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

DUNTI NETWORK TECHNOLOGIES, LLC, ***Plaintiff,*** **v.** **CISCO SYSTEMS, INC.,** ***Defendant.***	**Civil Action No._____** **JURY TRIAL DEMANDED**

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Dunti Network Technologies, LLC ("Dunti"), is the owner and assignee of patents critical to the efficiency, security, and scalability of modern communications networks. In recent years, defendant Cisco Systems, Inc. ("Cisco") has adopted Dunti's patented technologies—developed more than a decade ago right here in Texas—*en masse*. Cisco has profited handsomely from its use of Dunti's patented inventions, and Dunti deserves to be compensated for this use. But Cisco has not paid Dunti its fair share. This lawsuit, which alleges infringement of Dunti's U.S. Patent Nos. 6,587,462 ("the '462 patent"); 6,788,701 ("the '701 patent"); 6,804,235 ("the '235 patent"); 6,643,286 ("the '286 patent"); 7,778,259 ("the '259 patent"); 6,912,196 ("the '196 patent") and 6,754,214 ("the '214 patent") (collectively, "the patents-in-suit"), is brought to ensure that Cisco pays Dunti what it fairly owes.

THE PARTIES

1. Dunti, based in Longview, Texas, is committed to advancing the current state of innovation in the field of secure, optimized data transmission across communication networks. In addition to the ongoing efforts of the lead inventor, Dunti employs a resident of Longview, Texas as a Technology Analyst. Dunti is a Texas limited liability company with its principal place of business at 911 NW Loop 281, Suite 211-44, Longview, TX 75604.

2. Dunti is a small, Texas-based company. Dunti depends on patent protection to effectively license its innovative technologies and build its business. Like Defendant Cisco, Dunti relies on its intellectual property.

3. On information and belief, Defendant Cisco Systems, Inc. is a California corporation with its principal office at 170 West Tasman Drive, San Jose, California 95134. Cisco can be served through its registered agent, Prentice Hall Corporation System, 211 E. 7th Street, Suite 620, Austin, TX 78701-3218.

4. On information and belief, Cisco maintains offices throughout the State of Texas[1] and is registered with the Texas Secretary of State to do business in Texas.

5. On information and belief, according to Cisco's website, Cisco offers infringing products for sale throughout the United States and Canada, including in the Eastern District of Texas. Further, Cisco advertises its infringing products throughout the Eastern District of Texas and claims financial benefits through its conducting of business in Texas.

JURISDICTION AND VENUE

6. This action arises under the patent laws of the United States, Title 35 of the United States Code. Accordingly, this Court has exclusive subject matter jurisdiction over this action under 28 U.S.C. §§ 1331 and 1338(a).

[1] *See* Cisco's "US Sales Offices" Page, http://www.cisco.com/web/siteassets/contacts/offices/us.html (indicating that Cisco has offices in Houston, San Antonio, Austin, Richardson, and Irving, Texas) (accessed Sept. 6, 2016).

7. Upon information and belief, this Court has personal jurisdiction over Cisco in this action because Cisco has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Cisco would not offend traditional notions of fair play and substantial justice. Defendant Cisco, directly and/or through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the patents-in-suit. Moreover, Cisco is registered to do business in the State of Texas, and has appointed Prentice Hall Corporation System, 211 E. 7th Street, Suite 620, Austin, TX 78701-3218, as its agent for service of process.

8. Venue is proper in this district under 28 U.S.C. §§ 1391(b)-(d) and 1400(b). Defendant Cisco is registered to do business in Texas, has several offices in Texas, and, upon information and belief, has transacted business in the Eastern District of Texas and has committed acts of direct and indirect infringement in the Eastern District of Texas.

DUNTI'S LANDMARK NETWORK COMMUNICATION SYSTEMS

9. Dunti is the owner and assignee of ten patents on pioneering network technologies, including the seven patents-in-suit (collectively, "the Dunti patents").

10. Electrical engineer and entrepreneur Rupaka Mahalingaiah is a named inventor on each of the Dunti Patents and the founder of Dunti Corp. and Dunti LLC. For more than 30 years, Rupaka has worked at the cutting edge of computing and networking technologies.

11. Even today, female engineers are rare in the American workforce, comprising just over ten percent of all engineers in recent government surveys.[2] When Rupaka began her career in the 1980s, female engineers were rarer still—and *foreign-born, female, computer* engineers were almost inconceivable. Yet through many years of hard work, creativity, and innovation,

[2] According to the Bureau of Labor Statistics Current Population Survey, women comprised just 10.3% of American engineers in 2003, and 11.7% in 2011. *See, e.g.*, http://www.nsf.gov/statistics/wmpd/2013/pdf/tab9-2_updated_2013_11.pdf (accessed Sept. 6, 2016).

Rupaka did more than just defy the odds (and overcome large-scale industry pushback and skepticism)—she became an American engineering success story by any measure.

12. After earning a Bachelor's Degree in Electronic Engineering from Bangalore University and a Master's Degree in Electrical Engineering from Virginia Tech, Rupaka began working at Concurrent Computer Corporation, a company that specialized in multi-processing systems used for real-time computing (i.e., computer systems that are subject to strict time constraints and must respond to inputs within milliseconds). While real-time computing performance is common today, real-time systems were state of the art at that time.

13. After several years at Concurrent, Rupaka joined Teradata, a hardware/software company built around research conducted at the California Institute of Technology (Caltech) specializing in database and parallel processor computing. At Teradata, Rupaka was responsible for architecting a next-generation, database supercomputer.

14. After briefly working at a networking startup in Austin, Rupaka joined Advanced Micro Devices ("AMD"), where she was one of the lead architects on K7/K7+, which became AMD's wildly successful Athlon processor. The original Athlon processor was the first desktop processor to reach speeds of one gigahertz. The Athlon processor's revolutionary architecture and design made these unprecedented speeds possible by allowing the processor to achieve substantially higher clocking speeds and to keep the processing pipeline full. The result was a faster, more efficient chip design.

15. Although she was only at AMD for three years, her contributions during that time were enduring, helping to generate billions of dollars in revenue and resulting in over 30 patents.[3] Her innovations at AMD have inspired others and been cited by nearly one-thousand United States patents and published patent applications as prior art before the United States Patent and Trademark Office, including by:

- International Business Machines Corporation;
- Oracle Corporation;

[3] In total, Rupaka is a named inventor on nearly 50 issued U.S. patents.

- Fujitsu Ltd.;
- Sun Microsystems, Inc.;
- Intel Corporation;
- Qualcomm Inc.;
- Cisco Technology, Inc.;
- Texas Instruments Inc.;
- ARM Holdings, PLC;
- Samsung Electronics Co. Ltd.;
- Freescale Semiconductor, Inc.;
- SK Hynix, Inc.;
- Rambus, Inc.;
- Hitachi, Ltd.; and
- Apple, Inc.

16. Rupaka left AMD in 1997 to become an entrepreneur, shifting her focus from architecting fast, efficient processors to architecting fast, efficient networks. She recognized the inefficiencies, lack of fault tolerance, and security vulnerabilities in then-state-of-the-art network designs, so she set out to solve the separate but related problems of (1) network inefficiency and (2) the lack of network security. It was at this time that Rupaka began to develop the technologies that would be the foundation of Dunti's next-generation networking systems.

17. In early 1999, Rupaka and Viren Kapadia began working together to perfect and expand on her network security and efficiency innovations.

18. Combining Rupaka's expertise in processor design and Viren's expertise in network communications, they created a new holistic network architecture that solved many of the problems inherent to computer networks of that time and that would become widely used in modern data centers. This new architecture combined efficient addressing schemes with built-in security and priority mechanisms to allow for faster, more efficient, and more secure networks that were backwards compatible with the networks of the time.

19. Recognizing the importance of what they had developed, Rupaka set out to build and commercialize this new network architecture, hiring a team of engineers to create several operational prototypes of the Dunti network module—the Dunti Trupta.[4]

[4] "Trupta" means "complete" in Sanskrit.

20. With the working module prototypes in hand, Rupaka hired PricewaterhouseCoopers ("PWC") to audit the Dunti Trupta system and design. PWC engineers used the prototypes to set up a metropolitan area network and spent days running tests on the Dunti Trupta module prototypes and the network to verify their designs. At the end of the audit, PWC provided an audit report verifying the viability of the new network architectures and the modules for implementing those architectures.

21. Unfortunately, Rupaka set out to fund her technical innovations at the worst possible time—at the height of the dot-com and telecom crashes in late 2000 and early 2001. With venture capital all but extinct marketwide, Rupaka was unable to widely commercialize her Dunti inventions in this period.

22. But Rupaka's groundbreaking innovations in network architecture and module design did not go unnoticed, gaining the attention of the Department of Defense, the Department of Energy, and the Department of Homeland Security—all of which awarded her Small Business Innovation Research ("SBIR") grants to develop other computing and networking technologies. In addition, in 2005, the Department of Defense asked Rupaka to present her technological innovations to the Defense Advanced Research Projects Agency ("DARPA") to further the agency's mission—to transform revolutionary concepts and even seeming impossibilities into practical capabilities.

23. The Dunti patents and applications have been cited by 418 United States patents and published patent applications as prior art before the United States Patent and Trademark Office. Companies whose patents cite the Dunti patents include:

- Avaya, Inc.;
- Hitachi Ltd.;
- Advanced Micro Devices, Inc.;
- Microsoft Corporation;
- Hewlett Packard Enterprise Development LP.;
- Cisco Technology, Inc.;
- F5 Networks, Inc.;
- AT&T Corporation;

- CA, Inc.;
- Brocade Communication Systems, Inc.;
- Intel Corporation;
- International Business Machines Corporation;
- Alcatel Lucent S.A.;
- Apple, Inc.;
- Marvell International, Ltd.;
- ZTE Corporation;
- Broadcom Corporation;
- Vodafone Group PLC;
- Nokia Corporation;
- NEC Corporation;
- Terascale Supercomputing, Inc.;
- Siemens AG;
- British Telecommunications PLC;
- Fujitsu, Ltd.;
- Ciena Corporation; and
- Texas Instruments, Inc.

TECHNOLOGY BACKGROUND

24. A communication network is generally regarded as an interconnected set of subnetworks that uses various networking protocols at various networking layers to communicate information—in the form of data packets—across the network. Each networking layer provides some particular functionality using layer-specific networking protocols, such as the well-known IP and Ethernet protocols.

25. For example, the IP protocol is generally considered a layer 3 protocol. The IP protocol uses IP addresses—which are 32-bit addresses—to send and receive data over the internet by delivering packets from a sending (i.e., source) device to a receiving (i.e., destination) device.

26. As another example, the Ethernet protocol is generally considered a layer 2 protocol. The Ethernet protocol uses MAC addresses—which are 48-bit addresses that are unique to every internet-connected device—to send and receive data over the physical network.

27. Data is, therefore, sent from a source device to a destination device using IP addresses at layer 3 and MAC addresses at layer 2. But before that data is sent, the various

networking layers divide the data into packets and wrap the data by placing the packets into datagrams that include additional control information, such as a header containing IP and MAC addresses. Data can be wrapped multiple times before being sent across the network.

28. Links of a network are connected by various hardware components, such as routers and switches.

29. Traditionally, routers operate at layer 3 and direct traffic across the internet by looking at the destination IP address in the IP-addressed packet, determining the best route for the packet, and then sending the packet to the next hop along the path to the destination. To determine the best route for a packet, a router compares the destination address against an internal routing table. Routing tables are dynamic and can accommodate multiple modules having IP addresses that change as the network is reconfigured with new routers, switches, or other network components. Thus, routers can adapt to network conditions by using complex routing algorithms and by updating the routing tables accordingly.

30. Unlike routers, switches traditionally operate at layer 2 and use MAC addresses to forward packets to the next hop without first determining the best route. Switches receive data packets on a particular input port and then send them to a particular output port (or ports). This operation can be quickly repeated each time a packet is received. Because of this, data travels faster through switches than it does through routers.

LIMITATIONS OF THEN-STATE-OF-THE-ART SYSTEMS

31. The next-generation technologies described in the Dunti patents addressed a number of limitations of then-state-of-the-art systems.

32. First, the next-generation technologies described in the Dunti patents addressed problems associated with using a single addressing domain, such as IP addressing, for all internet-connected devices.

33. For example, as explained in the Dunti patents, using a common IP addressing domain for every node in a network made up of hundreds, thousands, or even more sub-networks

can pose several problems. The first major version of IP, called IPv4, uses 32-bit IP addresses; thus, the maximum number of possible IPv4 addresses in the IP addressing domain is approximately 4.3 billion. Given the explosive growth of the Internet and the constantly increasing number of internet-connected devices, the inventors of the Dunti patents recognized that the IPv4 addressing domain would soon become insufficient, and by 2011, this was indeed the case. They also recognized that simply increasing the size of the IP addressing domain (and therefore, the number of available IP addresses) by adding bits to the addressing domain would increase the amount of decoding required and, as a result, the amount of time required for routing.

34. Second, the next-generation technologies described in the Dunti patents addressed problems associated with slow routing-table lookups.

35. For example, a packet can travel through many hops before arriving at its destination, with each hop requiring a complex address-translation operation. As described above, because of the complex routing-table lookups required at each hop to make routing decisions, routing can be a relatively slow process. Switches, on the other hand, are relatively fast, but, unlike routers, they are not able to adapt to changes in traffic conditions.

36. Third, the next-generation technologies described in the Dunti patents addressed problems associated with security and prioritization of data packets as they traverse a network.

37. For example, common network security mechanisms have traditionally included firewalls implemented in hardware and software, and authentication systems implemented in software, such as encryption and passwords. Firewalls, which analyze incoming packets to determine if a packet should be placed on the internal network, add latency at the interface between the external and internal networks and generally operate at a single point in the communication path rather than over the entire communication path. In addition, they can be difficult to configure because each firewall must be updated and configured separately as needs change.

38. Encryption adds overhead to the packet and involves time-consuming decryption at the receiving end. Using passwords takes up less transmission bandwidth than encryption, but passwords can sometimes be broken either because of a user's improper choice of password or through a brute-force attack.

DUNTI'S NEXT-GENERATION NETWORKING SOLUTIONS

39. The next-generation networking technology described in the Dunti patents covers various aspects of networking systems that work together to provide networks that are faster, more efficient, more scalable, and more secure.

40. For example, some of the Dunti patents describe, among other things, using multiple separate and independent addressing domains to overcome the mathematical and practical limitations of the traditional IP packet addressing domains to allow for the transmission of data packets more quickly and efficiently than was possible with any prior art systems. They describe architectures, systems, and methods for transparently mapping addresses across multiple addressing domains, as shown, for example, in the figure below. Because an addressing domain in one network is separate from an addressing domain in another network, a module in the first network and a module in the second network can each have the same identifier, which allows addressing (such as IP addresses) to be reused among networks. These new designs allow for the segmentation of a given network, permitting multiple networks and/or multiple services to share the same infrastructure.

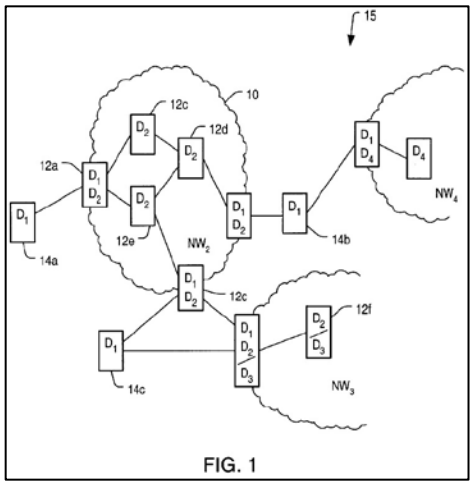

'462 Patent, Fig. 1.

41. As another example, some of the Dunti patents describe using intelligent network infrastructure and hierarchical networks to more efficiently transfer data packets across a network, as shown, for example, in the figure below. By structuring a network and informing each module of its relative location within the network, modules internal to a particular network can operate as switches, quickly forwarding packets towards their final destination. As a result, only modules at the edges of a given network are required to analyze or decode the destination address of the packet.

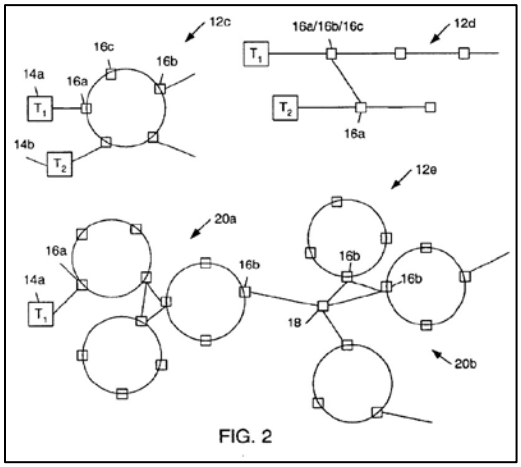

'286 Patent, Fig. 2.

42. As another example, other Dunti patents describe using a packet-based security mechanism that prevents decryption techniques and password attacks. This packet-level security

mechanism also enables hosts that are connected across a public network to be connected to form secure virtual networks by modifying the contents of packets in such a way that only the destined host that is part of the same secure virtual network will be able to restore the contents of the packet when it is received.

43. The continued growth of the number of internet-connected devices and internet-based services, as well as a recent shift toward cloud-based services, has led to wide adoption of Dunti's next-generation networking technology in the industry. For example, Dunti's next-generation networking technology has particular applicability to data-center networking and has been widely implemented by many major networking companies as part of their data center fabric solutions to provide faster, more efficient, more scalable, and more secure data centers. Dunti's next-generation networking technology also applies to the backbone ring networks that connect multiple data center physical locations into a single virtual data center.

CISCO'S INFRINGING PRODUCTS AND SERVICES

44. On information and belief, Cisco offers a high-performance data center networking solution called Unified Fabric, which, as Cisco describes it, is a "holistic network architecture comprising switching, security, and services that is designed for physical, virtual, and cloud environments. It uniquely integrates with servers, storage, and orchestration platforms ***for more efficient operations and greater scalability***." *See* Overview, *Data Center Switches*, CISCO.COM, available at http://www.cisco.com/c/en/us/products/switches/data-center-switches/index.html - ~tab-overview (accessed July 20, 2016) (emphasis added).

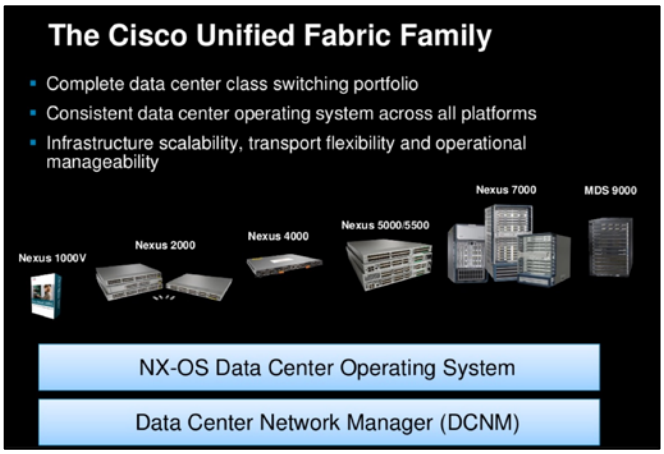

Timur Muminov, *Cisco Unified Fabric*, CISCO PRESENTATION (2010), at 4, available at http://www.slideshare.net/xKinAnx/presentation-cisco-unified-fabric (accessed Aug. 26, 2016).

45. On information and belief, Unified Fabric incorporates multiple Cisco data center technologies, including, for example, Cisco FabricPath. *See* Technologies, *Data Center Switches*, CISCO.COM, available at http://www.cisco.com/c/en/us/products/switches/data-center-switches/index.html - ~tab-technologies (accessed July 20, 2016). Cisco FabricPath technology is, according to Cisco, a "Cisco NX-OS software innovation combining the plug-and-play simplicity of Ethernet with the reliability and scalability of Layer 3 routing. Using FabricPath, you can build highly scalable Layer 2 multipath networks without Spanning Tree Protocol. Such networks are particularly suitable for large virtualization deployments, private clouds, and high-performance computing (HPC) environments." *See* Overview, *Cisco FabricPath*, CISCO.COM, available at http://www.cisco.com/c/en/us/solutions/data-center-virtualization/fabricpath/index.html (accessed July 20, 2016).

46. On information and belief, and according to Cisco, FabricPath is (1) "simple" because it is "transparent to [layer 3] protocols" and does not use MAC addressing tables within the network; (2) "efficient" because it uses an optimal path between any two nodes; and (3) "scalable."

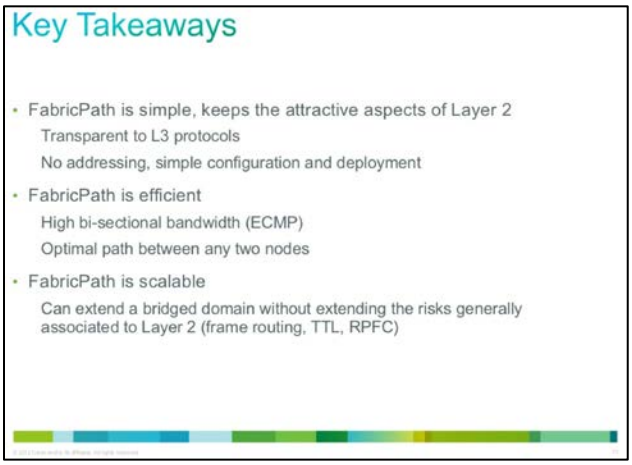

Babi Seal and Patrick Warichet, *Efficient Data Center Design with FabricPath*, CISCO IOS ADVANTAGE WEBINAR PRESENTATION (2012), at 72, available at http://www.slideshare.net/getyourbuildon/fabric-path-webinar (accessed Aug. 26, 2016).

47. On information and belief, FabricPath was recognized by eWEEK Labs as a Top Product of 2010 for building next-generation high-performance networking systems.

eWEEK Labs, *eWEEK Labs Names the Top Products of 2010*, EWEEK LABS PRESENTATION (Dec. 10, 2010), at 4, available at http://www.eweek.com/c/a/IT-Infrastructure/eWEEK-Labs-Names-the-Top-Products-of-2010-203236 (accessed July 20, 2016).

48. And on information and belief, in 2010, NetworkWorld tested FabricPath's "ability to boost bandwidth, reroute around trouble, and simplify network management. In all three areas, FabricPath delivered: Cisco's pre-standard take on the IETF's forthcoming TRILL specification showed real improvement over STP-based designs." *See* David Newman, *Cisco FabricPath Enables Faster, Simpler, Flatter Data Center Networks*, NETWORK WORLD (Oct. 25,

2010), available at http://www.networkworld.com/article/2191595/router/cisco-fabricpath-enables-faster--simpler--flatter-data-center-networks.html (accessed July 20, 2016).

49. On information and belief, Cisco's Unified Fabric solution is implemented using Cisco's Nexus family of products and Cisco's MDS family of products. *See, e.g.*, Products & Services, *Data Center Switches*, CISCO.COM, available at http://www.cisco.com/c/en/us/products/switches/data-center-switches/index.html - ~tab-products-services (accessed July 20, 2016).

50. On information and belief, Cisco's Nexus family of products includes both hardware and software. Such hardware includes at least the Nexus 9000 Series Switches and Nexus 7000 Series Switches. *See, e.g.*, Products & Services, *Data Center Switches*, CISCO.COM, available at http://www.cisco.com/c/en/us/products/switches/data-center-switches/index.html - ~tab-products-services (accessed July 20, 2016). Such software includes the Prime Data Center Network Manager and the NX-OS operating software. *Id.*

51. Similarly, on information and belief, Cisco's MDS family of products includes both hardware and software. Such hardware includes at least the MDS 9700 Series Multilayer Directors, the MDS 9500 Series Multilayer Directors, the MDS 9200 Series Multiservice Switches, and the MDS 9100 Series Multilayer Fabric Switches. *See* Products & Services, *Data Center Switches*, CISCO.COM, available at http://www.cisco.com/c/en/us/products/switches/data-center-switches/index.html - ~tab-products-services (accessed July 20, 2016). Such software includes the MDS 9000 NX-OS and the SAN-OS Software. *Id.*

52. On information and belief, in addition to these products, Cisco offers services called Cisco Nexus Services to assist its customers with "incorporat[ing] the Cisco Nexus switches into your data center environment to create a next-generation data center architecture." *See Cisco Nexus Services Service Overview*, CISCO TECHNICAL DOCUMENTATION (2011), at 1. Cisco Nexus Services include Cisco Data Center Plan and Build Services for Nexus, Cisco SMARTnet Service, and Cisco Data Center Optimization Service for Unified Fabric Nexus.

53. On information and belief, with Cisco Data Center Plan and Build Services for Nexus, Cisco assists its customers with "develop[ing] a more scalable, efficient, and resilient

data center architecture based on the Cisco Nexus hardware platform," which includes activities such as pilot support, Nexus assessment, high- and low-level design reviews, proof of concept and pilot program, migration and implementation plan and support, and remote knowledge transfer sessions. *See Cisco Nexus Services Service Overview*, CISCO TECHNICAL DOCUMENTATION (2011), at 2.

54. On information and belief, with Cisco SMARTnet Service, Cisco provides technical support service, which includes direct, anytime access to Cisco engineers, Cisco.com knowledge base and tools, ongoing operating system software updates and upgrades, Cisco IOS Software support, and proactive diagnostics and real-time alerts. *See Cisco Nexus Services Service Overview*, CISCO TECHNICAL DOCUMENTATION (2011), at 3.

55. On information and belief, with Cisco Data Center Optimization Service for Unified Fabric Nexus, Cisco assists its customers with "analyz[ing], improv[ing], and optimiz[ing] your Cisco Nexus solution while supporting your planned data center transformation initiatives," which includes data center assessment, data center support, and data center learning. *See Cisco Nexus Services Service Overview*, CISCO TECHNICAL DOCUMENTATION (2011), at 3.

56. In addition to Unified Fabric, on information and belief, Cisco also offers a number of products and services that implement Cisco's proprietary Resilient Packet Ring ("RPR") network and/or an IEEE 802.17 RPR network. For example, Cisco offers optical services modules, including, for example, the Cisco ML-Series networking cards and the Cisco ONS 15454 series of networking cards, that implement either Cisco's proprietary RPR network or the IEEE standard RPR network.

57. RPR networks, which transport data traffic over widespread optical fiber rings, can be used to connect data centers that are spread across multiple physical locations or to connect smaller networks to larger and/or backbone networks.

58. RPR networks include dual, counter-rotating rings that are implemented using a series of switches located around the rings where traffic enters and exits the network.

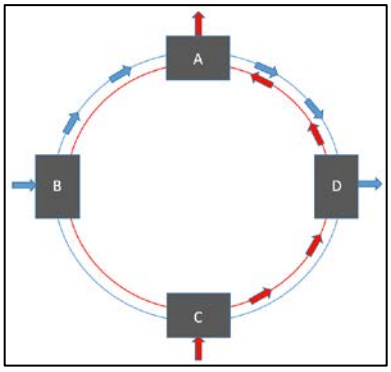

59. The dual-ring topology provides robustness by including the capability of automatic reconfiguration after a link failure. If a node on the ring or a link between two nodes is disrupted or fails, traffic can be looped back around the ring in the opposite direction to the destination node and avoid the disruption/failure.

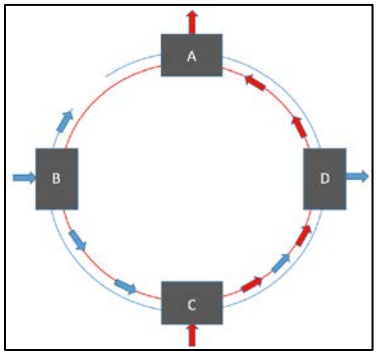

COUNT I
INFRINGEMENT OF U.S. PATENT NO. 6,587,462

60. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

61. U.S. Patent No. 6,587,462 ("the '462 patent"), entitled "Address Mapping Mechanism Enabling Multi-Domain Addressing in Communication Networks, was filed on February 16, 2001. Dunti is the owner by assignment of the '462 patent. A true and correct copy of the '462 patent is attached hereto as Exhibit A. The '462 patent claims a specific architecture, systems, and methods for transparently mapping addresses across multiple addressing domains and/or protocols.

62. The ’462 patent has been cited by at least fifteen United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ’462 patent as relevant prior art:

- Hewlett Packard Enterprise Development LP;
- International Business Machines Corporation;
- Terascale Supercomputing, Inc.;
- NEC Corporation; and
- Microsoft Corporation.

63. The ’462 patent teaches, for example, a networking system with multiple independent addressing domains. Because an addressing domain in a first network is separate from an addressing domain in a second network, the first and second networks need not have a common addressing mechanism in which each module of both the first and second networks requires a unique identification number. Instead, a module in the first network and a module in the second network can each have the same identifier, which allows addressing to be reused among networks.

64. The end modules and termination devices, however, must have a common addressing scheme, in which each end module and termination device has its own unique identifier. Thus, while the end modules and termination devices connected to the end modules have unique and corresponding lower layer addresses, the intermediate modules in the networks can have an independent set of identifiers separate from those of the end modules and termination devices.

65. Set up in this way, sending a data packet from a termination device to another termination device, separated by a network with an internal addressing domain that is different from external addressing domains, uses a simple mapping function. The entry end module adds to the data packet the separate addressing protocols unique to the internal modules, such that the packet includes the IP source and destination addresses, the Ethernet source and destination addresses, and the internal source and destination addresses of the network. The internal

addresses are added when the data packet enters the network and are stripped when the data packet leaves the network.

66. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

67. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch, Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus 9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, “the Cisco Nexus Accused Products”).

68. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, “the Cisco MDS Accused Products”).

69. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

70. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, “the Cisco ’462 Accused Products”).

71. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '462 Accused Products ("a Cisco '462 Accused Product Network").

72. On information and belief, a Cisco '462 Accused Product Network implements at least Cisco's FabricPath technology.

73. On information and belief, a Cisco '462 Accused Product Network comprises a communication system.

74. On information and belief, a Cisco '462 Accused Product Network comprises an entry end module, an exit end module, and at least one intermediate module between the entry end module and the exit end module. For example, the figures below show an entry end module and an exit end module at the edges of a FabricPath network. The second figure below shows multiple switches coupled between two end switches within a FabricPath network.

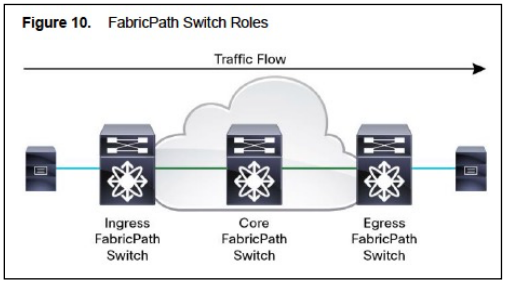

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 12.

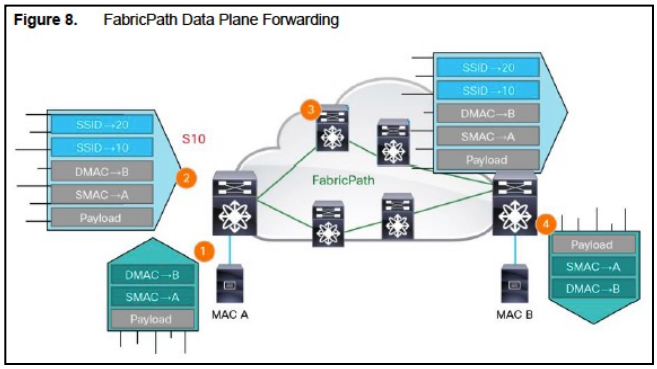

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 9.

75. On information and belief, a Cisco ’462 Accused Product Network comprises a first addressing domain for identifying each of the end modules and the intermediate module. For example, each switch within a FabricPath network is assigned a unique Switch ID.

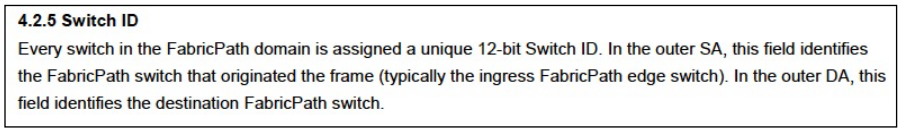

4.2.5 Switch ID

Every switch in the FabricPath domain is assigned a unique 12-bit Switch ID. In the outer SA, this field identifies the FabricPath switch that originated the frame (typically the ingress FabricPath edge switch). In the outer DA, this field identifies the destination FabricPath switch.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

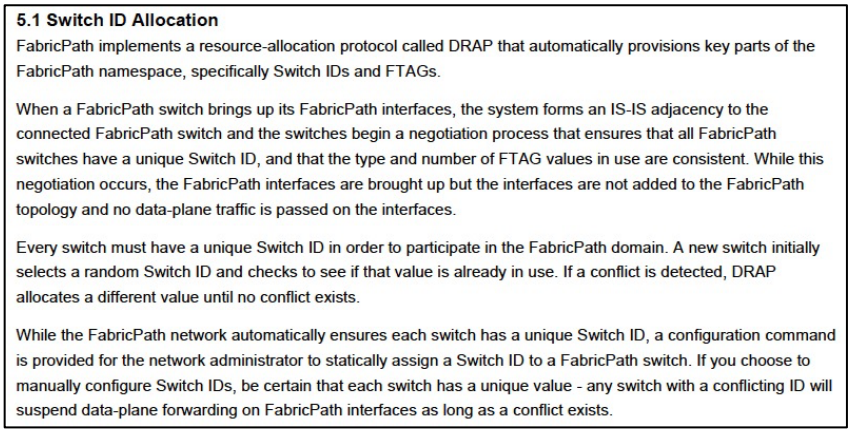

5.1 Switch ID Allocation

FabricPath implements a resource-allocation protocol called DRAP that automatically provisions key parts of the FabricPath namespace, specifically Switch IDs and FTAGs.

When a FabricPath switch brings up its FabricPath interfaces, the system forms an IS-IS adjacency to the connected FabricPath switch and the switches begin a negotiation process that ensures that all FabricPath switches have a unique Switch ID, and that the type and number of FTAG values in use are consistent. While this negotiation occurs, the FabricPath interfaces are brought up but the interfaces are not added to the FabricPath topology and no data-plane traffic is passed on the interfaces.

Every switch must have a unique Switch ID in order to participate in the FabricPath domain. A new switch initially selects a random Switch ID and checks to see if that value is already in use. If a conflict is detected, DRAP allocates a different value until no conflict exists.

While the FabricPath network automatically ensures each switch has a unique Switch ID, a configuration command is provided for the network administrator to statically assign a Switch ID to a FabricPath switch. If you choose to manually configure Switch IDs, be certain that each switch has a unique value - any switch with a conflicting ID will suspend data-plane forwarding on FabricPath interfaces as long as a conflict exists.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 13.

76. On information and belief, a Cisco ’462 Accused Product Network comprises a second addressing domain, separate and independent from the first addressing domain, for identifying each of the end modules exclusive of identifying the intermediate module. For example, edge switches in a Cisco ’462 Accused Product Network can be addressed using IP addresses, but IP addresses are not used to address intermediate switches when forwarding packets within a FabricPath network.

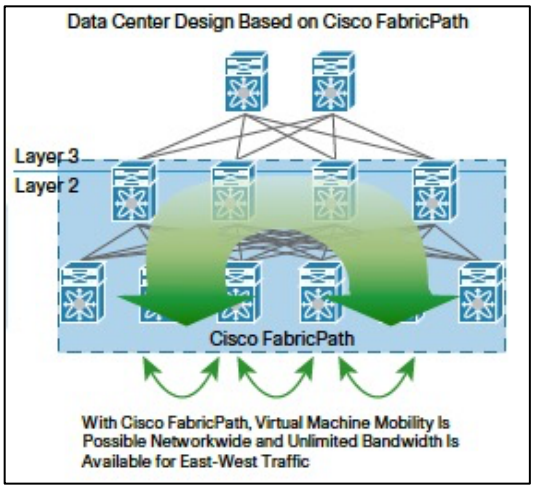

Cisco FabricPath At-A-Glance, CISCO TECHNICAL DOCUMENTATION (2010), at 1.

77. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco '462 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the '462 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

78. On information and belief, Cisco also indirectly infringes the '462 patent by actively inducing infringement under 35 U.S.C. § 271(b).

79. On information and belief, Cisco has had knowledge of the '462 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the '462 patent and knew of its infringement, including by way of this lawsuit.

80. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco '462 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the '462 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '462 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco '462 Accused Products, which are capable of operating in a manner that infringes one or more claims of the '462 patent, including at least claim 1, and

Cisco further provides documentation and training materials that cause customers of the Cisco '462 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the '462 patent. By providing instruction and training to customers on how to use the Cisco '462 Accused Products, Cisco specifically intended to induce infringement of the '462 patent, including at least claim 1. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco '462 Accused Products and to actively induce its customers to infringe the '462 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '462 patent, knowing that such use constitutes infringement of the '462 patent.

81. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the '462 patent.

82. As a result of Cisco's infringement of the '462 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco's infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 6,788,701

83. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

84. U.S. Patent No. 6,788,701 ("the '701 patent"), entitled "Communication Network Having Modular Switches that Enhance Data Throughput," was filed on May 14, 1999. Dunti is the owner by assignment of the '701 patent. A true and correct copy of the '701 patent is attached hereto as Exhibit B. The '701 patent claims a specific architecture, system, and method for efficiently transferring packets of data across a communication network.

85. The '701 patent has been cited by at least fifteen United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '701 patent as relevant prior art:

- Alcatel Lucent S.A.;
- Terascale Supercomputing, Inc.;
- Arbor Networks, Inc.;
- Apple, Inc.;
- International Business Machines Corporation;
- Marvell International, Ltd.; and
- Ericsson.

86. The ’701 patent teaches, for example, an addressing and distributed routing mechanism used by forwarding modules (i.e., switches) that are topologically related to one another based on their position within a network. The modules, due to an awareness of their position or location with respect to the network, enable adaptive fast forwarding of packets across the network. Instead of statically routing packets in the same manner each time, as in conventional switches, the modules include some features of conventional routers, but without the detriments of routers. The modules can forward packets of data relatively quickly (similar to conventional switches), and can dynamically change the forwarding path based on activity within the network (similar to conventional routers).

87. The switches described in the ’701 patent can be used to forward or route incoming packets received on an input port to one or more output ports. Each switch within the network is assigned a unique identification number that is used for routing within the network. When a switch within the network receives an incoming packet on an input port, it decodes part of the packet to direct the packet to the appropriate output port, as shown in Figure 6 below. The switches are aware of their position relative to the network and their neighboring modules, and they use that knowledge to determine which output port to use for forwarding the packet.

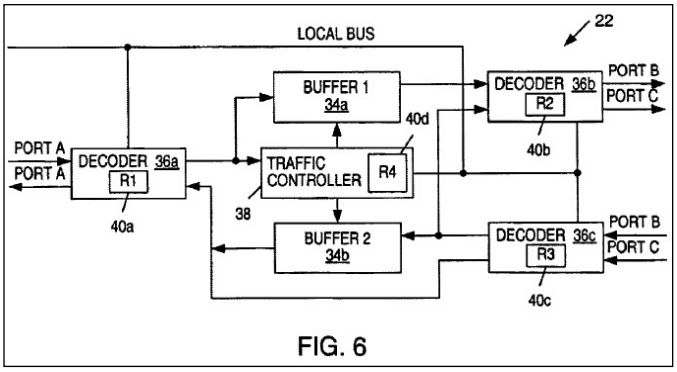

FIG. 6

’701 Patent, Fig. 6.

88. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

89. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch, Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus 9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, “the Cisco Nexus Accused Products”).

90. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, “the Cisco MDS Accused Products”).

91. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

92. Cisco makes, uses, sells, and/or offers to sell the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, "the Cisco '701 Accused Products").

93. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '701 Accused Products ("a Cisco '701 Accused Product Network").

94. On information and belief, a Cisco '701 Accused Product Network implements at least Cisco's FabricPath technology.

95. On information and belief, the Cisco '701 Accused Products comprise a switch.

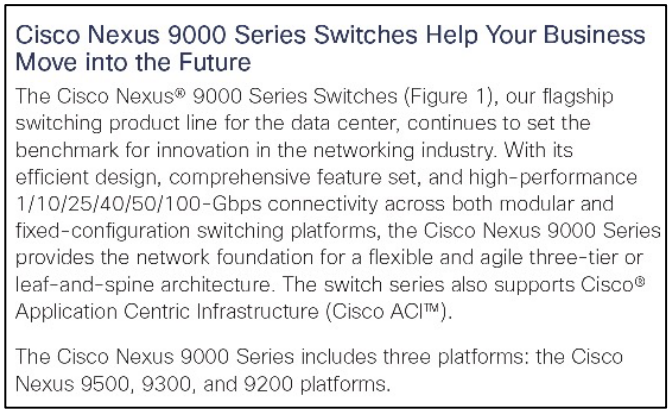

Cisco Nexus 9000 Series Switches Help Your Business Move into the Future

The Cisco Nexus® 9000 Series Switches (Figure 1), our flagship switching product line for the data center, continues to set the benchmark for innovation in the networking industry. With its efficient design, comprehensive feature set, and high-performance 1/10/25/40/50/100-Gbps connectivity across both modular and fixed-configuration switching platforms, the Cisco Nexus 9000 Series provides the network foundation for a flexible and agile three-tier or leaf-and-spine architecture. The switch series also supports Cisco® Application Centric Infrastructure (Cisco ACI™).

The Cisco Nexus 9000 Series includes three platforms: the Cisco Nexus 9500, 9300, and 9200 platforms.

Cisco Nexus 9000 Series Switches At-A-Glance, CISCO TECHNICAL DOCUMENTATION (2016), at 1.

96. On information and belief, the Cisco '701 Accused Products within a FabricPath network comprise a traffic manager which dispatches a series of read operations to a memory coupled within a data flow path. For example, the FabricPath system architecture includes a Supervisor Engine, and the Cisco '701 Accused Products include memory and at least one processor.

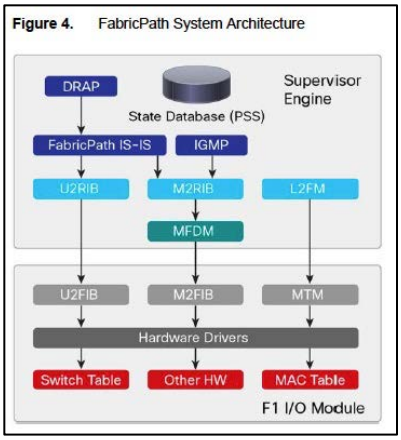

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 6.

97. On information and belief, the Cisco '701 Accused Products within a FabricPath network include a Switch Table and MAC Table comprised in memory, which include a source address and a destination address of a pair of network nodes routably coupled within the data flow path.

Each I/O module has hardware specifically designed to perform FabricPath forwarding lookups and other functions. The primary hardware tables associated with FabricPath forwarding include:

- **Switch table** - Contains Switch IDs and next-hop interfaces
- **MAC address table** - Contains local and remote MAC address entries
- **Other hardware** - Variety of other table memories, hardware registers, etc. required for FabricPath forwarding

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 7.

98. On information and belief, the Cisco '701 Accused Products comprise an input port and an output port.

3.2 FabricPath Interfaces

Every interface involved in FabricPath switching falls into one of two categories (see Figure 5):

- **FabricPath edge port** - FabricPath edge ports are interfaces at the edge of the FabricPath domain. These interfaces run Classic Ethernet, or CE, and behave exactly like normal Ethernet ports. You can attach any CE device to the FabricPath fabric by connecting it to a FabricPath edge port. FabricPath switches perform MAC learning on edge ports and frames transmitted on edge ports are standard IEEE 802.3 Ethernet frames. You can configure an edge port as an access port or as an IEEE 802.1Q trunk.
- **FabricPath core port** - FabricPath core ports always forward Ethernet frames encapsulated in a FabricPath header. Ethernet frames transmitted on a FabricPath interface always carry an IEEE 802.1Q tag and therefore can conceptually be thought of as a trunk port.

Figure 5. FabricPath Edge Ports and Fabricpath Core Ports

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 7.

99. On information and belief, the memory in the Cisco ’701 Accused Products comprises packets of data dispatched from the input port. For example, the Cisco ’701 Accused Products encapsulate incoming data packets within a FabricPath header. Incoming data packets are comprised in memory within an ingress FabricPath switch as they are encapsulated within a FabricPath header as forwarding decisions are made.

100. On information and belief, the Cisco ’701 Accused Products comprise a decoder coupled to the input port for decoding only a single field of bits within a plurality of fields which comprise the destination address. For example, a FabricPath header includes an Outer DA field, which is made up of multiple fields. One of the fields within the Outer DA field is the Switch ID field, which is decoded as forwarding decisions are made.

4.2 FabricPath Encapsulation

FabricPath switches encapsulate all Ethernet frames traversing the fabric inside a 16-byte FabricPath header. Figure 9 shows the details of the FabricPath header.

Figure 9. FabricPath Encapsulation

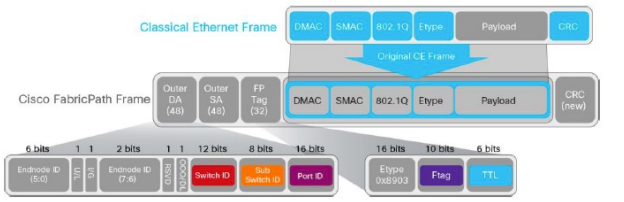

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

4.2.5 Switch ID

Every switch in the FabricPath domain is assigned a unique 12-bit Switch ID. In the outer SA, this field identifies the FabricPath switch that originated the frame (typically the ingress FabricPath edge switch). In the outer DA, this field identifies the destination FabricPath switch.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

101. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco ’701 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the ’701 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

102. On information and belief, Cisco also indirectly infringes the ’701 patent by actively inducing infringement under 35 U.S.C. § 271(b).

103. On information and belief, Cisco has had knowledge of the ’701 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the ’701 patent and knew of its infringement, including by way of this lawsuit.

104. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco ’701 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the ’701 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’701 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’701 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’701 patent, including at least claim 1, and Cisco further provides documentation and training materials that cause customers of the Cisco ’701 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’701 patent. By providing instruction and training to customers on how to use the Cisco ’701 Accused Products, Cisco specifically intended to induce infringement of the ’701 patent, including at least claim 1. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’701 Accused Products and to actively induce its customers to infringe the ’701 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’701 patent, knowing that such use constitutes infringement of the ’701 patent.

105. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’701 patent.

106. As a result of Cisco’s infringement of the ’701 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT III
INFRINGEMENT OF U.S. PATENT NO. 6,804,235

107. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

108. U.S. Patent No. 6,804,235 ("the '235 patent"), entitled "Address Mapping Mechanism Enabling Multi-Domain Addressing in Communication Networks," was filed on February 27, 2003 and claims priority as a continuation of U.S. Patent Application No. 09/785,899, filed on February 16, 2001. Dunti is the owner by assignment of the '235 patent. A true and correct copy of the '235 patent is attached hereto as Exhibit C.

109. The '235 patent has been cited by six United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '235 patent as relevant prior art:

- Texas Instruments, Inc.; and
- International Business Machines Corporation.

110. The '235 patent teaches, for example, a communication system that transparently maps addresses across multiple addressing domains and/or protocols. The communication system described in the '235 patent operates using a scalable addressing domain of an independent identification layer that is different from the addressing domain interfacing with the network. This independent identification layer is an improvement to the OSI reference model and can be considered an even lower layer addressing domain within the OSI reference model because the existing lower-level layer addressing information is further wrapped with the independent identification layer addressing information.

111. The independent identification layer can be used to represent, for example, unique identification numbers of intermediate modules within the communication system of the '235 patent. The networking modules described in the '235 patent can be classified as either end modules (i.e., entry and exit end modules) or as intermediate modules. End modules are coupled to other networks, addressing domains, or devices outside of the network. Entry end modules

perform protocol wrapping functions as data packets enter the network, and exit end modules strip protocol used by the network as data packets exit the network. Identification addresses for the intermediate modules and end modules of a given network can utilize that network's unique and independent identification layer.

112. As described in the '235 patent, sending a data packet from a source device to a destination device, where the devices are separated by a network with an internal addressing domain that is different from the external addressing domains, requires only a simple mapping function. One addressing domain can be used to forward data from a source device to a unique entry end module and from an exit end module to the destination device. Within the network, among the intermediate modules, a separate and independent addressing domain can be used.

113. When data packets enter a network from a device external to the network, the IP address and Ethernet address within the network layer and the lower-level data/physical layer addressing domains are further wrapped with the independent identification layer source address and corresponding destination addresses unique to that addressing domain. The wrapped information indicates where the data came from external to the network and, due to the wrapped independent identification layer, where within the network the data enters the network and exits the network. When data packets exit the network, an end module strips the wrapped information from the packets.

114. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

115. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch, Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus

9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, "the Cisco Nexus Accused Products").

116. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, "the Cisco MDS Accused Products").

117. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

118. Cisco makes, uses, sells, and/or offers to sell the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, "the Cisco '235 Accused Products").

119. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '235 Accused Products ("a Cisco '235 Accused Product Network").

120. On information and belief, a Cisco '235 Accused Product Network implements at least Cisco's FabricPath technology.

121. On information and belief, a Cisco '235 Accused Product Network comprises a communication network.

122. On information and belief, a Cisco '235 Accused Product Network comprises a plurality of interconnected modules adapted to direct packets of data through the network.

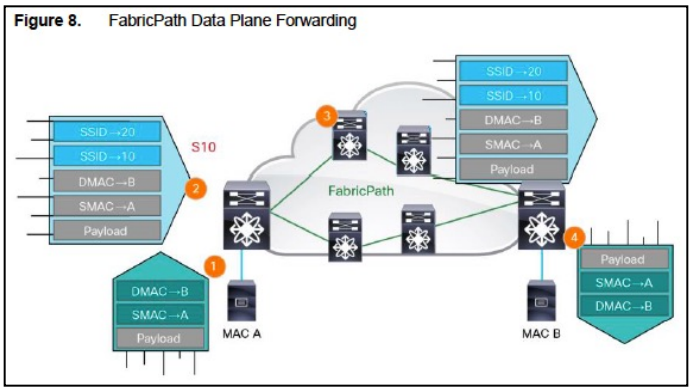

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 9.

123. On information and belief, modules within a Cisco '235 Accused Product Network are identified according to identification numbers contained within a first addressing domain of a first model layer independent and separate from a second addressing domain of a second model layer used to identify modules which forward and receive the packets of data outside the network. For example, each switch within a FabricPath network is assigned a unique Switch ID, which is a unique identification number that is independent of the MAC address, and can be assigned to different FabricPath topologies.

4.2.5 Switch ID

Every switch in the FabricPath domain is assigned a unique 12-bit Switch ID. In the outer SA, this field identifies the FabricPath switch that originated the frame (typically the ingress FabricPath edge switch). In the outer DA, this field identifies the destination FabricPath switch.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

5.6 FabricPath Multitopology

Prior to NX-OS release 6.2(2), Nexus 7000 switches support a single (default) FabricPath topology. All FabricPath VLANs map to the default topology and all FabricPath core ports carry all FabricPath-mode VLANs.

With FabricPath multitopology, introduced in NX-OS release 6.2(2), you can configure additional topologies and map FabricPath-mode VLANs to those topologies. Multitopology controls the scope of VLAN traffic based on which FabricPath core ports carry which VLANs, allowing you to restrict traffic to specific subsections of the network fabric or to independently engineer traffic flows in the network on a per-VLAN basis.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 17.

124. On information and belief, the first model layer used in a Cisco '235 Accused Product Network is an improvement to, and is lower than, a physical layer of the OSI reference model. For example, data packets entering a FabricPath network, which already include headers

from higher layers, are further wrapped/encapsulated within a FabricPath header that includes the Switch ID of the FabricPath egress switch.

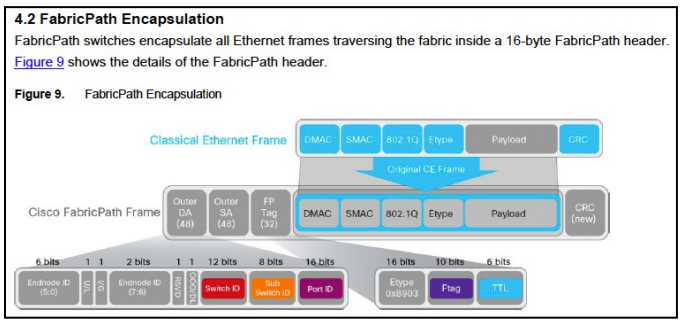

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

125. On information and belief, the second model layer used in a Cisco ’235 Accused Product Network is higher than a physical layer of the OSI reference model. For example, the edge switches in a Cisco ’235 Accused Product Network can use IP addresses to route data packets outside of a FabricPath network, and the IP address layer is higher than a physical layer of the OSI model.

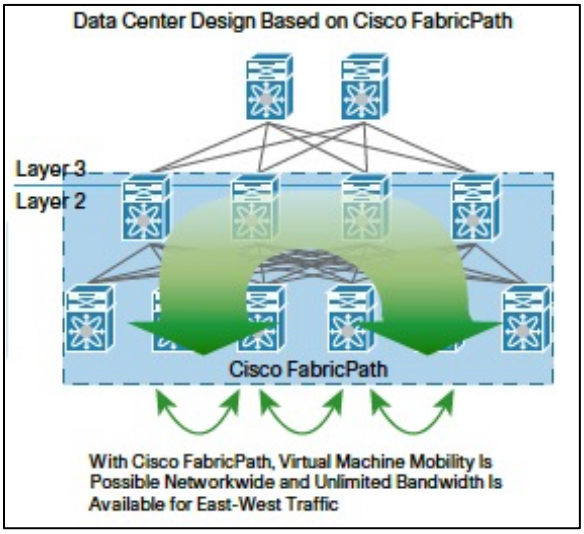

Cisco FabricPath At-A-Glance, CISCO TECHNICAL DOCUMENTATION (2010), at 1.

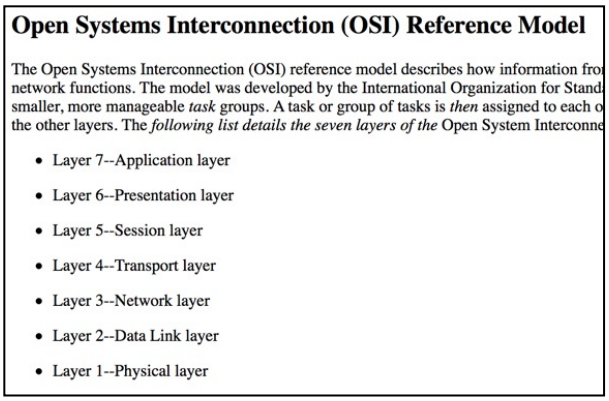

Open Systems Interconnection (OSI) Reference Model

The Open Systems Interconnection (OSI) reference model describes how information fro
network functions. The model was developed by the International Organization for Stand
smaller, more manageable *task* groups. A task or group of tasks is *then* assigned to each o
the other layers. The *following list details the seven layers of the* Open System Interconne

- Layer 7--Application layer
- Layer 6--Presentation layer
- Layer 5--Session layer
- Layer 4--Transport layer
- Layer 3--Network layer
- Layer 2--Data Link layer
- Layer 1--Physical layer

Internetworking Basics, CISCO PRESS, available at http://www.cisco.com/cpress/cc/td/cpress/fund/ith/ith01gb.htm (accessed July 19, 2016).

126. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco ’235 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the ’235 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

127. On information and belief, Cisco also indirectly infringes the ’235 patent by actively inducing infringement under 35 U.S.C. § 271(b).

128. On information and belief, Cisco has had knowledge of the ’235 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the ’235 patent and knew of its infringement, including by way of this lawsuit.

129. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco ’235 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the ’235 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’235 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’235 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’235 patent, including at least claim 1, and Cisco further provides documentation and training materials that cause customers of the Cisco

’235 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’235 patent. By providing instruction and training to customers on how to use the Cisco ’235 Accused Products, Cisco specifically intended to induce infringement of the ’235 patent, including at least claim 1. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’235 Accused Products and to actively induce its customers to infringe the ’235 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’235 patent, knowing that such use constitutes infringement of the ’235 patent.

130. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’235 patent.

131. As a result of Cisco’s infringement of the ’235 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT IV
INFRINGEMENT OF U.S. PATENT NO. 6,643,286

132. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

133. U.S. Patent No. 6,643,286 (“the ’286 patent”), entitled “Modular Switches Interconnected Across a Communication Network to Achieve Minimal Address Mapping or Translation Between Termination Devices,” was filed on May 14, 1999. Dunti is the owner by assignment of the ’286 patent. A true and correct copy of the ’286 patent is attached hereto as Exhibit D. The ’286 patent claims a specific architecture, system, and method for efficiently transferring packets of data across a communication network with hierarchical levels of high speed switches throughout the network.

134. The ’286 patent has been cited by fourteen issued United States patents and published patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ’286 patent as relevant prior art.

- Google, Inc.;
- Ciena Corporation;
- Advanced Micro Devices, Inc.; and
- Fujitsu Ltd.

135. The ’286 patent teaches, for example, an addressing and distributed routing mechanism used by forwarding modules within a network that perform fast decoding to forward data packets, thereby reducing the number of full network address mapping/translation operations as the packet traverses the network. It claims a technical solution to a problem unique to computer networks—quickly and efficiently transmitting data packets through a computer network without needing to perform a full network address mapping/translation operation at every intermediate node.

136. The forwarding modules of the ’286 patent are topologically related to one another based on their position within the network and can perform adaptive fast forwarding of packets across the network due to an awareness of their position or location with respect to the network.

137. The adaptive fast forwarding occurs through decoding operations using a series of comparisons within only select switches. An entry end switch wraps entering data packets with internal control information that includes an originating identification number of the entry end switch and an identification number of the exit end switch. The wrapped packet can then be forwarded through the structured network without performing full network address translation operations at each hop. When the packet arrives at the exit end switch, the internal control information of the network is stripped from the packet, and a mapping table is used to forward the packet to a destination termination device connected to the exit end switch. This full network

address translation at the exit end switch bridges the gap between the structured network and any external protocol or domain.

138. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

139. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch, Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus 9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, “the Cisco Nexus Accused Products”).

140. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, “the Cisco MDS Accused Products”).

141. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

142. Cisco makes, uses, sells, and/or offers to sell the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, “the Cisco ’286 Accused Products”).

143. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '286 Accused Products ("a Cisco '286 Accused Product Network").

144. On information and belief, a Cisco '286 Accused Product Network implements at least Cisco's FabricPath technology.

145. On information and belief, a Cisco '286 Accused Product Network comprises a communication network.

146. On information and belief, a Cisco '286 Accused Product Network comprises an entry end switch.

147. On information and belief, a Cisco '286 Accused Product Network comprises an exit end switch, which is selectably coupled to multiple termination devices including at least one exit termination device.

148. On information and belief, a Cisco '286 Accused Product Network comprises multiple intermediate switches coupled between the entry end switch and the exit end switch. For example, the figure below shows an entry end switch (i.e., FabricPath ingress switch), an exit end switch (i.e., FabricPath egress switch), and multiple intermediate switches in between them (i.e., FabricPath core switches). The exit end switch can be connected to multiple hosts.

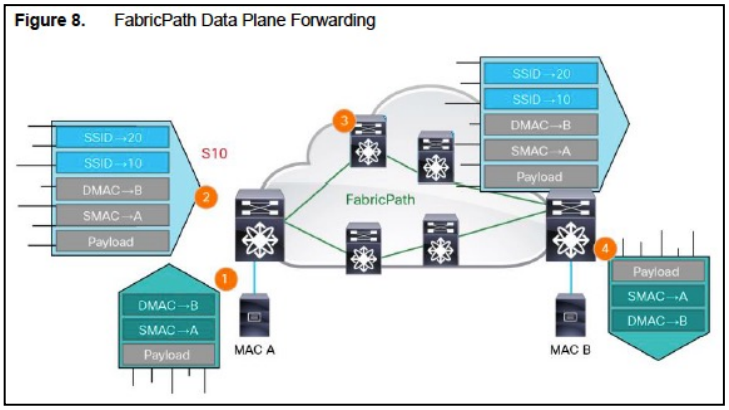

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 9.

149. On information and belief, an entry end switch in a Cisco '286 Accused Product Network compiles a packet that contains a destination address of the exit end switch. For example, an entry end switch (i.e., FabricPath ingress switch) encapsulates an incoming data

packet within a FabricPath header. The FabricPath header includes an "Outer DA" field, which contains the unique Switch ID of the exit end switch (i.e., FabricPath egress switch).

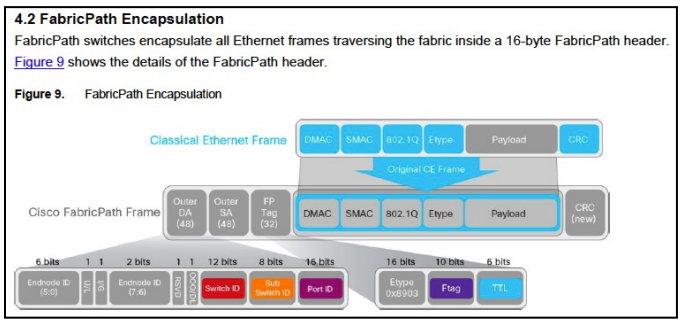

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

150. On information and belief, in a Cisco '286 Accused Product Network, the packet is forwarded through the plurality of intermediate switches with each intermediate switch having an identification number which points the packet to a successive one of the plurality of intermediate switches and finally to the exit end switch which performs the entirety of all translation needed by the communication network to route the packet from the exit end switch to the exit termination device. For example, each intermediate switch (i.e., FabricPath core switch) uses the Switch ID within the FabricPath header to point the packet to the next FabricPath switch.

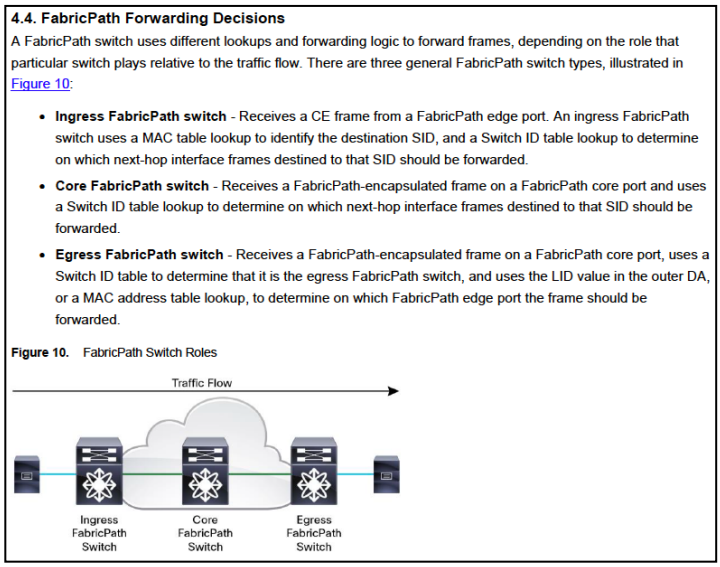

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 12.

151. In addition, on information and belief, the exit end switch (i.e., FabricPath egress switch) performs the entirety of all translation needed by the FabricPath network to route the packet from the FabricPath egress switch to the exit termination device (i.e., the packet's final destination). For example, FabricPath switches do not populate MAC tables within the FabricPath network: "the first general rule of FabricPath MAC learning is that only FabricPath edge switches populate the MAC table and use MAC table lookups to forward frames."

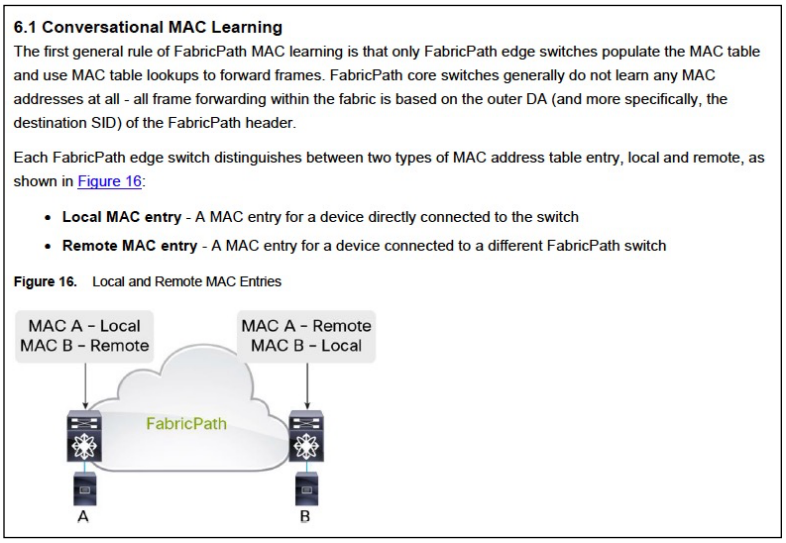

6.1 Conversational MAC Learning

The first general rule of FabricPath MAC learning is that only FabricPath edge switches populate the MAC table and use MAC table lookups to forward frames. FabricPath core switches generally do not learn any MAC addresses at all - all frame forwarding within the fabric is based on the outer DA (and more specifically, the destination SID) of the FabricPath header.

Each FabricPath edge switch distinguishes between two types of MAC address table entry, local and remote, as shown in Figure 16:

- **Local MAC entry** - A MAC entry for a device directly connected to the switch
- **Remote MAC entry** - A MAC entry for a device connected to a different FabricPath switch

Figure 16. Local and Remote MAC Entries

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 19.

152. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco '286 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the '286 patent, including at least claim 6, pursuant to 35 U.S.C. § 271(a).

153. On information and belief, Cisco also indirectly infringes the '286 patent by actively inducing infringement under 35 U.S.C. § 271(b).

154. On information and belief, Cisco has had knowledge of the '286 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the '286 patent and knew of its infringement, including by way of this lawsuit.

155. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco ’286 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the ’286 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’286 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’286 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’286 patent, including at least claim 6, and Cisco further provides documentation and training materials that cause customers of the Cisco ’286 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’286 patent. By providing instruction and training to customers on how to use the Cisco ’286 Accused Products, Cisco specifically intended to induce infringement of the ’286 patent, including at least claim 6. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’286 Accused Products and to actively induce its customers to infringe the ’286 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’286 patent, knowing that such use constitutes infringement of the ’286 patent.

156. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’286 patent.

157. As a result of Cisco’s infringement of the ’286 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT V
INFRINGEMENT OF U.S. PATENT NO. 7,778,259

158. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

159. U.S. Patent No. 7,778,259 ("the '259 patent"), entitled "Network Packet Transmission Mechanism," was filed on June 11, 2004. Dunti is the owner by assignment of the '259 patent. A true and correct copy of the '259 patent is attached hereto as Exhibit E.

160. The '259 patent has been cited by ten United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the '259 patent as relevant prior art:

- International Business Machines Corporation;
- Toshiba Corporation;
- Nicira, Inc.; and
- The University of Zurich.

161. The '259 patent teaches, for example, a communication network that efficiently transfers data packets by using an independent numbering mechanism with distinct identification addresses, referred to as transport IDs, for transporting packets across a network. This solution eliminates complex lookup operations at intermediate modules, resulting in faster transmission across the network.

162. Each packet in the network of the '259 patent is embedded with unique destination transport ID information when the packet enters the network and carries this routing information along with the data. This transport ID-based packet transmission mechanism utilizes the logical structure in the network, which enables simple distributed packet direction operations as the packet traverses the network.

163. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

164. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch,

Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus 9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, "the Cisco Nexus Accused Products").

165. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, "the Cisco MDS Accused Products").

166. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

167. Cisco makes, uses, sells, and/or offers to sell the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, "the Cisco '259 Accused Products").

168. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '259 Accused Products ("a Cisco '259 Accused Product Network").

169. On information and belief, a Cisco '259 Accused Product Network implements at least Cisco's FabricPath technology.

170. On information and belief, the Cisco '259 Accused Products perform a method of transporting packets across a network.

171. On information and belief, the Cisco ’259 Accused Products embed a destination transport identification to a data packet when the data packet enters the network. For example, data packets entering a FabricPath network are encapsulated within a FabricPath header, which includes, for example, an “Outer DA” field that contains the Switch ID of the exit end switch.

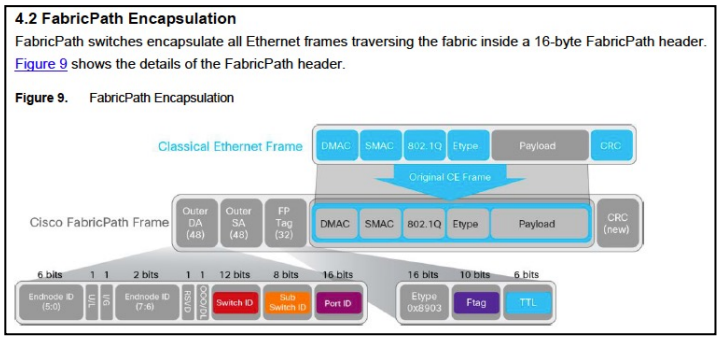

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 10.

172. On information and belief, the Cisco ’259 Accused Products connect a plurality of routing switches within a network with the routing switches grouped into two or more groups within the network based on network topology. For example, in a FabricPath network, the Cisco ’259 Accused Products are grouped into FabricPath ingress switches, FabricPath core switches, and FabricPath egress switches based on whether they are entry switches, intermediate switches, or exit switches.

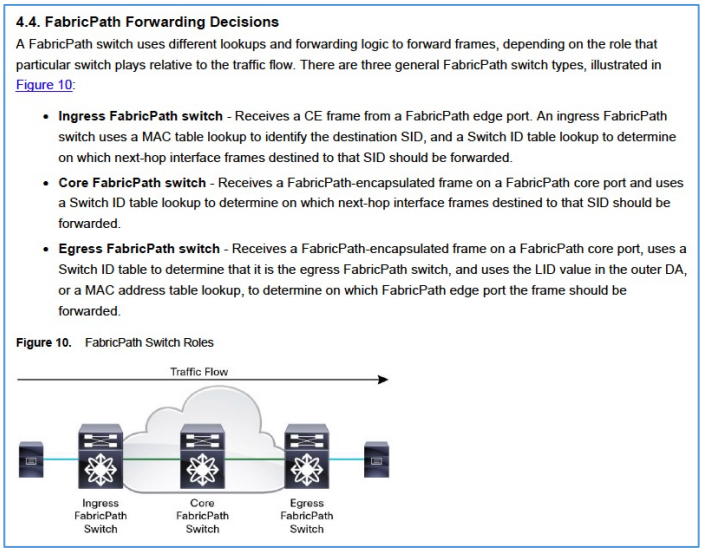

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 12.

173. On information and belief, the Cisco '259 Accused Products assign a unique transport identification number to each routing switch indicative, at least in part, of the network topology. For example, each FabricPath switch is assigned a unique Switch ID, which is part of a network topology database and adjacency tables that are built and maintained by the FabricPath Layer 2 IS-IS protocol and, when used with the topology database and/or the adjacency tables, indicates, at least in part, the network topology.

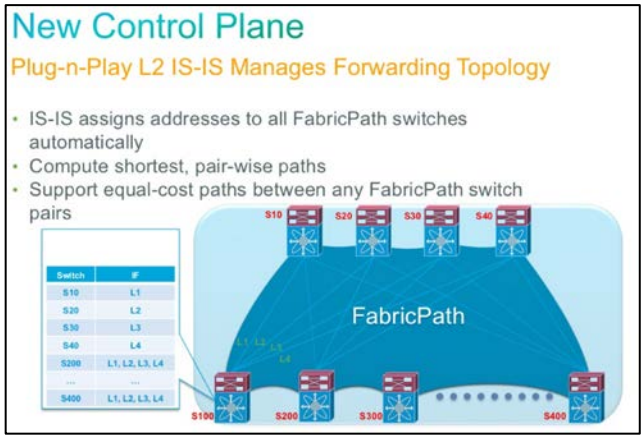

Babi Seal and Patrick Warichet, *Efficient Data Center Design with FabricPath*, CISCO IOS ADVANTAGE WEBINAR PRESENTATION (2012), at 23, available at http://www.slideshare.net/getyourbuildon/fabric-path-webinar (accessed Aug. 26, 2016).

Default IS-IS Behavior with FabricPath

The interfaces in a FabricPath network run only the FabricPath Layer 2 IS-IS protocol; you do not need to run STP in the FabricPath network because FabricPath Layer 2 IS-IS discovers topology information dynamically.

FabricPath Layer 2 IS-IS is a dynamic link-state routing protocol that detects changes in the network topology and calculates loop-free paths to other nodes in the network. Each FabricPath device maintains a link-state database (LSDB) that describes the state of the network; each device updates the status of the links that are adjacent to the device. The FabricPath device sends advertisements and updates to the LSDB through all the existing adjacencies. FabricPath Layer 2 IS-IS protocol packets do not conflict with standard Layer 3 IS-IS packets because the FabricPath packets go to a different Layer 2 destination MAC address than that used by standard IS-IS for IPv4/IPv6 address families.

Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Jan. 28, 2015), at 11.

show fabricpath isis switch-id	Displays the switch IDs and reachability information for the topology.
show fabricpath isis topology summary	Displays the FabricPath Layer 2 IS-IS topology database.

Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Jan. 28, 2015), at 95.

174. On information and belief, the Cisco ’259 Accused Products compare the destination transport identification of a packet with the transport identification of a routing switch. For example, data packets entering a FabricPath network are encapsulated within a FabricPath header, which includes an “Outer DA” field that contains the Switch ID of the FabricPath egress switch. FabricPath switches, such as the Cisco ’259 Accused Products, compare the values in the OuterDA field to values in the switches’ memory.

4.4. FabricPath Forwarding Decisions

A FabricPath switch uses different lookups and forwarding logic to forward frames, depending on the role that particular switch plays relative to the traffic flow. There are three general FabricPath switch types, illustrated in Figure 10:

- **Ingress FabricPath switch** - Receives a CE frame from a FabricPath edge port. An ingress FabricPath switch uses a MAC table lookup to identify the destination SID, and a Switch ID table lookup to determine on which next-hop interface frames destined to that SID should be forwarded.
- **Core FabricPath switch** - Receives a FabricPath-encapsulated frame on a FabricPath core port and uses a Switch ID table lookup to determine on which next-hop interface frames destined to that SID should be forwarded.
- **Egress FabricPath switch** - Receives a FabricPath-encapsulated frame on a FabricPath core port, uses a Switch ID table to determine that it is the egress FabricPath switch, and uses the LID value in the outer DA, or a MAC address table lookup, to determine on which FabricPath edge port the frame should be forwarded.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 12.

175. On information and belief, the Cisco ’259 Accused Products forward data packets through a network based on the comparison of destination transport identification. For example, FabricPath switches, such as the Cisco ’259 Accused Products, forward encapsulated data packets using the Switch ID of the FabricPath egress switch.

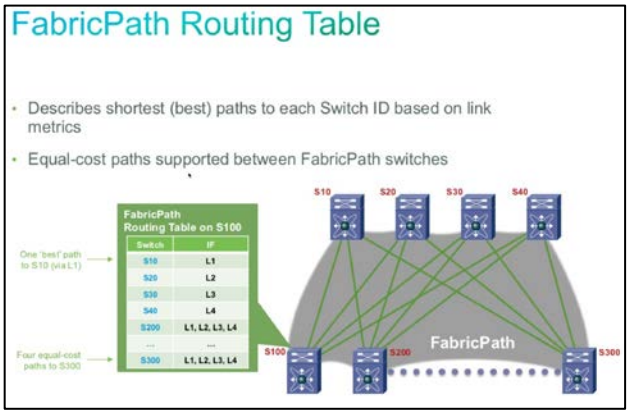

Babi Seal and Patrick Warichet, *Efficient Data Center Design with FabricPath*, CISCO IOS ADVANTAGE WEBINAR PRESENTATION (2012), at 23, available at http://www.slideshare.net/getyourbuildon/fabric-path-webinar (accessed Aug. 26, 2016).

4.1 Basic Data Plane Operation

Figure 8 illustrates the basics of FabricPath data-plane forwarding.

Figure 8. FabricPath Data Plane Forwarding

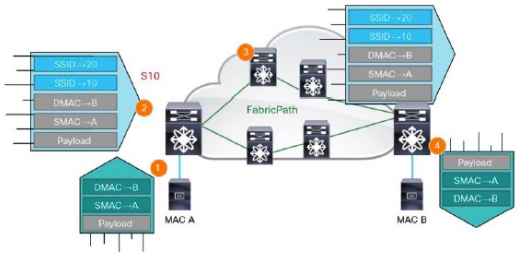

For the purposes of this example, assume that all control-plane routing information already exists, and data-plane MAC address learning for Host A and Host B has already occurred.

The following steps provide the basics of FabricPath data-plane operation:

1. Host A sends a CE frame, sourced from MAC A and destined to MAC B. The ingress FabricPath switch receives the frame in a FabricPath-mode VLAN on a FabricPath edge port and performs a MAC address table lookup on the DMAC.
2. The MAC lookup result indicates the frame should be sent to FabricPath switch S20. The ingress FabricPath switch encapsulates the frame in a FabricPath header for forwarding on a FabricPath core port, with a destination Switch ID of S20.

   S10 has two equal-cost paths to reach S20. One of the paths is selected using a hash function, with input from the packet data (source and destination IP plus Layer 4 ports by default).
3. Other FabricPath switches simply forward the frame based on their routing table ("what is the best path to S20?").
4. When S20 receives the frame, it removes the FabricPath header and forwards the frame as a Classic Ethernet frame on the FabricPath edge port where Host B is connected.

Nexus 7000 FabricPath White Paper, CISCO TECHNICAL DOCUMENTATION (Oct. 9, 2013), at 9.

176. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco '259 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the '259 patent, including at least claim 9, pursuant to 35 U.S.C. § 271(a).

177. On information and belief, Cisco also indirectly infringes the '259 patent by actively inducing infringement under 35 U.S.C. § 271(b).

178. On information and belief, Cisco has had knowledge of the '259 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the '259 patent and knew of its infringement, including by way of this lawsuit.

179. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco '259 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and

customary use of the accused products would infringe the ’259 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’259 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’259 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’259 patent, including at least claim 9, and Cisco further provides documentation and training materials that cause customers of the Cisco ’259 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’259 patent. By providing instruction and training to customers on how to use the Cisco ’259 Accused Products, Cisco specifically intended to induce infringement of the ’259 patent, including at least claim 9. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’259 Accused Products and to actively induce its customers to infringe the ’259 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’259 patent, knowing that such use constitutes infringement of the ’259 patent.

180. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’259 patent.

181. As a result of Cisco’s infringement of the ’259 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT VI
INFRINGEMENT OF U.S. PATENT NO. 6,912,196

182. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

183. U.S. Patent No. 6,912,196 (“the ’196 patent”), entitled “Communication Network and Protocol Which Can Efficiently Maintain Transmission Across a Disrupted Network,” was filed on May 15, 2000. Dunti is the owner by assignment of the ’196 patent. A true and correct

copy of the ’196 patent is attached hereto as Exhibit F. The ’196 patent claims a specific packet architecture, communication system, and method for determining the location at which a network is disrupted, disabled, and/or severed.

184. The ’196 patent has been cited by thirty-seven United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ’196 patent as relevant prior art:

- Alcatel Lucent S.A.;
- Fujitsu, Ltd.;
- Google, Inc.;
- International Business Machines Corporation;
- Samsung Electronics Co., Ltd.;
- Terascale Supercomputing, Inc.;
- Siemens AG; and
- NEC Corporation.

185. The ’196 patent teaches, for example, an improved packet protocol and communication system that can determine where within a network a transmission error exists.

186. Figure 3 of the ’196 patent, shown below, depicts a ring topology with multiple modules labeled “A” through “F” that communicate over two transmission channels that form a single transmission path. One channel is used for counter-clockwise data transmission, while the other channel is used for clockwise data transmission.

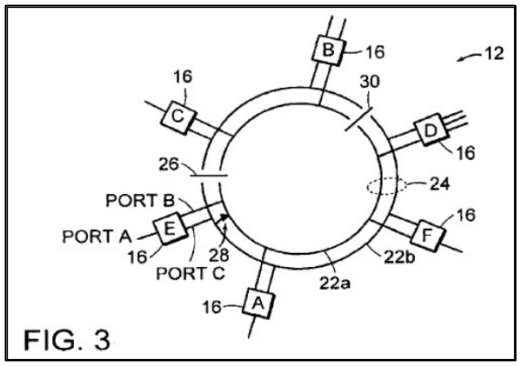

’196 Patent, Fig. 3.

187. Figure 3 shows an example where the transmission path has been severed between modules C and E. In an attempted transmission from module A to module C (clockwise through

module E), module E will detect the severance and notify the other modules by employing a loop-back path of the packet sent from module A back to module A. When module E detects the downstream severance, it sends control bits to the originating module A indicating the downstream error. Since receiving module E was the last module in the path before the severance, module E sends both control bits and error identification bits. The control bits are set to indicate a disruption immediately downstream of receiving module E. The error identification bits identify the receiving module E by the identification number assigned to that module.

188. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to network communications.

189. Cisco makes, uses, sells, offers to sell, and/or imports the ML-Series of networking cards, including but not limited to the Cisco ONS 15454 ML-Series Ethernet cards (collectively, "the Cisco '196 Accused Products").

190. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '196 Accused Products ("a Cisco '196 Accused Product Network").

191. On information and belief, a Cisco '196 Accused Product Network implements at least the IEEE 802.17 Resilient Packet Ring ("RPR") protocol and/or Cisco's proprietary RPR protocol.

192. On information and belief, a Cisco '196 Accused Product Network comprises a communication system.

ML-Series Feature List

The ML100T-12, ML100X-8, and the ML1000-2 cards have the following features:

* * *

- Resilient packet ring (RPR)
  - Ethernet FCS preservation for customers
  - CRC error alarm generation
  - FCS detection and threshold configuration
  - Shortest path determination
  - Keep alives

Ethernet Card Software Feature and Configuration Guide for the Cisco ONS 15454 SDH, Cisco ONS 15454, and Cisco ONS 15327, CISCO TECHNICAL DOCUMENTATION (Aug. 2012), at 1-2.

Understanding RPR-IEEE

RPR, as described in IEEE 802.17, is a metropolitan area network (MAN) technology supporting data transfer among stations interconnected in a dual-ring configuration. The IEEE 802.17b spatially aware sublayer amendment is not yet ratified but is expected to add support for bridging to IEEE 802.17. Since the amendment is not yet ratified, no equipment is currently IEEE 802.17b compliant. The ML-Series card's RPR-IEEE is based on the expected IEEE 802.17b based standard.

The ML-Series card supports RPR-IEEE. RPR-IEEE is well suited for transporting Ethernet over a SONET/SDH ring topology and enables multiple ML-Series cards to become one functional network segment. When used in this role, RPR-IEEE overcomes the limitations of earlier schemes, such as IEEE 802.1D Spanning Tree Protocol (STP), IEEE 802.1W Rapid Spanning Tree Protocol (RSTP), and SONET/SDH.

Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Aug. 2012), at 26-1.

Understanding Cisco Proprietary RPR

Cisco proprietary RPR is a MAC protocol operating at the Layer 2 level. It is well suited for transporting Ethernet over a SONET/SDH ring topology and it enables multiple ML-Series cards to become one functional network segment or shared packet ring (SPR). Cisco proprietary RPR overcomes the limitations of earlier schemes, such as IEEE 802.1D Spanning Tree Protocol (STP), IEEE 802.1W Rapid Spanning Tree Protocol (RSTP), and SONET/SDH when used in this role.

In Software Release 7.2 and later, the ML-Series card supports IEEE 802.17b based RPR (RPR-IEEE) in addition to Cisco proprietary RPR. Throughout this book, Cisco proprietary RPR is referred to as Cisco proprietary RPR, and IEEE 802.17 based RPR is referred to as RPR-IEEE. This chapter covers Cisco proprietary RPR. Chapter 26, "Configuring IEEE 802.17b Resilient Packet Ring" covers IEEE 802.17b based RPR.

Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Aug. 2012), at 17-2.

193. On information and belief, a Cisco '196 Accused Product Network comprises at least two transmission channels. For example, an RPR ring is made up of dual counter-rotating rings that are used to transmit data, as shown below:

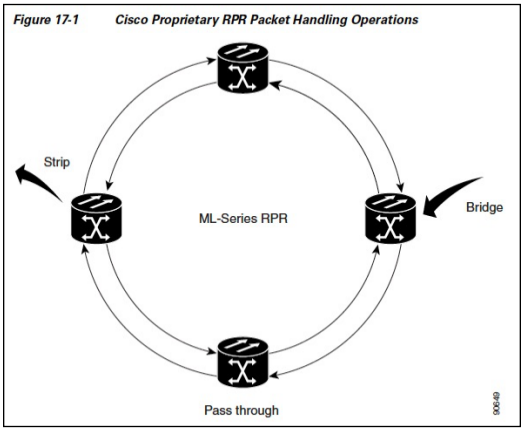

Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Aug. 2012), at 17-3.

194. On information and belief, a Cisco '196 Accused Product Network comprises a receiving module connected to the transmission channels, which can be seen, for example, in Figure 17-1 shown directly above.

195. On information and belief, a receiving module in a Cisco '196 Accused Product Network comprises a loop-back conductor that connects the two transmission channels. For example, each receiving module can receive packets on either the west or the east interface and can transmit packets out of either the west or the east interface, such that a packet received on the east interface can be looped back and sent out on the east interface. RPR networks use wrapping protection to avoid disrupted links.

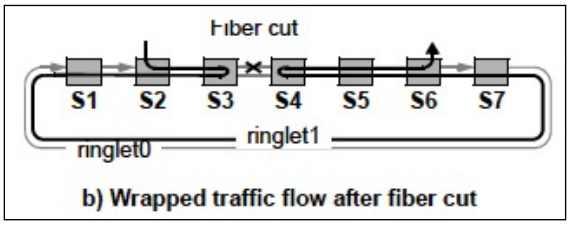

IEEE Std 802.17-2011, *Part 17: Resilient Packet Ring (RPR) Access Method and Physical Layer Specifications*, IEEE STANDARDS ASSOCIATION (Sept. 20, 2011), at 279.

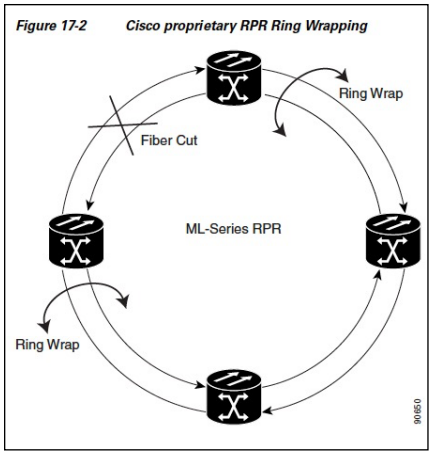

Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Aug. 2012), at 17-4.

196. On information and belief, a receiving module in a Cisco '196 Accused Product Network returns a packet containing error bits if one or more of the transmission channels

downstream of the receiving module is disturbed. For example, when a span is determined to be an edge (i.e., the transmission channel downstream of the module is disturbed), topology and protection ("TP") frames, which are control frames, are sent to report the edge. The payload of a TP frame includes two fields that indicate whether an edge is present on either the west span or the east span of a station.

11.3.1 Topology and protection (TP) frame format

The fixed-length TP frame is identified by the *controlType* field value of CT_TOPO_PROT and is broadcast to minimize transmission delays. The TP frame contents include information for signaling span protection state, for discovery of the physical topology of the ring, and for reporting station preferences information, as specified in Figure 11.6.

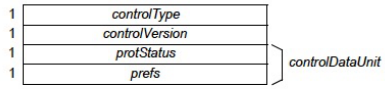

Figure 11.6—TP payload

IEEE Std 802.17-2011, *Part 17: Resilient Packet Ring (RPR) Access Method and Physical Layer Specifications*, IEEE STANDARDS ASSOCIATION (Sept. 20, 2011), at 292.

11.3.1.3 *protStatus*

The subfields of the 8-bit *protStatus* (protection status) field are shown in Figure 11.7 and described thereafter.

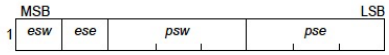

Figure 11.7—*protStatus* field format

11.3.1.3.1 *esw*: A (edge state, west) bit that indicates whether an edge is present on the west span of a station. A value of 0 indicates that there is no edge, whereas a value of 1 indicates that an edge is present.

11.3.1.3.2 *ese*: A (edge state, east) bit that indicates whether an edge is present on the east span of a station. A value of 0 indicates that there is no edge, whereas a value of 1 indicates that an edge is present.

IEEE Std 802.17-2011, *Part 17: Resilient Packet Ring (RPR) Access Method and Physical Layer Specifications*, IEEE STANDARDS ASSOCIATION (Sept. 20, 2011), at 292-93.

197. On information and belief, the error bits sent by a receiving module in a Cisco '196 Accused Product Network comprise a unique identification number assigned to the receiving module to note the receiving module was the last of a plurality of modules that received the packet destined for a destination module dissimilar from and located downstream of the receiving module. For example, a control frame in an RPR network includes the source

address of the originating module, which is a unique identifier that identifies the module sending the error bits indicating that an edge is present at the module.

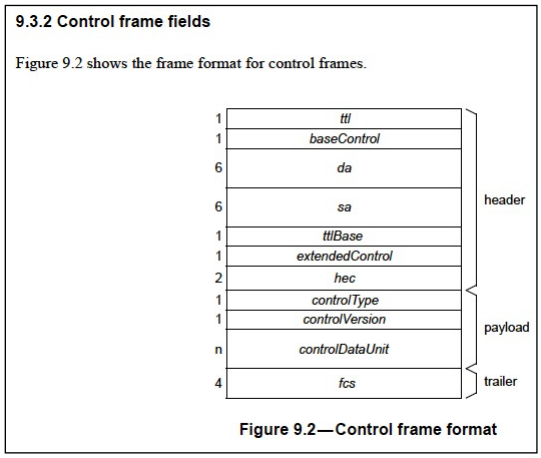

IEEE Std 802.17-2011, *Part 17: Resilient Packet Ring (RPR) Access Method and Physical Layer Specifications*, IEEE STANDARDS ASSOCIATION (Sept. 20, 2011), at 206.

198. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco ’196 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the ’196 patent, including at least claim 10, pursuant to 35 U.S.C. § 271(a).

199. On information and belief, Cisco also indirectly infringes the ’196 patent by actively inducing infringement under 35 U.S.C. § 271(b).

200. On information and belief, Cisco had knowledge of the ’196 patent since at least July 2, 2013. Cisco cited the ’196 patent in the following issued United States patents and published patent applications of its wholly-owned subsidiary Cisco Technology, Inc.:

- U.S. Patent No. 8,477,638
- U.S. Patent Publication No. 2010/0135295

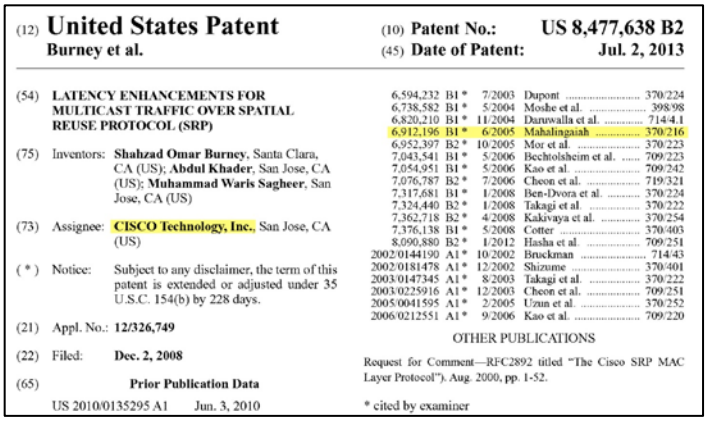

(12) **United States Patent**
Burney et al.

(10) **Patent No.: US 8,477,638 B2**
(45) **Date of Patent: Jul. 2, 2013**

(54) **LATENCY ENHANCEMENTS FOR MULTICAST TRAFFIC OVER SPATIAL REUSE PROTOCOL (SRP)**

(75) Inventors: **Shahzad Omar Burney**, Santa Clara, CA (US); **Abdul Khader**, San Jose, CA (US); **Muhammad Waris Sagheer**, San Jose, CA (US)

(73) Assignee: **CISCO Technology, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 228 days.

(21) Appl. No.: **12/326,749**

(22) Filed: **Dec. 2, 2008**

(65) **Prior Publication Data**

US 2010/0135295 A1 Jun. 3, 2010

6,594,232	B1 *	7/2003	Dupont	370/224
6,738,582	B1 *	5/2004	Moshe et al.	398/98
6,820,210	B1 *	11/2004	Daruwalla et al.	714/4.1
6,912,196	B1 *	6/2005	Mahalingaiah	370/216
6,952,397	B2 *	10/2005	Mor et al.	370/223
7,043,541	B1 *	5/2006	Bechtolsheim et al.	709/223
7,054,951	B1 *	5/2006	Kao et al.	709/242
7,076,787	B2 *	7/2006	Cheon et al.	719/321
7,317,681	B1 *	1/2008	Ben-Dvora et al.	370/224
7,324,440	B2 *	1/2008	Takagi et al.	370/222
7,362,718	B2 *	4/2008	Kakivaya et al.	370/254
7,376,138	B1 *	5/2008	Cotter	370/403
8,090,880	B2 *	1/2012	Hasha et al.	709/251
2002/0144190	A1 *	10/2002	Bruckman	714/43
2002/0181478	A1 *	12/2002	Shizume	370/401
2003/0147345	A1 *	8/2003	Takagi et al.	370/222
2003/0225916	A1 *	12/2003	Cheon et al.	709/251
2005/0041595	A1 *	2/2005	Uzun et al.	370/252
2006/0212551	A1 *	9/2006	Kao et al.	709/220

OTHER PUBLICATIONS

Request for Comment—RFC2892 titled "The Cisco SRP MAC Layer Protocol"). Aug. 2000, pp. 1-52.

* cited by examiner

U.S. Pat. No. 8,477,638, at 1 (highlighting added).

201. Alternatively, on information and belief, Cisco has had knowledge of the ’196 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the ’196 patent and knew of its infringement, including by way of this lawsuit.

202. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco ’196 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the ’196 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’196 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’196 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’196 patent, including at least claim 10, and Cisco further provides documentation and training materials that cause customers of the Cisco ’196 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’196 patent. By providing instruction and training to customers on how to use the Cisco ’196 Accused Products, Cisco specifically intended to induce infringement of the ’196 patent, including at least claim 10. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’196 Accused Products and to actively induce its

customers to infringe the ’196 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’196 patent, knowing that such use constitutes infringement of the ’196 patent.

203. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’196 patent.

204. As a result of Cisco’s infringement of the ’196 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

COUNT VII
INFRINGEMENT OF U.S. PATENT NO. 6,754,214

205. Dunti restates and incorporates by reference the preceding paragraphs of this Complaint as if fully set forth herein.

206. U.S. Patent No. 6,754,214 (“the ’214 patent”), entitled “Communication Network Having Packetized Security Codes and a System for Detecting Security Breach Locations Within the Network,” was filed on July 19, 1999. Dunti is the owner by assignment of the ’214 patent. A true and correct copy of the ’214 patent is attached hereto as Exhibit G.

207. The ’214 patent has been cited by at least one hundred and four United States patents and patent applications as relevant prior art. Specifically, patents issued to the following companies have cited the ’214 patent as relevant prior art:

- Bank of America Corporation;
- Sony Corporation
- Hitachi, Ltd.;
- Nokia Corporation;
- International Business Machines Corporation;
- Ntt Docomo, Inc.;
- Cisco Technology, Inc.
- EMC Corporation;
- AT&T Mobility II, LLC; and
- AT&T Intellectual Property I, L.P.

208. Cisco has repeatedly cited the ’214 patent in the issued patents and published patent applications of its wholly-owned subsidiary Cisco Technology, Inc. As of August 2, 2016, Cisco has cited the ’214 patent as prior art to at least nineteen issued Cisco patents and at least three published Cisco patent applications, enumerated below:

- U.S. Patent No. 7,669,244
- U.S. Patent No. 7,721,323
- U.S. Patent No. 7,827,402
- U.S. Patent No. 7,836,490
- U.S. Patent No. 7,840,708
- U.S. Patent No. 7,877,601
- U.S. Patent No. 7,877,796
- U.S. Patent No. 7,886,145
- U.S. Patent No. 7,954,163
- U.S. Patent No. 8,301,882
- U.S. Patent No. 8,302,157
- U.S. Patent No. 8,539,571
- U.S. Patent No. 8,555,056
- U.S. Patent No. 8,561,140
- U.S. Patent No. 8,621,596
- U.S. Patent No. 8,661,556
- U.S. Patent No. 8,713,201
- U.S. Patent No. 9,237,158
- U.S. Patent No. 9,407,604
- U.S. Patent Application Publication No. 2006/0117058
- U.S. Patent Application Publication No. 2011/0119753
- U.S. Patent Application Publication No. 2014/0269759

209. The ’214 patent was cited by the United States Patent and Trademark Office as relevant prior art in examination of several of the foregoing Cisco patents and patent applications, including at least Cisco’s U.S. Patent No. 7,721,323 (“Method and System for Including Network Security Information in a Frame”), U.S. Patent Application Publication No. 2006/0117058 (“Method and apparatus for ingress filtering using security group information”), U.S. Patent Application Publication No. 2011/0119753 (“Method and apparatus for best effort propagation of security group information”), and U.S. Patent Application Publication No.

2014/0269759 (“Schedule-based prioritization in contention-based shared-media computer networks”).

210. The ’214 patent discloses and claims a specific architecture and system for securing and prioritizing packets of data sent through a communication network. By assigning security and priority codes to packets as they enter the network through a specially configured entry module, maximum bandwidth allocation can be achieved among linked, configurable entry, exit, and intermediate modules in a hierarchical, packet-switched environment. Advantageously over prior art networks and conventional solutions, the communication network claimed in the ’214 patent dynamically ensures data path security, QoS-related packet forwarding priority in the presence of congestion at a shared network resource, and modular configurability. As described in the foregoing and in the paragraphs below, the ’214 patent claims a technical solution to a problem unique to computer networks.

211. In the communication network of the ’214 patent, distinct entry and exit modules are coupled in a structured, hierarchical network by linked, configurable intermediate modules, as shown below in Figure 6 from the ’214 patent. Modules 16a and 16b are the distinct entry and exit modules of the claimed communication network, while the intermediate ellipses represent the configurable nature of linked intermediate modules.

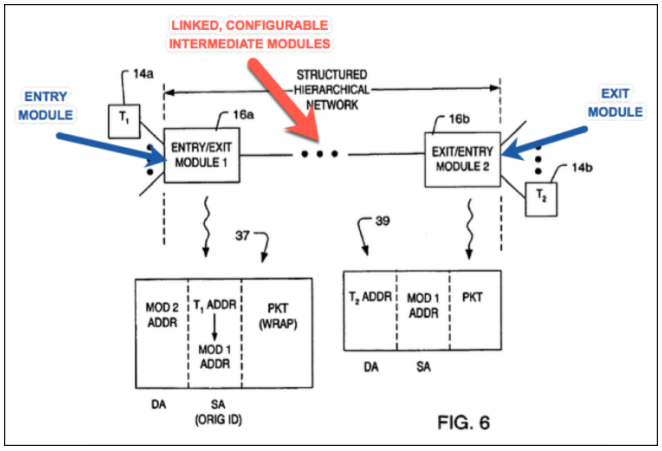

’214 Patent, Fig. 6 (blue and red illustrations added).

212. The entry module of the communication network claimed in the ’214 patent is coupled to assign and transfer a security code and an identification number to a packet of data. The exit module (through an included exit compare unit) is coupled to compare this security code and identification number before transferring the packet from the claimed communication network. As shown below, the green-boxed portion of Fig. 6 of the ’214 patent, as well as Figs. 11 and 16 of the ’214 patent, illustrate exemplary security code and identification number labels assigned, used, and compared by the configurable modules (including distinct entry and exit modules) of the communication network claimed in the ’214 patent.

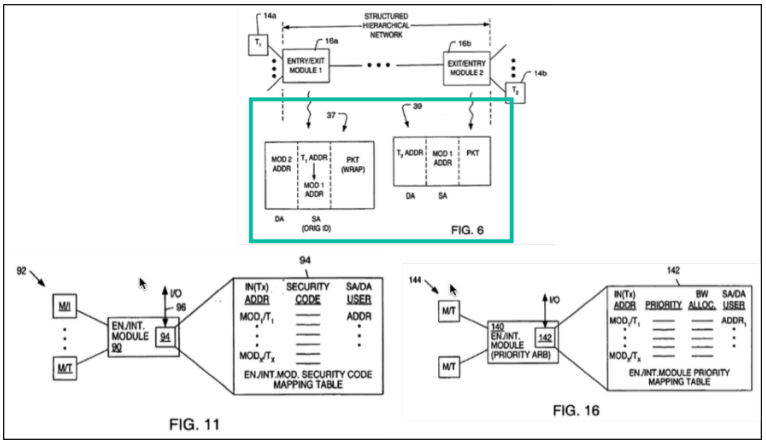

’214 Patent, Figs. 6, 11, and 16 (green illustration added).

213. Advantageously, the entry module claimed in the ’214 patent comprises a decoder, a specially-configured storage device, and a coupled entry compare unit configured to evaluate and route entering packets according to (among other things) the aforementioned security code and identification number.

214. Cisco makes, uses, sells, and/or offers for sale in the United States products and/or services relating to secure, prioritized network communications.

215. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco Nexus 7000 and Nexus 9000 series switches, including but not limited to the Cisco Nexus 93180YC-EX Switch, Cisco Nexus 93128TX Switch, the Cisco Nexus 93120TX Switch, the Cisco Nexus 93108TC-EX Switch, the Cisco Nexus 92304QC Switch, the Cisco Nexus 92160YC-X Switch, the Cisco

Nexus 9516 Switch, the Cisco Nexus 9508 Switch, the Cisco Nexus 9504 Switch, the Cisco Nexus 9396PX Switch, the Cisco Nexus 9396TX Switch, the Cisco Nexus 9372PX Switch, the Cisco Nexus 9372TX Switch, the Cisco Nexus 9336PQ ACI Spine Switch, the Cisco Nexus 9332PQ Switch, the Cisco Nexus 9272Q Switch, the Cisco Nexus 9236C Switch, the Cisco Nexus 7700 18-Slot Switch, the Cisco Nexus 7700 10-Slot Switch, the Cisco Nexus 7700 6-Slot Switch, the Cisco Nexus 7700 2-Slot Switch, the Cisco Nexus 7000 18-Slot Switch, the Cisco Nexus 7000 10-Slot Switch, the Cisco Nexus 7000 9-Slot Switch, and the Cisco Nexus 7000 4-Slot Switch (collectively, "the Cisco Nexus Accused Products").

216. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco MDS Family of products, including but not limited to the Cisco MDS 9718 Multilayer Director, the Cisco MDS 9710 Multilayer Director, the Cisco MDS 9706 Multilayer Director, the Cisco MDS 9513 Multilayer Director, the Cisco MDS 9506 Multilayer Director, the Cisco MDS 9250i Multiservice Fabric Switch, the Cisco MDS 9148S 16G Multilayer Fabric Switch, and the Cisco MDS 9124 Multilayer Fabric Switch (collectively, "the Cisco MDS Accused Products").

217. Cisco makes, uses, sells, offers to sell, and/or imports the Cisco NX-OS operating system, which runs on the Cisco Nexus Accused Products and the Cisco MDS Accused Products.

218. Cisco makes, uses, sells, and/or offers to sell the Cisco Nexus Accused Products, the Cisco MDS Accused Products, and the Cisco NX-OS operating system (collectively, "the Cisco '214 Accused Products").

219. Cisco makes, uses, sells, and/or offers to sell networks comprised of the Cisco '214 Accused Products ("a Cisco '214 Accused Product Network").

220. On information and belief, a Cisco '214 Accused Product Network implements at least Cisco's Virtual Extensible LAN (VXLAN) technology.

221. On information and belief, Cisco VXLAN is an overlay technology for network virtualization. It provides Layer-2 extension over a shared Layer-3 underlay infrastructure network by using MAC address in IP User Datagram Protocol (MAC in IP/UDP) tunneling encapsulation.

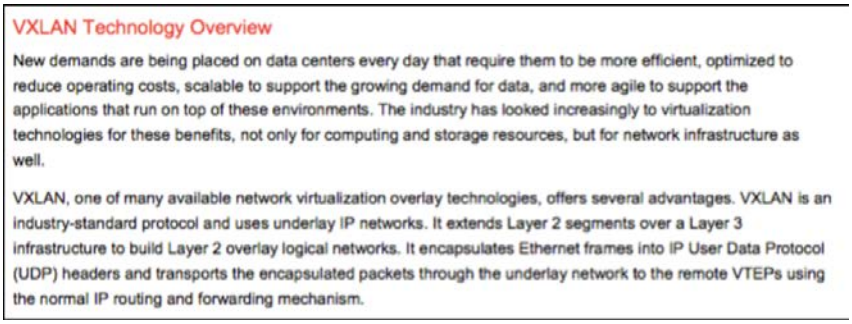

VXLAN Technology Overview

New demands are being placed on data centers every day that require them to be more efficient, optimized to reduce operating costs, scalable to support the growing demand for data, and more agile to support the applications that run on top of these environments. The industry has looked increasingly to virtualization technologies for these benefits, not only for computing and storage resources, but for network infrastructure as well.

VXLAN, one of many available network virtualization overlay technologies, offers several advantages. VXLAN is an industry-standard protocol and uses underlay IP networks. It extends Layer 2 segments over a Layer 3 infrastructure to build Layer 2 overlay logical networks. It encapsulates Ethernet frames into IP User Data Protocol (UDP) headers and transports the encapsulated packets through the underlay network to the remote VTEPs using the normal IP routing and forwarding mechanism.

VXLAN Design with Cisco Nexus 9300 Platform Switches Guide, CISCO TECHNICAL DOCUMENTATION (Oct. 2014), at 3.

222. On information and belief, Cisco's VXLAN technology differs from both conventional communication network technology and from VLAN network segmentation technology and offers at least the following benefits over VLAN network segmentation technology:

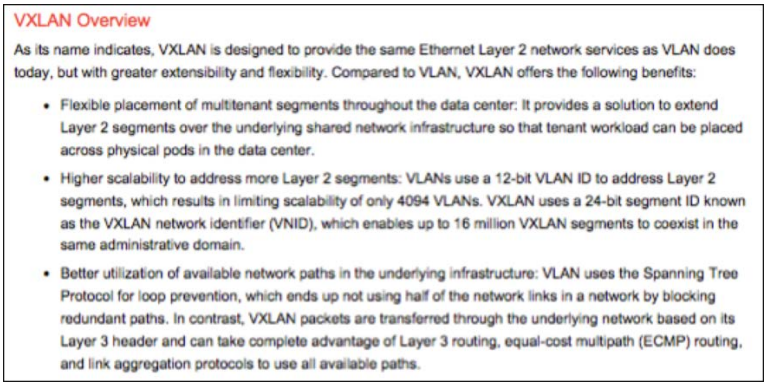

VXLAN Overview

As its name indicates, VXLAN is designed to provide the same Ethernet Layer 2 network services as VLAN does today, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- Flexible placement of multitenant segments throughout the data center: It provides a solution to extend Layer 2 segments over the underlying shared network infrastructure so that tenant workload can be placed across physical pods in the data center.
- Higher scalability to address more Layer 2 segments: VLANs use a 12-bit VLAN ID to address Layer 2 segments, which results in limiting scalability of only 4094 VLANs. VXLAN uses a 24-bit segment ID known as the VXLAN network identifier (VNID), which enables up to 16 million VXLAN segments to coexist in the same administrative domain.
- Better utilization of available network paths in the underlying infrastructure: VLAN uses the Spanning Tree Protocol for loop prevention, which ends up not using half of the network links in a network by blocking redundant paths. In contrast, VXLAN packets are transferred through the underlying network based on its Layer 3 header and can take complete advantage of Layer 3 routing, equal-cost multipath (ECMP) routing, and link aggregation protocols to use all available paths.

VXLAN Overview: Cisco Nexus 9000 Series Switches, CISCO TECHNICAL DOCUMENTATION (2015), at 2.

223. On information and belief, the structured, hierarchical MAC-in-UDP encapsulated packet structure and Layer 2-over-Layer 3 network tunneling technology defined and used in Cisco VXLAN enables secure, prioritized packet forwarding over dynamic, configurable networks through Cisco's integration of technologies such MPLS-IP VPN, Segment Routing, and MP-BGP-EVPN into the base VXLAN network architecture. On information and belief, a Cisco '214 Accused Product Network with VXLAN implements each of the aforementioned technologies, including Cisco's Segment Routing-based MPLS-IP VPN ("SR-MPLS-IP VPN").

224. On information and belief, a Cisco ’214 Accused Product Network comprises a communication network. For example, a Cisco ’214 Accused Product Network comprises a packetized communication network that securely routes packets in a dynamic, configurable, QoS-aware manner using specially configured, SR-MPLS-IP-aware VXLAN entry, exit, and intermediate forwarding modules.

225. On information and belief, a Cisco ’214 Accused Product Network comprises an entry module; an exit module; and at least one intermediate module.

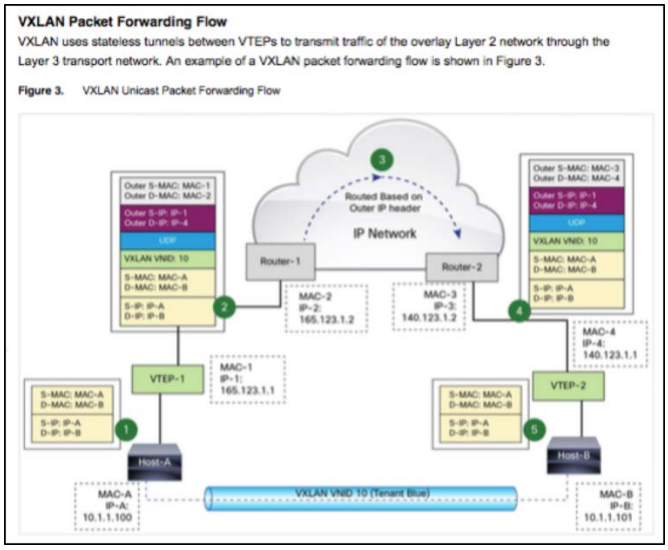

VXLAN Overview: Cisco Nexus 9000 Series Switches, CISCO TECHNICAL DOCUMENTATION (2015), at 4.

226. On information and belief, a Cisco ’214 Accused Product Network entry module is coupled to assign and transfer a security code and an identification number to a packet of data. For example, the entry VTEP/SRTE headend/MPLS VPN PE-configured Cisco 9000 series platform router is coupled to assign VXLAN/SR-MPLS-IP/MPLS VPN security and priority labels within the MAC-in-UDP encapsulated header that the VTEP/STRE headend/MPLS VPN PE entry module assigns to a data packet entering the Cisco ’214 Accused Product Network.

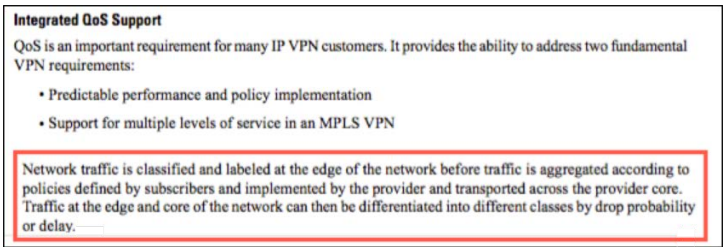

Integrated QoS Support

QoS is an important requirement for many IP VPN customers. It provides the ability to address two fundamental VPN requirements:

- Predictable performance and policy implementation
- Support for multiple levels of service in an MPLS VPN

Network traffic is classified and labeled at the edge of the network before traffic is aggregated according to policies defined by subscribers and implemented by the provider and transported across the provider core. Traffic at the edge and core of the network can then be differentiated into different classes by drop probability or delay.

MPLS: Layer 3 VPNs Configuration Guide, CISCO TECHNICAL DOCUMENTATION (Mar. 15, 2013), at 8-9.

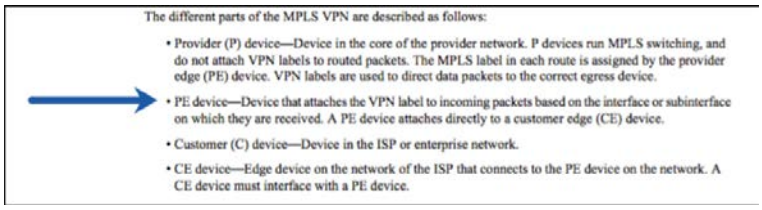

The different parts of the MPLS VPN are described as follows:

- Provider (P) device—Device in the core of the provider network. P devices run MPLS switching, and do not attach VPN labels to routed packets. The MPLS label in each route is assigned by the provider edge (PE) device. VPN labels are used to direct data packets to the correct egress device.
- PE device—Device that attaches the VPN label to incoming packets based on the interface or subinterface on which they are received. A PE device attaches directly to a customer edge (CE) device.
- Customer (C) device—Device in the ISP or enterprise network.
- CE device—Edge device on the network of the ISP that connects to the PE device on the network. A CE device must interface with a PE device.

Id. at 4.

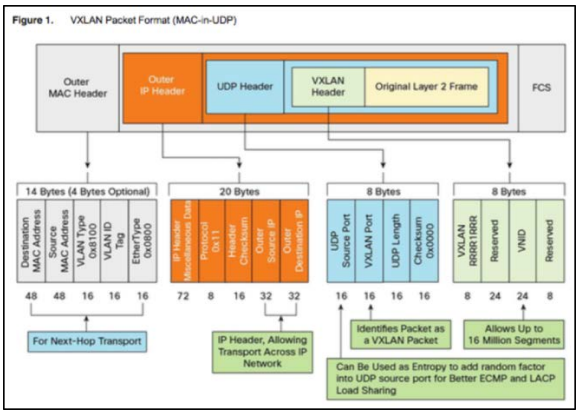

VXLAN Design with Cisco Nexus 9300 Platform Switches Guide, CISCO TECHNICAL DOCUMENTATION (Oct. 2014), at 4.

227. On information and belief, a Cisco '214 Accused Product Network entry module comprises a decoder (e.g., a VXLAN/SR-MPLS-IP/MPLS VPN-aware hardware and/or software decoder) and a storage device configured with a set of bits (e.g., a Nexus hardware and/or NX-OS software configuration register storing base and/or context-specific VXLAN/SR-MPLS-IP/MPLS VPN settings and identifiers for the entry VTEP/SRTE/PE node).

228. On information and belief, an entry compare unit (e.g., a VTEP/SRTE entry module) is coupled between the decoder and storage device within a Cisco '214 Accused Product Network entry module to facilitate MPLS VPN, MPLS-IP, and VXLAN-aware tunneling and

traffic management through comparison of VXLAN/SR-MPLS-IP/MPLS VPN identifiers in the UDP-in-MAC packet header with the stored base and/or context-specific VXLAN/SR-MPLS-IP/MPLS VPN settings and identifiers for the entry VTEP/SRTE/MPLS VPN node.

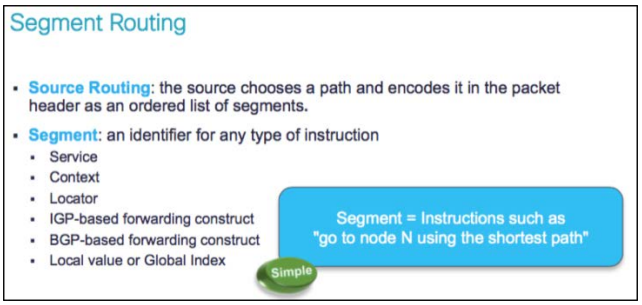

Segment Routing: Key Concepts, CISCO TECHNICAL DOCUMENTATION (2014), at 3.

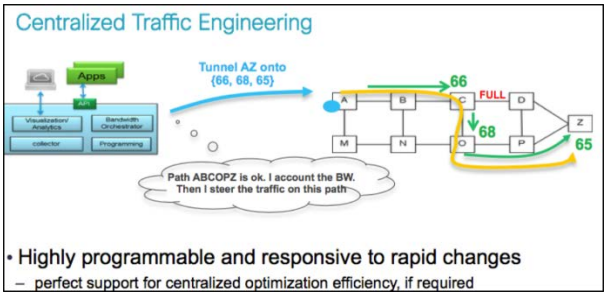

Id. at 22.

229. On information and belief, secure traffic management and packet tunneling in a Cisco ’214 Accused Product Network is achieved through (among other things) packet classification on the basis of security.

230. On information and belief, such security-based classification uses at least the identification number and security code described in the foregoing paragraphs, as well as trailer bit groupings identifying a count number for intermediate modules traversed by a packet and an identification of traversed intermediate modules.

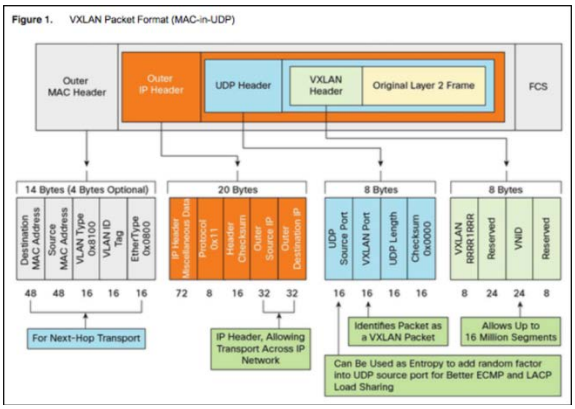

VXLAN Design with Cisco Nexus 9300 Platform Switches Guide, CISCO TECHNICAL DOCUMENTATION (Oct. 2014), at 4.

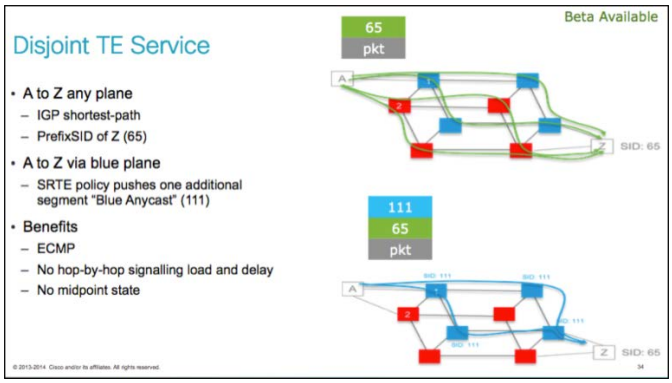

Segment Routing: Key Concepts, CISCO TECHNICAL DOCUMENTATION (2014), at 33.

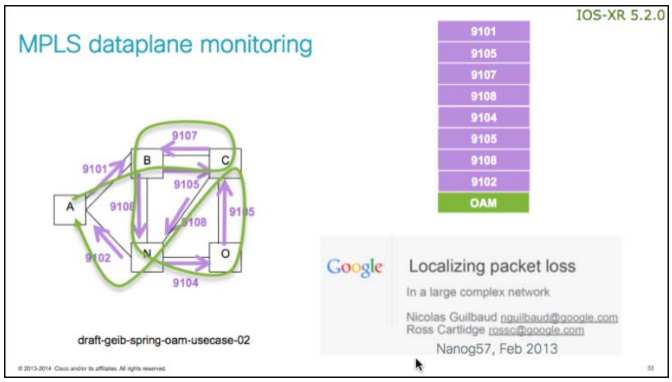

Segment Routing: Key Concepts, CISCO TECHNICAL DOCUMENTATION (2014), at 34.

231. On information and belief, a Cisco ’214 Accused Product Network exit module comprises an exit compare unit coupled to compare the security code and identification number before transferring the packet of data from the Cisco ’214 Accused Product Network. For example, the exit VTEP/SRTE-configured Cisco 9000 series platform router is coupled to

compare the aforementioned VXLAN and SR-MPLS-IP security and priority labels within the MAC-in-UDP encapsulated header that the VTEP/STRE headend entry module assigned to the data packet to a (at least temporarily) stored value before transferring the packet of data from the Cisco VXLAN communication network.

232. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Cisco ’214 Accused Products, Cisco has injured Dunti and is liable to Dunti for directly infringing one or more claims of the ’214 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

233. On information and belief, Cisco also indirectly infringes the ’214 patent by actively inducing infringement under 35 U.S.C. § 271(b).

234. On information and belief, Cisco has had knowledge of the ’214 patent since at least March 10, 2009, when Cisco filed a response to a December 10, 2008 Office Action by the USPTO that cited the ’214 patent as relevant prior art to Cisco’s U.S. Patent Application No. 10/996,101 (now Cisco’s U.S. Patent No. 7,721,323), as shown below:

Notice of References Cited	Application/Control No. 10/996,101	Applicant(s)/Patent Under Reexamination FINN ET AL.	
	Examiner AMARE TABOR	Art Unit 2439	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,754,214	06-2004	Mahalingaiah, Rupaka	370/392
*	B	US-2003/0154400	08-2003	Pirttimaa et al.	713/201
	C	US-			
	D	US-			

235. Additionally, as noted earlier in this Count, Cisco has affirmatively cited the ’214 patent as relevant prior art to approximately two dozen distinct Cisco patent applications, stretching across several years, multiple Cisco patent families, and numerous Cisco inventors and attorneys resident in multiple locations, including Texas.

236. Alternatively, on information and belief, Cisco has had knowledge of the ’214 patent since at least the date of service of this Complaint or shortly thereafter, and on information and belief, Cisco knew of the ’214 patent and knew of its infringement, including by way of this lawsuit.

237. On information and belief, Cisco intended to induce patent infringement by third-party customers and users of the Cisco ’214 Accused Products and had knowledge that the inducing acts would cause infringement or was willfully blind to the possibility that its inducing acts would cause infringement. Cisco specifically intended and was aware that the normal and customary use of the accused products would infringe the ’214 patent. Cisco performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’214 patent and with the knowledge that the induced acts would constitute infringement. For example, Cisco provides the Cisco ’214 Accused Products, which are capable of operating in a manner that infringes one or more claims of the ’214 patent, including at least claim 1, and Cisco further provides documentation and training materials that cause customers of the Cisco ’214 Accused Products to utilize the products and services in a manner that directly infringes one or more claims of the ’214 patent. By providing instruction and training to customers on how to use the Cisco ’214 Accused Products, Cisco specifically intended to induce infringement of the ’214 patent, including at least claim 1. On information and belief, Cisco engaged in such inducement to promote the sales of the Cisco ’214 Accused Products and to actively induce its customers to infringe the ’214 patent. Accordingly, Cisco has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’214 patent, knowing that such use constitutes infringement of the ’214 patent.

238. To the extent applicable, the requirements of 35 U.S.C. § 287(a) have been met with respect to the ’214 patent.

239. As a result of Cisco’s infringement of the ’214 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Cisco’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Cisco together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff Dunti respectfully requests that this Court enter:

A. A judgment in favor of Plaintiff Dunti that Cisco has infringed, either literally and/or under the doctrine of equivalents, the ’462 patent, the ’701 patent, the ’235 patent, the ’286 patent, the ’259 patent, the ’196 patent, and/or the ’214 patent;

B. An award of damages resulting from Cisco’s acts of infringement in accordance with 35 U.S.C. § 284;

C. A judgment and order requiring Cisco to provide accountings and to pay supplemental damages to Dunti, including, without limitation, prejudgment and post-judgment interest; and

D. Any and all other relief to which Dunti may show itself to be entitled.

JURY TRIAL DEMANDED

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, Dunti requests a trial by jury of any issues so triable by right.

Dated: September 19, 2016

Respectfully submitted,

/s/ Matt Olavi_____

Elizabeth L. DeRieux (TX Bar No. 05770585)
D. Jeffrey Rambin (TX Bar No. 00791478)
CAPSHAW DERIEUX, LLP
114 E. Commerce Ave.
Gladewater, Texas 75647
Telephone: 903-845-5770
E-mail: ederieux@capshawlaw.com
E-mail: jrambin@capshawlaw.com

Matt Olavi (TX Bar No. 24095777)
Brian J. Dunne (CA SB No. 275689)
Douglas W. Meier (TX Bar No. 24100889)
OLAVI DUNNE LLP
816 Congress Ave., Ste. 1620
Austin, Texas 78701
Telephone: 512-717-4485
Facsimile: 512-717-4495
E-mail: molavi@olavidunne.com
E-mail: bdunne@olavidunne.com
E-mail: dmeier@olavidunne.com

Attorneys for Dunti Network Technologies, LLC