Cisco IOS Software Release 15 M and T Features and Hardware Support

PB561938

Last Updated: October 2009

This Product Bulletin introduces Cisco IOS Software Release 15 M and T, and includes the following sections:

1. Introduction

Cisco IOS® Software is the world's premiere network infrastructure software, delivering seamless integration of technology innovations, business-critical services, and hardware support. Currently operating on millions of active systems, from small home office routers to the core systems of the world's largest service provider networks, Cisco IOS Software is the most widely leveraged network operating system software in the world.

Developed for wide deployment in the world's most demanding enterprise, access, and service provider aggregation networks, Cisco IOS Software Release 15 M and T provides a comprehensive portfolio of Cisco technologies, including the leading-edge functionality and hardware support from Releases 12.4 and 12.4T, These integrated technologies are delivered on the broadest range of hardware in the industry, including Cisco's Integrated Services Routers, and Cisco 7200 and 7301 Series.

Release 15 M and T key innovations span multiple technology areas, including Security, Voice, High Availability, IP Routing and Multicast, Quality of Service (QoS), IP Mobility, Multiprotocol Label Switching (MPLS), VPNs, and Embedded Management.

Release 15 M and T provides customers new feature release delivery and hardware support in a shorter amount of time, broadened feature consistency, more reliable new feature release and rebuild schedules, proactive release support lifecycle policies, and easier software selection, deployment, and migration guidelines.

Key highlights of Release 15 M and T, illustrated in Figure 1 below, include the following:

- Feature inheritance from Cisco IOS Software Releases 12.4T and 12.4 Mainline[1]
- M (extended maintenance) releases every 20 months - allows customers to qualify/deploy/remain on releases longer with active bug fix support
- Standard maintenance 15 T releases - provides the latest features and hardware support before the next M release becomes available on Cisco.com[2]
- Rebuilds of Release 15 M and T releases for ongoing bug fixes
- **Cisco IOS Software Release 15.0(1)M is the first release**

Figure 1. Release 15 M & T Overview

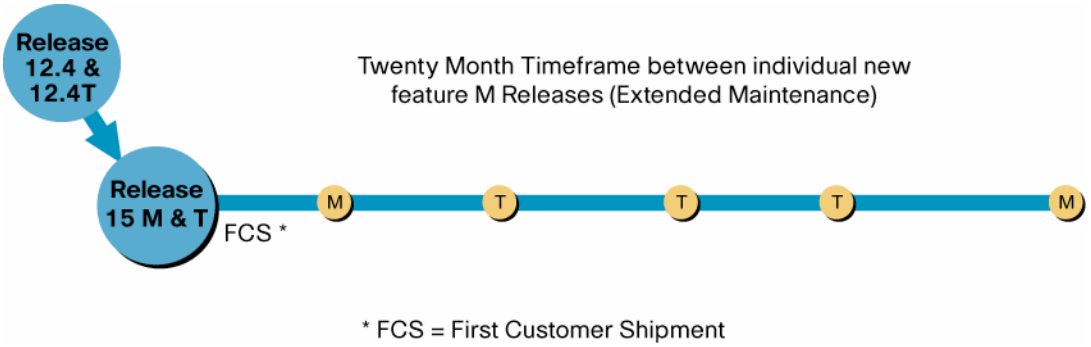

1.1 Migration Guide

Cisco recommends that customers running Release 12.4T, Release 12.4, or prior Cisco IOS M and T releases upgrade to Release 15.0(1)M where possible. Several Cisco hardware platforms ended support on Releases 12.4(15)T and 12.4 Mainline. These platforms will not be supported in 15 M and T. Refer to the following bulletin for more information:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps8802/ps6968/ps6441/product_bulletin_c25_466578.html

Customers should determine their functionality needs and choose the appropriate release.

Figure 2 illustrates the current migration path from Cisco IOS Release 12.4, Release 12.4T, and prior releases to Release 15.0(1)M.

Figure 2. Release 12.4T Migration Path

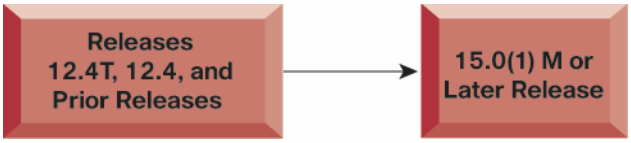

Hardware Requirements for Release 15.0(1)M

Please refer to the following web sites for Integrated Services Router hardware platform requirements for Release 15.0(1)M:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5460/prod_bulletins_list.html

Information on hardware requirements for other Release 15.0(1)M supported hardware platforms can be found via the Cisco.com products and services and support web sites:

http://www.cisco.com/en/US/products/index.html

[1] AppleTalk Phase I and II, and Service Selection Gateway (SSG) features are not supported in Release 15 M and T.
[2] The first Cisco IOS 15 T release is planned for the 2nd quarter of 2010.

http://www.cisco.com/go/support

AppleTalk Support Discontinuation

Due to a significant decrease in AppleTalk usage and demand among its customer base, and given the fact that Apple now fully supports the TCP/IP family of protocols, Cisco has reached the decision to discontinue AppleTalk support in Cisco IOS releases 15.0(1)M and later releases. Refer to the following product bulletin for more details:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps8802/ps5460/product_bulletin_c25-520459.html

Cisco Service Selection Gateway (SSG) Feature Discontinuation

The Cisco Service Selection Gateway (SSG) feature is not supported from Cisco IOS Software Release 15.0(1)M and onward. Refer to the following product bulletin for more information:

http://www.cisco.com/en/US/prod/collateral/routers/ps341/end_of_life_notice_c51-501483.html

1.2 Release 15.0(1)M Support Lifecycle Guidelines

Cisco IOS Software Release 15.0(1)M is an extended maintenance 15 M Release. The software support lifecycle for Release 15.0(1)M is illustrated in Figure 3 below.

Figure 3. Support lifecycle for Release 15.0(1)M

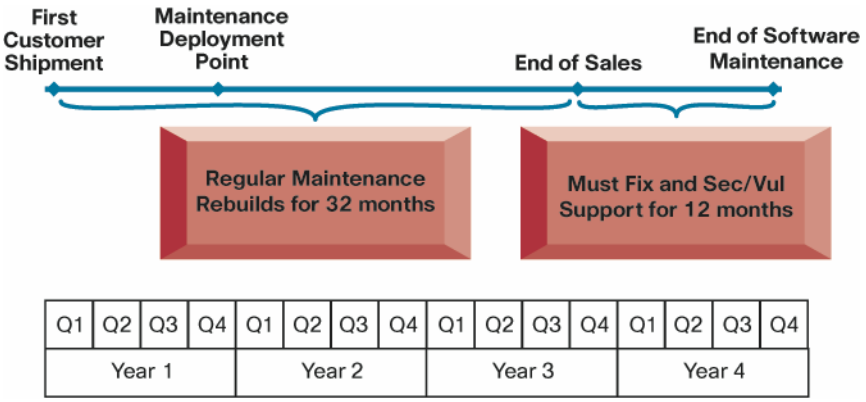

Support milestones for Release 15.0(1)M are defined in Table 1 below.

Table 1. Cisco IOS Software Release 15.0(1)M Support Milestones

Milestone	Definition	Date
First Customer Shipment (FCS)	The date the software release is first available to customers on Cisco Software Download Center http://www.cisco.com/public/sw-center/	Day 0
Maintenance Deployment Point	Maintenance Deployment designation is based on extensive Cisco internal and customer-focused testing efforts, assessment of the release stability in customer networks, and ongoing Cisco IOS Software improvements. A set of stringent criteria is used to analyze the quality of the release, such as high customer satisfaction, expansive customer deployment, demonstrated software reliability, and rigorous ongoing testing.	Typically 9-12 months (3rd or 4th maintenance rebuild) after the initial posting (FCS date) of the release on Cisco.com Software Download Center
End of Sale Date	The day the release is no longer orderable or included in manufacturing shipments of Cisco hardware.	32 months after FCS date
End of Software Maintenance	This is the last date that Cisco Engineering may release any final software maintenance releases or software fixes for the release.	44 months after FCS date

1.3 Cisco IOS Packaging Consideration

Cisco IOS Packaging for Cisco Routers Previously Supported in Release 12.4T

Figure 4. Cisco IOS Release 15.0(1)M Packaging Framework for Cisco Routers Previously Supported in 12.4T[3]

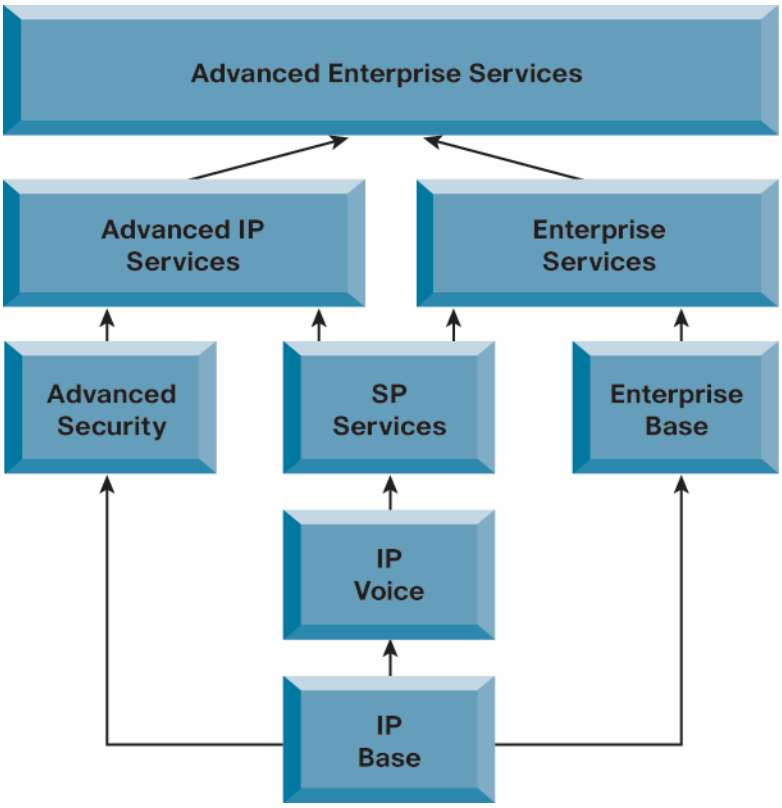

Cisco IOS Packaging for Cisco Integrated Services Routers Generation 2

Cisco Integrated Services Routers Generation 2 (ISR G2) 1900, 2900, and 3900 Series support services on demand through the use of software licensing which enables customers to realize operational savings through ease of software ordering and management. When you order a new ISR G2 platform, the router is shipped with a single universal IOS software image and corresponding feature set packages as shown in Figure 5 below.

Figure 5. Cisco IOS Release 15.0(1)M Packaging for ISR G2 1900, 2900, and 3900 Series

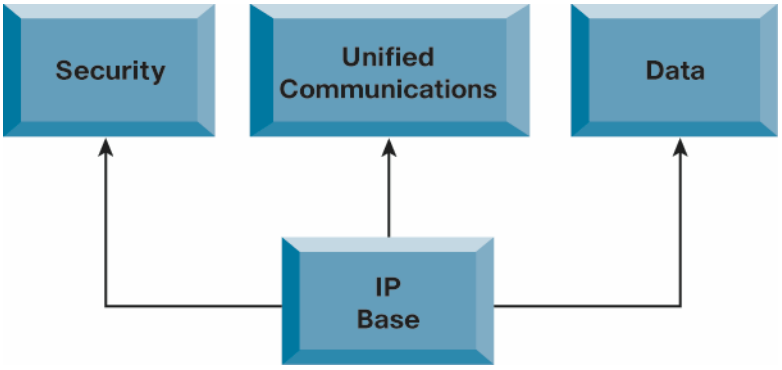

For more information on Cisco IOS Software packaging for ISR G2 visit

http://www.cisco.com/en/US/products/sw/iosswrel/ps5460/prod_bulletins_list.html

[3] Hardware-specific feature sets may affect respective Cisco IOS Software Package availability. Refer to the Cisco Product & Services support page on Cisco.com to obtain additional information specific to your hardware platform: http://www.cisco.com/en/US/products/index.html. For more information on Cisco IOS Software packaging visit: http://www.cisco.com/en/US/products/sw/iosswrel/ps5460/prod_bulletins_list.html

1.4. Cisco IOS Software Activation for Integrated Services Routers Generation 2

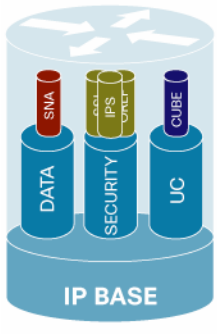

Software Activation enables an 'On Demand Service' activation model for Cisco ISR G2 and supported network modules with consistent licensing policies and manageability as available on other Cisco Software Activation platforms. Customers will realize operational savings through reduced truck rolls, ease of software ordering, simplified software management and greater flexibility in deploying new features.

Each ISR G2 will ship with a universal Cisco IOS Software image and will use Software Activation Licensing keys to activate Cisco IOS functionality. Features ordered with the router will be pre-activated at the factory or by your Cisco partner.

For more information on Cisco IOS Software Activation visit http://www.cisco.com/go/sa.

1.5 Cisco IOS IPv6 Repackaging for Integrated Services Routers Generation 2

For years, Cisco IOS Software has expanded support of IP Version 6 (IPv6) to the majority of its technology areas and hardware platforms Due to market trends such as available IPv4 address pool exhaustion and regional registries issuing advisory to the Internet community to adopt IPv6 and national mandates, Cisco IOS packaging for IPv6 is now evolving.

Cisco now supports packaging parity for IPv6 with IPv4 for Integrated Services Routers Generation 2 starting with Cisco IOS Software Release 15.0(1)M. IPv6 feature support for a technology will now be packaged in the same feature set as IPv4 for these platforms.

This new Cisco IOS Software packaging for IPv6 is already available on the following Cisco IOS Software Releases:

- Catalyst® 3700 and 3560 Series images starting with Release 12.2(50)SE
- Catalyst 4500 Series images starting with Release 12.2(52)SG
- Catalyst 6500 Series images starting with Release 12.2(33)SXI

Cisco will expand the package parity effort to additional Cisco IOS Software releases and hardware platforms in the near future.

Product Management Contact

Steve Koretsky (skoretsk@cisco.com)

1.6 Additional Information

- Cisco IOS Software Page
  http://www.cisco.com/go/ios
- Additional details on features in Cisco IOS Software Release 15.0(1)M
  http://www.cisco.com/en/US/docs/ios/15_0/release/notes/150MFEAT.html

- Release notes for Cisco IOS Software
  http://www.cisco.com/cisco/web/psa/default.html
- Cisco IOS Software Download Center
  Download Cisco IOS Software releases and access software upgrade planners.
  http://www.cisco.com/public/sw-center/
- Cisco Feature Navigator
  A web-based application that allows you to quickly match Cisco IOS Software releases to features to hardware.
  http://www.cisco.com/go/fn/
- Cisco Software Advisor
  Determine the minimum supported software for selected hardware.
  http://tools.cisco.com/Support/Fusion/FusionHome.do

2. Release 15.0(1)M Highlights

Release 15.0(1)M adds support for Service Advertisement Framework (SAF) to enable dynamic discovery of network applications and services, Flexible NetFlow and NBAR integration for layer 2 through 7 per-application-flow visibility and statistics, Embedded Event Manager Version 3.1 enhancements to event detection, notification, and command execution capabilities, and support for the Cisco Integrated Services Routers Generation 2 (ISR G2) 1900, 2900, and 3900 Series.

Table 2 below lists the key features and hardware supported delivered in Release 15.0(1)M.

Table 2. Release 15.0(1)M New Feature Highlights

New Hardware	IP Routing	IP Multicast	Call Admissions Control
• Cisco Integrated Services Routers Generation Two (ISR G2) 1900, 2900 and 3900 Series	• Graceful OSPF Restart (RFC 3623) (Helper Mode Only) • Graceful Restart for OSPFv3 (RFC 5187) (Helper Mode Only) • OSPF Graceful Shutdown • OSPF Generic Time to Live (TTL) Security Check (GTSM) • Performance Routing (PfR) Protocol Independent Route Control (PIRO) • Performance Routing (PfR) EIGRP mGRE DMVPN Hub-and-Spoke Support • Service Advertisement Framework (SAF) • BGP Graceful Restart per Neighbor • Intermediate System-to-Intermediate System (IS-IS) BFD Support • IS-IS VRF Support • MPLS VPN - Inter-AS Option AB • BGP Route Target Changes Without PE-CE Impact • IS-IS MIB Support • MPLS VPN—BGP Local Convergence	• IGMP Static Group Range Support • IP Multicast Load Splitting - Equal Cost Multipath (ECMP) using S, G and Next-hop • IPv4 and IPv6 Multicast Address Group Range Support • Multicast MIB VRF Support • Multicast VPN Extranet Support • Multicast VPN VRF Select • PIM Triggered Joins	• RSVP Interface-based Receiver Proxy • RSVP Fast Link Repair • RSVP VRF Lite Aware Admission Control

High Availability	Embedded Management	IOS Security	Voice
• BFD client for IPv4 Static Routes • BFD VRF support • BFD Support for WAN Interfaces	• Flexible NetFlow and NBAR Integration • EEM Version 3.1 Enhancements	• Lightweight IPS Engines for Signatures • New Default IOS IPS Category signatures • Chaining of Traffic Scanning (Regular Expression) Tables for IPS • Configurable Threshold Limits for IPS Signatures • GET VPN VRF-Aware GDOI on GM • Ability to Disable Volume-based IPSec Lifetime Rekey • DMVPN Enhancements	• Packet Voice, Video DSP Module-3 • Transcoding and Codec Enhancements • Cisco Unified Border Element (CUBE) Support for SRTP-RTP Internetworking • Cisco Unified Border Element (CUBE) Support for Out-of-dialog SIP OPTIONS Ping Messages to Monitor SIP Servers • UC Trusted Firewall Control Version 2 • Service Advertisement Framework (SAF) Support for UC Manager Express, Cisco Unified Border Element and Voice Gateways

3. New Hardware

3.1 Cisco Integrated Services Routers Generation Two (ISR G2)

Cisco's 1900, 2900 and 3900 Series Integrated Services Routers build on 25 years of innovation and product leadership. Generation Two of the Integrated Services Routers are architected to enable the next wave of branch office evolution enabling richer media experiences with video, local compute resources for local application flexibility and survivability and Wide Area Network evolution. Adding to the rich breadth of services available on the existing integrated services routers, such as security, unified communications, wireless, and application optimization services, Generation Two of the Integrated Services Routers enhance the reduction in overall branch office expenses with the introduction of pay-as-you-grow software licensing and simplified Cisco IOS software packaging.

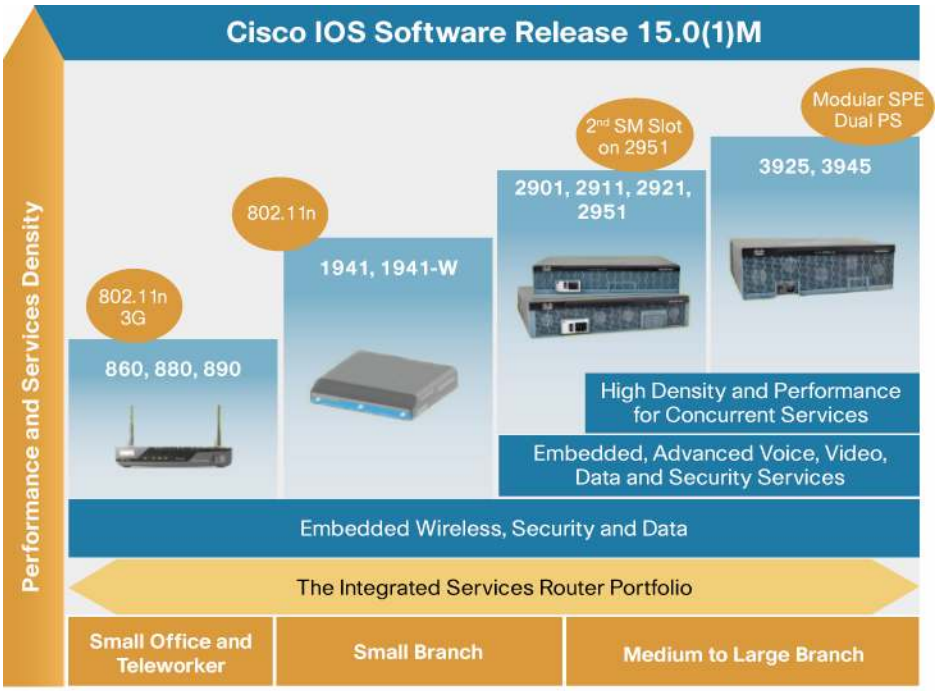

Generation Two of the Integrated Services Router portfolio consists of three product families similar to the current generation of ISR: the Cisco 1900, 2900 and 3900 Series Integrated Services Routers. From the Cisco 1941 through the Cisco 3945, the routers in portfolio provide increasing performance, module slot density and features to match the needs of diverse branch offices running varying degrees of rich services.

Additional Information

For more information on ISR G2 visit http://www.cisco.com/go/isrg2

Product Management Contact

Aseem Srivastava, asesriva@cisco.com

4. IP Routing

4.1 Graceful OSPF Restart (RFC 3623) (Helper Mode Only)

This feature enables nonstop forwarding (NSF) helper mode for OSPFv2 in Cisco IOS software using the IETF standardized Graceful Restart (GR) helper mode functionality described in RFC 3623. Graceful Restart allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a processor switchover. When GR is used, peer networking devices are informed, via protocol extensions prior to the restart event, of the GR-capable router's ability to perform graceful restart. This helper mode capability enables the peer routers to understand this messaging. When a restart occurs, the enabled peer will continue to forward to the switching-over router, although in most cases the OSPF peering relationships will need to be rebuilt.

Essentially, this feature enables the peer router to give the switching-over router a "grace" period to re-establish the OSPF neighbor relationship while continuing to forward to the routes from that peer.

Benefits

- Enables IETF standard OSPF RFC 3623 Graceful Restart helper mode in multi-vendor networks.
- Increases network availability by allowing the original OSPF router to stay on the forwarding path even as their OSPF software is restarted.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6629/products_ios_protocol_option_home.html

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.2 Graceful Restart for OSPFv3 (RFC 5187) (Helper Mode Only)

This feature enables nonstop forwarding (NSF) helper mode for OSPFv3 in Cisco IOS software, using the IETF standardized Graceful Restart (GR) helper mode functionality described in RFC 5187. GR allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a processor switchover. When GR is used, peer networking devices are informed, via protocol extensions prior to the restart event, of the GR-capable router's ability to perform graceful restart. This helper mode capability enables the peer routers to understand this messaging. When a restart occurs, the enabled peer will continue to forward to the switching-over router, although in most cases the OSPFv3 peering relationships will need to be rebuilt. Essentially, this feature enables the peer router to give the switching-over router a "grace" period to re-establish the OSPFv3 neighbor relationship, all the while continuing to forward to the routes from that peer.

Benefits

- Enables IETF standard OSPFv3 RFC 5187 Graceful Restart helper mode in multi-vendor networks.
- Expands availability of your network by allowing OSPFv3 routers to stay on the forwarding path even as their OSPFv3 software is restarted.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6629/products_ios_protocol_option_home.html

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.3 OSPF Graceful Shutdown

The OSPF Graceful Shutdown feature allows network administrators to remove a router from the network gracefully without impacting data traffic. When a user issues an OSPF shutdown command, the router informs all its neighbors that it is going offline by sending OSPF messages indicating that all links originating from the router are not useful for data forwarding. In addition, it also sends an empty hello message to bring down any OSPF adjacency relationships with its neighbors. Note that the router is reachable after the graceful shutdown for troubleshooting or upgrading router software or hardware.

Benefits

- Enables software and hardware upgrades in a single route processor device—Users can gracefully shutdown the router from the network, and then upgrade software or hardware in the router as needed.
- Enables trouble shooting and debugging of a router without impacting data traffic— After shutting down the router gracefully, users can login to the router for debugging or trouble shooting.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6629/products_ios_protocol_option_home.html

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.4 OSPF Generic Time to Live (TTL) Security Check (GTSM)

An OSPF router exchanges topology information with neighboring routers for building routing tables. The OSPF mechanism allows a user to configure MD5 or a simple password, for authenticating an adjacent router before information is exchanged. The OSPF Generic TTL Security Mechanism (GTSM) provides an additional security mechanism by ensuring that the neighboring OSPF router is exactly the number of hops away as specified in the configuration. When an OSPF router receives a message from another OSPF neighbor, it compares the TTL in IP header with the TTL that is configured for that neighbor. Only when the TTL matches will the OSPF router process the message from the neighbor.

The flexible configuration of this feature allows the user to configure the TTL per OSPF process or per interface. When the TTL is configured per OSPF process, this TTL is used to validate all neighbors on all interfaces on that router. If a TTL is configured on an interface, it will override the TTL configured at the process level.

Note: This feature needs to be configured on all neighboring routers in order to set the appropriate TTL value in the IP header.

Benefits

- Simplfied OSPF Security— This feature provides an additional security mechanism that is enabled by simply configuring the required TTL value between two OSPF routers. This ensures that a remote hacker cannot form an adjacency with any OSPF router in the network.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6629/products_ios_protocol_option_home.html

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.5 Performance Routing (PfR) Protocol Independent Route Control (PIRO)

Cisco Performance Routing (PfR) is a solution that addresses network performance problems by enabling the network to intelligently choose a path that meets application performance requirements. In addition, PfR enables the network to choose resources appropriately to reduce enterprise operational costs. This feature complements classic routing technologies by adding intelligence to best path selection, meeting performance requirements of applications.

PfR selects an egress or ingress WAN path based on parameters that impact application performance. These application parameters include reachability, delay, cost, jitter, and Mean Opinion Score (MOS). PfR can also select the best egress or ingress WAN path to intelligently load balance traffic based upon utilization and/or circuit cost to reduce costs incurred by enterprises. To achieve this, PfR selects a WAN path based interface parameters such as reachability, load, throughput, and link cost of using a path.

The current implementation of PfR provides explicit route control for static routes and BGP routes. The PfR Protocol Independent Route Optimization (PIRO) feature enables PfR to be used with any IP routing protocol. PIRO enables PfR to perform route control by performing a route lookup in the IP Routing Information Base (RIB) and by applying dynamic policy-based routing (PBR), rather than by injecting more specific routes into the routing table. This enables PfR to be deployed with any routing protocol (including EIGRP and OSPF) provided that parent routes are available in the RIB (via OSPF ECMP configuration for example). This feature is enabled by default and requires no additional configuration.

Figure 6. Cisco Performance Routing (PfR) use case

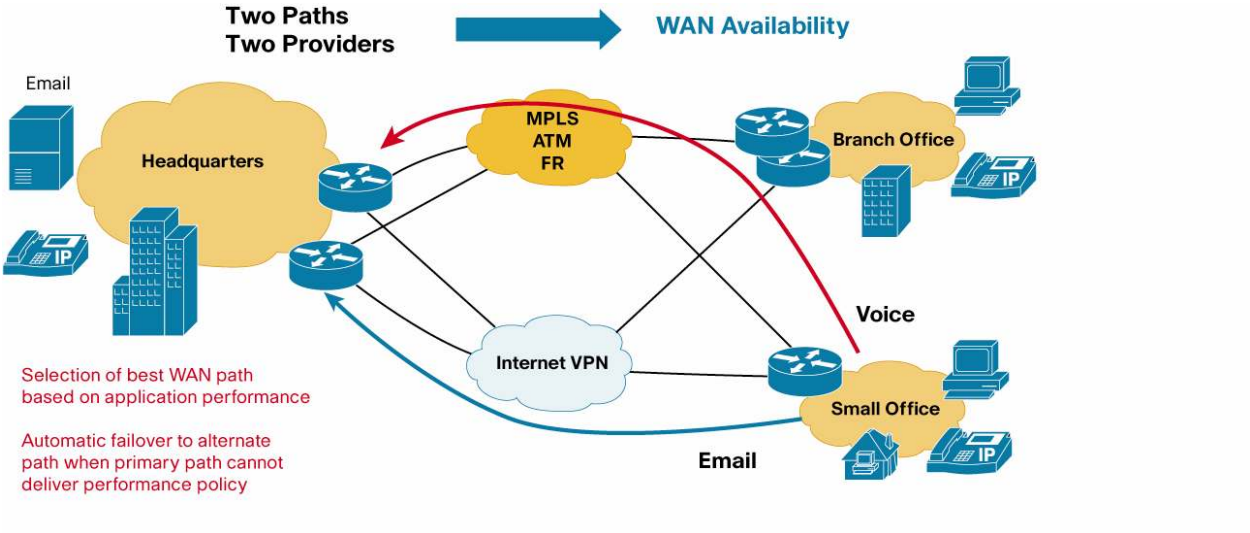

As shown in Figure 6 above, PfR is deployed at both branch offices and at the headquarters. In this example, PfR can be used to enforce application-specific WAN path policies. Suppose the administrator wishes to send voice traffic from the small office over the path currently providing the lowest delay, and send email traffic from the small office over the path providing the most throughput. PfR automatically measures application latency and throughput via active and passive measurement techniques, enabling it to automatically route email traffic over the Internet VPN and voice traffic over the MPLS cloud in this example.

Benefits

- Extends PfR deployability for customers who want to implement PfR in EIGRP or OSPF networks
- Enables optimal path selection for critical applications with performance requirements:
  - Advanced load balancing techniques based upon utilization, throughput, and cost
  - Transactional Traffic—Low Latency SQL Transactions, Automated B2B, ERP
  - Time/Delay sensitive—Voice, Vertical Apps (trading floor)
  - Loss Sensitive—Voice, Video, Circuit Emulation
  - Performs soft error detection and mitigation undetectable by traditional routing protocols, such as, network black-holes, brownouts, suboptimal performance
- PfR monitoring-only mode can be used for capacity planning and policy testing. In monitoring only mode, PfR can determine the network dynamics based upon policies or SLAs without actively controlling routes.
- Enhances network visibility into performance problems.
- Embedded in Cisco IOS Software, no additional network appliance is required for PfR.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/go/pfr

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.6 Performance Routing (PfR) EIGRP mGRE Dynamic Multipoint VPN (DMVPN) Hub-and-Spoke Support

This feature enables Cisco Performance Routing (PfR) to provide intelligent route control for EIGRP networks. Specifically, it enables the ability to inject routes into the EIGRP routing table in order to control prefixes and applications over EIGRP routes. This feature also adds PfR support for mGRE tunnel interfaces in Dynamic Multipoint Virtual Private Network (DMVPN) deployments.

Figure 7. PfR EIGRP mGRE DMVPN Hub-and-Spoke use case.

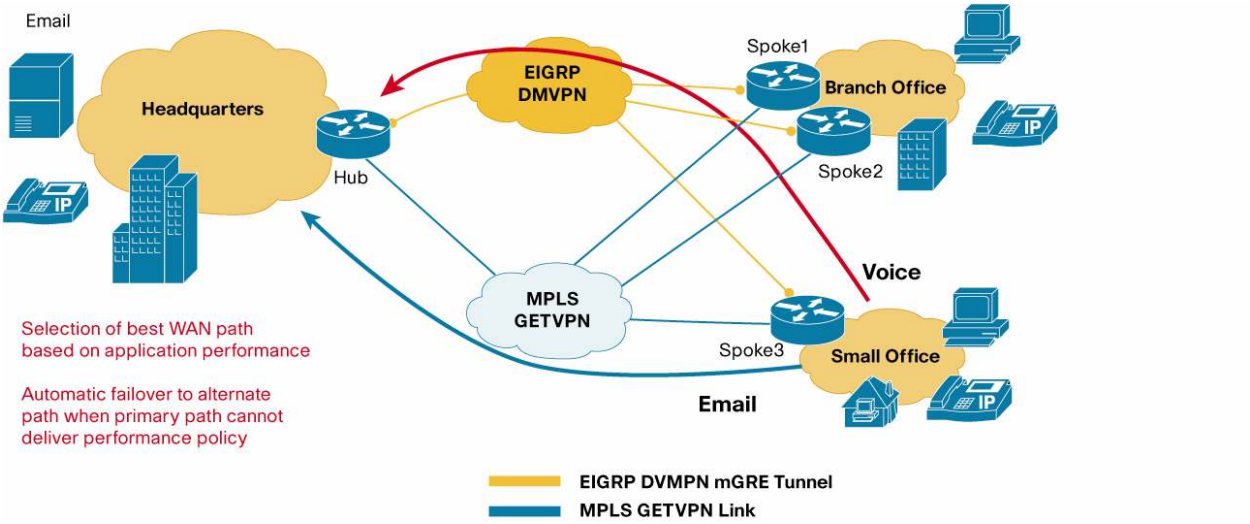

As shown in Figure 7 above PfR is deployed at both branch offices and at the headquarters. In this example, PfR can be used to enforce application-specific path policies over the Internet and EIGRP DMVPN networks. Suppose the administrator wishes to send voice traffic from the small office over the path currently providing the lowest delay, and send email traffic from the small office over the path providing the most throughput. PfR automatically measures application latency and throughput via active and passive measurement techniques, enabling it to automatically route email traffic over the MPLS GETVPN cloud and voice traffic over the Internet EIGRP DMVPN cloud in our example.

Benefits

- Enables PfR to provide performance-based optimal path selection for critical applications running on EIGRP networks
- Extends PfR applicability in the enterprise by enabling PfR deployment in EIGRP-based DMVPN networks

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/go/pfr

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.7 Service Advertisement Framework (SAF)

Service Advertisement Framework (SAF) enables network applications to discover the existence and configuration and presence of other networked application services dynamically.

The Cisco Service Advertisement Framework (SAF) provides a foundation for network-based applications and devices to dynamically distribute and discover service information across a Cisco network infrastructure enabling real-time service discovery, availability, auto-configuration and more. This is accomplished by coupling the dynamic properties of routing protocols and messaging services to create a high-performance resilient network-based messaging service.

SAF consists of several components:

- **SAF Client**—An SAF Client is an application that wishes to advertise a service to the network or request a service from the network or both. The client communicates with a SAF Forwarder in order to advertise or discover services. The SAF Client application may be embedded in a Cisco IOS device or, an external (non-IOS) appliance, host or other network device.
- **SAF Forwarder**—The SAF Forwarder node provides the relationship between SAF clients and the SAF-enabled network and is typically deployed in Cisco routers and switches. The SAF Forwarder receives service advertisements from SAF Clients and forwards the advertisements to its neighbor SAF Forwarders and Clients if needed. In the case that a SAF Client removes a service or the service becomes unreachable to the SAF Forwarder, the SAF Forwarder will inform the rest of the network that the service is no longer available. When the SAF Forwarder node receives advertisements from other SAF Forwarders, it keeps a copy of the advertisement and forwards it to other SAF peers. In this way, the SAF network can dynamically discover the availability and configuration of networked services. The SAF Forwarder is a Cisco IOS Software IP networking capability.

SAF leverages Cisco EIGRP technology to distribute service advertisements but operates independently of the IP routing infrastructure. As a separate Cisco IOS capability, SAF can be implemented on networks using EIGRP, OSPF, IS-IS, BGP, or other IP routing protocols.

Cisco Unified Communications (UC) Release 8.0 SAF Support

Cisco Unified Communications (UC) Release 8.0 is the first application to leverage SAF to improve large scale UC systems while significantly reducing time to deploy and operational costs.

With UC Release 8.0, Cisco Unified Communications Manager (CUCM) and Call Manager Express (CME) use SAF for dynamic Call Agent discovery and automated dial plan configuration.

UC 8.0 also includes SAF Client support for Cisco Integrated Services Router (ISR) based Cisco Unified Border Element (CUBE), TDM Gateways and Cisco Survivable Remote Site Telephony (SRST) services.

Figure 8. SAF for Cisco Unified Communications use case.

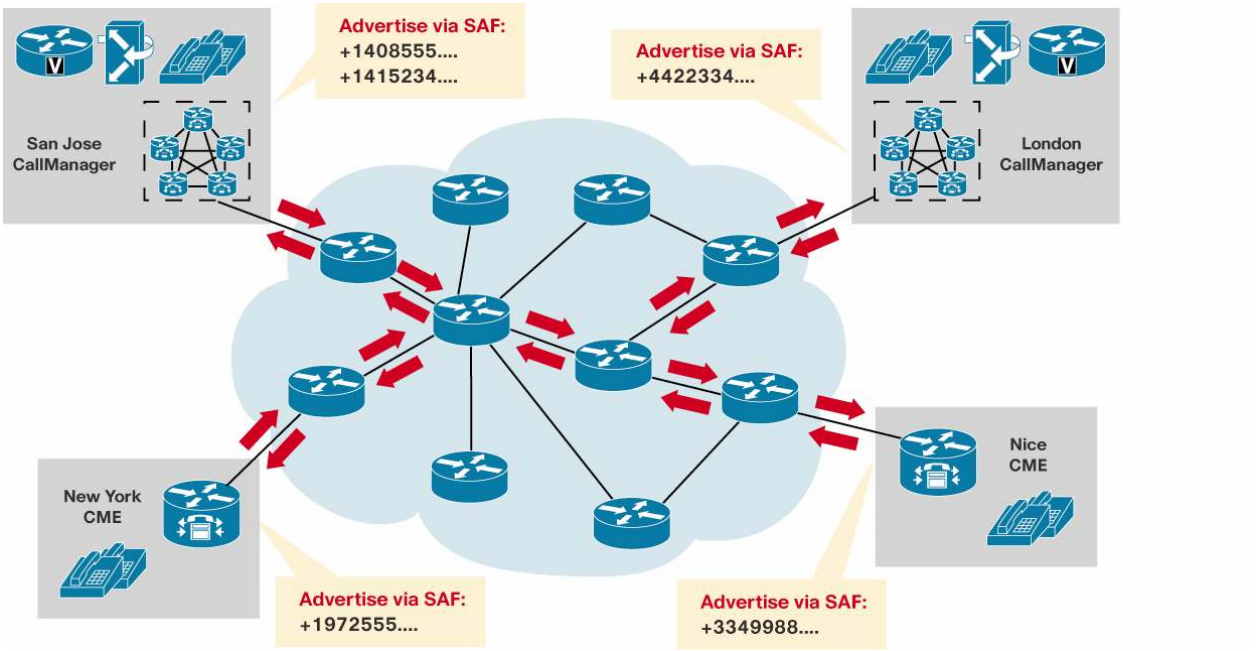

Cisco Unified Communications is the first of many network applications and services to utilize SAF. Future Cisco IOS Software releases will include support for additional UC products, such as Medianet, IP Video Surveillance, TelePresence, Performance Routing (PfR), and Digital Media Processing applications.

Benefits

- SAF's dynamic discovery and auto-configuration capabilities increase service scalability while simplifying deployment tasks.
- SAF enables fast notification of new network services and the accessibility of existing services, resulting in increased service availability and resiliency.
- SAF reduces operational costs by reducing the time to deploy and reconfigure a network based service.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6599/products_ios_protocol_group_home.html

Product Management Contact

Ernie Mikulic, emikulic@cisco.com

4.8 BGP Graceful Restart per Neighbor

BGP Graceful Restart per Neighbor introduces the ability to enable or disable BGP graceful restart for every individual BGP neighbor. Three new methods of configuring BGP graceful restart for BGP peers, in addition to the existing global BGP graceful restart configuration, are now available. Graceful restart can be enabled or disabled for a BGP peer or a BGP peer group using the neighbor ha-mode graceful-restart command, or a BGP peer can inherit a graceful restart configuration from a BGP peer-session template using the ha-mode graceful-restart command.

Although BGP graceful restart is disabled by default, the existing global command enables graceful restart for all BGP neighbors regardless of their capabilities. The ability to enable or disable BGP graceful restart for individual BGP neighbors provides a greater level of control for a network administrator.

When the BGP graceful restart capability is configured for an individual neighbor, each method of configuring graceful restart has the same priority, and the last configuration instance is applied to the neighbor. For example, if global graceful restart is enabled for all BGP neighbors but an individual neighbor is subsequently configured as a member of a peer group for which the graceful restart is disabled, graceful restart is disabled for that neighbor.

The configuration of the restart and stale-path timers is available only with the global bgp graceful-restart command, but the default values are set when the neighbor ha-mode graceful-restart or ha-mode graceful-restart commands are configured. The default values are optimal for most network deployments, and these values should be adjusted only by an experienced network operator.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series • Cisco 10000 Series • Cisco ASR 1000 Routers

Additional Information

http://www.cisco.com/en/US/docs/ios/iproute/configuration/guide/irp_bgp_adv_features.html - wp1056291

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.9 Intermediate System-to-Intermediate System (IS-IS) BFD Support

BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and re-convergence time will be consistent and predictable.

Cisco IOS Software Release 15.0(1)M adds IS-IS BFD support.

4.10 Benefits of Using BFD for Failure Detection

When any feature is deployed, it is important to consider the alternatives and not the potential trade-off. The closest alternative to BFD in conventional IS-IS deployments is the use of modified failure detection mechanisms for IS-IS routing protocols. If fast hellos for IS-IS are used, Interior Gateway Protocol (IGP) protocols reduce their failure detection mechanisms to a minimum of one second.

Benefits to implementing BFD over reduced timer mechanisms for routing protocols include:

- Provides failure detection in less than one second, although reducing the EIGRP, IS-IS, and OSPF timers can result in minimum detection timer of one to two seconds.
- Enables generic and consistent failure detection mechanism for IS-IS, as BFD is not tied to any particular routing protocol.
- Reduces CPU cycles, over control-plane-based IS-IS timers, because some parts of BFD can be distributed to the data plane.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series • Cisco 10000 Series • Cisco ASR 1000 Router

Additional Information

http://www.cisco.com/en/US/docs/ios/iproute/configuration/guide/irp_bfd.html - wp1054179

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.11 IS-IS VRF Support

This feature provides multiple VPN routing and forwarding VRF-aware IS-IS instances. This allows network providers to separate routing protocol information and propagate it to the appropriate routing table and network neighbors. Using one router with VRF functionality is more cost-effective than using separate routers to separate and forward the routing information.

Configuration Example

Router(config)# router isis
Router(config-router)# vrf first
Router(config-router)# net 49.000b.0000.0001.ffff.00
Router(config-router)# is-type level-1

Router(config)# interface POS 6/1
Router(config-if)# ip vrf forwarding first
Router(config-if)# ip address 172.16.2.1 255.255.255.0
Router(config-if)# ip router isis
Router(config-if)# no shutdown

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series • Cisco 10000 Series • Cisco ASR 1000 Series

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.12 MPLS VPN—Inter-AS Option AB

The MPLS VPN—Inter-AS Option AB feature combines the best functionality of an Inter-AS Option (10) A and Inter-AS Option (10) B network to allow a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) service provider to interconnect different autonomous systems to provide VPN services. These networks are defined in RFC 4364 section 10 "Multi-AS Backbones," option "a" and option "b" respectively.

When different autonomous systems are interconnected in an MPLS VPN—Inter-AS Option AB configuration, the entire network's configuration is scaled, simplified, and maintains IP Quality of Service (QoS) functions between Autonomous System Boundary Router (ASBR) peers.

In an Inter-AS Option A network, ASBR peers are connected by multiple sub-interfaces with at least one interface VPN that spans the two autonomous systems. These ASBRs associate each sub-interface with a VRF and a BGP session to signal unlabeled IP prefixes. As a result, traffic between the back-to-back VRFs is IP-based. In this

scenario, the VPNs are isolated from each other and since the traffic is IP, QoS mechanisms that operate on IP traffic can be applied to achieve customer Service Level Agreements (SLAs). The downside of this configuration is that there needs to be one BGP session for each sub-interface (and at least one sub-interface for each VPN), which causes scalability concerns as this network grows.

In an Inter-AS Option B network, ASBR peers are connected by one or more sub-interfaces that are enabled to receive MPLS traffic. A Multi-protocol Border Gateway Protocol (MP-BGP) session is used to distribute labeled VPN prefixes between the ASBR. As a result, the traffic that flows between them is labeled. The downside of this configuration is that because the traffic is MPLS, QoS mechanisms that can only be applied to IP traffic cannot be applied and the VRFs cannot be isolated.

Figure 9. MPLS VPN Inter-AS Option AB Topology

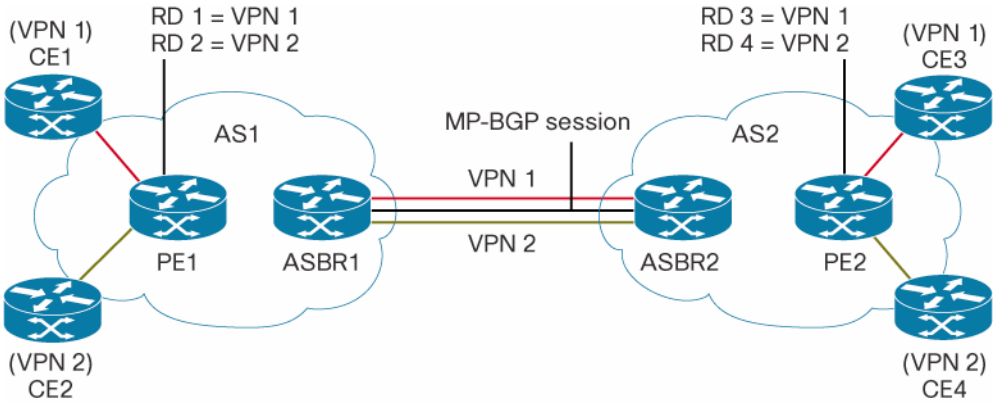

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series

Additional Information

http://www.cisco.com/en/US/docs/ios/mpls/configuration/guide/mp_vpn_ias_optab.html

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.13 BGP Route Target (RT) Changes without PE-CE Impact

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by Border Gateway Protocol (BGP) extended communities. Distribution of VPN routing information works as follows:

1. When a VPN route learned from a CE router is injected into BGP, a list of VPN route target extended community attributes is associated with it. Typically the list of route target community values is set from an export list of route targets associated with the VRF from which the route was learned.
2. An import list of route target extended communities is associated with each VRF. The import list defines route target extended community attributes a route must have for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, or C—is imported into the VRF.

This feature allows changes to the RT value without impacting PE-CE routing or Interior Gateway Protocol (IGP) routes.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series • Cisco 7600 Series • Cisco 10,000 Series • Cisco ASR 1000 Router

Additional Information

http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsmvpns.html - wp1037067

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.14 IS-IS MIB

This feature introduces MIB support for the Intermediate System-to-Intermediate System (IS-IS) link-state routing protocol.

The IS-IS MIB feature offers network providers improved capabilities to continuously monitor the changing state of an IS-IS network by use of MIB objects to gather information relating to protocol parameters and trap notification objects that can signal the occurrence of significant protocol events such as an authentication failure or a mismatch in area addresses between Intermediate Systems (IS). The protocol information collected by the IS-IS MIB objects and trap objects can be used by the network manager to derive statistics that can help monitor and improve overall network performance.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series • Cisco 10,000 Series • Cisco ASR 1000 Router

Additional Information

http://www.cisco.com/en/US/docs/ios/iproute/configuration/guide/is_mib.html

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

4.15 MPLS VPN—BGP Local Convergence

The Duration of a Link Failure

Within a Layer 3 VPN network, the failure of a PE-CE link can cause a loss of connectivity to a customer site, which is detrimental to time-sensitive applications. Several factors contribute to the duration of such an outage:

- The time to detect the failure
- Programming the forwarding
- The convergence of BGP (particularly in large networks: the restored traffic's arrival time at its destination varies according to the prefix)

Figure 10. Sample PE-CE topology, showing original ("Primary") traffic path

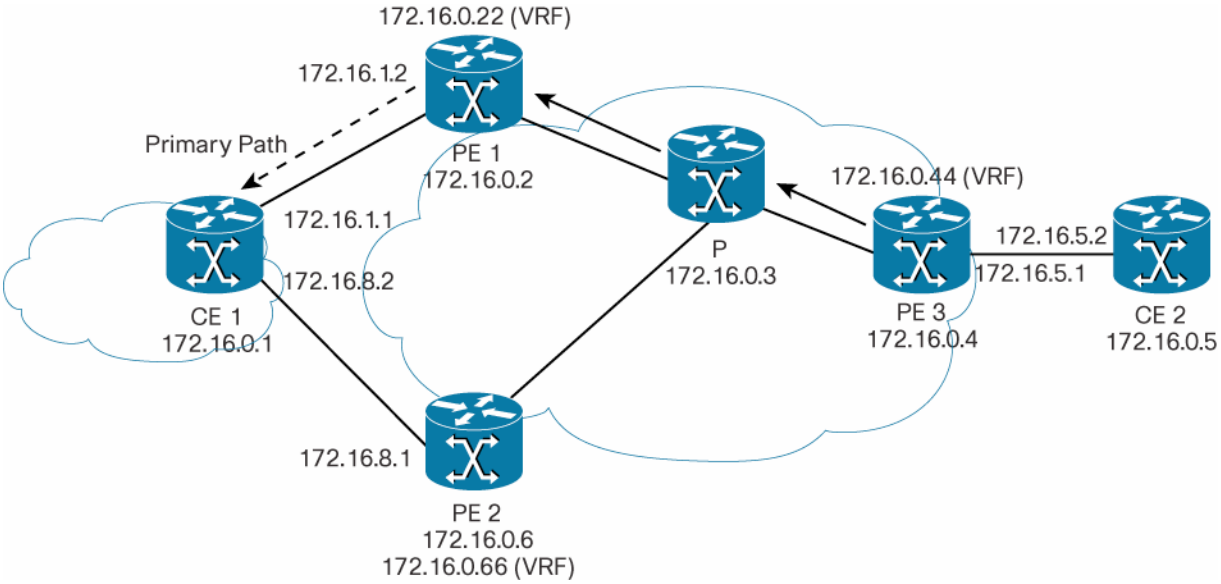

BGP's Usual Response

Under default BGP behavior, after a PE-CE link failure and its detection in routing, all the BGP paths via the failing link are removed. BGP runs the bestpath algorithm on the affected prefixes, and selects alternate paths for each. These new paths (which typically include a remote PE) are installed into forwarding. The local labels are removed and BGP withdrawals are sent to all BGP neighbors. As each BGP neighbor receives the withdrawal messages (typically indirectly via route-reflectors), the bestpath algorithm is called and the prefixes are switched to an alternate path. Only then is connectivity restored.

MPLS VPN—BGP Local Convergence Feature Solution: Preserving and Programming the Outgoing Label

The MPLS VPN—BGP Local Convergence feature reduces loss of connectivity time by sending the broken link's traffic over a backup path (shown in Figure 11 below) instead of waiting for total network convergence. The feature accomplishes this by maintaining the local label (for 5 minutes) for the prefixes which are switching from the failing local path to the backup path. Because the label is not freed as had been the usual practice, forwarding can continue to take place, but now to the alternate (backup) path chosen by the bestpath algorithm. Thus, the local label has been applied in place of the failed BGP bestpath label; (this is sometimes called "label swapping"). Traffic thereby is locally restored, while the network propagation of the BGP withdrawal messages takes place. Eventually, the egress PE router converges and bypasses the local repair.

Figure 11. Sample PE-CE topology, showing replacement ("Backup") traffic path

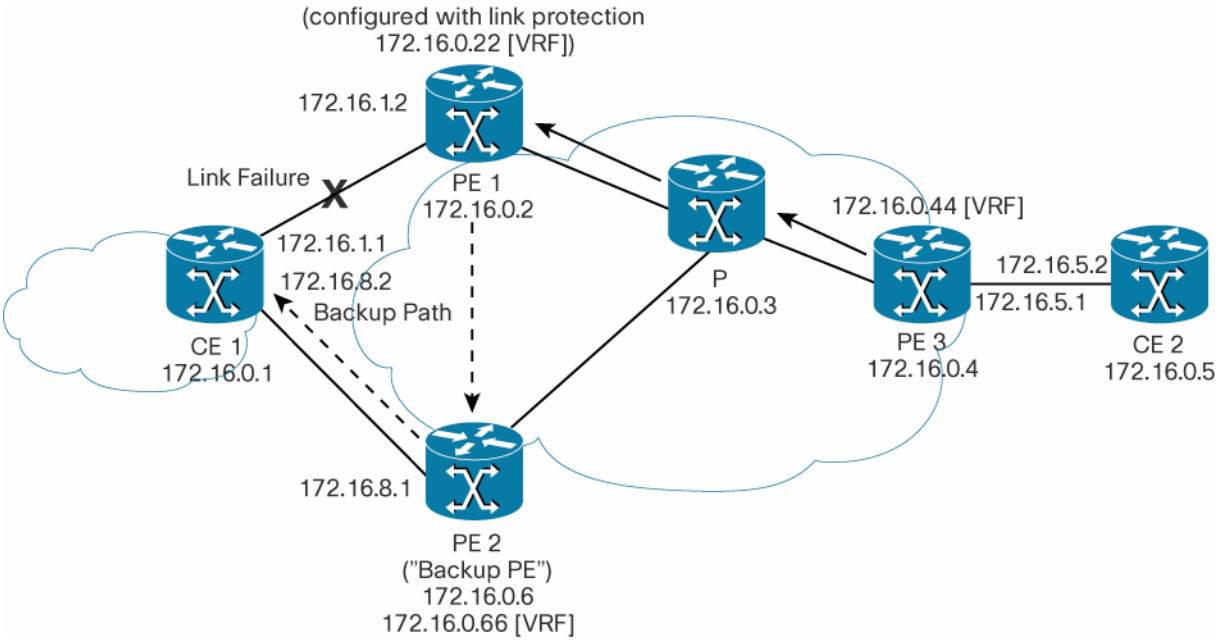

The delay in local label deletion does not modify normal BGP addition and deletion of BGP paths and nets. Rather, BGP re-programs the new backup bestpath into forwarding as usual.

Freeing the Local Label

After the 5-minute label preservation, the preserved local labels are freed. Any BGP prefix that is remote and not Carrier-Supporting-Carrier should not have a local label and thus is cleaned.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7600 Series • Cisco 10,000 Series • Cisco ASR 1000 Router

Additional Information

http://www.cisco.com/en/US/docs/ios/mpls/configuration/guide/mp_vpn_pece_lnk_prot.html

Product Management Contact

Bertrand Duvivier, bduvivie@cisco.com

5. IP Multicast

5.1 Internet Group Management Protocol (IGMP) Static Group Range Support

The IGMP Static Group Range Support feature introduces the capability to configure group ranges in class maps and attach class maps to the **ip igmp static-group** command.

Prior to the introduction of the IGMP Static Group Range Support feature, the **ip igmp static-group** command did not provide an option to specify group ranges. Administering devices that required many **ip igmp static-group** command configurations was challenging in some network environments, because each static group had to be configured individually with the **ip igmp static-group** command, which resulted in configurations that were excessively long and difficult to manage.

Figure 12. IGMP Static Group Enhancement

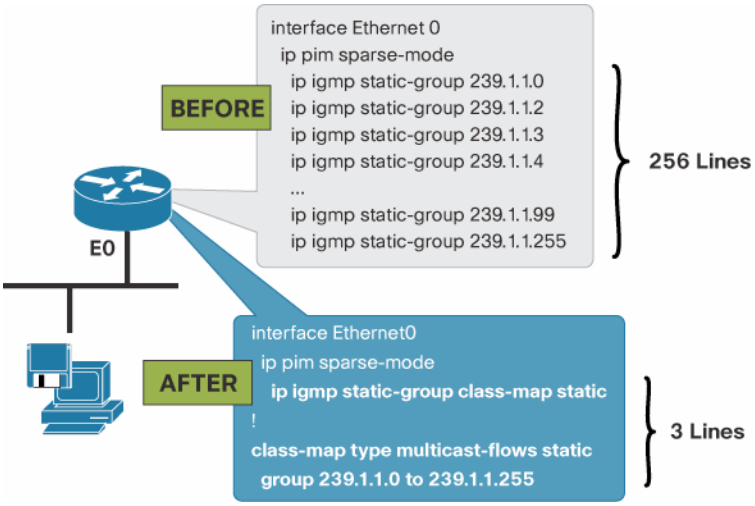

This feature is an enhancement that simplifies the administration of networks with devices that require many interfaces to be configured with many different **ip igmp static-group** command configurations. It introduces a class map that is used to define group ranges, group addresses, Source Specific Multicast (SSM) channels, and SSM channel ranges. Once created, the class map can be attached to interfaces.

Benefits

- Simplifies the administration of devices that require many interfaces to be configured with many different **ip igmp static-group** command configurations by introducing the capability to configure group ranges in class maps and attach class maps to the **ip igmp static-group** command.
- Reduces the number of commands required to administer devices that require many **ip igmp static-group** command configurations

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.2 IP Multicast Load Splitting—Equal Cost Multipath (ECMP) using S, G and Next-hop

The IP Multicast Load Splitting—Equal Cost Multipath (ECMP) Using S, G and Next Hop feature introduces more flexible support for ECMP multicast load splitting by adding support for load splitting based on source and group address and on source, group, and next-hop address. This feature enables multicast traffic from devices that send many streams to groups or that broadcast many channels, such as IPTV servers or MPEG video servers, to be more effectively load split across equal-cost paths. Prior to the introduction of this feature, the Cisco IOS software only supported ECMP multicast load splitting based on source address, which restricted multicast traffic sent by a single source to multiple groups from being load split across equal-cost paths.

Figure 13. Enhanced Multicast Load Splitting

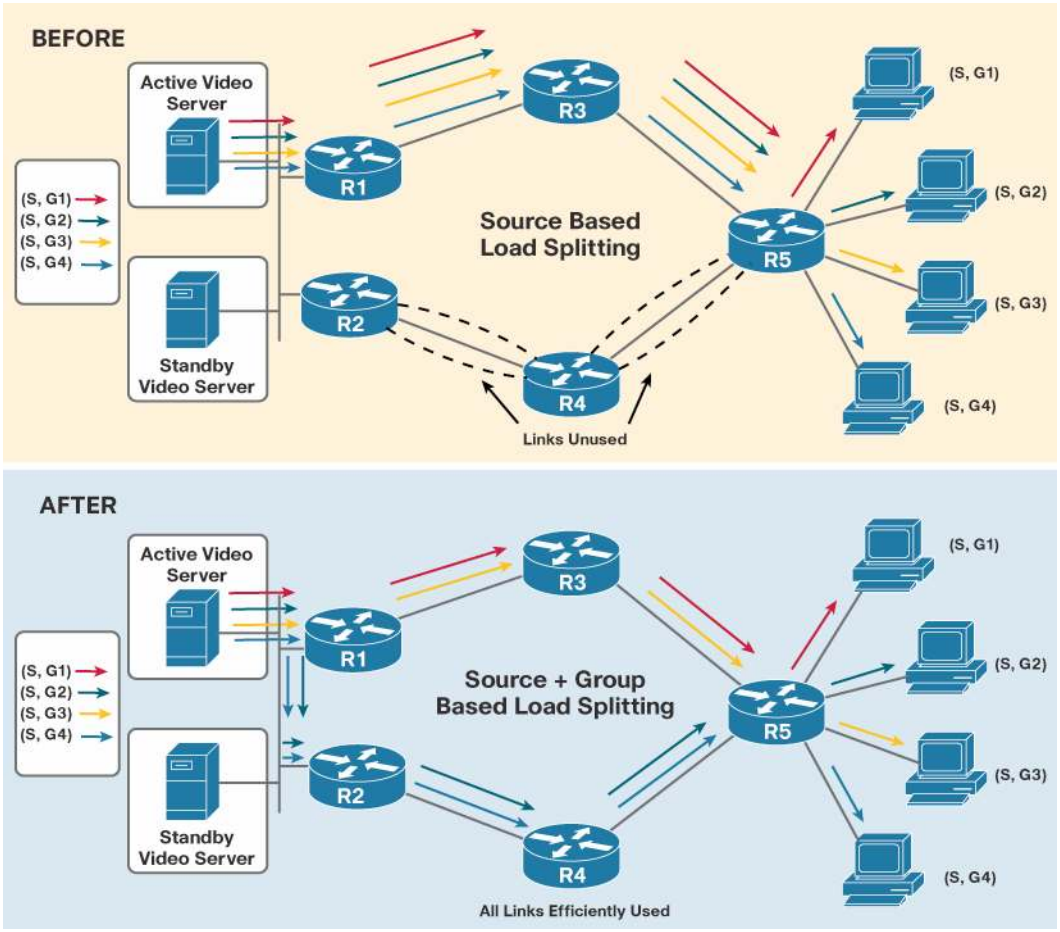

Benefits

- Multicast traffic from different sources, or from different sources and groups, are load split across equal-cost paths to take advantage of multiple paths through the network.
- Optimum use of network resources for multicast traffic

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.3 IPv4 Multicast Address Group Range Support

This feature allows the router to avoid receiving multicast traffic from unauthenticated groups or unauthorized channels. The Multicast Address Group Range Support feature disables IPv4 multicast protocol actions and traffic forwarding for unauthorized groups or channels on all the interfaces in a router for a range of groups.

Benefits

- Provides ability to disable PIM, IGMP, and Multicast Source Discovery Protocol (MSDP) control plane actions
- No IGMP (cache), Protocol Independent Multicast (PIM), Multicast Routing Information Base (MRIB)/Multicast Forwarding Information Base (MFIB) state created for denied groups
- Drops all data packets for denied groups

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.4 IPv6 Multicast Address Group Range Support

The Multicast Address Group Range Support feature disables IPv6 multicast protocol actions and traffic forwarding for unauthorized groups or channels on all the interfaces in a router for a range of groups. It can be used to provide access control to prevent multicast traffic from being received unless the subscriber is authenticated and the channels are authorized as per access control profiles.

Benefits

- Provides ability to disable PIMv6, Multicast Listener Discovery (MLD), and MSDP control plane actions
- No MLD (cache), PIMv6, MRIB)/MFIB state created for denied groups
- Drops all data packets for denied groups

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.5 Multicast MIB VRF Support

The set of Multicast MIBs are made to be Multicast VRF (mVRF)-aware to help customers manage their Cisco devices in a multicast VPN environment using Simple Network Management Protocol (SNMP).

mVRF awareness of Multicast MIBs means, the MIB objects can be queried and set for individual VRF configuration. It has the ability to detect conditions for a trap inside of a VRF and lookup the correct information for that VRF. The traps would be sent to the manager which is configured for that VRF.

The following set of Multicast MIBs is made to be mVRF aware.

- PIM-MIB
- CISCO-PIM-MIB
- MSDP-MIB

- IGMP-STD-MIB
- IPMROUTE-STD-MIB
- CISCO-IPMROUTE-MIB

SNMP contexts provide VPN users with a secure way of accessing MIB data. When a VPN is associated with a context, that VPN's specific MIB data exists in that context. Associating a VPN with a context enables service providers to manage networks with multiple VPNs. By creating and associating a context with a VPN, a provider can prevent the users of one VPN from accessing information about users of other VPNs on the same networking device. VPN-aware SNMP requires an agreement between SNMP manager and agent entities operating in a VPN environment on a mapping between the SNMP security name and the VPN ID.

When an SNMP request is received by the agent, the agent will use the context-string to identify the mVRF to which this context-string is mapped. As per the access policy configured, MIB objects specific to that mVRF will get returned. If the SNMP request is made without context-string, MIB objects specific to default mVRF will get returned.

When a trap is enabled, the notification would be sent to the manager which is configured for that VRF. If SNMP manager is not configured for a VRF, the traps for that VRF would not be sent. MIBs that are not VRF aware will not be able to report on an event that occurs in a VRF. They will only report on events in the default/global routing tables.

Benefits

- Enables SNMP gets and sets to be made to the individual VRFs
- Allows the MIB to have the ability to detect conditions for a trap inside of a VRF and lookup the additional information in the VRF context
- Sends traps to a manager located inside a VRF

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.6 Multicast VPN Extranet Support

This feature introduces extranet capabilities in Multicast VPN. Extranet allows VPN closed user groups to share information and common multicast information to be distributed across multiple VPN customers.

An extranet can be viewed as part of a company's intranet that is extended to users outside the company. An extranet is a VPN connecting the corporate site or sites to external business partners or suppliers, to securely share part of a business's information or operations among them. MPLS VPNs inherently provide security, ensuring that users access only appropriate information. The MPLS VPN Extranet service offers users unicast connectivity without comprising the integrity of their corporate data. Multicast VPN Extranet extends this service offering to include multicast connectivity to the extranet community of interest. It allows service providers to offer next generation of flexible extranet services that enable business partnerships between different enterprises.

As shown in Figure 14 below, the Multicast VPN Extranet feature allows network providers to source multicast content from VPN Green into VPN Red.

Figure 14. Two Extranet Scenarios with Ingress and Egress Edge Router Packet Replication across VPNs

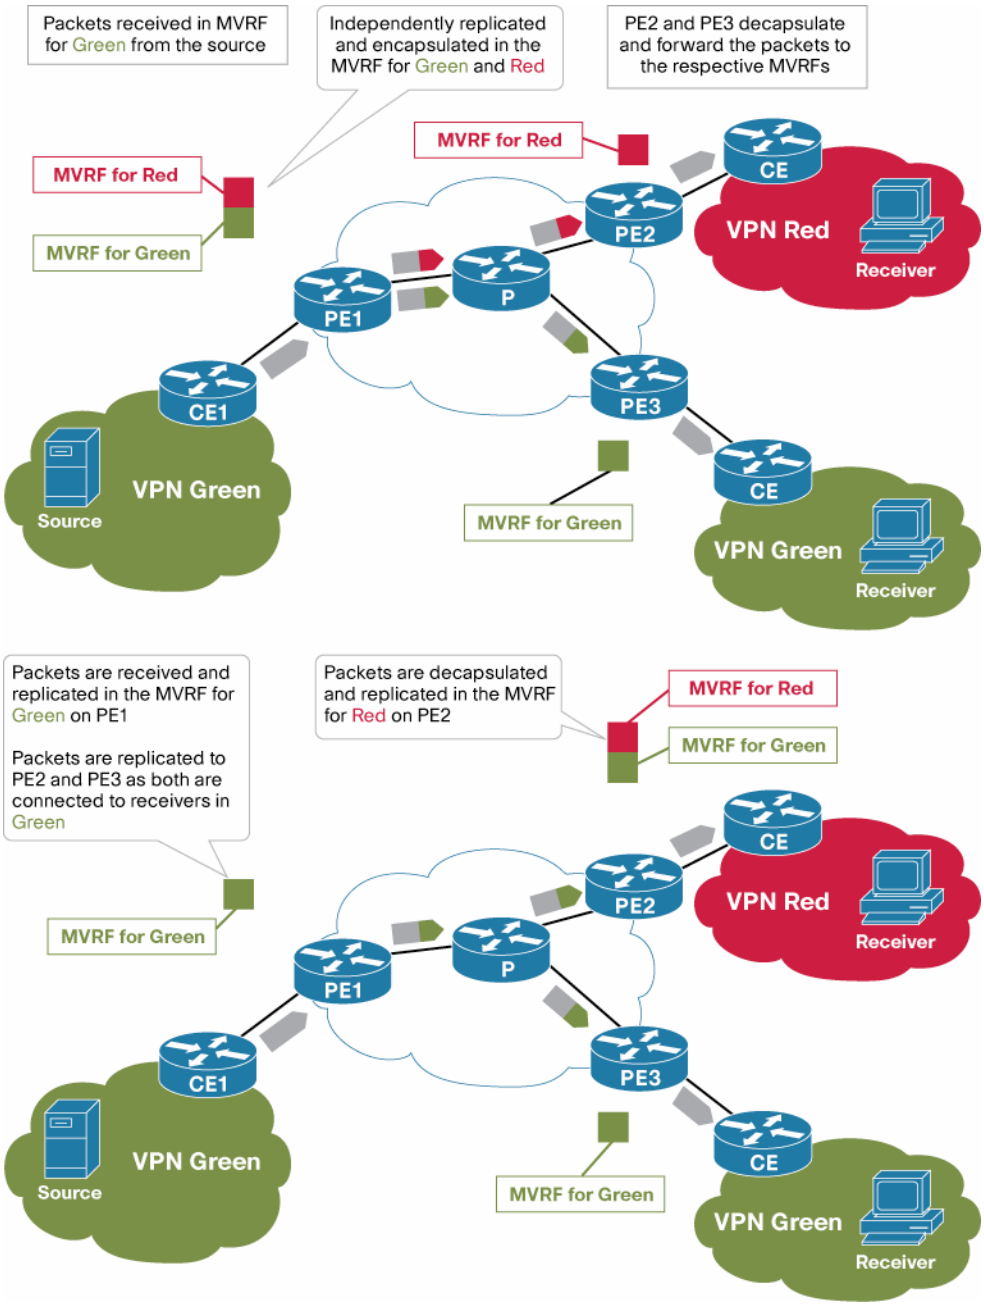

Benefits

- Scalable and efficient method to transport and replicate customer multicast information across an MPLS network between different VPNs
- Extranet Multicast VPN (MVPN) solves these business needs
  - Efficient content distribution between Enterprises
  - Efficient sharing of multicast resources with external or business partners
  - Efficient content distribution from Service Providers or content provider to its different VPN customers
  - Integrated transparently with unicast MPLS VPN services

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

http://www.cisco.com/en/US/technologies/tk648/tk828/tk363/technologies_white_paper0900aecd802aea84.html

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.7 Multicast VPN VRF Select

The Multicast VPN Extranet VRF Select feature provides the capability for Reverse Path Forwarding (RPF) lookups to be performed to the same source address in different VPN routing and forwarding (VRF) instances using the group address as the VRF selector.

Prior to the introduction of the Multicast VPN Extranet VRF Select feature, RPF lookups for a source address could only be performed in a single VRF, that is, in the VRF where Internet Group Management Protocol (IGMP) or PIM joins are received, in the VRF learned from BGP imported routes, or in the VRF specified in IP multicast static routes (mroutes), when RPF for an extranet MVPN is configured using static mroutes. In those cases, the source VRF is solely determined by the source address or the way the source address was learned.

Figure 15 illustrates an extranet MVPN topology with the Multicast VPN VRF Select feature configured. In this topology, (S, G1) and (S, G2) PIM joins originating from VPN-Green, the receiver VRF, are forwarded to PE1, the receiver PE. Based on the group-based VRF selection policies configured, PE1 sends the PIM joins to VPN-Red and VPN-Blue for groups G1 and G2, respectively.

Figure 15. RPF Lookups Using Group-Based VRF Selection Policies

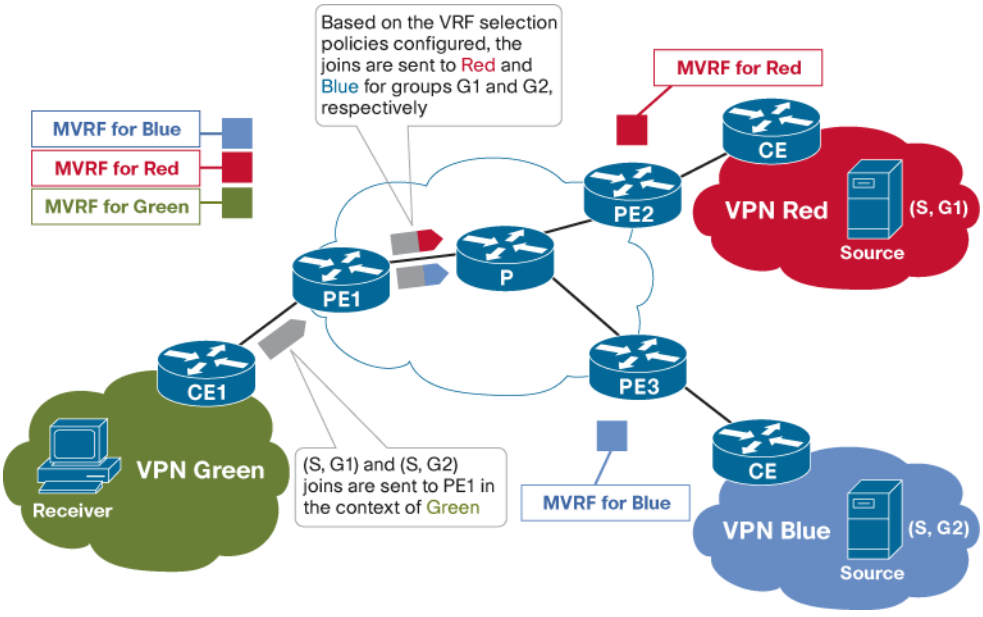

Benefits

- Enhances extranet MVPNs by enabling service providers to distribute content streams coming in from multiple MVPNs and then redistributing them

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

5.8 PIM Triggered Joins

The PIM Triggered Joins feature is a high availability (HA) multicast enhancement that improves the re-convergence of multicast routes (mroutes). This feature utilizes the Generation ID (GenID) value as a mechanism to trigger adjacent Protocol Independent Multicast (PIM) neighbors on an interface to send PIM join messages for all (*, G) and (S, G) mroutes that use that interface as a reverse path forwarding (RPF) interface, immediately reestablishing those states.

Benefits

- Helps prevent mroute state in the control plane from timing out
- Prevents temporary blackouts of multicast traffic

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series Routers

Additional Information

http://www.cisco.com/go/multicast

Product Management Contact

Ritesh Mukherjee, ritmukhe@cisco.com

6. Call Admissions Control

6.1 CISCO IOS RSVP Interface Based Receiver Proxy

The RSVP Interface-Based Receiver Proxy feature enables the use of RSVP to signal reservations and guarantee bandwidth on behalf of a receiver that does not support RSVP. This is performed by terminating the PATH message and generating a RESV message on the RSVP-capable router closest to the endpoint. An example is a video-on-demand (VoD) flow from a video server to a set-top box, which is a computer that acts as a receiver and decodes the incoming signal from a satellite dish, a cable network, or a telephone line.

Because set-top boxes may not support RSVP natively, end-to-end RSVP reservations between a video server and a set-top box cannot be configured. Instead, a proxy on the router that is closest to that set-top box may be configured.

The router terminates the end-to-end sessions for many set-top boxes and performs admission control on the outgoing interface of the PATH message, where the receiver proxy is configured, as a proxy for Call Admission Control (CAC) on the router-to-set-top link. The interface-based receiver proxy determines which PATH messages to terminate by looking at the interfaces used by the traffic flow.

An RSVP interface-based receiver proxy to terminate PATH messages going out a specified interface with a specific action (reply with RESV, or reject) may be configured. The most common application is to configure the receiver

proxy on the edge of an administrative domain on interdomain interfaces. The router then terminates PATH messages going out the administrative domain, while still permitting PATH messages transitioning through the router within the same administrative domain to continue downstream.

In the video-on-demand example described above, the last-hop Layer 3 router supporting RSVP implements the receiver proxy, which is then configured on the interfaces facing the Layer 2 distribution network (for example, Digital Subscriber Line access (DSLAM) or cable distribution). Also, since RSVP is running and performing CAC on the router with the receiver proxy, you can configure RSVP enhancements such as local policy and Common Open Policy Service (COPS) for more fine-grained control on video flow CAC.

The router terminates the end-to-end sessions for many set-top boxes, with the assumption that the link between the router and the set-top box never becomes congested or, more likely, in the case of congestion, that the voice and video traffic from the router gets the highest priority and access to the bandwidth.

The router determines which sessions to terminate by looking at the interfaces used by the traffic flow. An RSVP interface-based receiver proxy to terminate PATH messages going out a specific interface, with a specific action (reply with RESV or reject) may be configured. The most common application of this is to configure the receiver proxy on the edge of an administrative domain, on interdomain interfaces. The router then terminates PATH messages going out the administrative domain, while still permitting PATH messages transitioning through the router within the same administrative domain to continue downstream.

Benefits

Previously, a receiver proxy for every separate RSVP stream or set-top box had to be configured. Now the proxy by outbound interface may be configured. For example, if there are 100 set-top boxes downstream from the proxy router, 100 proxies had to be configured. With this enhancement, only the outbound interface(s) have to be configured. In addition, the receiver proxy is guaranteed to terminate the reservation only on the last hop within the core network. Nodes that may function as transit nodes for some PATH messages but should proxy others, depending on their placement in the network, can perform the correct functions on a flow-by-flow basis.

In the video-on-demand example described above, a PATH message that transits through an edge router to another edge router (around the edge) is not terminated, whereas an otherwise identical PATH message that actually exits the core and transitions to the distribution network is terminated. This allows for more accurate CAC in the network, and also simplifies and reduces configuration requirements.

Hardware

Routers	• Cisco Integrated Services Routers
	• Cisco 7200 Series
	• Cisco 7600 Series
	  ◦ SUP2/MSFC2
	  ◦ SUP720/MSFC3
	  ◦ CMM
	  ◦ SUP32/MSFC2A

Additional Information

http://www.cisco.com/en/US/products/ps6550/products_ios_technology_home.html

Product Management Contact

Bertrand Duvivier (bduvivie@cisco.com)

6.2 Resource Reservation Protocol (RSVP) Fast Link Repair

RSVP is used in some applications to make resource reservations for IPv4 unicast flows. RSVP provides for dynamic adaptation when routing changes occur: when a route changes, the next Path and Resv message refreshes

establish path and reservation states along the new route. Depending on the configured refresh interval this re-route will happen in the order of tens of seconds. During this time, the QoS of flows may not be honored as congestion may occur while data packets travel over links where reservations are not in place yet.

The RSVP Fast Link Repair capability addresses this issue. In order to provide faster adaptation to routing changes, without the overhead of refresh period, RSVP registers and receives notifications when routes change and triggers state refreshes for the affected destinations. These triggered refreshes use the new route information and therefore will install reservations over the new path.

Benefits

- Allows sub-second response time to routing changes, which is critically important for voice and video deployments.

Hardware

Routers	• Cisco 7600 Series
	  ◦ SUP2/MSFC2
	  ◦ SUP720/MSFC3
	  ◦ SUP32/MSFC2A
	• Cisco Integrated Services Routers
	• Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6550/products_ios_technology_home.html

Product Management Contact

Bertrand Duvivier (bduvivie@cisco.com)

6.3 RSVP VRF Lite Aware Admission Control

The RSVP VRF Lite Aware Admission Control feature provides the benefits of RSVP in a VRF-lite environment to include the following:

- Guaranteed QoS through explicit admission control
- Virtualization
- Security
- Separation of routing contexts
- Overlapping of IP addresses

An RSVP flow is identified by its tuple, which includes its destination IP address, its destination port, and its protocol. This tuple should be unique on all the nodes along the path from the sender to the receiver. In the context of the global routing domain, each flow can be uniquely identified through its tuple.

However, with the implementation of virtual routing and forwarding (VRF), a separate instance of the routing and forwarding table for each VRF routing domain can exist. Each of the VRF instances has its own address pool range, which could overlap between VRF routing domains. This poses a problem to the existing implementation of RSVP, where sessions are identified by the tuple. Sessions with the same tuple can exist in the context of different VRF domains. To solve the problem, the tuple has to be extended to take into account the VRF instance. The new tuple has a VRF ID, a destination IP address, a destination port, and a protocol.

Figure 16. VRF-lite deployment scenario

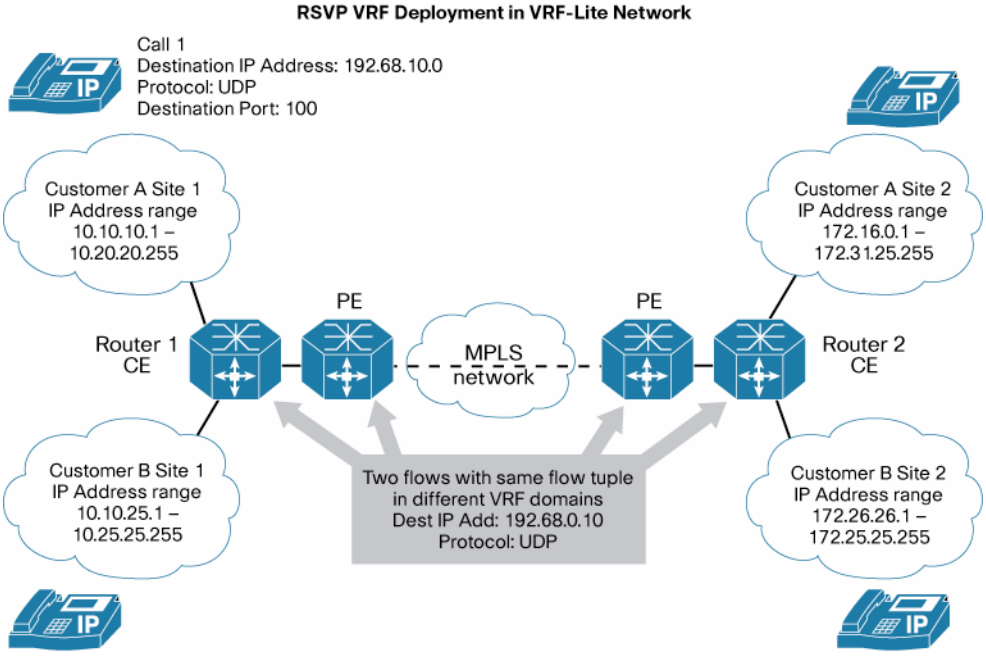

Figure 16 above illustrates VRF lite configured on Router 1 customer edge (CE) and Router 2 CE, and MPLS-VPN configured between the provider edge (PE) routers. In such a deployment scenario, the RSVP implementation needs to be VRF aware in the CE routers and also in the PE routers; that is, the flows must be recognized in the context of the VRF domain in which the sender and receiver of the flow reside.

On the CE routers, with VRF lite functionality, VRF is identified based on the VRF configured on the incoming interface; that is, on the interface facing the customer site and the interface facing the PE.

On the PE routers, on the interface facing the CE, the VRF is identified by the incoming interface on which the RSVP message arrives. However, for RSVP messages coming from the core-facing interface, RSVP requires a protocol extension in order to be able to communicate that the VRF domain of the flow is out-of-band.

Figure 17 below shows a deployment scenario in which the customer site is directly connected to the PE router. The deployment scenario in Figure 17 is a subset of the one depicted in Figure 16 above.

Figure 17. VRF-lite deployment scenario

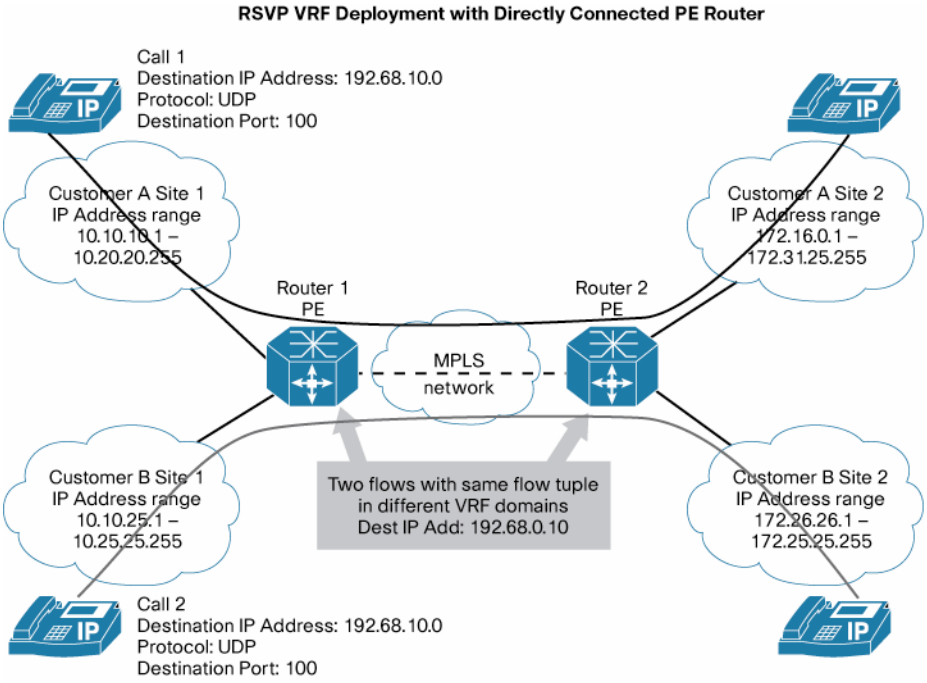

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

http://www.cisco.com/en/US/products/ps6550/products_ios_technology_home.html

Product Management Contact

Bertrand Duvivier (bduvivie@cisco.com)

7. High Availability

7.1 Bidirectional Forwarding Detection (BFD) Enhancements

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. This detection is typically accomplished through hardware detection mechanisms. However, not all of the hardware mechanisms have the capability to detect failures, for example Ethernet failures.

BFD also provides a consistent failure detection method for network administrators. This enables network administrators to use BFD for detecting forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, and network profiling. This makes planning easier, and re-convergence times more consistent and predictable.

Initial BFD support was introduced in Release 12.2(4)T. Release 12.4(15)T added support for Cisco 1800, 2800, and 3800 Series Routers.

The following BFD feature enhancements are being introduced in Cisco IOS Release 15.0(1)M:

- BFD—Client for IPv4 static route
- BFD—VRF aware support,
- BFD—WAN interfaces support,

Bidirectional Forwarding Detection (BFD) client for IPv4 Static Routes

BFD client for IPv4 static routes extends the mechanisms that are available at Layer 3 for detecting failures in the forwarding path between two adjacent peers when no routing protocol is enabled between those two peers but instead static route is used.

In the case of static routes, scenarios exist where the next-hop in a static route is down but the interface remains up. BFD client for static route provides fast failure detection to facilitate the rerouting of traffic via an alternate path so that traffic will not be black-holed.

Note: Release 15.0(1)M only supports BFD static routes for IPv4. BFD support for IPv6, and IPv6 static routes BFD client, is planned for future Cisco IOS Release 15 M and T releases.

Bidirectional Forwarding Detection (BFD) VRF support

BFD VRF support extends BFD failure detection capability within a VRF context. It is now possible to run BFD from a VRF based interface, so that any failure in the forwarding path between PE and CE devices can be detected even though the physical link might still be up.

The combination of BFD VRF support along with the different embedded OAM MPLS tools such as MPLS Ping and Traceroute provides network operators a comprehensive end-to-end solution to address overall network reliability, and enhance L3VPN service availability.

Bidirectional Forwarding Detection (BFD) WAN interface Support

For failure detection on WAN interfaces, customers have typically relied on physical layer characteristics, such as Loss of Signal (LOS) for Packet-over-SONET (POS) interfaces. For IP traffic over WAN interfaces, situations exist where the next-hop is not reachable but the interface remains up and hence the lack of reachability in the forwarding path is not detected.

BFD for WAN interface enables the use of BFD as fast failure detection in the forwarding path for interfaces such as: ATM, POS, and Frame Relay. In addition, BFD for VLAN interface (802.1q) is also supported.

Deployment Example

Many network providers are launching ADSL2+ services aggregated on IP-DSLAMs and carried over Metro Ethernet networks towards the PE. The CPE is connected to the IP-DSLAM via ADSL2+ which is via ATM interfaces. BFD can provide a standard failure detection mechanism in this case for the ATM interfaces.

Supported Interfaces	• ATM interface with AAL5 MUX, AAL5 SNAP, AAL0 encapsulations • ATM sub interface • POS interface with HDLC and PPP Encapsulations • POS sub interface • Serial interface, Serial interfaces with FR Encapsulation • Serial sub interface with FR Encapsulation • VLAN interface (802.1q)

Product Management contact

Hari Rakotoranto (hrakotor@cisco.com)

8. Embedded Management

8.1 Flexible NetFlow and Network Based Application Recognition (NBAR) Integration

Flexible NetFlow is the next-generation in flow technology. It allows optimization of the network infrastructure, reducing operation costs, improved capacity planning and security incident detection with increased NetFlow flexibility and scalability beyond other flow based technologies available today.

Benefits

Flexible NetFlow advantages include:

- Flexibility, scalability, and customization of flow data
- The ability to monitor a wider range of packet information
- Enhanced network anomaly and security detection
- User configurable flow information to perform customized traffic identification and the ability to focus and monitor specific network behavior
- Convergence of multiple accounting technologies into one accounting mechanism
- Multiple configurable flow caches

NBAR is an intelligent classification engine in Cisco IOS Software that can recognize a wide variety of applications, including Web-based and client/server applications. Once the applications are recognized, the network can invoke required services for that particular application.

NBAR can classify applications that use:

- Statically assigned Transfer Control Protocol (TCP) and User Datagram Protocol (UDP) port numbers
- Non-UDP and non-TCP IP protocols
- Dynamically assigned TCP and UCP port numbers negotiated during connection establishment; Stateful inspection is required for classification of applications and protocols. This is the ability to discover data connections that will be classified, by passing the control connections over the data connection port where assignments are made.
- Sub-port classification; Classification of HTTP (URLs, mime or host names) and Citrix applications Independent Computing Architecture (ICA) traffic based on published application name)
- Classification based on deep packet inspection and multiple application-specific attributes. Real-Time Transport Protocol (RTP) Payload Classification is based on this algorithm, in which the packet is classified as RTP, based on multiple attributes in the RTP header.

Figure 18. Differences between Flexible NetFlow and NBAR

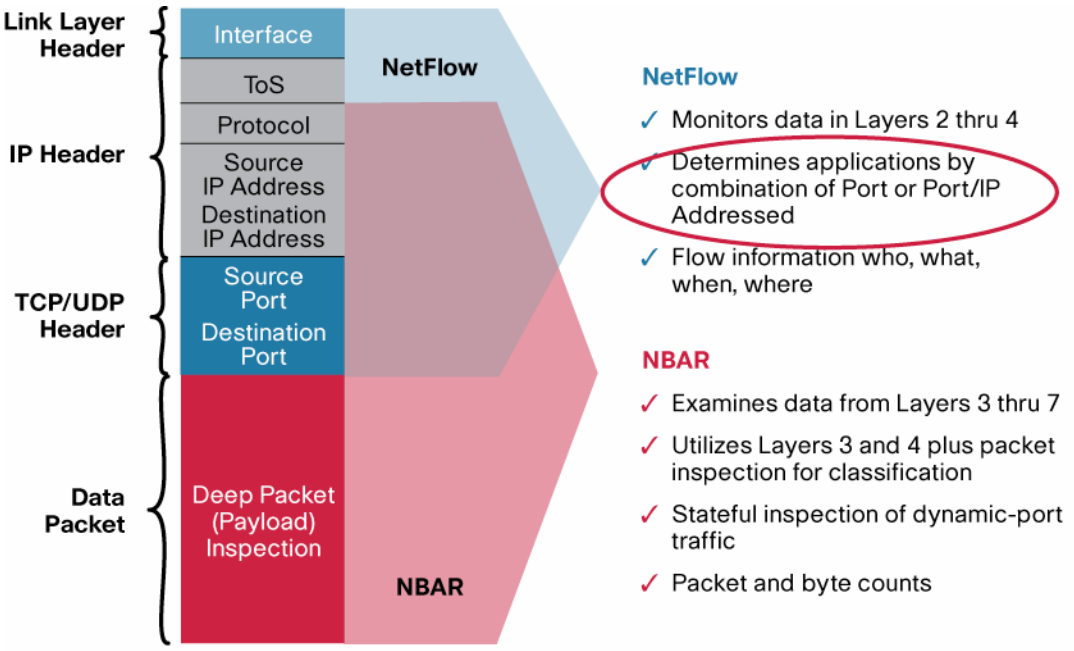

Flexible NetFlow and NBAR Integration combines strengths of both features to provide per Flow application visibility. Flexible NetFlow currently provides L2/L3/L4 Flow statistics while NBAR provide application-level statistics per interface using the protocol discovery MIB. The integration of both features enables a new “Application name” field discovered by NBAR to be introduced into the Flexible NetFlow flow record definition. This enables Layer 2 through 7 per-flow application visibility in the network.

Benefits

- Enabling application visibility per Flow

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 and 7300 Series Routers

Additional Information

http://www.cisco.com/go/netflow

http://www.cisco.com/go/fnf

http://www.cisco.com/go/nbar

Product Management Contact

Jean-Charles Griviaud, jgriviau@cisco.com

8.2 Cisco IOS Embedded Event Manager (EEM) Version 3.1

EEM is a powerful and flexible subsystem that provides real-time network event detection and onboard automation. Customers can use EEM to create and run programs or scripts directly on a router or switch. The scripts are referred to as EEM Policies and can be programmed using a simple CLI-based interface or using a scripting language called Tool Command Language (Tcl). EEM allows customers to harness the significant intelligence within Cisco IOS Software to respond to real-time events, automate tasks, create customer commands and take local automated action based on conditions detected by the Cisco IOS Software itself.

Due to its powerful onboard automation engine and the wide integration with various IOS subsystems, EEM can be used in many different areas including automatic fault detection and alert, automatic data collection and reporting, automated troubleshooting, high availability, zero-touch deployment, and more.

EEM is supported on a wide range of Cisco products and provides a consistent onboard programmability and automation interface. Each new version of the EEM feature introduces new event detectors or new capabilities. Please refer to the Cisco documentation for detailed information.

Product Architecture

The Cisco IOS Embedded Event Manager is a primarily product independent software feature consisting of a series of Event Detectors, an Embedded Event Manager Server, and interfaces to allow action routines called Policies to be invoked. The diagram in Figure 19 illustrates the EEM components.

Figure 19. EEM Architecture

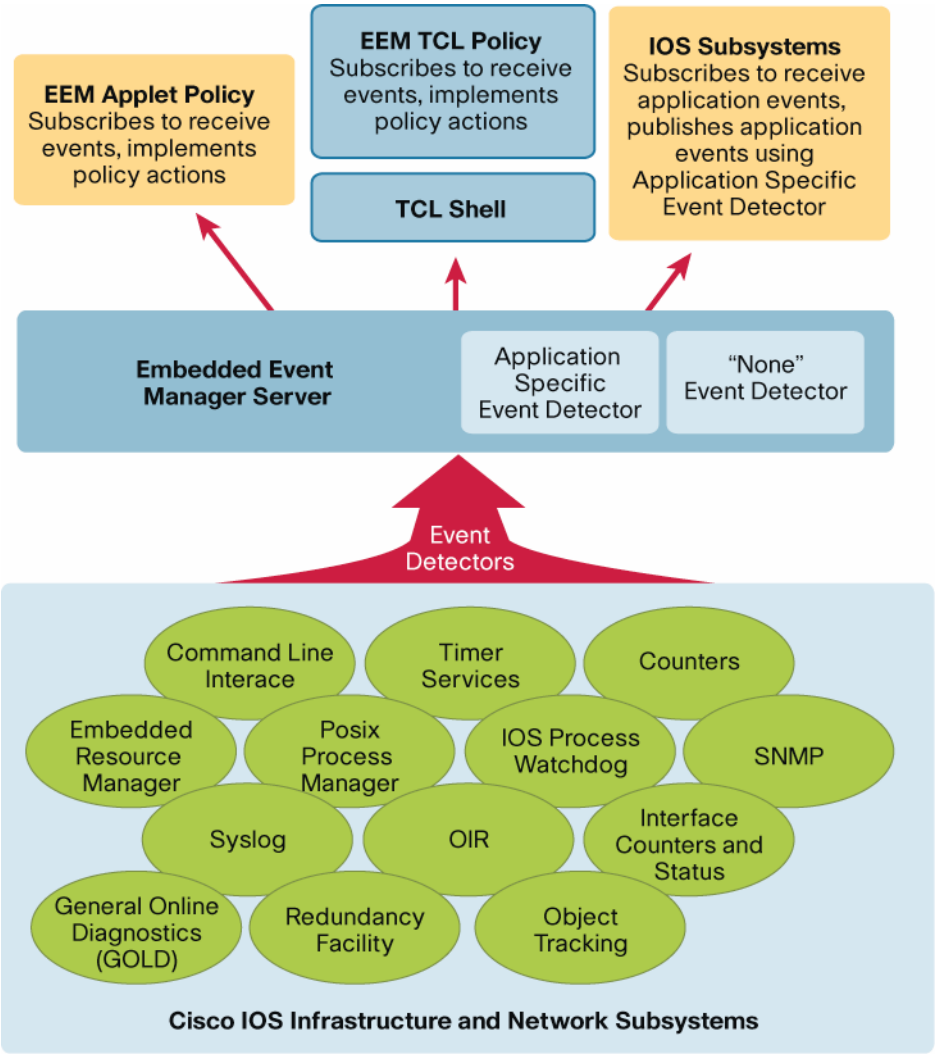

Notice there are two types of EEM Policies:

- Applet Policies—Easy-to-use interface, defined using the configuration CLI
- Tcl Policies—More flexible and extensive capabilities, defined using the Tcl programming language

Once one or more policies are defined, the Event Detector software will watch for the conditions that match those defined by the policy. When a condition occurs, the event is passed to the Event Manager Server. The server then

invokes any policy that has registered for that particular event. The actions defined within the policy are then carried out.

Each type of event has specific options, parameters and detailed information that is available to the policy when it is invoked. All of these details are described in the Cisco IOS documentation.

8.3 EEM Version 3.1 Enhancements

The following enhancements for event detection, notification mechanism, and command execution capabilities are being introduced in Release 15.0(1)M:

- SNMP Object Event Detector
  - Intercepts incoming SNMP Get and Set requests and generate event
  - Allow user to build EEM policies and provide customized responses to requests for existing SNMP objects
  - Users can also use this event detector to simulate responses to requests for missing SNMP objects
- Outgoing Trap Interception Enhancement for SNMP Notification Event Detector
  - Detect outgoing SNMP traps and generate corresponding event from SNMP Notification event detector.
  - With this enhancement, users will be able to register EEM policies to listen to both incoming and outgoing SNMP traps and provide customizable logic to handle them.
- EEM Policy Description Enhancement
  - Both TCL-based and Applet-based EEM policies are enhanced to support a one-line description provided by user
  - “show event manager policy” command is enhanced to display the description for EEM policies
- Logging Action Enhancement
  - The logging actions for both TCL-based and CLI-based EEM policies are enhanced to support user defined facility so that EEM policies can generate Syslog messages with their own unique facility names. For example, some diagnostic scripts will generate their own Syslog messages that get mingled with the original error messages. The ability to specify its own facility will allow user to easily separate the original error messages from the additional diagnostic information.
- EEM Policy AAA Bypass
  - A new configuration command “event manager applet <applet-name> authorization <authorization type> ” has been added in this release. It allows user to bypass AAA checking for commands executed as part of an EEM policy
- EEM TCL Library Enhancement
  - Two new TCL procedures are added in this release to allow the execution of multiple CLI commands in one function call with a combined result.

Hardware

Routers	• Cisco Integrated Services Routers • Cisco 7200 Series

Additional Information

For more information about Cisco IOS Embedded Event Manager, visit http://cisco.com/go/eem, or contact your local account representative or askabouteem@cisco.com for additional questions.

9. IOS Security

9.1 Cisco IOS Intrusion Prevention System (IPS) Enhancements

Lightweight IPS Engines for existing and new signatures optimized for HTTP, SMTP and FTP protocols

- Memory efficient traffic scanning for attack signatures consuming up to 40% less memory on the router.

New Default IOS IPS Category signatures, with frequent updates by Cisco Signature Team

- More comprehensive and effective attack coverage by default.
- Much quicker inclusion of most relevant new threat signatures within the default set (category).

Chaining of Traffic Scanning (Regular Expression) Tables

- Capability to load more signatures simultaneously and provide protection for larger number of threats and vulnerabilities (particularly advantageous with the larger default memory sizes available on Integrated Services Router Generation Two).

Configurable Threshold (Upper Limit) to be dedicated to IPS feature

- Avoid unintended allocation of large amount of router memory by IPS signature tables
- Prevents performance and other operational problems

Additional Information

For more information visit http://www.cisco.com/go/ips.

9.2 Additional Cisco IOS Security Feature Enhancements in Cisco IOS Software Release 15.0(1)M

GET VPN VRF-Aware GDOI on GM

- Adding to the existing “GET VPN VRF-Lite” support that provides data plane traffic separation, this provides control plane traffic separation (group registration and rekeys) for a dedicated management VRF so different SP customers or Enterprise departments can share the same key servers or alternatively have separate key servers.

Ability to Disable Volume-based IPSec Lifetime Rekey

- Provides options to disable IPSec volume-based lifetimes that may waste system resources in high throughput environments, or increase the upper limit value for IPSec volume lifetime.

DMVPN Enhancements

- Enhancements to management and reporting on DMVPN tunnels such as tunnel health and recovery, interface status, syslog, and MIB/

Additional Information

For more information visit http://www.cisco.com/go/security.

10. Voice

10.1 Packet Voice, Video DSP Module-3

The Cisco® High-Density Packet Voice, Video Digital Signal Processor Module (PVDM3) enables Cisco 2900 and 3900 Series Integrated Services Routers Generation 2 to provide rich-media capabilities such as high-density voice video connectivity, conferencing, transcoding, transrating, and secure voice in Cisco Unified Communications Solutions.

The next-generation high-density packet voice digital signal processor (DSP) modules are available in six densities: PVDM3-16, PVDM3-32, PVDM3-64, PVDM3-128, PVDM3-192, and PVDM3-256, with 16, 32, 64, 128, 192, and 256 channels, respectively.

These products provide higher scale and investment protection for customers that need voice services today and video services in the future.

Additional Information

For more information please visit http://www.cisco.com/go/voice.

10.2 Transcoding and Codec Enhancements

Cisco IOS Release 15.0(1)M adds a number of features which address transcoding and codec treatment for the voice gateway, UCME, CUBE and SRST.

Cisco Unified Border Element enhancements include codec transrating, IP leg call level adjustment and IP-to-IP gateway voice call quality monitoring.

This release also adds G.722 support for SRST environments.

Additional Information

For more information please visit http://www.cisco.com/go/voice.

Product Management contact

Li Shen (lishen@cisco.com)

10.3 Cisco Unified Border Element (CUBE) Support for SRTP-RTP Internetworking

Critical to securing Unified Communications deployments, secure-RTP allows customers to deploy security and interwork with CUCM. This feature allows customers to create a secure voice extranet over any link, including Internet, to lower cost and add flexibility to deployment options.

Additional Information

For more information visit http://www.cisco.com/go/uc and http://www.cisco.com/go/voice.

Product Management contact

Darryl Sladden (dsladden@cisco.com)

10.4 Cisco Unified Border Element (CUBE) Support for Out-of-dialog SIP OPTIONS Ping Messages to Monitor SIP Servers

This feature allows for faster determination of when a SIP trunk that is used for PSTN access is down. It lowers post dial delay in failover.

This feature can be used to enforce SP SLA for SIP trunks for PSTN access, and also allow for the monitoring and management of SIP trunk capacity and status.

The result is lower customer operations and support costs.

Additional Information

For more information visit http://www.cisco.com/go/uc and http://www.cisco.com/go/voice.

Product Management contact

Darryl Sladden (dsladden@cisco.com)

10.5 UC Trusted Firewall Control Version 2

The UC trusted firewall control enhancements offers expanded support through Unified Communications Manager Express, Cisco Unified Border Element and Unified Communications Manager to an external Firewall to indicate on which UDP ports to expect the media.

Additional Information

For more information visit http://www.cisco.com/go/voice.

Service Advertisement Framework (SAF) Support for UC Manager Express, Cisco Unified Border Element and Voice Gateways

This feature adds support for SAF, which is designed to allow networking applications to discover the existence, location, and configuration of networked services within networks.

It supports SAF advertising and listening on voice gateways, Cisco Unified CallManager Express (CME), and Cisco Unified Survivable Remote Site Telephony (SRST).

It provides interoperation for dynamic discovery between CME, SRST, gateways, and Cisco Unified CallManager (CCM) for dynamic discovery.

Product Management contact

David Sauerhaft (dsauerha@cisco.com)

Additional Information

For more information visit http://www.cisco.com/go/uc.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C25-561938-00 10/09