

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

ORCKIT CORPORATION,

Plaintiff,

v.

CISCO SYSTEMS, INC.,

Defendant.

Civil Action No. 2:25-cv-181

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Orckit Corporation (“Orckit” or “Plaintiff”) submits this Complaint for patent infringement against Defendant Cisco Systems, Inc. (“Cisco” or “Defendant”), requests a trial by jury, and alleges the following upon actual knowledge with respect to itself and its own acts and upon information and belief as to all other matters:

NATURE OF ACTION

1. This is an action for patent infringement. Orckit alleges that Cisco infringes U.S. Patents Nos. 6,680,904 (“the ’904 Patent”), 8,830,821 (“the ’821 Patent”), and 10,652,111 (“the ’111 Patent”) (collectively, “the Asserted Patents”), copies of which are attached hereto.

2. Orckit alleges that Cisco: (1) directly and indirectly infringes the Asserted Patents by making, using, offering for sale, selling, and importing certain networking hardware and software; (2) induces infringement of the Asserted Patents and contributes to others’ infringement of the Asserted Patents; and (3) infringes the Asserted Patents willfully. Orckit seeks damages and other relief for Cisco’s wrongful conduct.

PARTIES

3. Orckit is a Delaware corporation and owns the Asserted Patents by assignment.

4. Cisco is a Delaware corporation with its principal place of business at 300 East Tasman Drive, Building 10, San Jose, California 95134.

5. Cisco is registered to do business in Texas, maintains places of business in Texas, and conducts business in Texas. Cisco has at least one place of business in this district, including a 162,000 square foot data center at 2260 Chelsea Boulevard, Allen, Texas 75013. The Collin County Appraisal District appraised this facility at a value of nearly \$100,000,000.

6. Cisco has a permanent and continuous presence in Texas and a regular and established place of business in the Eastern District of Texas.

JURISDICTION AND VENUE

7. This action arises under the patent laws of the United States, 35 U.S.C. § 271 *et seq.* The Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

8. The Court has personal jurisdiction over Cisco. As alleged above, Cisco has sufficient minimum contacts with Texas so that this action does not offend due process or the traditional notions of fair play and substantial justice and so that Texas's long-arm statute is satisfied. Among other factors, Cisco is registered in Texas, is domiciled in this district, and has a continuous presence in and systematic contact with this district. Specifically, Cisco regularly conducts business at its facilities in Richardson and Allen and derives substantial revenue from the goods and services that it provides to its customers in Texas. Cisco also undertakes a portion of its infringing activities in Texas—including by making, using, importing, offering for sale, and selling products and services that infringe the Asserted Patents—directly and through its distributors, retailers, and other intermediaries.

9. Venue is proper in this judicial district pursuant to 28 U.S.C. §§1391(b), (c), (d) and 1400(b) because Cisco has a permanent and continuous presence in, has committed acts of infringement in, and maintains a regular and established place of business in this district.

BACKGROUND

10. This is Orckit's second action against Cisco related to the Asserted Patents.

11. On July 22, 2022, Orckit filed a complaint against Cisco in *Orckit Corporation v. Cisco Systems, Inc.*, 2:22-cv-00276-JRG-RSP (E.D. Tx July 22, 2022), alleging that Cisco infringed U.S. Patents Nos. 6,680,904 ("the '904 Patent"), 7,545,740 ("the '740 Patent"), 8,830,821 ("the '821 Patent"), and 10,652,111 ("the '111 Patent"). *See* Dkt. No. 1 (the "Original Complaint").

12. Cisco filed petitions seeking inter partes review (IPR) of the patents that Orckit asserted in its Original Complaint, and, before the issuance of any institution decisions by the Patent Trial and Appeal Board ("PTAB"), Cisco filed a motion to stay the case, which the Court denied. Dkt. Nos. 55, 56.

13. In the months that followed, pursuant to the Court's Amended Docket Control Order, the parties engaged in fact discovery (*e.g.*, exchanged and responded to interrogatories, produced a substantial amount of documents, completed depositions of the inventors, were in the process of scheduling depositions of Cisco witnesses, and were meeting and conferring productively to resolve their discovery disputes) and presented their respective positions on claim construction in the Markman hearing that took place on September 7, 2023.

14. Between September 11, 2023 and September 20, 2023, the PTAB issued decisions instituting IPR proceedings with respect to three of four of the patents that Orckit asserted in its Original Complaint.¹

15. Given the PTAB's decisions to institute IPR proceedings, and to conserve party and Court resources, Orckit and Cisco entered into a joint stipulation that Orckit's case would be dismissed without prejudice, with each party bearing its own costs and attorneys' fees, which the Court accepted and acknowledged, ordering that the case be dismissed without prejudice on October 18, 2023. Dkt Nos. 99, 103.

16. On September 17, 2024, the PTAB issued its final written decision that claims 1-9, 12-24, and 27-31 of the '111 Patent were not unpatentable. *Cisco Systems, Inc. and Juniper Networks, Inc. v. Orckit Corporation*, IPR2023-00554, Paper 43 at 24 (PTAB Sept. 17, 2024).

17. On October 8, 2024, the PTAB issued its final written decision that claims 6, 17, and 23 of the '904 Patent were not unpatentable. *Cisco Systems, Inc. v. Orckit Corporation*, IPR2023-00714, Paper 35 at 45 (PTAB Oct. 8, 2024).

18. On October 22, 2024, the PTAB issued its final written decision that claims 14-16 of the '821 Patent were not unpatentable. *Cisco Systems, Inc. and Juniper Networks, Inc. v. Orckit Corporation*, IPR2023-00402, Paper 40 at 68 (PTAB Oct. 22, 2024).²

19. In view of the PTAB's Final Written Decisions as to Cisco's IPRs, Orckit now brings this new action based on the following allegations (which largely mirror those which Orckit made in the Original Complaint). Given the significant progress the parties had made toward

¹ The PTAB issued a fourth institution decision on October 10, 2023 with respect to the fourth of Orckit's four patents.

² Orckit has appealed a separate final written decision from the PTAB concerning the patentability of the claims in the '740 patent, which the PTAB issued on October 22, 2024.

completion of fact discovery and claim construction in Orckit's first action, Orckit intends to seek an expedited schedule for this second action, subject to meeting and conferring with Cisco and the Court's approval.

FACTUAL ALLEGATIONS

Orckit Communications Ltd. and Its Breakthrough Communications Technology

20. The patented technology is rooted in research by Orckit Communications Ltd. (later reorganized and renamed Orckit-Corrigent Ltd.), a company founded in Israel in 1990 by Izhak Tamir. The company was a pioneer in the development of infrastructure-level networking products, and in its first decade became the market leader in Asymmetric Digital Subscriber Line (ADSL) technology, winning a client base that included some of the world's pre-eminent telecommunications providers. The company went public, and in 1996 was listed on the Nasdaq Stock Exchange in the United States.

21. Building on that initial success, Orckit Communications Ltd. turned its attention to overcoming significant limitations in Ethernet technology, the predominant technology used for local area networks used in offices, schools and other local environments. With the proliferation of data and the development of the Internet, demand for data transmission skyrocketed. While Ethernet could be used to connect a limited number of computers, it was not well-suited for the delivery of video, voice, and other applications with higher bandwidth requirements for a larger number of users. The existing standard for delivering voice communications, known as the Synchronous Optical Network ("SONET") protocol, was not a viable alternative because it was not designed to process data in an efficient and scalable way. As a result, providers like cable companies were required to develop and install their own infrastructure to deliver services and could not rely on a single network to provide different services in parallel.

22. Orckit Communications Ltd.’s solutions addressed those shortcomings. It quickly recognized that existing solutions could accommodate network traffic only so long as data occupied only a small portion of overall network traffic. The company’s technology overcame those limitations by enhancing Ethernet switching and routing to optimize the transmission of data, voice and video, including those using Internet Protocol (“IP”) telecommunications networks. The capacity, reliability, and resilience offered by Orckit Communications Ltd.’s inventions opened the possibility of the transmission of data, voice, and video services on the same network—the hugely valuable “bundled services” or “triple-play services” sought by both telecommunications companies and their customers.

23. Between 2000 and 2010, Orckit Communications Ltd. invested hundreds of millions of US dollars in research and development of those solutions. It earned recognition around the world for those innovations and won contracts to rebuild national telecommunications infrastructure systems along with hundreds of patents—including those at issue in this lawsuit.

24. With the economic downturn of 2007 and 2008, many of Orckit Communications Ltd.’s most significant potential customers dramatically reduced their infrastructure spending. Even with its superior technology the company was unable to weather the global recession and ultimately went into liquidation.

25. Plaintiff Orckit Corporation obtained all rights to the Asserted Patents.

The Asserted Patents


U.S. Patent No. 6,680,904

26. Orckit is the lawful owner of all rights, title and interest in U.S. Patent No. 6,680,904 (“the ’904 Patent”) entitled “BI-DIRECTIONAL CHAINING OF NETWORK ACCESS PORTS” (attached as Exhibit 1), including the right to sue and recover for infringement

thereof. The '904 Patent was duly and legally issued on January 20, 2004, naming Menachem Kaplan, David Zelig, Roy Kinamon, Eli Aloni, Ron Sdayoor, Eric Paneth and Eli Magal as the inventors.

27. The '904 Patent has 26 claims: six independent claims and 20 dependent claims.

28. The '904 Patent presented novel and unconventional apparatuses and methods for (among other things) “efficient, high-speed transfer of data packets within an access multiplexer system.” Ex. 1, '904 Patent at 1:65-67. The inventions patented in the '904 Patent include, for example, “slave” and “master” units that are “connected in one or more daisy chains between the active and standby masters and are configured so that both downstream and upstream packets can be transmitted in either direction along each of the chains.” *Id.* at 2:11-14. Thus, “if a failure occurs in any one of the slaves or in a link between them, the traffic direction in the chain in which the failure has occurred is simply reversed so as to run through the standby master.” *Id.* at 2:15-18. “An advantage of the architecture of system 31 is that additional slaves may be added to the chains as needed, without having to change the number of interfaces associated with masters 30, and 32.” *Id.* at 6:33-36. One embodiment of the inventions of the '904 Patent is shown in Fig. 3, reproduced below:



29. The claims of the '904 Patent, including claim 6 (reproduced below, together with claims 4, and 5, from which it depends), recite at least these inventive concepts of the '904 Patent:

.4. Network access apparatus, comprising:

first and second master units, each comprising a physical interface to a packet-switched network;

a plurality of slave units, each slave unit comprising one or more ports to respective subscriber lines; and

a plurality of physical interface lines, which link the slave units in one or more daisy chains, in which the slave units are mutually connected in series by the physical interface lines therebetween, each daisy chain comprising at least a first slave unit connected by one of the physical interface lines to the first master unit and a last slave unit connected by another of the physical interface lines to the second master unit,

wherein in normal operation, downstream data packets received from the network are passed from the first master unit to each of the daisy chains via the first slave unit in each chain, and upstream data packets received by the slaves in each chain from the subscriber lines are passed via the first slave unit in the chain to the first master unit for transmission over the network.

5. Apparatus according to claim 4, and comprising a protection interface, which couples the second master unit to the first master unit, and over which interface data packets are conveyed between the first and second master units in case of a fault.

6. Apparatus according to claim 5, wherein the first master unit multicasts the upstream data packets that it receives from the slave units to the network and, via the protection interface, to the second master unit, which transmits the upstream data packets to the network.

Id. at 11:61-12:23 (claims 4, 5, 6).

30. The subject matter described and claimed in the '904 Patent, including the subject matter of claim 6, represented an improvement in computer and communications functionality, performance, and efficiency, and was novel and not well-understood, routine, or conventional at the time of the invention of the '904 Patent.


31. Cisco had knowledge of the '904 Patent, including at least as of March 2017 when Orckit IP LLC (“Orckit IP”)—a prior owner of the Asserted Patents—initiated discussions with Cisco about its patent portfolio, including the Asserted Patents, as described and alleged below, and at least as of the filing of the Original Complaint.

U.S. Patent No. 8,830,821

32. Orckit is the lawful owner of all right, title, and interest in U.S. Patent No. 8,830,821 (“the '821 Patent”) entitled “METHOD FOR SUPPORTING MPLS TRANSPORT PATH RECOVERY WITH MULTIPLE PROTECTION ENTITIES” (attached as Exhibit 2), including the right to sue and recover for infringement thereof. The '821 Patent was duly and legally issued on September 9, 2014, naming Daniel Cohn and Rafi Ram as the inventors.

33. The '821 Patent has 20 claims: three independent claims and 17 dependent claims.

34. The '821 Patent presented novel and unconventional apparatuses and methods for (among other things) selecting network transport entities between a first and second endpoint, using working and protection entities to minimize simultaneous failure and/or a cost function. Ex. 2, '821 Patent, at Abstract; 2:5-21. The inventions patented in the '821 Patent include, for example, switching between working and protection entities, determining a probability of concurrent failure of both entities, and reselecting an entity pair. *Id.* at 2:32-43. One embodiment of the inventions of the '821 Patent is shown in Fig. 1, reproduced below:



35. The claims of the '821 Patent, including claim 14 (reproduced below), recite at least these inventive concepts of the '821 Patent:

14. A system for selecting entities within an MPLS network, comprising:

a data structure comprising a plurality of transport entity descriptors;

an entity protection switch configured to switch between a working entity and a protection entity; and

digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity;

logic configured to determine an entity cost of said plurality of transport entity descriptors; and

logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,

wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in over all cost for one of said plurality of transport entities.

Id. at 8:42-63 (claim 14).

36. The subject matter described and claimed in the '821 Patent, including the subject matter of claim 14, represented an improvement in computer and communications functionality, performance and efficiency, and was novel and not well-understood, routine, or conventional at the time of the invention of the '821 Patent.

37. Cisco had knowledge of the '821 Patent, including at least as of March 2017 when Orckit IP LLC initiated discussions with Cisco about its patent portfolio, including the Asserted Patents, as described and alleged below, and at least as of the filing of the Original Complaint.

U.S. Patent No. 10,652,111

38. Orckit is the lawful owner of all right, title, and interest in U.S. Patent No. 10,652,111 (“the '111 Patent”) entitled “METHOD AND SYSTEM FOR DEEP PACKET INSPECTION IN SOFTWARE DEFINED NETWORKS” (attached as Exhibit 3), including the right to sue and recover for infringement thereof. The '111 Patent was duly and legally issued on May 12, 2020, naming Yossi Barsheshet, Simhon Doctori and Ronen Solomon as the inventors.

39. The '111 Patent has 54 claims: two independent claims and 52 dependent claims.

40. The '111 Patent presented novel and unconventional methods for (among other things) “deep packet inspection (DPI) in a software defined network (SDN), wherein the method is performed by a central controller of the SDN.” Ex. 3, '111 Patent at 2:28-30. As an example, unlike the prior art, the inventions patented in the '111 Patent enable the inspection or extraction of content from data packets belonging to a specific flow or session, thereby enabling security threat detection. *Id.* at 1:61-67. The patented inventions also decrease traffic delays between a client and server, avoid overflowing the controller with data, and prevent the concentration of data

traffic through a single point of potential failure. *Id.* at 2:1-7. One embodiment of the inventions of the '111 Patent is shown in Fig. 1, reproduced below:

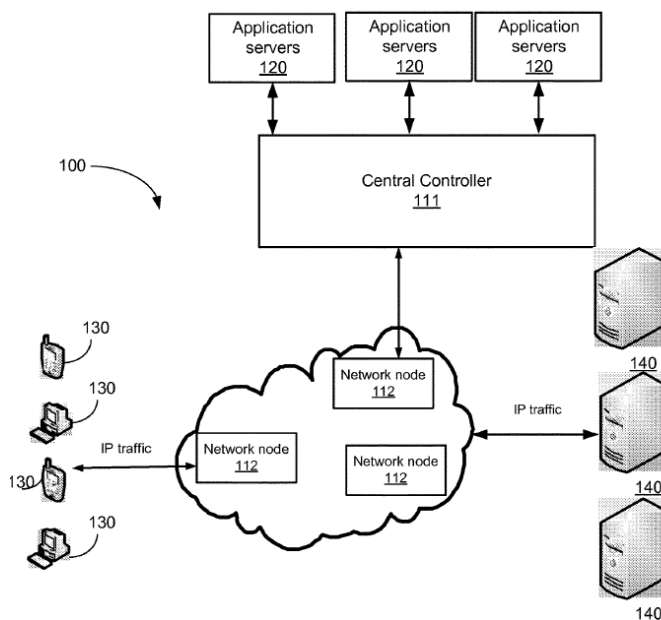


FIG. 1

41. The claims of the '111 Patent, including claim 1 (reproduced below), recite at least these inventive concepts of the '111 Patent:

1. A method for use with a packet network including a network node for transporting packets between first and second entities under control of a controller that is external to the network node, the method comprising:

sending, by the controller to the network node over the packet network, an instruction and a packet-applicable criterion;

receiving, by the network node from the controller, the instruction and the criterion;
receiving, by the network node from the first entity over the packet network, a packet addressed to the second entity;

checking, by the network node, if the packet satisfies the criterion;

responsive to the packet not satisfying the criterion, sending, by the network node over the packet network, the packet to the second entity; and

responsive to the packet satisfying the criterion, sending the packet, by the network node over the packet network, to an entity that is included in the instruction and is other than the second entity.

Id. at 10:52-11:4 (claim 1).

42. The subject matter described and claimed in the '111 Patent, including the subject matter of claim 1, represented an improvement in computer and communications functionality, performance and efficiency, and was novel and not well-understood, routine, or conventional at the time of the invention of the '111 Patent.

43. Cisco had knowledge of the '111 Patent, including at least as of the filing of the Original Complaint.

BACKGROUND OF CISCO'S INFRINGING CONDUCT

44. Defendant Cisco Systems Inc. is a computer networking company that makes, uses, sells, offers for sale in the United States, and/or imports into the United States, or has otherwise made, used, sold, offered for sale in the United States, and/or imported in the United States, routers, switches, and other networking equipment and software that infringe the Asserted Patents, and also has induced and contributed to and continues to induce and contribute to infringement of others who have made, used, sold, offered for sale in the United States, and/or imported in the United States, routers, switches, and other networking equipment and software that infringe the Asserted Patents.³

45. A non-comprehensive list of products that infringe the Asserted Patents is set out in Appendices A-C hereto ("the Accused Products"). Cisco's infringement includes the making, using, selling, offering for sale and/or importing the listed products, and Cisco's active inducement of infringement, including by supplying the listed products to third parties that use those products

³ Orckit notes that the '904 Patent expired on December 27, 2019 and does not therefore allege that Cisco's infringement of that patent continued beyond that date.

to practice the claimed methods of the Asserted Patents. Orckit reserves the right to supplement and amend the list of Accused Products recited in Appendices A-C as permitted by the Court.

46. Cisco infringes and continues to infringe the Asserted Patents by making, using, selling, offering to sell, and/or importing, without license or authority, the Accused Products as alleged herein.

47. Cisco markets, advertises, offers for sale, and/or otherwise promotes the Accused Products and does so to induce, encourage, instruct, and aid one or more persons in the United States to make, use, sell, and/or offer to sell their Accused Products. For example, Cisco advertises, offers for sale, and/or otherwise promotes the Accused Products on its web site. Cisco further publishes and distributes data sheets, manuals, and guides for the Accused Products, as set forth in detail below. Therein, Cisco describes and touts the use of the subject matter claimed in the Asserted Patents, as described and alleged below.

**BACKGROUND OF CISCO’S KNOWLEDGE OF THE INVENTIONS DESCRIBED
AND CLAIMED IN THE ASSERTED PATENTS**

48. Cisco has had knowledge of the Asserted Patents and the inventions described and claimed therein since at least around March 2017, when Orckit IP—a prior owner of the Asserted Patents—initiated discussions with Cisco about the Asserted Patents and the Accused Products. On March 20, 2017 Orckit IP sent a letter to Cisco concerning its “Patent Portfolio.” Ex. 4 (“March 2017 Letter from Orckit IP to Cisco”). In that letter, Orckit IP notified Cisco that it:

...owns a patent portfolio related to certain communications technologies developed by Orckit Communications Ltd. and Corrigent Systems Ltd. (f/k/a Orckit-Corrigent Ltd.). Orckit IP’s patent portfolio includes over 100 patents and pending patent applications. One or more of these patents and patent applications may be of interest to Cisco and require your company’s attention.

Ex. 4 at 1.

49. Orckit IP further identified several “Cisco switches and routers,” including certain of the Accused Products, which are accused of infringing the Asserted Patents. *Id.* Orckit IP concluded that “Cisco may be interested in obtaining a license to (or acquiring) the ’983 Patent and/or other patent assets from Orckit IP’s patent portfolio.” *Id.* at 2.

50. On April 10, 2017, Cisco responded by letter and requested additional information. Ex. 5 (“April 2017 Letter from Cisco to Orckit IP”). On July 11, 2018, Orckit IP sent a second notice letter to Cisco, again concerning its “Patent Portfolio.” Ex. 6 (“July 2018 Letter from Orckit IP to Cisco”). Orckit IP again notified Cisco that Orckit IP’s patent portfolio relates to Cisco’s switch and router products and concluded that “Cisco may be interested in obtaining a license to (or acquiring) the ’821 Patent, the ’928 Patent, and/or other patent assets from Orckit IP’s patent portfolio (in addition to the ’983 Patent, discussed above).” Ex. 6 at 2.

51. On July 25, 2018, Cisco responded by letter and requested additional information. Ex. 7 (“July 2018 Letter from Cisco to Orckit IP”).

52. On November 20, 2018, Orckit IP identified additional patents within its patent portfolio, including the asserted ’904 Patent. Ex. 8 (“November 2018 Email from Orckit IP to Cisco”). Orckit IP offered to send Cisco exemplary “evidence of use charts” relating to any of the patents, including the asserted ’904 Patent. Ex. 8 at 2.

53. Cisco has also had knowledge of the Asserted Patents and the inventions described and claimed therein since at least as of the filing of the Original Complaint.

COUNT ONE: INFRINGEMENT OF U.S. PATENT 6,680,904

54. Cisco directly infringes at least claim 6 of the ’904 Patent by making, using, offering for sale, selling, and/or importing products, including at least the Accused Products, which include but are not limited to the products set forth in Appendix A (“the ’904 Accused Products”),

that meet every limitation, either literally or under the doctrine of equivalents, of at least claim 6 of the '904 Patent, in violation of 35 U.S.C. § 271(a).

55. The '904 Accused Products, including the Cisco 550X Series Stackable Managed Switches (“Cisco 550X”), which is exemplary of all of the '904 Accused Products, constitute network access apparatuses. *See, e.g.*, “Cisco 550X Series Stackable Managed Switches” Data Sheet (available at <https://www.cisco.com/c/en/us/products/collateral/switches/550x-series-stackable-managed-switches/datasheet-c78-735874.pdf>) at 3 (“The Cisco® 550X Series (Figure 1) are the next-generation stackable managed Ethernet switches that provide the advanced capabilities and superior performance you need to support a more demanding network environment at an affordable price.”):

Cisco 550X Series Stackable Managed Switches

The Cisco® 550X Series (Figure 1) are the next-generation stackable managed Ethernet switches that provide the advanced capabilities and superior performance you need to support a more demanding network environment at an affordable price. These switches incorporate fan and power hardware redundancy, increasing overall network availability. The SG550X and SF550X models provide 24 or 48 ports of Gigabit Ethernet and Fast Ethernet connectivity with 10 Gigabit uplinks. The SX550X models provide 12, 16, 24, or 48 ports of 10 Gigabit Ethernet with both copper and fiber connection options, providing a solid foundation for your current business applications, as well as those you are planning for the future. At the same time, these switches are easy to deploy and manage, without a large IT staff.




Figure 1.
Cisco 550X Series Stackable Managed Switches

Cisco 550X Series switches are designed to protect your technology investment as your business grows. Unlike switches that claim to be stackable but have elements that are administered and troubleshot separately, the Cisco 550X Series provides true stacking capability, allowing you to configure, manage, and troubleshot multiple physical switches as a single device and more easily expand your network.

A true stack delivers a unified data and control plane, in addition to a management plane, providing flexibility, scalability, and ease of use because the stack of units operate as a single entity constituting all the ports of the stack members. The switches also protect your technology investment with an enhanced warranty, dedicated technical support, and the ability to upgrade equipment in the future and receive credit for your Cisco 550X Series switch. Overall, the Cisco 550X Series provides the ideal technology foundation for a growing business.


For example, the '904 Accused Products, including the Cisco 550X, contain one or more Ethernet switches, *i.e.*, network access apparatuses.

56. The '904 Accused Products, including the Cisco 550X, are network apparatuses and comprise first and second master units, each comprising a physical interface to a packet-switched network. See Cisco YouTube Video entitled “What Is Stacking” (available at https://www.youtube.com/watch?v=bUGRT_ncDMU) at 1:55:



For example, the '904 Accused Products, including the Cisco 550X, contain “MASTER” and “MASTER BACKUP” units with ports, *i.e.*, first and second master units, each comprising a physical interface to a packet-switched network.

57. The '904 Accused Products, including the Cisco 550X, comprise a plurality of slave units, each comprising one or more ports to respective subscriber lines. See also Cisco YouTube Video entitled “What Is Stacking” (available at https://www.youtube.com/watch?v=bUGRT_ncDMU) at 1:55:



For example, the '904 Accused Products, including the Cisco 550X, include several “SLAVE” units with ports, *i.e.*, a plurality of slave units, each slave unit comprising one or more ports to respective subscriber lines.

58. The '904 Accused Products, including the Cisco 550X, comprise a plurality of physical interface lines, which link the slave units in one or more daisy chains, in which the slave units are mutually connected in series by the physical interface lines therebetween, each daisy chain comprising at least a first slave unit connected by one of the physical interface lines to the first master unit and a last slave unit connected by another of the physical interface lines to the second master unit. *See, e.g.*, “Chain and Ring Topologies on the SG550XG and SG350XG Switches” ([available at https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-350x-series-stackable-managed-switches/smb5237-chain-and-ring-topologies-on-the-sg550xg-and-sg350xg-switches.pdf](https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-350x-series-stackable-managed-switches/smb5237-chain-and-ring-topologies-on-the-sg550xg-and-sg350xg-switches.pdf)) at 1-2. (“A chain topology is a linear connection between all units via stacking links. Starting with one switch, each unit connects to its next, neighboring switch through a single link between their stack ports, until the last unit has been linked with the one before it... In a Ring topology, all units in the stack are connected in a loop, creating failover capability. It is similar to a chain, except the last unit connects back to the first unit providing additional redundancy in the case of a failed stack link.”); *see also id.* at 2:

Setting Up Chain and Ring Topologies

To physically set up the two stack topologies in this demonstration, we will use 4 SG550XG Switches.

Chain Topology

Step 1. Take a cable and connect the first and second switch together. To connect units to each other with the stacking links, you can use any network port on the switch as a stack port.

Note: Take note of the port numbers you use to connect the switches. You will need to designate these ports as stack ports in the Graphical User Interface Configuration for the stack topology.

Step 2. Connect the second and third switch together using a stacking cable.

Step 3. Connect the third and fourth switch together using a stacking cable.

Note: If you have more than four units in your stack, repeat this process for every subsequent switch until the last unit is connected to the one before it.


Ring Topology

Step 1. Follow the Chain Topology Physical Configuration steps to connect your switches into a chain topology. A ring topology uses the same configuration as a chain, except the last unit connects back to the first.

Step 2. Connect the last switch back to the first switch using a stacking cable.

See also, Cisco YouTube Video entitled “What Is Stacking” (available at

https://www.youtube.com/watch?v=bUGRT_ncDMU) at 1:55:



For example, in one illustration, the '904 Accused Products, including the Cisco 550X, contain cables that connect two master units and three slave units, *i.e.*, a plurality of physical interface lines, which link the slave units in one or more daisy chains, in which the slave units are mutually connected in series by the physical interface lines therebetween, each daisy chain comprising at

least a first slave unit connected one of the physical interface lines to the first master unit and a last slave unit connected by another of the physical interface lines to the second master unit.

59. The '904 Accused Products, including the Cisco 550X, operate such that in normal operation, downstream data packets received from the network are passed from the first master unit to each of the daisy chains via the first slave unit in each chain, and upstream data packets received by the slaves in each chain from the subscriber lines are passed via the first slave unit in the chain to the first master unit for transmission over the network. *See, e.g.*, “Cisco 550X Series Stackable Managed Switches” Data Sheet (available at <https://www.cisco.com/c/en/us/products/collateral/switches/550x-series-stackable-managed-switches/datasheet-c78-735874.html>) at 7.

- Storm control can be applied to broadcast, multicast, and unknown unicast traffic.

For example, the Cisco 550X is configured in normal operation to broadcast data to each of the daisy chains, as shown above, with which data packets received from the network by the first master unit are transmitted to each daisy chain via the first slave unit in each chain. For further example, the Cisco 550X is configured in normal operation to transmit data bidirectionally, as is evident from its performance of Unidirectional Link Detection, *i.e.*, upstream data packets received by the slaves in each chain from the subscriber lines are passed via the first slave unit in the chain to the first master unit for transmission over the network. *Id* at 11.

Unidirectional Link Detection (UDLD)

UDLD monitors physical connection to detect unidirectional links caused by incorrect wiring or port faults to prevent forwarding loops and blackholing of traffic in switched networks

60. The '904 Accused Products, including the Cisco 550X, are network apparatuses that comprise a protection interface, which couples the second master unit to the first master unit,

and over which interface data packets are conveyed between the first and second master units in case of a fault. *See, e.g., id* at 4.

The Cisco 550X Series provides an additional layer of resiliency with support for the Virtual Router Redundancy Protocol (VRRP). VRRP lets you extend the same resiliency that stacking provides for individual switches to complete network domains. By running VRRP between two stacks, you can instantly cut over from one stack to another in the event of a problem and continue operating even after a failure.

For example, in the event of a fault, the first master is disabled and data is transmitted from the first master to the second master. *See id* at 4.

High reliability and resiliency

In a growing business in which availability 24 hours a day, 7 days a week is critical, you need to assure that employees and customers can always access the data and resources whenever they need. In these environments, stackable switches can play an important role in minimizing downtime and improving network resiliency. For example, if the master switch within a Cisco 550X Series stack fails, another switch takes over, keeping your network up and running. You can also replace individual devices in the stack without taking your whole network offline or affecting employee productivity.

See also “Cisco 350 & 550 Series Managed Switches Administration Guide” (https://www.cisco.com/c/en/us/td/docs/switches/lan/csbms/350xseries/2_5_7/Administration/tesla-350-550.pdf) at 418.

Unit Failure in Stack

If the active unit fails, then the standby unit will take over the primary role and continues to operate the stack normally.

For the standby switch to be able to take the place of the active switch, both units remain on reserve at all times. When on reserve mode, the active switch and its standby switches are synchronized with a static configuration (contained in both the Startup and Running configuration files). The standby switch configuration file remains on the previous active switch.

61. The '904 Accused Products, including the Cisco 550X, are network apparatuses wherein the first master unit multicasts the upstream data packets that it receives from the slave units to the network and, via the protection interface, to the second master unit, which transmits the upstream data packets to the network. *Id.*

Unit Failure in Stack

If the active unit fails, then the standby unit will take over the primary role and continues to operate the stack normally.

For the standby switch to be able to take the place of the active switch, both units remain on reserve at all times. When on reserve mode, the active switch and its standby switches are synchronized with a static configuration (contained in both the Startup and Running configuration files). The standby switch configuration file remains on the previous active switch.

For example, if the first master fails to send data packets from the slave units to the network, then the second master unit sends these packets to the network, and the standby and active units both “remain on reserve at all times” and, when on reserve, “the active switch and its standby switches are synchronized,” *i.e.*, there is a connection between both master units via the protection interface.

62. The '904 Accused Products, including the Cisco 550X, further support port mirroring. *See* “Configuration of Port and VLAN Mirroring on the Sx500 Series Stackable Switches” Product Support (available at <https://www.cisco.com/c/en/us/support/docs/smb/switches/cisco-small-business-500-series-stackable-managed-switches/smb2997-configuration-of-port-and-vlan-mirroring-on-the-sx500-series.html>).

Port Mirroring is a method used to monitor network traffic. With Port Mirroring, copies of incoming and outgoing packets at the ports (Source Ports) of a network device are forwarded to another port (Target Port) where the packets are studied. This is used as a diagnostic tool by the network administrator.

63. With knowledge of the '904 Patent, Cisco has actively induced the direct infringement of one or more claims of the '904 Patent, including claim 6 and claim 17, in violation of 35 U.S.C. § 271(b) by its customers and/or end users of their products, including at least the '904 Accused Products, by selling products with a particular design, providing support for, providing instructions for use of, and/or otherwise encouraging its customers and/or end-users to directly infringe, either literally and/or under the doctrine of equivalents, one or more claims of

the '904 Patent, including claim 6 and claim 17, with the intent to encourage those customers and/or end-users to infringe the '904 Patent.

64. By way of example, Cisco has actively induced infringement of the '904 Patent by encouraging, instructing, and aiding one or more persons in the United States, including but not limited to customers and end users who purchase, test, operate, and use Cisco's products, including at least the '904 Accused Products, to make, use, sell, and/or offer to sell Cisco's products, including at least the '904 Accused Products, in a manner that infringes at least one claim of the '904 Patent, including claim 6 and claim 17.

65. As a result of Cisco's inducement of infringement, its customers and/or end users made, used, sold, offered for sale, or imported Cisco's products, including the '904 Accused Products, in ways that directly infringe one or more claims of the '904 Patent, including claim 6 and claim 17, such as in the manner described above with respect to the Cisco 550X. Cisco had knowledge of its customers' and/or end users' direct infringement at least by virtue of its sales, instruction, and/or promotion of Cisco's products, including the Accused Products, at least as of March 2017 when Orckit IP initiated discussions with Cisco about its patent portfolio, including the Asserted Patents.

66. Cisco has also contributed to the infringement by others, including its customers and/or the end users of its products, of at least claim 6 and claim 17 of the '904 Patent under 35 U.S.C. § 271(c) by, among other things, selling, offering for sale within the United States and/or importing into the United States or otherwise making available the '904 Accused Products for use in practicing the patented inventions of the '904 Patent, knowing that the '904 Accused Products are especially made or adapted for use in infringement of the '904 Patent, are used in practicing the method and process claims of the '904 Patent, embody a material part of the inventions claimed

in the '904 Patent, and are not staple articles of commerce suitable for substantial non-infringing use. Cisco's customers and/or the end users of the '904 Accused Products directly infringed the '904 Patent by using the '904 Accused Products.

67. With knowledge of the '904 Patent, Cisco has willfully, deliberately, and intentionally infringed the '904 Patent. Cisco had actual knowledge of the '904 Patent and Cisco's infringement of the '904 Patent as set forth above. After acquiring that knowledge, Cisco directly and indirectly infringed the '904 Patent as set forth above. Cisco knew or should have known that its conduct amounted to infringement of the '904 Patent at least because Orckit IP notified Cisco of the '904 Patent and its infringement of the '904 Patent as set forth above.

68. Cisco, by way of its infringing activities, has caused Orckit to suffer damages in an amount to be determined.

69. Orckit is entitled to recover from Cisco damages at least in an amount adequate to compensate for its infringement of the '904 Patent, which amount has yet to be determined, together with interest and costs determined by the Court.

70. Orckit has complied with the requirements of 35 U.S.C. § 287 with respect to the '904 Patent.

COUNT TWO: INFRINGEMENT OF U.S. PATENT 8,830,821

71. Cisco directly infringes at least claim 14 of the '821 Patent by making, using, offering for sale, selling, and/or importing products, including at least the Accused Products, which include but are not limited to the products set forth in Appendix B ("the '821 Accused Products"), that meet every limitation, either literally or under the doctrine of equivalents, of at least claim 14 of the '821 Patent, in violation of 35 U.S.C. § 271(a).

72. The '821 Accused Products, including the Cisco Network Convergence System 4000 Series ("Cisco NCS 4000"), which is exemplary of all of the '821 Accused Products, constitute systems for selecting entities within an MPLS network. *See, e.g.*, "Cisco Network Convergence System 4000 Series" Data Sheet (available at <https://www.cisco.com/c/en/us/products/collateral/optical-networking/network-convergence-system-4000-series/datasheet-c78-729222.html>) at 3:

Product overview

The Cisco® Network Convergence System 4000 (NCS 4000) Series is a converged optical service platform providing Dense Wavelength-Division Multiplexing (DWDM), Optical Transport Network (OTN), Multiprotocol Label Switching (MPLS), Carrier Ethernet, and Label Switch Router (LSR) or IP multiservice capabilities (Figure 1). It delivers massive scale through a state-of-the-art silicon and system design, while offering dramatic network efficiency and simplification led by innovations in usability, automation, service management, turn-up, and monitoring.



Figure 1.
Cisco NCS 4016 Chassis (Right) and NCS 4009 Chassis (Left)

See also, "Configuration Guide for Cisco NCS 4000 Series" (available at

https://www.cisco.com/c/en/us/td/docs/routers/ncs4000/software/configuration_guide/configurationguide.pdf) at 343:

Flex LSP also known as Associated Bidirectional LSPs is the combination of static bidirectional MPLS-TP and dynamic MPLS-TE. Flex LSP provides bidirectional label switched paths (LSPs) set up dynamically through Resource Reservation Protocol–Traffic Engineering (RSVP-TE). It does not support non-co routed LSPs.

Flex Label Switched Paths are LSP instances where the forward and the reverse direction paths are setup, monitored and protected independently and associated together during signaling. You use a RSVP Association object to bind the two forward and reverse LSPs together to form a co-routed associated bidirectional TE tunnel.

You can associate a protecting MPLS-TE tunnel with either a working MPLS-TE LSP, protecting MPLS-TE LSP, or both. The working LSP is the primary LSP backed up by the protecting LSP. When a working LSP goes down, the protecting LSP is automatically activated. You can configure a MPLS-TE tunnel to operate without protection as well.

For example, the '821 Accused Products, including Cisco NCS 4000, are MPLS networking platforms, *i.e.*, systems for selecting entities within an MPLS network.


73. The '821 Accused Products, including Cisco NCS 4000, comprise a data structure comprising a plurality of transport entity descriptors and an entity protection switch configured to switch between a working entity and a protection entity. *See, e.g., id.* at 199 (Configuration Guide for Cisco NCS 4000 Series includes configurations using IOS XR); *see also id.* at 344:

Associated Bidirectional Co-routed LSPs

This section provides an overview of associated bidirectional co-routed LSPs. Establishment of MPLS TE-LSP involves computation of a path between a head-end node to a tail-end node, signaling along the path, and modification of intermediate nodes along the path. The signaling process ensures bandwidth reservation (if signaled bandwidth is lesser than 0 and programming of forwarding entries).

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.

Associated Bidirectional Co-routed LSPs: A co-routed bidirectional TE LSP denotes a bidirectional tunnel where the forward direction LSP and reverse direction LSP must follow the same path, for example, the same nodes and paths. Here is an illustration.



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 green) 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 green working LSP to Router 1.
- Router 3 sets up protecting LSP to Router 1 (in red) after performing bidirectional CSPF and sends reverse ERO to Router 1. Node Router 1 uses the received reverse ERO to set up reverse red protecting LSP to Router 3.

For example, the '821 Accused Products, including Cisco NCS 4000, include label-switched paths (“LSP’s”) that employ constrained shortest-path first (“CSPF”) protocols that include “working co-routed LSP pairs” and “protecting co-routed LSP pairs,” *i.e.*, they comprise a data structure comprising a plurality of transport entity descriptors.

74. The '821 Accused Products, including Cisco NCS 4000, comprise digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity. *See, e.g., id.* at 280, 346, 397:

OCH Mutual Circuit Diversity

The OCH Mutual Circuit Diversity feature is an interoperability feature between a NCS 4000 series router and a NCS 2000 series router.

This feature enables the user to create two separate circuits whose paths use a different set of nodes.

Consider a DWDM circuit carrying a service. In order to provide protection and reduce the probability of simultaneous connection failures, the user can create a new circuit by defining a different set of nodes. In case of failure, the service is seamlessly carried forward by the other circuit, which has a different path. Typically, nodes dynamically choose the shortest path, where a circuit is created to reach the destination using minimum number of hops. This might result in network congestion if the same nodes are used by many circuits. Mutual circuit diversity enables the user to allocate different network paths for two circuits. Both the circuits are defined in such a way that there are no overlapping nodes (except the source node), and the paths are independent of each other.

- **SRLG-Aware Path Protection** : This feature specifies that a protecting LSP should be SRLG-diverse from the primary LSP. The user can also specify node-diversity.

```
RP/0/# configure
RP/0/(config)#interface tunnel-te 100
RP/0/(config-if)#path-protection srlg-diverse
RP/0/(config-if)#
```

Bidirectional Forwarding Detection

Bidirectional forwarding detection (BFD) provides low-overhead, short-duration detection of failures in the path between adjacent forwarding engines. BFD allows a single mechanism to be used for failure detection over any media and at any protocol layer, with a wide range of detection times and overhead. The fast detection of failures provides immediate reaction to failure in the event of a failed link or neighbor.

For example, the '821 Accused Products, including Cisco NCS 4000, detect failures in the paths between nodes, *i.e.*, they comprise digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity.


75. The '821 Accused Products, including Cisco NCS 4000, comprise logic configured to determine an entity cost of said plurality of transport entity descriptors. *See, e.g., id.* at 344, 362:

Associated Bidirectional Co-routed LSPs

This section provides an overview of associated bidirectional co-routed LSPs. Establishment of MPLS TE-LSP involves computation of a path between a head-end node to a tail-end node, signaling along the path, and modification of intermediate nodes along the path. The signaling process ensures bandwidth reservation (if signaled bandwidth is lesser than 0 and programming of forwarding entries).

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.

Associated Bidirectional Co-routed LSPs: A co-routed bidirectional TE LSP denotes a bidirectional tunnel where the forward direction LSP and reverse direction LSP must follow the same path, for example, the same nodes and paths. Here is an illustration.



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 green) 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 green working LSP to Router 1.
- Router 3 sets up protecting LSP to Router 1 (in red) after performing bidirectional CSPF and sends reverse ERO to Router 1. Node Router 1 uses the received reverse ERO to set up reverse red protecting LSP to Router 3.

Multiprotocol Label Switching Traffic Engineering

The MPLS TE feature enables an MPLS backbone to replicate and expand the traffic engineering capabilities of Layer 2 ATM and Frame Relay networks. MPLS is an integration of Layer 2 and Layer 3 technologies.

For IS-IS, MPLS TE automatically establishes and maintains MPLS TE label-switched paths across the backbone by using Resource Reservation Protocol (RSVP). The route that a label-switched path uses is determined by the label-switched paths resource requirements and network resources, such as bandwidth. Available resources are flooded by using special IS-IS TLV extensions in the IS-IS. The label-switched paths are explicit routes and are referred to as traffic engineering (TE) tunnels.

For example, the '821 Accused Products, including Cisco NCS 4000, determine entity costs of the entities, such as traffic engineering ("TE") and bandwidth data, *i.e.*, they comprise logic configured to determine an entity cost of said plurality of transport entity descriptors.

76. The '821 Accused Products, including Cisco NCS 4000, comprise logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. *See, e.g., id.* at 428, 590-91:

Multicast-Intact Support for OSPF

The multicast-intact feature provides the ability to run multicast routing (PIM) when IGP shortcuts are configured and active on the router. Both OSPFv2 and IS-IS support the multicast-intact feature.

You can enable multicast-intact in the IGP when multicast routing protocols (PIM) are configured and IGP shortcuts are configured on the router. IGP shortcuts are MPLS tunnels that are exposed to IGP. The IGP routes IP traffic over these tunnels to destinations that are downstream from the egress router of the tunnel (from an SPF perspective). PIM cannot use IGP shortcuts for propagating PIM joins, because reverse path forwarding (RPF) cannot work across a unidirectional tunnel.

When you enable multicast-intact on an IGP, the IGP publishes a parallel or alternate set of equal-cost next hops for use by PIM. These next hops are called *mcast-intact* next hops. The mcast-intact next hops have the following attributes:

- They are guaranteed not to contain any IGP shortcuts.
- They are not used for unicast routing but are used only by PIM to look up an IPv4 next-hop to a PIM source.
- They are not published to the FIB.
- When multicast-intact is enabled on an IGP, all IPv4 destinations that were learned through link-state advertisements are published with a set equal-cost mcast-intact next hops to the RIB. This attribute applies even when the native next hops have no IGP shortcuts.

In OSPF, the max-paths (number of equal-cost next hops) limit is applied separately to the native and mcast-intact next hops. The number of equal cost mcast-intact next hops is the same as that configured for the native next hops.

MPLS-TE Automatic Bandwidth Overview

MPLS-TE automatic bandwidth is configured on individual Label Switched Paths (LSPs) at every head-end. MPLS-TE monitors the traffic rate on a tunnel interface. Periodically, MPLS-TE resizes the bandwidth on the tunnel interface to align it closely with the traffic in the tunnel. MPLS-TE automatic bandwidth can perform these functions:

- Monitors periodic polling of the tunnel output rate
- Resizes the tunnel bandwidth by adjusting the highest rate observed during a given period

For every traffic-engineered tunnel that is configured for an automatic bandwidth, the average output rate is sampled, based on various configurable parameters. Then, the tunnel bandwidth is readjusted automatically based upon either the largest average output rate that was noticed during a certain interval, or a configured maximum bandwidth value.

When reoptimizing the LSP with the new bandwidth, a new path request is generated. If the new bandwidth is not available, the last good LSP continues to be used. This way, the network experiences no traffic interruptions.

If minimum or maximum bandwidth values are configured for a tunnel, the bandwidth, which the automatic bandwidth signals, stays within these values.

For example, the '821 Accused Products, including Cisco NCS 4000, resizes, readjusts, and reoptimizes LSPs and calculates “next hops” when necessary to align the LSP with network traffic,

i.e., they comprise logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event.

77. The '821 Accused Products, including Cisco NCS 4000, comprise said reselection event being selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See, e.g., id.* For example, the '821 Accused Products, including Cisco NCS 4000, resizes, readjusts, and reoptimizes LSPs and calculates “next hops” when necessary to align the LSP with network traffic, including when an operational status change or overall cost change occurs, *i.e.*, said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

78. With knowledge of the '821 Patent, Cisco has actively induced and continues to induce the direct infringement of one or more claims of the '821 Patent, including claim 14, in violation of 35 U.S.C. § 271(b) by its customers and/or end users of their products, including at least the '821 Accused Products, by selling products with a particular design, providing support for, providing instructions for use of, and/or otherwise encouraging its customers and/or end-users to directly infringe, either literally and/or under the doctrine of equivalents, one or more claims of the '821 Patent, including claim 14, with the intent to encourage those customers and/or end-users to infringe the '821 Patent.

79. By way of example, Cisco actively induces infringement of the '821 Patent by encouraging, instructing, and aiding one or more persons in the United States, including but not

limited to customers and end users who purchase, test, operate, and use Cisco's products, including at least the '821 Accused Products, to make, use, sell, and/or offer to sell Cisco's products, including at least the '821 Accused Products, in a manner that infringes at least one claim of the '821 Patent, including claim 14.

80. As a result of Cisco's inducement of infringement, its customers and/or end users made, used, sold, offered for sale, or imported, and continue to make, use, sell, offer to sell, or import Cisco's products, including the '821 Accused Products, in ways that directly infringe one or more claims of the '821 Patent, including claim 14, such as in the manner described above with respect to the Cisco NCS 4000. Cisco had knowledge of its customers' and/or end users' direct infringement at least by virtue of its sales, instruction, and/or promotion of Cisco's products, including the '821 Accused Products, at least as of March 2017 when Orkit IP initiated discussions with Cisco about its patent portfolio, including the Asserted Patents, and no later than the filing of the Original Complaint.

81. Cisco has also contributed to and continues to contribute to the infringement by others, including its customers and/or the end users of its products, of at least claim 14 of the '821 Patent under 35 U.S.C. § 271(c) by, among other things, selling, offering for sale within the United States and/or importing into the United States or otherwise making available the '821 Accused Products for use in practicing the patented inventions of the '821 Patent, knowing that the '821 Accused Products are especially made or adapted for use in infringement of the '821 Patent, are used in practicing the method and process claims of the '821 Patent, embody a material part of the inventions claimed in the '821 Patent, and are not staple articles of commerce suitable for substantial non-infringing use. Cisco's customers and/or the end users of the '821 Accused Products directly infringe the '821 Patent by using the '821 Accused Products.

82. With knowledge of the '821 Patent, Cisco has willfully, deliberately, and intentionally infringed the '821 Patent, and continues to willfully, deliberately, and intentionally infringe the '821 Patent. Cisco had actual knowledge of the '821 Patent and Cisco's infringement of the '821 Patent as set forth above. After acquiring that knowledge, Cisco directly and indirectly infringed the '821 Patent as set forth above. Cisco knew or should have known that its conduct amounted to infringement of the '821 Patent at least because Orckit IP notified Cisco of the '821 Patent and its infringement of the '821 Patent as set forth above.

83. Cisco will continue to infringe the '821 Patent unless and until it is enjoined by this Court. Cisco, by way of its infringing activities, has caused and continues to cause Orckit to suffer damages in an amount to be determined, and has caused and is causing Orckit irreparable harm. Orckit has no adequate remedy at law against Cisco's acts of infringement and, unless it is enjoined from its infringement of the '821 Patent, Orckit will continue to suffer irreparable harm.

84. Orckit is entitled to recover from Cisco damages at least in an amount adequate to compensate for its infringement of the '821 Patent, which amount has yet to be determined, together with interest and costs determined by the Court.

85. Orckit has complied with the requirements of 35 U.S.C. § 287 with respect to the '821 Patent.

COUNT THREE: INFRINGEMENT OF U.S. PATENT 10,652,111

86. Cisco directly infringes at least claim 1 of the '111 Patent by using the Accused Products, which include but are not limited to the products set forth in Appendix C ("the '111 Accused Products"), in a manner that meets every limitation, either literally or under the doctrine of equivalents, of at least claim 1 of the '111 Patent, in violation of 35 U.S.C. § 271(a). For

example, Cisco directly infringes at least claim 1 of the '111 Patent, including by its own use of the '111 Accused Products in the infringing manner set forth below.

87. The '111 Accused Products are designed and operate in such manner that Cisco's customers and/or end users of the Accused Products directly infringe every element of at least claim 1 of the '111 Patent when they follow the instructions described in various materials with which Cisco induces its users to use the Accused Products. Induced by Cisco's sale of the '111 Accused Products, its promotion and advertising of them for their intended infringing use, its instructions on their use in the infringing manner, and other inducing activities, Cisco's customers and/or the end users of the Accused Products directly infringe through that use at least claim 1 of the '111 Patent by using the '111 Accused Products in a manner that practices every element of at least claim 1 of the '111 Patent.

88. For example, Cisco induces its customers and/or end users of its products to use the '111 Accused Products, including the Cisco ASR 1000 Series Aggregation Services Router ("Cisco ASR 1000"), which is exemplary of all of the '111 Accused Products, to practice a method for use with a packet network including a network node for transporting packets between first and second entities under control of a controller that is external to the network node. *See, e.g.*, "Cisco ASR 1000 Series Aggregation Services Routers" Data Sheet (available at <https://www.cisco.com/c/en/us/products/collateral/routers/asr-1000-series-aggregation-services-routers/datasheet-c78-731632.pdf>) at 22:

Support for Cisco Software-Defined WAN

The ASR 1000 series is optimized for Cisco Software Defined WAN (SD-WAN). For enterprises, this means that business critical applications run faster, with more reliability and reduced Operational Expenditure (OpEx). Cisco SD-WAN achieves this by making all branches and Data Centers have the ability to monitor, control, move and report on streams of application data such as specific web (HTTP) traffic for example. The ASR 1000 series has deep packet inspection capability and can accurately identify and control thousands of different applications including custom in-house enterprise applications.

The entire SD-WAN implementation on the ASR 1000 is implemented by managing the end device either from the Cloud or On-Premise through ascending levels of throughput based licenses. All licenses that support Cisco SD-WAN, whether On-Premise or on Cloud are all enabled using Subscription Licenses. These subscription licenses enable all customers to seamlessly transition between On-Premise and Cloud management as needed. The license tiers are structured to support the growth in business needs through simple subscriptions that help simplify the journey to intent-based networking for the WAN.

Cisco SD-WAN subscriptions are aligned across three subscription licenses of **Cisco DNA Essentials**, **Cisco DNA Advantage** and **Cisco DNA Premier**, each expanding functionally. The **Cisco DNA Essentials on ISR 1000 and ISR 4000** covers all types of connectivity and router life cycle management, support for Network and application visibility coupled with basic premise and transport security. ASR 1000 series support two Cisco DNA tiers, Cisco DNA Advantage and Cisco DNA Premier. The **Cisco DNA Advantage** provides for Advanced WAN topologies, Application aware policies supported by enhanced network security. The **Cisco DNA Premier** provides for Cloud connectivity with unlimited segmentation, Advanced Application optimization and Network Analytics, secured by advanced threat protection.


See also “Cisco Catalyst SD-WAN Getting Started Guide” (available at <https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book.pdf>) at

11:

Step 4: Influence Reachability through Centralized Policy

Policy configured on a centralized controller strongly influences how prefixes are advertised among the routers. For example, if all traffic between routers P3 and P4 in the figure here has to make a U-turn at router vEdge-1, the network administrator can apply a simple route policy on the centralized controller. The controller then passes the policy to the affected edge routers. The network administrator does not need to provision the policy on each individual router.

Figure 4: Policy Configured on a Centralized Controller



For example, the '111 Accused Products, including the Cisco ASR 1000, employ a Controller to control a number of entities that communicate data packets over a network, *i.e.*, they are used by

an end user to perform method for use with a packet network including a network node for transporting packets between first and second entities under control of a controller that is external to the network node.


89. Cisco induces its customers and/or end users of its products to use the '111 Accused Products, including the Cisco ASR 1000, in such manner as to (i) send, by the controller to the network node over the packet network, an instruction and a packet-applicable criterion, (ii) receive, by the network node from the controller, the instruction and the criterion, and (iii) receive, by the network node from the first entity over the packet network, a packet addressed to the second entity.

See, e.g., id. at 11-12:

Step 4: Influence Reachability through Centralized Policy

Policy configured on a centralized controller strongly influences how prefixes are advertised among the routers. For example, if all traffic between routers P3 and P4 in the figure here has to make a U-turn at router vEdge-1, the network administrator can apply a simple route policy on the centralized controller. The controller then passes the policy to the affected edge routers. The network administrator does not need to provision the policy on each individual router.

Figure 4: Policy Configured on a Centralized Controller



This approach has many benefits:

- The controller centrally influences access control, that is, which prefixes are allowed to talk to each other inside a VPN.

- The controller optimizes user experience by influencing transport link choice based on SLA or other attributes. The network administrator can color transport links (such as gold and bronze), and allow applications to map the colors to appropriate transport links.
- The network administrator can map business logic from a single centralized point.
- The network can react faster to planned and unexpected situations, such as routing all traffic from high-risk countries through an intermediate point.
- The network can centralize services such as firewalls, IDPs, and IDSs. Instead of distributing these services throughout the network at every branch and campus, the network administrator can centralize these functions, achieving efficiencies of scale and minimizing the number of touch points for provisioning.

See also “Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE Catalyst

SD-WAN Release 17.x” (available at

<https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/ios-xe-17/policies-book-xe.pdf>) at 27-28:

Configure and Execute Cisco SD-WAN Controller Policies

All Cisco SD-WAN Controller policies are configured on the Cisco IOS XE Catalyst SD-WAN devices, using a combination of policy definition and lists. All Cisco SD-WAN Controller policies are also applied on the Cisco IOS XE Catalyst SD-WAN devices, with a combination of apply-policy and lists. However, where the actual Cisco SD-WAN Controller policy executes depends on the type of policy, as shown in this figure:

Figure 11: Cisco SD-WAN Controller Policy

| | Action | App-route Policy | Cflowd Template | Control Policy | Data Policy | VPN Membership Policy |
|------------|-----------|------------------|-----------------|----------------|-------------|-----------------------|
| Controller | Configure | • | • | • | • | • |
| | Apply | • | • | • | • | • |
| | Execute | | | • | | • |
| Device | Configure | | | | | |
| | Apply | | | | | |
| | Execute | • | • | | • | |

For control policy and VPN membership policy, the entire policy configuration remains on the Cisco SD-WAN Controller, and the actions taken as a result of routes or VPNs that match a policy are performed on the Cisco SD-WAN Controller.

For the other three policy types—application-aware routing, cflowd templates, and data policy—the policies are transmitted in OMP updates to the Cisco IOS XE Catalyst SD-WAN devices, and any actions taken as a result of the policies are performed on the devices.

See also “Cisco Catalyst SD-WAN Getting Started Guide” (available at

<https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book.pdf>) at 16.

(“The Cisco SD-WAN Controller maintains a centralized route table that stores the route information, called OMP routes, that it learns from the edge routers and from any other Cisco SD-WAN Controllers in the Cisco Catalyst SD-WAN overlay network. Based on the configured

policy, the Cisco SD-WAN Controller shares this route information with the Cisco edge network devices in the network so that they can communicate with each other.”). For example, the ’111 Accused Products, including the Cisco ASR 1000, execute “policies” that constitute the claimed instruction and packet-applicable criteria and send them by the controller to the network node, *i.e.*, they are used by an end user for (i) sending by the controller to the network node over the packet network, an instruction and a packet-applicable criterion, (ii) receiving, by the network node from the controller, the instruction and the criterion; and (iii) receiving, by the network node from the first entity over the packet network, a packet addressed to the second entity.


90. Cisco induces its customers and/or end users of its products to use the ’111 Accused Products, including the Cisco ASR 1000, in such manner as to check, by the network node, if the packet satisfies the criterion. *See, e.g.*, “Cisco SD-WAN Policies Configuration Guide, Cisco IOS XE Release 17.x” (available at <https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/ios-xe-17/policies-book-xe.pdf>) at 133 (“When data traffic matches the conditions in the match portion of a centralized data policy, the packet can be accepted or dropped, and it can be counted.”); *see also, e.g., id.* at 14, 24-26:

Configure Centralized Policy Based on Prefixes and IP Headers

A centralized data policy based on source and destination prefixes and on headers in IP packets consists of a series of numbered (ordered) sequences of match-action pair that are evaluated in order, from lowest sequence number to highest sequence number. When a packet matches one of the match conditions, the associated action is taken and policy evaluation on that packets stops. Keep this in mind as you design your policies to ensure that the desired actions are taken on the items subject to policy.

Data Policy

Data policy influences the flow of data traffic traversing the network based either on fields in the IP header of packets or the router interface on which the traffic is being transmitted or received. Data traffic travels over the IPsec connections between Cisco IOS XE Catalyst SD-WAN devices, shown in purple in the adjacent figure.



The Cisco IOS XE Catalyst SD-WAN architecture implements two types of data policy:

- Centralized data policy controls the flow of data traffic based on the source and destination addresses and ports and DSCP fields in the packet's IP header (referred to as a 5-tuple), and based on network segmentation and VPN membership. These types of data policy are provisioned centrally, on the Cisco SD-WAN Controller, and they affect traffic flow across the entire network.

Data policy examines fields in the headers of data packets, looking at the source and destination addresses and ports, and the protocol and DSCP values, and for matching packets, it can modify the next hop in a variety of ways or apply a policer to the packets. Data policy is configured and applied on the Cisco SD-WAN Controller, and then it is carried in OMP updates to the Cisco IOS XE Catalyst SD-WAN devices in the site-list that the policy is applied to. The match operation and any resultant actions are performed on the devices as it transmits or receives data traffic.

In the Data Policy Topology figure, a data policy named “change_next_hop” is applied to a list of sites that includes Site 3. The OMP update that the Cisco SD-WAN Controller sends to the devices at Site 3 includes this policy definition. When the device sends or receives data traffic that matches the policy, it changes the next hop to the specified TLOC. Non-matching traffic is forwarded to the original next-hop TLOC.

For example, the '111 Accused Products, including the Cisco ASR 1000, examines data packets pursuant to the “policies,” *i.e.*, they are used by an end user for checking, by the network node, if the packet satisfies the criterion.

91. Cisco induces its customers and/or the end users of its products to use the '111 Accused Products, including the Cisco ASR 1000, such that responsive to the packet not satisfying the criterion, send, by the network node over the packet network, the packet to the second entity. *See, e.g., id.* at 135 (“If a data packet being evaluated does not match any of the match conditions in a data policy, a default action is applied to the packet.”); *see also., id.* at 148-149, 26:

Configure Application-Aware Routing

Table 30: Feature History

| Feature Name | Release Information | Description |
|------------------------------------|--|--|
| Application-Aware Routing for IPv6 | Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco vManage Release 20.9.1 | This feature enables you to configure application-aware routing (AAR) policies to operate with IPv6 application traffic. |

This topic provides general procedures for configuring application-aware routing. Application-aware routing policy affects only traffic that is flowing from the service side (the local/WAN side) to the tunnel (WAN) side of the Cisco IOS XE Catalyst SD-WAN device.

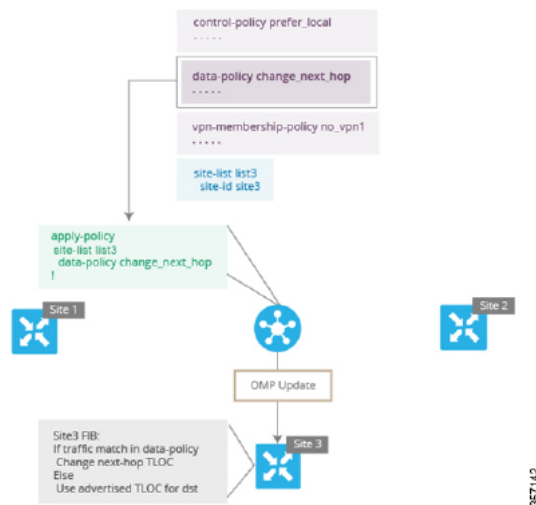
An application-aware routing policy matches applications with an SLA, that is, with the data plane tunnel performance characteristics that are necessary to transmit the applications' data traffic. The primary purpose of application-aware routing policy is to optimize the path for data traffic being transmitted by Cisco IOS XE Catalyst SD-WAN devices.

An application-aware routing policy is a type of centralized data policy: you configure it on the Cisco SD-WAN Controller, and the controller automatically pushes it to the affected Cisco IOS XE Catalyst SD-WAN devices. As with any policy, an application-aware routing policy consists of a series of numbered (ordered) sequences of match-action pairs that are evaluated in order, from lowest sequence number to highest sequence number. When a data packet matches one of the match conditions, an SLA action is applied to the packet to determine the data plane tunnel to use to transmit the packet. If a packet matches no parameters in any of the policy sequences, and if no SLA class is configured for the default-action, the packet is accepted and forwarded with no consideration of SLA. Because application-aware routing policy accepts nonmatching traffic by default,

it is considered as a positive policy. Other types of policies in the Cisco IOS XE Catalyst SD-WAN software are negative policies, because by default they drop nonmatching traffic.

In the Data Policy Topology figure, a data policy named "change_next_hop" is applied to a list of sites that includes Site 3. The OMP update that the Cisco SD-WAN Controller sends to the devices at Site 3 includes this policy definition. When the device sends or receives data traffic that matches the policy, it changes the next hop to the specified TLOC. Non-matching traffic is forwarded to the original next-hop TLOC.

Figure 9: Data Policy Topology



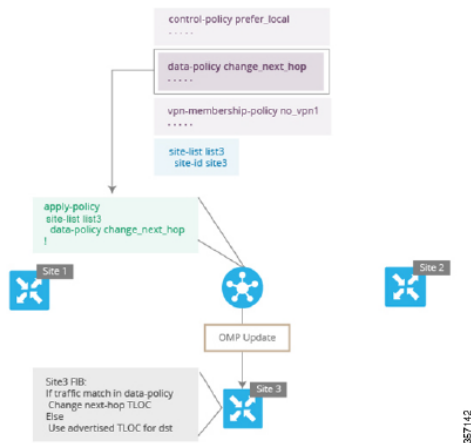
In the `apply-policy` command for a data policy, specify a direction from the perspective of the device. The "all" direction in the figure applies the policy to incoming and outgoing data traffic transiting the tunnel interface. You can limit the span of the policy to only incoming traffic with a `data-policy change_next_hop from-tunnel` command or to only outgoing traffic with a `data-policy change_next_hop from-service` command.

See also “Cisco SD-WAN Policies Configuration Guide, Cisco IOS XE Release 17.x” (available at <https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/policies/ios-xe-17/policies-book-xe.pdf>) at 72 (“Restrict Traffic - This examples illustrates how to disallow certain types of data traffic from being sent from between VPNs. This policy drops data traffic on port 25, which carries SMTP mail traffic, that originates in 209.165.201.0/27. However, the policy accepts all other data traffic, including non-SMTP traffic from 209.165.201.0/27.”). For example, the ’111 Accused Products, including the Cisco ASR 1000, drop or redirect packets that do not satisfy the “policies,” *i.e.*, they are used by an end user for, responsive to the packet not satisfying the criterion, sending, by the network node over the packet network, the packet to the second entity.

92. Cisco induces its customers and/or the end users of its products to use the ’111 Accused Products, including the Cisco ASR 1000, such that responsive to the packet satisfying the criterion, send the packet, by the network node over the packet network, to an entity that is included in the instruction and is other than the second entity. *See, e.g., id.* at 131 (“The Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow provides the ability to look into the packet past the basic header information. The SAIE flow determines the contents of a particular packet, and then either records that information for statistical purposes or performs an action on the packet.”); *see also id.* at 26, 131, 133-34:

In the Data Policy Topology figure, a data policy named "change_next_hop" is applied to a list of sites that includes Site 3. The OMP update that the Cisco SD-WAN Controller sends to the devices at Site 3 includes this policy definition. When the device sends or receives data traffic that matches the policy, it changes the next hop to the specified TLOC. Non-matching traffic is forwarded to the original next-hop TLOC.

Figure 8: Data Policy Topology



In the `apply-policy` command for a data policy, specify a direction from the perspective of the device. The "all" direction in the figure applies the policy to incoming and outgoing data traffic transiting the tunnel interface. You can limit the span of the policy to only incoming traffic with a `data-policy change_next_hop from-tunnel` command or to only outgoing traffic with a `data-policy change_next_hop from-service` command.

Cisco Catalyst SD-WAN Application Intelligence Engine Flow Overview

The Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow provides the ability to look into the packet past the basic header information. The SAIE flow determines the contents of a particular packet, and then either records that information for statistical purposes or performs an action on the packet.



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

Action Parameters for Configuring SD-WAN Application Intelligence Engine Flow

When data traffic matches the conditions in the match portion of a centralized data policy, the packet can be accepted or dropped, and it can be counted. Then, you can associate parameters with accepted packets.

From the Cisco SD-WAN Manager menu, you can configure match parameters from:

- **Configuration > Policies > Centralized Policy > Add Policy > Configure Traffic Rules > (Application-Aware Routing | Traffic Data | Cflowd) > Sequence Type > Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Centralized Policy > Traffic Policy > (Application-Aware Routing | Traffic Data | Cflowd) > Sequence Type > Sequence Rule > Action.**

In the CLI, you configure the action parameters under the `policy data-policy vpn-list sequence action` command.

Each sequence in a centralized data policy can contain one action condition.

In the action, you first specify whether to accept or drop a matching data packet, and whether to count it:

Table 22:

| Description | Cisco SD-WAN Manager | CLI Command | Value or Range |
|---|---|---------------------------|--|
| Accept the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the policy configuration. | Click Accept . | accept | — |
| Count the accepted or dropped packets. | Action Counter Click Accept , then action Counter | count counter-name | Name of a counter. Use the show policy access-lists counters command on the Cisco device. |
| Discard the packet. This is the default action. | Click Drop | drop | — |

For example, the '111 Accused Products, including the Cisco ASR 1000, “accept[]” the packets or direct them to the designated destination if they satisfy the “Policies,” *i.e.*, they are used by an end user for, responsive to the packet satisfying the criterion, sending the packet, by the network node over the packet network, to an entity that is included in the instruction and is other than the second entity.

93. With knowledge of the '111 Patent, Cisco has actively induced and continues to induce the direct infringement of one or more claims of the '111 Patent, including claim 1, in violation of 35 U.S.C. § 271(b) by its customers and/or end users of its products, including at least the '111 Accused Products, by selling products with a particular design, providing support for, providing instructions for use of, and/or otherwise encouraging its customers and/or end-users to directly infringe, either literally and/or under the doctrine of equivalents, one or more claims of the '111 Patent, including claim 1, with the intent to encourage those customers and/or end-users to infringe the '111 Patent.

94. By way of example, Cisco knowingly and actively aided and abetted the direct infringement of the '111 Patent by encouraging, instructing, and aiding one or more persons in the United States, including but not limited to customers and end users who purchase, test, operate, and use Cisco's products, including at least the '111 Accused Products, to use Cisco's products, including at least the '111 Accused Products, in a manner that infringes at least one claim of the '111 Patent, including claim 1.

95. For example, Cisco updates and maintains a website with various materials addressed to end users of its products, including its customers, which instruct its customers on how to use the '111 Accused Products, which are designed in such manner as to infringe at least claim 1 of the '111 Patent when used in the manner shown in such materials. Said materials include, without limitation, quick-start guides, administration guides, user guides, operating instructions, blogs, white papers, data sheets, how-to videos, and other like materials, which cover in depth aspects of how to operate Cisco routers/switches and/or other products, including the '111 Accused Products, and instruct end users how to operate these products in a manner that infringes at least claim 1 of the '111 Patent. *See, e.g.*, “Cisco DNA Software for SD-WAN and Routing” Guide (available at <https://www.cisco.com/c/en/us/products/collateral/software/one-wan-subscription/nb-06-sdwan-migration-quickstart-guide-cte.html>); *see also, e.g.*, “Cisco ASR 1000 Series Aggregation Services Routers” At-a-Glance (available at <https://www.cisco.com/c/en/us/products/collateral/routers/asr-1000-series-aggregation-services-routers/at-a-glance-c45-612993.html#:~:text=Cisco%20%2C%20AE%20ASR%201000%20Series%20Aggregated%20Services%20Routers,application%20performance%20among%20enterprise%20sites%20and%20cloud%20locations>); *see also, e.g.*, “Cisco 4000 Family Integrated Services Router” Data Sheet (available at https://www.cisco.com/c/en/us/products/collateral/routers/4000-series-integrated-services-routers-isr/data_sheet-c78-732542.html).

96. As a result of Cisco’s inducement of infringement, its customers and/or end users used and continue to use Cisco’s products, including the '111 Accused Products, in ways that directly infringe one or more claims of the '111 Patent, including claim 1, such as the ways described above with respect to the Cisco ASR 1000. Cisco had knowledge of its customers’

and/or end users' direct infringement at least by virtue of its design, sales, instruction, and/or otherwise promotion of Cisco's products, including the '111 Accused Products, at least as of March 2017 when Orckit IP initiated discussions with Cisco about its patent portfolio, including the Asserted Patents, and no later than the filing of the Original Complaint.

97. Cisco has also contributed to and continues to contribute to the infringement by others, including its customers and/or the end users of its products, of at least claim 1 of the '111 Patent under 35 U.S.C. § 271(c) by, among other things, selling, offering for sale within the United States and/or importing into the United States or otherwise making available the '111 Accused Products for use in practicing the patented inventions of the '111 Patent, knowing that the '111 Accused Products are especially made or adapted for use in infringement of the '111 Patent, are used in practicing the method and process claims of the '111 Patent, embody a material part of the inventions claimed in the '111 Patent, and are not staple articles of commerce suitable for substantial non-infringing use. Cisco's customers and/or the end users of the '111 Accused Products directly infringe the '111 Patent by using the '111 Accused Products.

98. With knowledge of the '111 Patent, Cisco has willfully, deliberately, and intentionally infringed the '111 Patent, and continues to willfully, deliberately, and intentionally infringe the '111 Patent. Cisco had actual knowledge of the '111 Patent and Cisco's infringement of the '111 Patent as set forth above. After acquiring that knowledge, Cisco directly and indirectly infringed the '111 Patent as set forth above. Cisco knew or should have known that its conduct amounted to infringement of the '111 Patent at least because Orckit IP notified Cisco of the '111 Patent and its infringement of the '111 Patent as set forth above.

99. Cisco will continue to infringe the '111 Patent unless and until it is enjoined by this Court. Cisco, by way of its infringing activities, has caused and continues to cause Orckit to suffer

damages in an amount to be determined, and has caused and is causing Orckit irreparable harm. Orckit has no adequate remedy at law against Cisco's acts of infringement and, unless it is enjoined from its infringement of the '111 Patent, Orckit will continue to suffer irreparable harm.

100. Orckit is entitled to recover from Cisco damages at least in an amount adequate to compensate for its infringement of the '111 Patent, which amount has yet to be determined, together with interest and costs determined by the Court.

101. Orckit has complied with the requirements of 35 U.S.C. § 287 with respect to the '111 Patent.

DEMAND FOR JURY TRIAL

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, Orckit hereby demands a jury trial on all issues triable to a jury.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff respectfully prays for entry of judgment for Orckit and against Cisco and enter the following relief:

a) A judgment that Cisco has infringed (directly and/or indirectly) one or more claims of the Asserted Patents, namely U.S. Patents Nos. 6,680,904 ("the '904 Patent"), 8,830,821 ("the '821 Patent"), and 10,652,111 ("the '111 Patent"), and continues to do so with respect to the '821 and '111 patents.

b) That Orckit recover all damages to which it is entitled under 35 U.S.C. § 284, but in no event less than a reasonable royalty;

c) That Cisco be permanently enjoined from further infringement of the '821 and '111 Asserted Patents;

- d) That Orckit, as the prevailing party, shall recover from Cisco all taxable costs of court;
- e) That Orckit shall recover from Cisco all pre- and post-judgment interest on the damages award, calculated at the highest interest rates allowed by law;
- f) That Orckit shall recover from Cisco an ongoing royalty in an amount to be determined for continued infringement after the date of judgment; and
- g) That Cisco's conduct was willful and that Orckit should therefore recover treble damages, including attorneys' fees, expenses, and costs incurred in this action, and an increase in the damage award pursuant to 35 U.S.C. § 284;
- h) That this case is exceptional and that Orckit shall therefore recover its attorneys' fees and other recoverable expenses, under 35 U.S.C. § 285; and
- i) That Orckit shall recover from Cisco such other and further relief as the Court deems appropriate.

Dated: February 14, 2025

Respectfully submitted,

/s/ Michael Ng by permission Andrea

Fair

Michael Ng

California State Bar No. 237915 (Lead Attorney)

Daniel A. Zaheer

California State Bar No. 237118

Kim A. Kennedy*

California State Bar No. 305499

michael.ng@kobrekim.com

daniel.zaheer@kobrekim.com

kim.kennedy@kobrekim.com

KOBRE & KIM LLP

150 California Street, 19th Floor

San Francisco, CA 94111

Telephone: 415-582-4800
Facsimile: 415-582-4811

George Stamatopoulos
KOBRE & KIM LLP
800 3rd Avenue
New York, NY 10022
Telephone: 415-582-4800
Facsimile: 415-582-4811

Zachary R. Ritz
California State Bar No. 301281
Zachary.ritz@kobrekim.com
KOBRE & KIM LLP
201 S. Biscayne Boulevard, Suite 1900
Miami, FL 33131
Telephone: 415-582-4800
Facsimile: 415-582-4811

Hangcheng (Robert) Zhou
California State Bar No. 320038
Robert.Zhou@kobrekim.com
KOBRE & KIM LLP
43RD Floor, 4302-4304 HKRI Centre One,
HKRI Taikoo Hui
288 Shimen Yi Road
Shanghai, PRC, 200041
Telephone: +86 21-3210-2100

Of Counsel:

Andrea L. Fair (TX Bar No. 24078488)
andrea@millerfairhenry.com
MILLER FAIR HENRY, PLLC
1507 Bill Owens Parkway
Longview, Texas 75604
Telephone: (903) 757-6400
Facsimile: (903) 757-2323

Attorneys for Plaintiff Orckit Corporation

**Application for Admission forthcoming*