IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

NETSOCKET, INC., Plaintiff, v. **CISCO SYSTEMS, INC.,** Defendant.	**JURY TRIAL DEMANDED** Case No. 2:22-cv-00172-JRG

SECOND AMENDED COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff NetSocket, Inc. ("NetSocket") for its Second Amended Complaint against Defendant Cisco Systems, Inc. ("Cisco"), alleges as follows:

NATURE OF THE ACTION

1. This is an action brought by NetSocket for infringement of U.S. Patent No. 7,616,601 (the "'601 Patent"), U.S. Patent No. 7,190,698 (the "'698 Patent"), U.S. Patent No. 7,734,796 (the "'796 Patent"), U.S. Patent No. 7,827,284 (the "'284 Patent"), U.S. Patent No. 7,720,966 (the "'966 Patent"), U.S. Patent No. 7,606,885 (the "'885 Patent"), and U.S. Patent No. 7,885,286 (the "'286 Patent") (collectively, the "Asserted Patents"), arising under the patent laws of the United States, 35 U.S.C. §§ 271 and 281.

PARTIES

2. NetSocket is a corporation organized and existing under the laws of the State of Delaware with its principal place of business at 7208 Chardonnay Drive, Frisco, Texas 75035.

3. NetSocket was founded in 2006 by three individuals who formerly worked at Chiaro Networks. Chiaro Networks developed certain routing software that it licensed to businesses. NetSocket initially licensed, but later purchased, the Chiaro routing software. The

product and business strategy was to leverage the substantial investment at Chiaro in its routing software and repurpose it as a service assurance provider for real-time services, e.g., voice, video, and data, running over internet protocol (IP) networks.

4. NetSocket's strategy was to passively peer with service provider routers and correlate real-time sessions with the path traveled in the IP network. By doing so, NetSocket could see problematic data links with congestion, packet drop, high latency, etc., making it easier to troubleshoot call or session quality. Such a capability did not exist previously and attracted attention from the likes of Microsoft, who became an early partner with NetSocket. See below:

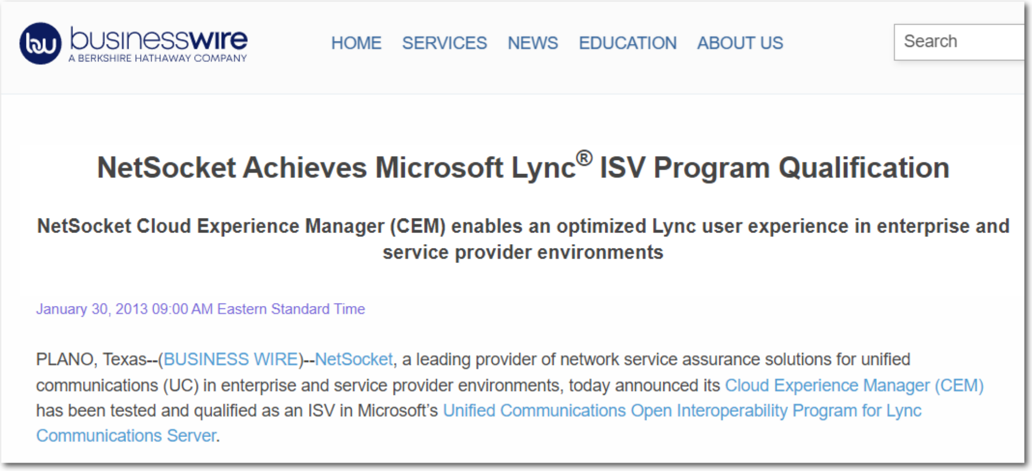

NetSocket Achieves Microsoft Lync® ISV Program Qualification

NetSocket Cloud Experience Manager (CEM) enables an optimized Lync user experience in enterprise and service provider environments

January 30, 2013 09:00 AM Eastern Standard Time

PLANO, Texas--(BUSINESS WIRE)--NetSocket, a leading provider of network service assurance solutions for unified communications (UC) in enterprise and service provider environments, today announced its Cloud Experience Manager (CEM) has been tested and qualified as an ISV in Microsoft's Unified Communications Open Interoperability Program for Lync Communications Server.

https://www.businesswire.com/news/home/20130130005800/en/NetSocket-Achieves-Microsoft-Lync%C2%AE-ISV-Program-Qualification

5. Approximately two years later, in 2008, NetSocket acquired Operax AB, a Swedish company that had accomplished several industry firsts. In particular, Operax was very aggressive in working on partnerships with the likes of Ericsson, Bridgewater Systems, Siemens Nokia Systems, and several others. Operax also conducted trials with many global service providers.

6. After acquiring Operax, NetSocket's go-forward plan was to merge the two teams and technologies, and select the strongest functionality from each product in areas of overlap, in order to have the industry's first and only service assurance and admission control platform so that

service providers could deliver strong service-level agreements (SLA) for emerging real-time services being delivered over IP networks.

7. On information and belief, Cisco is a corporation organized and existing under the laws of Delaware with its principal place of business located at 170 West Tasman Drive, San Jose, California 95134.

JURISDICTION AND VENUE

8. This action arises under 35 U.S.C. §§ 100, *et seq.*, and this Court has jurisdiction over the subject matter of this action under 28 U.S.C. §§ 1331 and 1338(a).

9. This Court has personal jurisdiction over Cisco in this action because Cisco has committed acts within this District giving rise to this action.

10. This Court has personal jurisdiction over Cisco in this action because Cisco has established minimum contacts with this forum such that the exercise of jurisdiction over Cisco would not offend traditional notions of fair play and substantial justice.

11. Cisco, directly and/or through subsidiaries or intermediaries, has conducted business in this District, the State of Texas, and elsewhere in the United States.

12. Cisco, directly and/or through subsidiaries or intermediaries has committed and continues to commit acts of infringement in this District by, among other things, making, using, importing, offering to sell, and selling products and providing services that infringe the Asserted Patents, and/or has induced acts of patent infringement by others in this judicial district, the State of Texas, and elsewhere in the United States.

13. Venue is proper in this Court under 28 U.S.C. §§ 1391 and 1400(b).

14. On information and belief, Cisco has regular and established physical presences in this District, including, but not limited to, ownership of or control over property, inventory, or infrastructure.

15. On information and belief, Cisco maintains several places of business within the State of Texas.

16. On information and belief, Cisco maintains a place of business at 2250 E President George Bush Highway, Richardson, Texas 75082.

17. On information and belief, Cisco maintains a data center at 2260 Chelsea Blvd., Allen, Texas 75013.

18. On information and belief, in 2019 the Collin County Appraisal District assessed the property located at 2250 E President George Bush Highway and 2260 Chelsea Boulevard at a combined value of over $300,000,000.

19. On information and belief, Cisco is registered to do business in the State of Texas.

20. On information and belief, Cisco may be served with process through its registered agent, Corporation Service Company dba CSC – Lawyers Incorporating Service Company, 211 E. 7th Street, Suite 620, Austin, Texas 78701-3218.

THE ASSERTED PATENTS

21. The ’601 Patent, entitled “Network Resource Manager In A Mobile Telecommunication System,” was duly and legally issued by the U.S. Patent and Trademark Office on November 10, 2009. A true and correct copy of the ’601 Patent is attached hereto as Exhibit A.

22. The ’698 Patent, entitled “Network Optimisation [sic] Method,” was duly and legally issued by the U.S. Patent and Trademark Office on March 13, 2007. A true and correct copy of the ’698 Patent is attached hereto as Exhibit B.

23. The ’796 Patent, entitled “Method and Arrangement For Reserving Resources To Obtain A Predetermined Quality Of Service In An IP Network,” was duly and legally issued by the U.S. Patent and Trademark Office on June 8, 2010. A true and correct copy of the ’796 Patent is attached hereto as Exhibit C.

24. The ’284 Patent, entitled “Method And Arrangement In A Communication System,” was duly and legally issued by the U.S. Patent and Trademark Office on November 2, 2010. A true and correct copy of the ’284 Patent is attached hereto as Exhibit D.

25. The ’966 Patent, entitled “Arrangements And Method For Hierarchical Resource Management In A Layered Network Architecture,” was duly and legally issued by the U.S. Patent and Trademark Office on May 18, 2010. A true and correct copy of the ’966 Patent is attached hereto as Exhibit E.

26. The ’885 Patent, entitled “Method For, And A Topology Aware Resource Manager In An IP-Telephony System,” was duly and legally issued by the U.S. Patent and Trademark Office on October 20, 2009. A true and correct copy of the ’284 Patent is attached hereto as Exhibit F.

27. The ’286 Patent, entitled “Method And Arrangement In An IP Network,” was duly and legally issued by the U.S. Patent and Trademark Office on February 8, 2011. A true and correct copy of the ’284 Patent is attached hereto as Exhibit G.

28. NetSocket is the current owner and assignee of the Asserted Patents.

29. The claims of the Asserted Patents are valid and enforceable.

30. As described in more detail below, Cisco infringes the Asserted Patents by making, using, selling, importing, and offering to sell routers running Cisco IOS 15 and above and related networking components, and all like products, collectively “the Accused Products,” in Texas and throughout the United States.

COUNT 1

(Infringement of U.S. Pat. No. 7,616,601)

31. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

32. Cisco has infringed and continues to infringe one or more claims of the ’601 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the ’601 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the ’601 Patent and continues to contribute to the infringement of the ‘601 Patent in violation of 35 U.S.C. §271(c).

33. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 1 of the ’601 Patent at least by making, using, offering to sell, importing, and/or selling end-to-end Quality of Service (QoS) within a mobile telecommunication system.

34. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell a Multiprocessor WAN Application Module (MWAM) and Catalyst 6500 and 7600 Series routing platforms, which when used together provide Quality of Service (QoS) within a mobile telecommunication system (the “Cisco QoS System”) in the United States and encourages distributors to sell, offer to sell, and use, and encourages Cisco’s customers to use, the Cisco QoS System.

35. Cisco has had knowledge of and notice of the ’601 patent and its infringement since before the filing of the original Complaint. NetSocket provided direct notice of Cisco’s infringement of the ’601 Patent to Cisco by letter dated May 23, 2022 and delivered prior to the filing of the original Complaint. Cisco offers to sell the Cisco QoS System in this District and does

so with knowledge that the sale and use of the Cisco QoS System infringes and with the intent for its customers to use the Cisco QoS System in an infringing manner.

36. The Cisco QoS System provides end-to-end Quality of Service (QoS) within a mobile telecommunication system that satisfies each of the limitations of at least claim 1 of the ’601 Patent.

37. For example, the Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, provides end-to-end Quality of Service (QoS) within a mobile telecommunications system such as the one shown in Figure 1-2 below:

Figure 1-2 GPRS/UMTS Network Components with GGSNs Implemented on the Cisco MWAM in the Catalyst 6500 / Cisco 7600 Platform

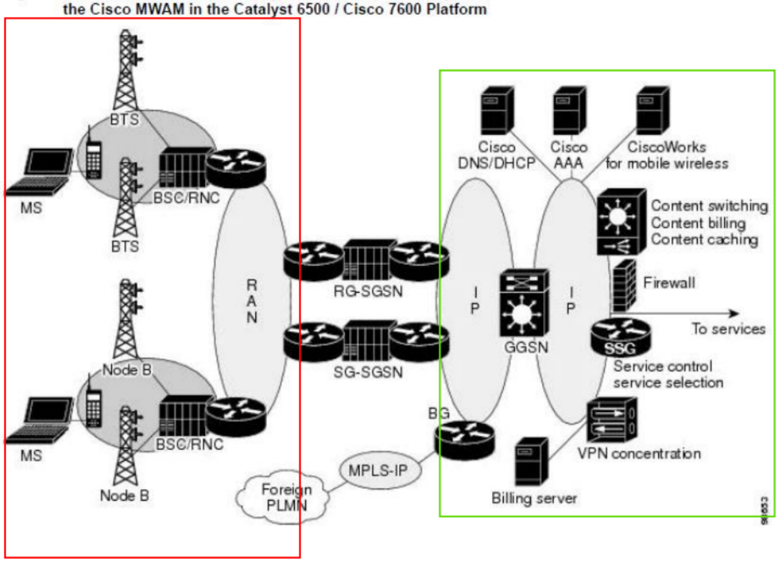

[Overview of GPRS and UMTS - Cisco]

38. As shown above, the mobile telecommunication system in Figure 1-2 comprises a Core Network outlined in green that contains Cisco routers (shown as short cylinders with four arrows) and a Gateway GPRS Support Node (GGSN). The system also includes a Radio Access Network (RAN) outlined in red.

39. The Core Network is connected to the RAN via a Serving GPRS Support Node (SGSN), which uses an IP based transmission:

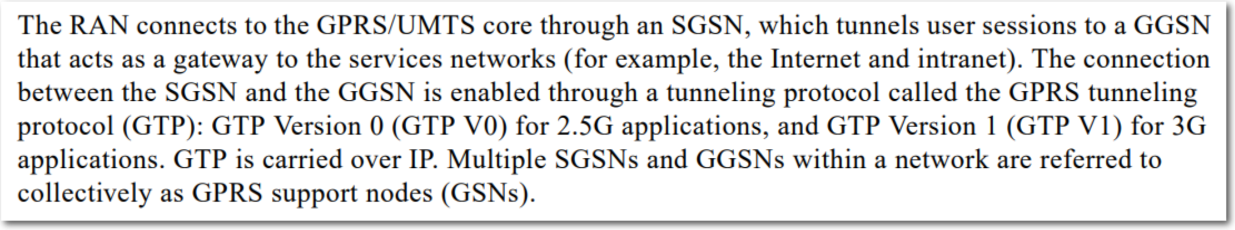

The RAN connects to the GPRS/UMTS core through an SGSN, which tunnels user sessions to a GGSN that acts as a gateway to the services networks (for example, the Internet and intranet). The connection between the SGSN and the GGSN is enabled through a tunneling protocol called the GPRS tunneling protocol (GTP): GTP Version 0 (GTP V0) for 2.5G applications, and GTP Version 1 (GTP V1) for 3G applications. GTP is carried over IP. Multiple SGSNs and GGSNs within a network are referred to collectively as GPRS support nodes (GSNs).

ggsnover.pdf (cisco.com)

40. Cisco implements Quality of Service (QoS) protocols in the Gateway GPRS Support Node (GGSN):

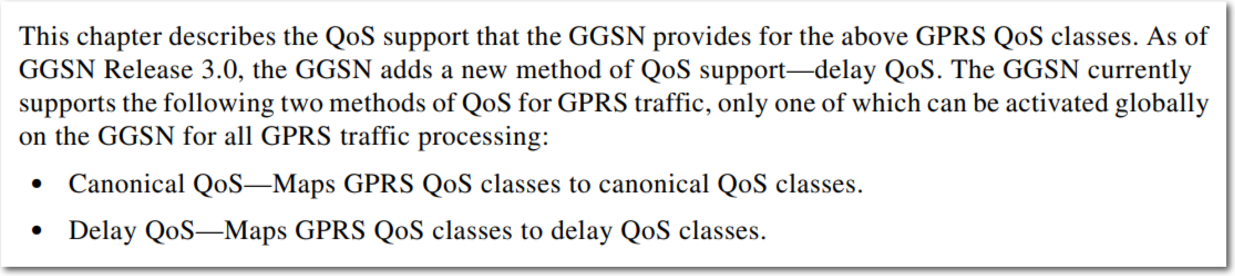

This chapter describes the QoS support that the GGSN provides for the above GPRS QoS classes. As of GGSN Release 3.0, the GGSN adds a new method of QoS support—delay QoS. The GGSN currently supports the following two methods of QoS for GPRS traffic, only one of which can be activated globally on the GGSN for all GPRS traffic processing:

- Canonical QoS—Maps GPRS QoS classes to canonical QoS classes.
- Delay QoS—Maps GPRS QoS classes to delay QoS classes.

ggsnqos.pdf (cisco.com)

41. Cisco implements Quality of Service (QoS) on an end-to-end basis:

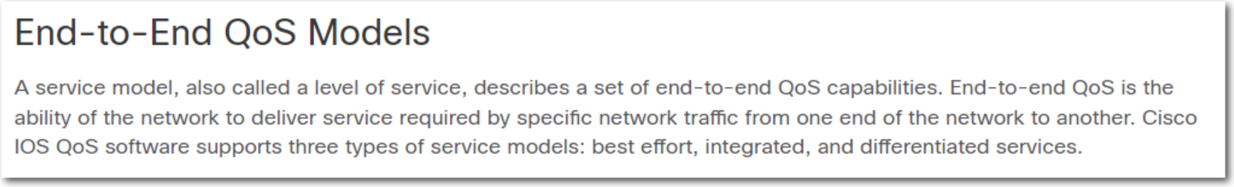

End-to-End QoS Models

A service model, also called a level of service, describes a set of end-to-end QoS capabilities. End-to-end QoS is the ability of the network to deliver service required by specific network traffic from one end of the network to another. Cisco IOS QoS software supports three types of service models: best effort, integrated, and differentiated services.

https://www.cisco.com/c/en/us/td/docs/ios/qos/configuration/guide/12_2sr/qos_12_2sr_book/qos_overview.html#wp1000918

42. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, handles radio resources within the RAN by using a radio resource manager.

43. The network shown in Figure 1-2 below contains two Serving GPRS Support Nodes (SGSN) circled in blue:

Figure 1-2 GPRS/UMTS Network Components with GGSNs Implemented on the Cisco MWAM in the Catalyst 6500 / Cisco 7600 Platform

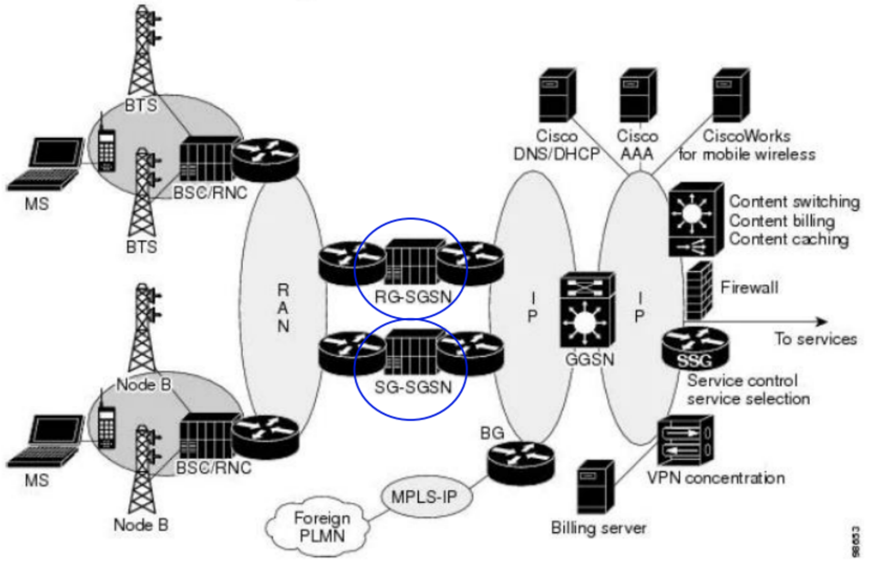

ggsnover.pdf (cisco.com)

44. The Cisco Serving GPRS Support Nodes (SGSN) are responsible for, among other things, RAB (Radio Access Bearer) Assignment Request and RAB Release Request. (*See, e.g.*, Serving GPRS Support Node (SGSN) Overview (cisco.com)).

45. Because the Cisco Serving GPRS Support Nodes (SGSN) are responsible for the Radio Access Bearer (RAB) interface, that functionality is the radio resource manager.

46. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, controls IP network resources by using a resource map in a Network Resource Manager (NRM) in order to provide end-to-end Quality of Service (QoS), where the Network Resource Manager (NRM) performs path-sensitive call admission control by using the resource map in the NRM, the NRM checking that resources are available along a path, and the NRM pre-allocating resources in an IP network.

47. In the network shown in Figure 1-2 below, the Network Resource Manager (NRM) functionality is performed by the Cisco Gateway GPRS Support Node (GGSN) circled in orange:

Figure 1-2 GPRS/UMTS Network Components with GGSNs Implemented on the Cisco MWAM in the Catalyst 6500 / Cisco 7600 Platform

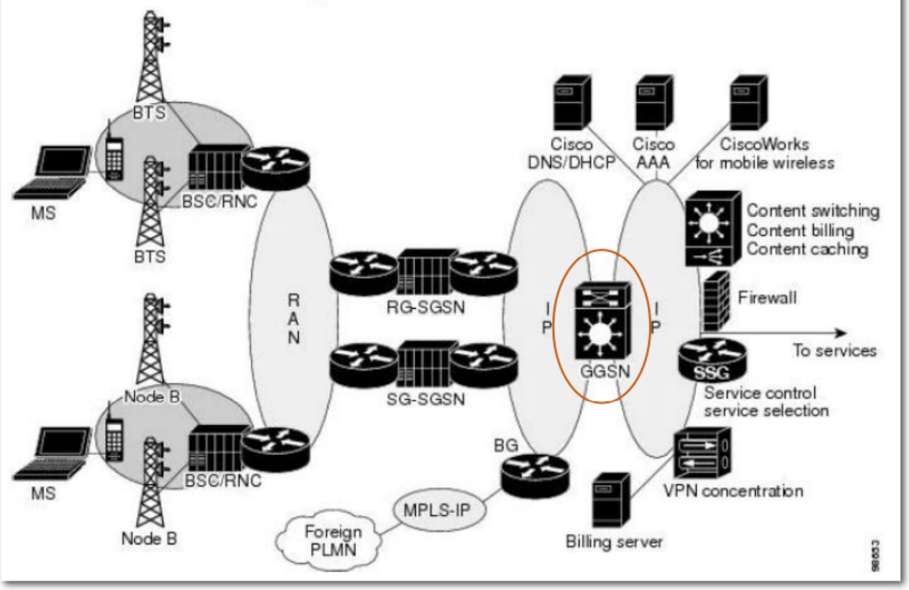

ggsnover.pdf (cisco.com)

48. The Cisco Gateway GPRS Support Node (GGSN) IOS provides Server Load Balancing functionality using Dynamic Feedback Protocol (IOS SLB DFP):

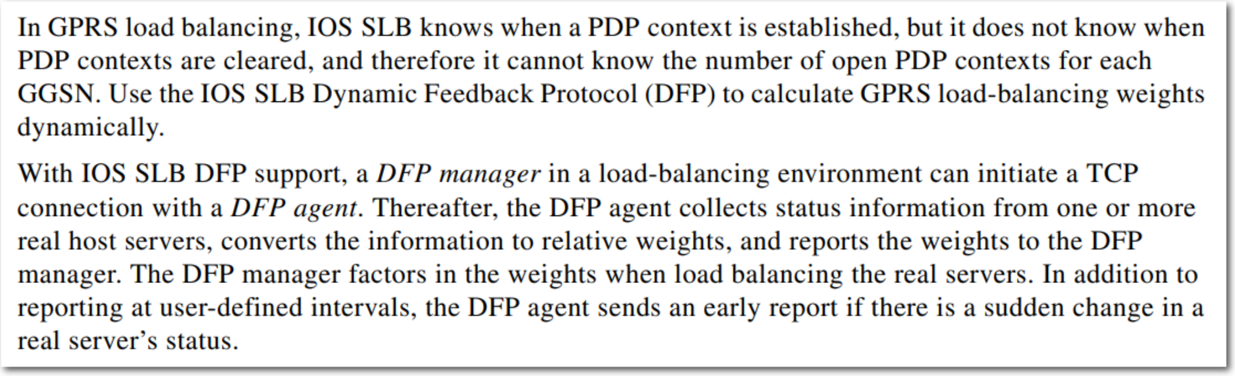

In GPRS load balancing, IOS SLB knows when a PDP context is established, but it does not know when PDP contexts are cleared, and therefore it cannot know the number of open PDP contexts for each GGSN. Use the IOS SLB Dynamic Feedback Protocol (DFP) to calculate GPRS load-balancing weights dynamically.

With IOS SLB DFP support, a *DFP manager* in a load-balancing environment can initiate a TCP connection with a *DFP agent*. Thereafter, the DFP agent collects status information from one or more real host servers, converts the information to relative weights, and reports the weights to the DFP manager. The DFP manager factors in the weights when load balancing the real servers. In addition to reporting at user-defined intervals, the DFP agent sends an early report if there is a sudden change in a real server's status.

ggsnslb.pdf (cisco.com)

49. The Cisco Gateway GPRS Support Node (GGSN) obtains information as part of its role in controlling Quality of Service (QoS):

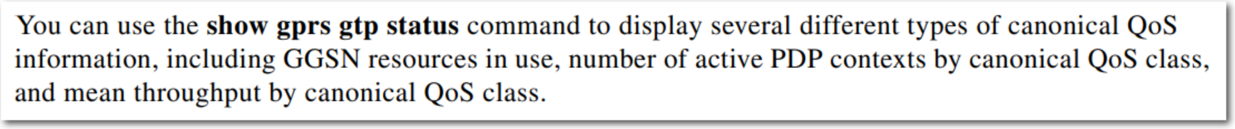

You can use the **show gprs gtp status** command to display several different types of canonical QoS information, including GGSN resources in use, number of active PDP contexts by canonical QoS class, and mean throughput by canonical QoS class.

ggsnqos.pdf (cisco.com)

50. The Cisco Gateway GPRS Support Node (GGSN) also performs path-sensitive call admission control:

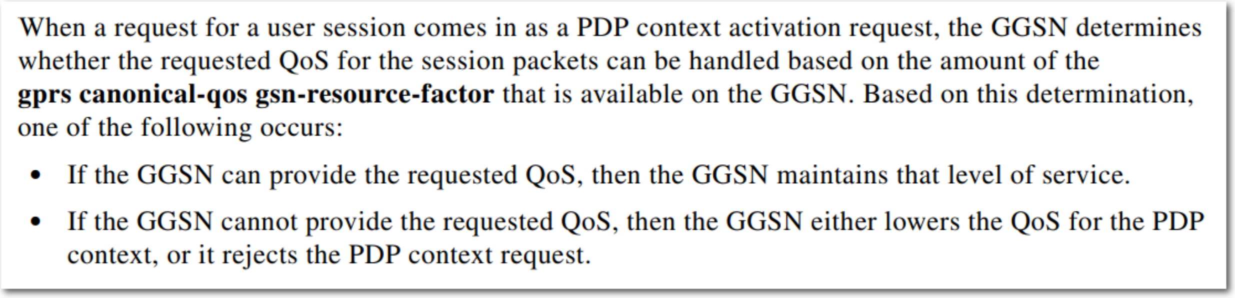

When a request for a user session comes in as a PDP context activation request, the GGSN determines whether the requested QoS for the session packets can be handled based on the amount of the **gprs canonical-qos gsn-resource-factor** that is available on the GGSN. Based on this determination, one of the following occurs:

- If the GGSN can provide the requested QoS, then the GGSN maintains that level of service.
- If the GGSN cannot provide the requested QoS, then the GGSN either lowers the QoS for the PDP context, or it rejects the PDP context request.

ggsnqos.pdf (cisco.com)

51. The Cisco Gateway GPRS Support Node (GGSN) also pre-allocates resources by recognizing different classes of Quality of Service (QoS) functionality and allocating more resources to the higher classes:

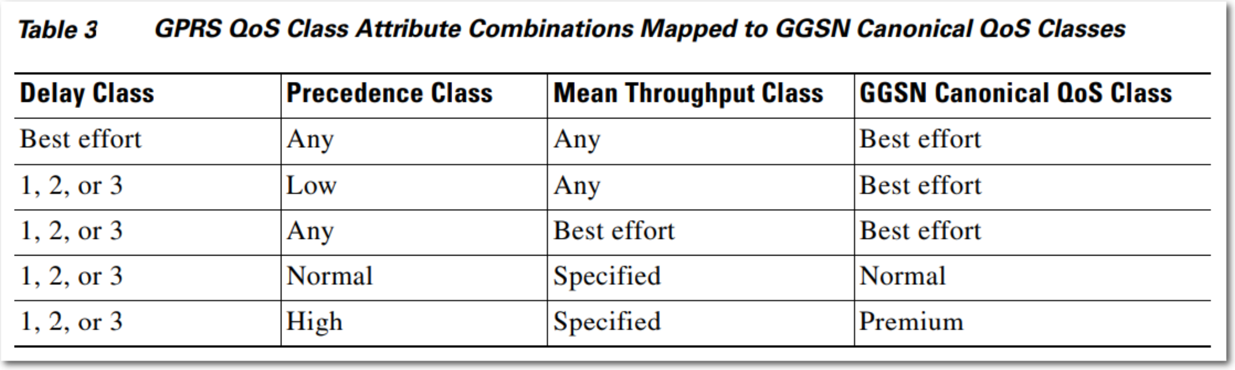

Table 3 GPRS QoS Class Attribute Combinations Mapped to GGSN Canonical QoS Classes

Delay Class	**Precedence Class**	**Mean Throughput Class**	**GGSN Canonical QoS Class**
Best effort	Any	Any	Best effort
1, 2, or 3	Low	Any	Best effort
1, 2, or 3	Any	Best effort	Best effort
1, 2, or 3	Normal	Specified	Normal
1, 2, or 3	High	Specified	Premium

ggsnqos.pdf (cisco.com)

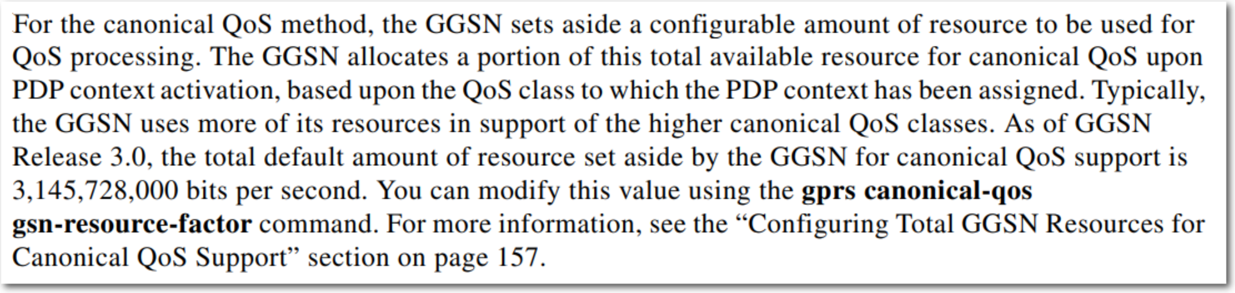

For the canonical QoS method, the GGSN sets aside a configurable amount of resource to be used for QoS processing. The GGSN allocates a portion of this total available resource for canonical QoS upon PDP context activation, based upon the QoS class to which the PDP context has been assigned. Typically, the GGSN uses more of its resources in support of the higher canonical QoS classes. As of GGSN Release 3.0, the total default amount of resource set aside by the GGSN for canonical QoS support is 3,145,728,000 bits per second. You can modify this value using the **gprs canonical-qos gsn-resource-factor** command. For more information, see the "Configuring Total GGSN Resources for Canonical QoS Support" section on page 157.

ggsnqos.pdf (cisco.com)

52. The Cisco Gateway GPRS Support Node (GGSN) is responsible for implementing Quality of Service (QoS) along the entire pathway ("end-to-end"):

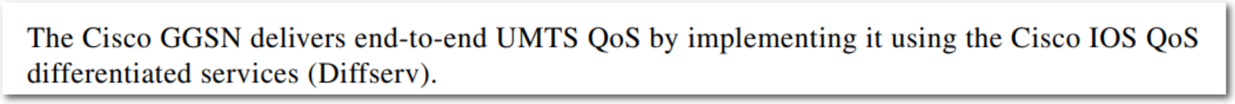

The Cisco GGSN delivers end-to-end UMTS QoS by implementing it using the Cisco IOS QoS differentiated services (Diffserv).

ggsnqos.pdf (cisco.com)

53. Additionally, adaptive or intelligent routing decisions are performed by the Cisco Serving GPRS Support Nodes (SGSN), which uses a process called Evolved ARP (E-ARP) to preempt certain traffic based on priority levels:

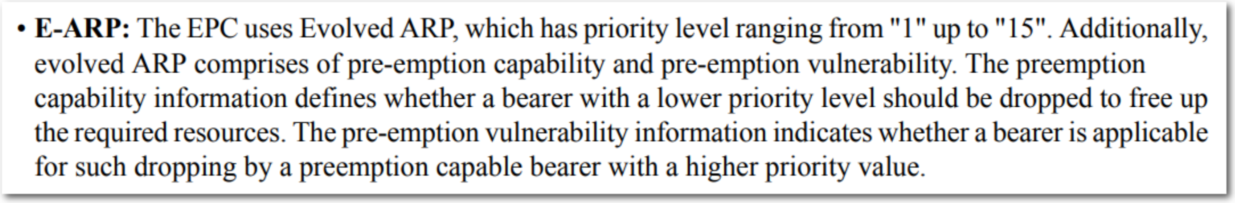

- **E-ARP:** The EPC uses Evolved ARP, which has priority level ranging from "1" up to "15". Additionally, evolved ARP comprises of pre-emption capability and pre-emption vulnerability. The preemption capability information defines whether a bearer with a lower priority level should be dropped to free up the required resources. The pre-emption vulnerability information indicates whether a bearer is applicable for such dropping by a preemption capable bearer with a higher priority value.

Quality of Service (QoS) Management for SGSN (cisco.com)

54. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, schedules resources over time by introducing a start and a stop time as a parameter in a resource request handled by the NRM.

55. Cisco's Quality of Service (QoS) policies can be applied locally through the use of start and stop times:

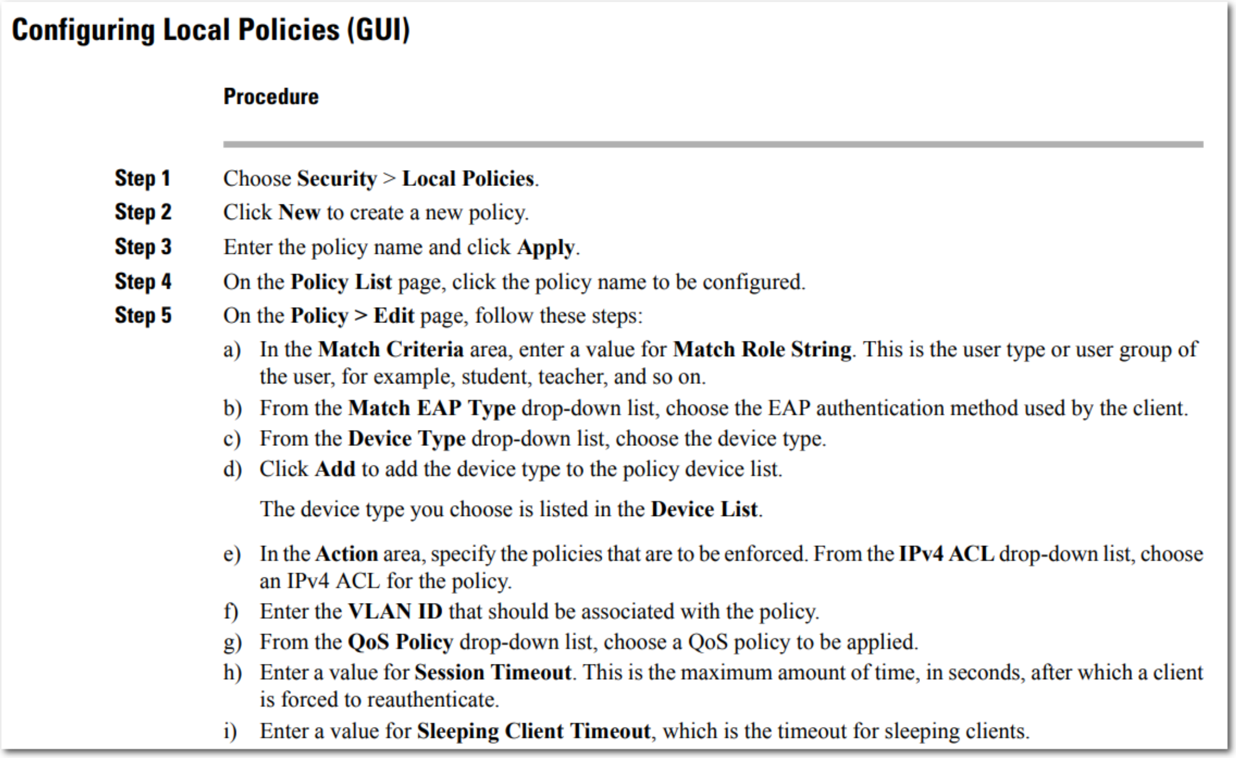

Configuring Local Policies (GUI)

Procedure

Step 1 Choose **Security** > **Local Policies**.

Step 2 Click **New** to create a new policy.

Step 3 Enter the policy name and click **Apply**.

Step 4 On the **Policy List** page, click the policy name to be configured.

Step 5 On the **Policy > Edit** page, follow these steps:

a) In the **Match Criteria** area, enter a value for **Match Role String**. This is the user type or user group of the user, for example, student, teacher, and so on.

b) From the **Match EAP Type** drop-down list, choose the EAP authentication method used by the client.

c) From the **Device Type** drop-down list, choose the device type.

d) Click **Add** to add the device type to the policy device list.

The device type you choose is listed in the **Device List**.

e) In the **Action** area, specify the policies that are to be enforced. From the **IPv4 ACL** drop-down list, choose an IPv4 ACL for the policy.

f) Enter the **VLAN ID** that should be associated with the policy.

g) From the **QoS Policy** drop-down list, choose a QoS policy to be applied.

h) Enter a value for **Session Timeout**. This is the maximum amount of time, in seconds, after which a client is forced to reauthenticate.

i) Enter a value for **Sleeping Client Timeout**, which is the timeout for sleeping clients.

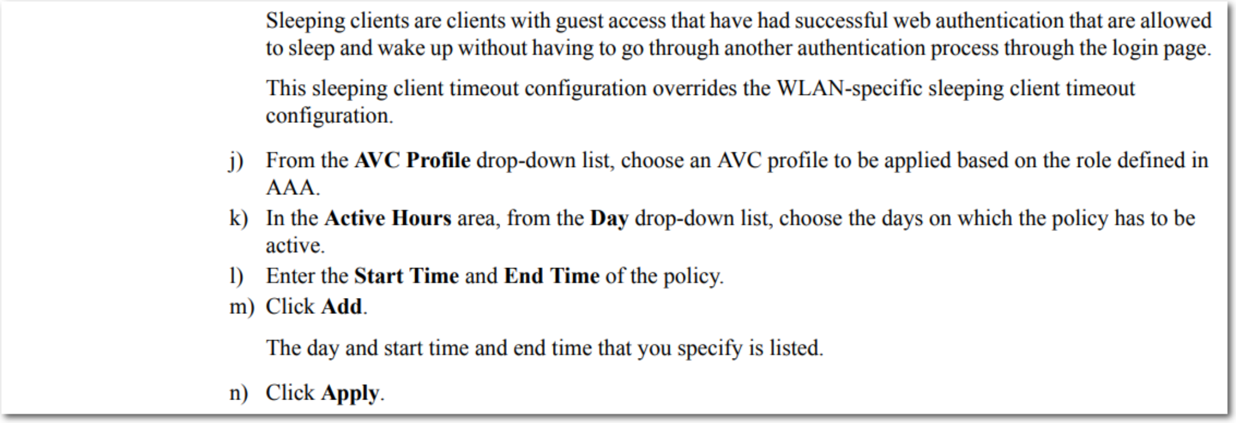

Sleeping clients are clients with guest access that have had successful web authentication that are allowed to sleep and wake up without having to go through another authentication process through the login page.

This sleeping client timeout configuration overrides the WLAN-specific sleeping client timeout configuration.

j) From the **AVC Profile** drop-down list, choose an AVC profile to be applied based on the role defined in AAA.

k) In the **Active Hours** area, from the **Day** drop-down list, choose the days on which the policy has to be active.

l) Enter the **Start Time** and **End Time** of the policy.

m) Click **Add**.

The day and start time and end time that you specify is listed.

n) Click **Apply**.

https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-3/config-guide/b_cg83.pdf

56. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, communicates resource information between the NRM and the radio resource manager.

57. In the network shown below in Figure 1-2, the Cisco Serving GPRS Support Nodes (SGSN), which contains the radio resource manager functionality as described above, also includes

functionality that communicates with the NRM functionality in the Cisco Gateway GPRS Support Node (GGSN):

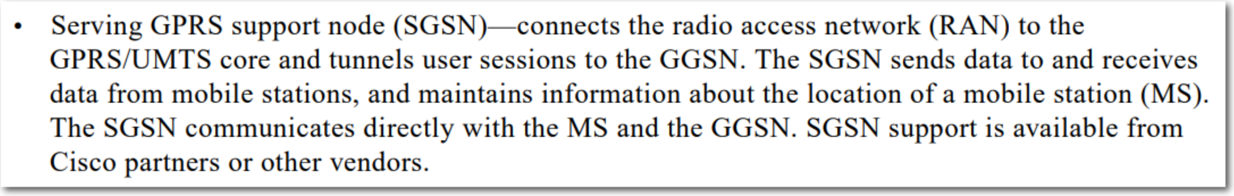

- Serving GPRS support node (SGSN)—connects the radio access network (RAN) to the GPRS/UMTS core and tunnels user sessions to the GGSN. The SGSN sends data to and receives data from mobile stations, and maintains information about the location of a mobile station (MS). The SGSN communicates directly with the MS and the GGSN. SGSN support is available from Cisco partners or other vendors.

ggsnover.pdf (cisco.com)

58. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, reserves the IP network resources along the path by the NRM to fulfill the end-to-end Quality of Service (QoS).

59. The Cisco Gateway GPRS Support Node (GGSN), acting as the NRM, allows the user to reserve resources for certain Quality of Service (QoS) classes:

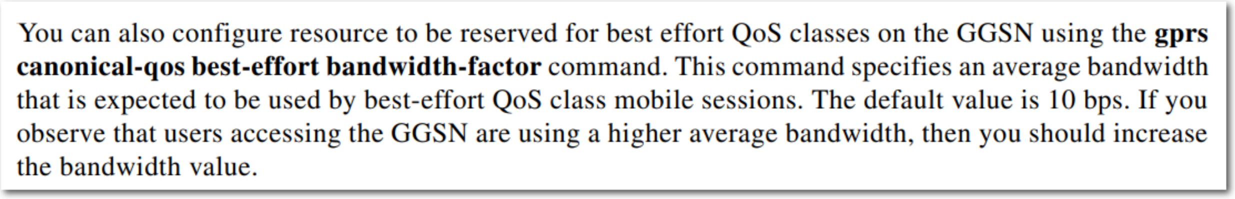

You can also configure resource to be reserved for best effort QoS classes on the GGSN using the **gprs canonical-qos best-effort bandwidth-factor** command. This command specifies an average bandwidth that is expected to be used by best-effort QoS class mobile sessions. The default value is 10 bps. If you observe that users accessing the GGSN are using a higher average bandwidth, then you should increase the bandwidth value.

ggsnqos.pdf (cisco.com)

60. Cisco markets, offers to sell, sells, and distributes the Cisco QoS Systems, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’601 Patent. The Cisco QoS Systems are not staple articles or commodities of commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

61. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the ’601 Patent. Cisco’s infringement of the ’601 Patent has been and continues to be deliberate and willful,

entitling NetSocket to an award of treble damages, reasonable attorney's fees, and costs in bringing this action.

62. Cisco's direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary and permanent injunctive relief and damages as a result of Cisco's infringement of the '601 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 2

(Infringement of U.S. Pat. No. 7,190,698)

63. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

64. Cisco has infringed and continues to infringe one or more claims of the '698 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the '698 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the '698 Patent and continues to contribute to the infringement of the '698 Patent in violation of 35 U.S.C. §271(c).

65. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 1 of the '698 Patent at least by making, using, offering to sell, importing, and/or selling Quality of Service (QoS) as part of its Internetwork Operating System (IOS) versions 12 and above.

66. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell the Cisco Internetworking Operating System (IOS) versions 12 and above (the "Cisco IOS") in the Unites States, and encourages distributors to sell, offer to sell and use, and encourages

Cisco's customers to use, the Cisco IOS in the United States with knowledge that the Cisco IOS infringes the '698 Patent.

67. Cisco has had knowledge of and notice of the '698 Patent and its infringement since before the filing of the original Complaint. NetSocket provided direct notice of Cisco's infringement of the '698 Patent to Cisco by letter dated May 23, 2022 and delivered prior to the filing of the original Complaint. Cisco offers to sell the Cisco IOS in this district and does so with the knowledge that the use of the Cisco IOS infringes and with the intent for its distributors and/or customers to use the Cisco IOS in an infringing manner.

68. The Cisco IOS satisfies each of the limitations of at least claim 1 of the '698 Patent.

69. For example, the Cisco IOS generates a request from an entity for a Virtual Leased Line, VLL, having a predefined Quality of Service, QoS, from a source network (SRC) to a destination network (DST) and applies the request to a Bandwidth Broker (BB) of a domain A, BBA, associated to the source network.

70. In Figure 1 below, CE1 (the "SRC") is an ingress router and CE3 (the "DST") is an egress router:

Figure 1 ***MPLS VPN Inter-AS Option AB Topology***

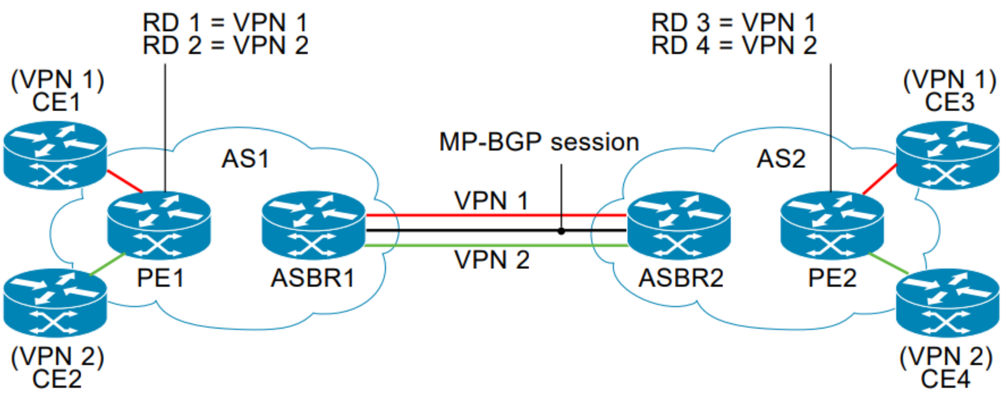

https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book.pdf

71. The Cisco IOS generates a request by having its router request RSVP policy decisions, which reserve network resources, primarily bandwidth (*i.e.*, comprise a Virtual Leased Line using routers that act as Bandwidth Brokers to manage bandwidth resources) and include end-to-end Quality of Service (QoS) guarantees (*i.e.*, a predefined Quality of Service (QoS) from a source network to a destination network), as shown above.

72. The Cisco IOS includes having domain BBA establish the different domains involved to reach the destination network (DST):

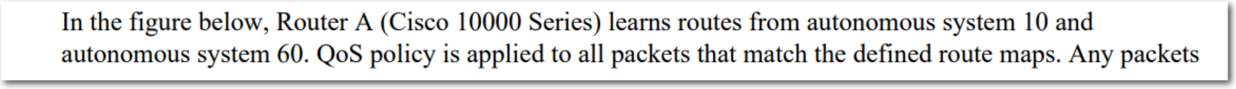

In the figure below, Router A (Cisco 10000 Series) learns routes from autonomous system 10 and autonomous system 60. QoS policy is applied to all packets that match the defined route maps. Any packets

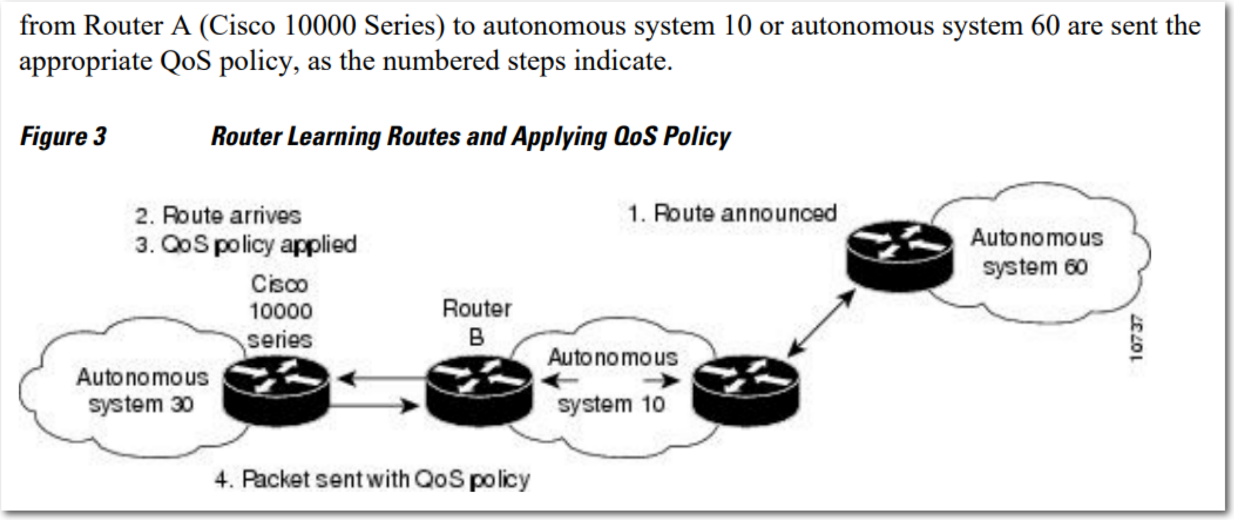

from Router A (Cisco 10000 Series) to autonomous system 10 or autonomous system 60 are sent the appropriate QoS policy, as the numbered steps indicate.

Figure 3 Router Learning Routes and Applying QoS Policy

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-1mt/iri-15-1mt-book.pdf

73. The Cisco IOS includes having domain BBA directly or indirectly pass requests to all Bandwidth Brokers (BB) of the involved domains regarding a Virtual Leased Line (VLL) of the predefined Quality of Service (QoS) from ingress to egress of each domain:

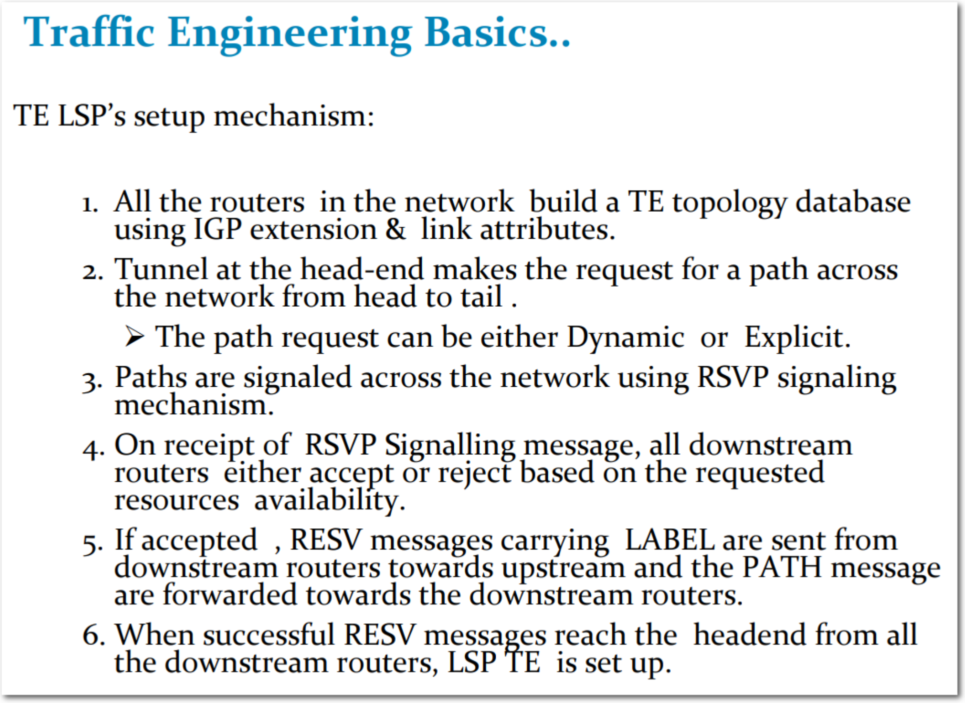

Traffic Engineering Basics..

TE LSP's setup mechanism:

1. All the routers in the network build a TE topology database using IGP extension & link attributes.
2. Tunnel at the head-end makes the request for a path across the network from head to tail .
   - ➢ The path request can be either Dynamic or Explicit.
3. Paths are signaled across the network using RSVP signaling mechanism.
4. On receipt of RSVP Signalling message, all downstream routers either accept or reject based on the requested resources availability.
5. If accepted , RESV messages carrying LABEL are sent from downstream routers towards upstream and the PATH message are forwarded towards the downstream routers.
6. When successful RESV messages reach the headend from all the downstream routers, LSP TE is set up.

https://archive.nanog.org/meetings/nanog49/presentations/Sunday/P2MP.pdf

74. The Cisco IOS includes having each involved bandwidth broker (BB) perform admission control in its domain:

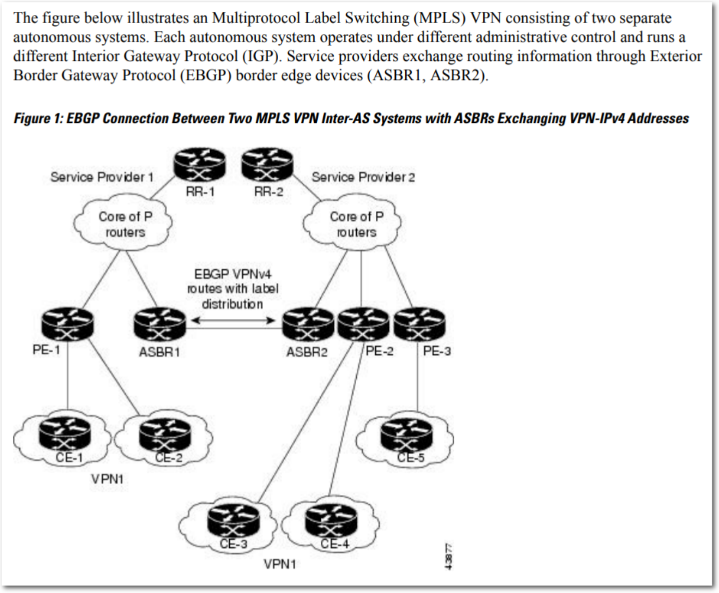

The figure below illustrates an Multiprotocol Label Switching (MPLS) VPN consisting of two separate autonomous systems. Each autonomous system operates under different administrative control and runs a different Interior Gateway Protocol (IGP). Service providers exchange routing information through Exterior Border Gateway Protocol (EBGP) border edge devices (ASBR1, ASBR2).

Figure 1: EBGP Connection Between Two MPLS VPN Inter-AS Systems with ASBRs Exchanging VPN-IPv4 Addresses

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-3s/mp-ias-and-csc-xe-3s-book/mp-vpn-connect-asbr.pdf

75. The Cisco IOS includes having each involved bandwidth broker (BB) return a result of the admission control to BBA that passes it back to the requesting entity and if the request was admitted along all domains between the source (SRC) and the destination (DST) a Virtual Leased Line (VLL) of the predefined Quality of Service (QoS) is granted:

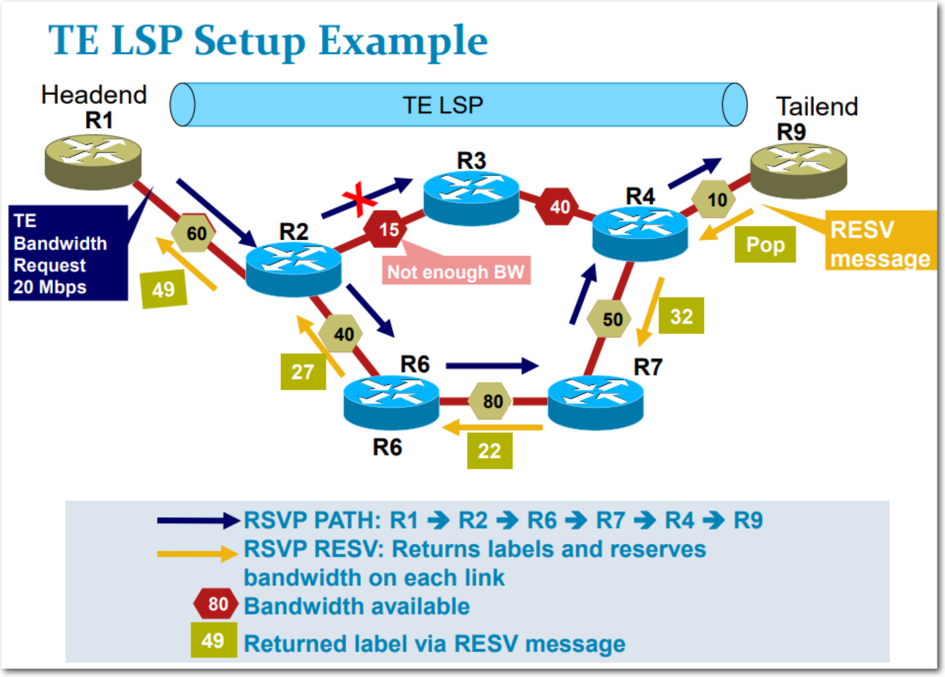

TE LSP Setup Example

RSVP PATH: R1 → R2 → R6 → R7 → R4 → R9
RSVP RESV: Returns labels and reserves bandwidth on each link
80 Bandwidth available
49 Returned label via RESV message

https://archive.nanog.org/meetings/nanog49/presentations/Sunday/P2MP.pdf

76. The Cisco IOS includes performing Label Switched Path (LSP), setup and passing resource requests from the bandwidth broker (BB) to at least one intra-domain (1DB), wherein each intra-domain (1DB) is responsible for admission control and the Label Switched Path (LSP) setup:

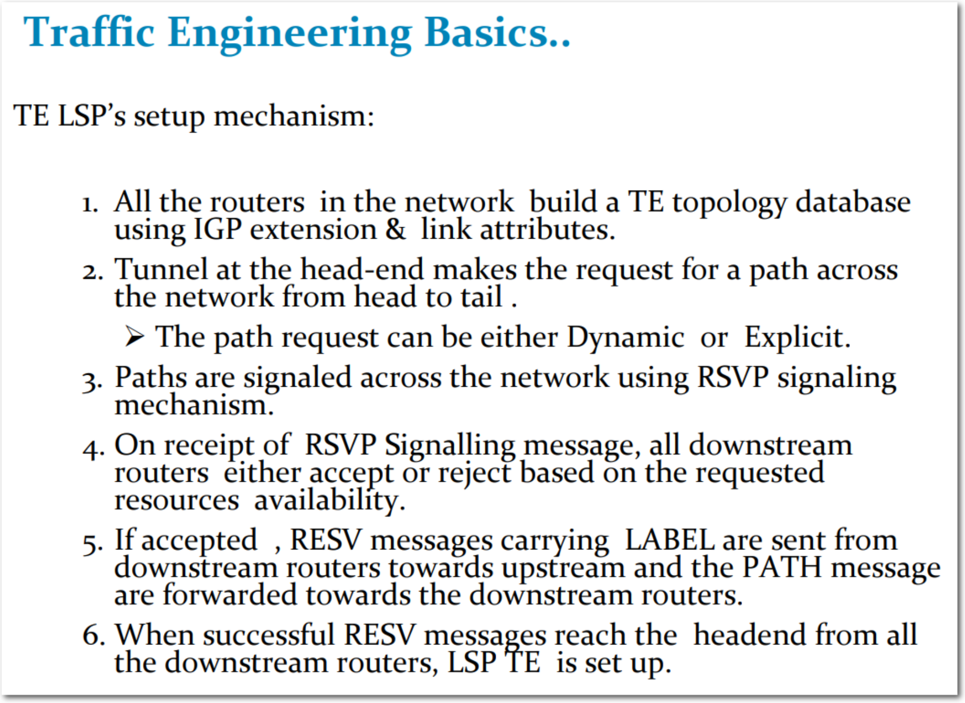

Traffic Engineering Basics..

TE LSP's setup mechanism:

1. All the routers in the network build a TE topology database using IGP extension & link attributes.
2. Tunnel at the head-end makes the request for a path across the network from head to tail .
   - ➢ The path request can be either Dynamic or Explicit.
3. Paths are signaled across the network using RSVP signaling mechanism.
4. On receipt of RSVP Signalling message, all downstream routers either accept or reject based on the requested resources availability.
5. If accepted , RESV messages carrying LABEL are sent from downstream routers towards upstream and the PATH message are forwarded towards the downstream routers.
6. When successful RESV messages reach the headend from all the downstream routers, LSP TE is set up.

https://archive.nanog.org/meetings/nanog49/presentations/Sunday/P2MP.pdf

77. Cisco markets, offers to sell, sells, and distributes the Cisco IOS, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the '698 Patent. The Cisco IOS is not a staple article or commodity of commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

78. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the '698 Patent. Cisco's infringement of the '698 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney's fees, and costs in bringing this action.

79. Cisco's direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary

and permanent injunctive relief and damages as a result of Cisco's infringement of the '698 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 3

(Infringement of U.S. Pat. No. 7,734,796)

80. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

81. Cisco has infringed and continues to infringe one or more claims of the '796 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the '796 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the '796 Patent and continues to contribute to the infringement of the '796 Patent in violation of 35 U.S.C. §271(c).

82. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 25 of the '796 Patent at least by making, using, offering to sell, importing, and/or selling end-to-end Quality of Service (QoS) as part of its Internetwork Operating System (IOS) versions 12 and above.

83. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell the Cisco Internetworking Operating System (IOS) versions 12 and above (the "Cisco IOS") in the Unites States, and encourages distributors to sell, offer to sell and use, and encourages Cisco's customers to use, the Cisco IOS in the United States with knowledge that the Cisco IOS infringes the '796 Patent.

84. Cisco has had knowledge of and notice of the '796 Patent and its infringement since before the filing of the original Complaint. For example, on information and belief, Cisco had knowledge of and notice of the '796 Patent as a result of a partnership with NetSocket since at

least 2012. Cisco offers to sell the Cisco QoS System in this District and does so with knowledge that the sale and use of the Cisco QoS System infringes and with the intent for its customers to use the Cisco QoS System in an infringing manner.

85. The Cisco QoS System satisfies each of the limitations of at least claim 25 of the '796 Patent.

86. For example, the Cisco QoS provides for reserving resources within an IP network to obtain a predetermined Quality of Service (QoS) between a source terminal within a source domain and a destination terminal within a destination domain within a Multi-Protocol Label Switch (MPLS) Virtual Private Network (VPN) Inter-Autonomous System (AS) system such as the one shown in Figure 1 below:

Figure 1. MPLS VPN Inter-AS Option AB Topology

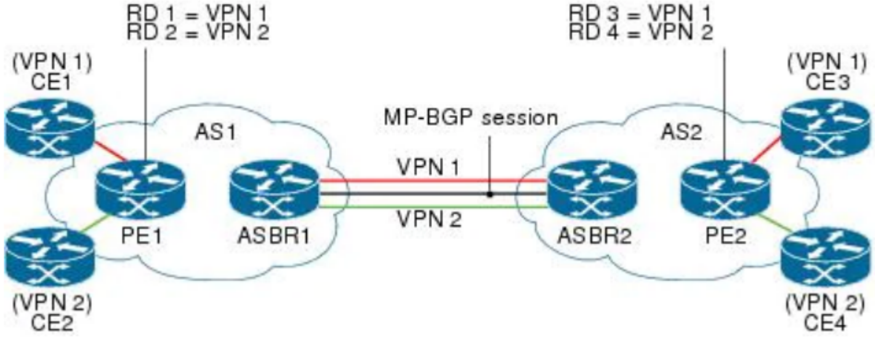

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

87. As shown above, the CE1 is a source terminal within a source domain ("AS1") and CE3 is a destination terminal within a destination domain ("AS2").

88. The MPLS-VPN Inter-AS system provided by Cisco IOS comprises a first network resource manager (NRM) located within the source domain, requesting from a second NRM, located within the destination domain, a resource required for fulfilling the QoS, the resource being

intended for transmission of IP packets, one or more intermediate domains being interposed between the source domain and the destination domain, as shown in Figure 1 below:

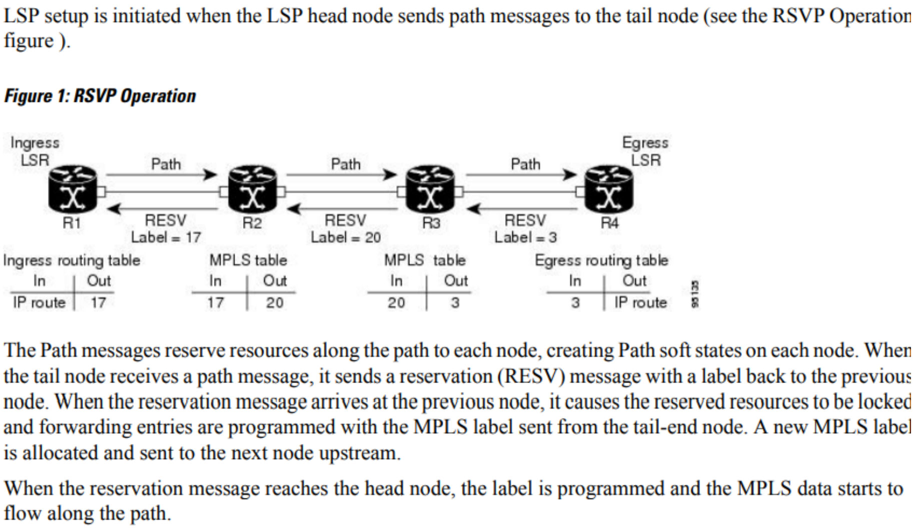

LSP setup is initiated when the LSP head node sends path messages to the tail node (see the RSVP Operation figure).

Figure 1: RSVP Operation

Ingress routing table

In	Out
IP route	17

MPLS table

In	Out
17	20

MPLS table

In	Out
20	3

Egress routing table

In	Out
3	IP route

The Path messages reserve resources along the path to each node, creating Path soft states on each node. When the tail node receives a path message, it sends a reservation (RESV) message with a label back to the previous node. When the reservation message arrives at the previous node, it causes the reserved resources to be locked and forwarding entries are programmed with the MPLS label sent from the tail-end node. A new MPLS label is allocated and sent to the next node upstream.

When the reservation message reaches the head node, the label is programmed and the MPLS data starts to flow along the path.

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

89. The Cisco IOS performs reservation requests utilizing Resource Reservation Protocol, as described below:

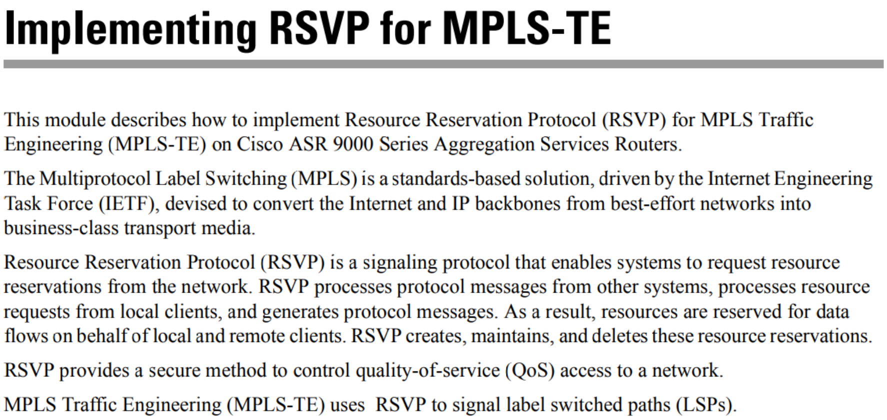

Implementing RSVP for MPLS-TE

This module describes how to implement Resource Reservation Protocol (RSVP) for MPLS Traffic Engineering (MPLS-TE) on Cisco ASR 9000 Series Aggregation Services Routers.

The Multiprotocol Label Switching (MPLS) is a standards-based solution, driven by the Internet Engineering Task Force (IETF), devised to convert the Internet and IP backbones from best-effort networks into business-class transport media.

Resource Reservation Protocol (RSVP) is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations.

RSVP provides a secure method to control quality-of-service (QoS) access to a network.

MPLS Traffic Engineering (MPLS-TE) uses RSVP to signal label switched paths (LSPs).

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

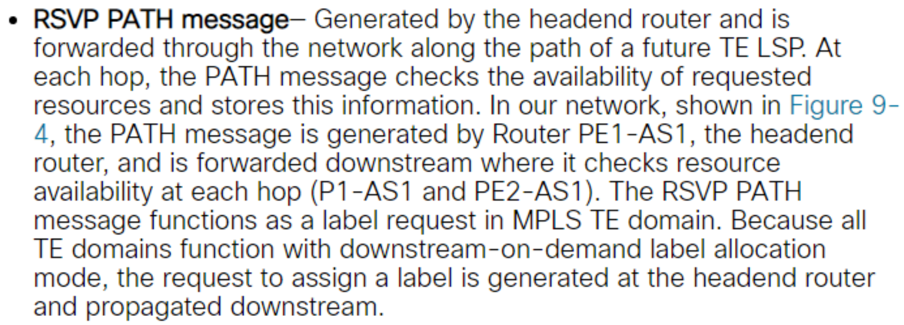

- **RSVP PATH message**– Generated by the headend router and is forwarded through the network along the path of a future TE LSP. At each hop, the PATH message checks the availability of requested resources and stores this information. In our network, shown in Figure 9-4, the PATH message is generated by Router PE1-AS1, the headend router, and is forwarded downstream where it checks resource availability at each hop (P1-AS1 and PE2-AS1). The RSVP PATH message functions as a label request in MPLS TE domain. Because all TE domains function with downstream-on-demand label allocation mode, the request to assign a label is generated at the headend router and propagated downstream.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

90. As shown below, the Cisco IOS MPLS-VPN Inter-AS system may include one or more intermediate domains between a source domain and a destination domain:

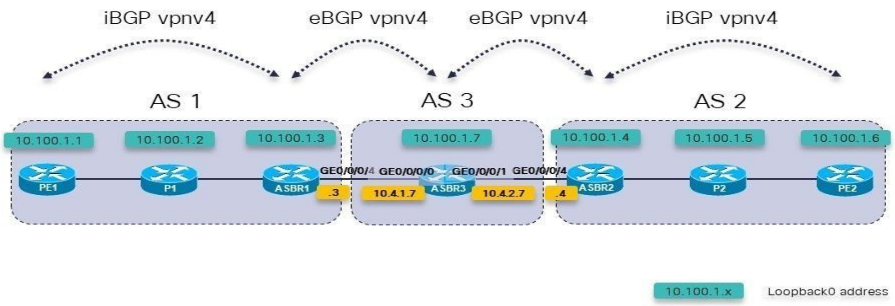

https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-xr-software/217202-cisco-ios-xr-bgp-with-mpls-designs.html

91. As described below, the Cisco MPLS-VPN Inter-AS system includes the second NRM announcing a domain property label of the destination domain to the first NRM, the domain property label characterizing the destination domain:

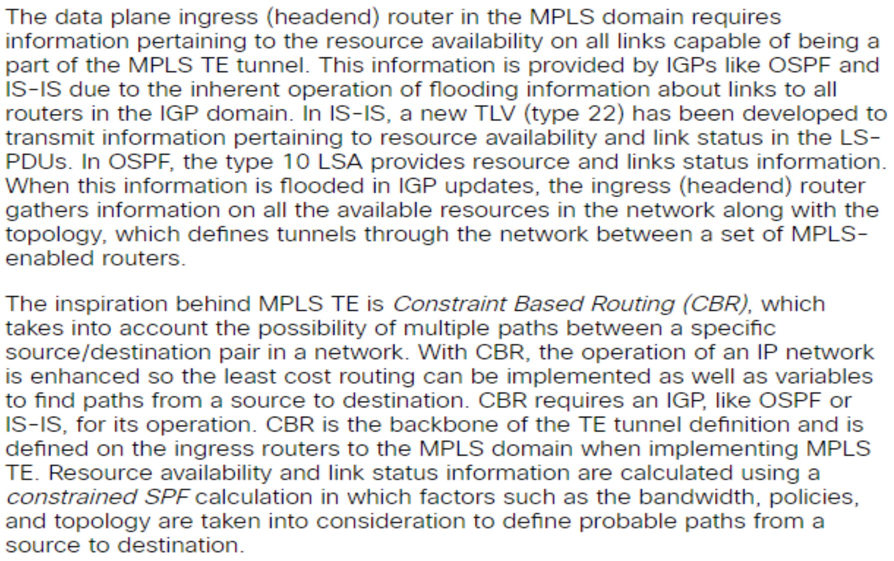

The data plane ingress (headend) router in the MPLS domain requires information pertaining to the resource availability on all links capable of being a part of the MPLS TE tunnel. This information is provided by IGPs like OSPF and IS-IS due to the inherent operation of flooding information about links to all routers in the IGP domain. In IS-IS, a new TLV (type 22) has been developed to transmit information pertaining to resource availability and link status in the LS-PDUs. In OSPF, the type 10 LSA provides resource and links status information. When this information is flooded in IGP updates, the ingress (headend) router gathers information on all the available resources in the network along with the topology, which defines tunnels through the network between a set of MPLS-enabled routers.

The inspiration behind MPLS TE is *Constraint Based Routing (CBR)*, which takes into account the possibility of multiple paths between a specific source/destination pair in a network. With CBR, the operation of an IP network is enhanced so the least cost routing can be implemented as well as variables to find paths from a source to destination. CBR requires an IGP, like OSPF or IS-IS, for its operation. CBR is the backbone of the TE tunnel definition and is defined on the ingress routers to the MPLS domain when implementing MPLS TE. Resource availability and link status information are calculated using a *constrained SPF* calculation in which factors such as the bandwidth, policies, and topology are taken into consideration to define probable paths from a source to destination.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

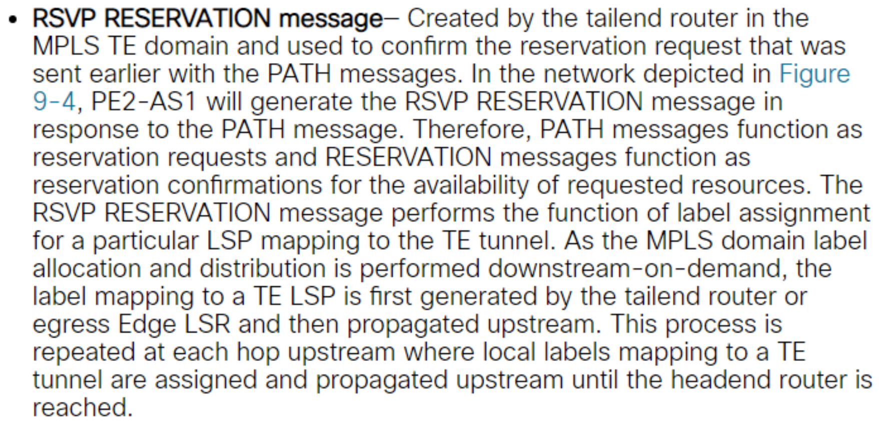

- **RSVP RESERVATION message**– Created by the tailend router in the MPLS TE domain and used to confirm the reservation request that was sent earlier with the PATH messages. In the network depicted in Figure 9-4, PE2-AS1 will generate the RSVP RESERVATION message in response to the PATH message. Therefore, PATH messages function as reservation requests and RESERVATION messages function as reservation confirmations for the availability of requested resources. The RSVP RESERVATION message performs the function of label assignment for a particular LSP mapping to the TE tunnel. As the MPLS domain label allocation and distribution is performed downstream-on-demand, the label mapping to a TE LSP is first generated by the tailend router or egress Edge LSR and then propagated upstream. This process is repeated at each hop upstream where local labels mapping to a TE tunnel are assigned and propagated upstream until the headend router is reached.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

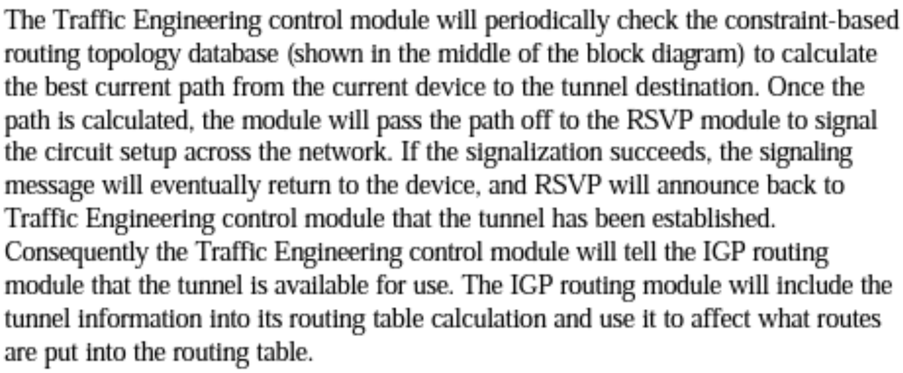

The Traffic Engineering control module will periodically check the constraint-based routing topology database (shown in the middle of the block diagram) to calculate the best current path from the current device to the tunnel destination. Once the path is calculated, the module will pass the path off to the RSVP module to signal the circuit setup across the network. If the signalization succeeds, the signaling message will eventually return to the device, and RSVP will announce back to Traffic Engineering control module that the tunnel has been established. Consequently the Traffic Engineering control module will tell the IGP routing module that the tunnel is available for use. The IGP routing module will include the tunnel information into its routing table calculation and use it to affect what routes are put into the routing table.

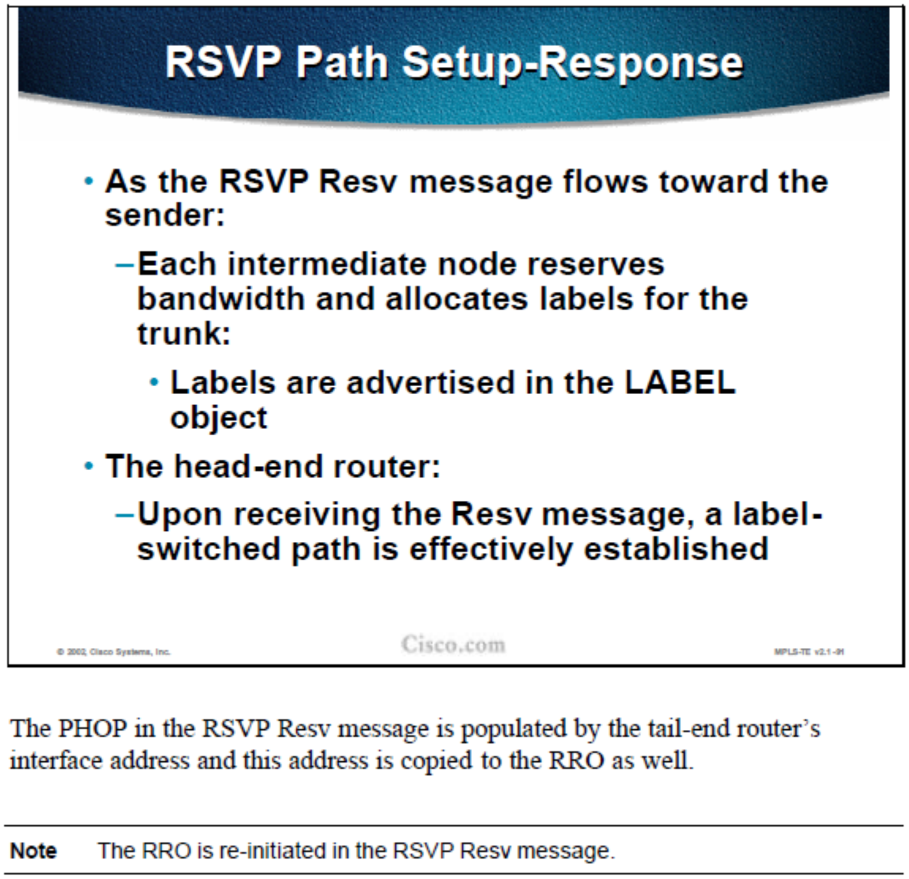

RSVP Path Setup-Response

- As the RSVP Resv message flows toward the sender:
  - –Each intermediate node reserves bandwidth and allocates labels for the trunk:
    - Labels are advertised in the LABEL object
- The head-end router:
  - –Upon receiving the Resv message, a label-switched path is effectively established

© 2002, Cisco Systems, Inc. Cisco.com MPLS-TE v2.1 -91

The PHOP in the RSVP Resv message is populated by the tail-end router's interface address and this address is copied to the RRO as well.

Note The RRO is re-initiated in the RSVP Resv message.

https://www.racf.bnl.gov/Facility/TechnologyMeeting/Archive/06-30-04-CISCO/CISCO-MPLS-TE.pdf

92. As described below, the Cisco MPLS-VPN Inter-AS system includes the first NRM and the second NRM performing one of several actions for transmitting IP packets with QoS between the source terminal and the destination terminal according to the announced domain property label, at least one of the several actions causing the first NRM to perform additional communication with the second NRM to reserve resources in the destination domain:

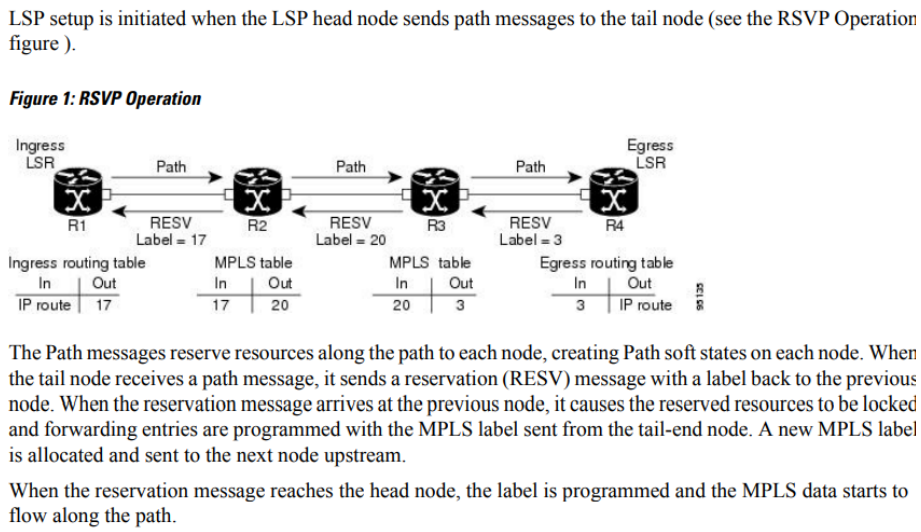

LSP setup is initiated when the LSP head node sends path messages to the tail node (see the RSVP Operation figure).

Figure 1: RSVP Operation

The Path messages reserve resources along the path to each node, creating Path soft states on each node. When the tail node receives a path message, it sends a reservation (RESV) message with a label back to the previous node. When the reservation message arrives at the previous node, it causes the reserved resources to be locked and forwarding entries are programmed with the MPLS label sent from the tail-end node. A new MPLS label is allocated and sent to the next node upstream.

When the reservation message reaches the head node, the label is programmed and the MPLS data starts to flow along the path.

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

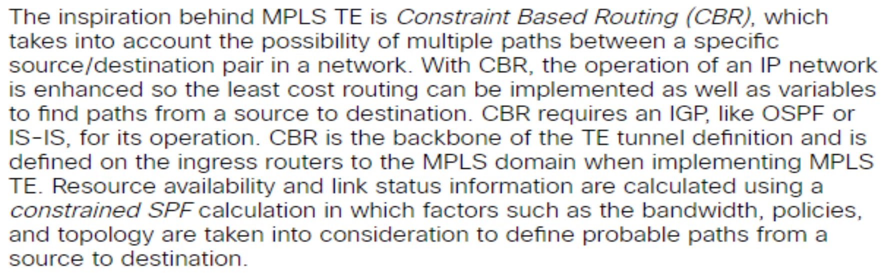

The inspiration behind MPLS TE is *Constraint Based Routing (CBR)*, which takes into account the possibility of multiple paths between a specific source/destination pair in a network. With CBR, the operation of an IP network is enhanced so the least cost routing can be implemented as well as variables to find paths from a source to destination. CBR requires an IGP, like OSPF or IS-IS, for its operation. CBR is the backbone of the TE tunnel definition and is defined on the ingress routers to the MPLS domain when implementing MPLS TE. Resource availability and link status information are calculated using a *constrained SPF* calculation in which factors such as the bandwidth, policies, and topology are taken into consideration to define probable paths from a source to destination.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

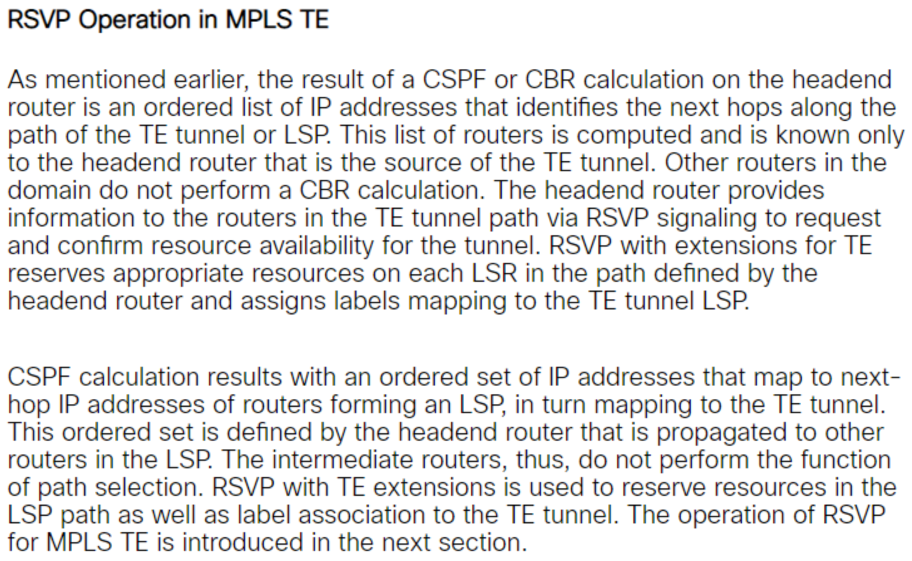

RSVP Operation in MPLS TE

As mentioned earlier, the result of a CSPF or CBR calculation on the headend router is an ordered list of IP addresses that identifies the next hops along the path of the TE tunnel or LSP. This list of routers is computed and is known only to the headend router that is the source of the TE tunnel. Other routers in the domain do not perform a CBR calculation. The headend router provides information to the routers in the TE tunnel path via RSVP signaling to request and confirm resource availability for the tunnel. RSVP with extensions for TE reserves appropriate resources on each LSR in the path defined by the headend router and assigns labels mapping to the TE tunnel LSP.

CSPF calculation results with an ordered set of IP addresses that map to next-hop IP addresses of routers forming an LSP, in turn mapping to the TE tunnel. This ordered set is defined by the headend router that is propagated to other routers in the LSP. The intermediate routers, thus, do not perform the function of path selection. RSVP with TE extensions is used to reserve resources in the LSP path as well as label association to the TE tunnel. The operation of RSVP for MPLS TE is introduced in the next section.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

Figure 1. MPLS VPN Inter-AS Option AB Topology

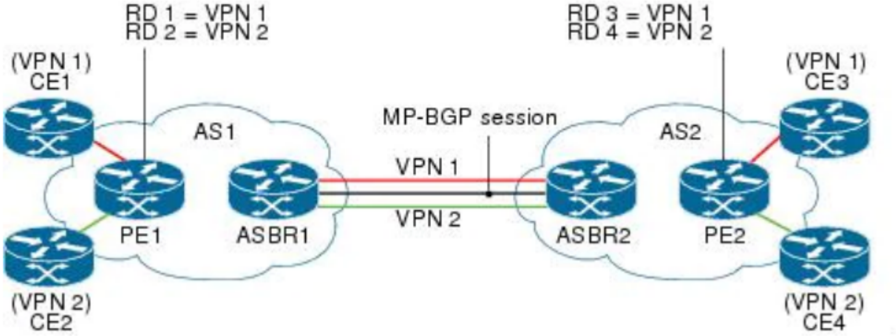

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

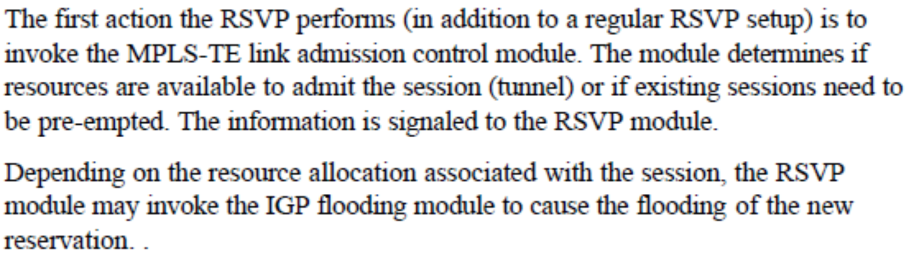
The first action the RSVP performs (in addition to a regular RSVP setup) is to invoke the MPLS-TE link admission control module. The module determines if resources are available to admit the session (tunnel) or if existing sessions need to be pre-empted. The information is signaled to the RSVP module.

Depending on the resource allocation associated with the session, the RSVP module may invoke the IGP flooding module to cause the flooding of the new reservation. .

https://www.racf.bnl.gov/Facility/TechnologyMeeting/Archive/06-30-04-CISCO/CISCO-MPLS-TE.pdf

In the Path and Resv messages of the RSVP there are five objects that are traffic engineering related:

- A Label_Request object is carried in the Path message and requests the label assignment. A request to bind labels to a specific LSP tunnel is initiated by an ingress node through the RSVP Path message.
- A Label object is returned with the Resv message. Labels are allocated downstream and distributed (propagated upstream – from tail-end to the head-end) by means of the RSVP Resv message.
- An Explicit_Route object (ERO) is carried in the Path message to request or suggest a specific route for the traffic tunnel (in the form of a concatenation of hops which constitutes the explicitly routed path). The object is used if the sender node has knowledge of a route that has a high likelihood of meeting the tunnel's QoS requirements, or that makes efficient use of network resources.
- A Record_Route object (RRO) is added to the Path and Resv message to enable the sender node to receive information about the actual route that the LSP tunnel traverses.
- A Session Attribute object can be added to Path messages to aid in session identification and diagnostics. Additional control information, such as setup and hold priorities, resource affinities, and local-protection, are also included in this object.

Record Route object in the Path message is used by the sender to receive information about the actual route that the LSP tunnel traverses. Since the Record Route object is analogous to a path vector, it can be used for loop detection as well.

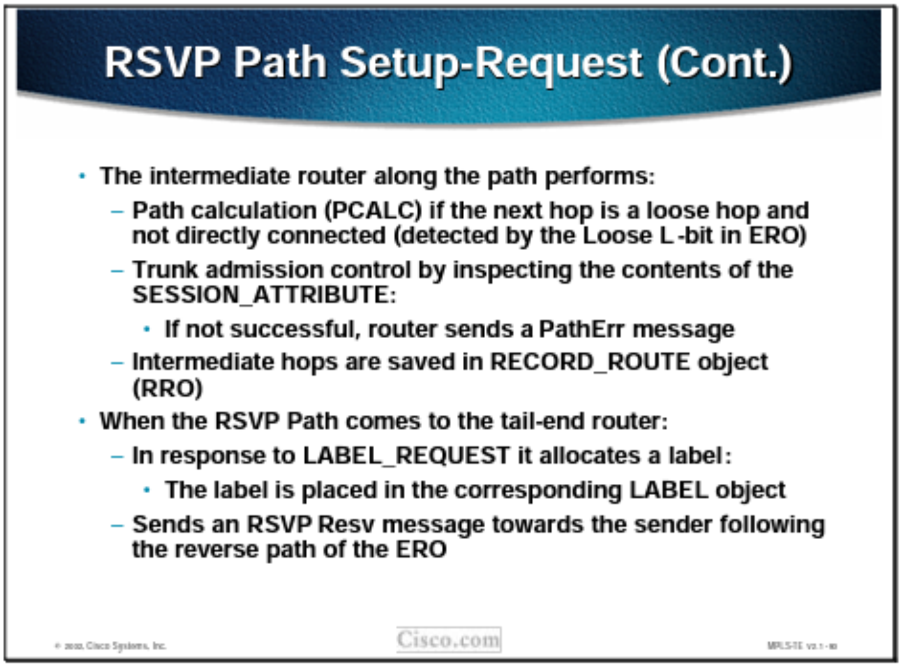

RSVP Path Setup-Request (Cont.)

- The intermediate router along the path performs:
  - Path calculation (PCALC) if the next hop is a loose hop and not directly connected (detected by the Loose L-bit in ERO)
  - Trunk admission control by inspecting the contents of the SESSION_ATTRIBUTE:
    - If not successful, router sends a PathErr message
  - Intermediate hops are saved in RECORD_ROUTE object (RRO)
- When the RSVP Path comes to the tail-end router:
  - In response to LABEL_REQUEST it allocates a label:
    - The label is placed in the corresponding LABEL object
  - Sends an RSVP Resv message towards the sender following the reverse path of the ERO

© 2002, Cisco Systems, Inc. Cisco.com MPLS-TE v2.1-10

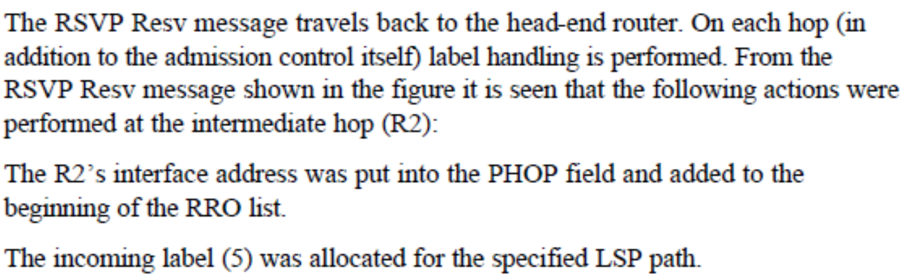

The RSVP Resv message travels back to the head-end router. On each hop (in addition to the admission control itself) label handling is performed. From the RSVP Resv message shown in the figure it is seen that the following actions were performed at the intermediate hop (R2):

The R2's interface address was put into the PHOP field and added to the beginning of the RRO list.

The incoming label (5) was allocated for the specified LSP path.

https://www.racf.bnl.gov/Facility/TechnologyMeeting/Archive/06-30-04-CISCO/CISCO-MPLS-TE.pdf

93. As described below, the Cisco MPLS-VPN Inter-AS system includes using an NRM path vector to identify NRMs along a path from the source terminal to an end-terminal:

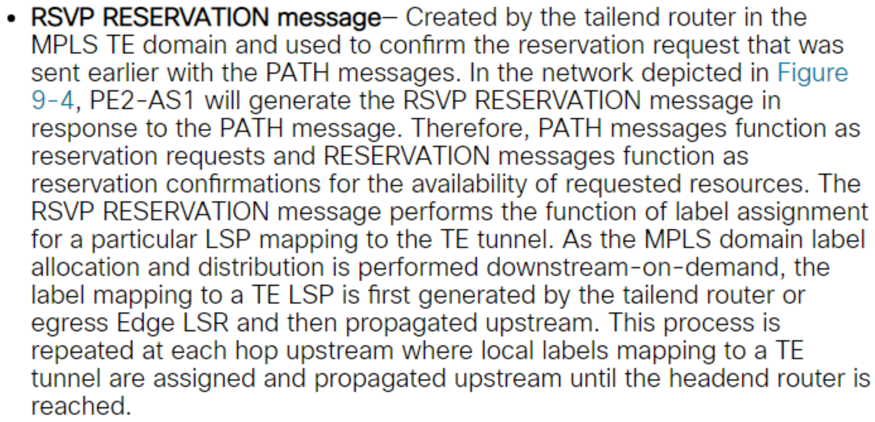

- **RSVP RESERVATION message**– Created by the tailend router in the MPLS TE domain and used to confirm the reservation request that was sent earlier with the PATH messages. In the network depicted in Figure 9-4, PE2-AS1 will generate the RSVP RESERVATION message in response to the PATH message. Therefore, PATH messages function as reservation requests and RESERVATION messages function as reservation confirmations for the availability of requested resources. The RSVP RESERVATION message performs the function of label assignment for a particular LSP mapping to the TE tunnel. As the MPLS domain label allocation and distribution is performed downstream-on-demand, the label mapping to a TE LSP is first generated by the tailend router or egress Edge LSR and then propagated upstream. This process is repeated at each hop upstream where local labels mapping to a TE tunnel are assigned and propagated upstream until the headend router is reached.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

In the Path and Resv messages of the RSVP there are five objects that are traffic engineering related:

- A Label_Request object is carried in the Path message and requests the label assignment. A request to bind labels to a specific LSP tunnel is initiated by an ingress node through the RSVP Path message.
- A Label object is returned with the Resv message. Labels are allocated downstream and distributed (propagated upstream – from tail-end to the head-end) by means of the RSVP Resv message.
- An Explicit_Route object (ERO) is carried in the Path message to request or suggest a specific route for the traffic tunnel (in the form of a concatenation of hops which constitutes the explicitly routed path). The object is used if the sender node has knowledge of a route that has a high likelihood of meeting the tunnel's QoS requirements, or that makes efficient use of network resources.
- A Record_Route object (RRO) is added to the Path and Resv message to enable the sender node to receive information about the actual route that the LSP tunnel traverses.
- A Session Attribute object can be added to Path messages to aid in session identification and diagnostics. Additional control information, such as setup and hold priorities, resource affinities, and local-protection, are also included in this object.

Record Route object in the Path message is used by the sender to receive information about the actual route that the LSP tunnel traverses. Since the Record Route object is analogous to a path vector, it can be used for loop detection as well.

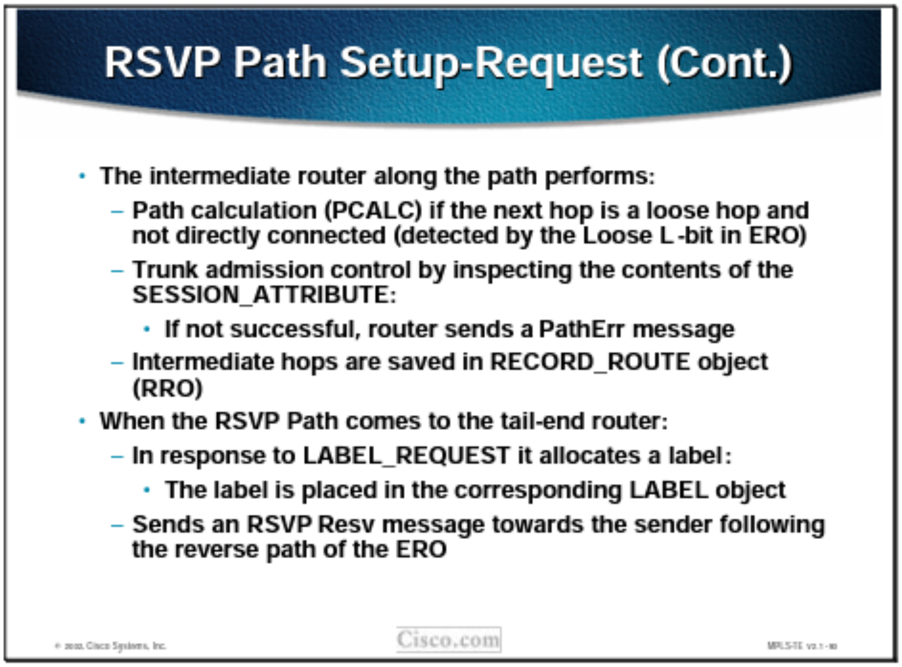

RSVP Path Setup-Request (Cont.)

- The intermediate router along the path performs:
  - Path calculation (PCALC) if the next hop is a loose hop and not directly connected (detected by the Loose L-bit in ERO)
  - Trunk admission control by inspecting the contents of the SESSION_ATTRIBUTE:
    - If not successful, router sends a PathErr message
  - Intermediate hops are saved in RECORD_ROUTE object (RRO)
- When the RSVP Path comes to the tail-end router:
  - In response to LABEL_REQUEST it allocates a label:
    - The label is placed in the corresponding LABEL object
  - Sends an RSVP Resv message towards the sender following the reverse path of the ERO

© 2002, Cisco Systems, Inc. Cisco.com MPLS-TE v2.1-30

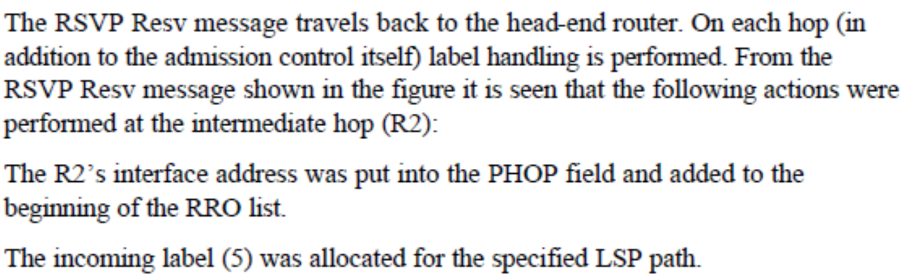

The RSVP Resv message travels back to the head-end router. On each hop (in addition to the admission control itself) label handling is performed. From the RSVP Resv message shown in the figure it is seen that the following actions were performed at the intermediate hop (R2):

The R2's interface address was put into the PHOP field and added to the beginning of the RRO list.

The incoming label (5) was allocated for the specified LSP path.

https://www.racf.bnl.gov/Facility/TechnologyMeeting/Archive/06-30-04-CISCO/CISCO-MPLS-TE.pdf

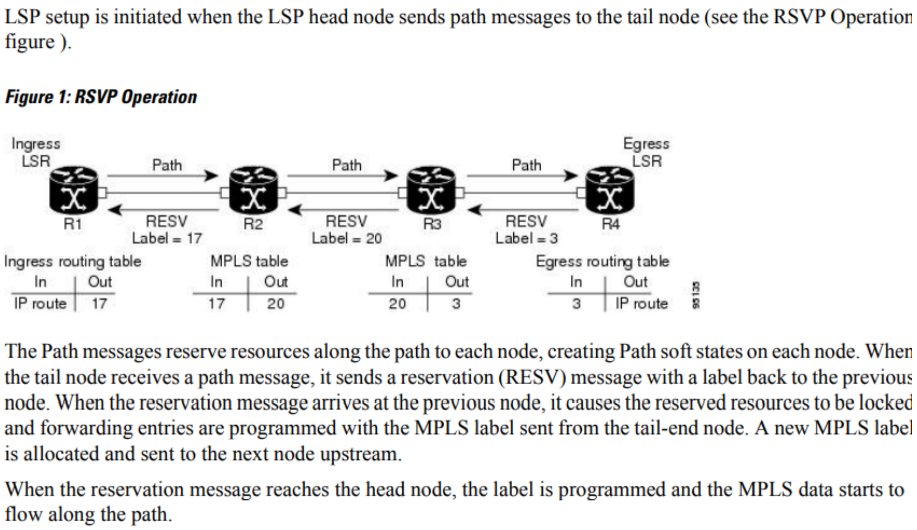

LSP setup is initiated when the LSP head node sends path messages to the tail node (see the RSVP Operation figure).

Figure 1: RSVP Operation

Ingress LSR
Path
Path
Path
Egress LSR
R1
RESV Label = 17
R2
RESV Label = 20
R3
RESV Label = 3
R4
Ingress routing table
In | Out
IP route | 17
MPLS table
In | Out
17 | 20
MPLS table
In | Out
20 | 3
Egress routing table
In | Out
3 | IP route
98135

The Path messages reserve resources along the path to each node, creating Path soft states on each node. When the tail node receives a path message, it sends a reservation (RESV) message with a label back to the previous node. When the reservation message arrives at the previous node, it causes the reserved resources to be locked and forwarding entries are programmed with the MPLS label sent from the tail-end node. A new MPLS label is allocated and sent to the next node upstream.

When the reservation message reaches the head node, the label is programmed and the MPLS data starts to flow along the path.

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

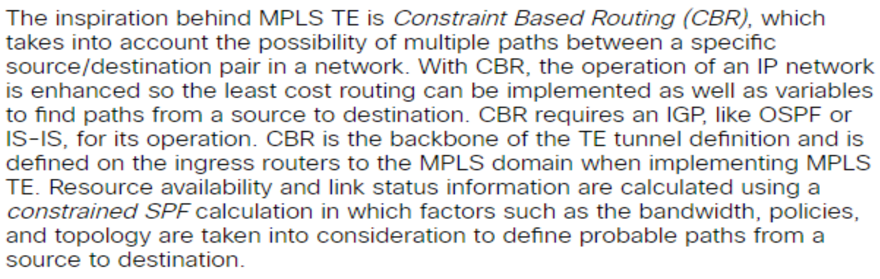

The inspiration behind MPLS TE is *Constraint Based Routing (CBR)*, which takes into account the possibility of multiple paths between a specific source/destination pair in a network. With CBR, the operation of an IP network is enhanced so the least cost routing can be implemented as well as variables to find paths from a source to destination. CBR requires an IGP, like OSPF or IS-IS, for its operation. CBR is the backbone of the TE tunnel definition and is defined on the ingress routers to the MPLS domain when implementing MPLS TE. Resource availability and link status information are calculated using a *constrained SPF* calculation in which factors such as the bandwidth, policies, and topology are taken into consideration to define probable paths from a source to destination.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

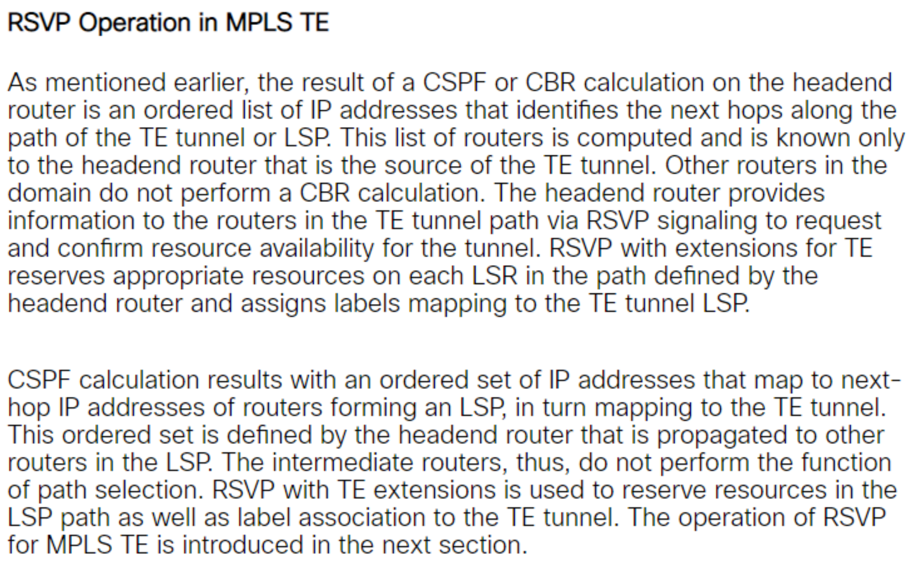

RSVP Operation in MPLS TE

As mentioned earlier, the result of a CSPF or CBR calculation on the headend router is an ordered list of IP addresses that identifies the next hops along the path of the TE tunnel or LSP. This list of routers is computed and is known only to the headend router that is the source of the TE tunnel. Other routers in the domain do not perform a CBR calculation. The headend router provides information to the routers in the TE tunnel path via RSVP signaling to request and confirm resource availability for the tunnel. RSVP with extensions for TE reserves appropriate resources on each LSR in the path defined by the headend router and assigns labels mapping to the TE tunnel LSP.

CSPF calculation results with an ordered set of IP addresses that map to next-hop IP addresses of routers forming an LSP, in turn mapping to the TE tunnel. This ordered set is defined by the headend router that is propagated to other routers in the LSP. The intermediate routers, thus, do not perform the function of path selection. RSVP with TE extensions is used to reserve resources in the LSP path as well as label association to the TE tunnel. The operation of RSVP for MPLS TE is introduced in the next section.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

94. Cisco markets, offers to sell, sells, and distributes the Cisco QoS Systems, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’796 Patent. The Cisco QoS Systems are not staple articles or commodities of commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

95. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the ’796 Patent. Cisco’s infringement of the ’796 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney’s fees, and costs in bringing this action.

96. Cisco’s direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary

and permanent injunctive relief and damages as a result of Cisco's infringement of the '796 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 4

(Infringement of U.S. Pat. No. 7,827,284)

97. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

98. Cisco has infringed and continues to infringe one or more claims of the '284 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the '284 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the '284 Patent and continues to contribute to the infringement of the '284 Patent in violation of 35 U.S.C. §271(c).

99. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 1 of the '284 Patent at least by making, using, offering to sell, importing, and/or selling end-to-end Quality of Service (QoS) within a mobile telecommunication system.

100. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell a Multiprocessor WAN Application Module (MWAM) and Catalyst 6500 and 7600 Series routing platforms, which when used together provide Quality of Service (QoS) within a mobile telecommunication system (the "Cisco QoS System") in the United States and encourages distributors to sell, offer to sell, and use, and encourages Cisco's customers to use, the Cisco QoS System.

101. Cisco has had knowledge of and notice of the '284 Patent and its infringement since before the filing of the original Complaint. Cisco offers to sell the Cisco QoS System in this

District and does so with knowledge that the sale and use of the Cisco QoS System infringes and with the intent for its customers to use the Cisco QoS System in an infringing manner.

102. The Cisco QoS System provides end-to-end Quality of Service (QoS) within a mobile telecommunication system that satisfies each of the limitations of at least claim 1 of the ’284 Patent.

103. For example, the Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, provides for performing admission control in order to offer assurances on forwarding quality in networks such as the one shown in Figure 1-1 below:

Figure 1-1 GPRS/UMTS Network Components with GGSNs Implemented on the Cisco SAMI in the Cisco 7600 Series Router

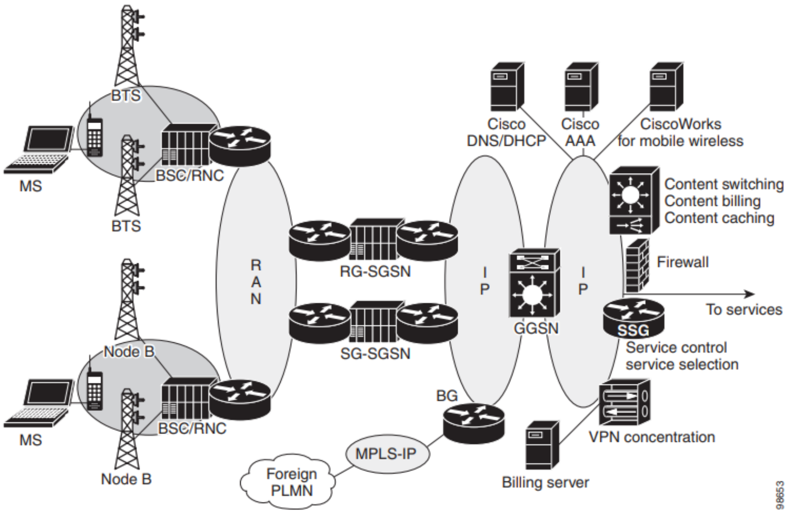

As Figure 1-1 shows, the RAN is made up of different components for 2.5G and 3G.

https://www.cisco.com/c/en/us/td/docs/ios/12_4/12_4y/12_4_24ye3/cfg/12424ye3cfg/ggsnover.pdf

104. As described in greater detail below, the mobile telecommunication system shown above in Figure 1-1 provides for admission control in order to offer assurances on forwarding quality in networks via quality of service (QoS) configured on the MWAM:

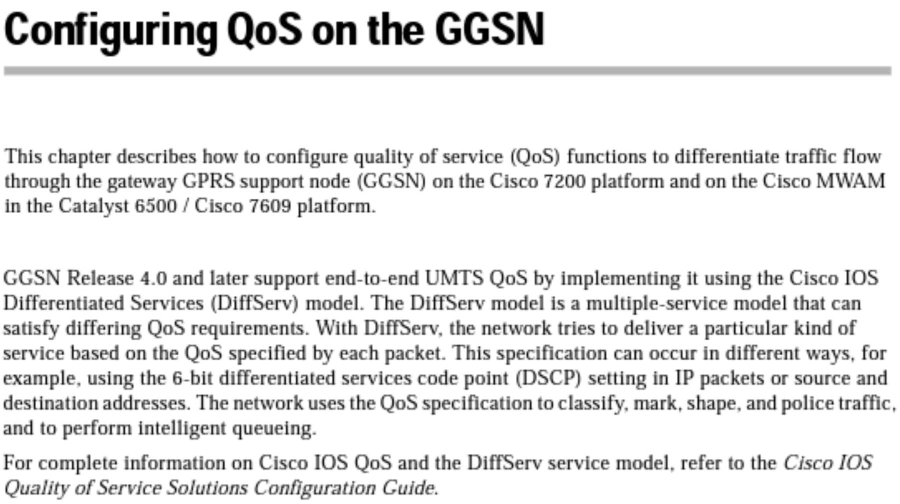

Configuring QoS on the GGSN

This chapter describes how to configure quality of service (QoS) functions to differentiate traffic flow through the gateway GPRS support node (GGSN) on the Cisco 7200 platform and on the Cisco MWAM in the Catalyst 6500 / Cisco 7609 platform.

GGSN Release 4.0 and later support end-to-end UMTS QoS by implementing it using the Cisco IOS Differentiated Services (DiffServ) model. The DiffServ model is a multiple-service model that can satisfy differing QoS requirements. With DiffServ, the network tries to deliver a particular kind of service based on the QoS specified by each packet. This specification can occur in different ways, for example, using the 6-bit differentiated services code point (DSCP) setting in IP packets or source and destination addresses. The network uses the QoS specification to classify, mark, shape, and police traffic, and to perform intelligent queueing.

For complete information on Cisco IOS QoS and the DiffServ service model, refer to the *Cisco IOS Quality of Service Solutions Configuration Guide*.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos

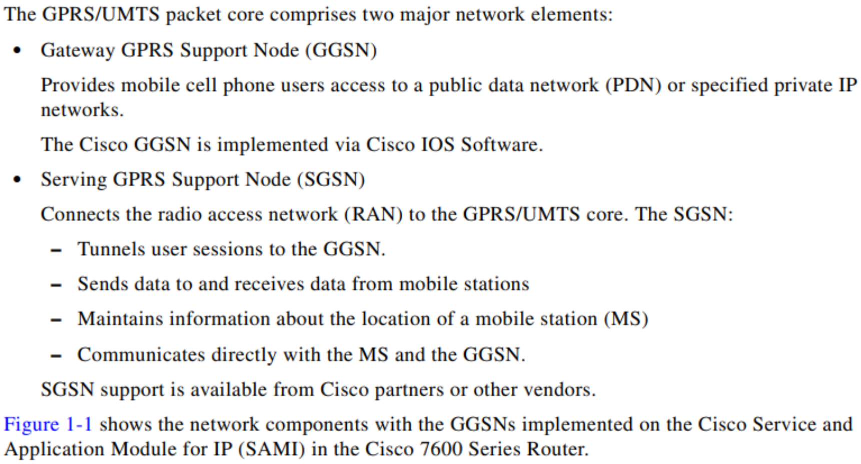

The GPRS/UMTS packet core comprises two major network elements:

- Gateway GPRS Support Node (GGSN)

  Provides mobile cell phone users access to a public data network (PDN) or specified private IP networks.

  The Cisco GGSN is implemented via Cisco IOS Software.

- Serving GPRS Support Node (SGSN)

  Connects the radio access network (RAN) to the GPRS/UMTS core. The SGSN:

  - Tunnels user sessions to the GGSN.
  - Sends data to and receives data from mobile stations
  - Maintains information about the location of a mobile station (MS)
  - Communicates directly with the MS and the GGSN.

  SGSN support is available from Cisco partners or other vendors.

Figure 1-1 shows the network components with the GGSNs implemented on the Cisco Service and Application Module for IP (SAMI) in the Cisco 7600 Series Router.

https://www.cisco.com/c/en/us/td/docs/ios/12_4/12_4y/12_4_24ye3/cfg/12424ye3cfg/ggsnover.pdf

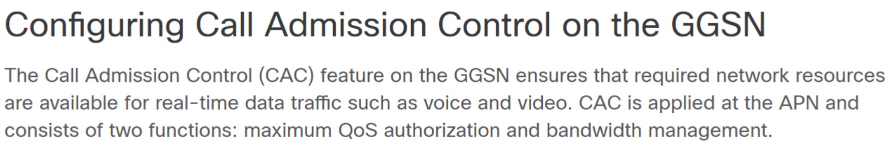

Configuring Call Admission Control on the GGSN

The Call Admission Control (CAC) feature on the GGSN ensures that required network resources are available for real-time data traffic such as voice and video. CAC is applied at the APN and consists of two functions: maximum QoS authorization and bandwidth management.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

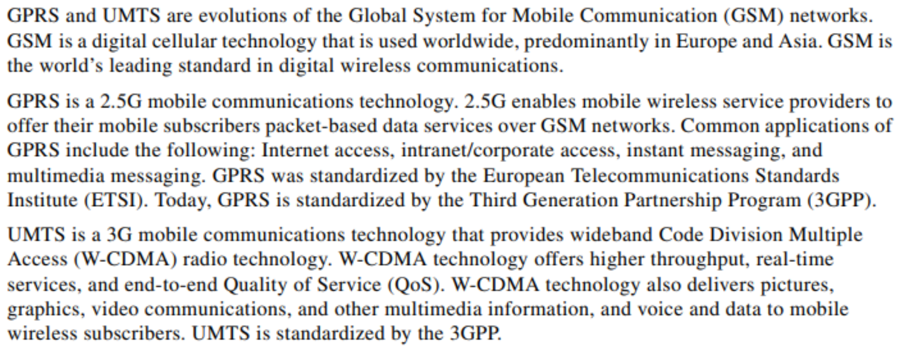

GPRS and UMTS are evolutions of the Global System for Mobile Communication (GSM) networks. GSM is a digital cellular technology that is used worldwide, predominantly in Europe and Asia. GSM is the world's leading standard in digital wireless communications.

GPRS is a 2.5G mobile communications technology. 2.5G enables mobile wireless service providers to offer their mobile subscribers packet-based data services over GSM networks. Common applications of GPRS include the following: Internet access, intranet/corporate access, instant messaging, and multimedia messaging. GPRS was standardized by the European Telecommunications Standards Institute (ETSI). Today, GPRS is standardized by the Third Generation Partnership Program (3GPP).

UMTS is a 3G mobile communications technology that provides wideband Code Division Multiple Access (W-CDMA) radio technology. W-CDMA technology offers higher throughput, real-time services, and end-to-end Quality of Service (QoS). W-CDMA technology also delivers pictures, graphics, video communications, and other multimedia information, and voice and data to mobile wireless subscribers. UMTS is standardized by the 3GPP.

https://www.cisco.com/c/en/us/td/docs/ios/12_4/12_4y/12_4_24ye3/cfg/12424ye3cfg/ggsnover.pdf

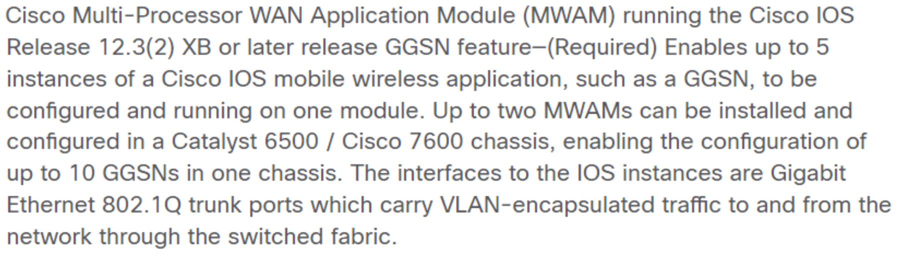

Cisco Multi-Processor WAN Application Module (MWAM) running the Cisco IOS Release 12.3(2) XB or later release GGSN feature—(Required) Enables up to 5 instances of a Cisco IOS mobile wireless application, such as a GGSN, to be configured and running on one module. Up to two MWAMs can be installed and configured in a Catalyst 6500 / Cisco 7600 chassis, enabling the configuration of up to 10 GGSNs in one chassis. The interfaces to the IOS instances are Gigabit Ethernet 802.1Q trunk ports which carry VLAN-encapsulated traffic to and from the network through the switched fabric.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnplan.html

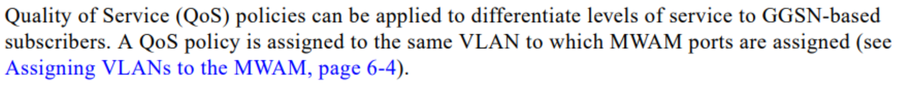

Quality of Service (QoS) policies can be applied to differentiate levels of service to GGSN-based subscribers. A QoS policy is assigned to the same VLAN to which MWAM ports are assigned (see Assigning VLANs to the MWAM, page 6-4).

https://www.cisco.com/c/en/us/td/docs/wireless/mwam/user/guide/mwam1.pdf

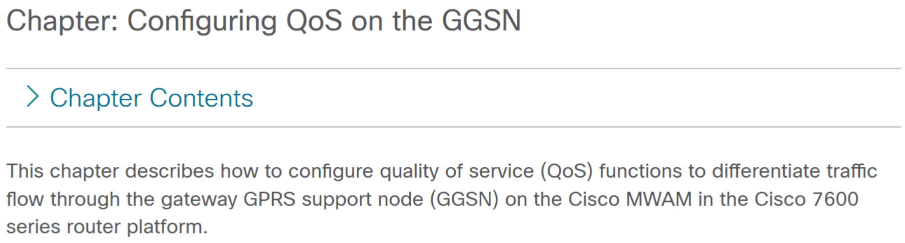

Chapter: Configuring QoS on the GGSN

> Chapter Contents

This chapter describes how to configure quality of service (QoS) functions to differentiate traffic flow through the gateway GPRS support node (GGSN) on the Cisco MWAM in the Cisco 7600 series router platform.

105. https://www.cisco.com/c/en/us/td/docs/ios/12_4/12_4x/12_4_9xg/ggsn7_0/cfg/ggsn70_c/ggsnqos.htmlThe Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, provides for setting a resources requested threshold for each link where the resources requested threshold defines a maximum sum

of forwarding resources requested by applications for their application data flows, ADFs, on the link as described below:

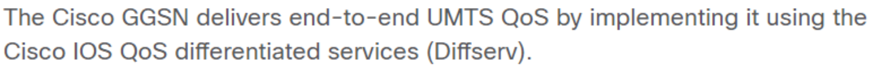

The Cisco GGSN delivers end-to-end UMTS QoS by implementing it using the Cisco IOS QoS differentiated services (Diffserv).

https://www.cisco.com/c/en/us/td/docs/ios/12_4/12_4y/12_4_22ye/ggsn9_0/cfg/12422ye_cfgbk/ggsnqos.html

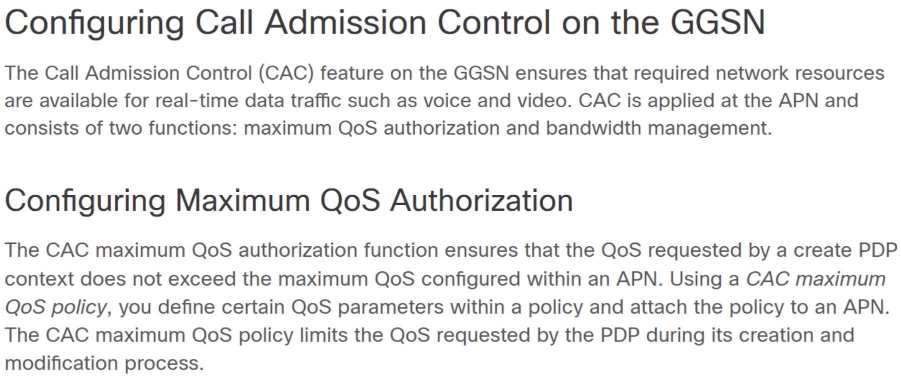

Configuring Call Admission Control on the GGSN

The Call Admission Control (CAC) feature on the GGSN ensures that required network resources are available for real-time data traffic such as voice and video. CAC is applied at the APN and consists of two functions: maximum QoS authorization and bandwidth management.

Configuring Maximum QoS Authorization

The CAC maximum QoS authorization function ensures that the QoS requested by a create PDP context does not exceed the maximum QoS configured within an APN. Using a *CAC maximum QoS policy*, you define certain QoS parameters within a policy and attach the policy to an APN. The CAC maximum QoS policy limits the QoS requested by the PDP during its creation and modification process.

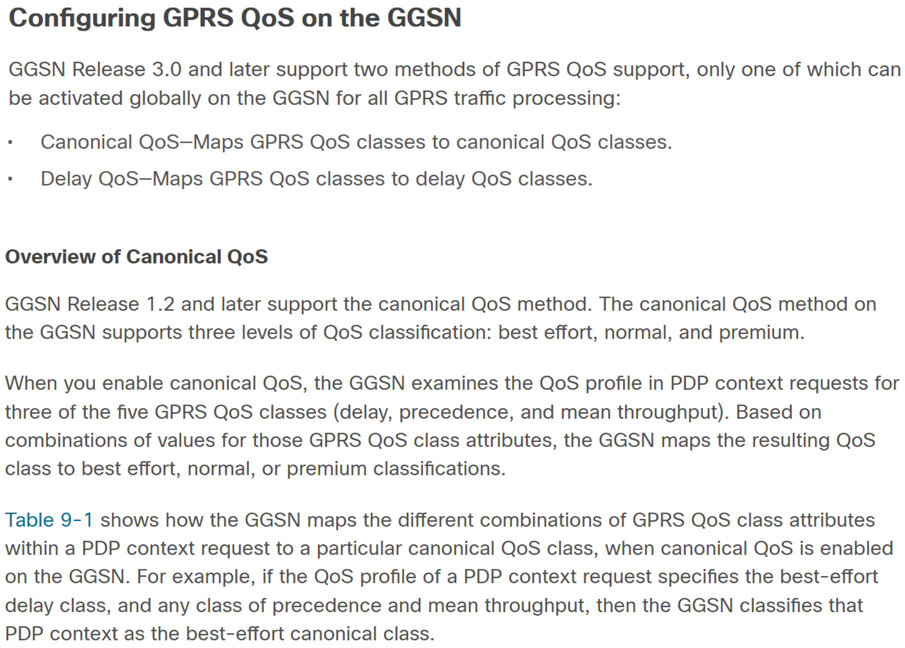

Configuring GPRS QoS on the GGSN

GGSN Release 3.0 and later support two methods of GPRS QoS support, only one of which can be activated globally on the GGSN for all GPRS traffic processing:

- Canonical QoS–Maps GPRS QoS classes to canonical QoS classes.
- Delay QoS–Maps GPRS QoS classes to delay QoS classes.

Overview of Canonical QoS

GGSN Release 1.2 and later support the canonical QoS method. The canonical QoS method on the GGSN supports three levels of QoS classification: best effort, normal, and premium.

When you enable canonical QoS, the GGSN examines the QoS profile in PDP context requests for three of the five GPRS QoS classes (delay, precedence, and mean throughput). Based on combinations of values for those GPRS QoS class attributes, the GGSN maps the resulting QoS class to best effort, normal, or premium classifications.

Table 9-1 shows how the GGSN maps the different combinations of GPRS QoS class attributes within a PDP context request to a particular canonical QoS class, when canonical QoS is enabled on the GGSN. For example, if the QoS profile of a PDP context request specifies the best-effort delay class, and any class of precedence and mean throughput, then the GGSN classifies that PDP context as the best-effort canonical class.

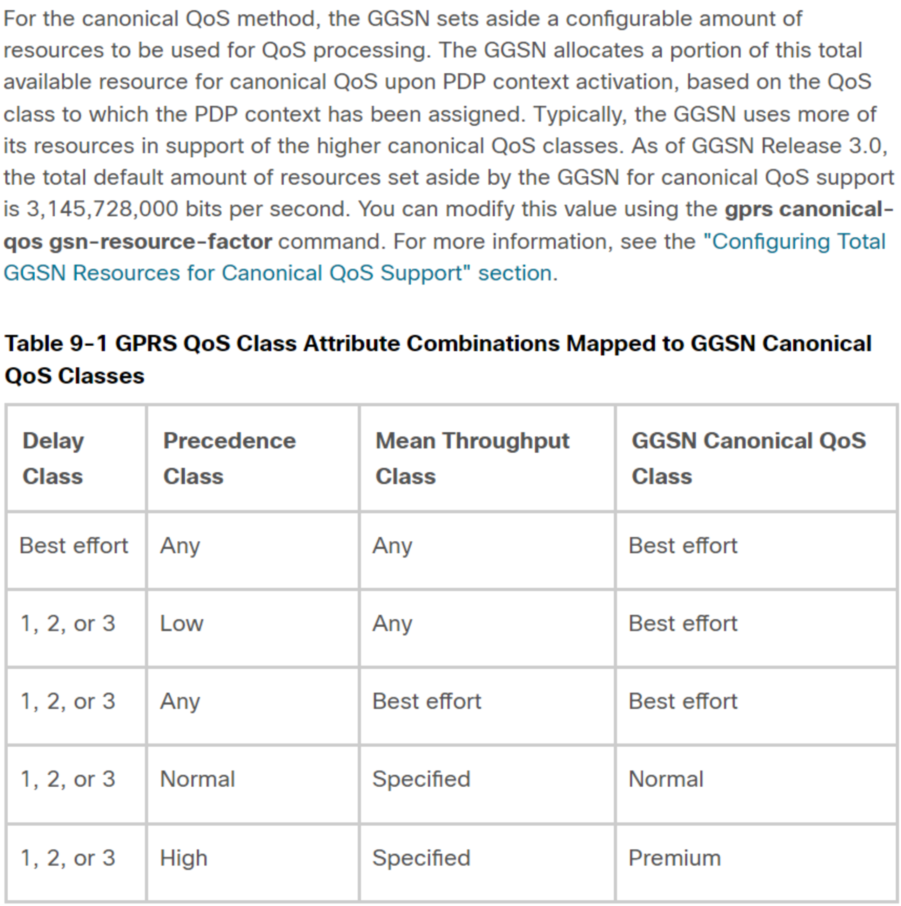

For the canonical QoS method, the GGSN sets aside a configurable amount of resources to be used for QoS processing. The GGSN allocates a portion of this total available resource for canonical QoS upon PDP context activation, based on the QoS class to which the PDP context has been assigned. Typically, the GGSN uses more of its resources in support of the higher canonical QoS classes. As of GGSN Release 3.0, the total default amount of resources set aside by the GGSN for canonical QoS support is 3,145,728,000 bits per second. You can modify this value using the **gprs canonical-qos gsn-resource-factor** command. For more information, see the "Configuring Total GGSN Resources for Canonical QoS Support" section.

Table 9-1 GPRS QoS Class Attribute Combinations Mapped to GGSN Canonical QoS Classes

Delay Class	Precedence Class	Mean Throughput Class	GGSN Canonical QoS Class
Best effort	Any	Any	Best effort
1, 2, or 3	Low	Any	Best effort
1, 2, or 3	Any	Best effort	Best effort
1, 2, or 3	Normal	Specified	Normal
1, 2, or 3	High	Specified	Premium

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

106. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, provides for calculating a measurement level threshold, the measurement level threshold being a rate lower than the resources requested threshold, as described below:

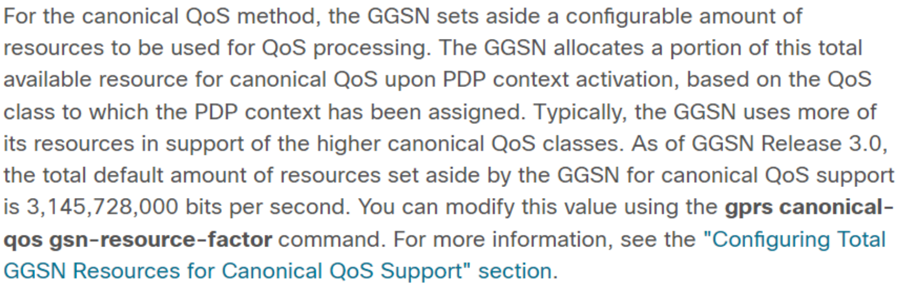

For the canonical QoS method, the GGSN sets aside a configurable amount of resources to be used for QoS processing. The GGSN allocates a portion of this total available resource for canonical QoS upon PDP context activation, based on the QoS class to which the PDP context has been assigned. Typically, the GGSN uses more of its resources in support of the higher canonical QoS classes. As of GGSN Release 3.0, the total default amount of resources set aside by the GGSN for canonical QoS support is 3,145,728,000 bits per second. You can modify this value using the **gprs canonical-qos gsn-resource-factor** command. For more information, see the "Configuring Total GGSN Resources for Canonical QoS Support" section.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

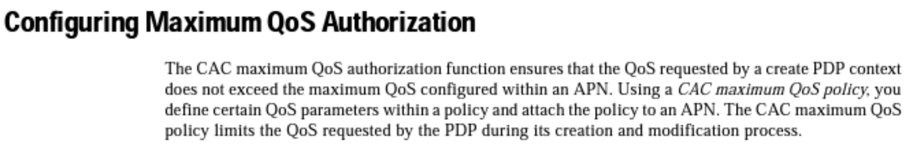

Configuring Maximum QoS Authorization

The CAC maximum QoS authorization function ensures that the QoS requested by a create PDP context does not exceed the maximum QoS configured within an APN. Using a *CAC maximum QoS policy*, you define certain QoS parameters within a policy and attach the policy to an APN. The CAC maximum QoS policy limits the QoS requested by the PDP during its creation and modification process.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.pdf

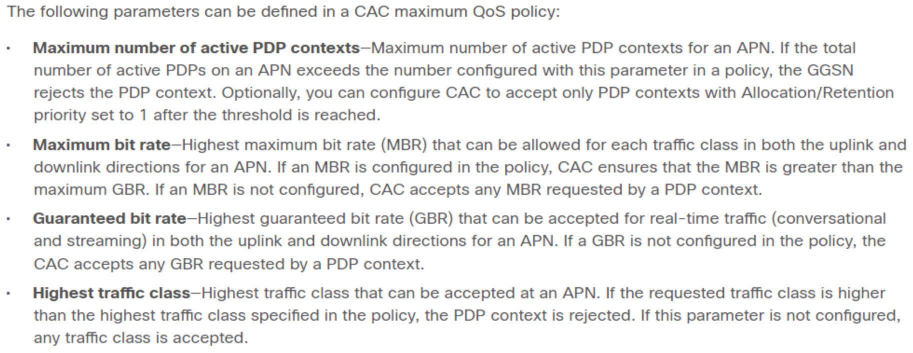

The following parameters can be defined in a CAC maximum QoS policy:

- **Maximum number of active PDP contexts**–Maximum number of active PDP contexts for an APN. If the total number of active PDPs on an APN exceeds the number configured with this parameter in a policy, the GGSN rejects the PDP context. Optionally, you can configure CAC to accept only PDP contexts with Allocation/Retention priority set to 1 after the threshold is reached.
- **Maximum bit rate**–Highest maximum bit rate (MBR) that can be allowed for each traffic class in both the uplink and downlink directions for an APN. If an MBR is configured in the policy, CAC ensures that the MBR is greater than the maximum GBR. If an MBR is not configured, CAC accepts any MBR requested by a PDP context.
- **Guaranteed bit rate**–Highest guaranteed bit rate (GBR) that can be accepted for real-time traffic (conversational and streaming) in both the uplink and downlink directions for an APN. If a GBR is not configured in the policy, the CAC accepts any GBR requested by a PDP context.
- **Highest traffic class**–Highest traffic class that can be accepted at an APN. If the requested traffic class is higher than the highest traffic class specified in the policy, the PDP context is rejected. If this parameter is not configured, any traffic class is accepted.

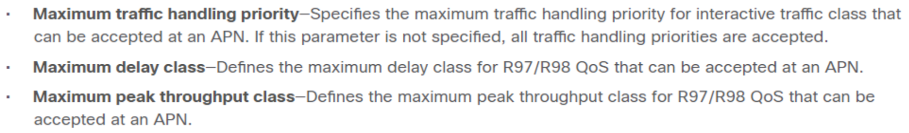

- **Maximum traffic handling priority**–Specifies the maximum traffic handling priority for interactive traffic class that can be accepted at an APN. If this parameter is not specified, all traffic handling priorities are accepted.
- **Maximum delay class**–Defines the maximum delay class for R97/R98 QoS that can be accepted at an APN.
- **Maximum peak throughput class**–Defines the maximum peak throughput class for R97/R98 QoS that can be accepted at an APN.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

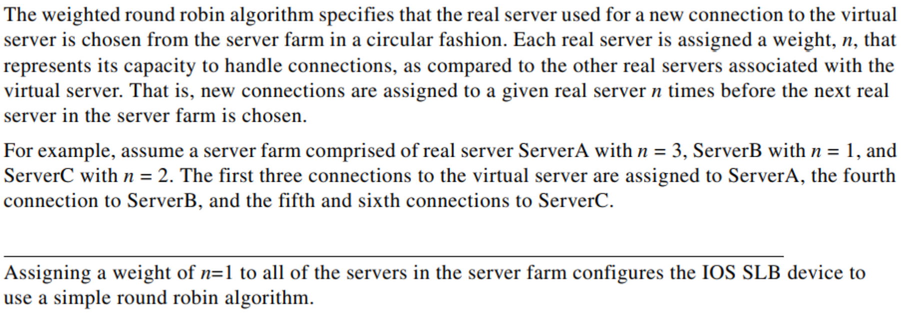

The weighted round robin algorithm specifies that the real server used for a new connection to the virtual server is chosen from the server farm in a circular fashion. Each real server is assigned a weight, *n*, that represents its capacity to handle connections, as compared to the other real servers associated with the virtual server. That is, new connections are assigned to a given real server *n* times before the next real server in the server farm is chosen.

For example, assume a server farm comprised of real server ServerA with *n* = 3, ServerB with *n* = 1, and ServerC with *n* = 2. The first three connections to the virtual server are assigned to ServerA, the fourth connection to ServerB, and the fifth and sixth connections to ServerC.

Assigning a weight of *n*=1 to all of the servers in the server farm configures the IOS SLB device to use a simple round robin algorithm.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

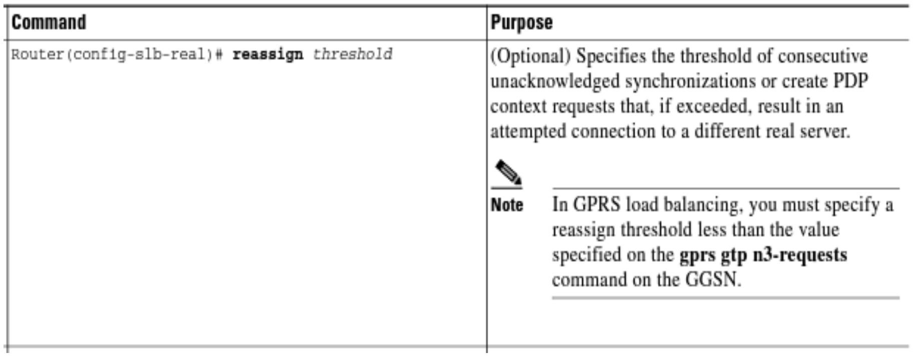

Command	Purpose
`Router(config-slb-real)# reassign threshold`	(Optional) Specifies the threshold of consecutive unacknowledged synchronizations or create PDP context requests that, if exceeded, result in an attempted connection to a different real server. **Note** In GPRS load balancing, you must specify a reassign threshold less than the value specified on the **gprs gtp n3-requests** command on the GGSN.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

107. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, handles determining whether a reservation level exceeds the measurement level threshold, as described below:

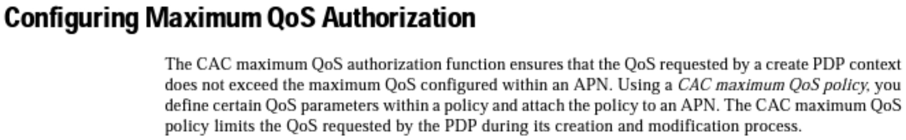

Configuring Maximum QoS Authorization

The CAC maximum QoS authorization function ensures that the QoS requested by a create PDP context does not exceed the maximum QoS configured within an APN. Using a *CAC maximum QoS policy*, you define certain QoS parameters within a policy and attach the policy to an APN. The CAC maximum QoS policy limits the QoS requested by the PDP during its creation and modification process.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.pdf

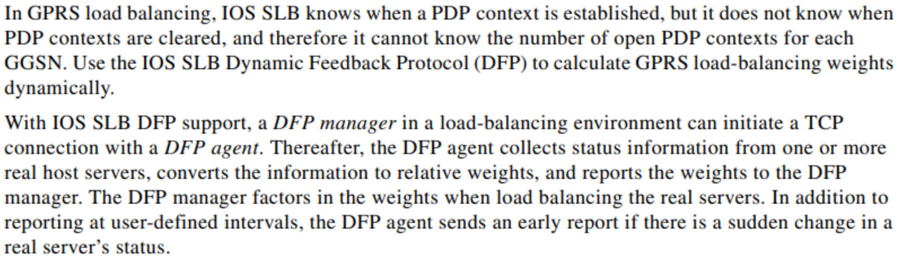

In GPRS load balancing, IOS SLB knows when a PDP context is established, but it does not know when PDP contexts are cleared, and therefore it cannot know the number of open PDP contexts for each GGSN. Use the IOS SLB Dynamic Feedback Protocol (DFP) to calculate GPRS load-balancing weights dynamically.

With IOS SLB DFP support, a *DFP manager* in a load-balancing environment can initiate a TCP connection with a *DFP agent*. Thereafter, the DFP agent collects status information from one or more real host servers, converts the information to relative weights, and reports the weights to the DFP manager. The DFP manager factors in the weights when load balancing the real servers. In addition to reporting at user-defined intervals, the DFP agent sends an early report if there is a sudden change in a real server's status.

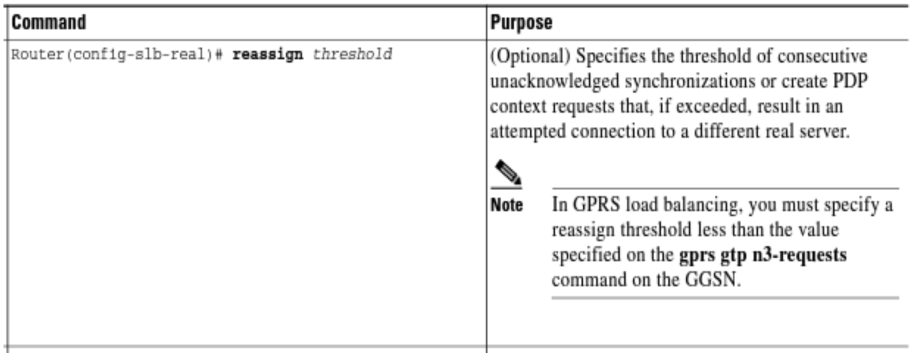

Command	Purpose
Router(config-slb-real)# **reassign** *threshold*	(Optional) Specifies the threshold of consecutive unacknowledged synchronizations or create PDP context requests that, if exceeded, result in an attempted connection to a different real server. **Note** In GPRS load balancing, you must specify a reassign threshold less than the value specified on the **gprs gtp n3-requests** command on the GGSN.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

108. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, upon determining that the reservation level exceeds the measurement level threshold, repeatedly measuring, during usage, multiplexing

properties of aggregated ADFs on each link, each measuring being performed over a period of time, as described below:

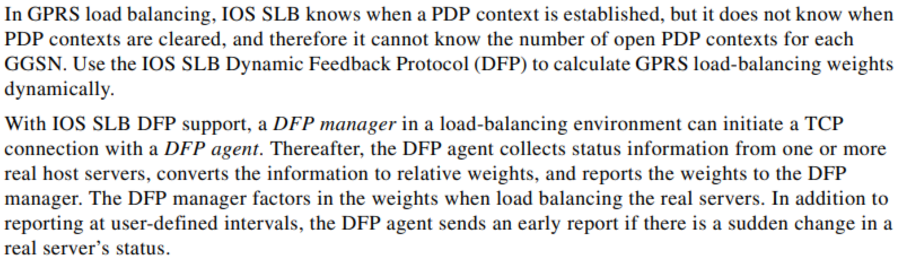

In GPRS load balancing, IOS SLB knows when a PDP context is established, but it does not know when PDP contexts are cleared, and therefore it cannot know the number of open PDP contexts for each GGSN. Use the IOS SLB Dynamic Feedback Protocol (DFP) to calculate GPRS load-balancing weights dynamically.

With IOS SLB DFP support, a *DFP manager* in a load-balancing environment can initiate a TCP connection with a *DFP agent*. Thereafter, the DFP agent collects status information from one or more real host servers, converts the information to relative weights, and reports the weights to the DFP manager. The DFP manager factors in the weights when load balancing the real servers. In addition to reporting at user-defined intervals, the DFP agent sends an early report if there is a sudden change in a real server's status.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

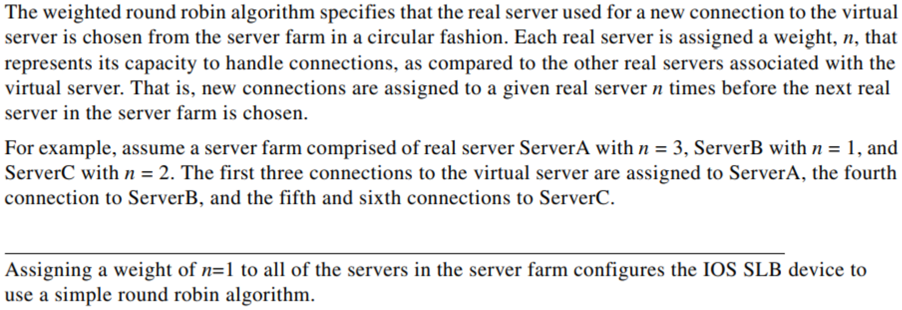

The weighted round robin algorithm specifies that the real server used for a new connection to the virtual server is chosen from the server farm in a circular fashion. Each real server is assigned a weight, *n*, that represents its capacity to handle connections, as compared to the other real servers associated with the virtual server. That is, new connections are assigned to a given real server *n* times before the next real server in the server farm is chosen.

For example, assume a server farm comprised of real server ServerA with *n* = 3, ServerB with *n* = 1, and ServerC with *n* = 2. The first three connections to the virtual server are assigned to ServerA, the fourth connection to ServerB, and the fifth and sixth connections to ServerC.

Assigning a weight of *n*=1 to all of the servers in the server farm configures the IOS SLB device to use a simple round robin algorithm.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

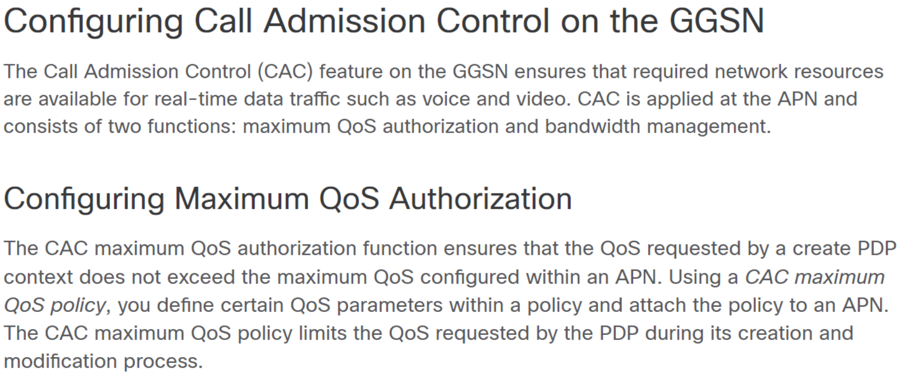

Configuring Call Admission Control on the GGSN

The Call Admission Control (CAC) feature on the GGSN ensures that required network resources are available for real-time data traffic such as voice and video. CAC is applied at the APN and consists of two functions: maximum QoS authorization and bandwidth management.

Configuring Maximum QoS Authorization

The CAC maximum QoS authorization function ensures that the QoS requested by a create PDP context does not exceed the maximum QoS configured within an APN. Using a *CAC maximum QoS policy*, you define certain QoS parameters within a policy and attach the policy to an APN. The CAC maximum QoS policy limits the QoS requested by the PDP during its creation and modification process.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

Configuring Call Admission Control

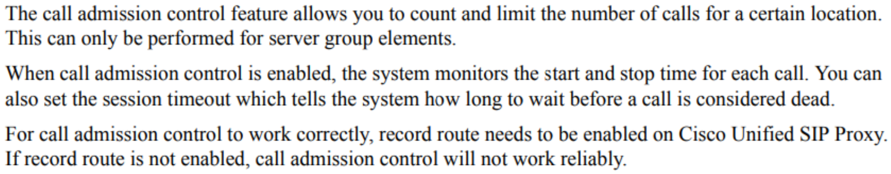

The call admission control feature allows you to count and limit the number of calls for a certain location. This can only be performed for server group elements.

When call admission control is enabled, the system monitors the start and stop time for each call. You can also set the session timeout which tells the system how long to wait before a call is considered dead.

For call admission control to work correctly, record route needs to be enabled on Cisco Unified SIP Proxy. If record route is not enabled, call admission control will not work reliably.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cusp/rel9_1/gui_configuration/en_US/cusp_m_admin_gui_olh_chapter_01101.pdf

109. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, dynamically adapts the resources requested threshold by utilizing the measured multiplexing properties of the ADFs on each link and by utilizing knowledge about the forwarding resources of the links, as described below:

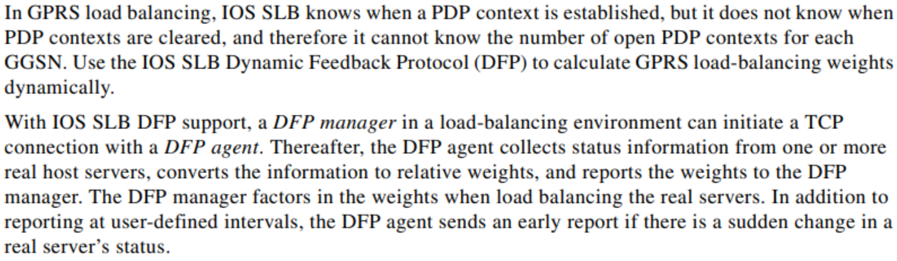

In GPRS load balancing, IOS SLB knows when a PDP context is established, but it does not know when PDP contexts are cleared, and therefore it cannot know the number of open PDP contexts for each GGSN. Use the IOS SLB Dynamic Feedback Protocol (DFP) to calculate GPRS load-balancing weights dynamically.

With IOS SLB DFP support, a *DFP manager* in a load-balancing environment can initiate a TCP connection with a *DFP agent*. Thereafter, the DFP agent collects status information from one or more real host servers, converts the information to relative weights, and reports the weights to the DFP manager. The DFP manager factors in the weights when load balancing the real servers. In addition to reporting at user-defined intervals, the DFP agent sends an early report if there is a sudden change in a real server's status.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

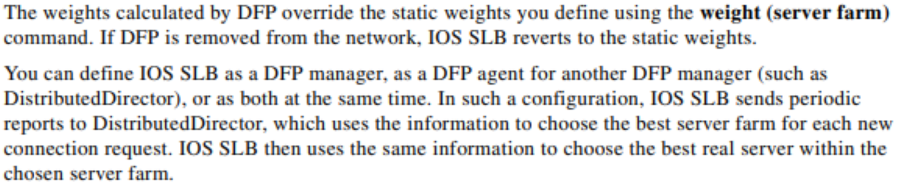

The weights calculated by DFP override the static weights you define using the **weight (server farm)** command. If DFP is removed from the network, IOS SLB reverts to the static weights.

You can define IOS SLB as a DFP manager, as a DFP agent for another DFP manager (such as DistributedDirector), or as both at the same time. In such a configuration, IOS SLB sends periodic reports to DistributedDirector, which uses the information to choose the best server farm for each new connection request. IOS SLB then uses the same information to choose the best real server within the chosen server farm.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

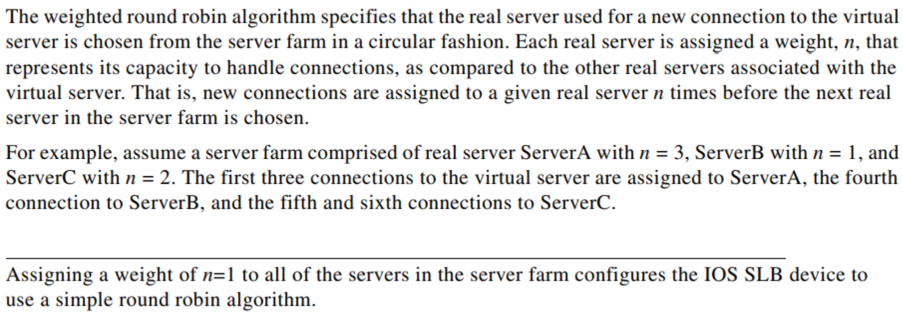

The weighted round robin algorithm specifies that the real server used for a new connection to the virtual server is chosen from the server farm in a circular fashion. Each real server is assigned a weight, *n*, that represents its capacity to handle connections, as compared to the other real servers associated with the virtual server. That is, new connections are assigned to a given real server *n* times before the next real server in the server farm is chosen.

For example, assume a server farm comprised of real server ServerA with *n* = 3, ServerB with *n* = 1, and ServerC with *n* = 2. The first three connections to the virtual server are assigned to ServerA, the fourth connection to ServerB, and the fifth and sixth connections to ServerC.

Assigning a weight of *n*=1 to all of the servers in the server farm configures the IOS SLB device to use a simple round robin algorithm.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

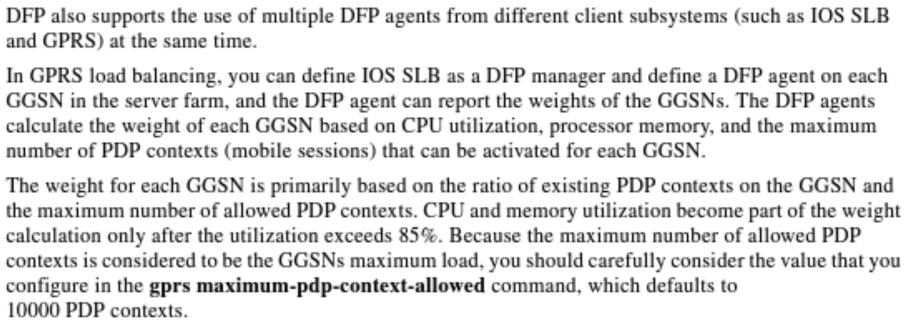

DFP also supports the use of multiple DFP agents from different client subsystems (such as IOS SLB and GPRS) at the same time.

In GPRS load balancing, you can define IOS SLB as a DFP manager and define a DFP agent on each GGSN in the server farm, and the DFP agent can report the weights of the GGSNs. The DFP agents calculate the weight of each GGSN based on CPU utilization, processor memory, and the maximum number of PDP contexts (mobile sessions) that can be activated for each GGSN.

The weight for each GGSN is primarily based on the ratio of existing PDP contexts on the GGSN and the maximum number of allowed PDP contexts. CPU and memory utilization become part of the weight calculation only after the utilization exceeds 85%. Because the maximum number of allowed PDP contexts is considered to be the GGSNs maximum load, you should carefully consider the value that you configure in the **gprs maximum-pdp-context-allowed** command, which defaults to 10000 PDP contexts.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

110. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, recalculates the measurement level threshold based on the dynamically adapted resources requested threshold, as described below:

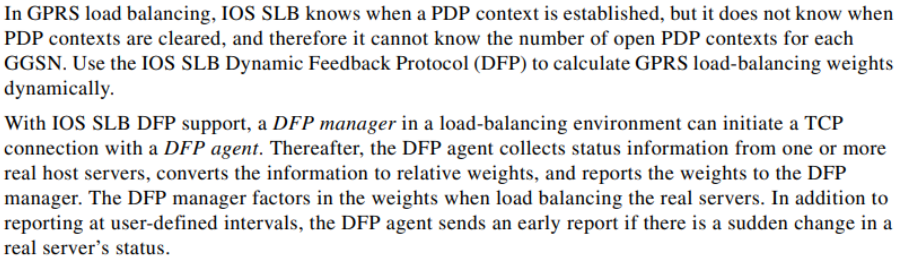

In GPRS load balancing, IOS SLB knows when a PDP context is established, but it does not know when PDP contexts are cleared, and therefore it cannot know the number of open PDP contexts for each GGSN. Use the IOS SLB Dynamic Feedback Protocol (DFP) to calculate GPRS load-balancing weights dynamically.

With IOS SLB DFP support, a *DFP manager* in a load-balancing environment can initiate a TCP connection with a *DFP agent*. Thereafter, the DFP agent collects status information from one or more real host servers, converts the information to relative weights, and reports the weights to the DFP manager. The DFP manager factors in the weights when load balancing the real servers. In addition to reporting at user-defined intervals, the DFP agent sends an early report if there is a sudden change in a real server's status.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

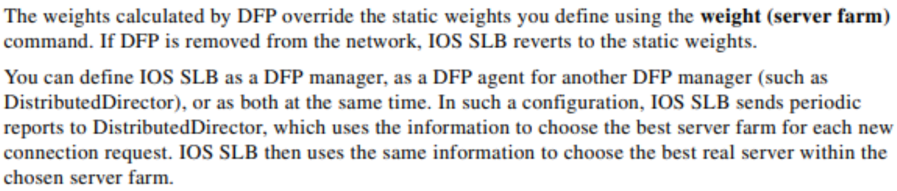

The weights calculated by DFP override the static weights you define using the **weight (server farm)** command. If DFP is removed from the network, IOS SLB reverts to the static weights.

You can define IOS SLB as a DFP manager, as a DFP agent for another DFP manager (such as DistributedDirector), or as both at the same time. In such a configuration, IOS SLB sends periodic reports to DistributedDirector, which uses the information to choose the best server farm for each new connection request. IOS SLB then uses the same information to choose the best real server within the chosen server farm.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

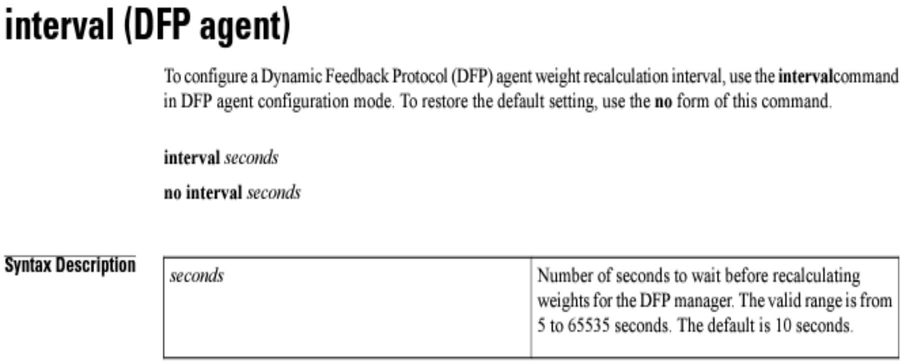

interval (DFP agent)

To configure a Dynamic Feedback Protocol (DFP) agent weight recalculation interval, use the **interval**command in DFP agent configuration mode. To restore the default setting, use the **no** form of this command.

interval *seconds*

no interval *seconds*

Syntax Description

seconds	Number of seconds to wait before recalculating weights for the DFP manager. The valid range is from 5 to 65535 seconds. The default is 10 seconds.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/slb/command/slb-cr-book.pdf

111. The Cisco Multiprocessor WAN Application Module (MWAM), when used with the Cisco Catalyst 6500 and 7600 Series routing platforms, prevents an overload before it occurs by controlling admission to each link based on the dynamically adapted resources requested threshold, as shown below:

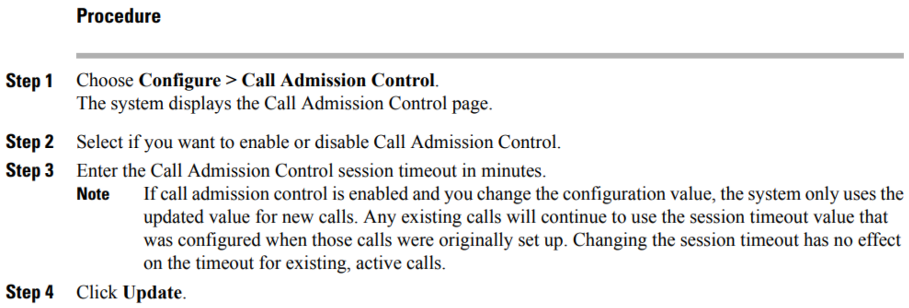

Procedure

Step 1 Choose **Configure > Call Admission Control**.
The system displays the Call Admission Control page.

Step 2 Select if you want to enable or disable Call Admission Control.

Step 3 Enter the Call Admission Control session timeout in minutes.

Note If call admission control is enabled and you change the configuration value, the system only uses the updated value for new calls. Any existing calls will continue to use the session timeout value that was configured when those calls were originally set up. Changing the session timeout has no effect on the timeout for existing, active calls.

Step 4 Click **Update**.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cusp/rel9_1/gui_configuration/en_US/cusp_m_admin_gui_olh_chapter_01101.pdf

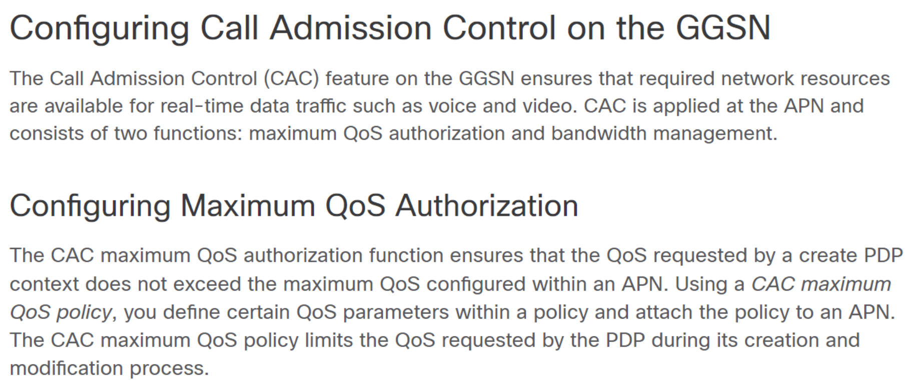

Configuring Call Admission Control on the GGSN

The Call Admission Control (CAC) feature on the GGSN ensures that required network resources are available for real-time data traffic such as voice and video. CAC is applied at the APN and consists of two functions: maximum QoS authorization and bandwidth management.

Configuring Maximum QoS Authorization

The CAC maximum QoS authorization function ensures that the QoS requested by a create PDP context does not exceed the maximum QoS configured within an APN. Using a *CAC maximum QoS policy*, you define certain QoS parameters within a policy and attach the policy to an APN. The CAC maximum QoS policy limits the QoS requested by the PDP during its creation and modification process.

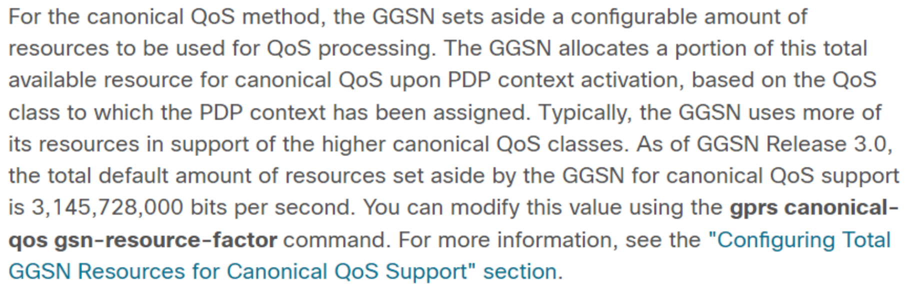

For the canonical QoS method, the GGSN sets aside a configurable amount of resources to be used for QoS processing. The GGSN allocates a portion of this total available resource for canonical QoS upon PDP context activation, based on the QoS class to which the PDP context has been assigned. Typically, the GGSN uses more of its resources in support of the higher canonical QoS classes. As of GGSN Release 3.0, the total default amount of resources set aside by the GGSN for canonical QoS support is 3,145,728,000 bits per second. You can modify this value using the **gprs canonical-qos gsn-resource-factor** command. For more information, see the "Configuring Total GGSN Resources for Canonical QoS Support" section.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

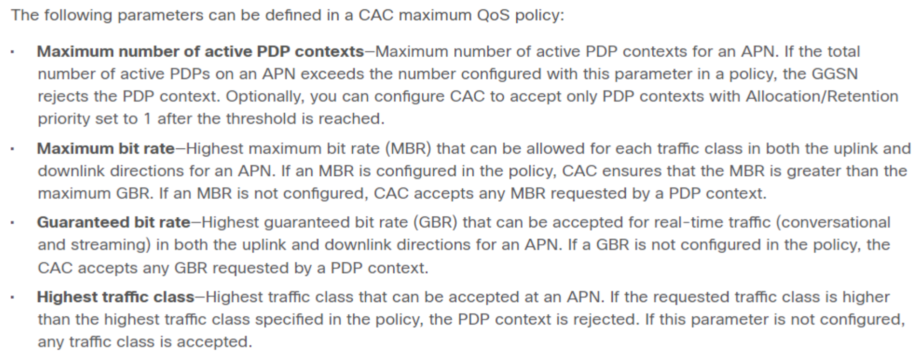

The following parameters can be defined in a CAC maximum QoS policy:

- **Maximum number of active PDP contexts**–Maximum number of active PDP contexts for an APN. If the total number of active PDPs on an APN exceeds the number configured with this parameter in a policy, the GGSN rejects the PDP context. Optionally, you can configure CAC to accept only PDP contexts with Allocation/Retention priority set to 1 after the threshold is reached.
- **Maximum bit rate**–Highest maximum bit rate (MBR) that can be allowed for each traffic class in both the uplink and downlink directions for an APN. If an MBR is configured in the policy, CAC ensures that the MBR is greater than the maximum GBR. If an MBR is not configured, CAC accepts any MBR requested by a PDP context.
- **Guaranteed bit rate**–Highest guaranteed bit rate (GBR) that can be accepted for real-time traffic (conversational and streaming) in both the uplink and downlink directions for an APN. If a GBR is not configured in the policy, the CAC accepts any GBR requested by a PDP context.
- **Highest traffic class**–Highest traffic class that can be accepted at an APN. If the requested traffic class is higher than the highest traffic class specified in the policy, the PDP context is rejected. If this parameter is not configured, any traffic class is accepted.

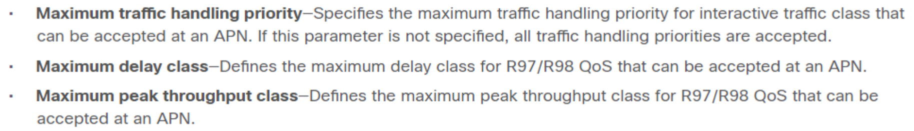

- **Maximum traffic handling priority**–Specifies the maximum traffic handling priority for interactive traffic class that can be accepted at an APN. If this parameter is not specified, all traffic handling priorities are accepted.
- **Maximum delay class**–Defines the maximum delay class for R97/R98 QoS that can be accepted at an APN.
- **Maximum peak throughput class**–Defines the maximum peak throughput class for R97/R98 QoS that can be accepted at an APN.

https://www.cisco.com/c/en/us/td/docs/ios/12_4t/mw_ggsn/configuration/guide/ggsnqos.html

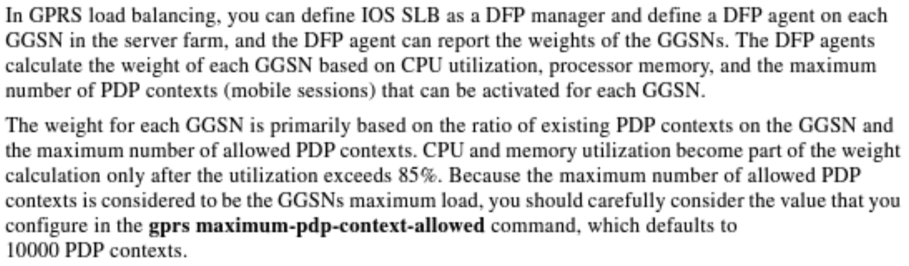

In GPRS load balancing, you can define IOS SLB as a DFP manager and define a DFP agent on each GGSN in the server farm, and the DFP agent can report the weights of the GGSNs. The DFP agents calculate the weight of each GGSN based on CPU utilization, processor memory, and the maximum number of PDP contexts (mobile sessions) that can be activated for each GGSN.

The weight for each GGSN is primarily based on the ratio of existing PDP contexts on the GGSN and the maximum number of allowed PDP contexts. CPU and memory utilization become part of the weight calculation only after the utilization exceeds 85%. Because the maximum number of allowed PDP contexts is considered to be the GGSNs maximum load, you should carefully consider the value that you configure in the **gprs maximum-pdp-context-allowed** command, which defaults to 10000 PDP contexts.

https://www.cisco.com/c/en/us/td/docs/ios/12_2/12_2y/12_2yy/ggsn31/31cfg/ggsn31_c/ggsnslb.pdf

112. Cisco markets, offers to sell, sells, and distributes the Cisco QoS Systems, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’284 Patent. The Cisco QoS Systems are not staple articles or commodities of commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

113. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the ’284 Patent. Cisco’s infringement of the ’284 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney’s fees, and costs in bringing this action.

114. Cisco’s direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary and permanent injunctive relief and damages as a result of Cisco’s infringement of the ’601 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 5

(Infringement of U.S. Pat. No. 7,720,966)

115. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

116. Cisco has infringed and continues to infringe one or more claims of the ’966 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the ’966 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the ’966 Patent and continues to contribute to the infringement of the ‘796 Patent in violation of 35 U.S.C. §271(c).

117. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 1 of the ’966 Patent at least by making, using, offering to sell, importing, and/or selling end-to-end Quality of Service (QoS) as part of its Internetwork Operating System (IOS) versions 12 and above.

118. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell the Cisco Internetworking Operating System (IOS) versions 12 and above (the “Cisco IOS”) in the Unites States, and encourages distributors to sell, offer to sell and use, and encourages Cisco’s customers to use, the Cisco IOS in the United States with knowledge that the Cisco IOS infringes the ’966 Patent.

119. Cisco has had knowledge of and notice of the ’966 Patent and its infringement since before the filing of the original Complaint. For example, on information and belief, Cisco had knowledge of and notice of the ’966 Patent as a result of a partnership with NetSocket during at least 2012. Cisco offers to sell the Cisco QoS System in this District and does so with knowledge that the sale and use of the Cisco QoS System infringes and with the intent for its customers to use the Cisco QoS System in an infringing manner.

120. The Cisco QoS System satisfies each of the limitations of at least claim 1 of the ’966 Patent.

121. For example, the Cisco QoS provides for controlling resources within a data network implemented by a first network level having a first addressing scheme and at least a second network level having a second addressing scheme, each network level providing connectivity over at least one network domain within a Multi-Protocol Label Switch (MPLS) Virtual Private Network (VPN) Inter-Autonomous System (AS) system such as the one shown in Figure 1 below:

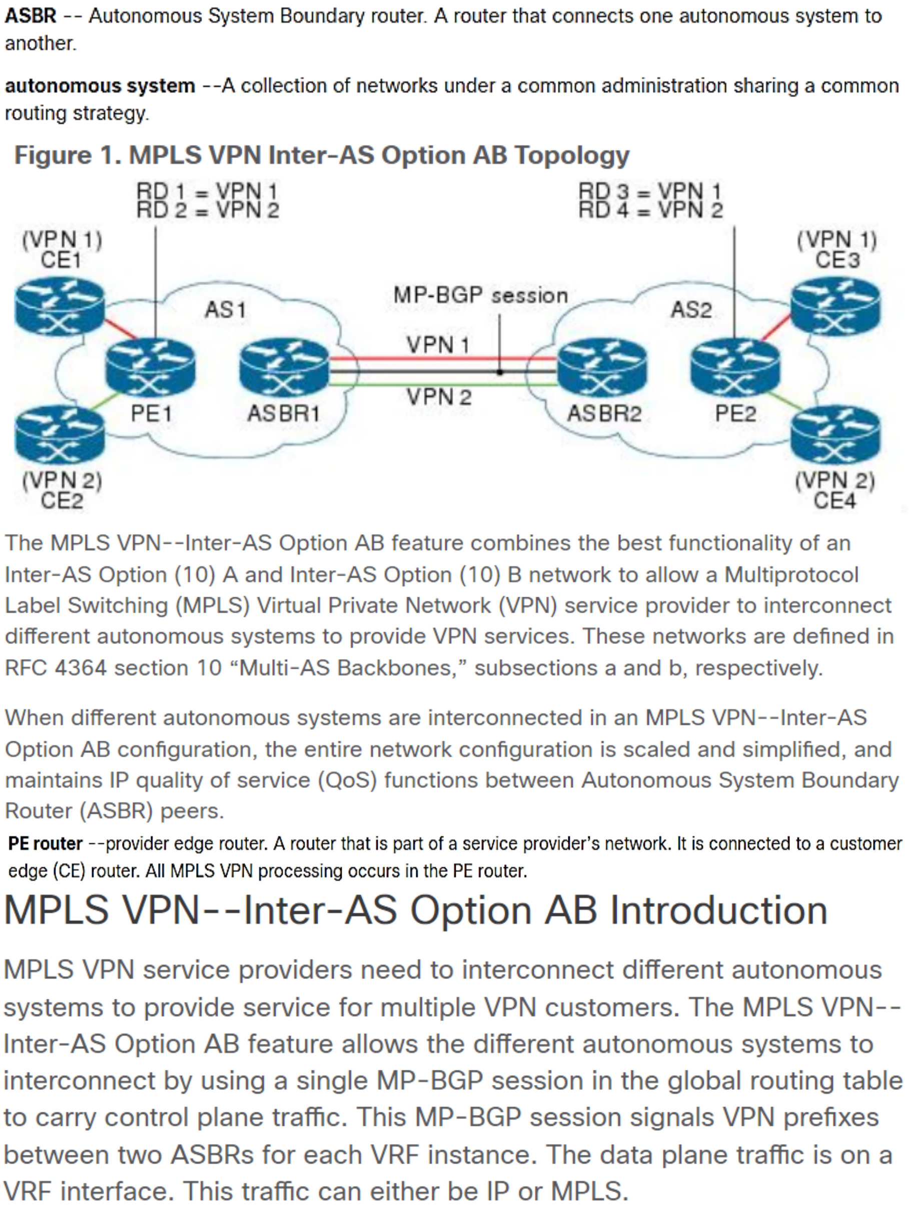

ASBR -- Autonomous System Boundary router. A router that connects one autonomous system to another.

autonomous system --A collection of networks under a common administration sharing a common routing strategy.

Figure 1. MPLS VPN Inter-AS Option AB Topology

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," subsections a and b, respectively.

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

PE router --provider edge router. A router that is part of a service provider's network. It is connected to a customer edge (CE) router. All MPLS VPN processing occurs in the PE router.

MPLS VPN--Inter-AS Option AB Introduction

MPLS VPN service providers need to interconnect different autonomous systems to provide service for multiple VPN customers. The MPLS VPN--Inter-AS Option AB feature allows the different autonomous systems to interconnect by using a single MP-BGP session in the global routing table to carry control plane traffic. This MP-BGP session signals VPN prefixes between two ASBRs for each VRF instance. The data plane traffic is on a VRF interface. This traffic can either be IP or MPLS.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

122. The MPLS-VPN Inter-AS system provided by Cisco IOS provides for controlling resources of the first network level by a first group of Network Resource Managers (NRMs), as shown below:

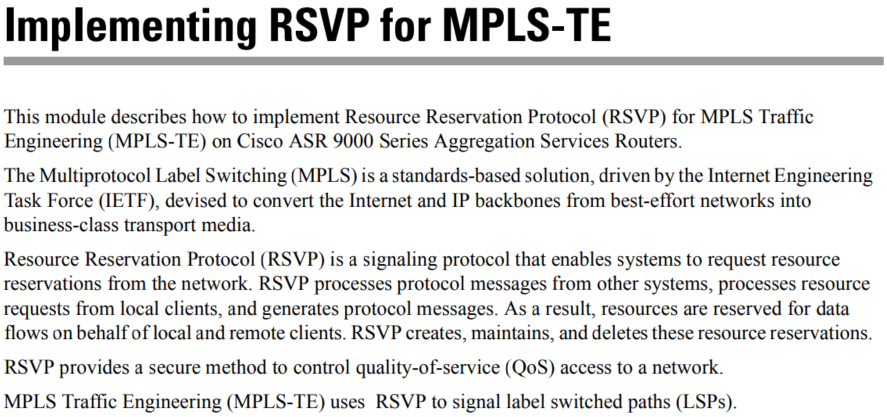

Implementing RSVP for MPLS-TE

This module describes how to implement Resource Reservation Protocol (RSVP) for MPLS Traffic Engineering (MPLS-TE) on Cisco ASR 9000 Series Aggregation Services Routers.

The Multiprotocol Label Switching (MPLS) is a standards-based solution, driven by the Internet Engineering Task Force (IETF), devised to convert the Internet and IP backbones from best-effort networks into business-class transport media.

Resource Reservation Protocol (RSVP) is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations.

RSVP provides a secure method to control quality-of-service (QoS) access to a network.

MPLS Traffic Engineering (MPLS-TE) uses RSVP to signal label switched paths (LSPs).

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

Figure 1. MPLS VPN Inter-AS Option AB Topology

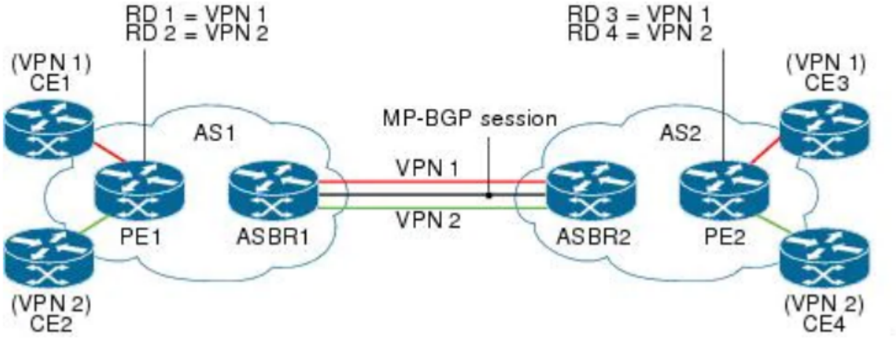

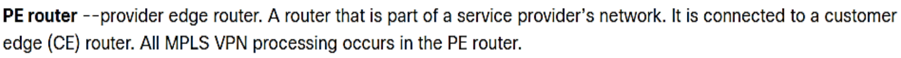

PE router --provider edge router. A router that is part of a service provider’s network. It is connected to a customer edge (CE) router. All MPLS VPN processing occurs in the PE router.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

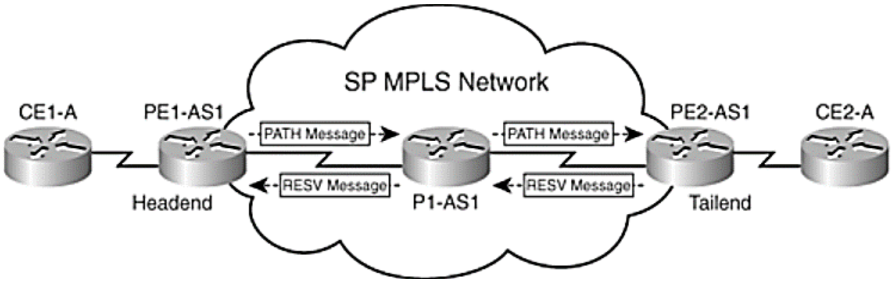

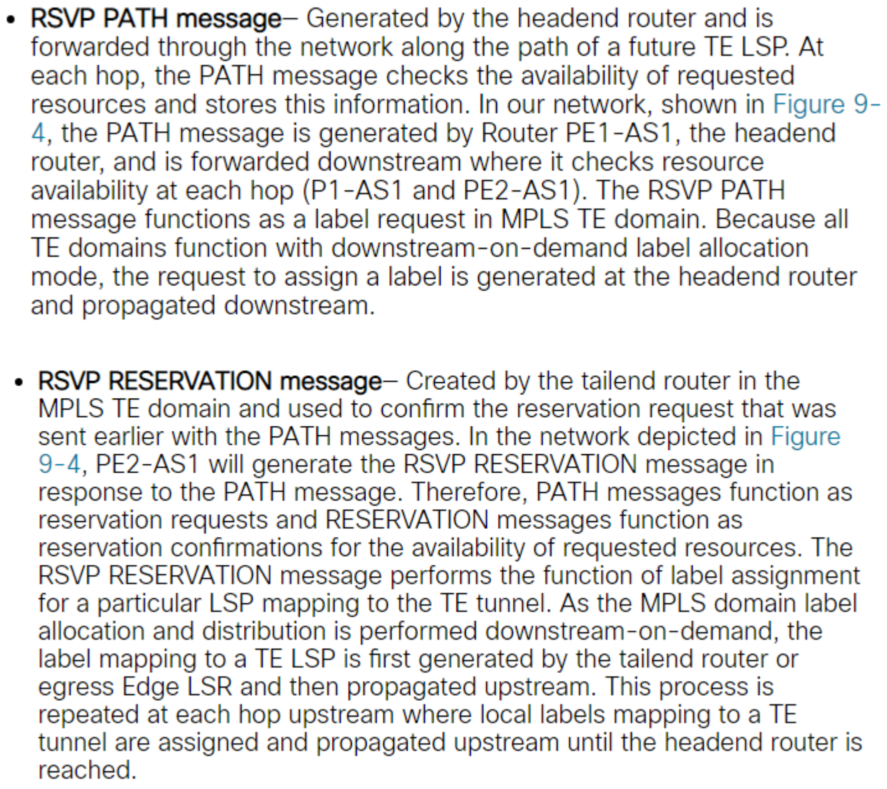

- **RSVP PATH message**– Generated by the headend router and is forwarded through the network along the path of a future TE LSP. At each hop, the PATH message checks the availability of requested resources and stores this information. In our network, shown in Figure 9-4, the PATH message is generated by Router PE1-AS1, the headend router, and is forwarded downstream where it checks resource availability at each hop (P1-AS1 and PE2-AS1). The RSVP PATH message functions as a label request in MPLS TE domain. Because all TE domains function with downstream-on-demand label allocation mode, the request to assign a label is generated at the headend router and propagated downstream.

- **RSVP RESERVATION message**– Created by the tailend router in the MPLS TE domain and used to confirm the reservation request that was sent earlier with the PATH messages. In the network depicted in Figure 9-4, PE2-AS1 will generate the RSVP RESERVATION message in response to the PATH message. Therefore, PATH messages function as reservation requests and RESERVATION messages function as reservation confirmations for the availability of requested resources. The RSVP RESERVATION message performs the function of label assignment for a particular LSP mapping to the TE tunnel. As the MPLS domain label allocation and distribution is performed downstream-on-demand, the label mapping to a TE LSP is first generated by the tailend router or egress Edge LSR and then propagated upstream. This process is repeated at each hop upstream where local labels mapping to a TE tunnel are assigned and propagated upstream until the headend router is reached.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

123. The MPLS-VPN Inter-AS system provided by Cisco IOS provides for controlling resources of the second network level by a second group of NRMs, wherein the first group and the second group of NRMs comprise means for communicating on a common network level, as described below:

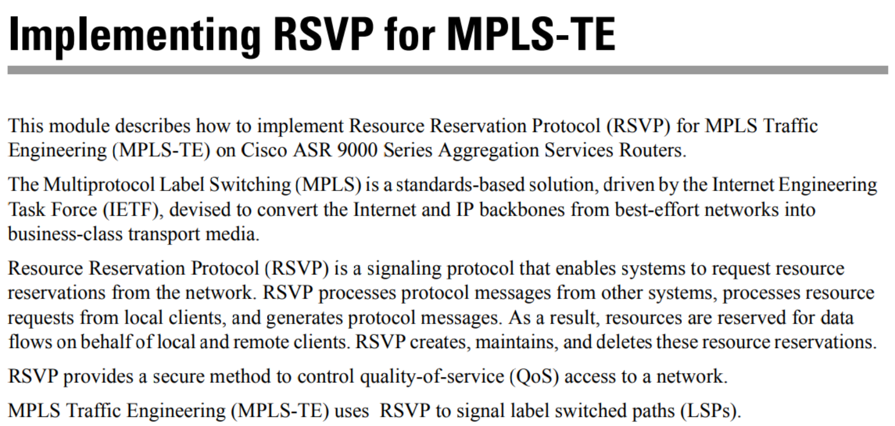

Implementing RSVP for MPLS-TE

This module describes how to implement Resource Reservation Protocol (RSVP) for MPLS Traffic Engineering (MPLS-TE) on Cisco ASR 9000 Series Aggregation Services Routers.

The Multiprotocol Label Switching (MPLS) is a standards-based solution, driven by the Internet Engineering Task Force (IETF), devised to convert the Internet and IP backbones from best-effort networks into business-class transport media.

Resource Reservation Protocol (RSVP) is a signaling protocol that enables systems to request resource reservations from the network. RSVP processes protocol messages from other systems, processes resource requests from local clients, and generates protocol messages. As a result, resources are reserved for data flows on behalf of local and remote clients. RSVP creates, maintains, and deletes these resource reservations.

RSVP provides a secure method to control quality-of-service (QoS) access to a network.

MPLS Traffic Engineering (MPLS-TE) uses RSVP to signal label switched paths (LSPs).

https://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.2/mpls/configuration/guide/b_mpls_cg42asr9k_chapter_010.pdf

Figure 1. MPLS VPN Inter-AS Option AB Topology

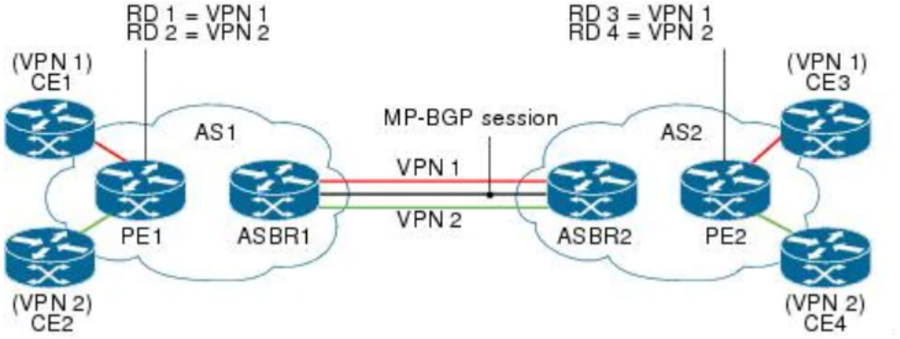

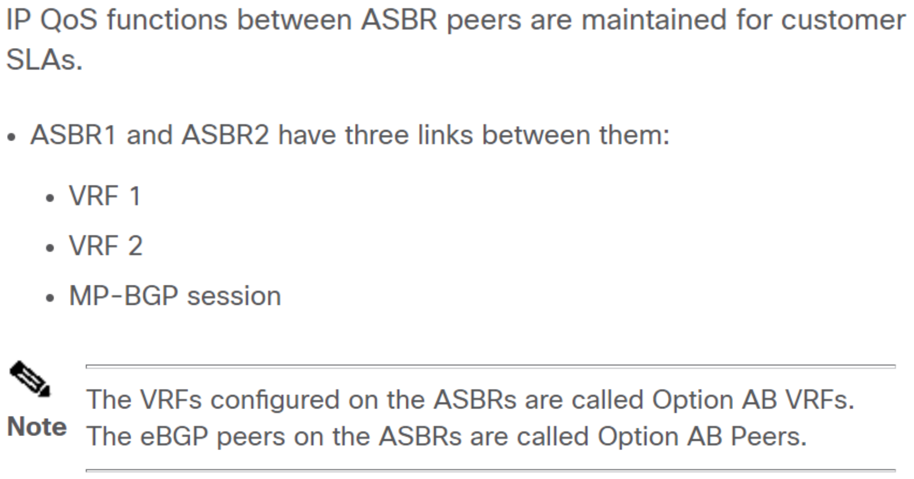

IP QoS functions between ASBR peers are maintained for customer SLAs.

- ASBR1 and ASBR2 have three links between them:
  - VRF 1
  - VRF 2
  - MP-BGP session

Note The VRFs configured on the ASBRs are called Option AB VRFs. The eBGP peers on the ASBRs are called Option AB Peers.

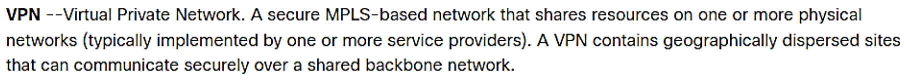

VPN --Virtual Private Network. A secure MPLS-based network that shares resources on one or more physical networks (typically implemented by one or more service providers). A VPN contains geographically dispersed sites that can communicate securely over a shared backbone network.

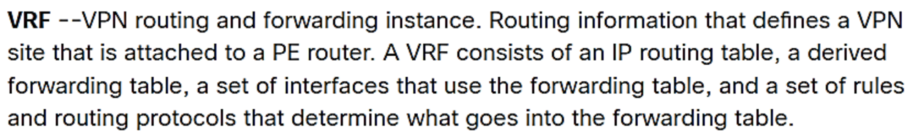

VRF --VPN routing and forwarding instance. Routing information that defines a VPN site that is attached to a PE router. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

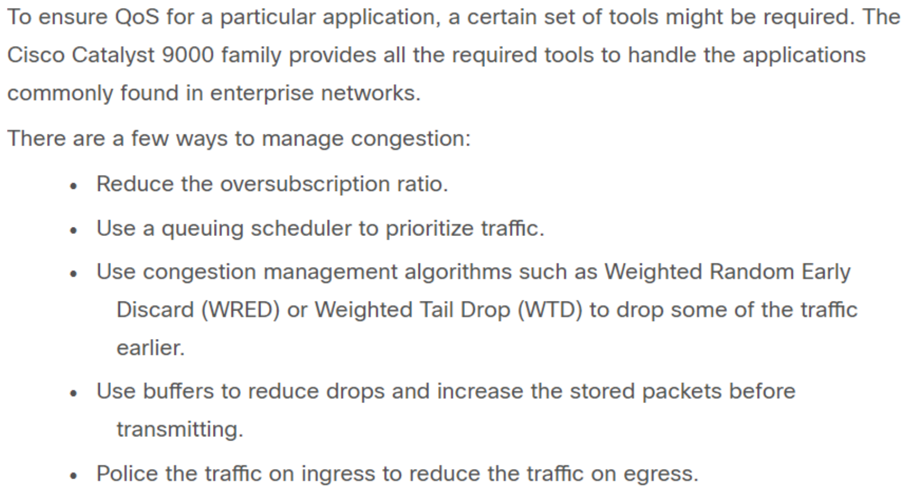

To ensure QoS for a particular application, a certain set of tools might be required. The Cisco Catalyst 9000 family provides all the required tools to handle the applications commonly found in enterprise networks.

There are a few ways to manage congestion:

- Reduce the oversubscription ratio.
- Use a queuing scheduler to prioritize traffic.
- Use congestion management algorithms such as Weighted Random Early Discard (WRED) or Weighted Tail Drop (WTD) to drop some of the traffic earlier.
- Use buffers to reduce drops and increase the stored packets before transmitting.
- Police the traffic on ingress to reduce the traffic on egress.

https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.html

124. As shown below, the Cisco IOS MPLS-VPN Inter-AS system provides for exchanging resource requests between the NRMs of the first and second groups using the first addressing scheme, the NRMs of the first group and the second group admitting new resource requests based at least in part on a total amount of available resources, an amount of resources currently reserved by previous reservations, and an amount of resources requested in the new resource requests:

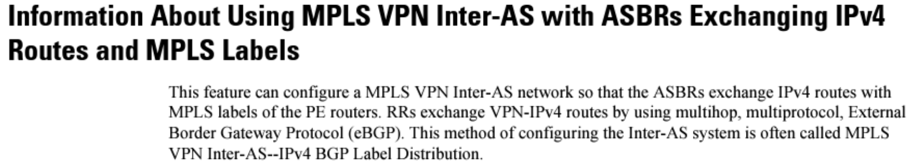

Information About Using MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels

This feature can configure a MPLS VPN Inter-AS network so that the ASBRs exchange IPv4 routes with MPLS labels of the PE routers. RRs exchange VPN-IPv4 routes by using multihop, multiprotocol, External Border Gateway Protocol (eBGP). This method of configuring the Inter-AS system is often called MPLS VPN Inter-AS--IPv4 BGP Label Distribution.

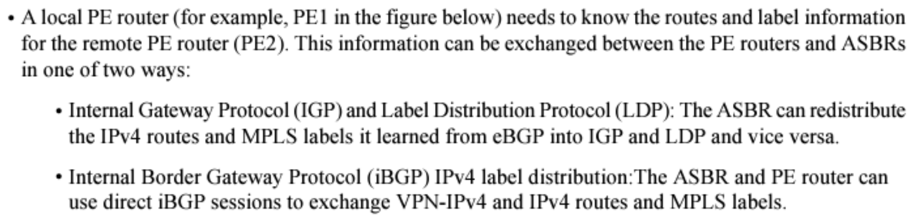

- A local PE router (for example, PE1 in the figure below) needs to know the routes and label information for the remote PE router (PE2). This information can be exchanged between the PE routers and ASBRs in one of two ways:
  - Internal Gateway Protocol (IGP) and Label Distribution Protocol (LDP): The ASBR can redistribute the IPv4 routes and MPLS labels it learned from eBGP into IGP and LDP and vice versa.
  - Internal Border Gateway Protocol (iBGP) IPv4 label distribution:The ASBR and PE router can use direct iBGP sessions to exchange VPN-IPv4 and IPv4 routes and MPLS labels.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-16-6/mp-ias-and-csc-xe-16-6-book/mpls-vpn-inter-as-with-asbrs-exchanging-ipv4-routes-and-mpls-labels.pdf

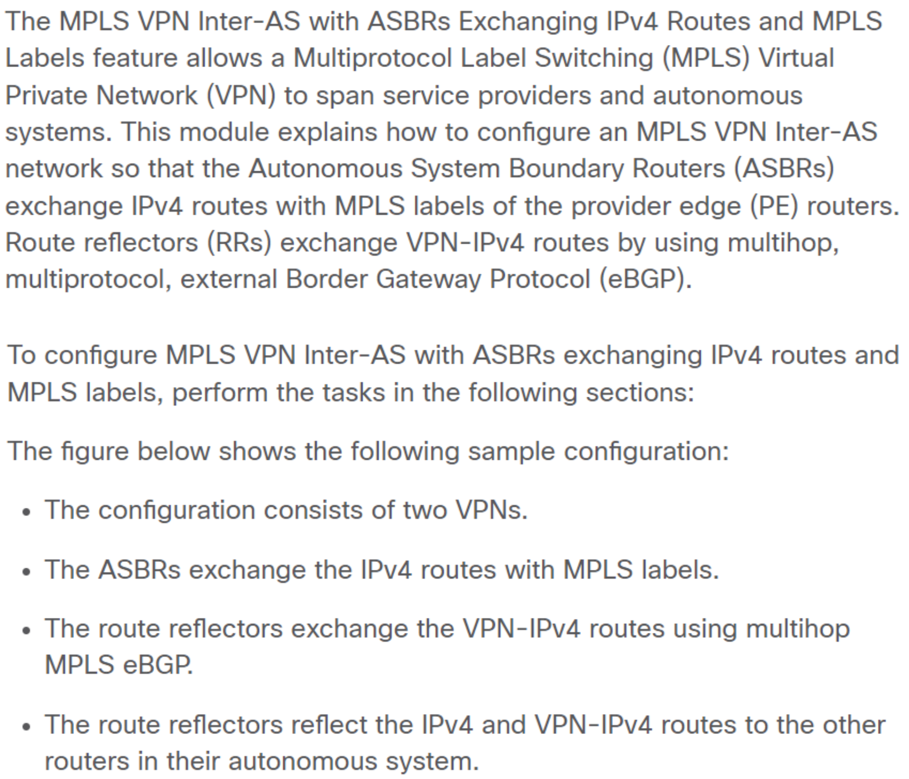

The MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels feature allows a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) to span service providers and autonomous systems. This module explains how to configure an MPLS VPN Inter-AS network so that the Autonomous System Boundary Routers (ASBRs) exchange IPv4 routes with MPLS labels of the provider edge (PE) routers. Route reflectors (RRs) exchange VPN-IPv4 routes by using multihop, multiprotocol, external Border Gateway Protocol (eBGP).

To configure MPLS VPN Inter-AS with ASBRs exchanging IPv4 routes and MPLS labels, perform the tasks in the following sections:

The figure below shows the following sample configuration:

- The configuration consists of two VPNs.
- The ASBRs exchange the IPv4 routes with MPLS labels.
- The route reflectors exchange the VPN-IPv4 routes using multihop MPLS eBGP.
- The route reflectors reflect the IPv4 and VPN-IPv4 routes to the other routers in their autonomous system.

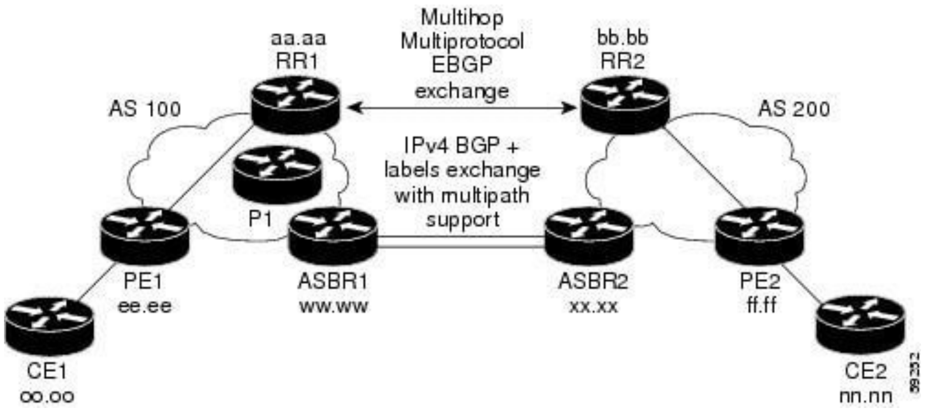

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-16/mp-ias-and-csc-xe-16-book/mpls-vpn-inter-as-with-asbrs-exchanging-ipv4-routes-and-mpls-labels.html

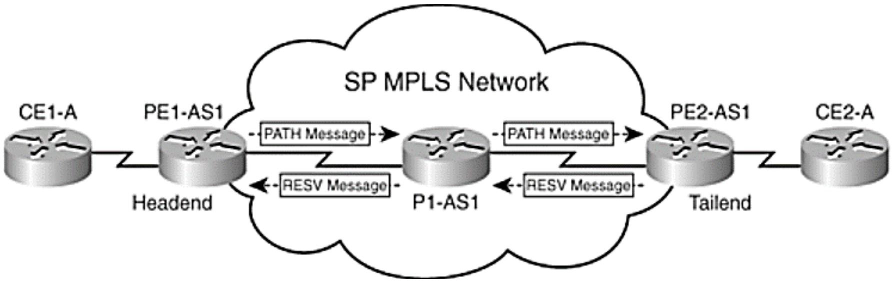

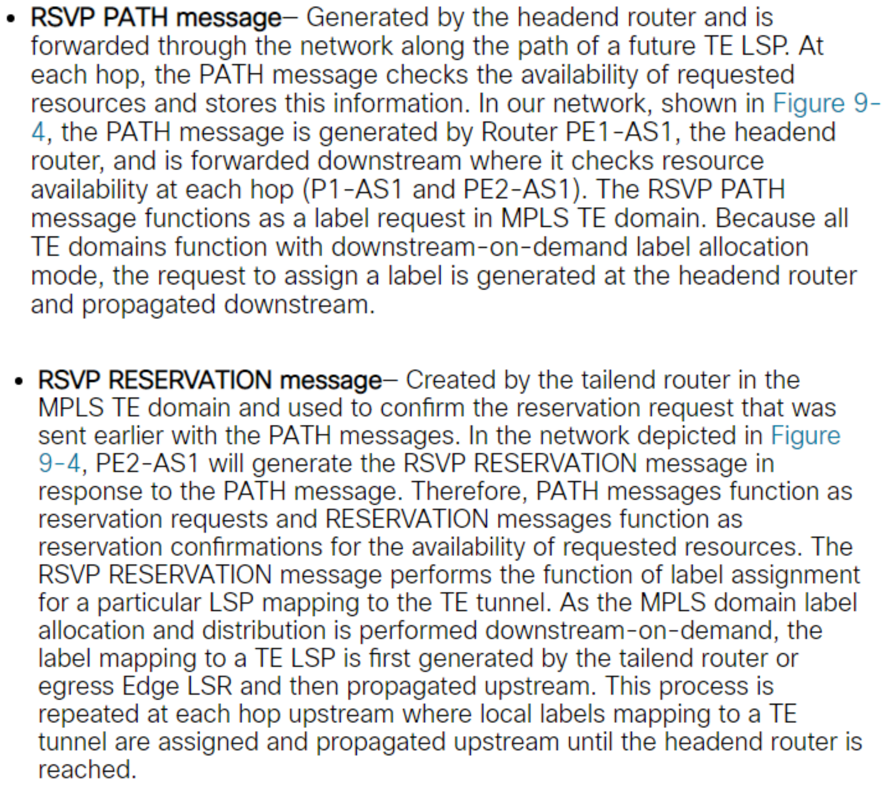

- **RSVP PATH message**– Generated by the headend router and is forwarded through the network along the path of a future TE LSP. At each hop, the PATH message checks the availability of requested resources and stores this information. In our network, shown in Figure 9-4, the PATH message is generated by Router PE1-AS1, the headend router, and is forwarded downstream where it checks resource availability at each hop (P1-AS1 and PE2-AS1). The RSVP PATH message functions as a label request in MPLS TE domain. Because all TE domains function with downstream-on-demand label allocation mode, the request to assign a label is generated at the headend router and propagated downstream.

- **RSVP RESERVATION message**– Created by the tailend router in the MPLS TE domain and used to confirm the reservation request that was sent earlier with the PATH messages. In the network depicted in Figure 9-4, PE2-AS1 will generate the RSVP RESERVATION message in response to the PATH message. Therefore, PATH messages function as reservation requests and RESERVATION messages function as reservation confirmations for the availability of requested resources. The RSVP RESERVATION message performs the function of label assignment for a particular LSP mapping to the TE tunnel. As the MPLS domain label allocation and distribution is performed downstream-on-demand, the label mapping to a TE LSP is first generated by the tailend router or egress Edge LSR and then propagated upstream. This process is repeated at each hop upstream where local labels mapping to a TE tunnel are assigned and propagated upstream until the headend router is reached.

https://www.ciscopress.com/articles/article.asp?p=426640&seqNum=2

125. As described below, the Cisco MPLS-VPN Inter-AS system includes performing an address mapping between the first and second addressing schemes so that a set of resources that is used by a reservation in the second group, controlled and known by the second group, is aggregated into a single resource in the first group of NRMs:

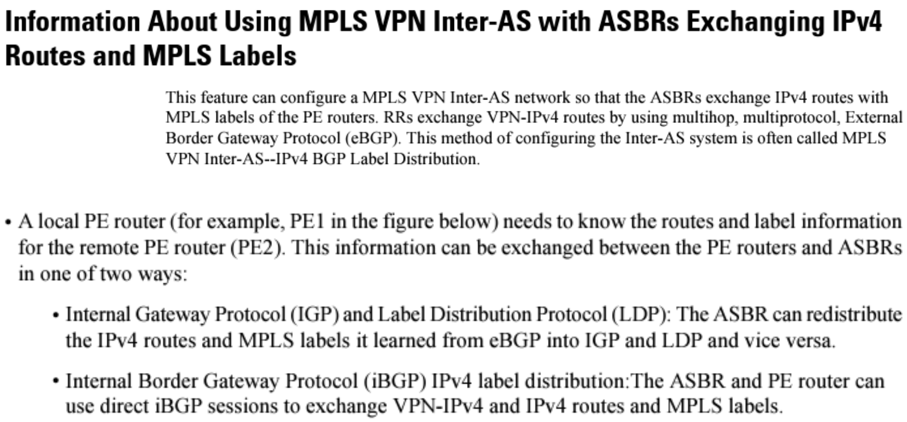

Information About Using MPLS VPN Inter-AS with ASBRs Exchanging IPv4 Routes and MPLS Labels

This feature can configure a MPLS VPN Inter-AS network so that the ASBRs exchange IPv4 routes with MPLS labels of the PE routers. RRs exchange VPN-IPv4 routes by using multihop, multiprotocol, External Border Gateway Protocol (eBGP). This method of configuring the Inter-AS system is often called MPLS VPN Inter-AS--IPv4 BGP Label Distribution.

- A local PE router (for example, PE1 in the figure below) needs to know the routes and label information for the remote PE router (PE2). This information can be exchanged between the PE routers and ASBRs in one of two ways:
  - Internal Gateway Protocol (IGP) and Label Distribution Protocol (LDP): The ASBR can redistribute the IPv4 routes and MPLS labels it learned from eBGP into IGP and LDP and vice versa.
  - Internal Border Gateway Protocol (iBGP) IPv4 label distribution:The ASBR and PE router can use direct iBGP sessions to exchange VPN-IPv4 and IPv4 routes and MPLS labels.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-16-6/mp-ias-and-csc-xe-16-6-book/mpls-vpn-inter-as-with-asbrs-exchanging-ipv4-routes-and-mpls-labels.pdf

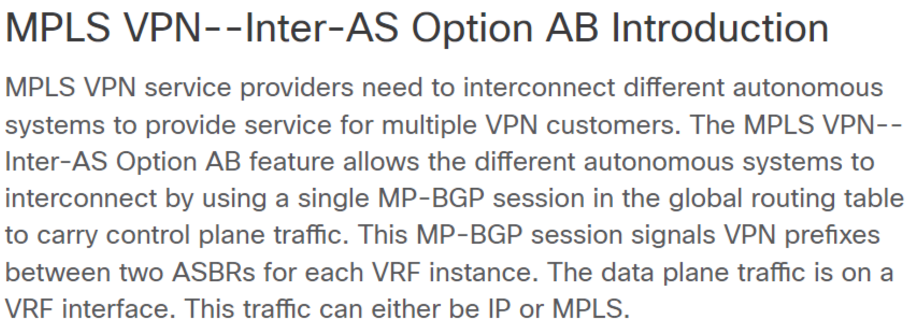

MPLS VPN--Inter-AS Option AB Introduction

MPLS VPN service providers need to interconnect different autonomous systems to provide service for multiple VPN customers. The MPLS VPN--Inter-AS Option AB feature allows the different autonomous systems to interconnect by using a single MP-BGP session in the global routing table to carry control plane traffic. This MP-BGP session signals VPN prefixes between two ASBRs for each VRF instance. The data plane traffic is on a VRF interface. This traffic can either be IP or MPLS.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/xe-17/mp-ias-and-csc-xe-17-book/mpls-vpn-inter-as-option-ab.html

126. Cisco markets, offers to sell, sells, and distributes the Cisco QoS Systems, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’966 Patent. The Cisco QoS Systems are not staple articles or commodities of commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

127. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the ’966 Patent.

Cisco's infringement of the '966 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney's fees, and costs in bringing this action.

128. Cisco's direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary and permanent injunctive relief and damages as a result of Cisco's infringement of the '966 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 6

(Infringement of U.S. Pat. No. 7,606,885)

129. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

130. Cisco has infringed and continues to infringe one or more claims of the '885 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the '885 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the '885 Patent and continues to contribute to the infringement of the '885 Patent in violation of 35 U.S.C. §271(c).

131. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 14 of the '885 Patent at least by making, using, offering to sell, importing, and/or selling Cisco Unified Communication as part of its Internetwork Operating System (IOS) versions 12 and above in conjunction with Cisco 2900 and 3900 series routers.

132. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell the Cisco Unified Communication and the Cisco Internetworking Operating System

(IOS) versions 12 and above (the “Cisco IOS”) in the Unites States, and encourages distributors to sell, offer to sell and use, and encourages Cisco’s customers to use, the Cisco Unified Communication and the Cisco IOS in the United States with knowledge that the Cisco Unified Communication and the Cisco IOS infringes the ’885 Patent.

133. Cisco has had knowledge of and notice of the ’885 Patent and its infringement since before the filing of the original Complaint. For example, on information and belief, Cisco had knowledge of and notice of the ’885 Patent as a result of a partnership with NetSocket during at least 2012. Cisco offers to sell the Cisco Unified Communication System in this District and does so with knowledge that the sale and use of the Cisco Unified Communication System infringes and with the intent for its customers to use the Cisco Unified Communication System in an infringing manner.

134. The Cisco Unified Communication System satisfies each of the limitations of at least claim 14 of the ’885 Patent.

135. For example, the Cisco Unified Communication System provides for handling performing resource management issues and admission control within an Internet Protocol IP telephony system adapted for transmission of multimedia over an IP network, the system comprising a topology aware resource manager (NRM) realized by a computer program deployed on a standalone server connected to the IP network and, via the IP network, connected to a gatekeeper, as described below:

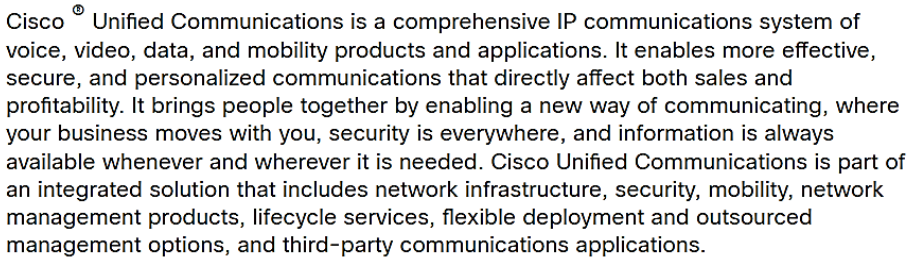

Cisco ® Unified Communications is a comprehensive IP communications system of voice, video, data, and mobility products and applications. It enables more effective, secure, and personalized communications that directly affect both sales and profitability. It brings people together by enabling a new way of communicating, where your business moves with you, security is everywhere, and information is always available whenever and wherever it is needed. Cisco Unified Communications is part of an integrated solution that includes network infrastructure, security, mobility, network management products, lifecycle services, flexible deployment and outsourced management options, and third-party communications applications.

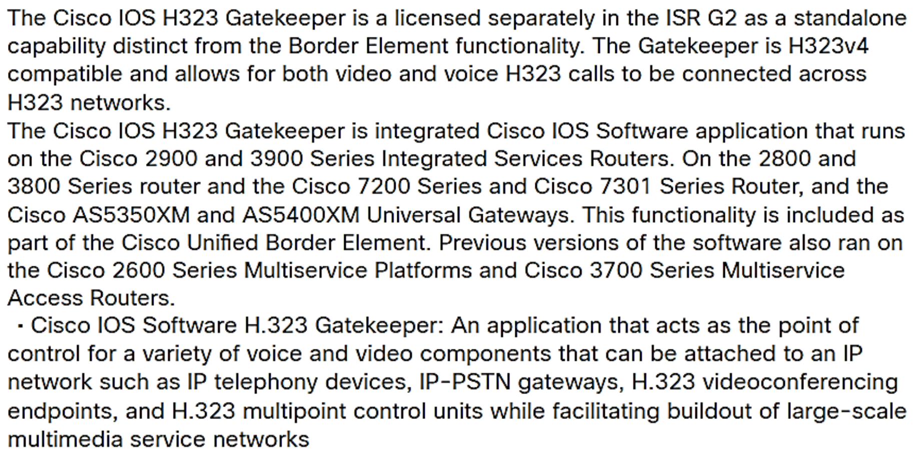

The Cisco IOS H323 Gatekeeper is a licensed separately in the ISR G2 as a standalone capability distinct from the Border Element functionality. The Gatekeeper is H323v4 compatible and allows for both video and voice H323 calls to be connected across H323 networks.

The Cisco IOS H323 Gatekeeper is integrated Cisco IOS Software application that runs on the Cisco 2900 and 3900 Series Integrated Services Routers. On the 2800 and 3800 Series router and the Cisco 7200 Series and Cisco 7301 Series Router, and the Cisco AS5350XM and AS5400XM Universal Gateways. This functionality is included as part of the Cisco Unified Border Element. Previous versions of the software also ran on the Cisco 2600 Series Multiservice Platforms and Cisco 3700 Series Multiservice Access Routers.

- Cisco IOS Software H.323 Gatekeeper: An application that acts as the point of control for a variety of voice and video components that can be attached to an IP network such as IP telephony devices, IP-PSTN gateways, H.323 videoconferencing endpoints, and H.323 multipoint control units while facilitating buildout of large-scale multimedia service networks

https://www.cisco.com/c/en/us/products/collateral/unified-communications/gatekeeper-multimedia-conference-manager/data_sheet_c78_561921.html

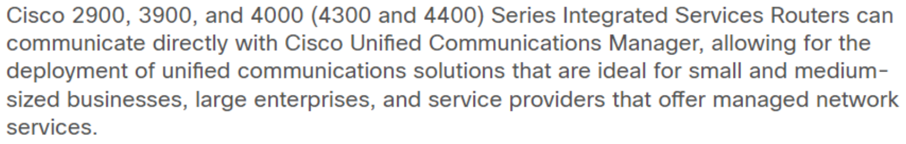

Cisco 2900, 3900, and 4000 (4300 and 4400) Series Integrated Services Routers can communicate directly with Cisco Unified Communications Manager, allowing for the deployment of unified communications solutions that are ideal for small and medium-sized businesses, large enterprises, and service providers that offer managed network services.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

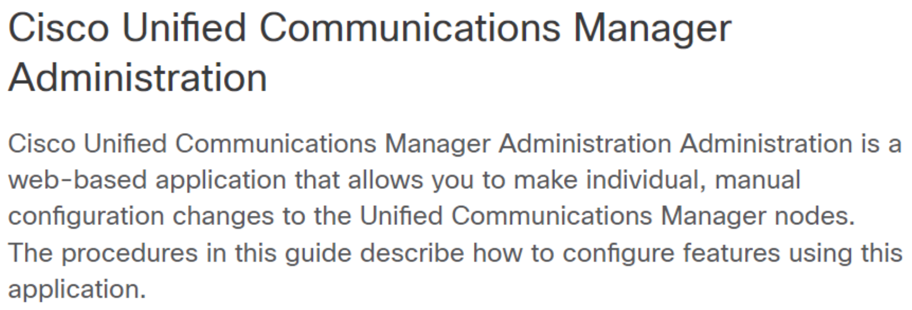

Cisco Unified Communications Manager Administration

Cisco Unified Communications Manager Administration Administration is a web-based application that allows you to make individual, manual configuration changes to the Unified Communications Manager nodes. The procedures in this guide describe how to configure features using this application.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/11_5_1/sysConfig/CUCM_BK_SE5DAF88_00_cucm-system-configuration-guide-1151/CUCM_BK_SE5DAF88_00_cucm-system-configuration-guide-1151_chapter_00.html

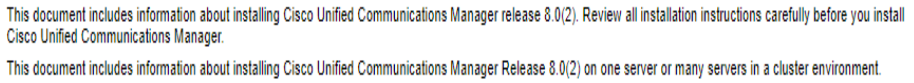

This document includes information about installing Cisco Unified Communications Manager release 8.0(2). Review all installation instructions carefully before you install Cisco Unified Communications Manager.

This document includes information about installing Cisco Unified Communications Manager Release 8.0(2) on one server or many servers in a cluster environment.

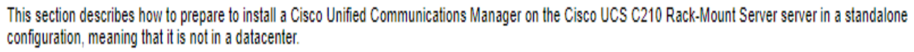

This section describes how to prepare to install a Cisco Unified Communications Manager on the Cisco UCS C210 Rack-Mount Server server in a standalone configuration, meaning that it is not in a datacenter.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/install/8_0_2/install/cmins802.html

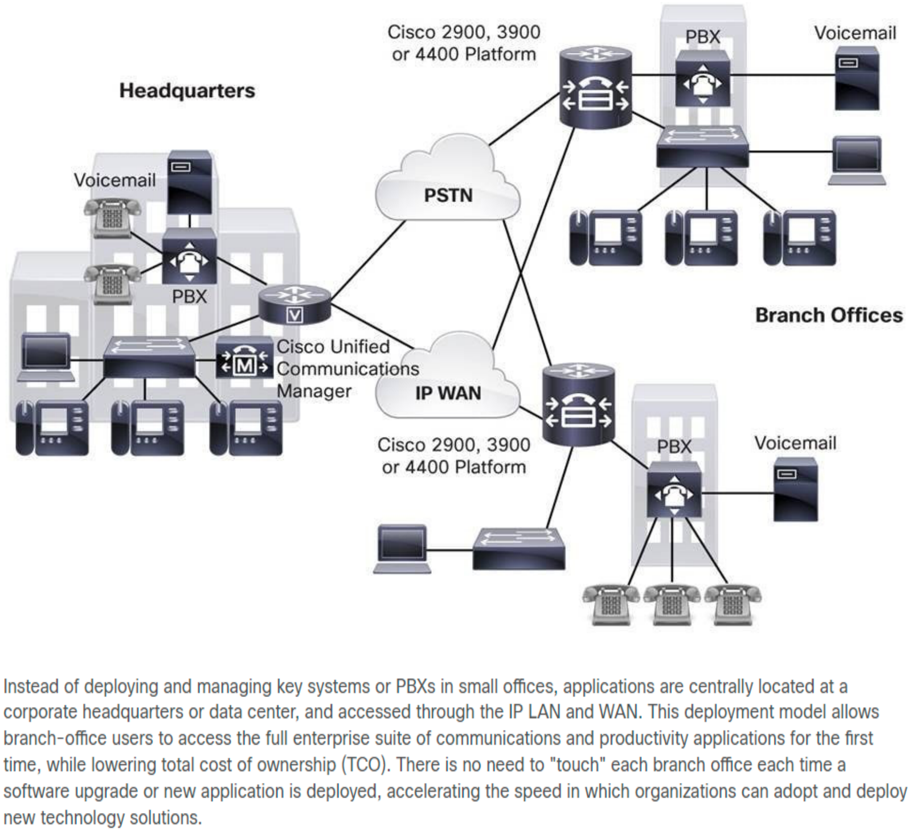
Headquarters
Voicemail
PBX
Cisco Unified Communications Manager
Cisco 2900, 3900 or 4400 Platform
PSTN
IP WAN
PBX
Voicemail
Branch Offices
Cisco 2900, 3900 or 4400 Platform
PBX
Voicemail

Instead of deploying and managing key systems or PBXs in small offices, applications are centrally located at a corporate headquarters or data center, and accessed through the IP LAN and WAN. This deployment model allows branch-office users to access the full enterprise suite of communications and productivity applications for the first time, while lowering total cost of ownership (TCO). There is no need to "touch" each branch office each time a software upgrade or new application is deployed, accelerating the speed in which organizations can adopt and deploy new technology solutions.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

136. The Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, provides for collecting routing information concerning the IP network using measurements including at least one of link-state routing, trace route and the Simple Network Management Protocol (SNMP) auto discovery by the resource manager (RM), as shown below:

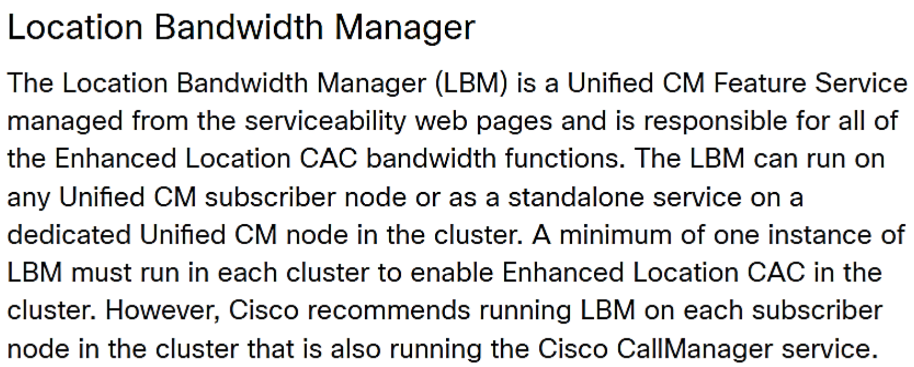

Location Bandwidth Manager

The Location Bandwidth Manager (LBM) is a Unified CM Feature Service managed from the serviceability web pages and is responsible for all of the Enhanced Location CAC bandwidth functions. The LBM can run on any Unified CM subscriber node or as a standalone service on a dedicated Unified CM node in the cluster. A minimum of one instance of LBM must run in each cluster to enable Enhanced Location CAC in the cluster. However, Cisco recommends running LBM on each subscriber node in the cluster that is also running the Cisco CallManager service.

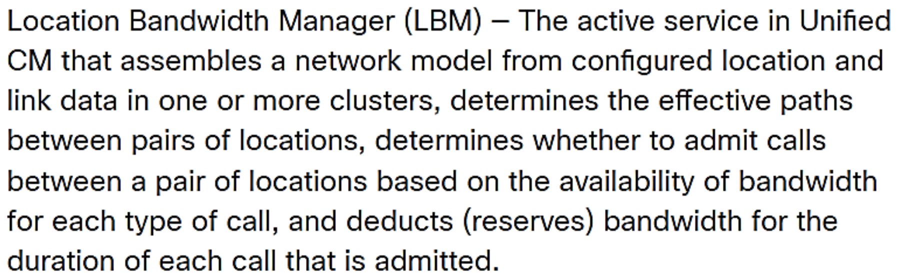

Location Bandwidth Manager (LBM) – The active service in Unified CM that assembles a network model from configured location and link data in one or more clusters, determines the effective paths between pairs of locations, determines whether to admit calls between a pair of locations based on the availability of bandwidth for each type of call, and deducts (reserves) bandwidth for the duration of each call that is admitted.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab11/collab11/cac.html

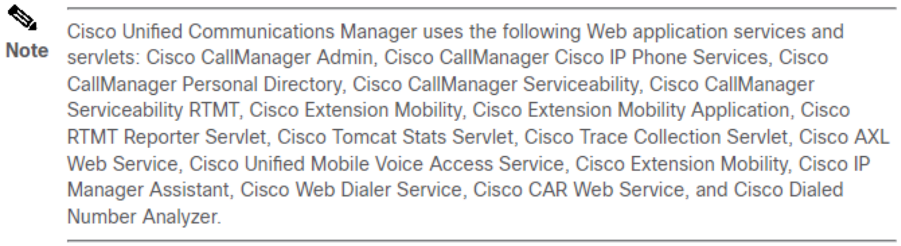

Note Cisco Unified Communications Manager uses the following Web application services and servlets: Cisco CallManager Admin, Cisco CallManager Cisco IP Phone Services, Cisco CallManager Personal Directory, Cisco CallManager Serviceability, Cisco CallManager Serviceability RTMT, Cisco Extension Mobility, Cisco Extension Mobility Application, Cisco RTMT Reporter Servlet, Cisco Tomcat Stats Servlet, Cisco Trace Collection Servlet, Cisco AXL Web Service, Cisco Unified Mobile Voice Access Service, Cisco Extension Mobility, Cisco IP Manager Assistant, Cisco Web Dialer Service, Cisco CAR Web Service, and Cisco Dialed Number Analyzer.

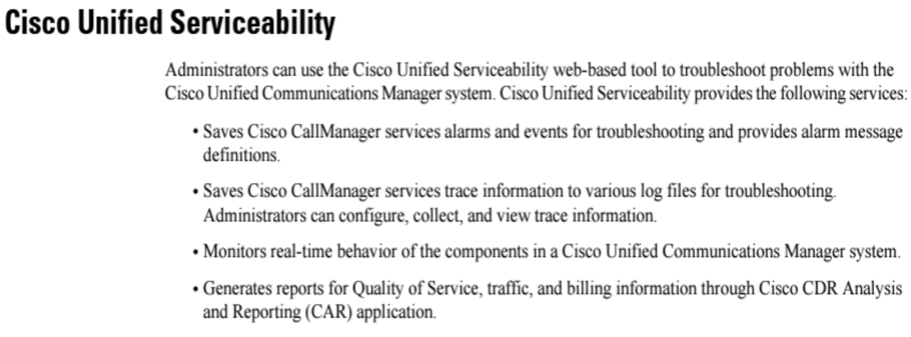

Cisco Unified Serviceability

Administrators can use the Cisco Unified Serviceability web-based tool to troubleshoot problems with the Cisco Unified Communications Manager system. Cisco Unified Serviceability provides the following services:

- Saves Cisco CallManager services alarms and events for troubleshooting and provides alarm message definitions.
- Saves Cisco CallManager services trace information to various log files for troubleshooting. Administrators can configure, collect, and view trace information.
- Monitors real-time behavior of the components in a Cisco Unified Communications Manager system.
- Generates reports for Quality of Service, traffic, and billing information through Cisco CDR Analysis and Reporting (CAR) application.

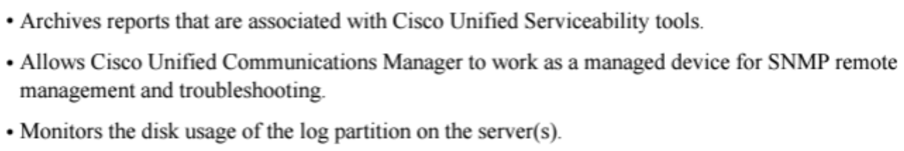

- Archives reports that are associated with Cisco Unified Serviceability tools.
- Allows Cisco Unified Communications Manager to work as a managed device for SNMP remote management and troubleshooting.
- Monitors the disk usage of the log partition on the server(s).

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/10_0_1/ccmsys/CUCM_BK_SE5FCFB6_00_cucm-system-guide-100.pdf

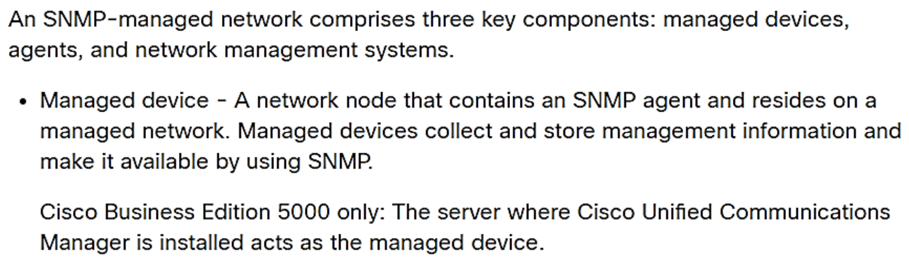

An SNMP-managed network comprises three key components: managed devices, agents, and network management systems.

- Managed device - A network node that contains an SNMP agent and resides on a managed network. Managed devices collect and store management information and make it available by using SNMP.

  Cisco Business Edition 5000 only: The server where Cisco Unified Communications Manager is installed acts as the managed device.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/11_5_1/Admin/CUCM_BK_CEF360A6_00_cisco-unified-serviceability-admin-guide_1151/CUCM_BK_CEF360A6_00_cisco-unified-serviceability-admin-guide_1151_chapter_0111.html

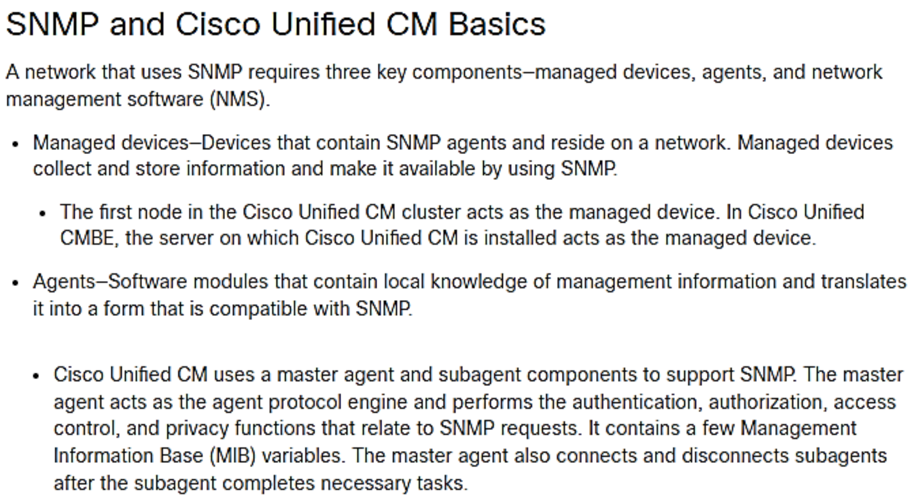

SNMP and Cisco Unified CM Basics

A network that uses SNMP requires three key components–managed devices, agents, and network management software (NMS).

- Managed devices–Devices that contain SNMP agents and reside on a network. Managed devices collect and store information and make it available by using SNMP.
  - The first node in the Cisco Unified CM cluster acts as the managed device. In Cisco Unified CMBE, the server on which Cisco Unified CM is installed acts as the managed device.
- Agents–Software modules that contain local knowledge of management information and translates it into a form that is compatible with SNMP.
  - Cisco Unified CM uses a master agent and subagent components to support SNMP. The master agent acts as the agent protocol engine and performs the authentication, authorization, access control, and privacy functions that relate to SNMP requests. It contains a few Management Information Base (MIB) variables. The master agent also connects and disconnects subagents after the subagent completes necessary tasks.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/service/10_0_1/rtmt/CUCM_BK_C8A0AF97_00_cucm-managed-service-guide-100/CUCM_BK_C8A0AF97_00_cucm-managed-service-guide-100_chapter_011.html

137. The Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, provides for obtaining resource information concerning

resources within the IP network, the resource information comprising available bandwidth of each resource, as described below:

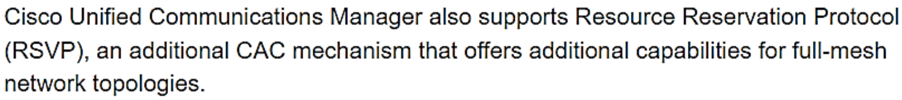

Cisco Unified Communications Manager also supports Resource Reservation Protocol (RSVP), an additional CAC mechanism that offers additional capabilities for full-mesh network topologies.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/10_0_1/ccmsys/CUCM_BK_SE5FCFB6_00_cucm-system-guide-100/CUCM_BK_SE5FCFB6_00_cucm-system-guide-100_chapter_01000.html

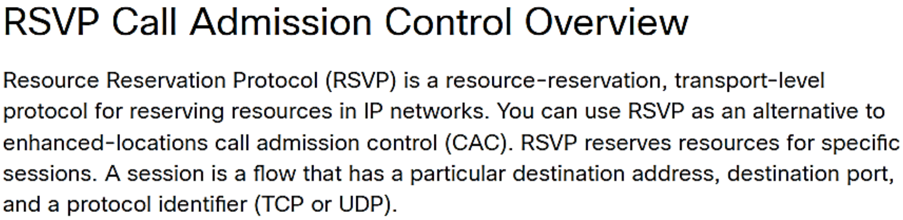

RSVP Call Admission Control Overview

Resource Reservation Protocol (RSVP) is a resource-reservation, transport-level protocol for reserving resources in IP networks. You can use RSVP as an alternative to enhanced-locations call admission control (CAC). RSVP reserves resources for specific sessions. A session is a flow that has a particular destination address, destination port, and a protocol identifier (TCP or UDP).

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/11_5_1/sysConfig/CUCM_BK_SE5DAF88_00_cucm-system-configuration-guide-1151/CUCM_BK_SE5DAF88_00_cucm-system-configuration-guide-1151_chapter_011110.html

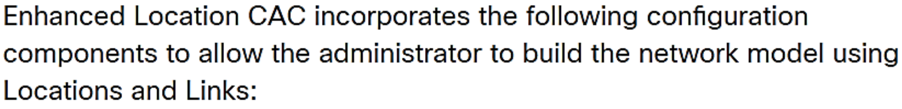

Enhanced Location CAC incorporates the following configuration components to allow the administrator to build the network model using Locations and Links:

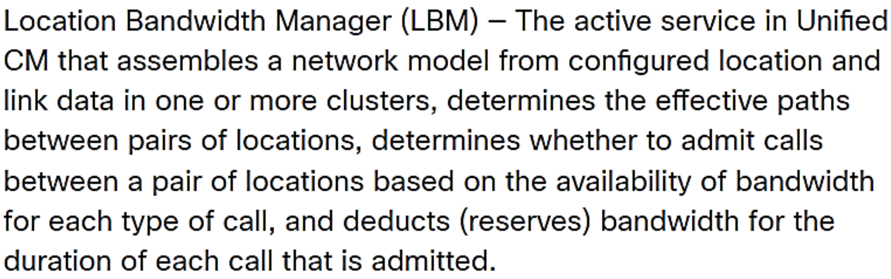

Location Bandwidth Manager (LBM) – The active service in Unified CM that assembles a network model from configured location and link data in one or more clusters, determines the effective paths between pairs of locations, determines whether to admit calls between a pair of locations based on the availability of bandwidth for each type of call, and deducts (reserves) bandwidth for the duration of each call that is admitted.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab11/collab11/cac.html

138. As shown below, the Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, provides for creating a resource map by means of combining routing information and resource information:

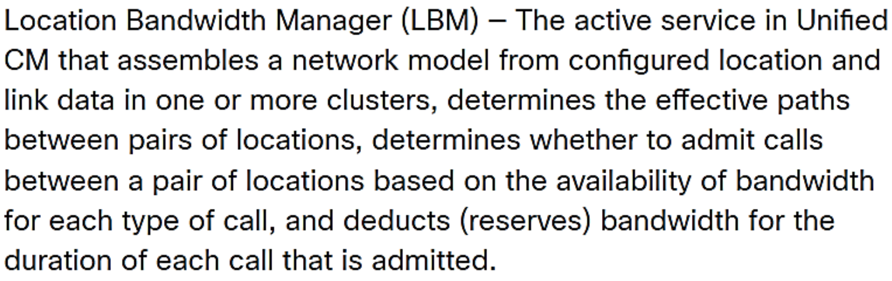

Location Bandwidth Manager (LBM) – The active service in Unified CM that assembles a network model from configured location and link data in one or more clusters, determines the effective paths between pairs of locations, determines whether to admit calls between a pair of locations based on the availability of bandwidth for each type of call, and deducts (reserves) bandwidth for the duration of each call that is admitted.

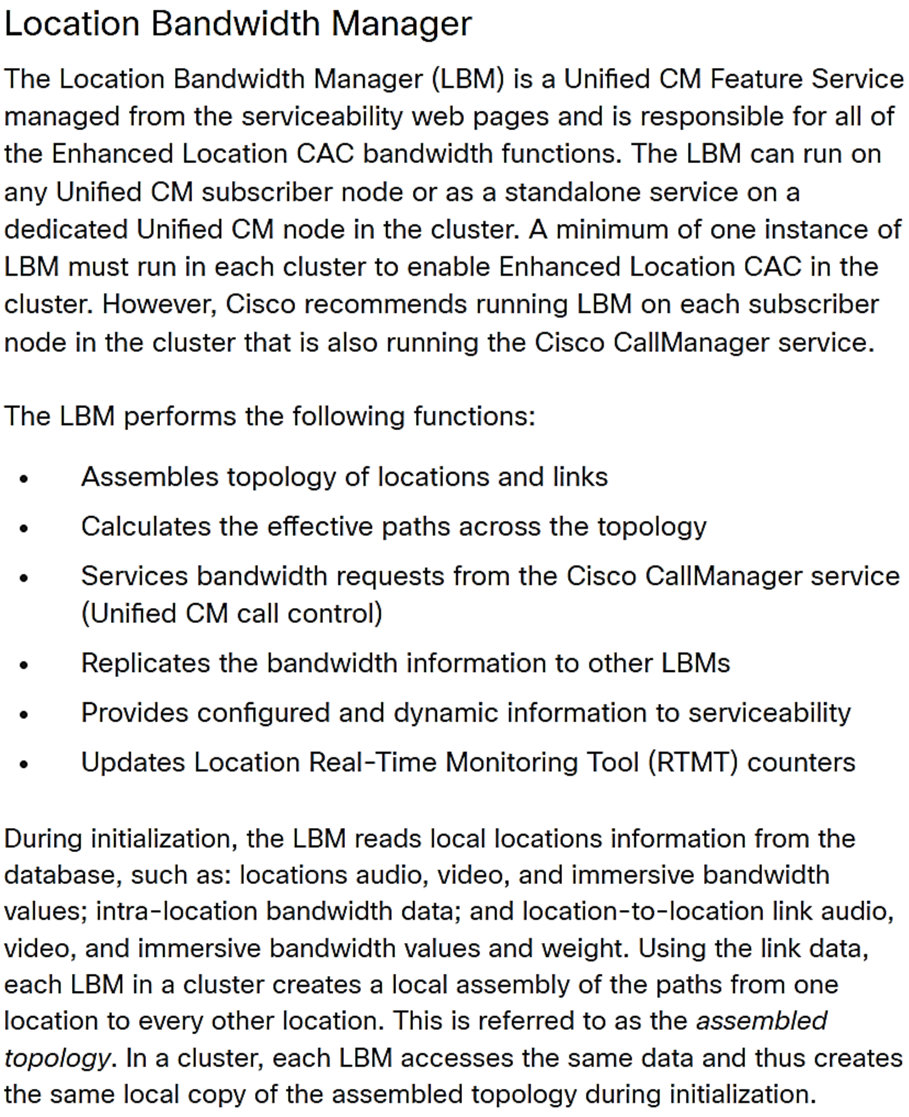

Location Bandwidth Manager

The Location Bandwidth Manager (LBM) is a Unified CM Feature Service managed from the serviceability web pages and is responsible for all of the Enhanced Location CAC bandwidth functions. The LBM can run on any Unified CM subscriber node or as a standalone service on a dedicated Unified CM node in the cluster. A minimum of one instance of LBM must run in each cluster to enable Enhanced Location CAC in the cluster. However, Cisco recommends running LBM on each subscriber node in the cluster that is also running the Cisco CallManager service.

The LBM performs the following functions:

- Assembles topology of locations and links
- Calculates the effective paths across the topology
- Services bandwidth requests from the Cisco CallManager service (Unified CM call control)
- Replicates the bandwidth information to other LBMs
- Provides configured and dynamic information to serviceability
- Updates Location Real-Time Monitoring Tool (RTMT) counters

During initialization, the LBM reads local locations information from the database, such as: locations audio, video, and immersive bandwidth values; intra-location bandwidth data; and location-to-location link audio, video, and immersive bandwidth values and weight. Using the link data, each LBM in a cluster creates a local assembly of the paths from one location to every other location. This is referred to as the *assembled topology*. In a cluster, each LBM accesses the same data and thus creates the same local copy of the assembled topology during initialization.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab11/collab11/cac.html

139. As described below, the Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, includes performing path-sensitive resource management issues and admission control within the system using the obtained

bandwidth information of the resource map and responsive to set-up requests received from a gatekeeper, wherein the collecting routing information concerning the IP network does not involve the gatekeeper:

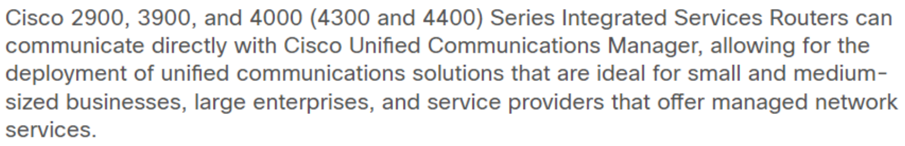

Cisco 2900, 3900, and 4000 (4300 and 4400) Series Integrated Services Routers can communicate directly with Cisco Unified Communications Manager, allowing for the deployment of unified communications solutions that are ideal for small and medium-sized businesses, large enterprises, and service providers that offer managed network services.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

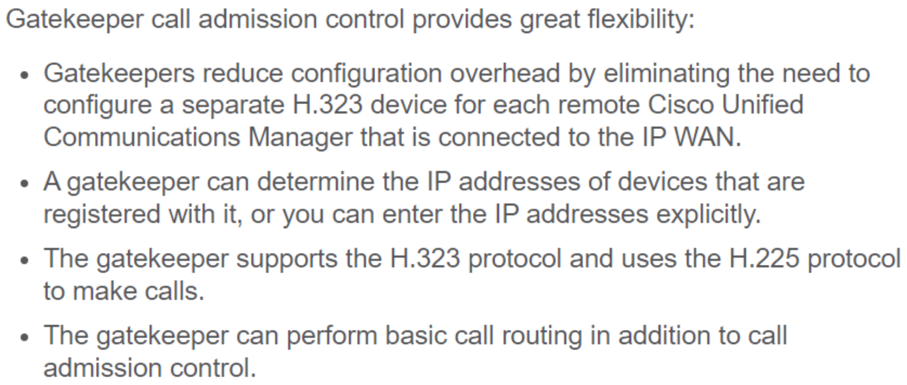

Gatekeeper call admission control provides great flexibility:

- Gatekeepers reduce configuration overhead by eliminating the need to configure a separate H.323 device for each remote Cisco Unified Communications Manager that is connected to the IP WAN.
- A gatekeeper can determine the IP addresses of devices that are registered with it, or you can enter the IP addresses explicitly.
- The gatekeeper supports the H.323 protocol and uses the H.225 protocol to make calls.
- The gatekeeper can perform basic call routing in addition to call admission control.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/admin/10_0_1/ccmsys/CUCM_BK_SE5FCFB6_00_cucm-system-guide-100/CUCM_BK_SE5FCFB6_00_cucm-system-guide-100_chapter_01000.html

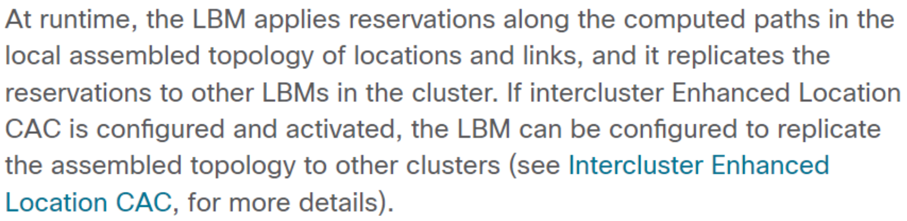

At runtime, the LBM applies reservations along the computed paths in the local assembled topology of locations and links, and it replicates the reservations to other LBMs in the cluster. If intercluster Enhanced Location CAC is configured and activated, the LBM can be configured to replicate the assembled topology to other clusters (see Intercluster Enhanced Location CAC, for more details).

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab11/collab11/cac.html

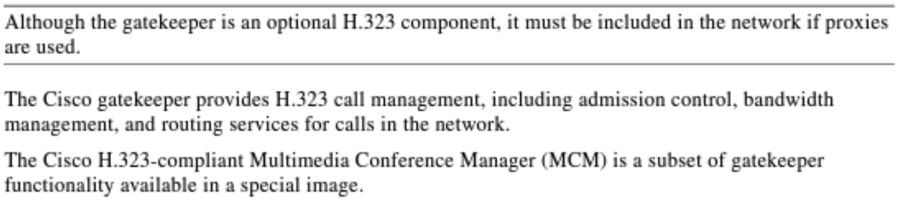

Although the gatekeeper is an optional H.323 component, it must be included in the network if proxies are used.

The Cisco gatekeeper provides H.323 call management, including admission control, bandwidth management, and routing services for calls in the network.

The Cisco H.323-compliant Multimedia Conference Manager (MCM) is a subset of gatekeeper functionality available in a special image.

https://www.cisco.com/c/en/us/td/docs/ios/voice/h323/configuration/guide/15_1/vh_15_1_book.pdf

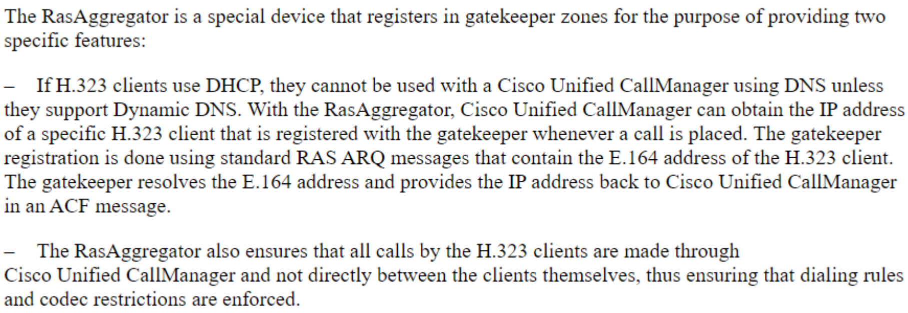

The RasAggregator is a special device that registers in gatekeeper zones for the purpose of providing two specific features:

– If H.323 clients use DHCP, they cannot be used with a Cisco Unified CallManager using DNS unless they support Dynamic DNS. With the RasAggregator, Cisco Unified CallManager can obtain the IP address of a specific H.323 client that is registered with the gatekeeper whenever a call is placed. The gatekeeper registration is done using standard RAS ARQ messages that contain the E.164 address of the H.323 client. The gatekeeper resolves the E.164 address and provides the IP address back to Cisco Unified CallManager in an ACF message.

– The RasAggregator also ensures that all calls by the H.323 clients are made through Cisco Unified CallManager and not directly between the clients themselves, thus ensuring that dialing rules and codec restrictions are enforced.

https://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/4x/42trunks.html

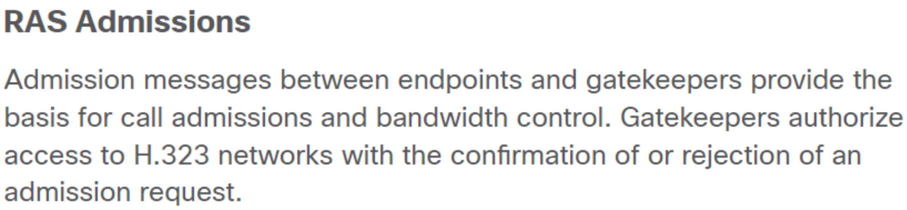

RAS Admissions

Admission messages between endpoints and gatekeepers provide the basis for call admissions and bandwidth control. Gatekeepers authorize access to H.323 networks with the confirmation of or rejection of an admission request.

https://www.cisco.com/c/en/us/support/docs/voice/h323/5244-understand-gatekeepers.html

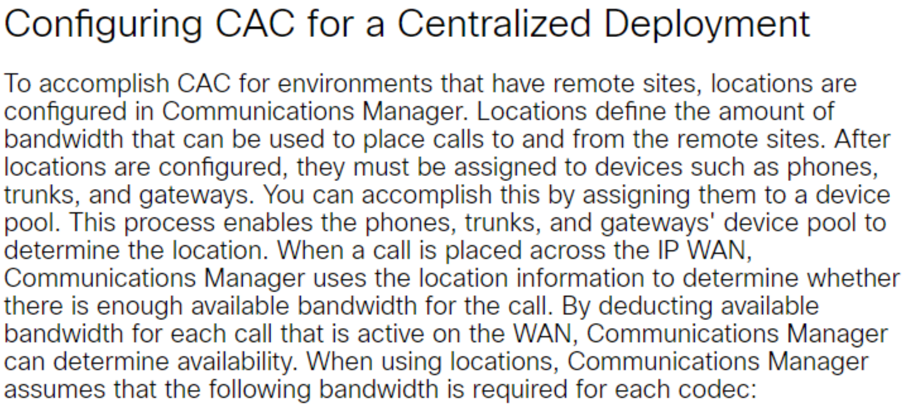

Configuring CAC for a Centralized Deployment

To accomplish CAC for environments that have remote sites, locations are configured in Communications Manager. Locations define the amount of bandwidth that can be used to place calls to and from the remote sites. After locations are configured, they must be assigned to devices such as phones, trunks, and gateways. You can accomplish this by assigning them to a device pool. This process enables the phones, trunks, and gateways' device pool to determine the location. When a call is placed across the IP WAN, Communications Manager uses the location information to determine whether there is enough available bandwidth for the call. By deducting available bandwidth for each call that is active on the WAN, Communications Manager can determine availability. When using locations, Communications Manager assumes that the following bandwidth is required for each codec:

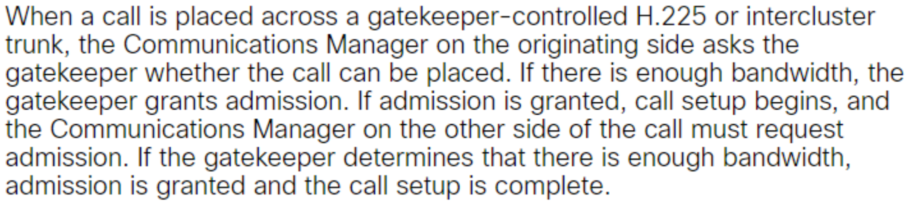

When a call is placed across a gatekeeper-controlled H.225 or intercluster trunk, the Communications Manager on the originating side asks the gatekeeper whether the call can be placed. If there is enough bandwidth, the gatekeeper grants admission. If admission is granted, call setup begins, and the Communications Manager on the other side of the call must request admission. If the gatekeeper determines that there is enough bandwidth, admission is granted and the call setup is complete.

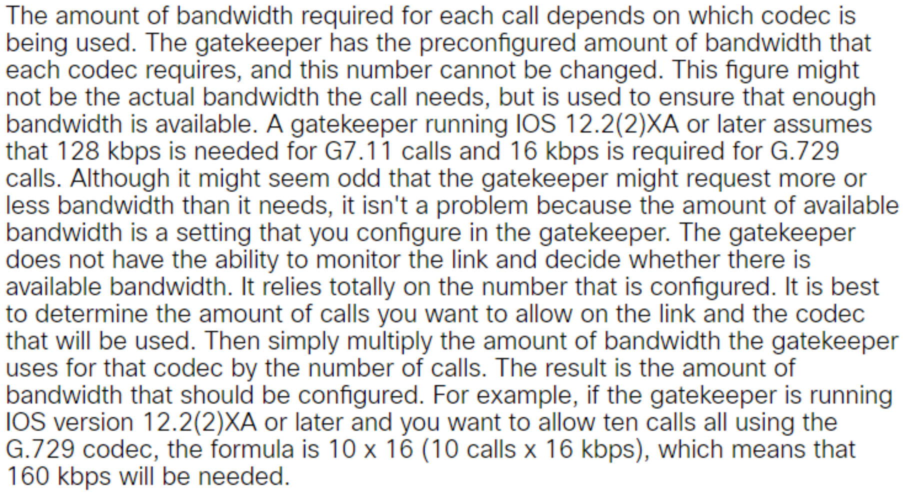

The amount of bandwidth required for each call depends on which codec is being used. The gatekeeper has the preconfigured amount of bandwidth that each codec requires, and this number cannot be changed. This figure might not be the actual bandwidth the call needs, but is used to ensure that enough bandwidth is available. A gatekeeper running IOS 12.2(2)XA or later assumes that 128 kbps is needed for G7.11 calls and 16 kbps is required for G.729 calls. Although it might seem odd that the gatekeeper might request more or less bandwidth than it needs, it isn't a problem because the amount of available bandwidth is a setting that you configure in the gatekeeper. The gatekeeper does not have the ability to monitor the link and decide whether there is available bandwidth. It relies totally on the number that is configured. It is best to determine the amount of calls you want to allow on the link and the codec that will be used. Then simply multiply the amount of bandwidth the gatekeeper uses for that codec by the number of calls. The result is the amount of bandwidth that should be configured. For example, if the gatekeeper is running IOS version 12.2(2)XA or later and you want to allow ten calls all using the G.729 codec, the formula is 10 x 16 (10 calls x 16 kbps), which means that 160 kbps will be needed.

https://www.ciscopress.com/articles/article.asp?p=1714577&seqNum=2

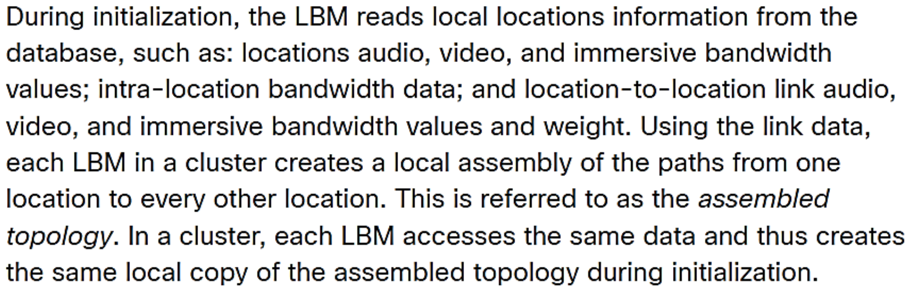

During initialization, the LBM reads local locations information from the database, such as: locations audio, video, and immersive bandwidth values; intra-location bandwidth data; and location-to-location link audio, video, and immersive bandwidth values and weight. Using the link data, each LBM in a cluster creates a local assembly of the paths from one location to every other location. This is referred to as the *assembled topology*. In a cluster, each LBM accesses the same data and thus creates the same local copy of the assembled topology during initialization.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/srnd/collab11/collab11/cac.html

140. As described below, the Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, provides for the resource manager in communication with the gatekeeper, performing, on a voice call link being initiated in response to voice call set-up requests received from the gatekeeper (Gk), the path-sensitive resource management issues and the admission control within the IP telephone system by using bandwidth

information from the resource map in a decision whether sufficient bandwidth resources are available for admitting a present call link being initiating:

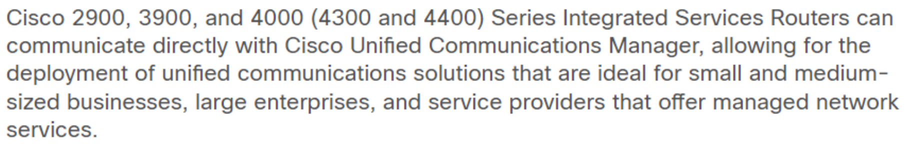

Cisco 2900, 3900, and 4000 (4300 and 4400) Series Integrated Services Routers can communicate directly with Cisco Unified Communications Manager, allowing for the deployment of unified communications solutions that are ideal for small and medium-sized businesses, large enterprises, and service providers that offer managed network services.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

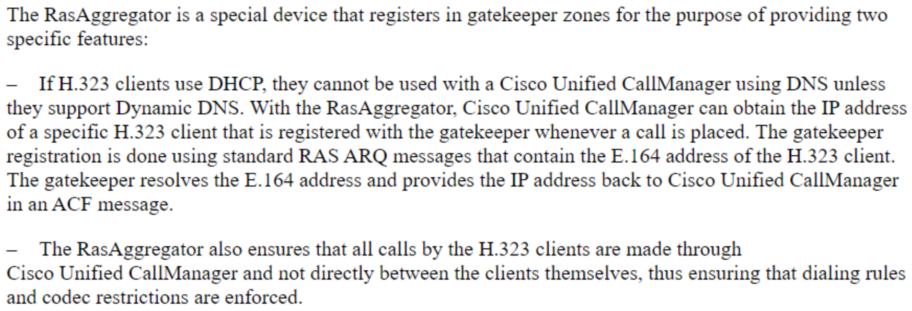

The RasAggregator is a special device that registers in gatekeeper zones for the purpose of providing two specific features:

– If H.323 clients use DHCP, they cannot be used with a Cisco Unified CallManager using DNS unless they support Dynamic DNS. With the RasAggregator, Cisco Unified CallManager can obtain the IP address of a specific H.323 client that is registered with the gatekeeper whenever a call is placed. The gatekeeper registration is done using standard RAS ARQ messages that contain the E.164 address of the H.323 client. The gatekeeper resolves the E.164 address and provides the IP address back to Cisco Unified CallManager in an ACF message.

– The RasAggregator also ensures that all calls by the H.323 clients are made through Cisco Unified CallManager and not directly between the clients themselves, thus ensuring that dialing rules and codec restrictions are enforced.

https://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/4x/42trunks.html

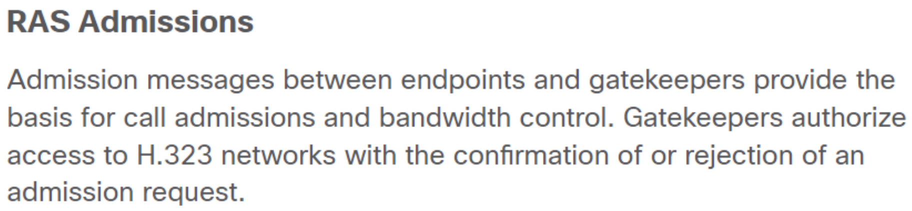

RAS Admissions

Admission messages between endpoints and gatekeepers provide the basis for call admissions and bandwidth control. Gatekeepers authorize access to H.323 networks with the confirmation of or rejection of an admission request.

https://www.cisco.com/c/en/us/support/docs/voice/h323/5244-understand-gatekeepers.html

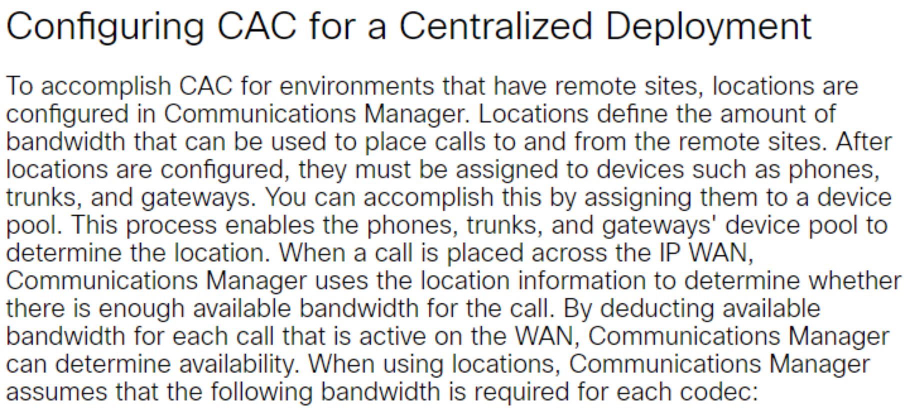

Configuring CAC for a Centralized Deployment

To accomplish CAC for environments that have remote sites, locations are configured in Communications Manager. Locations define the amount of bandwidth that can be used to place calls to and from the remote sites. After locations are configured, they must be assigned to devices such as phones, trunks, and gateways. You can accomplish this by assigning them to a device pool. This process enables the phones, trunks, and gateways' device pool to determine the location. When a call is placed across the IP WAN, Communications Manager uses the location information to determine whether there is enough available bandwidth for the call. By deducting available bandwidth for each call that is active on the WAN, Communications Manager can determine availability. When using locations, Communications Manager assumes that the following bandwidth is required for each codec:

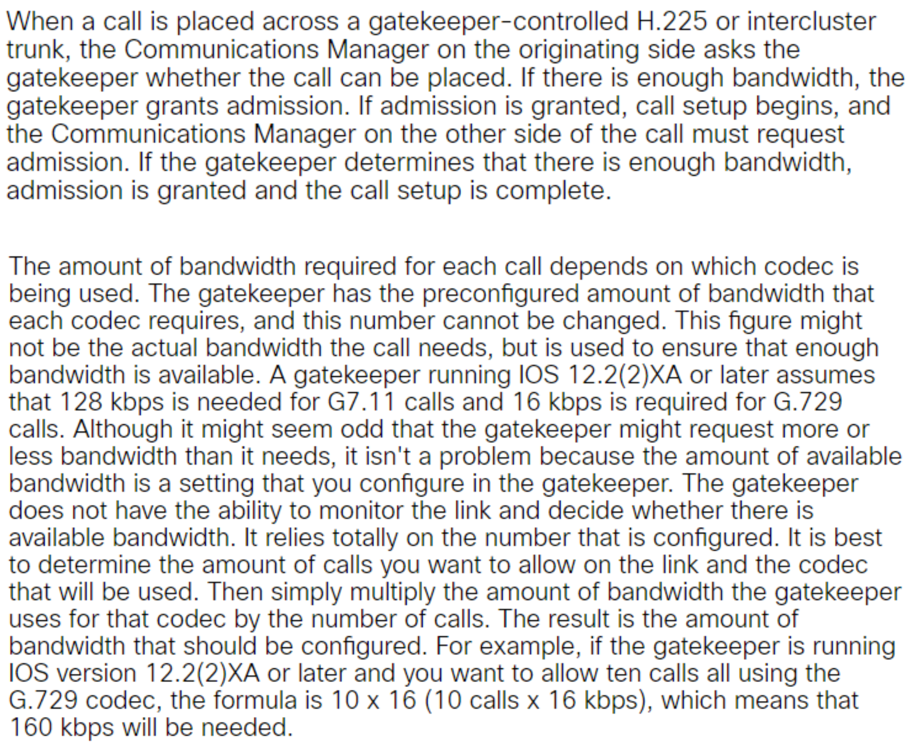

When a call is placed across a gatekeeper-controlled H.225 or intercluster trunk, the Communications Manager on the originating side asks the gatekeeper whether the call can be placed. If there is enough bandwidth, the gatekeeper grants admission. If admission is granted, call setup begins, and the Communications Manager on the other side of the call must request admission. If the gatekeeper determines that there is enough bandwidth, admission is granted and the call setup is complete.

The amount of bandwidth required for each call depends on which codec is being used. The gatekeeper has the preconfigured amount of bandwidth that each codec requires, and this number cannot be changed. This figure might not be the actual bandwidth the call needs, but is used to ensure that enough bandwidth is available. A gatekeeper running IOS 12.2(2)XA or later assumes that 128 kbps is needed for G7.11 calls and 16 kbps is required for G.729 calls. Although it might seem odd that the gatekeeper might request more or less bandwidth than it needs, it isn't a problem because the amount of available bandwidth is a setting that you configure in the gatekeeper. The gatekeeper does not have the ability to monitor the link and decide whether there is available bandwidth. It relies totally on the number that is configured. It is best to determine the amount of calls you want to allow on the link and the codec that will be used. Then simply multiply the amount of bandwidth the gatekeeper uses for that codec by the number of calls. The result is the amount of bandwidth that should be configured. For example, if the gatekeeper is running IOS version 12.2(2)XA or later and you want to allow ten calls all using the G.729 codec, the formula is 10 x 16 (10 calls x 16 kbps), which means that 160 kbps will be needed.

https://www.ciscopress.com/articles/article.asp?p=1714577&seqNum=2

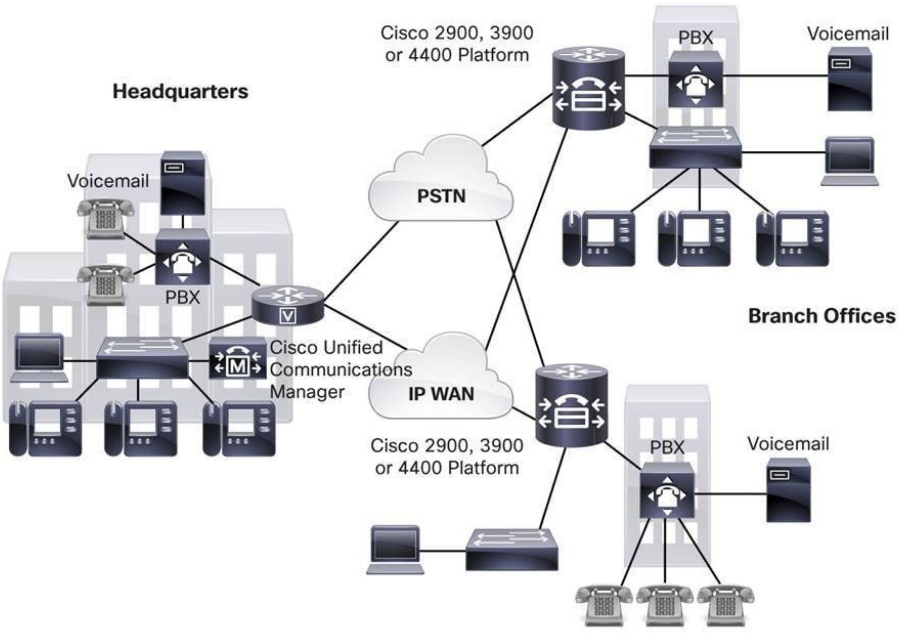
Headquarters
Voicemail
PBX
Cisco Unified Communications Manager
PSTN
IP WAN
Cisco 2900, 3900 or 4400 Platform
PBX
Voicemail
Branch Offices
Cisco 2900, 3900 or 4400 Platform
PBX
Voicemail

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

141. As described below, the Cisco Unified Communication system provided by Cisco IOS, in conjunction with Cisco 2900 and 3900 series routes, includes the gatekeeper residing on one of i) a second server connected to the IP network and ii) a router of the IP network and connected to the resource manager via the IP network, such that, via a second connection of the first server and via the IP network, the gatekeeper is connected to the first server:

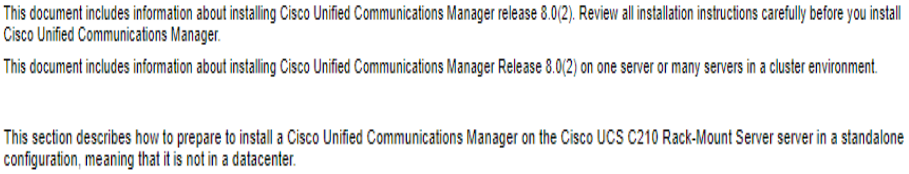

This document includes information about installing Cisco Unified Communications Manager release 8.0(2). Review all installation instructions carefully before you install Cisco Unified Communications Manager.

This document includes information about installing Cisco Unified Communications Manager Release 8.0(2) on one server or many servers in a cluster environment.

This section describes how to prepare to install a Cisco Unified Communications Manager on the Cisco UCS C210 Rack-Mount Server server in a standalone configuration, meaning that it is not in a datacenter.

https://www.cisco.com/c/en/us/td/docs/voice_ip_comm/cucm/install/8_0_2/install/cmins802.html

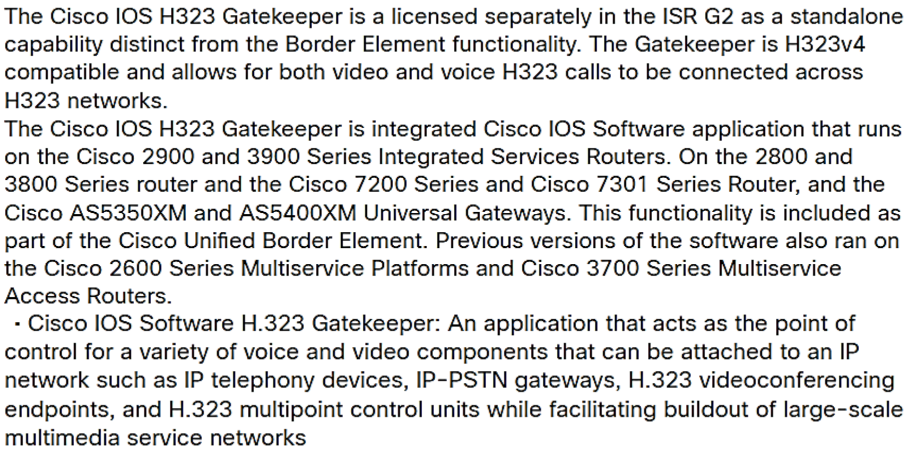

The Cisco IOS H323 Gatekeeper is a licensed separately in the ISR G2 as a standalone capability distinct from the Border Element functionality. The Gatekeeper is H323v4 compatible and allows for both video and voice H323 calls to be connected across H323 networks.
The Cisco IOS H323 Gatekeeper is integrated Cisco IOS Software application that runs on the Cisco 2900 and 3900 Series Integrated Services Routers. On the 2800 and 3800 Series router and the Cisco 7200 Series and Cisco 7301 Series Router, and the Cisco AS5350XM and AS5400XM Universal Gateways. This functionality is included as part of the Cisco Unified Border Element. Previous versions of the software also ran on the Cisco 2600 Series Multiservice Platforms and Cisco 3700 Series Multiservice Access Routers.

- Cisco IOS Software H.323 Gatekeeper: An application that acts as the point of control for a variety of voice and video components that can be attached to an IP network such as IP telephony devices, IP-PSTN gateways, H.323 videoconferencing endpoints, and H.323 multipoint control units while facilitating buildout of large-scale multimedia service networks

https://www.cisco.com/c/en/us/products/collateral/unified-communications/gatekeeper-multimedia-conference-manager/data_sheet_c78_561921.html

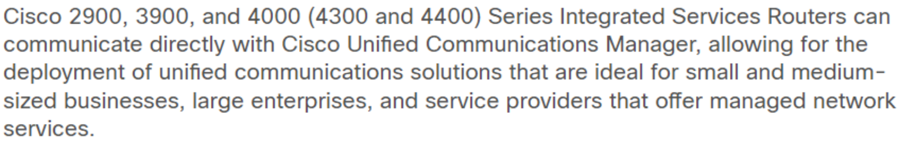

Cisco 2900, 3900, and 4000 (4300 and 4400) Series Integrated Services Routers can communicate directly with Cisco Unified Communications Manager, allowing for the deployment of unified communications solutions that are ideal for small and medium-sized businesses, large enterprises, and service providers that offer managed network services.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

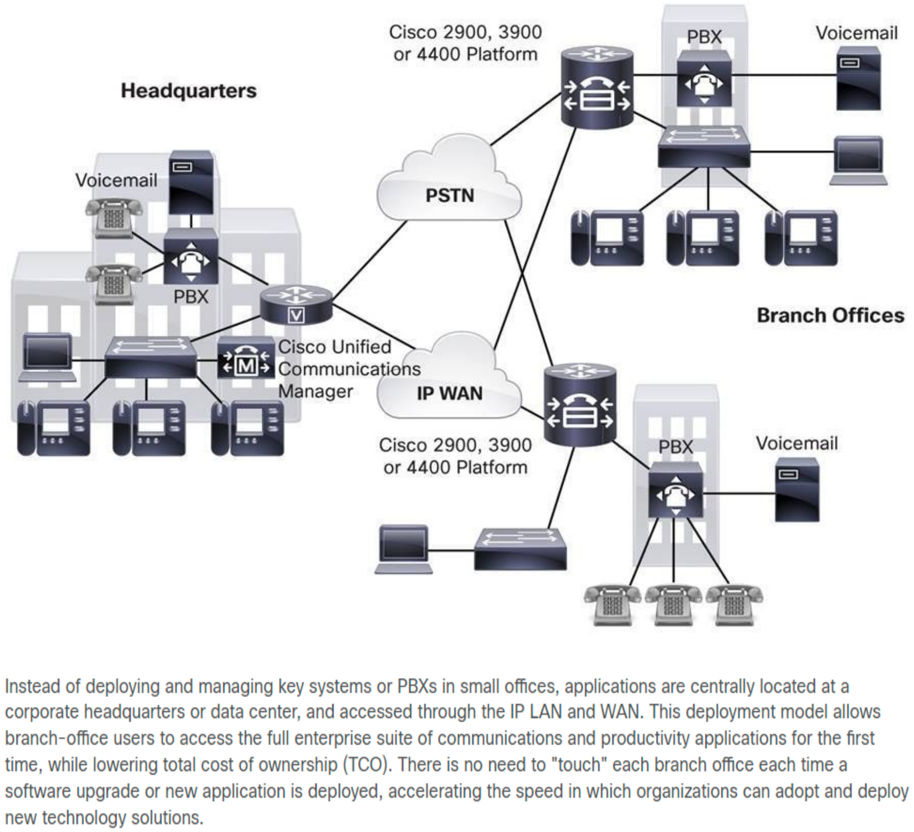

Instead of deploying and managing key systems or PBXs in small offices, applications are centrally located at a corporate headquarters or data center, and accessed through the IP LAN and WAN. This deployment model allows branch-office users to access the full enterprise suite of communications and productivity applications for the first time, while lowering total cost of ownership (TCO). There is no need to "touch" each branch office each time a software upgrade or new application is deployed, accelerating the speed in which organizations can adopt and deploy new technology solutions.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

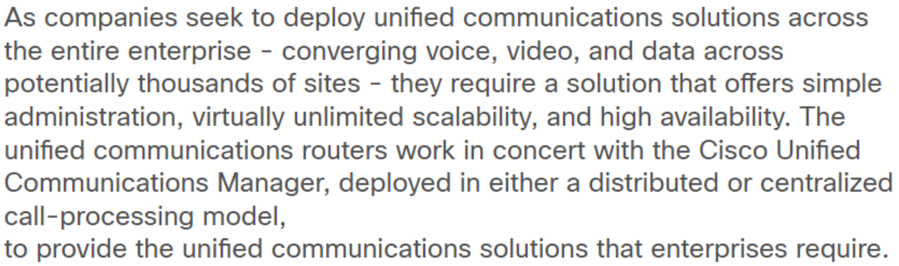

As companies seek to deploy unified communications solutions across the entire enterprise - converging voice, video, and data across potentially thousands of sites - they require a solution that offers simple administration, virtually unlimited scalability, and high availability. The unified communications routers work in concert with the Cisco Unified Communications Manager, deployed in either a distributed or centralized call-processing model,
to provide the unified communications solutions that enterprises require.

https://www.cisco.com/c/en/us/products/collateral/unified-communications/tdm-gateways/data-sheet-c78-729824.html

142. Cisco markets, offers to sell, sells, and distributes the Cisco Unified Communication System, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’885 Patent. The Cisco Unified Communication Systems are not staple articles or commodities of commerce suitable for any

substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

143. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the ’885 Patent. Cisco’s infringement of the ’885 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney’s fees, and costs in bringing this action.

144. Cisco’s direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary and permanent injunctive relief and damages as a result of Cisco’s infringement of the ’885 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

COUNT 7

(Infringement of U.S. Pat. No. 7,885,286)

145. NetSocket repeats and re-alleges all the allegations above as if fully set forth herein.

146. Cisco has infringed and continues to infringe one or more claims of the ’286 Patent by making, using, offering to sell, selling, and/or importing into the United States infringing devices without authority in violation of 35 U.S.C. § 271(a). Cisco has actively induced infringement of the ’286 Patent, and continues to induce infringement, without authority in violation of 35 U.S.C. § 271(b). Cisco has also contributed to the infringement of the ’286 Patent and continues to contribute to the infringement of the ’286 Patent in violation of 35 U.S.C. §271(c).

147. Cisco has and continues to infringe, directly and indirectly, literally and under the doctrine of equivalents, at least claim 1 of the ’286 Patent at least by making, using, offering to sell, importing, and/or selling Cisco CDA video quality experience (VQE) application system.

148. For example, upon information and belief, Cisco manufactures, imports, sells and offers to sell the Cisco CDA video quality experience (VQE) application system in the Unites States, and encourages distributors to sell, offer to sell and use, and encourages Cisco’s customers to use, the Cisco CDA video quality experience (VQE) application system in the United States with knowledge that the Cisco CDA video quality experience (VQE) application system infringes the ’286 Patent.

149. Cisco has had knowledge of and notice of the ’286 Patent and its infringement since before the filing of the original Complaint. For example, on information and belief, Cisco had knowledge of and notice of the ’286 Patent as a result of a partnership with NetSocket during at least 2012. Cisco offers to sell the Cisco CDA video quality experience (VQE) application system in this District and does so with knowledge that the sale and use of the Cisco CDA video quality experience (VQE) application system infringes and with the intent for its customers to use the Cisco CDA video quality experience (VQE) application system in an infringing manner.

150. The Cisco CDA video quality experience (VQE) application system satisfies each of the limitations of at least claim 1 of the ’286 Patent.

151. For example, the Cisco CDA video quality experience (VQE) application system provides for an application framework connectable to an IP network, as described below:

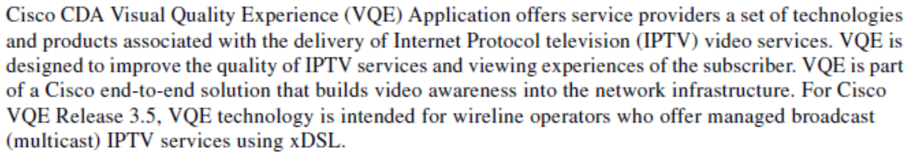

Cisco CDA Visual Quality Experience (VQE) Application offers service providers a set of technologies and products associated with the delivery of Internet Protocol television (IPTV) video services. VQE is designed to improve the quality of IPTV services and viewing experiences of the subscriber. VQE is part of a Cisco end-to-end solution that builds video awareness into the network infrastructure. For Cisco VQE Release 3.5, VQE technology is intended for wireline operators who offer managed broadcast (multicast) IPTV services using xDSL.

Figure 1-1 VQE Major Components

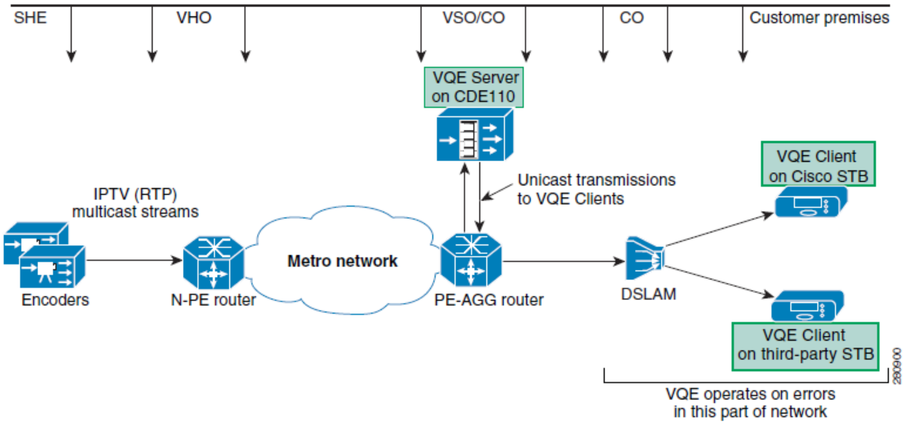
SHE
VHO
VSO/CO
CO
Customer premises
VQE Server on CDE110
Unicast transmissions to VQE Clients
VQE Client on Cisco STB
IPTV (RTP) multicast streams
Encoders
N-PE router
Metro network
PE-AGG router
DSLAM
VQE Client on third-party STB
280900
VQE operates on errors in this part of network

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

Figure 1-4 VQE Server Caches Multicast IPTV Packets for RCC

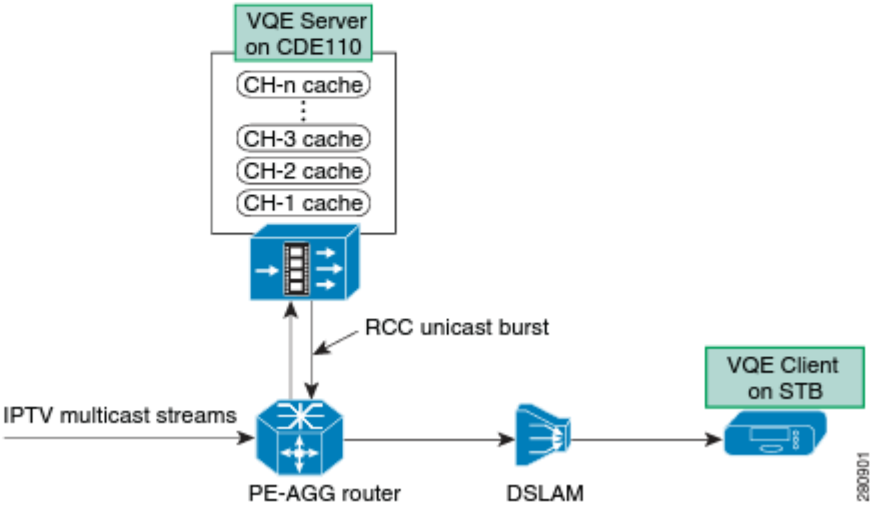

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

152. The Cisco CDA video quality experience (VQE) application system provides for receiving a request for a media distribution requiring a multicast distribution from a client in the IP network, as shown below:

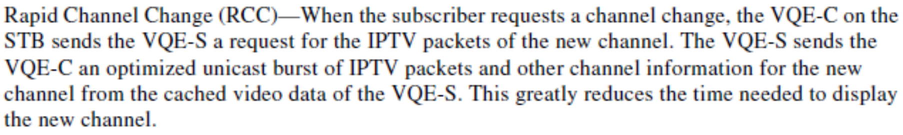

Rapid Channel Change (RCC)—When the subscriber requests a channel change, the VQE-C on the STB sends the VQE-S a request for the IPTV packets of the new channel. The VQE-S sends the VQE-C an optimized unicast burst of IPTV packets and other channel information for the new channel from the cached video data of the VQE-S. This greatly reduces the time needed to display the new channel.

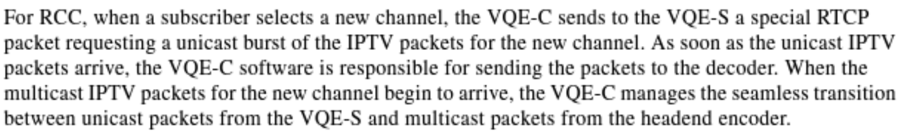

For RCC, when a subscriber selects a new channel, the VQE-C sends to the VQE-S a special RTCP packet requesting a unicast burst of the IPTV packets for the new channel. As soon as the unicast IPTV packets arrive, the VQE-C software is responsible for sending the packets to the decoder. When the multicast IPTV packets for the new channel begin to arrive, the VQE-C manages the seamless transition between unicast packets from the VQE-S and multicast packets from the headend encoder.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

Figure 1-1 VQE Major Components

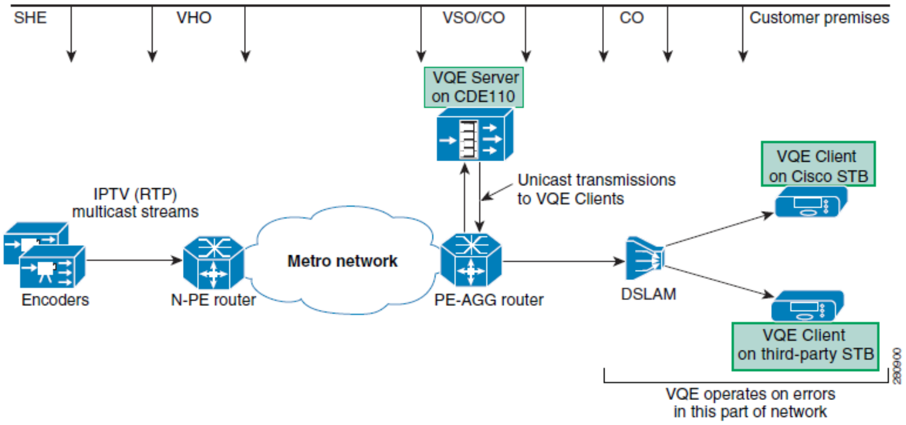
SHE
VHO
VSO/CO
CO
Customer premises
VQE Server on CDE110
Unicast transmissions to VQE Clients
VQE Client on Cisco STB
IPTV (RTP) multicast streams
Metro network
Encoders
N-PE router
PE-AGG router
DSLAM
VQE Client on third-party STB
280900
VQE operates on errors in this part of network

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

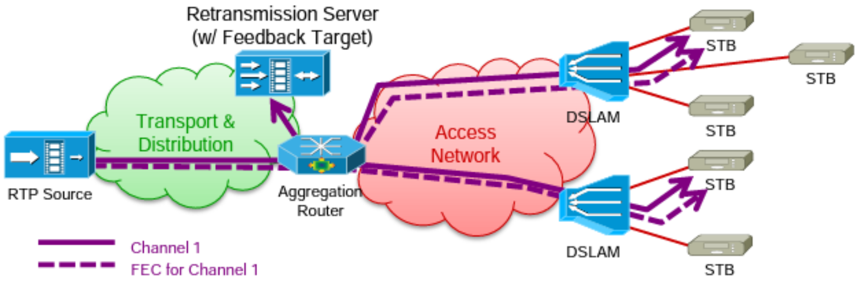

- **Each TV channel may be associated with one or more FEC streams**
  - FEC streams may have different repair capabilities
  - IP STBs may join the respective multicast sessions to receive FEC stream(s)

https://www.cisco.com/c/dam/global/en_ca/assets/plus/assets/pdf/IPTV-Internet-Video-ABEGEN.pdf

Figure 1-4 VQE Server Caches Multicast IPTV Packets for RCC

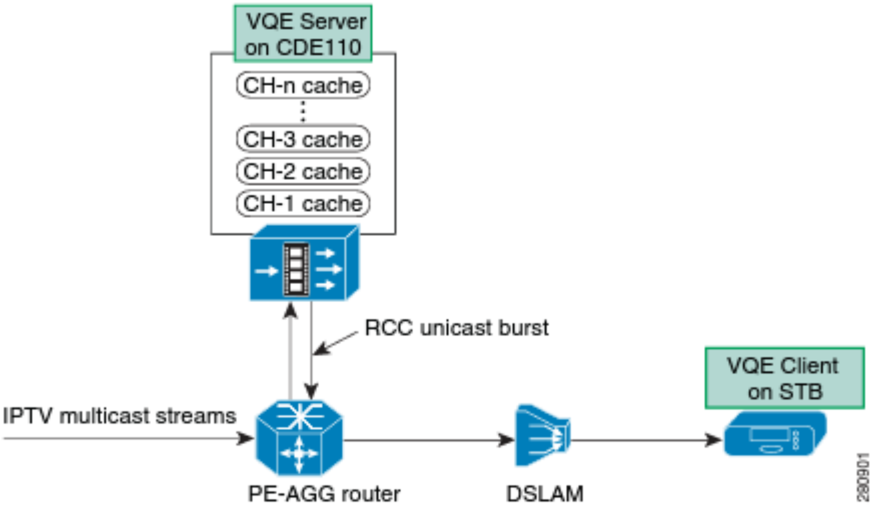

After a short period of time, multicast packets for the new channel start to arrive at the STB. The VQE-C monitors RTP sequence numbers from both unicast and multicast streams. It is likely that the VQE-C sees a few packets with duplicate RTP sequence numbers before the unicast stream ends. During this period, the VQE-C only forwards one copy of the RTP packet to the MPEG demultiplexing stage.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

153. The Cisco CDA video quality experience (VQE) application system provides for requesting network resources from a network resource manager for a Media Quick Start, to start the requested media distribution, as described below:

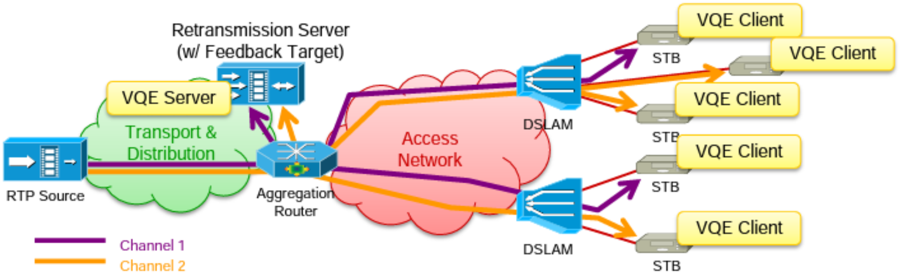

- **Each TV channel is served in a unique (SSM) multicast session**

  IP STBs join the respective multicast session for the desired TV channel

  Retransmission servers join all multicast sessions

- **Unicast feedback from IP STBs are collected by the feedback target**

  NACK messages reporting missing packets, rapid channel change requests

  RTCP receiver and extended reports reporting reception quality

https://www.cisco.com/c/dam/global/en_ca/assets/plus/assets/pdf/IPTV-Internet-Video-ABEGEN.pdf

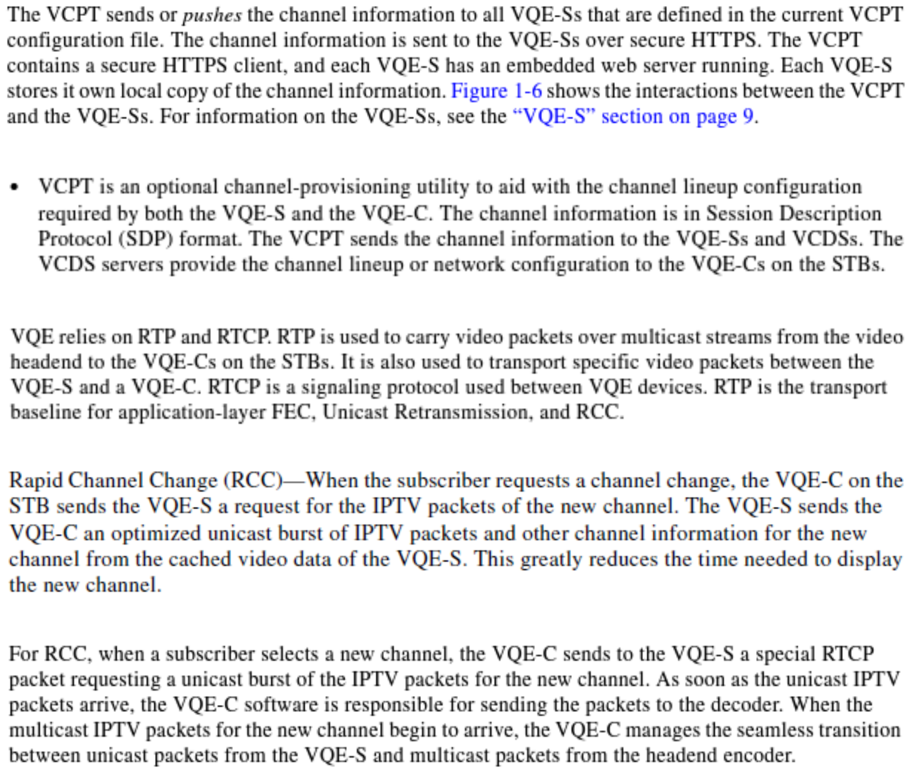

The VCPT sends or *pushes* the channel information to all VQE-Ss that are defined in the current VCPT configuration file. The channel information is sent to the VQE-Ss over secure HTTPS. The VCPT contains a secure HTTPS client, and each VQE-S has an embedded web server running. Each VQE-S stores it own local copy of the channel information. Figure 1-6 shows the interactions between the VCPT and the VQE-Ss. For information on the VQE-Ss, see the "VQE-S" section on page 9.

- VCPT is an optional channel-provisioning utility to aid with the channel lineup configuration required by both the VQE-S and the VQE-C. The channel information is in Session Description Protocol (SDP) format. The VCPT sends the channel information to the VQE-Ss and VCDSs. The VCDS servers provide the channel lineup or network configuration to the VQE-Cs on the STBs.

VQE relies on RTP and RTCP. RTP is used to carry video packets over multicast streams from the video headend to the VQE-Cs on the STBs. It is also used to transport specific video packets between the VQE-S and a VQE-C. RTCP is a signaling protocol used between VQE devices. RTP is the transport baseline for application-layer FEC, Unicast Retransmission, and RCC.

Rapid Channel Change (RCC)—When the subscriber requests a channel change, the VQE-C on the STB sends the VQE-S a request for the IPTV packets of the new channel. The VQE-S sends the VQE-C an optimized unicast burst of IPTV packets and other channel information for the new channel from the cached video data of the VQE-S. This greatly reduces the time needed to display the new channel.

For RCC, when a subscriber selects a new channel, the VQE-C sends to the VQE-S a special RTCP packet requesting a unicast burst of the IPTV packets for the new channel. As soon as the unicast IPTV packets arrive, the VQE-C software is responsible for sending the packets to the decoder. When the multicast IPTV packets for the new channel begin to arrive, the VQE-C manages the seamless transition between unicast packets from the VQE-S and multicast packets from the headend encoder.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

154. As shown below, the Cisco CDA video quality experience (VQE) application system provides for receiving feedback information relating to network resource availability from the network resource manager:

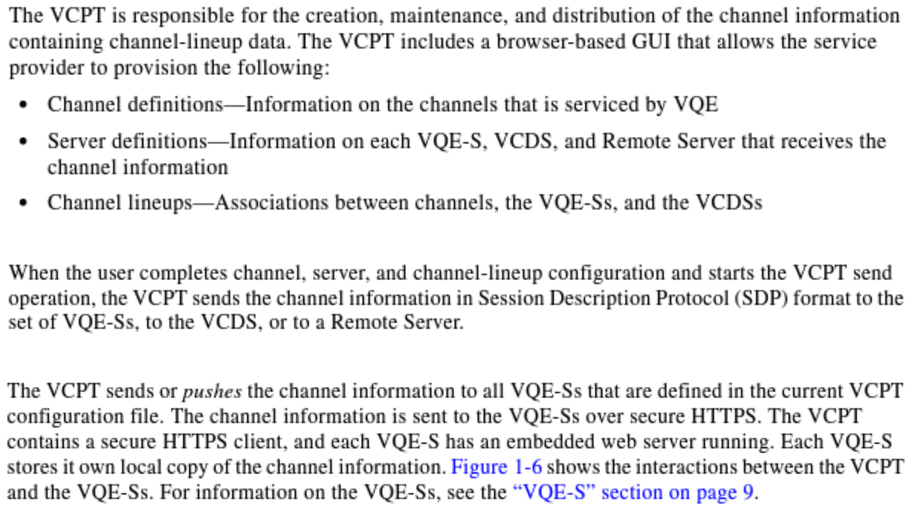

The VCPT is responsible for the creation, maintenance, and distribution of the channel information containing channel-lineup data. The VCPT includes a browser-based GUI that allows the service provider to provision the following:

- Channel definitions—Information on the channels that is serviced by VQE
- Server definitions—Information on each VQE-S, VCDS, and Remote Server that receives the channel information
- Channel lineups—Associations between channels, the VQE-Ss, and the VCDSs

When the user completes channel, server, and channel-lineup configuration and starts the VCPT send operation, the VCPT sends the channel information in Session Description Protocol (SDP) format to the set of VQE-Ss, to the VCDS, or to a Remote Server.

The VCPT sends or *pushes* the channel information to all VQE-Ss that are defined in the current VCPT configuration file. The channel information is sent to the VQE-Ss over secure HTTPS. The VCPT contains a secure HTTPS client, and each VQE-S has an embedded web server running. Each VQE-S stores it own local copy of the channel information. Figure 1-6 shows the interactions between the VCPT and the VQE-Ss. For information on the VQE-Ss, see the "VQE-S" section on page 9.

Like a regular IP host, the VQE-S joins multicast groups using Internet Group Management Protocol (IGMP). The VQE-S maintains a dedicated buffer for each channel. The VQE-S receives the multicast stream for each channel from upstream, caching a few seconds of the most recently received program content from each. The VQE-S can use the same cache of video to service both Unicast Retransmission requests and RCC requests.

Figure 1-6 VCPT: Sending Channel Information to VQE Servers

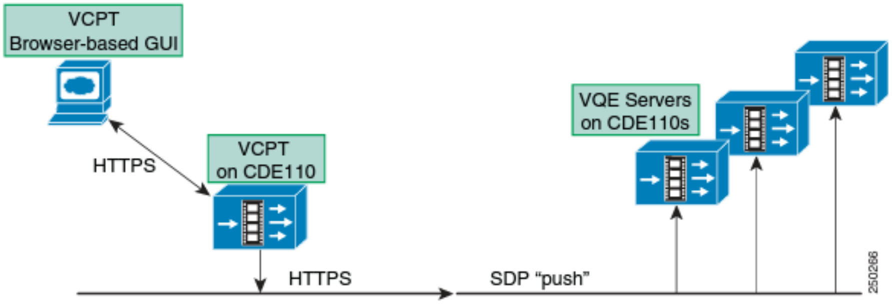

The VQE-S software is hosted on a Cisco CDE110 appliance running a standard Linux operating system. The Cisco CDE110 comes with the required software preinstalled: VQE-S, VQE-S AMT, Linux, Apache web server, and other software.

The VQE-S is responsible for the following functions:

- Creating a channel configuration database using the channel configuration information sent by the VCPT
- Maintaining per-channel and per-component state information
- Handling Unicast Retransmission by caching RTP data streams for channels and sending repair packets to the requesting VQE-Cs on the STBs
- Handling RCC by caching RTP data streams as well as PAT, PMT, ECM, PCR, and sequence information and sending RCC unicast bursts to requesting VQE-Cs when a channel change occurs
- Load balancing VQE-S services across the Cisco CDE110 Ethernet or bond interfaces
- Providing detailed statistics on IPTV delivery down to the STB VQE-C level
- Monitoring the health of VQE-S application processes

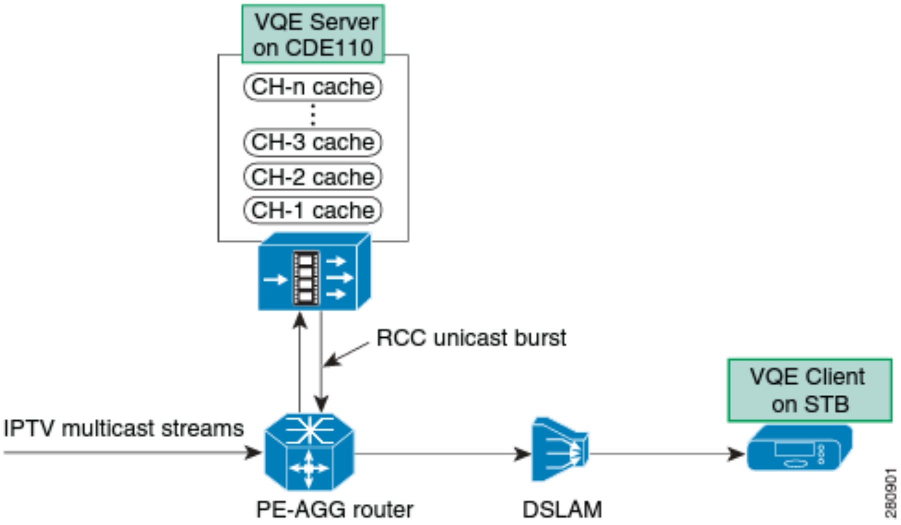

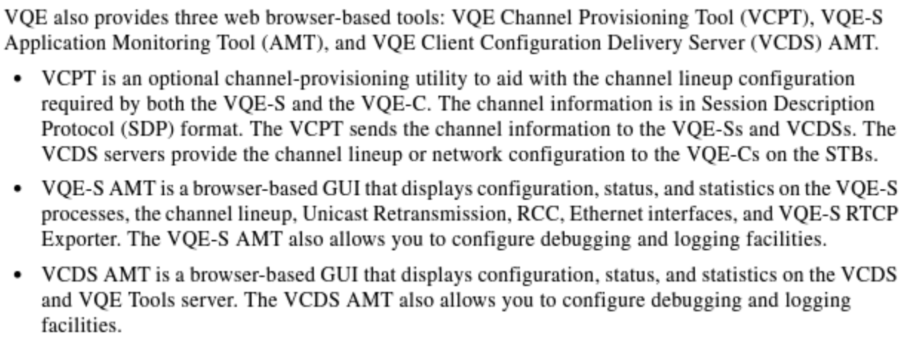

VQE also provides three web browser-based tools: VQE Channel Provisioning Tool (VCPT), VQE-S Application Monitoring Tool (AMT), and VQE Client Configuration Delivery Server (VCDS) AMT.

- VCPT is an optional channel-provisioning utility to aid with the channel lineup configuration required by both the VQE-S and the VQE-C. The channel information is in Session Description Protocol (SDP) format. The VCPT sends the channel information to the VQE-Ss and VCDSs. The VCDS servers provide the channel lineup or network configuration to the VQE-Cs on the STBs.
- VQE-S AMT is a browser-based GUI that displays configuration, status, and statistics on the VQE-S processes, the channel lineup, Unicast Retransmission, RCC, Ethernet interfaces, and VQE-S RTCP Exporter. The VQE-S AMT also allows you to configure debugging and logging facilities.
- VCDS AMT is a browser-based GUI that displays configuration, status, and statistics on the VCDS and VQE Tools server. The VCDS AMT also allows you to configure debugging and logging facilities.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

155. As described below, the Cisco CDA video quality experience (VQE) application system provides for allowing the Media Quick Start or using another behavior to start the requested media distribution based on the received feedback information wherein the another behavior is to use the feedback information provided by the network resource manager to signal to the client that a Media Quick Start unicast stream for the requested media distribution is not available due to resource shortage in the IP network but the client is still allowed to subscribe to the multicast distribution:

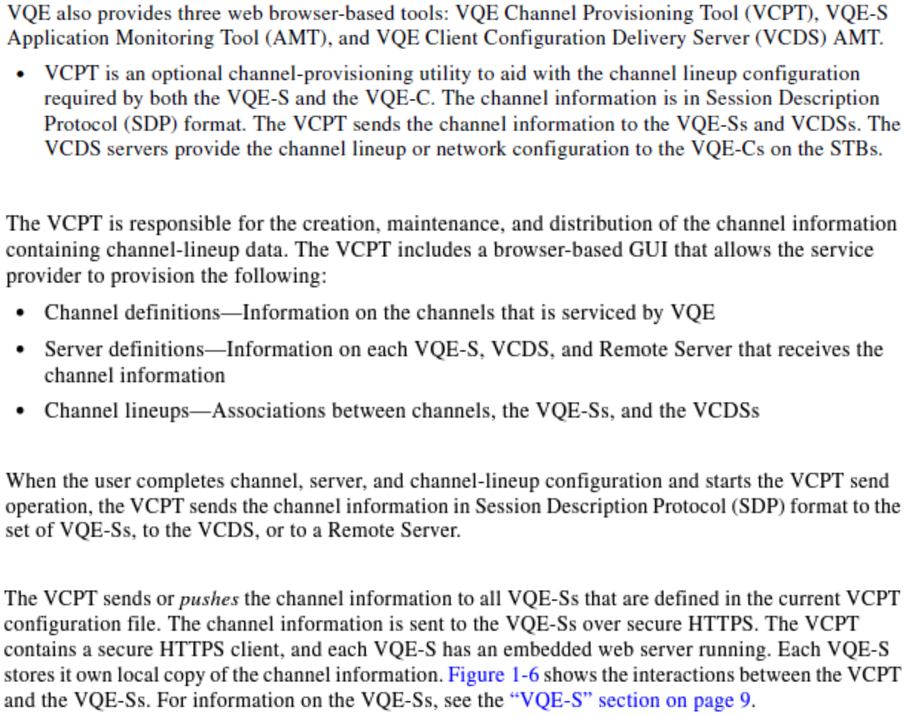

VQE also provides three web browser-based tools: VQE Channel Provisioning Tool (VCPT), VQE-S Application Monitoring Tool (AMT), and VQE Client Configuration Delivery Server (VCDS) AMT.

- VCPT is an optional channel-provisioning utility to aid with the channel lineup configuration required by both the VQE-S and the VQE-C. The channel information is in Session Description Protocol (SDP) format. The VCPT sends the channel information to the VQE-Ss and VCDSs. The VCDS servers provide the channel lineup or network configuration to the VQE-Cs on the STBs.

The VCPT is responsible for the creation, maintenance, and distribution of the channel information containing channel-lineup data. The VCPT includes a browser-based GUI that allows the service provider to provision the following:

- Channel definitions—Information on the channels that is serviced by VQE
- Server definitions—Information on each VQE-S, VCDS, and Remote Server that receives the channel information
- Channel lineups—Associations between channels, the VQE-Ss, and the VCDSs

When the user completes channel, server, and channel-lineup configuration and starts the VCPT send operation, the VCPT sends the channel information in Session Description Protocol (SDP) format to the set of VQE-Ss, to the VCDS, or to a Remote Server.

The VCPT sends or *pushes* the channel information to all VQE-Ss that are defined in the current VCPT configuration file. The channel information is sent to the VQE-Ss over secure HTTPS. The VCPT contains a secure HTTPS client, and each VQE-S has an embedded web server running. Each VQE-S stores it own local copy of the channel information. Figure 1-6 shows the interactions between the VCPT and the VQE-Ss. For information on the VQE-Ss, see the “VQE-S” section on page 9.

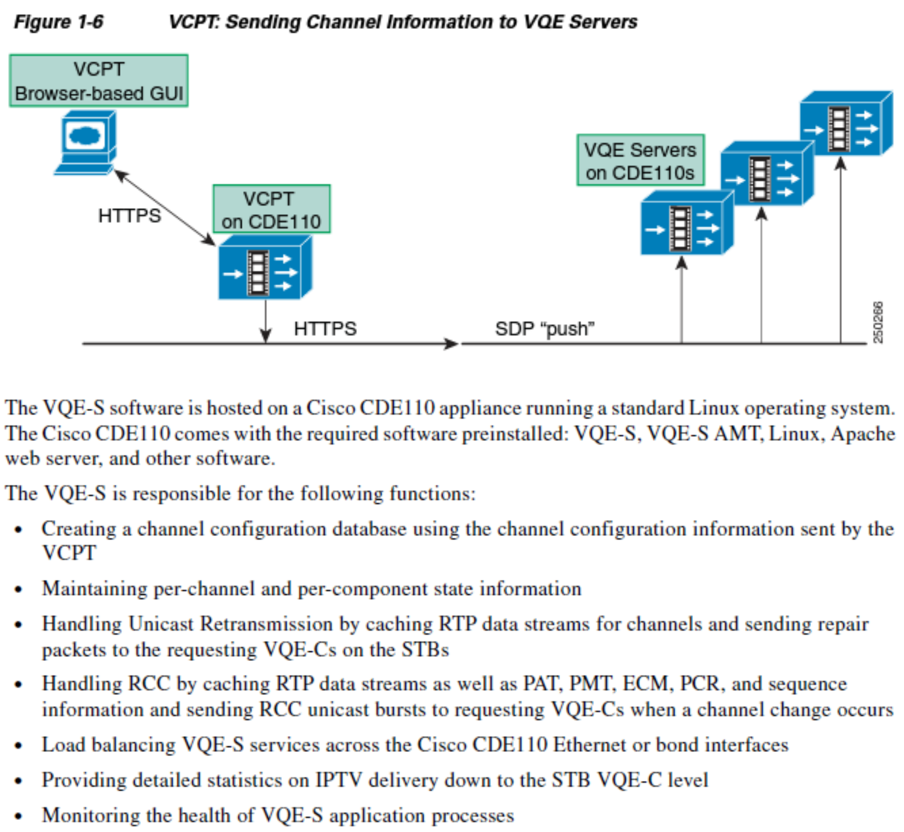

Figure 1-6 VCPT: Sending Channel Information to VQE Servers

VCPT
Browser-based GUI
HTTPS
VCPT
on CDE110
VQE Servers
on CDE110s
HTTPS
SDP "push"
250266

The VQE-S software is hosted on a Cisco CDE110 appliance running a standard Linux operating system. The Cisco CDE110 comes with the required software preinstalled: VQE-S, VQE-S AMT, Linux, Apache web server, and other software.

The VQE-S is responsible for the following functions:

- Creating a channel configuration database using the channel configuration information sent by the VCPT
- Maintaining per-channel and per-component state information
- Handling Unicast Retransmission by caching RTP data streams for channels and sending repair packets to the requesting VQE-Cs on the STBs
- Handling RCC by caching RTP data streams as well as PAT, PMT, ECM, PCR, and sequence information and sending RCC unicast bursts to requesting VQE-Cs when a channel change occurs
- Load balancing VQE-S services across the Cisco CDE110 Ethernet or bond interfaces
- Providing detailed statistics on IPTV delivery down to the STB VQE-C level
- Monitoring the health of VQE-S application processes

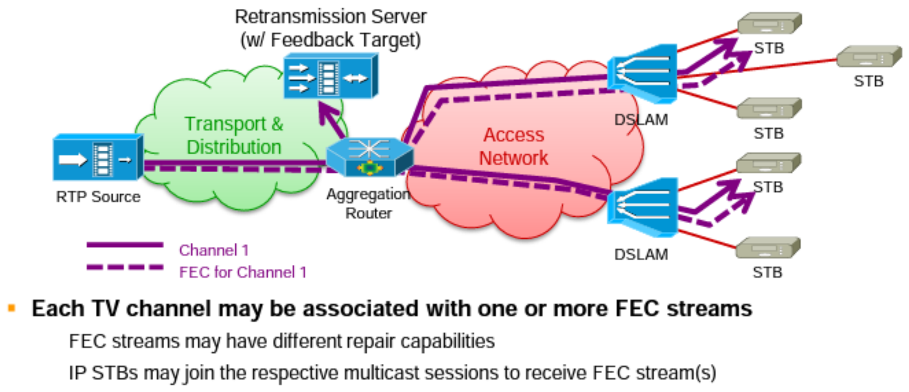

- **Each TV channel may be associated with one or more FEC streams**

  FEC streams may have different repair capabilities

  IP STBs may join the respective multicast sessions to receive FEC stream(s)

https://www.cisco.com/c/dam/global/en_ca/assets/plus/assets/pdf/IPTV-Internet-Video-ABEGEN.pdf

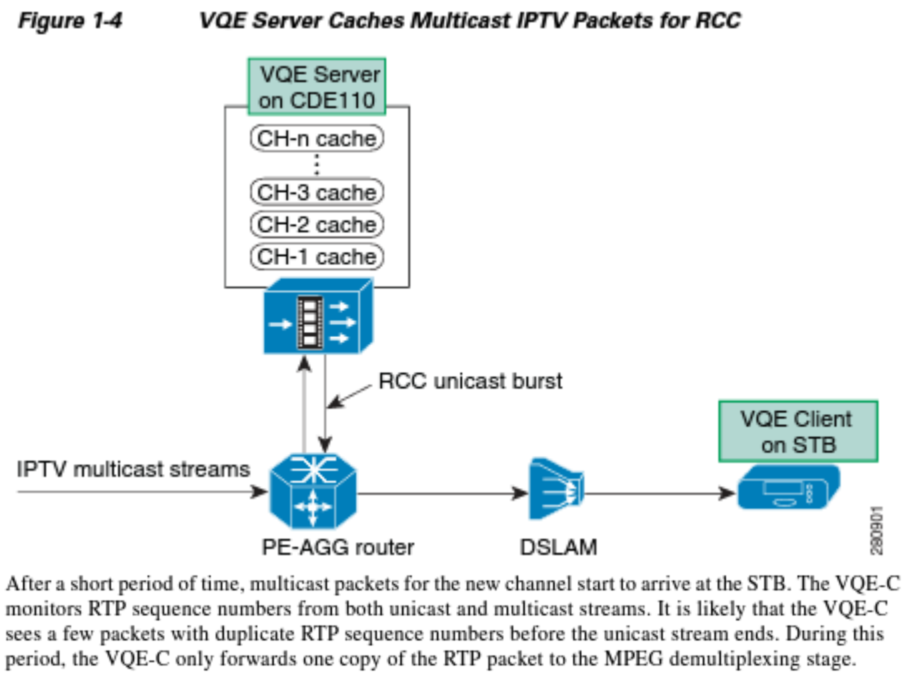

Figure 1-4 VQE Server Caches Multicast IPTV Packets for RCC

After a short period of time, multicast packets for the new channel start to arrive at the STB. The VQE-C monitors RTP sequence numbers from both unicast and multicast streams. It is likely that the VQE-C sees a few packets with duplicate RTP sequence numbers before the unicast stream ends. During this period, the VQE-C only forwards one copy of the RTP packet to the MPEG demultiplexing stage.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

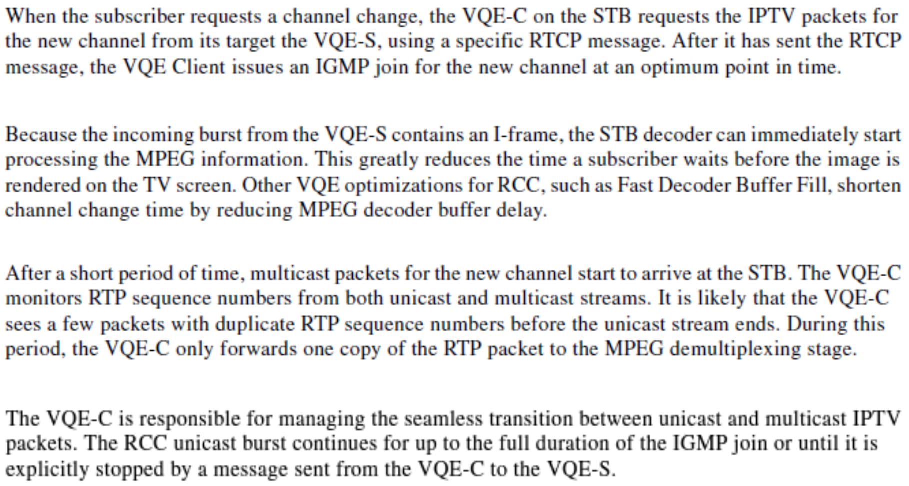

When the subscriber requests a channel change, the VQE-C on the STB requests the IPTV packets for the new channel from its target the VQE-S, using a specific RTCP message. After it has sent the RTCP message, the VQE Client issues an IGMP join for the new channel at an optimum point in time.

Because the incoming burst from the VQE-S contains an I-frame, the STB decoder can immediately start processing the MPEG information. This greatly reduces the time a subscriber waits before the image is rendered on the TV screen. Other VQE optimizations for RCC, such as Fast Decoder Buffer Fill, shorten channel change time by reducing MPEG decoder buffer delay.

After a short period of time, multicast packets for the new channel start to arrive at the STB. The VQE-C monitors RTP sequence numbers from both unicast and multicast streams. It is likely that the VQE-C sees a few packets with duplicate RTP sequence numbers before the unicast stream ends. During this period, the VQE-C only forwards one copy of the RTP packet to the MPEG demultiplexing stage.

The VQE-C is responsible for managing the seamless transition between unicast and multicast IPTV packets. The RCC unicast burst continues for up to the full duration of the IGMP join or until it is explicitly stopped by a message sent from the VQE-C to the VQE-S.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

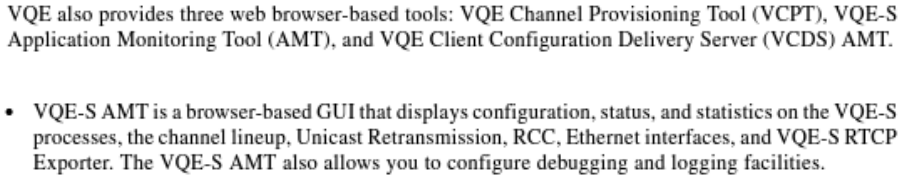

VQE also provides three web browser-based tools: VQE Channel Provisioning Tool (VCPT), VQE-S Application Monitoring Tool (AMT), and VQE Client Configuration Delivery Server (VCDS) AMT.

- VQE-S AMT is a browser-based GUI that displays configuration, status, and statistics on the VQE-S processes, the channel lineup, Unicast Retransmission, RCC, Ethernet interfaces, and VQE-S RTCP Exporter. The VQE-S AMT also allows you to configure debugging and logging facilities.

The VQE-S AMT is a browser-based GUI that allows the service-provider operator to do the following:

- Monitor the health of the VQE-S processes
- View channel configuration details, status, and statistics
- Monitor statistics for Unicast Retransmission and RCC
- Monitor statistics for STUN Server usage
- View configuration details, status, and statistics for:
  - Multicast Load Balancer
  - VQE-S RTCP Exporter
- Change VQE-S logging levels and debugging options

Figure 1-10 Monitoring VQE-S Processes

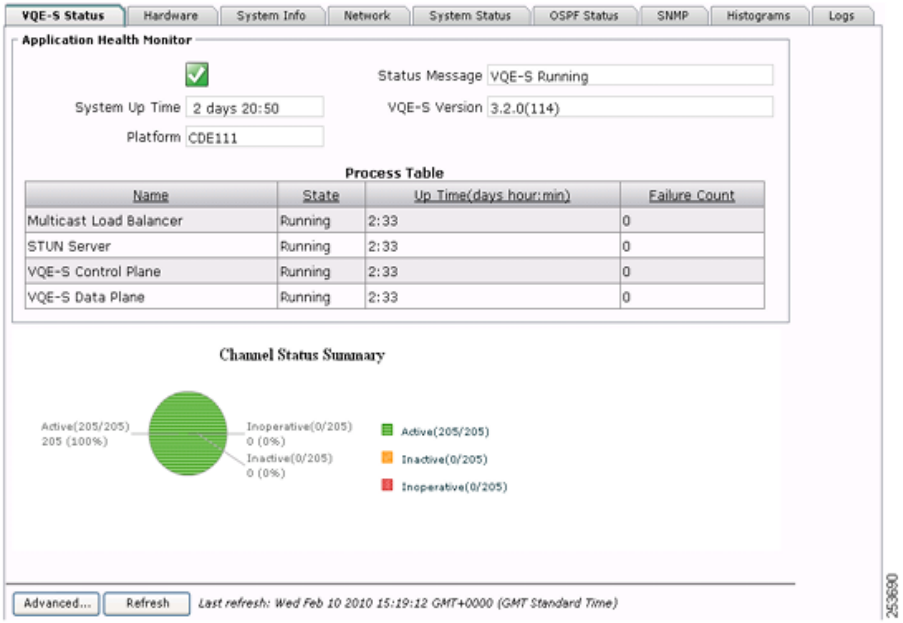

Application Health Monitor

✓

Status Message VQE-S Running

System Up Time 2 days 20:50

VQE-S Version 3.2.0(114)

Platform CDE111

Process Table

Name	State	Up Time(days hour:min)	Failure Count
Multicast Load Balancer	Running	2:33	0
STUN Server	Running	2:33	0
VQE-S Control Plane	Running	2:33	0
VQE-S Data Plane	Running	2:33	0

Channel Status Summary

Last refresh: Wed Feb 10 2010 15:19:12 GMT+0000 (GMT Standard Time)

- Status—Indicates channel status as follows:
  - Green with a checkmark—Channel is active. The VQE-S is receiving the multicast stream.
  - Yellow with exclamation mark (!)—Channel is inactive. The channel is successfully initialized, but the VQE-S is not receiving the multicast stream.
  - Red with an X—Channel is inoperative (for example, the channel is not configured correctly).
- Member Receiver Population—Provides the number of VQE-Cs that are currently receiving this multicast stream.

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/vqe/3_5/user/guide/vqe_guide3_5.pdf

156. Cisco markets, offers to sell, sells, and distributes the Cisco CDA video quality experience (VQE) application system, and will continue to do so, knowing the same to be especially made or especially adapted for use in an infringement of the ’286 Patent. The Cisco CDA video quality experience (VQE) application system is not a staple article or commodity of

commerce suitable for any substantial non-infringing uses. As shown above, the purpose and nature of the Accused System are directed to infringement and, as a result, there are no substantial non-infringing uses.

157. Cisco has committed and continues to commit acts of infringement that Cisco knew or should have known constituted an unjustifiably high risk of infringement of the '286 Patent. Cisco's infringement of the '286 Patent has been and continues to be deliberate and willful, entitling NetSocket to an award of treble damages, reasonable attorney's fees, and costs in bringing this action.

158. Cisco's direct and indirect infringement has caused and is continuing to cause damage and irreparable injury to NetSocket. NetSocket will continue to suffer damage and irreparable injury until that injury is enjoined by this Court. NetSocket is entitled to preliminary and permanent injunctive relief and damages as a result of Cisco's infringement of the '286 Patent in accordance with 35 U.S.C. §§ 271, 281, 283, 284, and 285.

PRAYER FOR RELIEF

WHEREFORE, NetSocket prays that this Court grant the following relief:

a) An order adjudging and decreeing that Cisco has infringed one or more claims of the Asserted Patents;

b) A permanent injunction pursuant to 35 U.S.C. § 283 against the continuing infringement of the claims of the Asserted Patents by Cisco, its officers, agents, employees, attorneys, representatives, and all others acting in concert therewith;

c) An order directing Cisco to account for and pay to NetSocket all damages caused to NetSocket by reason of Cisco's patent infringement, pursuant to 35 U.S.C. §§ 284 and 289, and that interest and costs be assessed against Cisco;

d) A declaration that Cisco's infringement was and is willful from the time it became aware of the infringing nature of its products and an award of treble damages for the period of such willful infringement of the Asserted Patents, pursuant to 35 U.S.C. § 284;

e) A declaration that this case is exceptional and an award of attorneys' fees and costs under 35 U.S.C. § 285 against Cisco; and

f) For all other relief the Court deems just and proper.

JURY DEMAND

Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, NetSocket hereby demands a trial by jury of all issues so triable.

DATED: January 11, 2022

Respectfully submitted,

By: */s/ Greg Love*___

Greg Love
State Bar No. 24013060
Email: greg@swclaw.com
STECKLER WAYNE CHERRY & LOVE, PLLC
107 East Main Street
Henderson, Texas 75652
Tel: (903) 212-4444

Mark D. Siegmund
State Bar No. 24117055
Email: mark@swclaw.com
Craig D. Cherry
State Bar No. 24012419
Email: craig@swclaw.com
Justin Allen
State Bar No. 24081977
Email: justin@swclaw.com
STECKLER WAYNE CHERRY & LOVE, PLLC
8416 Old McGregor Road
Waco, Texas 76712
Tel: (254) 651-3690
Fax: (254) 651-3689

Thomas G. Southard
Brian S. Seal
Shaun D. Gregory
(*Pro Hac Vice*)
TAFT STETTINIUS & HOLLISTER LLP
200 Massachusetts Ave., Suite 400
Washington, D.C. 20001

Tel: (202) 664-1537
Fax: (202) 664-1586
tsouthard@taftlaw.com
bseal@taftlaw.com
sgregory@taftlaw.com

Richard Eric Gaum
(*Pro Hac Vice*)
TAFT STETTINIUS & HOLLISTER LLP
200 Public Square, Suite 3500
Cleveland, Ohio 44114-2302
Tel: (216) 241-2838
Fax: (216) 241-3707
egaum@taftlaw.com

Mira Vats-Fournier
TAFT STETTINIUS & HOLLISTER LLP
(*Pro Hac Vice*)
2200 IDS Center
80 South 8th Street
Minneapolis, MN 55402
Tel: 612-977-8075
Fax: 612-977-8650
mvats-fournier@taftlaw.com

Shalu Maheshwari
(*Pro Hac Vice*)
DAIGNAULT IYER LLP
8618 Westwood Center Drive
Suite 150
Vienna, VA 22102
Tele: (202) 997-1925
smaheshwari@daignaultiyer.com

Attorneys for Plaintiff NetSocket, Inc.

CERTIFICATE OF SERVICE

The undersigned certifies that the foregoing document was filed electronically in compliance with Local Rules CV-5(a). As such, this document was served on all counsel who have consented to electronic service on this 11th day of January 2023.

/s/ *Greg Love*
Greg Love