffic between hosts (servers) within the same Layer 3 domain without the need for any change in the IP addressing scheme.

Figure 5 shows how the Cisco ACE Module can act as a private VLAN promiscuous port, allowing it to communicate with any server in any secondary VLAN while using the Cisco Catalyst 6500 Series hardware to block all communication between devices on isolated ports or belonging to a different community.

Figure 5. Private VLANs with the Cisco ACE Module and the Cisco Catalyst 6500 Series Switch



High Availability

Cisco ACE Modules can be deployed in redundant mode either within the same Cisco Catalyst 6500 Series chassis or, more commonly, in separate chassis. In both cases, the Cisco ACE Module supports the capability to preserve existing connections even after a module failure or shutdown, by synchronizing flows and client-to-server associations (stateful failover).

The Cisco ACE failover mechanism allows each virtual device to fail over independently, according to user-configurable rules that provide the flexibility to respond rapidly in a variety of failure scenarios. The basic heartbeat between the two modules can be sent at 100 milliseconds (ms) intervals, achieving a complete failover in a fraction of a second.

Using an internal secure communication channel, the Cisco Catalyst 6500 Supervisor Engine can immediately inform the Cisco ACE Module of link and interface failures or Hot Standby Router Protocol (HSRP) group state changes, thus enabling each Cisco ACE virtual device to take immediate action on events that the Cisco Catalyst 6500 Supervisor Engine is constantly monitoring. The Cisco ACE Module can then complement that information with health monitoring probes (including scripts) to check the status of upstream gateways or critical devices, thus providing a comprehensive set of simple and fast failover mechanisms that can adapt to the most complex scenarios.

Comprehensive, Powerful Solution

The combination of independent network management and software and tight communication with the most advanced features makes the integration of the Cisco ACE Module with the Cisco Catalyst 6500 Series Switches a unique and powerful solution. The virtualization and role-based management support of the Cisco ACE Module together with the depth and breadth of the Cisco Catalyst 6500 Series routing, switching, and security features set this solution apart from traditional appliance-based designs in meeting the application delivery requirements of the most demanding data centers.

Integration of Application Networking Services Throughout the Network

Not only high-performance data centers can benefit from the integration of application networking services within the switching and routing infrastructure. For this reason, Cisco has pioneered this approach in two other critical network locations to improve application performance and delivery:

- At the branch office, with Cisco Wide Area Application Services (WAAS) Software: Cisco WAAS Software 4.0 is a powerful application acceleration and WAN optimization solution for the branch office that improves the performance of any TCP-based application operating in a WAN environment. With Cisco WAAS, enterprises can consolidate costly branch-office servers and storage into centrally managed data centers, while still offering LAN-like service levels for remote users. The Cisco WAAS solution offers reduced total cost of ownership (TCO), high application performance, efficient WAN utilization, and transparent integration with the network with secure, centralized manageability and control in an easy-to-implement package.
- At the campus, with the Cisco Catalyst 6500 Supervisor Engine 32 Programmable Intelligent Services Accelerator (PISA): The Cisco Catalyst 6500 Supervisor Engine 32 PISA is an intelligent services supervisor for the Cisco Catalyst 6500 Series of modular switches. It delivers industry-leading deep packet inspection, application awareness, security, availability, and manageability services for the networks of small and medium-sized business, enterprises, and service providers. The Cisco Catalyst 6500 Supervisor Engine 32 PISA provides hardware acceleration of intelligent services such as Cisco Network-Based Application Recognition (NBAR) and Cisco IOS Flexible Packet Matching (FPM) at multigigabit speeds, in addition to the management and control plane functions traditionally provided by the Cisco Catalyst 6500 Series Multilayer Switch Feature Card (MSFC).

Conclusion

Companies in the process of consolidating, migrating, or upgrading their data center infrastructure and evaluating application delivery devices should consider the infrastructure benefits and operational savings of solutions that integrate such devices natively into the switching and routing infrastructure. Although standalone load balancers and application delivery appliances continue to serve well in many environments, high-demand data centers with dynamic application infrastructures can benefit greatly from an integrated approach, especially when this approach is combined with virtualization.

The Cisco ACE Module integrates natively with the Cisco Catalyst 6500 Series Switches and is currently the only application delivery product in the industry to combine integration with true virtualization, making it the best application delivery device to keep up with the requirements of complex and dynamic applications in high-performance data centers.

For More Information

For more information about the Cisco ACE Module and Cisco Catalyst 6500 Series Switches, visit:

- <http://www.cisco.com/go/ace/>
- <http://www.cisco.com/go/catalyst6500/>

To learn how Cisco can help companies build better data center solutions, review the technical papers and design guides at: <http://www.cisco.com/go/datacenter>



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

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.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0710R)