CISCO SYSTEMS

Technology Brief

Cisco Virtual Switch Architecture

The Cisco Virtual Switch Architecture is the foundation for the Cisco Advanced ATM Multiservice Portfolio. This architecture utilizes the Multiservice Switching Forum (MSF) architectural framework for a multiplane system to provide industry leading investment-protection with advanced service creation capabilities for service providers. The Cisco Virtual Switch Architecture (CVSA) provides a high degree of flexibility, simultaneously supporting Asynchronous Transfer Mode (ATM), voice, and Internet Protocol (IP) services, with separate control planes for maximum service efficiency.

SUPPORTS msf1 ARCHITECTURE

The Cisco Virtual Switch Architecture supports the framework established by the MSF, a group with broad industry support that is developing standards and architectures to support multiple services on a common networking infrastructure. The MSF framework calls for an open, standards-based architecture that allows flexibility in terms of the service adaptation, switching, and management and application control planes used to build a multiservice network. Cisco supports this architectural approach because it enables carriers to offer the greatest possible mix of services on a common forwarding architecture, providing the best possible support for voice, data, and emerging services on a single switch.

The multiplane architectural framework consists of five logical Planes: Adaptation, Switching, Control, Applications and Management. The Adaptation Plane supports the physical interface to a user or another network element. The Switching plane supports the actual switching fabric by which physical interfaces are connected. The Control Plane provides the capability to manage network service events and provides control over both the Adaptation and Switching Planes. The Application and Management Planes provide services that use the capabilities of the Control plane.

The MSF framework is designed to allow multiple controllers to share the resources of the Switching Plane and the Adaptation Plane. Using service and policy management to switch resources, the controllers can be separate or combined to support the network applications. In effect, virtual switches are created by each controller. For more information on the MSF architectural framework, go to: www.msforum.org.

Building on the MSF Framework—The Cisco Virtual Switch Architecture

The MSF framework forms the foundation of Cisco's next-generation architecture for multiservice switches. The Cisco Virtual Switch Architecture is fully compliant with the high-level guidelines specified in the *MSF Release 1 Implementation Guidelines* at the level defined as Support for Standardized Protocols. The Cisco Virtual Switch Architecture enhances the framework created by the MSF in two crucial ways—by supporting multiple simultaneous controllers and providing the ability to dynamically allocate queues for the multiple controllers. It is important to note that these enhancements extend the usability of the architecture in a manner consistent with the framework, and these extensions are under review for adoption (refer to Figure 1).

Figure 1
The Cisco Virtual Switch Architecture

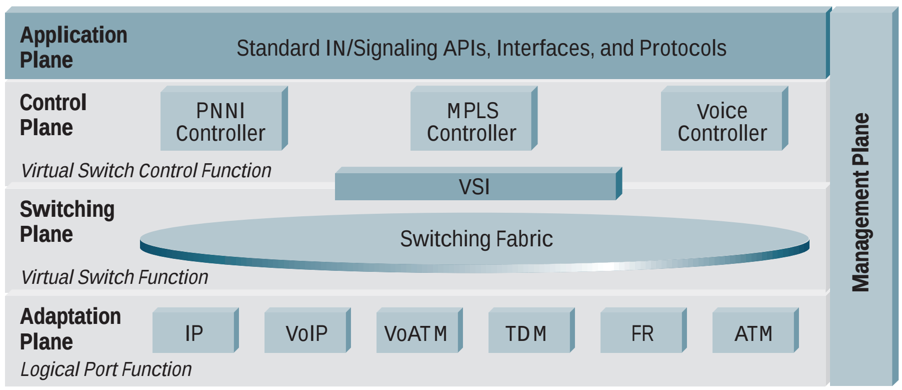

The Cisco Virtual Switch Architecture utilizes the five planes defined by the MSF architectural framework and uses a Virtual Switch Interface (VSI) as the protocol between the Control Plane and the Forwarding, or Switching, plane. The VSI enables the Cisco Virtual Switch Architecture to support *multiple, simultaneous* controllers for independent virtual switch functions, so that a single port can be controlled by multiple protocols such as PNNI and Multiprotocol Label Switching (MPLS). The VSI also provides a dynamic partitioning function that allows the in-service addition of new partitions, as well as the re-allocation of resources between active partitions.

Virtual Switch Interface: Support for Multiple, Simultaneous Controllers

The key to implementing the multiplane architecture within the Cisco Virtual Switch Architecture is the ability to communicate between the control and switching planes with a standard interface. Cisco created a standardized interface called the Virtual Switch Interface (VSI) for this purpose. The VSI is an open, published interface that can be used to create and implement new Control Planes. VSI defines the messages and associated functions that flow between the controller and the switch software in the Cisco Virtual Switch Architecture.

By sharing a common messaging interface, the controllers do not have to reside on the same card or chassis. The VSI is designed such that control-plane software can run on the processor cards in a system, on a separate card dedicated to the control software, or on an external device. Because the messaging interface is common, a single controller implementation may be used portably between various switching platforms. Each CVSA controller is a VSI master that talks to the VSI slave(s) residing on the platform software (Figure 2).

Figure 2
VSI Master and VSI Slaves Example

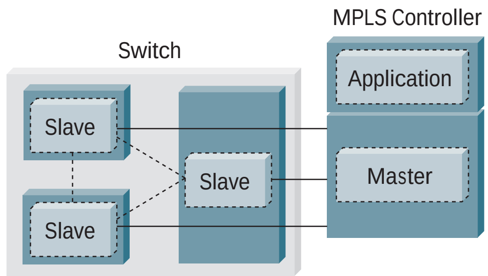

The VSI manages resource allocation so that controllers are completely independent and each service receives the quality of service (QoS) required. As a result, each control plane can be upgraded independently without any effect on the services using different control planes.

Currently, the Cisco Virtual Switch Architecture has implemented support for simultaneous usage of PNNI, MPLS, AutoRoute, and Media Gateway Control Protocol (MGCP) control planes. Additional control planes may also easily be added to the Cisco Virtual Switch Architecture. The General Switch Management Protocol (GSMP) has been selected by the MSF as the interface between the Control and Switching Planes. Unlike VSI, GSMP permits only one controller at a time to operate, with the capability to manage more than one service. With a single controller for multiple control planes, software upgrades and other controller configuration activities have the danger of impacting services in more than one control plane.

Flexible Dynamic Partitioning

Dynamic partitioning is another major attribute of the Cisco Virtual Switch Architecture for Adaptation Plane functionality. It allows individual ports on a switch to be partitioned for use by different control planes without affecting services.

Either ports and trunks can be added as new interfaces, or existing interfaces can be configured to allow traffic from additional control planes. When adding a new control plane to existing interfaces, the necessary bandwidth and a virtual path/channel identifier (VPI/VCI) range are made available to them. The biggest benefit of dynamic partitioning is that this in-service allocation of resources does not affect execution of the service in any way.

After partitioning, the respective control planes have a defined bandwidth available for their respective Connection Admission Control (CAC) algorithms for adding new connections, and each can access a pool of available VPIs/VCIs to assign to new connections. The total effective bandwidth assigned to all protocols in use on a port (AutoRoute, PNNI, MPLS, and so on) cannot exceed the physical port bandwidth. Effective bandwidth accounts for any overbooking allowed by each individual control plane. The allocated (or partitioned) bandwidth for each protocol impacts only the CAC bandwidth. Existing connections can use available bandwidth on each port, as dictated by their respective classes of service (CoSs), without regard to partitioning.

Dynamic partitioning enables the addition of a new Control Plane partition to existing interfaces (both ports and trunks) without affecting any of the existing Control Plane connections on that interface. The VPI/VCI range assigned to one Control Plane can contain existing connections from a separate Control Plane. However, no new connections may be made from the first Control Plane within the range taken over by the second Control Plane. Similarly, new connections from the second Control Plane can only be made in its own newly partitioned VPI/VCI range. This capability enables a simple conversion of end-user connections from one control plane to another.

Figure 3

Cisco Virtual Switch Architecture supports, ATM, voice and IP services simultaneously over a single switching plane.

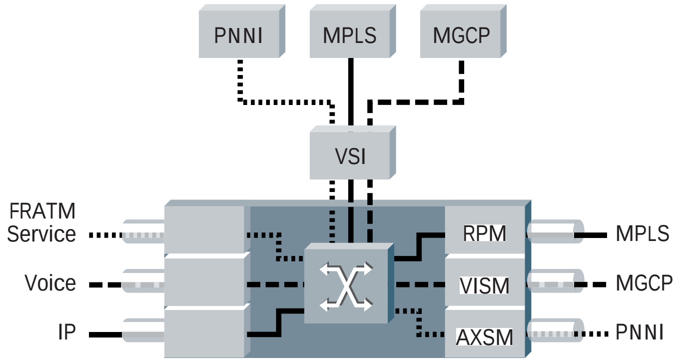

Enhanced QoS for Multiple Services

The Cisco Virtual Switch Architecture, with its VSI interface and dynamic partitioning capability, offers some exceptional benefits with respect to delivering QoS for different kinds of services. This is particularly important in the case when a carrier wants to offer both ATM and IP services using the same switching infrastructure.

A key feature is the separation of IP flows and ATM-based connections into separate queues. Each controller has its own resources within the switch. These resources are allocated to each controller, in this case PNNI and MPLS. The appropriate control plane can then manage each service individually. Carriers do not have to force-fit IP onto ATM, or ATM onto IP.

This feature is especially attractive for service providers who want to offer an end-to-end service for IP traffic across their networks, but need to have some of this traffic traverse a traditional ATM network. They will want to use IP Diff-Serv and other mechanisms for delivering QoS, but will have to rely on the ATM queues. This creates a challenge, because the provider must map the IP flows into ATM virtual circuits, by either pre-provisioning thousands of S-PVCs, or some form of IP-ATM SVC mapping-neither of which is very scalable. These IP flows are typically of short duration, using a different, less granular, QoS mechanism to deliver IP, and consequently do not map very easily into ATM classes of service.

The service provider can solve the flow set-up problem by adding an MPLS control plane to its ATM switch. However, if it has to map these multiple short IP/MPLS flows into ATM queues at random, it could have an adverse effect on that traffic. This is especially true if those queues are also being used to deliver very specific, per-VC-controlled ATM services, with tight service-level agreements (SLAs).

Contrast this approach to using a switch based on the Cisco Virtual Switch Architecture. Now it is possible to have *separate* IP queues under the control of MPLS, enabling a service provider to avoid all these issues. The IP flows will map directly across the cell-based forwarding plane, using the MPLS control plane, without interference to other ATM services (Figure 4).

Figure 4
With the Cisco Virtual Switch Architecture, full integration of IP and ATM service support is now possible.

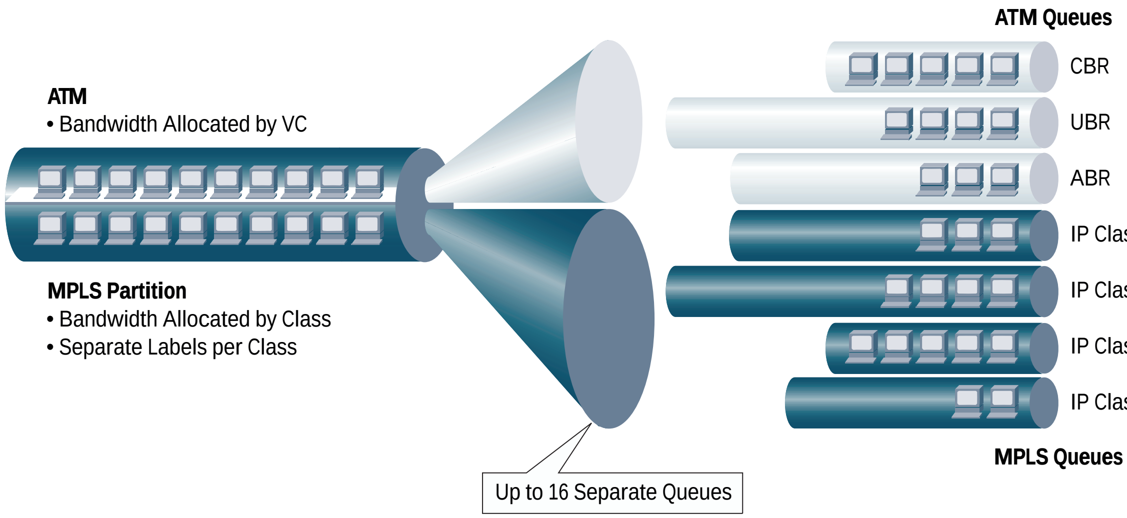

Using the Cisco Virtual Switch Architecture results in multiple benefits, including:

- Ease of new IP service creation-No mapping of IP flows to ATM SVC/PVCs is required.
- Native QoS support for each service type
- Ability to provide a native interface to a core network regardless of the core protocol

In addition, the cell-forwarding QoS mechanisms from ATM can be used additively to IP QoS in Cisco switches to produce the highest service quality control. As a result, the service provider gets the benefit of *both* cell-switching traffic management mechanisms and all the IP QoS features.

Now service providers can take advantage of the ability of ATM to control delay and delay variability across a statistical packet-based network for IP traffic. ATM CoS allows a service to offer multiple, standardized service offerings to customers by tightly controlling the way data is accepted into the network and shaped as it traverses the network. Cell forwarding can also be used to facilitate similar functionality for the MPLS control plane for services such as voice over IP (VoIP). The benefit to the service provider deploying the Cisco Virtual Switch Architecture is the ability to offer ATM CoSs, which are very granular, and IP-differentiated services, more suited to IP over the same forwarding plane. Because Cisco switches can separate IP and ATM traffic flows and support both PNNI and MPLS, they can support both types of traffic classes without impacting the SLAs of either. Additional, separate queues also provide the potential benefit of more IP queues, potentially needed for services such as Any Transport over MPLS (AToM), as the network evolves. AToM is a proposed standard for providing a single transport mechanism for carrier networks with multiple access services.

The benefits of the Cisco Virtual Switch Architecture for service providers include the following:

- Support for ATM, voice, and IP services simultaneously with separate control planes for maximum service efficiency
- Separate queues for delivering simultaneous use of ATM and IP QoS natively
- Can evolve a service mix using both PNNI and MPLS, without requiring new hardware
- ISAP—Service provides can IP enable any port, providing differentiated frame Frame Relay and ATM services
- Cost-effective manner of adding MPLS to ATM switches for new services, as well as controlling costs by consolidating network infrastructure
- Reliability—Allows data forwarding to occur independently from control functions

Summary

The Cisco Virtual Switch Architecture supports the MSF architecture framework and enhances the inherent benefits of a multiplane architecture by adding a Virtual Switch Interface and Dynamic Partitioning capabilities.

This enables multiple, simultaneous Control Planes to direct the Forwarding Plane, along with highly granular control of multiple Control Planes on a single port. Service providers using the Cisco Virtual Switch Architecture can support ATM, voice, and IP services simultaneously, with separate control planes for maximum service efficiency. They can provision separate queues for delivering ATM and IP QoS natively, allowing for simultaneous use of ATM CoSs and IP QoS.

Service providers can create a dynamic service mix, as desired, without requiring new hardware and enabling the gradual evolution to include new services or a common control plane. In short, the Cisco Virtual Switch Architecture offers the most cost-effective approach to adding new services as well as controlling costs by consolidating network infrastructure.

CISCO SYSTEMS

Corporate 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: 408 526-4100

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe