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

| | |
|---|---|
| **DUNTI NETWORK TECHNOLOGIES, LLC,**<br><br>***Plaintiff,***<br><br>**v.**<br><br>**INTEL CORPORATION,**<br><br>***Defendant.*** | **Civil Action No._____**<br><br>**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 Intel Corporation ("Intel") have adopted Dunti's patented technologies—developed more than a decade ago right here in Texas—*en masse*. Intel has profited handsomely from its use of Dunti's patented inventions, and Dunti deserves to be compensated for this use. But Intel 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"); and 7,778,259 ("the '259 patent") (collectively, "the patents-in-suit"), is brought to ensure that Intel 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.

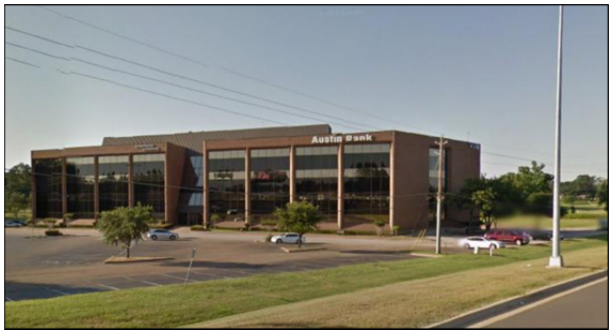

2. Dunti is a small, Texas-based company. Dunti depends on patent protection to effectively license its innovative technologies and build its business. Like Defendant Intel, Dunti relies on its intellectual property.

3. On information and belief, Defendant Intel is a Delaware corporation with its principal office at 2200 Mission College Blvd., Santa Clara, California 95054. Intel Corporation can be served through its registered agent, C T Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201-3136.

4. On information and belief, Defendant Intel has approximately 2,300 employees in Texas and maintains substantial offices in both Austin, Texas and Plano, Texas where it engages in sales, marketing, and research and development of its System on a Chip and microprocessor components.[1]

5. On information and belief, and according to Intel's website, Intel offers infringing products for sale throughout the United States and Canada, including in the Eastern District of Texas. Further, Intel 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] Intel in Texas, Intel Webpage, available at http://www.intel.com/content/www/us/en/corporate-responsibility/intel-in-texas.html (accessed Sept. 6, 2016).

7. Upon information and belief, this Court has personal jurisdiction over Intel in this action because Intel 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 Intel would not offend traditional notions of fair play and substantial justice. Defendant Intel, 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, Intel maintains substantial offices in Austin, Texas, and Plano, Texas, it is registered to do business in the State of Texas, and it has appointed C T Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201-3136, 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 Intel is registered to do business in Texas, has two substantial 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 five 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.

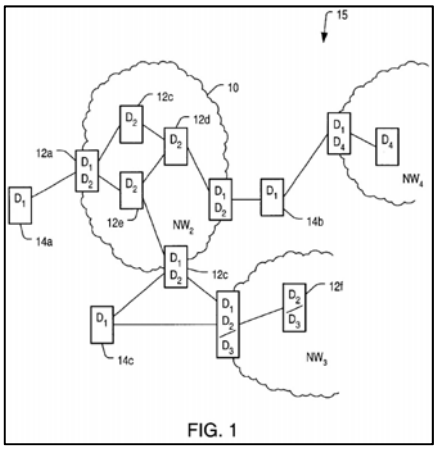


'462 Patent, Fig. 1.

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.

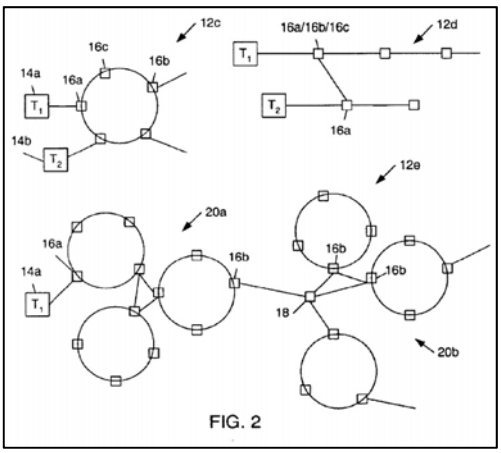


'286 Patent, Fig. 2.

42. 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.

## INTEL'S INFRINGING PRODUCTS AND SERVICES

43. On information and belief, Intel offers high-performance Ethernet Switch products that are designed for use in high performance "data center and communications infrastructure" applications. *Intel Ethernet Switch Silicon*, INTEL.COM, available at http://www.intel.com/content/www/us/en/ethernet-products/switch-silicon/overview.html (accessed Sept. 7, 2016).

44. On information and belief, Intel's Ethernet Switch products, such as at least the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet Switch, and the Intel FM6700 Ethernet Switch, implement the Transparent Interconnection of Lots of Links ("TRILL") protocol and can be used to build a TRILL network within a data center.

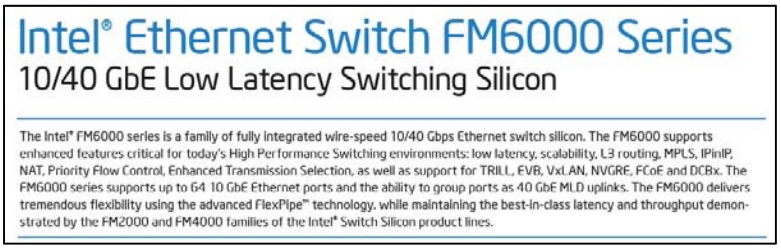
Intel® Ethernet Switch FM6000 Series

10/40 GbE Low Latency Switching Silicon

The Intel® FM6000 series is a family of fully integrated wire-speed 10/40 Gbps Ethernet switch silicon. The FM6000 supports enhanced features critical for today's High Performance Switching environments: low latency, scalability, L3 routing, MPLS, IPinIP, NAT, Priority Flow Control, Enhanced Transmission Selection, as well as support for TRILL, FVB, VxLAN, NVGRE, FCoE and DCBx. The FM6000 series supports up to 64 10 GbE Ethernet ports and the ability to group ports as 40 GbE MLD uplinks. The FM6000 delivers tremendous flexibility using the advanced FlexPipe™ technology, while maintaining the best-in-class latency and throughput demonstrated by the FM2000 and FM4000 families of the Intel® Switch Silicon product lines.

*Intel Ethernet Switch FM6000 Series*, INTEL PRODUCT BRIEF, at 1 (2013).

45. On information and belief, a TRILL network is implemented by interconnecting multiple TRILL-enabled switches (commonly referred to as RBridges), with edge devices connecting the TRILL network to an external network(s). When a data packet enters a TRILL network, an ingress RBridge encapsulates the packet within a TRILL header. Transit RBridges within a TRILL network then transport the packet across the network until it arrives at its

destination edge device, or egress RBridge, which strips off the headers and forwards the packet outside the TRILL network to its intended final destination.

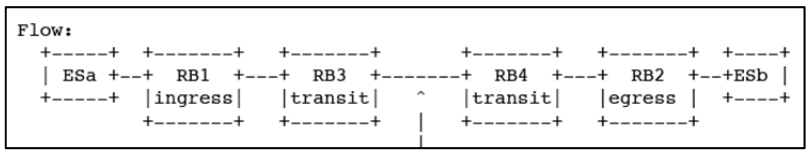


Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011).

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

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

47. 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.

48. 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.

49. 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.

50. 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.

51. 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.

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

53. Intel makes, uses, sells, offers to sell, and/or imports the Intel Ethernet switches, including but not limited to the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet Switch, and the Intel FM6700 Ethernet Switch (collectively, "the Intel Ethernet Switches" or "the Intel '462 Accused Products").

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

55. On information and belief, an Intel '462 Accused Product Network implements at least the TRILL protocol.

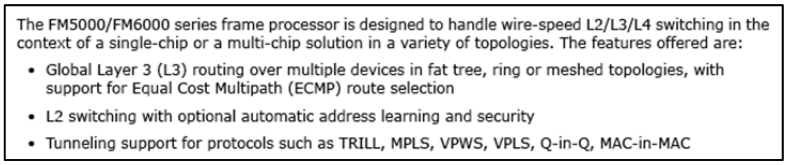
The FM5000/FM6000 series frame processor is designed to handle wire-speed L2/L3/L4 switching in the context of a single-chip or a multi-chip solution in a variety of topologies. The features offered are:

- Global Layer 3 (L3) routing over multiple devices in fat tree, ring or meshed topologies, with support for Equal Cost Multipath (ECMP) route selection
- L2 switching with optional automatic address learning and security
- Tunneling support for protocols such as TRILL, MPLS, VPWS, VPLS, Q-in-Q, MAC-in-MAC

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 43 (November 2014).

56. On information and belief, an Intel '462 Accused Product Network comprises a communication system.

57. On information and belief, an Intel '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 figure below shows an entry end module (i.e., ingress RBridge) and an exit end module (i.e., egress RBridge) at the edges of a TRILL network and multiple intermediate modules (i.e., transit RBridges) coupled between the entry and exit end modules.

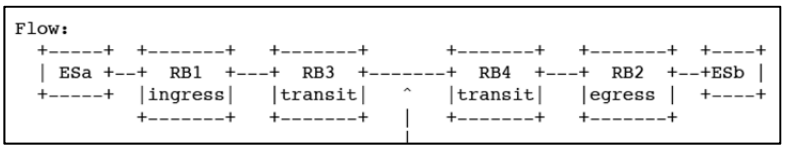


Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011).

58. On information and belief, an Intel '462 Accused Product Network comprises a first addressing domain for identifying each of the end modules and the intermediate module. For example, each RBridge within a TRILL network is assigned a unique RBridge Nickname.

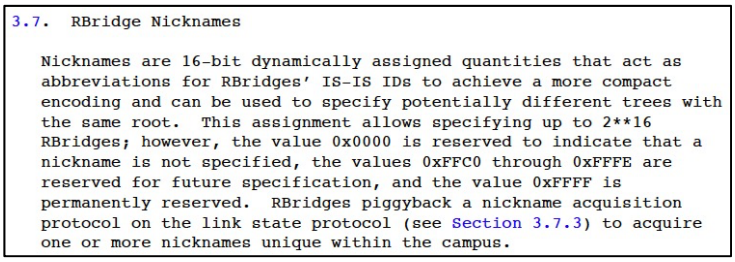
3.7. RBridge Nicknames

Nicknames are 16-bit dynamically assigned quantities that act as abbreviations for RBridges' IS-IS IDs to achieve a more compact encoding and can be used to specify potentially different trees with the same root. This assignment allows specifying up to 2**16 RBridges; however, the value 0x0000 is reserved to indicate that a nickname is not specified, the values 0xFFC0 through 0xFFFE are reserved for future specification, and the value 0xFFFF is permanently reserved. RBridges piggyback a nickname acquisition protocol on the link state protocol (see Section 3.7.3) to acquire one or more nicknames unique within the campus.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 23 (July 2011).

2. The following information in connection with the nickname or each of the nicknames of RBridge RBn:

   2.1. The nickname value (2 octets).

   2.2. The unsigned 8-bit priority for RBn to have that nickname (see Section 3.7.3).

   2.3. The 16-bit unsigned priority of that nickname to becoming a distribution tree root.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 38 (July 2011).

59. On information and belief, an Intel ’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 an Intel ’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 TRILL network.

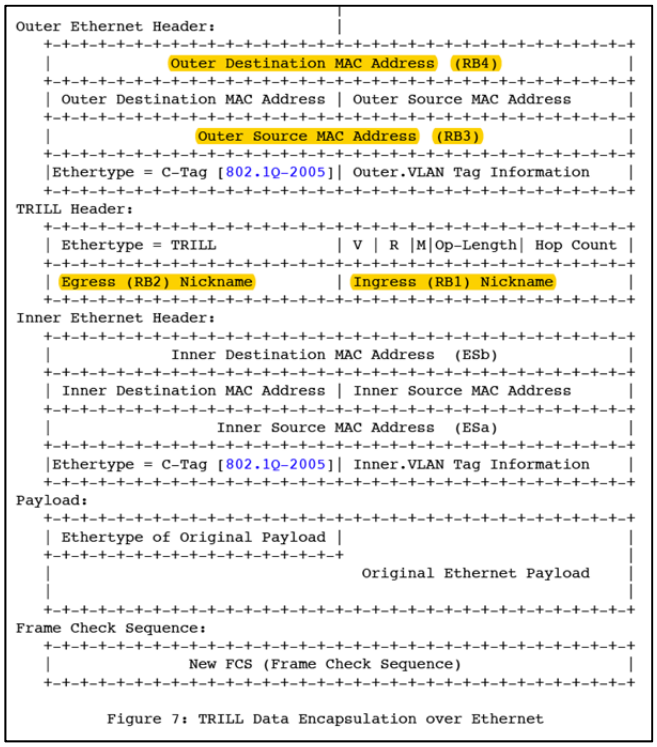


Figure 7: TRILL Data Encapsulation over Ethernet

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011) (highlighting added).

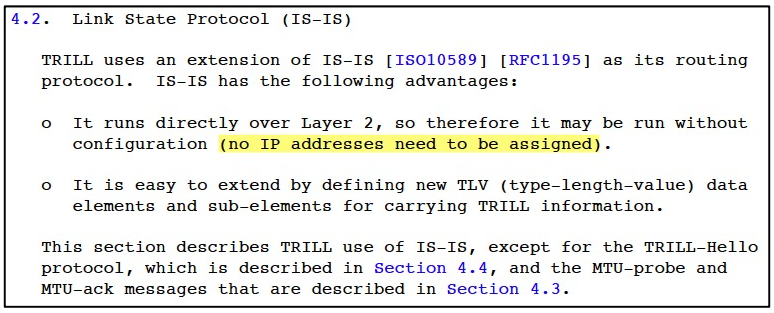

4.2. Link State Protocol (IS-IS)

TRILL uses an extension of IS-IS [ISO10589] [RFC1195] as its routing protocol. IS-IS has the following advantages:

- o It runs directly over Layer 2, so therefore it may be run without configuration (no IP addresses need to be assigned).
- o It is easy to extend by defining new TLV (type-length-value) data elements and sub-elements for carrying TRILL information.

This section describes TRILL use of IS-IS, except for the TRILL-Hello protocol, which is described in Section 4.4, and the MTU-probe and MTU-ack messages that are described in Section 4.3.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 32 (July 2011) (highlighting added).

60. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Intel ’462 Accused Products, Intel 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).

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

62. On information and belief, Intel 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, Intel knew of the ’462 patent and knew of its infringement, including by way of this lawsuit.

63. On information and belief, Intel intended to induce patent infringement by third-party customers and users of the Intel ’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. Intel specifically intended and was aware that the normal and customary use of the accused products would infringe the ’462 patent. Intel 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, Intel provides the Intel ’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 Intel further provides documentation and training materials that cause customers of the Intel ’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 Intel ’462 Accused Products, Intel specifically intended to induce infringement of the ’462 patent, including at least claim 1. On information and belief, Intel engaged in such inducement to promote the sales of the Intel ’462 Accused Products and to actively induce its customers to infringe the ’462 patent. Accordingly, Intel 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.

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

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

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

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

67. 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.

68. 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.

69. 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).

70. 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.

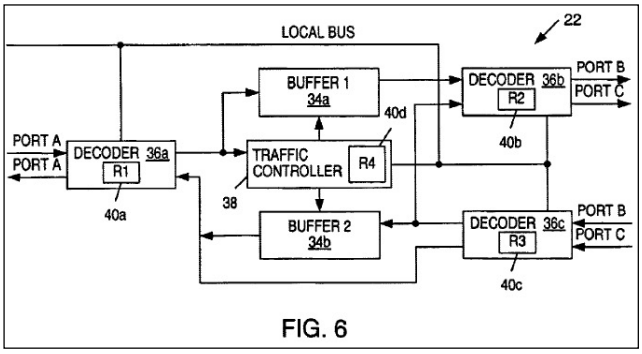

FIG. 6

’701 Patent, Fig. 6.

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

72. Intel makes, uses, sells, offers to sell, and/or imports the Intel Ethernet switches, including but not limited to the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet Switch, and the Intel FM6700 Ethernet Switch (collectively, "the Intel Ethernet Switches" or "the Intel '701 Accused Products").

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

74. On information and belief, an Intel '701 Accused Product Network implements at least the TRILL protocol.

> The FM5000/FM6000 series frame processor is designed to handle wire-speed L2/L3/L4 switching in the context of a single-chip or a multi-chip solution in a variety of topologies. The features offered are:
>
> - Global Layer 3 (L3) routing over multiple devices in fat tree, ring or meshed topologies, with support for Equal Cost Multipath (ECMP) route selection
> - L2 switching with optional automatic address learning and security
> - Tunneling support for protocols such as TRILL, MPLS, VPWS, VPLS, Q-in-Q, MAC-in-MAC

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 43 (November 2014).

75. On information and belief, the Intel '701 Accused Products comprise a switch.

> The FM5000/FM6000 operates as a one-armed IP router combined with an Ethernet L2 switch (see Figure 5-1). In this architecture, incoming packets are first associated with a VLAN (using the VTAG tag if present, or by associating a default VLAN), and are then either switched within their respective VLANs or routed across VLANs or both. The decision to switch, route or drop depends on tables stored in the FFU (which includes a large ternary CAM to store IP route entries and access control lists), other tables located in the switch (such as MAC Address table) or tables located in the router (such as ARP table).

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 42 (November 2014) (highlighting added).

76. On information and belief, the Intel '701 Accused Products within a TRILL network comprise a traffic manager which dispatches a series of read operations to a memory coupled within a data flow path. For example, the Intel '701 Accused Products include memory and at least one processor.

### 2.1.3 Packet Memory

The main memory stores incoming packets from ingress EPLs and forwards them to egress EPLs upon request from the scheduler. The main memory size is 7.5 MB. Additional memory in other functional blocks brings the total packet and frame header memory to approximately 9.5 MB.

### 2.1.4 Frame Handler

The frame handler makes forwarding decisions based on the frame header received from the EPL. The forwarding information (port mask and modification data) is sent to the scheduler.

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 18 (November 2014).

77. On information and belief, the Intel '701 Accused Products within a TRILL network include a TRILL unicast forwarding table comprised in memory, which includes a source address and a destination address of a pair of network nodes routably coupled within the data flow path.

> The design supports VLANs and the optimization of the distribution of multi-destination frames based on VLAN ID and based on IP-derived multicast groups. It also allows unicast forwarding tables at transit RBridges to be sized according to the number of RBridges (rather than the number of end nodes), which allows their forwarding tables to be substantially smaller than in conventional customer bridges.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 1 (July 2011) (highlighting added).

The FM5000/FM6000 operates as a one-armed IP router combined with an Ethernet L2 switch (see Figure 5-1). In this architecture, incoming packets are first associated with a VLAN (using the VTAG tag if present, or by associating a default VLAN), and are then either switched within their respective VLANs or routed across VLANs or both. The decision to switch, route or drop depends on tables stored in the FFU (which includes a large ternary CAM to store IP route entries and access control lists), other tables located in the switch (such as MAC Address table) or tables located in the router (such as ARP table).

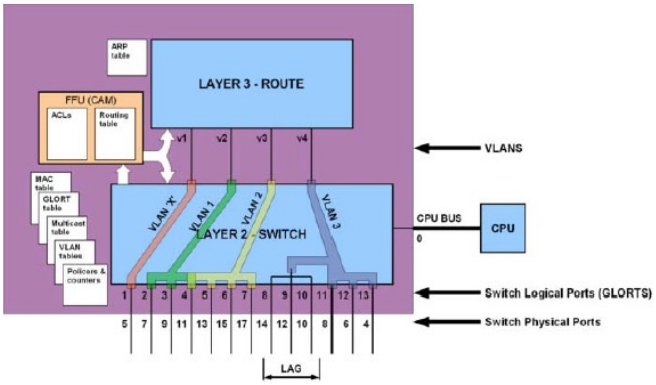


Figure 5-1 FM5000/FM6000 Switch/Router Concept

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 42 (November 2014).

78. On information and belief, the memory in the Intel ’701 Accused Products comprise an input port and an output port.

The FM5000/FM6000 main components are shown in Figure 2-1.

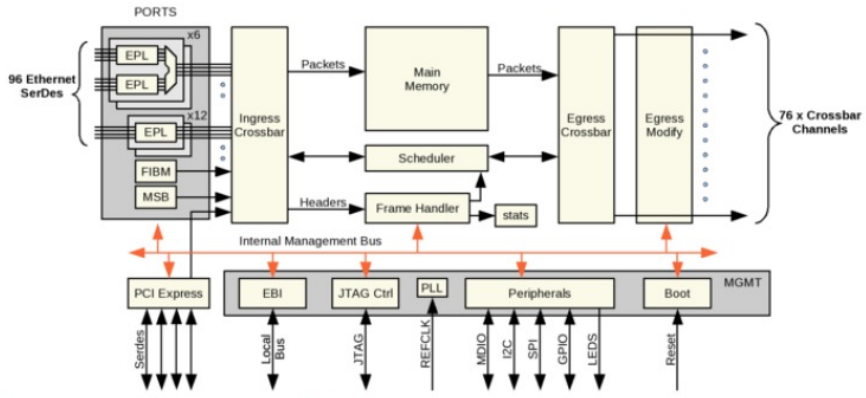


**Figure 2-1 FM5000/FM6000 Block Diagram**

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 17 (November 2014).

79. On information and belief, the memory in the Intel ’701 Accused Products comprises packets of data dispatched from the input port. For example, the Intel ’701 Accused Products encapsulate incoming data packets within a TRILL header. The incoming data packets are comprised in memory within an ingress RBridge as they are encapsulated within a TRILL header as forwarding decisions are made.

### 2.1.3 Packet Memory

The main memory stores incoming packets from ingress EPLs and forwards them to egress EPLs upon request from the scheduler. The main memory size is 7.5 MB. Additional memory in other functional blocks brings the total packet and frame header memory to approximately 9.5 MB.

### 2.1.4 Frame Handler

The frame handler makes forwarding decisions based on the frame header received from the EPL. The forwarding information (port mask and modification data) is sent to the scheduler.

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 18 (November 2014).

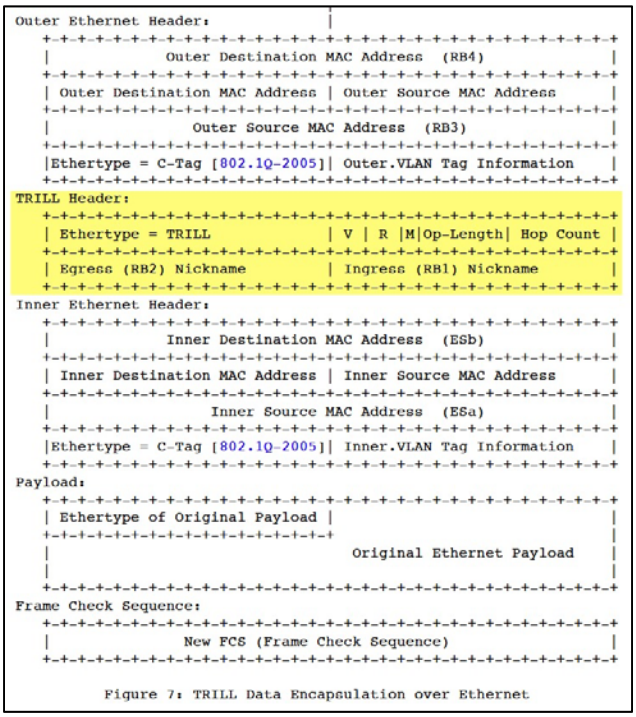


Figure 7: TRILL Data Encapsulation over Ethernet

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011) (highlighting added).

80. On information and belief, the Intel ’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, the TRILL header includes an Egress RBridge Nickname field and the Outer Ethernet Header includes an Outer Destination MAC Address field, both of which comprise a destination address. The Egress RBridge Nickname is decoded as forwarding decisions are made.

> These frames have a unicast inner MAC destination address (Inner.MacDA) and are those for which the ingress RBridge knows the egress RBridge for the destination MAC address in the frame’s VLAN.
>
> Such frames are forwarded Rbridge hop by Rbridge hop to their egress Rbridge.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 16 (July 2011) (highlighting added).

81. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Intel ’701 Accused Products, Intel

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).

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

83. On information and belief, Intel 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, Intel knew of the ’701 patent and knew of its infringement, including by way of this lawsuit.

84. On information and belief, Intel intended to induce patent infringement by third-party customers and users of the Intel ’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. Intel specifically intended and was aware that the normal and customary use of the accused products would infringe the ’701 patent. Intel 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, Intel provides the Intel ’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 Intel further provides documentation and training materials that cause customers of the Intel ’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 Intel ’701 Accused Products, Intel specifically intended to induce infringement of the ’701 patent, including at least claim 1. On information and belief, Intel engaged in such inducement to promote the sales of the Intel ’701 Accused Products and to actively induce its customers to infringe the ’701 patent. Accordingly, Intel 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.

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

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

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

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

88. 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.

89. 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.

90. 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.

91. 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.

92. 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.

93. 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.

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

95. Intel makes, uses, sells, offers to sell, and/or imports the Intel Ethernet switches, including but not limited to the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet

Switch, and the Intel FM6700 Ethernet Switch (collectively, "the Intel Ethernet Switches" or "the Intel '235 Accused Products").

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

97. On information and belief, an Intel '235 Accused Product Network implements at least the TRILL protocol.

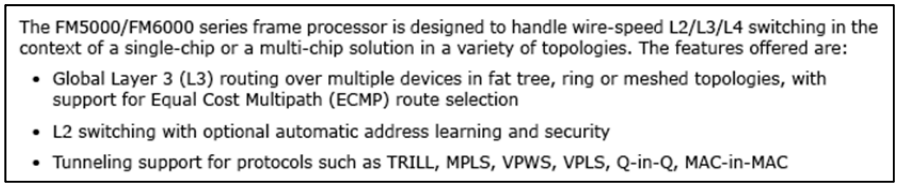

The FM5000/FM6000 series frame processor is designed to handle wire-speed L2/L3/L4 switching in the context of a single-chip or a multi-chip solution in a variety of topologies. The features offered are:

- Global Layer 3 (L3) routing over multiple devices in fat tree, ring or meshed topologies, with support for Equal Cost Multipath (ECMP) route selection
- L2 switching with optional automatic address learning and security
- Tunneling support for protocols such as TRILL, MPLS, VPWS, VPLS, Q-in-Q, MAC-in-MAC

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 43 (November 2014).

98. On information and belief, an Intel '235 Accused Product Network comprises a communication network.

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

100. On information and belief, modules within an Intel '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 RBridge within a TRILL network is assigned a unique RBridge Nickname, which is a unique identification number that is independent of the MAC address, and can be assigned to different topologies within the TRILL network.

2. The following information in connection with the nickname or each of the nicknames of RBridge RBn:

   2.1. The nickname value (2 octets).

   2.2. The unsigned 8-bit priority for RBn to have that nickname (see Section 3.7.3).

   2.3. The 16-bit unsigned priority of that nickname to becoming a distribution tree root.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 38 (July 2011).

101. On information and belief, the first model layer used in an Intel '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 TRILL network, which already include headers from higher layers, are further wrapped/encapsulated within a TRILL header that includes the RBridge Nickname of the egress RBridge.

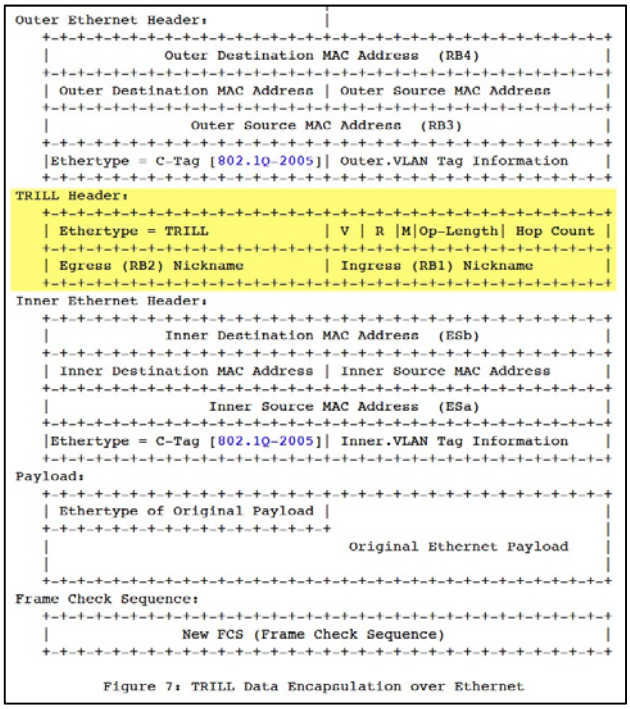


Figure 7: TRILL Data Encapsulation over Ethernet

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011) (highlighting added).

102. On information and belief, the second model layer used in an Intel '235 Accused Product Network is higher than a physical layer of the OSI reference model. For example, the

edge switches in an Intel ’235 Accused Product Network can use IP addresses to route data packets outside of a TRILL network, and the IP address layer is higher than a physical layer of the OSI model.

The FM5000/FM6000 operates as a one-armed IP router combined with an Ethernet L2 switch (see Figure 5-1). In this architecture, incoming packets are first associated with a VLAN (using the VTAG tag if present, or by associating a default VLAN), and are then either switched within their respective VLANs or routed across VLANs or both. The decision to switch, route or drop depends on tables stored in the FFU (which includes a large ternary CAM to store IP route entries and access control lists), other tables located in the switch (such as MAC Address table) or tables located in the router (such as ARP table).

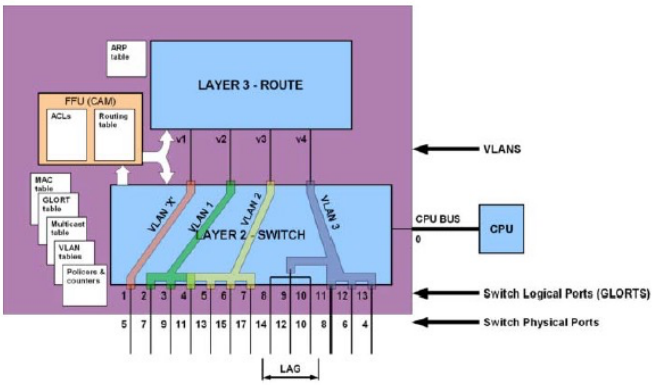


Figure 5-1 FM5000/FM6000 Switch/Router Concept

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 42 (November 2014).

103. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Intel ’235 Accused Products, Intel 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).

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

105. On information and belief, Intel 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, Intel knew of the ’235 patent and knew of its infringement, including by way of this lawsuit.

106. On information and belief, Intel intended to induce patent infringement by third-party customers and users of the Intel ’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. Intel specifically intended and was aware that the normal and customary use of the accused products would infringe the ’235 patent. Intel 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, Intel provides the Intel ’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 Intel further provides documentation and training materials that cause customers of the Intel ’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 Intel ’235 Accused Products, Intel specifically intended to induce infringement of the ’235 patent, including at least claim 1. On information and belief, Intel engaged in such inducement to promote the sales of the Intel ’235 Accused Products and to actively induce its customers to infringe the ’235 patent. Accordingly, Intel 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.

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

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

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

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

110. 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.

111. 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.

112. 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.

113. 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.

114. 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.

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

116. Intel makes, uses, sells, offers to sell, and/or imports the Intel Ethernet switches, including but not limited to the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet Switch, and the Intel FM6700 Ethernet Switch (collectively, "the Intel Ethernet Switches" or "the Intel '286 Accused Products").

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

118. On information and belief, an Intel '286 Accused Product Network implements at least the TRILL protocol.

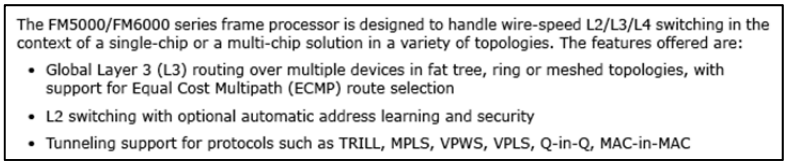
The FM5000/FM6000 series frame processor is designed to handle wire-speed L2/L3/L4 switching in the context of a single-chip or a multi-chip solution in a variety of topologies. The features offered are:

- Global Layer 3 (L3) routing over multiple devices in fat tree, ring or meshed topologies, with support for Equal Cost Multipath (ECMP) route selection
- L2 switching with optional automatic address learning and security
- Tunneling support for protocols such as TRILL, MPLS, VPWS, VPLS, Q-in-Q, MAC-in-MAC

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 43 (November 2014).

119. On information and belief, an Intel '286 Accused Product Network comprises a communication network.

120. On information and belief, an Intel '286 Accused Product Network comprises an entry end switch.

> The FM5000/FM6000 operates as a one-armed IP router combined with an Ethernet L2 switch (see Figure 5-1). In this architecture, incoming packets are first associated with a VLAN (using the VTAG tag if present, or by associating a default VLAN), and are then either switched within their respective VLANs or routed across VLANs or both. The decision to switch, route or drop depends on tables stored in the FFU (which includes a large ternary CAM to store IP route entries and access control lists), other tables located in the switch (such as MAC Address table) or tables located in the router (such as ARP table).

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 42 (November 2014) (highlighting added).

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

122. On information and belief, an Intel ’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., ingress RBridge RB1), an exit end switch (i.e., egress RBridge RB2), and intermediate switches (i.e., transit RBridges RB3 and RB4) in between them. The exit end switch (i.e., egress RBridge RB2) can be connected to multiple hosts (e.g., ESb).

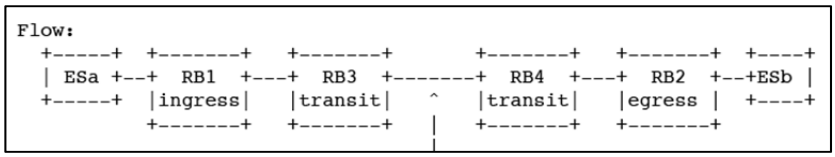


Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011).

123. On information and belief, an entry end switch in an Intel ’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., ingress RBridge) encapsulates an incoming data packet within a TRILL header and an Outer Ethernet header. The TRILL header includes an Egress RBridge Nickname field, which contains the unique RBridge Nickname of the exit end switch (i.e., egress RBridge).

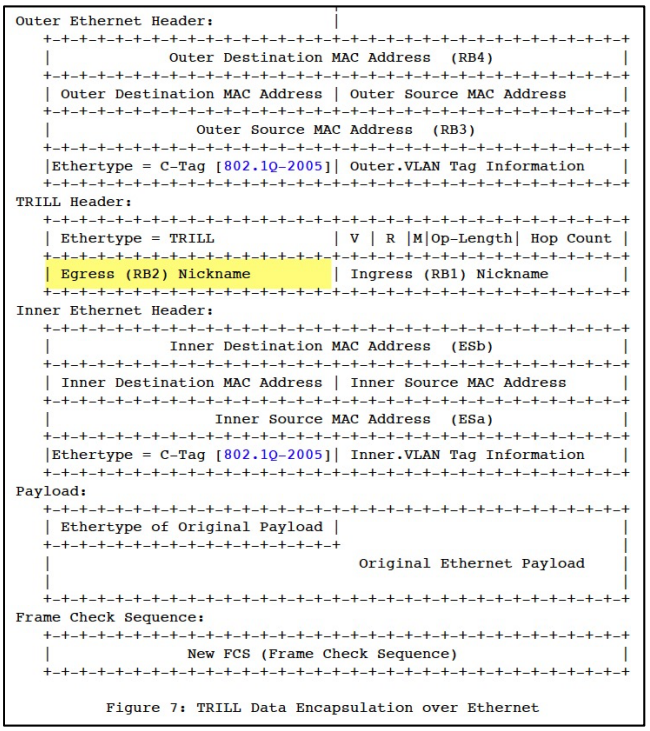


Figure 7: TRILL Data Encapsulation over Ethernet

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011) (highlighting added).

124. On information and belief, in an Intel ’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., transit RBridge) uses the Egress RBridge Nickname within the TRILL header to point the packet to the next RBridge.

These frames have a unicast inner MAC destination address (Inner.MacDA) and are those for which the ingress RBridge knows the egress RBridge for the destination MAC address in the frame’s VLAN.

Such frames are forwarded Rbridge hop by Rbridge hop to their egress Rbridge.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 16 (July 2011) (highlighting added).

> The first RBridge that a unicast frame encounters in a campus, RB1, encapsulates the received frame with a TRILL header that specifies the last RBridge, RB2, where the frame is decapsulated. RB1 is known as the "ingress RBridge" and RB2 is known as the "egress RBridge". To save room in the TRILL header and simplify forwarding lookups, a dynamic nickname acquisition protocol is run among the RBridges to select 2-octet nicknames for RBridges, unique within the campus, which are an abbreviation for the IS-IS ID of the RBridge. The 2-octet nicknames are used to specify the ingress and egress RBridges in the TRILL header.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 11 (July 2011) (highlighting added).

125. In addition, on information and belief, the exit end switch (i.e., egress RBridge) performs the entirety of all translation needed by the TRILL network to route the packet from the egress RBridge to the exit termination device (i.e., the packet's final destination) when it strips the TRILL header from the data packet and forwards it outside the TRILL network.

> If RBn is not a transit RBridge, that is, if the egress RBridge indicated is the RBridge performing the processing, the Inner.MacSA and Inner.VLAN ID are, by default, learned as associated with the ingress nickname unless that nickname is unknown or reserved or the Inner.MacSA is not unicast. Then the frame being forwarded is decapsulated to native form, and the following checks are performed:
>
> - The Inner.MacDA is checked. If it is not unicast, the frame is discarded.
> - If the Inner.MacDA corresponds to the RBridge doing the processing, the frame is locally delivered.
> - The Inner.VLAN ID is checked. If it is 0x0 or 0xFFF, the frame is discarded.
> - The Inner.MacDA and Inner.VLAN ID are looked up in RBn's local address cache and the frame is then either sent onto the link containing the destination, if the RBridge is appointed forwarder for that link for the frame's VLAN and is not inhibited (or discarded if it is inhibited), or processed as in one of the following two paragraphs.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 64 (July 2011).

126. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Intel '286 Accused Products, Intel 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).

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

128. On information and belief, Intel 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, Intel knew of the ’286 patent and knew of its infringement, including by way of this lawsuit.

129. On information and belief, Intel intended to induce patent infringement by third-party customers and users of the Intel ’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. Intel specifically intended and was aware that the normal and customary use of the accused products would infringe the ’286 patent. Intel 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, Intel provides the Intel ’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 Intel further provides documentation and training materials that cause customers of the Intel ’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 Intel ’286 Accused Products, Intel specifically intended to induce infringement of the ’286 patent, including at least claim 6. On information and belief, Intel engaged in such inducement to promote the sales of the Intel ’286 Accused Products and to actively induce its customers to infringe the ’286 patent. Accordingly, Intel 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.

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

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

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

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

133. 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.

134. 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.

135. 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.

136. 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.

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

138. Intel makes, uses, sells, offers to sell, and/or imports the Intel Ethernet switches, including but not limited to the Intel FM5224 Ethernet Switch, the Intel FM6300 Ethernet

Switch, and the Intel FM6700 Ethernet Switch (collectively, "the Intel Ethernet Switches" or "the Intel '259 Accused Products").

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

140. On information and belief, an Intel '259 Accused Product Network implements at least the TRILL protocol.

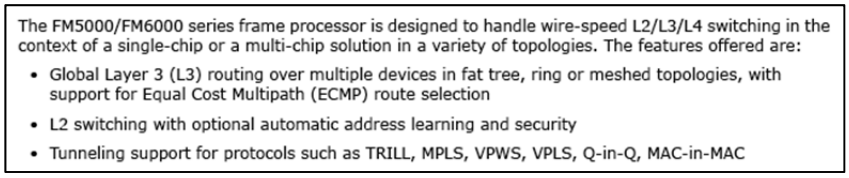

The FM5000/FM6000 series frame processor is designed to handle wire-speed L2/L3/L4 switching in the context of a single-chip or a multi-chip solution in a variety of topologies. The features offered are:

- Global Layer 3 (L3) routing over multiple devices in fat tree, ring or meshed topologies, with support for Equal Cost Multipath (ECMP) route selection
- L2 switching with optional automatic address learning and security
- Tunneling support for protocols such as TRILL, MPLS, VPWS, VPLS, Q-in-Q, MAC-in-MAC

*Intel Ethernet Switch FM5000/FM6000 Datasheet*, INTEL DATASHEET REVISION 3.3, at 43 (November 2014).

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

142. On information and belief, the Intel '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 TRILL network are encapsulated within a TRILL header, which includes, for example, an Egress RBridge Nickname field that contains the RBridge Nickname of the exit end switch.

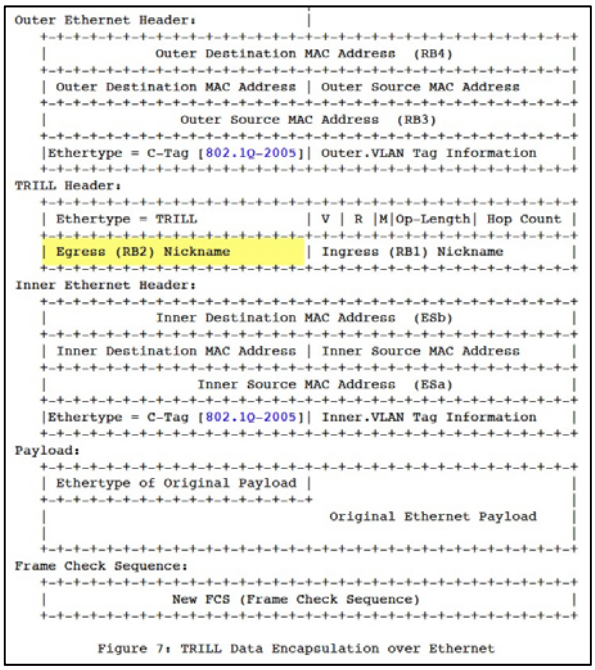


Figure 7: TRILL Data Encapsulation over Ethernet

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011) (highlighting added).

143. On information and belief, the Intel ’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 TRILL network, the Intel ’259 Accused Products are grouped into ingress RBridges, transit RBridges, and egress RBridges based on whether they are entry switches, intermediate switches, or exit switches.

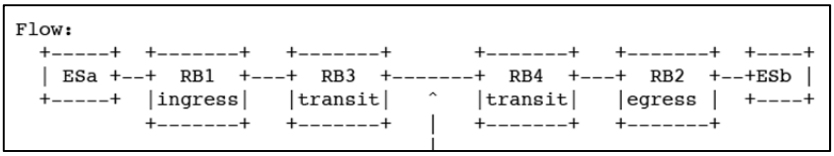


Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 29 (July 2011).

144. On information and belief, the Intel ’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 RBridge is assigned a unique RBridge Nickname, which includes a nickname priority value and a root priority value that are used in determining the topology of distribution trees within a TRILL network and are indicative, at least in part, of the network

topology. In addition, the RBridge Nickname is part of a network topology database and adjacency tables that are built and maintained by the IS-IS protocol and, when used with the topology database and/or adjacency tables, indicates, at least in part, the network topology.

> 3.7. RBridge Nicknames
>
> Nicknames are 16-bit dynamically assigned quantities that act as abbreviations for RBridges' IS-IS IDs to achieve a more compact encoding and can be used to specify potentially different trees with the same root. This assignment allows specifying up to 2**16 RBridges; however, the value 0x0000 is reserved to indicate that a nickname is not specified, the values 0xFFC0 through 0xFFFE are reserved for future specification, and the value 0xFFFF is permanently reserved. RBridges piggyback a nickname acquisition protocol on the link state protocol (see Section 3.7.3) to acquire one or more nicknames unique within the campus.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 23 (July 2011).

> 2. The following information in connection with the nickname or each of the nicknames of RBridge RBn:
>
> 2.1. The nickname value (2 octets).
>
> 2.2. The unsigned 8-bit priority for RBn to have that nickname (see Section 3.7.3).
>
> 2.3. The 16-bit unsigned priority of that nickname to becoming a distribution tree root.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 38 (July 2011).

> The first RBridge that a unicast frame encounters in a campus, RB1, encapsulates the received frame with a TRILL header that specifies the last RBridge, RB2, where the frame is decapsulated. RB1 is known as the "ingress RBridge" and RB2 is known as the "egress RBridge". To save room in the TRILL header and simplify forwarding lookups, a dynamic nickname acquisition protocol is run among the RBridges to select 2-octet nicknames for RBridges, unique within the campus, which are an abbreviation for the IS-IS ID of the RBridge. The 2-octet nicknames are used to specify the ingress and egress RBridges in the TRILL header.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 11 (July 2011) (highlighting added).

145. On information and belief, the Intel '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 TRILL network are encapsulated within a TRILL Header, which includes an Egress RBridge Nickname field that contains the RBridge Nickname

of the egress RBridge. TRILL switches, such as the Intel ’259 Accused Products, compare the value in the Egress RBridge Nickname field to values in the switches’ memory.

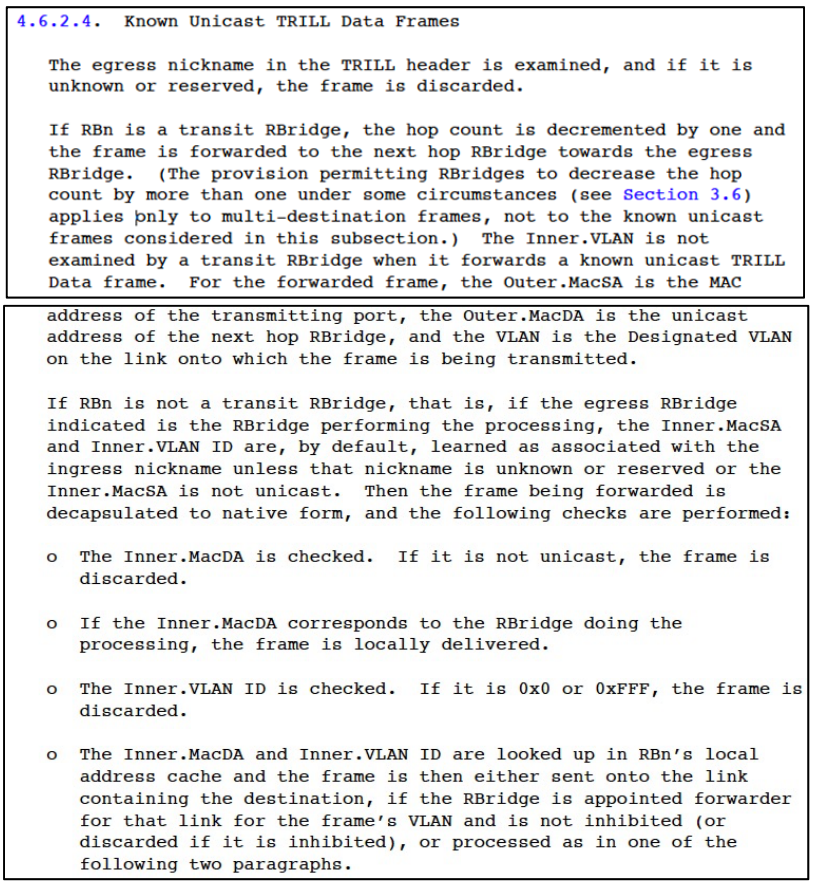

4.6.2.4. Known Unicast TRILL Data Frames

The egress nickname in the TRILL header is examined, and if it is unknown or reserved, the frame is discarded.

If RBn is a transit RBridge, the hop count is decremented by one and the frame is forwarded to the next hop RBridge towards the egress RBridge. (The provision permitting RBridges to decrease the hop count by more than one under some circumstances (see Section 3.6) applies only to multi-destination frames, not to the known unicast frames considered in this subsection.) The Inner.VLAN is not examined by a transit RBridge when it forwards a known unicast TRILL Data frame. For the forwarded frame, the Outer.MacSA is the MAC address of the transmitting port, the Outer.MacDA is the unicast address of the next hop RBridge, and the VLAN is the Designated VLAN on the link onto which the frame is being transmitted.

If RBn is not a transit RBridge, that is, if the egress RBridge indicated is the RBridge performing the processing, the Inner.MacSA and Inner.VLAN ID are, by default, learned as associated with the ingress nickname unless that nickname is unknown or reserved or the Inner.MacSA is not unicast. Then the frame being forwarded is decapsulated to native form, and the following checks are performed:

- The Inner.MacDA is checked. If it is not unicast, the frame is discarded.
- If the Inner.MacDA corresponds to the RBridge doing the processing, the frame is locally delivered.
- The Inner.VLAN ID is checked. If it is 0x0 or 0xFFF, the frame is discarded.
- The Inner.MacDA and Inner.VLAN ID are looked up in RBn’s local address cache and the frame is then either sent onto the link containing the destination, if the RBridge is appointed forwarder for that link for the frame’s VLAN and is not inhibited (or discarded if it is inhibited), or processed as in one of the following two paragraphs.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 63-64 (July 2011).

146. On information and belief, the Intel ’259 Accused Products forward data packets through a network based on the comparison of destination transport identification. For example, TRILL switches, such as the Intel ’259 Accused Products, forward encapsulated data packets using the Egress RBridge Nickname.

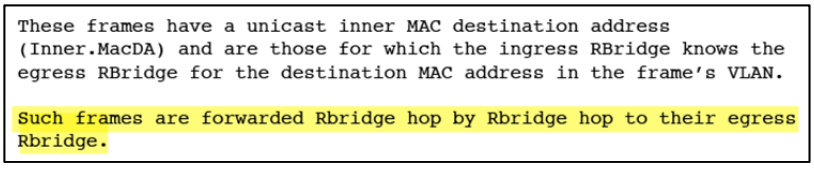

These frames have a unicast inner MAC destination address (Inner.MacDA) and are those for which the ingress RBridge knows the egress RBridge for the destination MAC address in the frame’s VLAN.

Such frames are forwarded Rbridge hop by Rbridge hop to their egress Rbridge.

Perlman, et al., *Routing Bridges (RBridges): Base Protocol Specification*, IETF RFC 6325, at 16 (July 2011) (highlighting added).

147. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Intel ’259 Accused Products, Intel 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).

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

149. On information and belief, Intel 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, Intel knew of the ’259 patent and knew of its infringement, including by way of this lawsuit.

150. On information and belief, Intel intended to induce patent infringement by third-party customers and users of the Intel ’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. Intel specifically intended and was aware that the normal and customary use of the accused products would infringe the ’259 patent. Intel 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, Intel provides the Intel ’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 Intel further provides documentation and training materials that cause customers of the Intel ’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 Intel ’259 Accused Products, Intel specifically intended to induce infringement of the ’259 patent, including at least claim 9. On information and belief, Intel engaged in such inducement to promote the sales of the Intel ’259 Accused Products and to actively induce its customers to infringe the ’259 patent. Accordingly, Intel 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.

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

152. As a result of Intel’s infringement of the ’259 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Intel’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Intel 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 Intel has infringed, either literally and/or under the doctrine of equivalents, the ’462 patent, the ’701 patent, the ’235 patent, the ’286 patent, and/or the ’259 patent;

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

C. A judgment and order requiring Intel 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 20, 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*