

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

WSOU INVESTMENTS d/b/a BRAZOS
LICENSING AND DEVELOPMENT,
Plaintiff,
v.

CISCO SYSTEMS, INC.,
Defendant.

Case No.

JURY TRIAL DEMANDED

COMPLAINT

Plaintiff WSOU Investments d/b/a Brazos Licensing and Development (“Brazos”) files this Complaint against Defendant Cisco Systems, Inc. (“Cisco”).

NATURE OF THE CASE

1. This is an action for the infringement of five United States Patents: (1) United States Patent No. 7,386,630 (“the ‘630 Patent”), (2) United States Patent No. 8,498,286 (“the ‘286 Patent”), (3) United States Patent No. 8,441,721 (“the ‘721 Patent”), (4) United States Patent No. 8,982,691 (“the ‘691 Patent”), and (5) United States Patent No. 9,450,884 (“the ‘884 Patent”), collectively referred to as the “Patents-in-Suit.”

2. Cisco has been infringing the Patents-in-Suit in violation of 35 U.S.C. § 271 by deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling its suite of networking products, including but not limited to at least the Catalyst 9000 Switching Platform, Ultra-M Platform, Cisco NCS 1010 Optical Line System, NCS 4200 Series Network Convergence Systems, and similar products (collectively, the “Infringing Products”).

3. Plaintiff Brazos seeks appropriate damages, injunctive relief, and prejudgment and post-judgment interest for Defendant's infringement of the Patents-in-Suit.

THE PARTIES

4. Plaintiff Brazos is a limited liability corporation organized and existing under the laws of Delaware, with its principal place of business at 605 Austin Avenue, Suite 6, Waco, Texas 76701.

5. Defendant Cisco is a corporation organized under the laws of the State of California, with its principal place of business at 170 W. Tasman Dr., San Jose, CA 95134.

6. On information and belief, Cisco's operations in the Eastern District of Texas are substantial and varied.

7. By registering to conduct business in Texas and by maintaining facilities in at least the city of Richardson, Cisco has multiple regular and established places of business within the Eastern District of Texas.

JURISDICTION AND VENUE

8. This is an action for patent infringement arising under the Patent Laws of the United States, Title 35 of the United States Code.

9. This Court has original subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

10. This Court has personal jurisdiction over Cisco in this action because Cisco has committed acts within this District giving rise to this action and 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. Cisco, directly and/or through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by,

among other things, deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling products that infringe the Asserted Patents.

11. Cisco, directly and/or through subsidiaries or intermediaries, has purposefully and voluntarily placed one or more products and/or services in the stream of commerce that practice the Asserted Patents with the intention and expectation that they will be purchased and used by consumers in the Eastern District of Texas. These products and/or services have been and continue to be purchased and used in the Eastern District of Texas.

12. Venue is proper in this District under 28 U.S.C. §§ 1391(b)-(c) and 1400(b). Cisco is registered to do business in Texas. Additionally, on information and belief, Cisco has transacted business in this District and has committed acts of direct and indirect infringement in this District by, among other things, deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling products that infringe the Asserted Patents. Moreover, on information and belief, Cisco has regular and established places of business in the Eastern District of Texas, including at 2250 and 2300 East President George Bush Turnpike, Richardson, Texas 75082.

13. Cisco's Richardson, Texas offices employ 2,000 individuals and have several open job postings. Cisco's physical presence in Richardson consists of a 78-acre campus including seven buildings. Cisco also spent \$80 million to retool an unused office building as a Tier III data center, firmly establishing itself in the city and this district.

14. Cisco's operations in the Eastern District of Texas are substantial and varied, and include at least some operations related to Routing and Switching. For example, Cisco certifies its own employees and other employees in the industry through its CCIE certification program, the largest subset of which is Routing and Switching. The certification process contains a lab

examination section, and the only non-virtual lab examination offered in the United States is based in Richardson, Texas.

15. Cisco also advertises that it is currently hiring for positions related to Routing and Switching in Richardson, Texas, including a “Technical Consulting Engineer” working with “virtualized switch[es]”.

THE PATENTS-IN-SUIT

16. The '630 Patent is titled “Using Policy-Based Management To Support Diffserv Over MPLS Network” and was issued by the United States Patent and Trademark Office to inventors Yin Ling Liong, Roberto Barnes, and Man Li on June 10, 2008, and assigned to Nokia Corporation. It was assigned to Brazos on July 22, 2017.

17. Brazos is the owner of all right, title, and interest in and to the '630 Patent with the full and exclusive right to bring suit to enforce the '630 Patent.

18. The '630 Patent is valid and enforceable under the United States Patent Laws.

19. The '286 Patent is titled “Radius Gateway on Policy Charging and Rules Function (PCRF) for Wireline/Wireless Converged Solution” and was issued by the United States Patent and Trademark Office to inventors Fan Mo, Satvinder Bawa, Lui Yeung, Justin Newcomb, Lay Been Tan, Felix Landry, and Ivaylo Tanouchev on July 30, 2013, and assigned to Alcatel Lucent. It was assigned to Brazos on November 26, 2019.

20. Brazos is the owner of all right, title, and interest in and to the '286 Patent with the full and exclusive right to bring suit to enforce the '286 Patent.

21. The '286 Patent is valid and enforceable under the United States Patent Laws.

22. The '721 Patent is titled “System and Method of Raman Amplifier Pump Control” and was issued by the United States Patent and Trademark Office to inventor Christopher Alan

White on May 14, 2013, and assigned to Alcatel Lucent. It was assigned to Brazos on November 26, 2019.

23. Brazos is the owner of all right, title, and interest in and to the '721 Patent with the full and exclusive right to bring suit to enforce the '721 Patent.

24. The '721 Patent is valid and enforceable under the United States Patent Laws.

25. The '691 Patent is titled "System and Method Providing Standby Bypass for Double Failure Protection in MPLS Network" and was issued by the United States Patent and Trademark Office to inventors Pradeep Jain, Kanwar Singh, and Srikrishnan Venkataraman on March 17, 2015, and assigned to Alcatel Lucent. It was assigned to Brazos on July 22, 2017.

26. Brazos is the owner of all right, title, and interest in and to the '691 Patent with the full and exclusive right to bring suit to enforce the '691 Patent.

27. The '691 Patent is valid and enforceable under the United States Patent Laws.

28. The '884 Patent is titled "Software Defined Networking Based Congestion Control" and was issued by the United States Patent and Trademark Office to inventors Jae Hyun Hwang and Thierry Klein on September 20, 2016, and assigned to Alcatel Lucent. It was assigned to Brazos on November 26, 2019.

29. Brazos is the owner of all right, title, and interest in and to the '884 Patent with the full and exclusive right to bring suit to enforce the '884 Patent.

30. The '884 Patent is valid and enforceable under the United States Patent Laws.

31. The Patents-in-Suit generally teach and claim novel improvements to computer and networking technology and recite inventive concepts, including elements that were not well-understood, routine, or conventional.

32. For example, the '630 Patent is directed to improvements to computer network technology, and more particularly, “software and methods for policy-based management of two combined functionalities (Diffserv over MPLS) in a single network.” '630 Patent at 4:19-21 (Ex. A). According to the '630 Patent, Diffserv functionality “simplifies the processing and storage requirements at core routers[,]” (*id.* at 1:62-63), and MPLS functionality “can relieve congestion and maximize bandwidth utilization by allowing multiple paths between source and destination.” *Id.* at 2:22-24. The inventors recognized that while combining the two would “enable a MPLS functionality that also performs with IP QoS support,” doing so involves “cumbersome mappings between Diffserv and MPLC policies.” *Id.* at 2:25-33. The Background of the '630 Patent noted that “there is no commercial solution addressing the policy management of Diffserv over MPLS with regards to the configuration of the E-LSP and tunneling modes” due perhaps to “different recommendations from the standards, and the limited capabilities supported at the network elements.” *Id.* at 2:39-45. The problem the inventors sought to address was thus rooted in computer networking technology, i.e., the combining of packet networking techniques for better quality of service and congestion/bandwidth optimization. *See id.*; Ex. F (Chrissan Declaration) at ¶ 25. This problem—combining packet networking techniques—is not something that arises in longstanding human activity, e.g., a fundamental economic practice or method of organizing human activity. *Id.* at ¶ 25.

33. To solve this problem, the inventors introduced a policy server that translates “policies into configurations of network resources and automates the configurations across multiple different network elements and different technologies (e.g. MPLS and Diffserv).” *Id.* at 4:26-31. The policy server is described and claimed to have a specific solution for combining MPLS and Diffserv functionalities, including being enabled to:

- configure a customer policy comprising a tunnel mode and a tunnel group identifier,
- configure a mapping policy that maps between an experimental field and a unique per-hop-behavior, and
- send the mapping policy and the customer policy to interfaces of devices of a network that includes multi-protocol label switching tunnels, corresponding to the tunnels, at least one of the network devices comprising an egress interface of one of said multi-protocol label switching tunnels, wherein the interfaces and the customer policy are associated with a same role name.

Id., cl. 1; Ex. F at ¶ 26.

34. The claimed solution “does not claim the result of combining MPLS and Diffserv but a specific how: a specific make up of customer policy, a specific mapping between fields, and a specific deployment of the aforementioned policies to MPLS device interfaces.” *Id.* at ¶ 27. Thus, claim 18 recites “a specific solution to the computer-rooted problem rather than a functional result.” *Id.* Accordingly, claim 18 is “directed to an improvement to computer technology and not an abstract idea.” *Id.*

35. Moreover, claim 18 recites additional elements that provide inventive concepts and “transform the claim into patent-eligible subject matter.” Ex. F at ¶ 26. Specifically, at least the following additional elements and their ordered combination were not “well-understood, routine, or conventional” as of 2008:

- configure a customer policy comprising a tunnel mode and a tunnel group identifier,
- configure a mapping policy that maps between an experimental field and a unique per-hop-behavior, and

- send the mapping policy and the customer policy to interfaces of devices of a network that includes multi-protocol label switching tunnels, corresponding to the tunnels, at least one of the network devices comprising an egress interface of one of said multi-protocol label switching tunnels, wherein the interfaces and the customer policy are associated with a same role name.

Ex. F at 28.

36. For example, during patent examination, the examiner allowed claim 18 because “the feature of configuring a customer policy comprising a tunnel mode and a tunnel group identifier is not expressly taught or suggested by the prior art, wherein Applicant's disclosure defines the tunnel mode as a method of ‘translating DiffServ information in the multi-protocol label switching header into DSCP values’ having two modes: uniform or pipe mode (see Applicant's Specification, page 4). Applicant's hardware embodiment of a policy server device having this functionality is not disclosed by the prior art.” *See* Chrissan Declaration Ex. 2 (’630 File History) at 6. That the Patent Office found the closest prior art did not disclose these additional elements supports their being non-routine and unconventional. Accordingly, claim 18 of the ’630 Patent recites additional elements that were not “well-understood, routine, or conventional,” which provide “something more” sufficient to transform any alleged abstract idea into eligible subject matter. Ex. F at ¶ 29.

37. The ’286 Patent is directed to improvements to computer network technology, and more particularly, a system comprising a Radius Gateway, a Diameter Proxy Agent, and a Policy Charging and Rules Function (PCRF) server (’286 Patent, cl. 6) that “recognize[s] and process[es] both Diameter messages and Radius messages.” *Id.* at Abstract; Ex. F at ¶ 30. PCRF servers perform important telecommunications functions such as “service data flow detection, policy enforcement, and flow-based charging.” ’630 Patent at 1:31-32. For wireless devices, Diameter

Proxy Agents, typically instantiated on the same platform as the PCRF server, “selects the most suitable PCRF on the platform, and thereafter allows communication between the selected PCRF and the wireless device by passing Diameter messages back and forth.” *Id.* at 1:37-40. In contrast to wireless, wireline services “such as wireline telephone services but also including internet and television services, are typically implemented using a different type of server, referred to as an Authentication, Authorization and Accounting (AAA) server.” *Id.* at 1:46-49.

38. The inventors of the '286 Patent recognized that a “communications company wishing to provide its customers with a variety of services usually has to use two servers, a server compliant with the 3GPP specifications in order to receive Diameter messages and an AAA server capable of processing Radius messages, and these servers are often managed by different divisions or sub-companies.” *Id.* at 1:66-2:5. Additionally, a “communication company may even use more than one AAA server for a given area, as AAA servers are typically not very scalable.” *Id.* at 2:5-7. “This problem is rooted in computer technology because it deals with the interoperability between wired and wireless communications servers, which is not some longstanding human activity.” Ex. F at ¶ 31.

39. To solve this problem, the inventors introduced a system that enables interoperability between wireless and wireline services, e.g., “a communication session for a wireline device [and corresponding Radius messages] in an Evolved Packet Core [wireless framework] [and corresponding Diameter messages].” *Id.* 2:19-20. The system is described and claimed as a specific combination of apparatuses and associated functions, including:

- a Radius gateway for translating a Radius message into a Network Access Server Request (NASReq) message and for transmitting the NASReq message;

- a Diameter Proxy Agent within a Policy and Charging Rules Function (PCRF) server for receiving a NASReq message from the Radius gateway, for selecting one of at least one PCRF cluster within the PCRF server, and for forwarding the NASReq message to the selected PCRF cluster; and
- at least one PCRF blade within the PCRF server, each PCRF blade belonging to one of the at least one PCRF cluster and configured to handle communication sessions for wireless devices, each PCRF blade for receiving a NASReq message from the Diameter Proxy Agent and for creating or updating a NASReq session object related to a communication session identified by the NASReq message received by the PCRF blade in response to receiving the NASReq message.

Id., cl. 6; Ex. F at ¶ 32.

40. The claimed solution does not claim the result of a system for interoperating between Radius and Diameter protocols but a specific how: a Radius gateway for translating the Radius message into a NASReq message, a Diameter Proxy Agent for selecting one PCRF cluster, and at least one PCRF blade to handle communication sessions for wireless devices and for updating a NASReq session object. Ex. F at ¶ 33. Thus, “claim 6 recites a specific solution to the computer-rooted problem rather than a functional result.” *Id.* Accordingly, claim 6 is directed to an improvement to computer technology and not an abstract idea. *Id.*

41. Moreover, claim 6 recites additional elements that provide inventive concepts and transform the claim into patent-eligible subject matter. *Id.* at ¶ 34. Specifically, at least the following additional elements and their ordered combination were not “well-understood, routine, or conventional” as of 2011:

- a Radius gateway for translating a Radius message into a Network Access Server Request (NASReq) message and for transmitting the NASReq message;
- a Diameter Proxy Agent within a Policy and Charging Rules Function (PCRF) server for receiving a NASReq message from the Radius gateway, for selecting one of at least one PCRF cluster within the PCRF server, and for forwarding the NASReq message to the selected PCRF cluster; and
- at least one PCRF blade within the PCRF server, each PCRF blade belonging to one of the at least one PCRF cluster and configured to handle communication sessions for wireless devices, each PCRF blade for receiving a NASReq message from the Diameter Proxy Agent and for creating or updating a NASReq session object related to a communication session identified by the NASReq message received by the PCRF blade in response to receiving the NASReq message.

Ex. F at ¶ 32.

42. For example, during patent examination, the examiner indicated allowance in the first office action as to the combination of additional elements, finding that the cited art did not disclose the ordered combination of a wireline system, wireless system, “in addition to NASReq and PCRF.” *See* Chrissan Declaration Ex. 3 at 8. That the Patent Office found the additional elements were not disclosed in the closest prior art evidences their non-routineness and unconventionality. Ex. F at ¶ 35. This provides “something more” sufficient to transform any alleged abstract idea into patent-eligible subject matter. *Id.*

43. The ’721 Patent is directed to improvements to optical transmission devices, and more specifically, “the control of Raman amplification via the adjustment of the power levels of Raman pump lasers which may be used to compensate for fiber loss in broadband optical

transmission systems.” ’721 Patent at 1:7-10 (Field of Invention); Ex. F at ¶ 36. “Optical systems are used to transmit information by sending pulses of light through an optical fiber.” Ex. F at ¶ 36. Their use is ubiquitous in the telecommunications industry, including for example, in telephony, internet, and cable television. The ’721 Patent describes that optical transmission over long distances requires signal amplification due to “signal distortion and noise growth do not allow for transmission over very long optical transmission links.” ’721 Patent at 1:25-27. Raman amplification involves transferring energy from a higher-power optical pump laser to lower-power payload signals. *Id.* at 1:34-49. Because of various factors, the Raman pump laser power must be dynamically and tightly controlled to be effective at amplification. *Id.* at 49-58. The ’721 Patent explained that existing Raman amplification techniques “can be slow to converge, may achieve non-optimal solutions at convergence and are sensitive to small (and large) measurement errors or noise.” *Id.* at 2:27-30. “These stated problems are inherently technological, i.e., rooted in optical communications technology, rather than human activity, mathematical concepts, or law of nature.” Ex. F at ¶ 37.

44. To solve these problems, the inventors disclosed an improved amplification technique “by projecting the current deviation from the target output onto the subspace of correctable changes so as to achieve a target output spectrum in a faster, more stable manner.” ’721 Patent at 2:64-67. The improved amplification technique is described and claimed as a specific solution, with claim 19 reciting a processor configured to perform the following specific operations:

- receive measured channel powers and to determine deviations of respective measured channel powers from respective target channel powers,

- wherein the processor is further configured to project the deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed, and
- wherein the processor is further configured to determine power setting values for the pump lasers based on the projected deviations.

Id., cl. 19; Ex. F at ¶ 38.

45. The claimed solution “does not claim the mere result of improved Raman amplification but a specific ‘how’ involving deviations from target channel powers, projections of such deviation into Raman gain profile space, and determining power setting values based on the projected deviations.” Ex. F at ¶ 39. Thus, “claim 19 [recites] a specific solution to the technology-rooted problem rather than a functional result.” *Id.* “While claim 19 involves Raman gain profiles, it does not attempt to claim Raman phenomena generally.” *Id.* Instead, “claim 19 sets forth a specific and practical application of Raman-enabled technology to improve such technology.” *Id.* Accordingly, “claim 19 is directed to an improvement to optical communications technology and not an abstract idea or law of nature.” *Id.*

46. Moreover, claim 19 recites additional elements that provide inventive concepts and transform the claim into patent-eligible subject matter. Ex. F at ¶ 40. Specifically, at least the following additional elements and their ordered combination were not “well-understood, routine, or conventional” as of 2009:

- wherein the processor is further configured to project the deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed, and

- wherein the processor is further configured to determine power setting values for the pump lasers based on the projected deviations.

Id.

47. For example, during patent examination, the examiner stated, “[t]he cited prior art does not teach or suggest a processor configured to project deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed.” *See* Chrissan Declaration Ex. 4 (’721 File History) at 4. That the Patent Office found that the additional elements were not disclosed by the closest cited art evidences their non-routineness and unconventionality. Ex. F at ¶ 41. This is sufficient to transform any alleged judicial exception into patent eligible subject matter. *Id.* Accordingly, claim 19 of the ’721 Patent is patent-eligible.

48. The ’691 Patent is directed to improvements to computer network technology, and more particularly, a technique for protecting Multiprotocol Label Switching (MPLS) against faults and associated downtime. Ex. F at ¶ 42. Even more specifically, the ’691 Patent is directed to “providing a standby Label Switched Path which operates to protect a Backup Label Switched Path.” ’691 Patent at 1:8-10. According to the ’691 Patent, MPLS Fast Reroute is a network resiliency mechanism that protects the network against failures at network nodes. *Id.* at 1:21-24. The inventors recognized, however, that MPLS Fast Reroute does not protect against double failures, which would result in network outages and downtime. *Id.* at 2:5-6. “This problem is rooted in communications network technology because it deals with vulnerabilities in network components rather than some longstanding human activity.” Ex. F at ¶ 42.

49. To solve this double-fault problem, the inventors introduced a technique that “for providing a Backup Label Switched Path (LSP) to an already established Bypass LSP.” *Id.*, cl. 6;

see also id. at 2:15-17. Claim 6 recites a specific solution to this double-fault problem, including instruction for:

- protecting the Primary LSP against dual failures, comprising:
- establishing the Bypass LSP for the Protected Primary LSP having a Point of Local Repair node and a Merge Point node;
- obtaining the nodes traversed by an end-to-end path of said Bypass LSP from said Point of Local Repair Node to said Merge Point node;
- generating a request to a path calculator using the nodes traversed by said end-to-end path of said Bypass LSP for a disjoint path connecting said Point of Local Repair Node to said Merge Point node;
- receiving a response from said path calculator; and
- in response to determining that a fully disjoint path connecting said Point of Local Repair Node to said Merge Point node is available, signaling, to at least one other MPLS label switch router, said fully disjoint path as the Backup LSP to said Bypass LSP.

Id., cl. 6; Ex. F at ¶ 43.

50. The claimed solution does not claim the result of protecting against double-faults but a specific how: establishing the Bypass LSP, obtaining the nodes traversed by the Bypass LSP from a Point of Local Repair and Merge Point, generating a request for a disjoint path, signaling to an MPLS label switch router the disjoint path as the Back LSP to the Bypass LSP. *Id.* at ¶ 44. Thus, claim 6 recites a specific solution to the computer-rooted problem rather than a functional result. *Id.* Accordingly, claim 6 is directed to an improvement to computer technology and not an abstract idea. *Id.*

51. Moreover, claim 6 recites additional elements that provide inventive concepts and transform the claim into patent-eligible subject matter. *Id.* at ¶ 45. Specifically, at least the following additional elements and their ordered combination were not “well-understood, routine, or conventional” as of 2012:

- protecting the Primary LSP against dual failures, comprising:
- establishing the Bypass LSP for the Protected Primary LSP having a Point of Local Repair node and a Merge Point node;
- obtaining the nodes traversed by an end-to-end path of said Bypass LSP from said Point of Local Repair Node to said Merge Point node;
- generating a request to a path calculator using the nodes traversed by said end-to-end path of said Bypass LSP for a disjoint path connecting said Point of Local Repair Node to said Merge Point node;
- receiving a response from said path calculator; and
- in response to determining that a fully disjoint path connecting said Point of Local Repair Node to said Merge Point node is available, signaling, to at least one other MPLS label switch router, said fully disjoint path as the Backup LSP to said Bypass LSP.

Ex. F at ¶ 45; *id.* at ¶ 45.

52. For example, during patent examination, the examiner indicated allowance because the closest prior art did not disclose “determining if a backup to a bypass LSP can be a fully disjoint path to prevent a dual/double failure in the MPLS network” or “establishing a bypass LSP that protects a primary LSP and calculating another backup LSP to act as a backup for the bypass LSP while determining if the backup LSP and the bypass LSP are fully disjoint paths connecting the

PLR and MP nodes to ensure protection against dual failures.” *See* Chrissan Declaration Ex. 5 at 8 (’691 File History). The Patent Office finding that the closest prior art did not disclose these additional elements evidences their non-routineness and unconventionality. Ex. F at ¶ 46. Accordingly, claim 6 recites additional elements that provide “something more” than the alleged abstract idea, transforming the claim into patent-eligible subject matter. *Id.*

53. The ’884 Patent is directed to improvements to computer network technology, and more particularly, a technique for “adjusting bandwidth allocation by a network element in a communications network.” ’884 Patent at 1:46-47. According to the ’884 Patent, “most data center Switches do not address network congestion, which may result in packet losses and may affect a quality of service (QoS) of some data flows.” *Id.* at 1:24-26. The inventors recognized that existing solutions, e.g., Explicit Congestion Notification (ECN) “may not be scalable in terms of the number of concurrent data flows.” *Id.* at 1:35-36. “This problem is rooted in computer network technology because it deals with the nature of ‘momentary data bursts’ at data centers and the technical limitations of data center switches. Ex. F at ¶ 47 (citing ’884 Patent. at 1:24-33). As a result, the problem addressed by the inventors is not some longstanding human activity. *See id.*

54. To solve this network congestion problem, the inventors introduced a specific technique that “adjusting the bandwidth allocation for the target port based on the fair-share bandwidth allocation.” ’884 Patent at 1:59-60. Claim 17 recites a specific apparatus with specific functions to solve the stated network congestion problem, including an edge switch configured to:

- monitor a data flow traversing the target port;
- determine a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow;

- determine a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and
- adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.

Id., cl. 17; Ex. F at ¶ 48.

55. The claimed solution does not claim the result of preventing packet loss and congestion but a specific how: monitoring data flow, determining a bandwidth allocation that is currently allocated to a target port for the data flow, determining a fair fair-share bandwidth allocation, which is proportional to the total bandwidth of the network switching element, and adjusting the bandwidth allocation for the target port based on the fair-share bandwidth allocation. Ex. F at ¶ 49. Thus, claim 17 recites a specific solution to the computer-rooted problem rather than a functional result. *Id.* Accordingly, claim 17 is directed to an improvement to computer technology and not an abstract idea.

56. Moreover, claim 17 recites additional elements that provide inventive concepts and “transform the claim into patent-eligible subject matter.” *Id.* at ¶ 50. Specifically, at least the following additional elements and their ordered combination were not “well-understood, routine, or conventional” as of 2014:

- determine a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow;
- determine a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and

- adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.

Id. at ¶ 50.

57. For example, during patent examination, the examiner indicated allowance because the closest prior art “fail to anticipate or render obvious a method comprising: "monitoring, by the network switching element, a data flow traversing the target port of the network switching element," and "determining, by the network switching element, a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element," in combination with all other limitations in the claim as claimed and defined by applicants.” *See* Chrissan Declaration Ex. 6 ('884 File History) at 4. The Patent Office finding that the closest prior art did not disclose these additional elements evidences their non-routineness and unconventionality. Ex. F at ¶ 51. Accordingly, claim 17 recites additional elements that provide “something more” than any alleged abstract idea, transforming the claim into patent-eligible subject matter. *Id.*

FIRST CAUSE OF ACTION
(PATENT INFRINGEMENT UNDER 35 U.S.C. § 271 OF THE
'630 PATENT BY CISCO)

58. Brazos re-alleges and incorporates by reference all of the foregoing paragraphs.

59. On information and belief, Cisco has directly infringed and continues to directly infringe either literally or under the doctrine of equivalents, one or more claims, including at least claim 18, of the '630 Patent in violation of 35 U.S.C. § 271, et seq., by deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling at least the Cisco Catalyst 9000 Switching Platforms, and similar products, that supports Multiprotocol Label Switching (MPLS) Quality of Service (QoS) and DiffServ tunneling, by use of Cisco IOS XE software and the Cisco

Unified Access Data Plane (UADP) Application-Specific Integrated Circuit (ASIC) (the “Accused DS-TE Products”).

60. Claim 18 of the ’630 Patent provides:

[Preamble] A method comprising:

[18A] defining a mapping policy configured to map between an experimental field and a unique per-hop-behavior;

[18B] defining a customer policy comprising a tunneling mode and a tunnel group identifier, the customer policy being configured to govern the treatment of individual customer traffic;

[18C] defining a network policy that is configured to define the Diffserv treatment of aggregated traffic;

[18D] translating the mapping policy, the network policy and the customer policy into device-specific commands; and

[18E] sending the device-specific commands to policy targets, wherein each policy target comprises a network device that includes an interface assigned a role name associated with the customer policy, at least one of the interfaces comprising an egress interface of one of multi-protocol label switching tunnels.

61. On information and belief, Cisco performs each and every limitation of at least claim 18 of the ’630 Patent as stated below.

62. On information and belief, and to the extent possible that the preamble of claim 18 is determined to be limiting, the Cisco performs a “method.” For example, Cisco documents provide:¹

¹ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 2.

Functional Description of MPLS

Label switching is a high-performance packet forwarding technology that integrates the performance and traffic management capabilities of data link layer (Layer 2) switching with the scalability, flexibility, and performance of network layer (Layer 3) routing.

Label Switching Functions

In conventional Layer 3 forwarding mechanisms, as a packet traverses the network, each router extracts all the information relevant to forwarding the packet from the Layer 3 header. This information is then used as an index for a routing table lookup to determine the next hop for the packet.

In the most common case, the only relevant field in the header is the destination address field, but in some cases, other header fields might also be relevant. As a result, the header analysis must be done independently at each router through which the packet passes. In addition, a complicated table lookup must also be done at each router.

In label switching, the analysis of the Layer 3 header is done only once. The Layer 3 header is then mapped into a fixed length, unstructured value called a *label*.

Many different headers can map to the same label, as long as those headers always result in the same choice of next hop. In effect, a label represents a *forwarding equivalence class* --that is, a set of packets which, however different they may be, are indistinguishable by the forwarding function.

The initial choice of a label need not be based exclusively on the contents of the Layer 3 packet header; for example, forwarding decisions at subsequent hops can also be based on routing policy.

Once a label is assigned, a short label header is added at the front of the Layer 3 packet. This header is carried across the network as part of the packet. At subsequent hops through each MPLS router in the network, labels are swapped and forwarding decisions are made by means of MPLS forwarding table lookup for the label carried in the packet header. Hence, the packet header does not need to be reevaluated during packet transit through the network. Because the label is of fixed length and unstructured, the MPLS forwarding table lookup process is both straightforward and fast.

63. Thus, to the extent the preamble of claim 18 is limiting, Cisco performs the preamble of claim 18 by using the Accused DS-TE Products.

64. On information and belief, Cisco performs claim element [18A] of claim 18 of the '630 Patent, "defining a mapping policy configured to map between an experimental field and a unique per-hop-behavior;" by using the Accused DS-TE Products. For example, Cisco documentation describes that the "Cisco IOS® XE network Operating System (OS) is the single OS for enterprise switching, routing, wired and wireless access."² Cisco uses Multiprotocol Label Switching (MPLS), as describes in the following Cisco documentation:³

² See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 9.

³ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 1.

MPLS Overview

Multiprotocol label switching (MPLS) combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. MPLS enables service providers to meet the challenges of explosive growth in network utilization while providing the opportunity to differentiate services without sacrificing the existing network infrastructure. The MPLS architecture is flexible and can be employed in any combination of Layer 2 technologies. MPLS support is offered for all Layer 3 protocols, and scaling is possible well beyond that typically offered in today's networks.

65. Cisco use MPLS Traffic Engineering, including DiffServ Aware (DS-TE), to define a mapping policy configured to map between an experimental field and a unique per-hop-behavior. For example, DS-TE provides differentiated service using DS-TE global pool tunnels.⁴ An administrative user can configure an accused device in many ways, including as follows:⁵

Providing Differentiated Service Using DS-TE Global Pool Tunnels

You can configure a tunnel using global pool bandwidth to carry best-effort as well as several other classes of traffic. Traffic from each class can receive differentiated service if you do all of the following:

1. Select a separate queue (a distinct diffserv PHB) for each traffic class. For example, if there are three classes (gold, silver, and bronze) there must be three queues (diffserv AF2, AF3, and AF4).
2. Mark each class of traffic using a unique value in the MPLS experimental bits field (for example gold = 4, silver = 5, bronze = 6).
3. Ensure that packets marked as Gold are placed in the gold queue, Silver in the silver queue, and so on. The tunnel bandwidth is set based on the expected aggregate traffic across all classes of service.

To control the amount of diffserv tunnel traffic you intend to support on a given link, adjust the size of the global pool on that link.


66. Furthermore, mapping policies can be configured using the Modular QoS Command Line Model (MQC), which allows a user to generate policy maps that define the

⁴ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 8.

⁵ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 8; see also *id.* at 22, 23, 25, 26, 34, 37, 51, 65-72, *passim* (showing additional examples of experimental field and a unique per-hop-behavior mapping policies).

relationship between an MPLS experimental field and a unique per-hop-behavior.⁶ As described in the “Cisco Catalyst 9000 Switching Platforms: QoS and Queuing” document:⁷

Every MQC policy is based on a class map, a policy map, and an interface target where the policy map will be attached. Figure 13 shows an MQC policy structure.



67. Therefore, Cisco performs element [18A] of claim 18 of the '630 patent by using the Accused DS-TE Products.

68. On information and belief, Cisco performs claim element [18B] of claim 18 of the '630 Patent, “defining a customer policy comprising a tunneling mode and a tunnel group identifier, the customer policy being configured to govern the treatment of individual customer traffic,” by using the Accused DS-TE Products. For example, the accused products feature “MPLS

⁶ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 15-16.

⁷ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 15-16.

DiffServ Tunneling Modes.”⁸ Cisco’s MPLS Traffic Engineering DiffServ Configuration Guide explains:⁹

MPLS DiffServ Tunneling Modes

MPLS DiffServ Tunneling Modes allows service providers to manage the quality of service (QoS) that a router will provide to a Multiprotocol Label Switching (MPLS) packet in an MPLS network. MPLS DiffServ Tunneling Modes conforms to the IETF draft standard for Uniform, Short Pipe, and Pipe modes. It also conforms to Cisco-defined extensions for scalable command line interface (CLI) management of those modes at customer edge, provider edge, and core routers.

The following features are supported on MPLS DiffServ Tunneling Modes:

- MPLS per-hop behavior (PHB) layer management.
- There is improved scalability of the MPLS layer management by control on managed customer edge (CE) routers.
- MPLS can “tunnel” a packet’s QoS (that is, the QoS is transparent from edge to edge).
- The MPLS experimental (MPLS EXP) field can be marked differently and independently of the PHB marked in the IP Precedence or differentiated services code point (DSCP) field.
- There are three MPLS QoS tunneling modes for the operation and interaction between the DiffServ marking in the IP header and the DiffServ marking in the MPLS header: Pipe mode with an explicit NULL LSP, Short Pipe mode, and Uniform mode. Pipe mode with an explicit NULL LSP and Short Pipe mode allow an MPLS network to transparently tunnel the DiffServ marking of packets.

MPLS DiffServ Tunneling Modes has the following benefits:

- Tunneling modes provide added QoS functionality by the creative manipulation of the MPLS EXP field during label imposition, forwarding, and label disposition.
- Tunneling modes provide a common set of PHBs to different service provider customers.
- Pipe mode provides transparency and customized edge service.
- Pipe mode with an explicit NULL LSP improves the scalability of management by performing per-customer packet metering and marking closer to the service provider’s customer networks.
- Pipe mode with an explicit NULL LSP provides QoS transparency by ensuring that customer’s packets will not be re-marked in the service provider’s network.
- In Pipe mode with an explicit NULL LSP, the explicit NULL LSP applies the service provider’s PHBs on the ingress CE-to-PE link.

⁸ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 51-96 (Chapter 3).

⁹ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 51.


69. Further, the Tunneling Modes include pipe mode, short pipe mode, and uniform mode:¹⁰

There are three ways to forward packets through a network:

- Pipe mode with an explicit NULL LSP
- Short Pipe mode
- Uniform mode

70. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains that the Pipe mode is associated with an explicit NULL LSP as follows:¹¹

Figure 2: Pipe Mode with an Explicit NULL LSP Operation with MPLS VPN Enabled



¹⁰ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 55.

¹¹ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 59-60.


Pipe mode with an explicit NULL LSP functions as follows. The circled numbers at the bottom of the illustration correspond to the step numbers.

- 1 IP packets arrive at the router CE1, the managed CE router, with a DSCP value of 1.
- 2 An explicit NULL label entry is imposed onto the packet that contains an EXP value of 5.
- 3 The packet is transmitted to PE1 on the explicit NULL LSP.
- 4 The PE1 router saves the value of the MPLS EXP field and removes the explicit NULL entry. The PE1 router then imposes new labels onto the IP packet. Each label entry is set to the saved MPLS EXP field 5.
- 5 The packet is transmitted to P1.
- 6 At P1, the received EXP value is copied into the swapped label entry.
- 7 The packet is transmitted to P2.
- 8 At P2, the topmost label is popped, exposing a label entry that also has an EXP value of 5.
- 9 The packet is transmitted to PE2.
- 10 PE2 stores the value of the MPLS EXP field in the qos-group and discard-class variables, and removes the label entry from the packet.
- 11 While transmitting the packet to CE2, PE2 does QoS on its egress interface based on the saved value of the MPLS EXP field (qos-group and discard-class).
- 12 The IP packet arrives at the CE2 router.

71. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains the short

pipe mode as follows:¹²

Figure 4: Short Pipe Mode Operation




¹² See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 62-63.

Short Pipe mode functions as follows. The circled numbers at the bottom of the illustration correspond to the step numbers.

- 1 C1, CE1 transmits an IP packet to PE1 with an IP DSCP value of 1.
- 2 PE1 sets the MPLS EXP field to 5 in the imposed label entries.
- 3 PE1 transmits the packet to P1.
- 4 P1 sets the MPLS EXP field value to 5 in the swapped label entry.
- 5 P1 transmits the packet to P2.
- 6 P2 pops the IGP label entry.
- 7 P2 transmits the packet to PE2.
- 8 PE2 pops the BGP label.
- 9 PE2 transmits the packet to C1, CE2, but does QoS based on the IP DSCP value.

72. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains the Uniform mode as follows:¹³

¹³ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 64-65.

Figure 5: Uniform Mode Operation

*In both the MPLS-to-MPLS and the MPLS-to-IP cases, the PHBs of the topmost popped label is copied into the new top label or the IP DSCP if no label remains

The procedure varies according to whether there are IP Precedence bit markings or DSCP markings.

The following actions occur if there are IP Precedence bit markings:


- 1 IP packets arrive in the MPLS network at PE1, the service provider edge router.
- 2 A label is copied onto the packet.
- 3 If the MPLS EXP field value is recolored (for example, if the packet becomes out-of-rate because too many packets are being transmitted), that value is copied to the IGP label. The value of the BGP label is not changed.
- 4 At the penultimate hop, the IGP label is removed. That value is copied into the next lower level label.
- 5 When all MPLS labels have been removed from the packet which is sent out as an IP packet, the IP Precedence or DSCP value is set to the last changed EXP value in the core.

Following is an example when there are IP precedence bit markings:

- 1 At CE1 (customer equipment 1), the IP packet has an IP Precedence value of 5.
- 2 When the packet arrives in the MPLS network at PE1 (the service provider edge router), the IP Precedence value of 5 is copied to the imposed label entries of the packet.
- 3 The MPLS EXP field in the IGP label header might be changed within the MPLS core (for example, at P1).
 - 1 At P2, when the IGP label is removed, the MPLS EXP field in this label entry is copied into the underlying BGP label.
 - 2 At PE2, when the BGP label is popped, the EXP field in this label header is copied into the IP Precedence field of the underlying IP header.

73. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains that the tunneling modes are configured in a variety of ways depending on customer policy.¹⁴

74. Further, MPLS allows for a plurality of tunnels to be created:¹⁵



75. MPLS tunnels can be grouped and given an identifier that references a collection of label switched paths (LSPs).¹⁶

¹⁴ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf 65-87 (“How to Configure MPLS DiffServ Tunneling Modes”) and 88-92 (“Configuration Examples for MPLS DiffServ Tunneling Modes”).

¹⁵ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 12-14.

¹⁶ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 12-14; https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 55-56; <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 56-57.

MPLS Transport Profile Links and Physical Interfaces

Multiprotocol Label Switching Transport Profile (MPLS-TP) link numbers may be assigned to physical interfaces only. Bundled interfaces and virtual interfaces are not supported for MPLS-TP link numbers.

The MPLS-TP link creates a layer of indirection between the MPLS-TP tunnel and midpoint LSP configuration and the physical interface. The **mplstp link** command is used to associate an MPLS-TP link number with a physical interface and next-hop node. The MPLS-TP out-links can be configured only on the ethernet interfaces, with either the next hop IPv4 address or next hop mac-address specified.

Multiple tunnels and LSPs may then refer to the MPLS-TP link to indicate that they are traversing that interface. You can move the MPLS-TP link from one interface to another without reconfiguring all the MPLS-TP tunnels and LSPs that refer to the link.

Link numbers must be unique on the router or node.

76. Further, packet traffic may be configured and identified in a number of ways and associated with a customer policy for that traffic. For example, the MQC (Modular Quality of Service (QoS) Command-Line Interface (CLI)) enables users to set packet classification and marking based on a QoS group value. MQC CLI allows users to create traffic classes and policies, enable a QoS feature (such as packet classification), and attach these policies to interfaces:¹⁷

The MQC Structure

The MQC (Modular Quality of Service (QoS) Command-Line Interface (CLI)) enables you to set packet classification and marking based on a QoS group value. MQC CLI allows you to create traffic classes and policies, enable a QoS feature (such as packet classification), and attach these policies to interfaces.

The MQC structure necessitates developing the following entities: traffic class, policy map, and service policy.

Elements of a Traffic Class

A traffic class contains three major elements: a traffic class name, a series of **match** commands, and, if more than one **match** command is used in the traffic class, instructions on how to evaluate these **match** commands.

77. Traffic classes, or class maps, can be defined based on a number of criteria, such as Access Control Lists or an MPLS Experimental Value:¹⁸

¹⁷ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mqc/configuration/xs/asr903/16-12-1/b-qos-mqc-cli-xe-16-12-asr900.pdf at 2.

¹⁸ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mqc/configuration/xs/asr903/16-12-1/b-qos-mqc-cli-xe-16-12-asr900.pdf at 2.

Available match Commands

The table below lists *some* of the available **match** commands that can be used with the MQC. The available **match** commands vary by Cisco IOS XE release. For more information about the commands and command syntax, see the *Cisco IOS Quality of Service Solutions Command Reference*.

Table 1: match Commands That Can Be Used with the MQC

Command	Purpose
match access-group	Configures the match criteria for a class map on the basis of the specified access control list (ACL).
match cos	Matches a packet based on a Layer 2 class of service (CoS) marking.
match discard-class	Matches packets of a certain discard class.
match [ip] dscp	Identifies a specific IP differentiated service code point (DSCP) value as a match criterion. Up to eight DSCP values can be included in one match statement.
match mpls experimental Note The match mpls experimental command is <i>not</i> supported on the Cisco RSP3 Module.	Configures a class map to use the specified value of the Multiprotocol Label Switching (MPLS) experimental (EXP) field as a match criterion.
match mpls experimental topmost	Matches the MPLS EXP value in the topmost label.

78. Further, the Cisco Catalyst 9000 family switches support classification using: Access Control Lists (ACLs) (source/destination IP, TCP/UDP ports, and more), DSCP, IP precedence, Traffic class, CoS, the MPLS EXP field, Network-Based Application Recognition (NBAR) protocols and VLANs:¹⁹

¹⁹ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 18.

Classification

Ingress classification is applied when the packet comes into the switch and policy is applied on the interface.

Cisco Catalyst 9000 family switches support classification using:

- Access Control Lists (ACLs) (source/destination IP, TCP/UDP ports, and more)
- DSCP
- IP precedence
- Traffic class
- CoS
- EXP (from 17.3.1)
- Network-Based Application Recognition (NBAR) protocols
- VLANs

79. Therefore, Cisco performs element [18B] of claim 18 of the '630 Patent by using the Accused DS-TE Products.

80. On information and belief, Cisco performs claim element [18C] of claim 18 of the '630 Patent, “defining a network policy that is configured to define the Diffserv treatment of aggregated traffic,” by using the Accused TS-DE Products.

81. For example, the accused products feature “MPLS DiffServ Tunneling Modes.”²⁰ Cisco’s MPLS Traffic Engineering DiffServ Configuration Guide explains:²¹

²⁰ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 51-96.

²¹ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 51.

MPLS DiffServ Tunneling Modes allows service providers to manage the quality of service (QoS) that a router will provide to a Multiprotocol Label Switching (MPLS) packet in an MPLS network. MPLS DiffServ Tunneling Modes conforms to the IETF draft standard for Uniform, Short Pipe, and Pipe modes. It also conforms to Cisco-defined extensions for scalable command line interface (CLI) management of those modes at customer edge, provider edge, and core routers.

The following features are supported on MPLS DiffServ Tunneling Modes:

- MPLS per-hop behavior (PHB) layer management.
- There is improved scalability of the MPLS layer management by control on managed customer edge (CE) routers.
- MPLS can “tunnel” a packet’s QoS (that is, the QoS is transparent from edge to edge).
- The MPLS experimental (MPLS EXP) field can be marked differently and independently of the PHB marked in the IP Precedence or differentiated services code point (DSCP) field.
- There are three MPLS QoS tunneling modes for the operation and interaction between the DiffServ marking in the IP header and the DiffServ marking in the MPLS header: Pipe mode with an explicit NULL LSP, Short Pipe mode, and Uniform mode. Pipe mode with an explicit NULL LSP and Short Pipe mode allow an MPLS network to transparently tunnel the DiffServ marking of packets.

MPLS DiffServ Tunneling Modes has the following benefits:

- Tunneling modes provide added QoS functionality by the creative manipulation of the MPLS EXP field during label imposition, forwarding, and label disposition.
- Tunneling modes provide a common set of PHBs to different service provider customers.
- Pipe mode provides transparency and customized edge service.
- Pipe mode with an explicit NULL LSP improves the scalability of management by performing per-customer packet metering and marking closer to the service provider’s customer networks.
- Pipe mode with an explicit NULL LSP provides QoS transparency by ensuring that customer’s packets will not be re-marked in the service provider’s network.
- In Pipe mode with an explicit NULL LSP, the explicit NULL LSP applies the service provider’s PHBs on the ingress CE-to-PE link.
- In Pipe mode with an explicit NULL LSP, the service provider’s PHBs are applied on the egress PE-to-CE link.

82. Cisco documentation explains that there are three Tunneling Modes for MPLS DiffServ, which correspond to three ways to forward packets through a network:²²


²² See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 55.

There are three ways to forward packets through a network:

- Pipe mode with an explicit NULL LSP
- Short Pipe mode
- Uniform mode

83. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains that the Pipe mode is associated with an explicit NULL LSP as follows:²³

Figure 2: Pipe Mode with an Explicit NULL LSP Operation with MPLS VPN Enabled



Pipe mode with an explicit NULL LSP functions as follows. The circled numbers at the bottom of the illustration correspond to the step numbers.


- 1 IP packets arrive at the router CE1, the managed CE router, with a DSCP value of 1.
- 2 An explicit NULL label entry is imposed onto the packet that contains an EXP value of 5.
- 3 The packet is transmitted to PE1 on the explicit NULL LSP.

²³ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 59-60.

- 4 The PE1 router saves the value of the MPLS EXP field and removes the explicit NULL entry. The PE1 router then imposes new labels onto the IP packet. Each label entry is set to the saved MPLS EXP field 5.
- 5 The packet is transmitted to P1.
- 6 At P1, the received EXP value is copied into the swapped label entry.
- 7 The packet is transmitted to P2.
- 8 At P2, the topmost label is popped, exposing a label entry that also has an EXP value of 5.
- 9 The packet is transmitted to PE2.
- 10 PE2 stores the value of the MPLS EXP field in the qos-group and discard-class variables, and removes the label entry from the packet.
- 11 While transmitting the packet to CE2, PE2 does QoS on its egress interface based on the saved value of the MPLS EXP field (qos-group and discard-class).
- 12 The IP packet arrives at the CE2 router.

84. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains the short pipe mode as follows:²⁴

Figure 4: Short Pipe Mode Operation




²⁴ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xr-16-9/mp-te-diffserv-xr-16-9-book.pdf at 62-63.

Short Pipe mode functions as follows. The circled numbers at the bottom of the illustration correspond to the step numbers.

- 1 C1, CE1 transmits an IP packet to PE1 with an IP DSCP value of 1.
- 2 PE1 sets the MPLS EXP field to 5 in the imposed label entries.
- 3 PE1 transmits the packet to P1.
- 4 P1 sets the MPLS EXP field value to 5 in the swapped label entry.
- 5 P1 transmits the packet to P2.
- 6 P2 pops the IGP label entry.
- 7 P2 transmits the packet to PE2.
- 8 PE2 pops the BGP label.
- 9 PE2 transmits the packet to C1, CE2, but does QoS based on the IP DSCP value.

85. Cisco's MPLS Traffic Engineering DiffServ Configuration Guide explains the Uniform mode as follows:²⁵

Figure 5: Uniform Mode Operation



²⁵ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 64-65.

The procedure varies according to whether there are IP Precedence bit markings or DSCP markings.

The following actions occur if there are IP Precedence bit markings:

- 1 IP packets arrive in the MPLS network at PE1, the service provider edge router.
- 2 A label is copied onto the packet.
- 3 If the MPLS EXP field value is recolored (for example, if the packet becomes out-of-rate because too many packets are being transmitted), that value is copied to the IGP label. The value of the BGP label is not changed.
- 4 At the penultimate hop, the IGP label is removed. That value is copied into the next lower level label.
- 5 When all MPLS labels have been removed from the packet which is sent out as an IP packet, the IP Precedence or DSCP value is set to the last changed EXP value in the core.

Following is an example when there are IP precedence bit markings:

- 1 At CE1 (customer equipment 1), the IP packet has an IP Precedence value of 5.
- 2 When the packet arrives in the MPLS network at PE1 (the service provider edge router), the IP Precedence value of 5 is copied to the imposed label entries of the packet.
- 3 The MPLS EXP field in the IGP label header might be changed within the MPLS core (for example, at P1).
- 1 At P2, when the IGP label is removed, the MPLS EXP field in this label entry is copied into the underlying BGP label.
- 2 At PE2, when the BGP label is popped, the EXP field in this label header is copied into the IP Precedence field of the underlying IP header.

86. The MPLS DiffServ Tunneling Modes are network policies that are configured to define the Diffserv treatment of aggregated traffic.²⁶


87. Therefore, Cisco performs element [18C] of claim 18 of the '630 Patent by using the Accused DS-TE Products.

²⁶ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 88-92 (“Configuration Examples for MPLS DiffServ Tunneling Modes”). See also <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-12-1/b-mpl-basic-16-12-1-ncs4200.pdf> at 1-10 (“Multiprotocol Label Switching (MPLS) on Cisco Routers”) and 11-40 (“MPLS Transport Profile”); <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 56-59; <https://www.cisco.com/c/dam/en/us/products/collateral/enterprise-networks/nb-06-ios-xe-prog-ebook-cte-en.pdf> at 68-79, 125-130.

88. On information and belief, Cisco performs claim element [18D] of claim 18 of the '630 Patent, "translating the mapping policy, the network policy and the customer policy into device-specific commands," by using the Accused TS-DE Products. The discussion for elements [18A]-[18C] are incorporated herein.

89. For example, the configured MPLS DiffServ Tunneling Modes result in translating the mapping policy, the network policy and the customer policy into device-specific commands, which are then sent over the control plane to MPLS devices. Examples of these MPLS devices are shown below, including Provider Edge routers PE1 and PE2, and core routers P1 and P2:²⁷

Figure 5: Uniform Mode Operation



90. The Cisco Modular QoS Command Line Model (MQC) uses a service-policy command to configure data plane interfaces according to traffic policies.²⁸

²⁷ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 64.

²⁸ See, e.g., <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 21-22; <https://www.cisco.com/c/en/us/td/docs/ios->

91. Therefore, Cisco performs element [18D] of claim 18 of the '630 Patent by using the Accused DS-TE Products.


92. On information and belief, Cisco performs claim element [18E] of claim 18 of the '630 Patent, "sending the device-specific commands to policy targets, wherein each policy target comprises a network device that includes an interface assigned a role name associated with the customer policy, at least one of the interfaces comprising an egress interface of one of multi-protocol label switching tunnels," by using the Accused TS-DE Products. The discussion for elements [18A]-[18D] are incorporated herein.

93. For example, device-specific commands are sent to MPLS devices to configure their data plane. Examples of these MPLS devices are shown in the diagrams below, including Provider Edge routers PE1 and PE2, and core routers P1 and P2:²⁹

[xml/ios/qos_mqc/configuration/xe-3s/asr903/16-12-1/b-qos-mqc-cli-xe-16-12-asr900.pdf](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf) at 6, 11-14; https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 22-23.

²⁹ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 64.

Figure 5: Uniform Mode Operation



*In both the MPLS-to-MPLS and the MPLS-to-IP cases, the PHBs of the topmost popped label is copied into the new top label or the IP DSCP if no label remains

94. The Cisco Modular QoS Command Line Model (MQC) uses a service-policy command to configure data plane interfaces according to traffic policies, which then configures the indicated interfaces.³⁰ On information and belief, the policy provided with a service policy command is associated with an interface and a customer policy, and therefore is a role name.

95. Therefore, Cisco performs element [18E] of claim 18 of the '630 Patent by using the Accused DS-TE Products.

96. Accordingly, Cisco's use of the Accused DS-TE Products satisfies each and every limitation of claim 18 of the '630 Patent.

97. On information and belief, Cisco has directly used the Accused DS-TE Products in Cisco data centers and has directly used the Accused DS-TE Products when setting up, running,

³⁰ See, e.g., <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 21-22; https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mqc/configuration/xe-3s/asr903/16-12-1/b-qos-mqc-cli-xe-16-12-asr900.pdf at 6, 11-14; https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_diffserv/configuration/xe-16-9/mp-te-diffserv-xe-16-9-book.pdf at 22-23.

and troubleshooting the Accused Products on behalf of customers. Support for these statements is located throughout Cisco websites.

98. On information and belief, the Accused DS-TE Products are used in data centers. Multiple pages on Cisco's website link the Accused DS-TE Products and data centers together. For example, the webpage titled "*Cisco Catalyst 9500X Series Switches Hardware Installation Guide*" states that "[t]he switch chassis must be installed in a cabinet or rack that is secured to the data center" in describing how to install the Accused DS-TE Product.³¹ Additionally, there is a "*Troubleshoot MACsec on Catalyst 9000*" webpage which lists use cases for the MACsec product.³² The list of use cases includes data centers; as the page is specific to *MACsec on Catalyst 9000*, this directly links the Accused DS-TE Products to being used in data centers.³³

99. On information and belief, Cisco runs several of its own data centers. Two of these data centers are known to be in Texas, with one in Allen, TX and the other in Richardson, TX.³⁴ On information and belief, the Cisco Data Center in Allen, TX was constructed in 2009.³⁵ A video produced by Cyclone Interactive for Cisco describes the build of the Data Center in Allen, TX, and shows that Cisco's products including but not limited to switches are in use in the Cisco Data

³¹ *Cisco Catalyst 9500X Series Switches Hardware Installation Guide*, CISCO, https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/hardware/install/b-c9500x-hig/9500x_installing-fru.html?dtid=ossdc000283 (last visited May 2, 2024).

³² *Troubleshoot MACsec on Catalyst 9000*, CISCO, <https://www.cisco.com/c/en/us/support/docs/switches/catalyst-9300-series-switches/216849-troubleshoot-macsec-on-catalyst-9000.html?dtid=ossdc000283> (last visited May 2, 2024).

³³ *Id.*

³⁴ *See Cisco Allen*, BAXTEL, <https://baxtel.com/data-center/cisco-allen> (last visited May 2, 2024); *also see Cisco Richardson TXDCI*, BAXTEL, <https://baxtel.com/data-center/cisco-richardson-txdc1> (last visited May 2, 2024).

³⁵ *Enterprise Thought Leadership*, CYCLONE INTERACTIVE, <https://www.cycloneinteractive.com/our-work/cisco-data-center/> (last visited May 2, 2024).

Center.³⁶ On information and belief, Cisco regularly updates equipment in its data centers at the end of its useable life and as such Cisco would maintain and update switches in their facilities to the current available generation.³⁷

100. A document titled “*Cisco Webex Contact Center Enterprise: Infrastructure as a Service (IaaS) Add-On Option Service Description*” describes the features and benefits of the Webex Contact Center Enterprise (“Webex CCE”) Infrastructure as a Service (IaaS) add-on to Cisco Data Center Networking (DCN) solutions.³⁸ This add-on allows customers to “run their IaaS applications in a geographically redundant manner,” as is further described in the screenshot below:³⁹

4. IaaS high availability and disaster recovery services

Customers can decide whether to run their IaaS applications in a geographically redundant manner or not. If a customer elects to run their IaaS applications in a single data center and connectivity to that data center goes down for any reason, the customer would not be able to access their IaaS applications until the data center comes back online.

For IaaS applications that run in a geographically redundant manner, Cisco will maintain the IaaS hardware such that every device is highly available within each of the Cisco data centers. Cisco will provide services on a continuously available basis by utilizing both data centers in production. Accordingly, at no time will both data centers be taken out of production at the same time.

Cisco agrees to maintain viable disaster recovery/business continuity plans for the IaaS services. Cisco will test, at least once every 12 months, plans to continue business and the provision of the IaaS services in the event of an interruption to the IaaS services or unavailability of any site from which IaaS services are being performed (the “disaster recovery/business continuity plans”).

³⁶ Cyclone Interactive, *Cisco DC2011-Texas: All in One Building*, YOUTUBE (June 14, 2012), <https://www.youtube.com/watch?v=F1IUTsMz8I4&t=32s> at 3:42-3:44, 4:00-4:02 (last visited May 2, 2024).

³⁷ *Cisco IT Data Center Sustainability*, CISCO, <https://www.cisco.com/c/en/us/solutions/collateral/enterprise/cisco-on-cisco/cisco-it-dc-sustainability-wp.html> (last visited May 2, 2024).

³⁸ *Cisco Webex Contact Center Enterprise: Infrastructure as a Service (IaaS) Add-On Option Service Description*, CISCO, <https://www.cisco.com/c/en/us/products/collateral/contact-center/webex-contact-center-enterprise/datasheet-c78-743858.pdf> (June, 2020).

³⁹ *Id.* at 4.

101. By duplicating IaaS applications in a geographically redundant manner, Cisco Systems maintains customers' products in its own data centers to ensure that customer products remain "highly available" even in the case of a customer's data center going offline. On information and belief, Cisco directly uses the Accused DS-TE Products in the Cisco data centers which are used to create this geographic redundancy.

102. Additionally, on the webpage for "Cisco Catalyst 9000 Switching Platform FAQ," there are several indications that Cisco directly uses the Accused DS-TE Products when supporting customers and troubleshooting products.⁴⁰ For instance, the "Services" portion of the FAQ informs customers and prospective clients that the services related to the Cisco Catalyst 9000 switches include the provision of "expert guidance to help [users] successfully plan, deploy, manage, and support [their] new switches."⁴¹ Examples from the FAQ page are shown in the screenshots below⁴²:

Q. What is the software support model with E-LLW on the Cisco Catalyst 9000 family?

A. The Enhanced Limited Lifetime Hardware Warranty is limited to hardware support. For the Cisco Catalyst 9000 family, Cisco Catalyst Software term-based subscription license-related issues are covered by SWSS, which is included as part of the Cisco Catalyst software subscription license. Also, for the first 90 days, you are entitled to configuration help for Cisco IOS XE software-related questions. Software updates to the Cisco IOS XE software will be supported by the business unit software policy.

⁴⁰ *Cisco Catalyst 9000 Switching Platform FAQ*, CISCO, <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat9k-swit-plat-faq-cte-en.html> (last visited May 2, 2024).

⁴¹ *Id.*

⁴² *Id.*

Services

Q. Are there any services available to support the Cisco Catalyst 9000 family of switches?

A. Yes. With Cisco Services, you can achieve infrastructure excellence faster with less risk. Our services for Cisco Catalyst 9000 switches provide expert guidance to help you successfully plan, deploy, manage and support your new switches. With unmatched networking expertise, best practices and innovative tools, Cisco Services can help you reduce overall upgrade, refresh, and migration costs as you introduce new hardware, software and protocols into the network. With a comprehensive lifecycle of services, Cisco experts will help you minimize disruption and improve operational efficiency to extract maximum value from your Cisco Catalyst infrastructure.

103. Cisco undertook these infringing actions despite an objectively high likelihood that such activities infringe the '630 Patent, which has been duly issued by the PTO and presumed valid. Cisco also had knowledge of the '630 Patent; Cisco's U.S. Patent No. 8,621,596, filed January 24, 2011 and issued February 13, 2013, cites the '630 Patent on its face.

104. Brazos has previously asserted patents against Cisco, and as such Cisco has been on notice since at least the filing of the previous complaint that Brazos's patent portfolio contains patents infringed by Cisco's products. *See WSOU Investments LLC d/b/a Brazos Licensing & Development v. Cisco Systems, Inc.*, No. 6:21-cv-00128-ADA (W.D. Tex. Feb. 5, 2021). As such, Cisco had knowledge that a number of Cisco's products had a high likelihood of infringement. Despite this, Cisco continued its infringing activities.

105. On information and belief, Cisco could not reasonably, subjectively believe that its actions do not constitute infringement of the '630 Patent. Despite that knowledge and subjective belief, and the objectively high likelihood that its actions constitute infringement, Cisco continued its infringing activities. As such, Cisco has willfully infringed the '630 Patent.

106. Since at least the date of first learning of the '630 Patent, through its actions, Cisco has also indirectly infringed and continued to indirectly infringe the '630 Patent in violation of 35 U.S.C. § 271(b). Cisco has actively induced product makers, distributors, retailers, and/or end users of the Accused DS-TE Products to directly infringe the '630 Patent by performing the method of claim 18 as detailed above throughout the United States, including within this Judicial District,

by, among other things, advertising and promoting the use of the Accused DS-TE Products in various websites, including providing and disseminating product descriptions, operating manuals, and other instructions on how to implement and configure the Accused DS-TE Products. Examples of such advertising, promoting, and/or instructing include the documents cited in the paragraphs above. Cisco did so knowing and intending that its customers and end users commit these infringing acts, despite its knowledge of the '630 Patent, thereby specifically intending for and inducing its customers to infringe the '630 Patent through the customers' normal and customary use of the Accused DS-TE Products.

107. As a result of Cisco's infringement of the '630 Patent, Brazos has suffered substantial injury and is entitled to recover all damages caused by Cisco's infringement to the fullest extent permitted by the Patent Act, together with prejudgment interest and costs for Cisco's wrongful conduct.

SECOND CAUSE OF ACTION
(PATENT INFRINGEMENT UNDER 35 U.S.C. § 271 OF THE
'286 PATENT BY CISCO)

108. Brazos re-alleges and incorporates by reference all of the foregoing paragraphs.

109. On information and belief, Cisco has directly infringed and continues to directly infringe either literally or under the doctrine of equivalents, one or more claims, including at least claim 7, of the '286 Patent in violation of 35 U.S.C. § 271, et seq., by deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling the Ultra-M Platform, which

includes Virtual Machines that run on UCS C240 Series servers⁴³, on which for example, the Cisco Prime Access Registrar (CPAR)⁴⁴ can be installed (collectively, the “Accused CPAR Products”).

110. Claim 7 of the ’286 Patent provides:

[Preamble] A system comprising:

[7A] a Radius gateway for translating a Radius message into a Network Access Server Request (NASReq) message and for transmitting the NASReq message;

[7B] a Diameter Proxy Agent within a Policy and Charging Rules Function (PCRF) server for receiving a NASReq message from the Radius gateway, for selecting one of at least one PCRF cluster within the PCRF server, and for forwarding the NASReq message to the selected PCRF cluster; and

[7C] at least one PCRF blade within the PCRF server, each PCRF blade belonging to one of the at least one PCRF cluster and configured to handle communication sessions for wireless devices, each PCRF blade for receiving a NASReq message from the Diameter Proxy Agent and for creating or updating a NASReq session object related to a communication session identified by the NASReq message received by the PCRF blade in response to receiving the NASReq message.


111. On information and belief, the Accused CPAR Products satisfy each and every limitation of at least claim 7 of the ’286 Patent as stated below.

112. On information and belief, and to the extent possible that the preamble of claim 7 is determined to be limiting, the Accused CPAR Products constitute a system. For example, Ultra-M is a pre-packaged and validated virtualized mobile packet core solution that is designed in order

⁴³ See https://www.cisco.com/c/en/us/td/docs/wireless/asr_5000/21-6_N6-0/Ultra-M-Solutions/N6-0-Ultra-M-Solution-Guide/N5-8-Ultra-M-Solution-Guide_chapter_010.pdf at 3.

⁴⁴ See <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/prime-home/213596-cpar-aaa-configuration.html>.

to simplify the deployment of VNFs.⁴⁵ Further, the Cisco Prime Access Registrar can be directly installed on a UCS-C Virtual Machine in the Ultra-M Platform shown below:⁴⁶



113. Thus, to the extent the preamble of claim 7 is limiting, the Accused CPAR Products satisfy the preamble of claim 7.

114. On information and belief, the Accused CPAR Products meet claim element [7A] of claim 7 of the '286 Patent, "a Radius gateway for translating a Radius message into a Network Access Server Request (NASReq) message and for transmitting the NASReq message." For example, the Prime Access Registrar is a 3GPP-compliant, 64-bit carrier-class RADIUS (Remote Authentication Dial-In User Service)/Diameter server that enables multiple dial-in Network

⁴⁵ See <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/prime-home/213596-cpar-aaa-configuration.html>.

⁴⁶ See <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/prime-home/213596-cpar-aaa-configuration.html>.

Access Server (NAS) devices to share a common authentication, authorization, and accounting database:⁴⁷

Prime Access Registrar is a 3GPP-compliant, 64-bit carrier-class RADIUS (Remote Authentication Dial-In User Service)/Diameter server that enables multiple dial-in Network Access Server (NAS) devices to share a common authentication, authorization, and accounting database.

115. Further, the Prime Access Registrar supports translation of an incoming RADIUS request to a Diameter request and vice versa:⁴⁸

Translation Framework for Diameter

Prime Access Registrar supports translation of an incoming RADIUS request to a Diameter request and vice versa.

The following services are created to set up the translation framework:

- Radius-Diameter—For translation of incoming RADIUS request to Diameter equivalent and then the Diameter response to RADIUS equivalent.

116. Therefore, the Accused CPAR Products meet element [7A] of claim 7.

117. On information and belief, the Accused CPAR Products meet claim element [7B] of claim 7 of the '286 Patent, “a Diameter Proxy Agent within a Policy and Charging Rules Function (PCRF) server for receiving a NASReq message from the Radius gateway, for selecting one of at least one PCRF cluster within the PCRF server, and for forwarding the NASReq message to the selected PCRF cluster.” For example, a Diameter Peer can be added in the Prime Access Registrar by configuring an entry in the PCRF client: ⁴⁹

For adding a Diameter peer in Prime Access Registrar, configure a new entry in the clients (including Policy and Charging Rules Functions (PCRF), Home Subscriber Servers (HSS), Mobility Management Entities (MME), Online Charging Systems (OCS), and others) and remote server object.

⁴⁷ See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 1-1.

⁴⁸ See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 2-27.

⁴⁹ See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 4-6.

118. Further, the Prime Access Registrar allows for creating two or more groups of Diameter remote servers in a Diameter proxy service configuration. Each of these groups will have a unique set of remote servers, i.e. no two groups will share the same remote server.⁵⁰

Group-Based Load Balancing in Diameter Proxy Server Configuration

Prime Access Registrar allows you to create two or more groups of Diameter remote servers in a Diameter proxy service configuration. Each of these groups will have a unique set of remote servers, i.e. no two groups will share the same remote server.

119. Further, the Accused CPAR Products provide group-based load-balancing.⁵¹

Group-Based Load Balancing in Diameter Proxy Server Configuration

Prime Access Registrar allows you to create two or more groups of Diameter remote servers in a Diameter proxy service configuration. Each of these groups will have a unique set of remote servers, i.e. no two groups will share the same remote server.

The traffic between each of these groups is load-balanced in failover mode; while traffic between remote servers within the same group is load-balanced based on round-robin or failover mode depending on the existing Diameter configuration. The priority of each of the groups is set with the help of metrics.

The workflow for group-based load balancing is as given below:

1. Traffic from Prime Access Registrar to remote server, via Diameter proxy service, is directed through the first group, till Prime Access Registrar has active communication channel with at least one remote server belonging to the first group.
2. When Prime Access Registrar loses connectivity with all the remote servers in the first group, it directs the rest of the Diameter traffic towards remote servers belonging to the second group.
3. Within a group, the load-balancing logic is chosen based on the configuration:
 - a. If the load-balancing logic is configured to be round-robin, the traffic is distributed across all the active remote servers.
 - b. If the load-balancing logic is configured to be failover, the traffic is directed towards first priority remote server. When Prime Access Registrar loses connectivity with the first priority remote server, it directs the subsequent traffic towards the second priority remote server. The priority of the Diameter remote servers, in case of failover logic, is set with the help of metrics.

120. Therefore, the Accused CPAR Products meet element [7B] of claim 7.

121. On information and belief, the Accused CPAR Products meet claim element [7C] of claim 7 of the '286 Patent, "at least one PCRF blade within the PCRF server, each PCRF blade belonging to one of the at least one PCRF cluster and configured to handle communication sessions

⁵⁰ See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 4-17.

⁵¹ See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 4-17.

for wireless devices, each PCRF blade for receiving a NASReq message from the Diameter Proxy Agent and for creating or updating a NASReq session object related to a communication session identified by the NASReq message received by the PCRF blade in response to receiving the NASReq message.” For example, as discussed in Limitation [7B], the Accused CPAR Products provide group-based load-balancing. Further, when a Diameter client issues an authentication request, Prime Access Registrar sends the packet with a Session-Id AVP, which can be used to correlate a Diameter message with a user-session.⁵²

Managing Diameter Sessions

Diameter provides two kinds of services namely authentication/authorization and accounting only (optional). Diameter sessions can be created when an authentication/authorization request comes from an access point or when an accounting start comes from an access point. When a Diameter client issues an authentication request, Prime Access Registrar sends the packet with a Session-Id AVP, which can be used to correlate a Diameter message with a user-session. When a Session Termination Request (STR) message is received from the Diameter client, Prime Access Registrar releases the sessions. Also Re-authentication requests must be mapped to the corresponding user session. In case of accounting packets, the session is created when the accounting start is received from the Diameter client. The session is deleted when the accounting stop message is received.

122. Therefore, the Accused CPAR Products meet element [7C] of claim 7.

123. Accordingly, the Accused CPAR Products satisfy each and every limitation of claim 7 of the '286 Patent.

124. Cisco undertook and continues its infringing actions despite an objectively high likelihood that such activities infringe the '286 Patent, which has been duly issued by the PTO and presumed valid. On information and belief, Cisco could not reasonably, subjectively believe that its actions do not constitute infringement of the '286 Patent. Despite that knowledge and subjective belief, and the objectively high likelihood that its actions constitute infringement, Cisco has

⁵² See https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/access_registrar/9-3/user/guide/user_guide.pdf at 4-17.

continued its infringing activities. As such, Cisco has willfully infringed and/or will continue to willfully infringe the '286 Patent at least as of the date of this Complaint.

125. As a result of Cisco's infringement of the '286 Patent, Brazos has suffered and continues to suffer substantial injury and is entitled to recover all damages caused by Cisco's infringement to the fullest extent permitted by the Patent Act, together with prejudgment interest and costs for Cisco's wrongful conduct.

126. Brazos has no adequate remedy at law to prevent future infringement of the '286 Patent. Brazos suffers and continues to suffer irreparable harm as a result of Cisco's patent infringement and is, therefore, entitled to injunctive relief to enjoin Cisco's wrongful conduct.

THIRD CAUSE OF ACTION
(PATENT INFRINGEMENT UNDER 35 U.S.C. § 271 OF THE
'721 PATENT BY CISCO)

127. Brazos re-alleges and incorporates by reference all of the foregoing paragraphs.

128. On information and belief, Cisco has directly infringed and continues to directly infringe either literally or under the doctrine of equivalents, one or more claims, including at least claim 19 of the '721 Patent in violation of 35 U.S.C. § 271, et seq., by deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling at least the Cisco NCS 1010 Optical Line System and other similar products.

129. Claim 19 of the '721 Patent provides:

[Preamble] An apparatus comprising:

[19A] a processor configured to receive measured channel powers and to determine deviations of respective measured channel powers from respective target channel powers,

[19A1] wherein the processor is further configured to project the deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed, and

[19A2] wherein the processor is further configured to determine power setting values for the pump lasers based on the projected deviations.

130. On information and belief, the Cisco NCS 1010 Optical Line System satisfies each and every limitation of at least claim 19 of the '721 Patent as stated below.

131. On information and belief, and to the extent possible that the preamble of claim 19 is determined to be limiting, the Cisco NCS 1010 Optical Line System is “an apparatus.”

132. For example, the Cisco NCS 1010 Optical Line System is described by Cisco documentation to have the following hardware features:⁵³

⁵³ See <https://www.cisco.com/c/en/us/products/collateral/optical-networking/network-convergence-system-1000-series/network-conver-system-1010-ds.pdf> at 5.




Figure 1.
NCS 1010 solution components

NCS 1010 hardware overview

The NCS 1010 is a 3RU 300-mm rack-compliant shelf with front-to-back straight-through cooling and is completely front-access. Rail kits and brackets are available to mount onto 19-, 21-, or 23-inch racks. Fixed brackets on the NCS 1010 have thumbscrews to allow for fast deployment or replacement of the shelf. Each NCS 1010 has two redundant and field-replaceable 1KW AC or DC power supply units and two field-replaceable fan trays that operate at 5+1 redundancy. It also provides a field-replaceable controller card (Figure 2). The NCS 1010 has one SSD on the shelf controller and a second SSD on the shelf to maintain redundant copies of XR images and system configuration. Each NCS 1010 chassis provides one line card slot, and the line card is field replaceable.

A range of management, USB 3.0, and timing input/output ports are provided on the top right the NCS 1010 faceplate that can be used for DCN management, daisy chaining of shelves, console access, passive device management, user data channel, and timing distribution. NCS 1010 is hardware designed to support class C-compliant 1588v2 PTP timing support.

133. Thus, to the extent the preamble of claim 19 is limiting, the Cisco NCS 1010 Optical Line System satisfies the preamble of claim 19.

134. On information and belief, the Cisco NCS 1010 Optical Line System meets claim element [19A] of claim 19 of the '721 Patent, "a processor configured to receive measured channel powers and to determine deviations of respective measured channel powers from respective target channel powers."

135. For example, the Cisco documentation explains the following about the Cisco NCS 1010 Optical Line System's processor, which is configured to run the preinstalled IOS XR software:⁵⁴

NCS 1010 is shipped with the Cisco IOS XR software preinstalled. Verify that the latest version of the software is installed. If a newer version is available, perform a [Upgrade Software, on page 51](#). This software upgrade installs the newer version of the software and provide the latest feature set on NCS 1010.

To verify the version of Cisco IOS XR Software running on NCS 1010, perform the following procedure.

show version

Displays the software version and details such as system uptime.

Example:

```
RP/0/RP0/CPU0:ios#show version
Thu Jul 28 09:49:34.374 UTC
Cisco IOS XR Software, Version 7.7.1 LNT
Copyright (c) 2013-2022 by Cisco Systems, Inc.

Build Information:
Built By      : ingunawa
Built On     : Mon Jul 25 06:07:25 UTC 2022
Build Host   : iox-lnx-109
Workspace    : /auto/srcarchive12/prod/7.7.1/ncs1010/ws
Version      : 7.7.1
Label        : 7.7.1

cisco NCS1010 (C3758 @ 2.20GHz)
cisco NCS1010-SA (C3758 @ 2.20GHz) processor with 32GB of memory
```

136. Furthermore, the Cisco NCS 1010 Optical Line System's processor is configured to receive measured channel powers. For example, Cisco documentation describes an Optical Line Terminal with built-in Raman amplification:⁵⁵

⁵⁴ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-system-setup-guide.pdf> at 39.

⁵⁵ See <https://www.cisco.com/c/en/us/products/collateral/optical-networking/network-convergence-system-1000-series/network-conver-system-1010-ds.pdf> at 6-7.

NCS 1010 Optical Line Terminal

The Optical Line Terminal (OLT) provides the following key functions integrated into a line card: 25dBm variable gain preamplifier with switchable gain range, 23dBm variable gain booster amplifier, 2x33 port twin flex-grid Wavelength Selective Switch (WSS), embedded EDFA for ASE loading, embedded Fixed Gain EDFA for -10dBm digital coherent optics sources, embedded OTDR, connection verification Transmit Optical Sub Assembly, embedded OCM, FE/GE Optical Supervisory Channel with a 2x2 switch, embedded C+L-Band combiner, and optional built-in Raman amplification.

Key benefits:

- 33-port twin flex-grid WSS provides scale of a multi-degree ROADM node both for colorless add-drop ports or express ports.
- Built-in ASE loading allows for easier turn-up of the network and consistency in performance from day 1 through the life of the network.
- 25dBm true variable-gain preamplifier with switchable gain range, 23dBm true variable-gain booster amplifier.
- Dedicated fixed-gain EDFAs at the WSS add for -10dBm low-power digital coherent optics.
- Connection verification on fiber patches from ROADM to passive breakout or add-drop modules and on ROADM-to-ROADM express connections. This allows for the user to verify if patching is set up correctly or allows patching to be discovered and reported to the user.
- The OLT design allows hitless C to C+L-Band upgrade.
- Bi-directional OSC support for in-fiber communication between devices, user data channel support, and class C-compliant IEEE 1588v2 PTP timing transport. OSC is available for both C-Band and L-Band separately.

© 2024 Cisco and/or its affiliates. All rights reserved.

Page 6 of 29

- Built in bi-directional OTDR that can run in-line with DWDM channels and take measurements in real time of loss and back reflection of both pairs of fiber in a link from both end points.
- Optical channel monitor with a 1x37 switch provides a spectrum-analyzer-like view of power distribution across the available spectrum at the WSS add ports and input and output of the PRE and BST amplifiers. This helps with troubleshooting the system and enables automatic turn-up of the ROADM through WSS attenuation and amplifier gain control loops.

OLT line card comes in three flavors: C-Band OLT, C-Band OLT with Raman, and L-Band OLT.




Figure 3.
NCS 1010 OLT-C functional layout

137. The Cisco NCS 1010 Optical Line System's processor is configured to receive measurements when Raman Tuning is performed:⁵⁶

Overview of Raman Tuning

Raman Tuning Algorithm calculates and sets the different pump power values across five Raman pumps to obtain the target Raman Gain on a span. Raman tuning runs in both directions of the span independently at the node level. Raman tuning requires communication between peer nodes. Hence, OSC communication between the two nodes is a prerequisite for Raman Tuning.

Raman tuning algorithm uses the following parameters to calculate the pump powers necessary to achieve the target Raman gain.

- Fiber type
- Fiber length
- Loss on the fiber at each pump wavelength
- Loss on the fiber at the signal wavelength

If you configure a span length value, Raman tuning uses this value.

Raman tuning is disabled by default. Raman tuning is enabled if automatic link bring up is enabled. You can manually trigger Raman tuning if necessary. The NCS 1010 initiates the tuning process under the following circumstances:

- During the initial link bring up
- After a fiber cut
- After a power cycle event
- After a line card cold reload event
- After a DFB shut or unshut event
- After an OTS controller shut or unshut event on near end or far end node
- After modification of span length configuration

⁵⁶ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-optical-apps-config-guide.pdf> at 9.

138. The following table lists and describes the different Raman Tuning statuses:⁵⁷

Raman Tuning Status	Description
WORKING – MEASUREMENT	The algorithm is measuring the span loss on the link.
WORKING – CALCULATION	The algorithm is calculating the gain target and required pump powers.
WORKING – OPTIMIZATION	The algorithm is optimizing the pump powers.
TUNED	Raman tuning is complete.
BLOCKED	The system is unable to perform Raman tuning. This status can occur because the link is down or the system detected high Raman Back Reflection.
DISABLED	Raman tuning is disabled.

Raman tuning works in the following three modes:

- Auto mode: Raman tuning defines the target gain and sets the pump powers and DFB VOA attenuation to achieve the target gain overwriting user configuration.
- Gain mode: User defines the gain target and Raman tuning sets the pump powers and DFB VOA attenuation to achieve the target gain.
- Manual mode: User disables Raman tuning and manually configures the Raman pumps and DFB VOA attenuation.

139. The Cisco NCS 1010 Optical Line System’s processor is also configured to determine deviations of respective measured channel powers from respective target channel powers. As shown in the above table, at least for Auto mode and/or Gain mode, the Raman tuning algorithm determines deviations of respective measured channel powers from respective target channel powers. Such an algorithm operates during the “WORKING – CALCULATION” and/or “WORKING – OPTIMIZATION” states listed in the table above.

140. Furthermore, Cisco documentation shows the following about the “show olc raman-tuning” command line command:⁵⁸

⁵⁷ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-optical-apps-config-guide.pdf> at 10.

⁵⁸ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-optical-apps-config-guide.pdf> at 11; see also *id.* at 12-14 (including “Configure Raman Tuning” section).

View Raman Tuning Status

You can view the Raman tuning status using **show olc raman-tuning** command. The following sample is an output of the **show olc raman-tuning** command.

```
RP/0/RP0/CPU0:ios#sh olc raman-tuning
Tue Mar 21 06:11:36.944 UTC
```

```
Controller : Ots0/0/0/0
Raman-Tuning Status : TUNED
Tuning Complete Timestamp : 2023-03-20 07:54:00
Estimated Max Possible Gain : 19.8 dB
Raman Gain Target : 16.0 dB
Gain Achieved on Tuning Complete : 15.7 dB
```

You can view the Raman tuning status for individual controllers using **show olc raman-tuning controller ots r/s/i/p** command. The following sample is an output of the **show olc raman-tuning controller ots r/s/i/p** command.

```
RP/0/RP0/CPU0:ios#sh olc raman-tuning controller ots 0/0/0/0
Tue Mar 21 06:13:26.535 UTC
```

```
Controller : Ots0/0/0/0
Raman-Tuning Status : TUNED
Tuning Complete Timestamp : 2023-03-20 07:54:00
Estimated Max Possible Gain : 19.8 dB
Raman Gain Target : 16.0 dB
Gain Achieved on Tuning Complete : 15.7 dB
```

To view the individual Raman pump information and other parameters, use the **show controllers ots r/s/i/p raman-info** command.

141. The command output shows a “Raman Gain Target” and a “Gain Achieved,” both of which are directly associated with a power level. On information and belief, these measurements are used to determine deviations of respective measured channel powers from respective target channel powers.

142. Therefore, the Cisco NCS 1010 Optical Line System meets element [19A] of claim 19.

143. On information and belief, the Cisco NCS 1010 Optical Line System meets claim element [19A1] of claim 19 of the '721 Patent, “wherein the processor is further configured to project the deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed.”

144. For example, the Cisco NCS 1010 Optical Line System's processor performs Raman Tuning, as shown below:⁵⁹

Raman Tuning Status	Description
WORKING – MEASUREMENT	The algorithm is measuring the span loss on the link.
WORKING – CALCULATION	The algorithm is calculating the gain target and required pump powers.
WORKING – OPTIMIZATION	The algorithm is optimizing the pump powers.
TUNED	Raman tuning is complete.
BLOCKED	The system is unable to perform Raman tuning. This status can occur because the link is down or the system detected high Raman Back Reflection.
DISABLED	Raman tuning is disabled.

Raman tuning works in the following three modes:

- Auto mode: Raman tuning defines the target gain and sets the pump powers and DFB VOA attenuation to achieve the target gain overwriting user configuration.
- Gain mode: User defines the gain target and Raman tuning sets the pump powers and DFB VOA attenuation to achieve the target gain.
- Manual mode: User disables Raman tuning and manually configures the Raman pumps and DFB VOA attenuation.

145. On information and belief, at least in Auto mode and/or Gain mode, the Raman tuning algorithm projects the deviations into a space that defines Raman gain profiles achievable with a set of channels and pump lasers whereby projected deviations are formed. Such an algorithm operates during the “WORKING – CALCULATION” and/or “WORKING – OPTIMIZATION” statuses listed in the table above.

146. Therefore, the Cisco NCS 1010 Optical Line System meets element [19A1] of claim 19.

⁵⁹ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-optical-apps-config-guide.pdf> at 10.

147. On information and belief, the Cisco NCS 1010 Optical Line System meets claim element [19A2] of claim 19 of the '721 Patent, “wherein the processor is further configured to determine power setting values for the pump lasers based on the projected deviations.”

148. For example, the Cisco NCS 1010 Optical Line System’s processor is configured to perform Raman Tuning:⁶⁰

Raman Tuning Status	Description
WORKING – MEASUREMENT	The algorithm is measuring the span loss on the link.
WORKING – CALCULATION	The algorithm is calculating the gain target and required pump powers.
WORKING – OPTIMIZATION	The algorithm is optimizing the pump powers.
TUNED	Raman tuning is complete.
BLOCKED	The system is unable to perform Raman tuning. This status can occur because the link is down or the system detected high Raman Back Reflection.
DISABLED	Raman tuning is disabled.

Raman tuning works in the following three modes:

- Auto mode: Raman tuning defines the target gain and sets the pump powers and DFB VOA attenuation to achieve the target gain overwriting user configuration.
- Gain mode: User defines the gain target and Raman tuning sets the pump powers and DFB VOA attenuation to achieve the target gain.
- Manual mode: User disables Raman tuning and manually configures the Raman pumps and DFB VOA attenuation.

149. At least in Auto mode and/or Gain mode, the Raman tuning algorithm determines power setting values for the pump lasers. On information and belief, these power setting values are based on the projected deviations. Determination of power setting values occurs during the “WORKING – CALCULATION” and/or “WORKING – OPTIMIZATION” states listed in the

⁶⁰ See <https://www.cisco.com/c/en/us/td/docs/optical/ncs1010/77x/configuration/guide/b-ncs1010-optical-apps-config-guide.pdf> at 10.

table above. The determined power setting values are deployed to the Raman amplifiers, and can be reported to an operator via a command line command as shown below:⁶¹

View Raman Tuning Status

You can view the Raman tuning status using **show olc raman-tuning** command. The following sample is an output of the **show olc raman-tuning** command.

```
RP/0/RP0/CPU0:ios#sh olc raman-tuning
Tue Mar 21 06:11:36.944 UTC

Controller : Ots0/0/0/0
Raman-Tuning Status : TUNED
Tuning Complete Timestamp : 2023-03-20 07:54:00
Estimated Max Possible Gain : 19.8 dB
Raman Gain Target : 16.0 dB
Gain Achieved on Tuning Complete : 15.7 dB
```

You can view the Raman tuning status for individual controllers using **show olc raman-tuning controller ots r/s/i/p** command. The following sample is an output of the **show olc raman-tuning controller ots r/s/i/p** command.

```
RP/0/RP0/CPU0:ios#sh olc raman-tuning controller ots 0/0/0/0
Tue Mar 21 06:13:26.535 UTC

Controller : Ots0/0/0/0
Raman-Tuning Status : TUNED
Tuning Complete Timestamp : 2023-03-20 07:54:00
Estimated Max Possible Gain : 19.8 dB
Raman Gain Target : 16.0 dB
Gain Achieved on Tuning Complete : 15.7 dB
```

To view the individual Raman pump information and other parameters, use the **show controllers ots r/s/i/p raman-info** command.

150. Therefore, the Cisco NCS 1010 Optical Line System meets element [19A2] of claim 19.

151. Accordingly, Cisco's NCS 1010 Optical Line System satisfies each and every limitation of claim 19 of the '721 Patent.

152. Cisco undertook and continues its infringing actions despite an objectively high likelihood that such activities infringe the '721 Patent, which has been duly issued by the PTO and presumed valid. On information and belief, Cisco could not reasonably, subjectively believe that its actions do not constitute infringement of the '721 Patent. Despite that knowledge and subjective

⁶¹ *Id.* at 11; *see also id.* at 12-14 (including "Configure Raman Tuning" section).

belief, and the objectively high likelihood that its actions constitute infringement, Cisco has continued its infringing activities. As such, Cisco has willfully infringed and/or will continue to willfully infringe the '721 Patent at least as of the date of this Complaint.

153. As a result of Cisco's infringement of the '721 Patent, Brazos has suffered and continues to suffer substantial injury and is entitled to recover all damages caused by Cisco's infringement to the fullest extent permitted by the Patent Act, together with prejudgment interest and costs for Cisco's wrongful conduct.

154. Brazos has no adequate remedy at law to prevent future infringement of the '721 Patent. Brazos suffers and continues to suffer irreparable harm as a result of Cisco's patent infringement and is, therefore, entitled to injunctive relief to enjoin Cisco's wrongful conduct.

FOURTH CAUSE OF ACTION
(PATENT INFRINGEMENT UNDER 35 U.S.C. § 271 OF THE
'691 PATENT BY CISCO)

155. Brazos re-alleges and incorporates by reference all of the foregoing paragraphs.

156. On information and belief, Cisco has directly infringed and continues to directly infringe either literally or under the doctrine of equivalents, one or more claims, including at least claim 6 of the '691 Patent in violation of 35 U.S.C. § 271, et seq., by deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling at least the Cisco NCS 4200 Series Network Convergence Systems and similar products.

157. Claim 6 of the '691 Patent provides:

[Preamble] A non-transitory machine readable storage medium encoded with instructions for execution by a network processor of a Multiprotocol Label Switching (MPLS) label switch for providing a Backup Label Switched Path (LSP) to a Bypass LSP already established for a Protected Primary LSP, the medium comprising:

[6A] instructions for protecting the Primary LSP against dual failures, comprising:

[6B] instructions for establishing the Bypass LSP for the Protected Primary LSP having a Point of Local Repair node and a Merge Point node;

[6C] instructions for obtaining the nodes traversed by an end-to-end path of said Bypass LSP from said Point of Local Repair Node to said Merge Point node;

[6D] instructions for generating a request to a path calculator using the nodes traversed by said end-to-end path of said Bypass LSP for a disjoint path connecting said Point of Local Repair Node to said Merge Point node;

[6E] instructions for receiving a response from said path calculator; and


[6F] in response to determining that a fully disjoint path connecting said Point of Local Repair Node to said Merge Point node is available, instructions for signaling, to at least one other MPLS label switch router, said fully disjoint path as the Backup LSP to said Bypass LSP.

158. On information and belief, Cisco's Network Convergence System 4200 Series satisfy each and every limitation of at least claim 6 of the '691 Patent as stated below.

159. On information and belief, and to the extent possible that the preamble of claim 6 is determined to be limiting, Cisco's Network Convergence System 4200 Series includes "a non-transitory machine readable storage medium encoded with instructions for execution by a network processor of a Multiprotocol Label Switching (MPLS) label switch for providing a Backup Label Switched Path (LSP) to a Bypass LSP already established for a Protected Primary LSP."

160. For example, the Cisco NCS 4200 Series delivers any-to-any connectivity using a packet-based network (MPLS/Flex LSP/SR). Further, the Cisco IOS XE Software available on the NCS 4200 Series portfolio shown below includes all features required MPLS:⁶²

⁶² See <https://www.cisco.com/c/en/us/products/collateral/optical-networking/network-convergence-system-4200-series/datasheet-c78-738104.pdf> at 3.



161. Further, Flex LSP provides a Restore Path Option that signals a restore LSP after the double failure of both, primary and protect LSPs.⁶³

Restore Path Option

The restore path option signals a restore LSP after the double failure of both, primary and protect LSPs. The restore LSP is signaled only after both, primary and protect LSPs fail or are administratively down and it is destroyed when the primary LSP comes back up. If the sticky option is configured, and both, primary and protect LSPs fail, restore LSP is destroyed when either the primary or protect LSP comes up. Also, restore LSP can be SRLG capable if it is configured.

162. Thus, to the extent the preamble of claim 6 is limiting, Cisco's Network Convergence System 4200 Series satisfies the preamble of claim 6.

163. On information and belief, Cisco's Network Convergence System 4200 Series meets claim element [6A] of claim 6 of the '691 Patent, "instructions for protecting the Primary

⁶³ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mpls/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 145.

LSP against dual failures, comprising” [6B]-[6F]. For example, Flex LSP provides a Restore Path Option that signals a restore LSP after the double failure of both, primary and protect LSPs:⁶⁴

Restore Path Option

The restore path option signals a restore LSP after the double failure of both, primary and protect LSPs. The restore LSP is signaled only after both, primary and protect LSPs fail or are administratively down and it is destroyed when the primary LSP comes back up. If the sticky option is configured, and both, primary and protect LSPs fail, restore LSP is destroyed when either the primary or protect LSP comes up. Also, restore LSP can be SRLG capable if it is configured.

164. Therefore, Cisco’s Network Convergence System 4200 Series meets element [6A] of claim 6.

165. On information and belief, Cisco’s Network Convergence System 4200 Series meets claim element [6B] of claim 6 of the ’691 Patent, “instructions for establishing the Bypass LSP for the Protected Primary LSP having a Point of Local Repair node and a Merge Point node.” For example, there are two types of LSPs: protect LSPs and working LSPs. The protect LSP acts as a backup for a working LSP:⁶⁵

MPLS-TP Path Protection

MPLS-TP label switched paths (LSPs) support 1-to-1 path protection. There are two types of LSPs: protect LSPs and working LSPs. You can configure the both types of LSPs when configuring the MPLS-TP tunnel. The working LSP is the primary LSP used to route traffic. The protect LSP acts as a backup for a working LSP. If the working LSP fails, traffic is switched to the protect LSP until the working LSP is restored, at which time forwarding reverts back to the working LSP.

166. Further, MPLS Link Protection provides backup tunnels that bypass only a single link. Backup tunnels protect LSPs if there is any failure/fault in a primary LSP:⁶⁶

⁶⁴ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mpls/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 145.


⁶⁵ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mpls/16-12-1/b-mp-basic-16-12-1-ncs4200.pdf> at 145.

⁶⁶ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_path_protect/configuration/12-4mt/mp-te-path-protect-12-4t-book.pdf at 13.

Link Protection

Backup tunnels that bypass only a single link of the LSP's path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. The figure below illustrates an NHOP backup tunnel.

Figure 1 NHOP Backup Tunnel



167. Further, as shown above, the backup tunnels should intersect with a primary tunnel at a minimum of two nodes: point of local repair (PLR) and merge point (MP). The PLR should be the headend LSR of the backup tunnel, and the MP should be the tail end LSR of the backup tunnel:⁶⁷

The backup tunnel must meet the following requirements:

- It should not pass through the element it protects.
- It should intersect with a primary tunnel at a minimum of two nodes: point of local repair (PLR) and merge point (MP). The PLR should be the headend LSR of the backup tunnel, and the MP should be the tailend LSR of the backup tunnel. The PLR is where FRR is triggered when a link, node, or SRLG failure occurs.

168. Therefore, Cisco's Network Convergence System 4200 Series meets element [6B] of claim 6.


169. On information and belief, Cisco's Network Convergence System 4200 Series meets claim element [6C] of claim 6 of the '691 Patent, "instructions for obtaining the nodes traversed by an end-to-end path of said Bypass LSP from said Point of Local Repair Node to said

⁶⁷ See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_path_protect/configuration/12-4mt/mp-te-path-protect-12-4t-book.pdf at 187.

Merge Point node.” For example, path computation for the participating LSP is performed using Constrained Shortest Path First (CSPF):⁶⁸

Path computation is performed by the head-end nodes of both the participating LSPs using Constrained Shortest Path First (CSPF). CSPF is the 'shortest path (measured in terms of cost) that satisfies all relevant LSP TE constraints or attributes, such as required bandwidth, priority and so on.

170. Further, as one example depicted below, after performing bidirectional CSPF, Router 1 sends reverse explicit route object (ERO) to Router 3:⁶⁹



171. Generally, the Explicit Route Object (ERO) limit LSP routing to a specified list of LSRs.⁷⁰

⁶⁸ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-6-1/b-mp-basic-16-6-1-ncs4200.pdf> at 118.

⁶⁹ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-6-1/b-mp-basic-16-6-1-ncs4200.pdf> at 119.

⁷⁰ See <https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r7-5/mps/configuration/guide/b-mps-cg-asr9000-75x/implementing-gmps-uni-74x.pdf> at 5.

172. Therefore, Cisco's Network Convergence System 4200 Series meets element [6C] of claim 6.

173. On information and belief, Cisco's Network Convergence System 4200 Series meets claim element [6D] of claim 6 of the '691 Patent, "instructions for generating a request to a path calculator using the nodes traversed by said end-to-end path of said Bypass LSP for a disjoint path connecting said Point of Local Repair Node to said Merge Point node." For example, Cisco IOS XE supports Shared Risk Link Groups (SRLGs) Protection. SRLGs indicate situations where links in a network share a common fiber. Further, if one shared link fails, the other links in the group may also fail. Accordingly, SRLG Protection avoids using links in the same SRLG as interfaces the backup tunnel is protecting.⁷¹

- Shared Risk Link Groups (SRLGs) Protection – SRLGs indicate situations where links in a network share a common fiber (or a common physical attribute). If one link fails, the other links in the group may also fail. Links in this group have a shared risk.

The MPLS-TE SRLG Protection feature enhances backup tunnel path selection so that a backup tunnel avoids using links that are in the same SRLG as interfaces the backup tunnel is protecting.

174. Further, enabling the MPLS Traffic Engineering (TE)—IP Explicit Address Exclusion feature with the "ip explicit-path" adds subcommands for excluding addresses:⁷²

The MPLS Traffic Engineering (TE)—IP Explicit Address Exclusion feature provides a means to exclude a link or node from the path for Multiprotocol Label Switching (MPLS) TE label switched path (LSP).

The feature is enabled through the **ip explicit-path** command that allows you to create an IP explicit path and enter a configuration submode for specifying the path. The feature adds to the submode commands the **exclude-address** command for specifying addresses to exclude from the path.

⁷¹ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-6-1/b-mp-basic-16-6-1-ncs4200.pdf> at 131.

⁷² See https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_te_path_protect/configuration/12-4mt/mp-te-path-protect-12-4t-book.pdf at 119.

175. Further, the “exclude-srlg” command specifies an address to get SRLGs from for exclusion:⁷³

Flex LSP SRLG and Exclude Option for Explicit Path

Use the following commands to configure SRLG on an interface:

```
Router(config)# interface Ethernet0/1
Router(config-if)# srlg gid <1-4294967295>
```

SRLG values configured on MPLS TE enabled interfaces are flooded through IGP (IS-IS or OSPF), and are used by MPLS TE in the following scenarios:

- Restrict protection path to avoid SRLGs of links in the working path.

```
Router(config)# interface Tunnel100
Router(config-if)# tunnel mpls traffic-eng path-option protect 1 diverse
srlg [node] lockdown
```

- Exclude SRLG in IP Explicit Path.

```
Router(config)# ip explicit-path name EXAMPLE
Router(cfg-ip-expl-path)# exclude-srlg A.B.C.D
```

The exclude-srlg command specifies an address to get SRLGs from for exclusion.

176. Further, SRLG supports preferring a disjoint repair path when there are two repair paths for a prefix:⁷⁴

- srlg-disjoint—Prefers the SRLG disjoint repair path

The SRLG ID can be configured for each interface. When there are two repair paths for a prefix, the configured SRLG ID for the repair path is compared with that of the primary path SRLG ID. If the SRLG IDs for the secondary path is different than that of the primary, that path is chosen as the repair path.


177. Therefore, Cisco’s Network Convergence System 4200 Series meets element [6D] of claim 6.

178. On information and belief, Cisco’s Network Convergence System 4200 Series meets claim element [6E] of claim 6 of the ’691 Patent, “instructions for receiving a response from

⁷³ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mps/16-6-1/b-mp-basic-16-6-1-ncs4200.pdf> at 132.

⁷⁴ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/segment-routing/17-1-1/b-segment-routing-17-1-ncs4200.pdf> at 30.

said path calculator.” For example, CSPF sends a reverse explicit route option to the MPLS Router as a response:⁷⁵



In the above topology:

- Paths at the top of the figure (in green) indicate working co-routed LSP pairs.
- Paths at the bottom of the figure (in red) indicate protecting co-routed LSP pairs.
- Router 1 sets up working LSP to Router 3 (in red) after performing bidirectional CSPF and sends reverse explicit route object (ERO) to Router 3. Node Router 3 uses the received reverse ERO to set up reverse red working LSP to Router 1.
- Router 3 sets up protecting LSP to Router 1 (in green) after performing bidirectional CSPF and sends reverse ERO to Router 1. Node Router 1 uses the received reverse ERO to set up reverse green protecting LSP to Router 3.

179. Therefore, Cisco’s Network Convergence System 4200 Series meets element [6E] of claim 6.

180. On information and belief, Cisco’s Network Convergence System 4200 Series meets claim element [6F] of claim 6 of the ’691 Patent, “in response to determining that a fully disjoint path connecting said Point of Local Repair Node to said Merge Point node is available, instructions for signaling, to at least one other MPLS label switch router, said fully disjoint path as the Backup LSP to said Bypass LSP.” For example, when there are two repair paths for a prefix, the configured SRLG ID for the repair path is compared with that of the primary path SRLG ID. If the SLRG IDs for the secondary path is different than that of the primary, that path is chosen as the repair path:⁷⁶

⁷⁵ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/mpls/16-6-1/b-mp-basic-16-6-1-ncs4200.pdf> at 119, Exhibit C.

⁷⁶ See <https://www.cisco.com/c/en/us/td/docs/routers/ncs4200/configuration/guide/segment-routing/17-1-1/b-segment-routing-17-1-ncs4200.pdf> at 30.

- srlg-disjoint—Prefers the SRLG disjoint repair path

The SRLG ID can be configured for each interface. When there are two repair paths for a prefix, the configured SRLG ID for the repair path is compared with that of the primary path SRLG ID. If the SRLG IDs for the secondary path is different than that of the primary, that path is chosen as the repair path.

181. Therefore, Cisco's Network Convergence System 4200 Series meets element [6F] of claim 6.

182. Accordingly, Cisco's Network Convergence System 4200 Series satisfies each and every limitation of claim 6 of the '691 Patent.

183. Cisco undertook and continues its infringing actions despite an objectively high likelihood that such activities infringe the '691 Patent, which has been duly issued by the PTO and presumed valid. On information and belief, Cisco could not reasonably, subjectively believe that its actions do not constitute infringement of the '691 Patent. Despite that knowledge and subjective belief, and the objectively high likelihood that its actions constitute infringement, Cisco has continued its infringing activities. As such, Cisco has willfully infringed and/or will continue to willfully infringe the '691 Patent at least as of the date of this Complaint.

184. As a result of Cisco's infringement of the '691 Patent, Brazos has suffered and continues to suffer substantial injury and is entitled to recover all damages caused by Cisco's infringement to the fullest extent permitted by the Patent Act, together with prejudgment interest and costs for Cisco's wrongful conduct.

185. Brazos has no adequate remedy at law to prevent future infringement of the '691 Patent. Brazos suffers and continues to suffer irreparable harm as a result of Cisco's patent infringement and is, therefore, entitled to injunctive relief to enjoin Cisco's wrongful conduct.

FIFTH CAUSE OF ACTION
(PATENT INFRINGEMENT UNDER 35 U.S.C. § 271 OF THE
'884 PATENT BY CISCO)

186. Brazos re-alleges and incorporates by reference all of the foregoing paragraphs.

187. On information and belief, Cisco has directly infringed and continues to directly infringe either literally or under the doctrine of equivalents, one or more claims, including at least claim 17 of the '884 Patent in violation of 35 U.S.C. § 271, et seq., by making, using, offering to sell, importing, and exporting Cisco's Catalyst® 9000 Switching Platforms ("Catalyst Switches").

188. Claim 17 of the '884 Patent provides:

[Preamble] An edge switch for adjusting bandwidth allocation in a communications network, the edge switch including a target port, the edge switch configured to:

[17A] monitor a data flow traversing the target port;

[17B] determine a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow;

[17C] determine a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation being a proportional allocation of a total bandwidth of the network switching element; and

[17D] adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.

189. On information and belief, Cisco's Catalyst Switches satisfy each and every limitation of at least claim 17 of the '884 Patent as stated below.

190. On information and belief, and to the extent possible that the preamble of claim 17 is determined to be limiting, Cisco's Catalyst Switch is "[an] edge switch for adjusting bandwidth

allocation in a communications network, the edge switch including a target port, the edge switch configured.” For example, a Cisco Catalyst 9400 are edge switches. Cisco Documentation states:⁷⁷

Cisco® Catalyst® 9400 Series switches are Cisco's lead modular enterprise access switching platform and as part of the Catalyst 9000 family, are built to transform your network to handle a hybrid world where the workplace is anywhere, endpoints could be anything, and applications are hosted all over the place. The Catalyst 9400 Series, including the Catalyst 9400 SUP-2/2XL supervisor and line cards, continues to shape the future with continued innovation that helps you reimagine connections, reinforce security and redefine the experience for your hybrid workforce big and small.

The Cisco Networking Cloud and Software-Defined Access (SD-Access) is the network fabric that powers business. Cisco Networking Cloud is an open and extensible, software-driven architecture that accelerates and simplifies your enterprise network operations. The programmable architecture frees your IT staff from time consuming, repetitive network configuration tasks so they can focus instead on innovation that positively transforms your business. SD-Access enables policy-based automation from edge to cloud with foundational capabilities. These include:

191. Catalyst Switches comprise a plurality of target ports.

<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9200-series-switches/nb-06-cat9200-ser-data-sheet-cte-en.pdf> (Cisco Catalyst 9200 Series Switches data sheet), pp. 6–7,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/nb-06-cat9300-ser-data-sheet-cte-en.pdf> (Cisco Catalyst 9300 Series Switches data sheet), pp. 7–9,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9400-series-switches/nb-06-cat9400-ser-data-sheet-cte-en.pdf> (Cisco Catalyst 9400 Series Switches data sheet), pp. 9–10,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9500-series-switches/nb-06-cat9500-ser-data-sheet-cte-en.pdf> (Cisco Catalyst 9500 Series Switches data sheet), pp. 12–13,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9600-series-switches/nb-06-cat9600-series-data-sheet-cte-en.pdf> (Cisco Catalyst 9600 Series Switches data sheet), pp. 8–11.

192. Cisco documentation states:⁷⁸

⁷⁷ <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9400-series-switches/nb-06-cat9400-ser-data-sheet-cte-en.pdf>, p. 3.

⁷⁸ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 3.

QoS helps a network provide guaranteed and predictable services to selected network traffic by adding the following techniques:

- Scheduling to support guaranteed bandwidth
- Reducing loss characteristics for specified traffic
- Avoiding and managing network congestion
- Shaping network traffic
- Setting traffic priorities across the network

Using the above techniques to implement QoS in your network has the following advantages:

- Control over resources such as bandwidth, rate-limiting, and so on. For example, you can limit bandwidth consumed over a link by FTP transfers or give priority to an important database access.

193. Cisco documentation describes the Catalyst 9200 series as follows:⁷⁹

The Cisco Catalyst 9200 Series Switches have a simple architecture. All of the front panel ports, including the network module ports, are connected to one UADP Mini ASIC for the 24- or 48-port models or two UADP Mini ASICs for the Multigigabit models. The QoS buffer is shared among all ports, as the UADP Mini has a single ASIC core. Every port supports individual queuing capabilities.

The Cisco Catalyst 9200 Series are stackable switches using StackWise®-160/80. Each switch comes with two stack interfaces that can connect the switch to two other switches in a stack. The stack interface is part of the ASIC, which pushes the packets onto the stack ring. A dedicated amount of buffer is allocated for the stack interface and is not user configurable.




Figure 3.
Cisco Catalyst 9200 Series architecture of dual ASIC switch

⁷⁹ *Id.*, p. 6.

194. Cisco documentation describes the Catalyst 9300 series as follows:⁸⁰

The Cisco Catalyst 9300 Series Switches have a simple architecture. All of the front panel ports, including the network module ports, are connected to the UADP 2.0 ASIC. Depending on the model, the switch can have one or more ASICs serving all ports. For switches with more than one ASIC, generally the ports are equally divided among the ASICs, and therefore all ports will have an equal amount of resources available from each of the ASICs. The QoS buffer is provided per ASIC core and is shared only among the ports connected to that ASIC core. Every port supports individual queuing capabilities.

The Cisco Catalyst 9300 Series are stackable switches. Each switch comes with two stack interfaces that can connect the switch to two other switches in a stack. The stack interface is part of the ASIC, which pushes the packets onto the stack ring. A dedicated amount of buffer is allocated for the stack interface and is not user configurable.




Figure 4.
Cisco Catalyst 9300 Series architecture of dual ASIC switch

⁸⁰ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 7; see also <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/nb-06-cat9300-ser-data-sheet-cte-en.pdf>.

195. Cisco documentation describes the Catalyst 9400 series as follows:⁸¹

The Cisco Catalyst 9400 Series Switches are based on a centralized architecture, which means all packets are processed on the supervisor, while line cards are considered transparent, containing only stub ASICs and PHYs. Therefore, all QoS resources reside on the supervisor, including the per-port buffer and other QoS resources. The simplicity of this centralized design allows easy feature upgrades just by upgrading the supervisor while keeping the existing line cards.




Figure 5.
Cisco Catalyst 9400 Series architecture

⁸¹ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 8; see also <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9400-series-switches/nb-06-cat9400-ser-data-sheet-cte-en.pdf>.

196. Cisco documentation describes the Catalyst 9500 series as follows:⁸²

Each model in the Cisco Catalyst 9500 Series offers different port speeds and port density, but from a QoS architecture point of view, the 9500 Series is similar to the 9300 Series. Depending on the model, the 9500 Series switches are based on either UADP 2.0 XL or UADP 3.0. Table 1 summarizes the models and the ASICs used in each.

Table 1. UADP versions in the Cisco Catalyst 9500 Series

Model	UADP 2.0 XL	UADP 3.0
C9500-32C		2 ASICs

⁸² See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 8–10; see also <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9500-series-switches/nb-06-cat9500-ser-data-sheet-cte-en.pdf>.

Model	UADP 2.0 XL	UADP 3.0
C9500-32QC		1 ASIC
C9500-48Y4C		1 ASIC
C9500-24Y4C		1 ASIC
C9500-16X	1 ASIC	
C9500-40X	2 ASICs	
C9500-12Q	2 ASICs	
C9500-24Q	3 ASICs	

Figures 6 and 7 show the Cisco Catalyst 9500 Series Switches' front panel-to-ASIC connections.




Figure 6.
Cisco Catalyst 9500 Series architecture for models based on UADP 2.0 XL




Figure 7.
Cisco Catalyst 9500 Series High Performance architecture for models based on UADP 3.0

All switches that use the UADP 3.0 ASIC support a unified buffer between the two ASIC cores. That will increase the burst absorption, as it is a single shared buffer between all ports on the switch.




Figure 8.
Difference between the UADP 2.0 (and 2.0 XL) and UADP 3.0 buffers

197. Cisco documentation describes the Catalyst 9600 series as follows:⁸³

Cisco Catalyst 9600 Series Switches are based on a centralized architecture, which means all packets are processed on the supervisor, while line cards are considered transparent, containing only PHYs. Therefore, all QoS resources reside on the supervisor, including the per-port buffer and other QoS resources. The simplicity of this centralized design allows easy feature upgrades just by upgrading the supervisor while keeping the existing line cards.




Figure 9.
Cisco Catalyst 9600 Series architecture

The supervisor architecture is based on UADP 3.0, which offers larger tables for QoS in comparison to UADP 2.0 XL used in the Cisco Catalyst 9500 Series, and a unified buffer between the ASIC cores.

⁸³ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 11–12; see also <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9600-series-switches/nb-06-cat9600-series-data-sheet-cte-en.pdf>.

198. Cisco documentation describes the packet flow in Catalyst Switches as follows:⁸⁴

The packet walk from a QoS operations point of view can be split into four main parts:

1. Ingress classification, policing, and marking performed by the Ingress Forwarding Controller (IFC)
2. Queuing to the stack interface performed by the Ingress Queue Scheduling (IQS) and Stack Queuing and Scheduling (SQS) blocks
3. Egress queueing and scheduling performed by Active Queue Management (AQM)
4. Egress classification, policing, and marking performed by the Egress Forwarding Controller (EFC)

Figure 12 depicts these four parts. Each step is described in detail in later sections.




Figure 12.
QoS tools per ASIC block

199. In the architecture above, Cisco documentation states:⁸⁵

⁸⁴ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14–15.

⁸⁵ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 13.

When a packet enters the UADP ASIC, it goes to the MACsec block for decryption. Next, the ingress FIFO is used to create a copy of the packet, and while the packet is stored unchanged in the Packet Buffer Complex (PBC), the Ingress Forwarding Controller (IFC) will do multiple parallel lookups and will store the lookup results in the descriptor.


- If the packet needs to go over the stack, it will be sent over the Ingress Queue Scheduling (IQS) block and received from the stack by the Stack Queuing and Scheduling (SQS) block.
- If the packet will be sent over the same ASIC, it will skip the stack interface.

When the packet is received either from the stack or from the local PBC, it is ready for egress processing. It is sent to the Egress Queue System (EQS) for queuing. The EQS is built on two sub-blocks: the SQS, which received the packet from the stack, and Active Queue Management (AQM), which manages the port queues. Then a copy of the packet and the information from the packet descriptor that was set on ingress is used by the Egress Forwarding Controller (EFC) to apply the features configured on egress. Once the egress lookup completes, the final result is stored in the descriptor.

The packet is rewritten based on the final value in the descriptor. Next, the packet will be encrypted and sent out of the ASIC.

200. A target port comprises at least a physical egress port and an Egress Queue System, including Stack Queuing and Scheduling (SQS), Active Queue Management (AQM) and Egress Forwarding Controller (EFC). As shown in Figure 11 below, the EQS operates on packets in the packet buffers complex:⁸⁶

⁸⁶ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14.



201. Thus, to the extent the preamble of claim 17 is limiting, Cisco’s Catalyst Switch satisfies the preamble of claim 17.

202. On information and belief, Cisco’s Catalyst Switch meets claim element [17A] of claim 17 of the ‘884 Patent, “monitor a data flow traversing the target port.” For example, Cisco documentation describes the packet flow in Catalyst Switches as follows:⁸⁷

The packet walk from a QoS operations point of view can be split into four main parts:

1. Ingress classification, policing, and marking performed by the Ingress Forwarding Controller (IFC)
2. Queueing to the stack interface performed by the Ingress Queue Scheduling (IQS) and Stack Queueing and Scheduling (SQS) blocks
3. Egress queueing and scheduling performed by Active Queue Management (AQM)
4. Egress classification, policing, and marking performed by the Egress Forwarding Controller (EFC)

⁸⁷ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14–15.

Figure 12 depicts these four parts. Each step is described in detail in later sections.




Figure 12.
QoS tools per ASIC block

203. According to Cisco documentation:⁸⁸

When a packet enters the UADP ASIC, it goes to the MACsec block for decryption. Next, the ingress FIFO is used to create a copy of the packet, and while the packet is stored unchanged in the Packet Buffer Complex (PBC), the Ingress Forwarding Controller (IFC) will do multiple parallel lookups and will store the lookup results in the descriptor.

- If the packet needs to go over the stack, it will be sent over the Ingress Queue Scheduling (IQS) block and received from the stack by the Stack Queuing and Scheduling (SQS) block.
- If the packet will be sent over the same ASIC, it will skip the stack interface.


When the packet is received either from the stack or from the local PBC, it is ready for egress processing. It is sent to the Egress Queue System (EQS) for queuing. The EQS is built on two sub-blocks: the SQS, which received the packet from the stack, and Active Queue Management (AQM), which manages the port queues. Then a copy of the packet and the information from the packet descriptor that was set on ingress is used by the Egress Forwarding Controller (EFC) to apply the features configured on egress. Once the egress lookup completes, the final result is stored in the descriptor.

The packet is rewritten based on the final value in the descriptor. Next, the packet will be encrypted and sent out of the ASIC.

204. A target port comprises at least a physical egress port and an Egress Queue System, including Stack Queuing and Scheduling (SQS), Active Queue Management (AQM) and Egress Forwarding Controller (EFC). As shown in Figure 11 below, the EQS operates on packets in the packet buffers complex.⁸⁹

⁸⁸ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 13.

⁸⁹ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14.



205. A target port has associated output queues. Cisco documentation describes:⁹⁰

⁹⁰ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 25.

To provide this queue structure, the UADP ASIC uses a block called Egress Queue System (EQS). The main components of EQS are as follows:

- **Port queues**

This type of queue allows packets to be grouped and processed in the same way. Having multiple queues allows the ASIC to select which queue to get the next packet from to send out. Queues are built inside of the ASIC per port. Using the grocery store analogy, every queue represents one cashier that processes the items purchased. The express queue shows how people can purchase items faster and get out of the grocery store by reducing the processing time.

Figure 20 shows a visual representation of the port queue structures.




Figure 20.
Front panel port queues

- The port is split into physical queues.
- Every queue has a buffer/memory that can be used to store packets.
- The number of queues used per port can vary from one to eight.

206. Catalyst switches monitor data flows traversing the target port by use of the Egress Forwarding Controller and related egress processing, such as shaping and policing.⁹¹ Furthermore, Cisco documentation states:⁹²

⁹¹ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 24-46 (“Egress Tool Set”), 46-49 (“Hierarchical QoS”).

⁹² See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, pp. 49-52 (“Policy-map counters”).

Policy-map counters

Ingress and egress QoS tools also provide statistics on and visibility into the packets that are processed on the UADP ASIC.

Policy maps support different counters based on the actions used:

- If marking or policing is used, the policy map will collect counters indicating how many packets or bytes are mapped to a class map or dropped by the policer.
- If a queuing action is configured under the interface, the policy map supports enqueue and drop counters.
- If the same policy is shared between two or more ports, the counters are not aggregated between all ports sharing the policy.
- QoS policy-map counters are available via:
 - Operational and configuration YANG models
 - SNMP MIB: **CiscoCBQoS**MIB

The following outputs provide sample counters based on different tools:

Note: The CLI outputs were collected with software release 16.9(2).

- Marking and policer policy counters (**applicable for ingress and egress**):
Notes:
 - If a policy with the same name is shared on multiple ports, the policy map counters will be aggregated, as the TCAM entries are shared between the ports. As result, an interface without traffic might see increments, but the TCAM usage will be reduced, as the entries are shared.
 - If policy-map counters are needed per port, create policy maps with different names and the TCAM entries will not be shared.


207. Therefore, Cisco's Catalyst Switch meets element [17A] of claim 17.

208. On information and belief, Cisco's Catalyst Switch meets claim element [17B] of claim 17 of the '884 Patent, "determine a bandwidth allocation for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow." For example, Cisco documentation states:⁹³

⁹³ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf> at 13.

The packet walk from a QoS operations point of view can be split into four main parts:

1. Ingress classification, policing, and marking performed by the Ingress Forwarding Controller (IFC)
2. Queueing to the stack interface performed by the Ingress Queue Scheduling (IQS) and Stack Queuing and Scheduling (SQS) blocks
3. Egress queueing and scheduling performed by Active Queue Management (AQM)
4. Egress classification, policing, and marking performed by the Egress Forwarding Controller (EFC)



209. In the architecture above, Cisco documentation describes:⁹⁴

⁹⁴ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 13.


When a packet enters the UADP ASIC, it goes to the MACsec block for decryption. Next, the ingress FIFO is used to create a copy of the packet, and while the packet is stored unchanged in the Packet Buffer Complex (PBC), the Ingress Forwarding Controller (IFC) will do multiple parallel lookups and will store the lookup results in the descriptor.

- If the packet needs to go over the stack, it will be sent over the Ingress Queue Scheduling (IQS) block and received from the stack by the Stack Queuing and Scheduling (SQS) block.
- If the packet will be sent over the same ASIC, it will skip the stack interface.

When the packet is received either from the stack or from the local PBC, it is ready for egress processing. It is sent to the Egress Queue System (EQS) for queuing. The EQS is built on two sub-blocks: the SQS, which received the packet from the stack, and Active Queue Management (AQM), which manages the port queues. Then a copy of the packet and the information from the packet descriptor that was set on ingress is used by the Egress Forwarding Controller (EFC) to apply the features configured on egress. Once the egress lookup completes, the final result is stored in the descriptor.

The packet is rewritten based on the final value in the descriptor. Next, the packet will be encrypted and sent out of the ASIC.

210. A target port comprises at least a physical egress port and an Egress Queue System, including Stack Queuing and Scheduling (SQS), Active Queue Management (AQM) and Egress Forwarding Controller (EFC). As shown in Figure 11 below, the EQS operates on packets in the packet buffers complex:⁹⁵



⁹⁵ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14.

211. Cisco documentation further states:⁹⁶

To provide this queue structure, the UADP ASIC uses a block called Egress Queue System (EQS). The main components of EQS are as follows:

- **Port queues**

This type of queue allows packets to be grouped and processed in the same way. Having multiple queues allows the ASIC to select which queue to get the next packet from to send out. Queues are built inside of the ASIC per port. Using the grocery store analogy, every queue represents one cashier that processes the items purchased. The express queue shows how people can purchase items faster and get out of the grocery store by reducing the processing time.

Figure 20 shows a visual representation of the port queue structures.




Figure 20.
Front panel port queues

- The port is split into physical queues.
- Every queue has a buffer/memory that can be used to store packets.
- The number of queues used per port can vary from one to eight.

212. At least through egress packet processing and the EFC, Catalyst switches continuously determine, allocate and monitor the data flow for the target port, the bandwidth allocation for the target port being a bandwidth that is currently allocated for the data flow.

213. Therefore, Cisco's Catalyst Switch meets element [17B] of claim 17.


214. On information and belief, Cisco's Catalyst Switch meets element [17C], "determine a fair-share bandwidth allocation for the target port, the fair-share bandwidth allocation

⁹⁶ *Id.*, p. 25.

being a proportional allocation of a total bandwidth of the network switching element.” For example, Cisco documentation describes the packet flow in Catalyst Switches as follows:⁹⁷

The packet walk from a QoS operations point of view can be split into four main parts:

1. Ingress classification, policing, and marking performed by the Ingress Forwarding Controller (IFC)
2. Queueing to the stack interface performed by the Ingress Queue Scheduling (IQS) and Stack Queuing and Scheduling (SQS) blocks
3. Egress queueing and scheduling performed by Active Queue Management (AQM)
4. Egress classification, policing, and marking performed by the Egress Forwarding Controller (EFC)



215. In the architecture above, Cisco documentation describes:⁹⁸

⁹⁷ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 14–15.

⁹⁸ *Id.*, p. 13.


When a packet enters the UADP ASIC, it goes to the MACsec block for decryption. Next, the ingress FIFO is used to create a copy of the packet, and while the packet is stored unchanged in the Packet Buffer Complex (PBC), the Ingress Forwarding Controller (IFC) will do multiple parallel lookups and will store the lookup results in the descriptor.

- If the packet needs to go over the stack, it will be sent over the Ingress Queue Scheduling (IQS) block and received from the stack by the Stack Queuing and Scheduling (SQS) block.
- If the packet will be sent over the same ASIC, it will skip the stack interface.

When the packet is received either from the stack or from the local PBC, it is ready for egress processing. It is sent to the Egress Queue System (EQS) for queuing. The EQS is built on two sub-blocks: the SQS, which received the packet from the stack, and Active Queue Management (AQM), which manages the port queues. Then a copy of the packet and the information from the packet descriptor that was set on ingress is used by the Egress Forwarding Controller (EFC) to apply the features configured on egress. Once the egress lookup completes, the final result is stored in the descriptor.

The packet is rewritten based on the final value in the descriptor. Next, the packet will be encrypted and sent out of the ASIC.

216. A target port comprises at least a physical egress port and an Egress Queue System, including Stack Queuing and Scheduling (SQS), Active Queue Management (AQM) and Egress Forwarding Controller (EFC). As shown in Figure 11 below, the EQS operates on packets in the packet buffers complex.



217. Cisco documentation further states:⁹⁹

To provide this queue structure, the UADP ASIC uses a block called Egress Queue System (EQS). The main components of EQS are as follows:

- **Port queues**

This type of queue allows packets to be grouped and processed in the same way. Having multiple queues allows the ASIC to select which queue to get the next packet from to send out. Queues are built inside of the ASIC per port. Using the grocery store analogy, every queue represents one cashier that processes the items purchased. The express queue shows how people can purchase items faster and get out of the grocery store by reducing the processing time.

Figure 20 shows a visual representation of the port queue structures.




Figure 20.
Front panel port queues

- The port is split into physical queues.
- Every queue has a buffer/memory that can be used to store packets.
- The number of queues used per port can vary from one to eight.

218. At least through egress packet processing and the EFC, Catalyst switches continuously determine, allocate and monitor the fair- share bandwidth allocation, as the total allocation of all the output queues. The fair- share bandwidth allocation is a proportional allocation of a total bandwidth of the network switching element, the latter being, e.g., the capacity of the Catalyst ASIC. See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9200-c11-742388.pdf>, p. 25.

⁹⁹ See <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/white-paper-c11-742388.pdf>, p. 25.

[series-switches/nb-06-cat9200-ser-data-sheet-cte-en.pdf](https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9200-series-switches/nb-06-cat9200-ser-data-sheet-cte-en.pdf), pp. 4–7,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/nb-06-cat9300-ser-data-sheet-cte-en.pdf>, pp. 5–9,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9400-series-switches/nb-06-cat9400-ser-data-sheet-cte-en.pdf>, pp. 4–10,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9500-series-switches/nb-06-cat9500-ser-data-sheet-cte-en.pdf>, pp. 4–13,
<https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9600-series-switches/nb-06-cat9600-series-data-sheet-cte-en.pdf>, pp. 4–11.


219. Therefore, Cisco’s Catalyst Switch meets claim element [17C] of claim 17.

220. On information and belief, Cisco’s Catalyst Switch meets claim element [17D] of claim 17 of the ’884 Patent, “adjust the bandwidth allocation for the target port based on the fair-share bandwidth allocation.” For example, the discussion for the preamble and claim elements [17A]–[C] are incorporated herein. Furthermore, Cisco documentation describes the packet flow in Catalyst Switches as follows:¹⁰⁰

The packet walk from a QoS operations point of view can be split into four main parts:

1. Ingress classification, policing, and marking performed by the Ingress Forwarding Controller (IFC)
2. Queueing to the stack interface performed by the Ingress Queue Scheduling (IQS) and Stack Queueing and Scheduling (SQS) blocks
3. Egress queueing and scheduling performed by Active Queue Management (AQM)
4. Egress classification, policing, and marking performed by the Egress Forwarding Controller (EFC)

¹⁰⁰ *Id.*, p. 14–15.



221. In the architecture above, Cisco documentation describes:¹⁰¹

When a packet enters the UADP ASIC, it goes to the MACsec block for decryption. Next, the ingress FIFO is used to create a copy of the packet, and while the packet is stored unchanged in the Packet Buffer Complex (PBC), the Ingress Forwarding Controller (IFC) will do multiple parallel lookups and will store the lookup results in the descriptor.

- If the packet needs to go over the stack, it will be sent over the Ingress Queue Scheduling (IQS) block and received from the stack by the Stack Queuing and Scheduling (SQS) block.
- If the packet will be sent over the same ASIC, it will skip the stack interface.


When the packet is received either from the stack or from the local PBC, it is ready for egress processing. It is sent to the Egress Queue System (EQS) for queuing. The EQS is built on two sub-blocks: the SQS, which received the packet from the stack, and Active Queue Management (AQM), which manages the port queues. Then a copy of the packet and the information from the packet descriptor that was set on ingress is used by the Egress Forwarding Controller (EFC) to apply the features configured on egress. Once the egress lookup completes, the final result is stored in the descriptor.

The packet is rewritten based on the final value in the descriptor. Next, the packet will be encrypted and sent out of the ASIC.

222. A target port comprises at least a physical egress port and an Egress Queue System, including Stack Queuing and Scheduling (SQS), Active Queue Management (AQM) and Egress

¹⁰¹ *Id.*, p. 13.

Forwarding Controller (EFC). As shown in Figure 11 below, the EQS operates on packets in the packet buffers complex:¹⁰²



223. Cisco documentation further states:¹⁰³

¹⁰² *Id.*, p. 14.

¹⁰³ *Id.*, p. 25.

To provide this queue structure, the UADP ASIC uses a block called Egress Queue System (EQS). The main components of EQS are as follows:

- **Port queues**

This type of queue allows packets to be grouped and processed in the same way. Having multiple queues allows the ASIC to select which queue to get the next packet from to send out. Queues are built inside of the ASIC per port. Using the grocery store analogy, every queue represents one cashier that processes the items purchased. The express queue shows how people can purchase items faster and get out of the grocery store by reducing the processing time.

Figure 20 shows a visual representation of the port queue structures.




Figure 20.
Front panel port queues

- The port is split into physical queues.
- Every queue has a buffer/memory that can be used to store packets.
- The number of queues used per port can vary from one to eight.

224. At least through egress packet processing and the EFC, Catalyst switches continuously adjust the bandwidth allocation for the target port (a bandwidth that is currently allocated for the data flow) based on the fair-share bandwidth allocation (the total allocation of all the output queues).

225. Therefore, Cisco's Catalyst Switch meets claim element [17D] of claim 17.

226. Accordingly, Cisco's Catalyst Switch satisfies each and every limitation of claim 17 of the '884 Patent.

227. Cisco undertook and continues its infringing actions despite an objectively high likelihood that such activities infringe the '884 Patent, which has been duly issued by the PTO and presumed valid. On information and belief, Cisco could not reasonably, subjectively believe that its actions do not constitute infringement of the '884 Patent. Despite that knowledge and subjective belief, and the objectively high likelihood that its actions constitute infringement, Cisco has continued its infringing activities. As such, Cisco has willfully infringed and/or will continue to willfully infringe the '884 Patent at least as of the date of this Complaint.

228. As a result of Cisco's infringement of the '884 Patent, Brazos has suffered and continues to suffer substantial injury and is entitled to recover all damages caused by Cisco's infringement to the fullest extent permitted by the Patent Act, together with prejudgment interest and costs for Cisco's wrongful conduct.

229. Brazos has no adequate remedy at law to prevent future infringement of the '884 Patent. Brazos suffers and continues to suffer irreparable harm as a result of Cisco's patent infringement and is, therefore, entitled to injunctive relief to enjoin Cisco's wrongful conduct.

PRAYER FOR RELIEF

230. WHEREFORE, Brazos respectfully requests judgment against Cisco as follows:

231. that this Court adjudge that Cisco infringes the '630 Patent, the '286 Patent, the '721 Patent, the '691 Patent, and the '884 Patent.

232. that the Court enter an injunction prohibiting Cisco, and its agents, officers, servants, employees and all persons in active concert or participation with Cisco from deploying, operating, maintaining, testing, using, making, offering to sell, and/or selling the Infringing Products, and from otherwise infringing any of the Patents-in-Suit;

233. that this Court adjudge that Cisco willfully infringes the '630 Patent, the '286 Patent, the '721 Patent, the '691 Patent, and the '884 Patent;

234. that this Court ascertain and award Brazos damages under 35 U.S.C. § 284 sufficient to compensate for Cisco's infringement, including but not limited to infringement occurring before the filing of this lawsuit;

235. that this Court ascertain and award Brazos any post-judgment ongoing royalties under 35 U.S.C. § 284 as may be appropriate;

236. that this Court award Brazos any applicable pre-judgment and post-judgment interest;

237. that this Court award Brazos such other relief at law or in equity as the Court deems just and proper.

JURY DEMAND

238. Brazos requests that all claims and causes of action raised in this Complaint against Cisco be tried to a jury to the fullest extent possible.

Date: May 6, 2024

Respectfully Submitted,

/s/ Joseph M. Abraham

Joseph M. Abraham, TX Bar No. 24088879 –

LEAD ATTORNEY

Timothy Dewberry, TX Bar No. 24090074

FOLIO LAW GROUP PLLC

13492 Research Blvd., Ste. 120, No. 177

Austin, TX 78750

Tel: (737) 234-0201

Email: joseph.abraham@foliolaw.com

timothy.dewberry@foliolaw.com

Gregory P. Love TX SB No. 24013060

Cherry Johnson Siegmund James

400 Austin Ave, Ste. 9th Floor

Waco, TX 76701

Tel: (254) 732-2242

Email: glove@cjsjlw.com

Attorneys for Plaintiff WSOU Investments

d/b/a Brazos Licensing and Development

