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

DUNTI NETWORK TECHNOLOGIES, LLC,

Plaintiff,

v.

JUNIPER NETWORKS, INC. AND JUNIPER NETWORKS (US), INC.,

***Defendants*.**

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, defendants Juniper Networks, Inc. and Juniper Networks (US), Inc. (collectively, "Juniper") has adopted Dunti's patented technologies—developed more than a decade ago right here in Texas—*en masse*. Juniper has profited handsomely from its use of Dunti's patented inventions, and Dunti deserves to be compensated for this use. But Juniper 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 Juniper 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 Juniper, Dunti relies on its intellectual property.

3. On information and belief, Defendant Juniper Networks, Inc. is a Delaware corporation with its principal office at 1133 Innovation Way, Sunnyvale, California 94089. Juniper Networks, Inc. can be served through its registered agent, The Corporation Trust Company, Corporation Trust Center, 1209 Orange St., Wilmington, DE 19801.

4. On information and belief, Defendant Juniper Networks (US), Inc. is a California corporation with its principal office at 1133 Innovation Way, Sunnyvale, California 94089. Juniper Networks (US), Inc. can be served through its registered agent, C T Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201-3136.

5. On information and belief, Juniper maintains one or more offices in the State of Texas, including one in Austin, Texas,[1] which, on information and belief, employ a number of technical, engineering, marketing, and sales employees. On information and belief, these employees have relevant information about the Juniper products at issue here.

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

[1] *See* Buchholz, Jan, *Real Estate Inc. Roundup: Brandmuscle, Juniper Networks Ink Leases*, AUSTIN BUSINESS JOURNAL ONLINE (published September 4, 2013), available at http://www.bizjournals.com/austin/blog/real-estate/2013/09/roundup-brandmuscle-juniper-networks.html.

JURISDICTION AND VENUE

7. 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).

8. Upon information and belief, this Court has personal jurisdiction over Juniper Networks, Inc. and Juniper Networks (US), Inc. in this action because they have committed acts within the Eastern District of Texas giving rise to this action and have established minimum contacts with this forum such that the exercise of jurisdiction over them would not offend traditional notions of fair play and substantial justice. Defendants Juniper Networks, Inc. and Juniper Networks (US), Inc., directly and/or through subsidiaries or intermediaries (including distributors, retailers, and others), have committed and continue 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, on information and belief, Juniper maintains at least one office in Texas, employs a number of technical, engineering, marketing, and sales employees in Texas, and Juniper Networks (US), Inc. has registered with the Texas Secretary of State to conduct business in Texas.

9. Venue is proper in this district under 28 U.S.C. §§ 1391(b)-(d) and 1400(b) because, on information and belief, Defendants Juniper Networks, Inc. and Juniper Networks (US), Inc. have at least one office in Texas, have transacted business in the Eastern District of Texas, and have committed acts of direct and indirect infringement in the Eastern District of Texas.

DUNTI’S LANDMARK NETWORK COMMUNICATION SYSTEMS

10. Dunti is the owner and assignee of ten patents on pioneering network technologies, including the five patents-in-suit (collectively, “the Dunti patents”).

11. 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.

12. 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, 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.

13. 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.

14. 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.

15. 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

[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).

substantially higher clocking speeds and to keep the processing pipeline full. The result was a faster, more efficient chip design.

16. 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;
- 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.

17. 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.

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

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

19. 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.

20. 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]

21. 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.

22. 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.

23. 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

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

agency's mission—to transform revolutionary concepts and even seeming impossibilities into practical capabilities.

24. 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

25. 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.

26. 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.

27. 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.

28. 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.

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

30. 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.

31. 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

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

33. 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.

34. 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.

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

36. 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.

37. 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.

38. 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.

39. 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

40. 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.

41. 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.

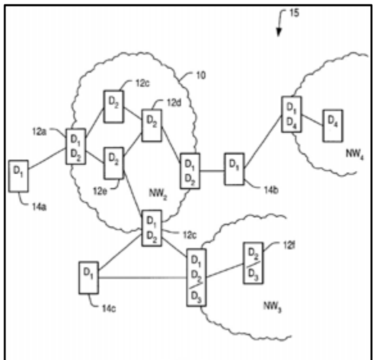

FIG. 1

’462 Patent, Fig. 1.

42. 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.

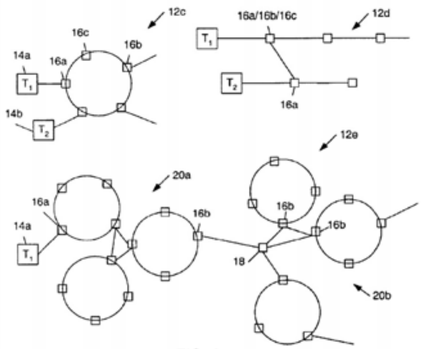

FIG. 2

’286 Patent, Fig. 2.

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.

JUNIPER'S INFRINGING PRODUCTS AND SERVICES

44. On information and belief, Juniper offers a high-performance data-center networking solution called the QFabric System architecture, which, as Juniper describes it, is "a revolutionary approach that provides a foundation for cloud-ready virtualized data center network environments. *Juniper QFabric System Data Sheet*, JUNIPER NETWORKS TECHNICAL DOCUMENTATION (2015), at 1.

45. On information and belief, the Juniper QFabric System is the "cornerstone of the QFabric architecture—a purpose-built solution that allows the creation of high-performance, scalable, cost-effective, dynamic, and easy to manage large enterprise IT or service provider cloud data centers." *Id*.

46. On information and belief, and according to Juniper, "[QFabric] is a scalable, high-performance, non-blocking, and easy-to-manage fabric that enables traditional Layer 2 and Layer 3 connectivity along with virtualization and convergence. The standards-based QFabric System is completely interoperable and seamlessly integrates with customers' existing data center environments, allowing them to easily migrate traditional tiered networks to a single tier QFabric architecture that connects compute, storage, network, and services resources as extensions of a low latency network." *Id*.

47. On information and belief, Juniper's QFabric System architecture consists of at least three separate but interdependent interconnect and control devices: (1) the QFabric Node, (2) the QFabric Interconnect, and (3) the QFabric Director.

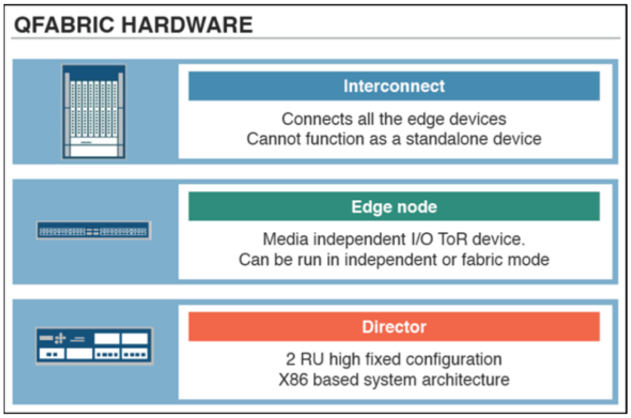

QFABRIC HARDWARE

Interconnect

Connects all the edge devices
Cannot function as a standalone device

Edge node

Media independent I/O ToR device.
Can be run in independent or fabric mode

Director

2 RU high fixed configuration
X86 based system architecture

Suptapa Bansal, *Data Center Fabric Architecture: Competitive Differentiators*, JUNIPER NETWORKS PRESENTATION (2011), at 15.

48. On information and belief, the QFabric Nodes are edge devices for the QFabric System and provide computing, storage, services, and network access for the QFabric System. Juniper's QFabric Nodes include at least the QFX3500, QFX3600, and QFX5100 switches.

49. On information and belief, the QFabric Interconnect represents the typical backplane of a modular switch, connecting the QFabric Node edge devices. Juniper's QFabric Interconnects include at least the QFX3600-I and QFX3008-I modules.

50. On information and belief, the QFabric Director manages the scalable data plane provided by the edge nodes and interconnects. Juniper's QFabric Directors include at least the QFX3100 device.

51. Exemplary implementations of a Juniper QFabric network are illustrated below:

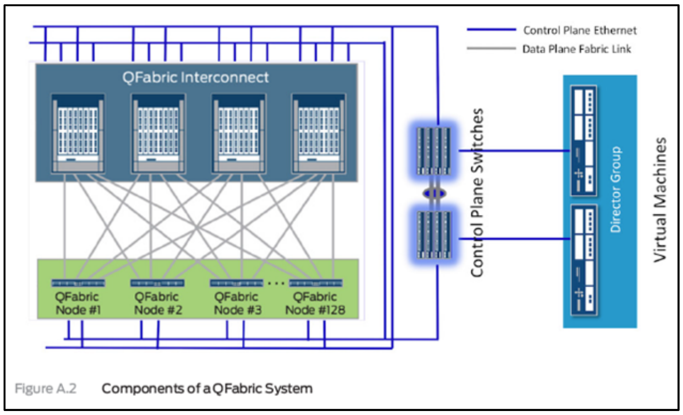

Figure A.2 Components of a QFabric System

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at vii.

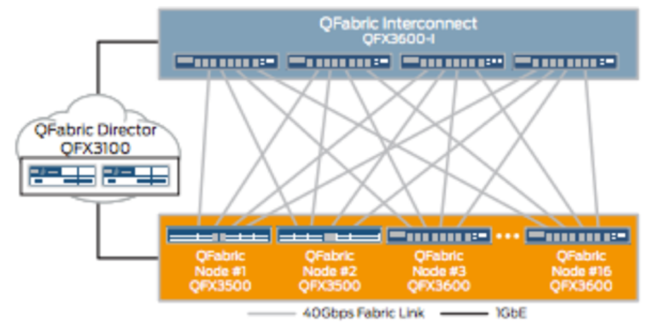

Figure 1: The QFX3000-M system topological view

Data Sheet, *Juniper QFabric System*, JUNIPER NETWORKS TECHNICAL DOCUMENTATION (2015), at 2.

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

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

53. 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.

54. 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.

55. 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.

56. 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.

57. 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.

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

59. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX series data center switches, including but not limited to the QFX3008-I switches, the QFX3100 switches, the QFX3500 switches, the QFX3600 switches, the QFX3600-I switches, and the QFX5100 switches (collectively, “the Juniper QFX Switches”).

60. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper EX Series switches, including but not limited to the EX2200 switches, the EX2300 switches, the EX3300

switches, the EX3400 switches, the EX4200 switches, the EX4300 switches, the EX4550 switches, the EX4600 switches, the EX8200 switches, and the EX9200 switches (collectively, “the Juniper EX Switches”).

61. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper Junos Network Operating System.

62. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX Switches, the Juniper EX Switches, and the Juniper Junos Network Operating System (collectively, “the Juniper ’462 Accused Products”).

63. Juniper makes, uses, sells, and/or offers to sell networks comprised of the Juniper ’462 Accused Products (“a Juniper ’462 Accused Product Network”).

64. On information and belief, a Juniper ’462 Accused Product Network implements at least Juniper’s QFabric technology.

65. On information and belief, a Juniper ’462 Accused Product Network comprises a communication system.

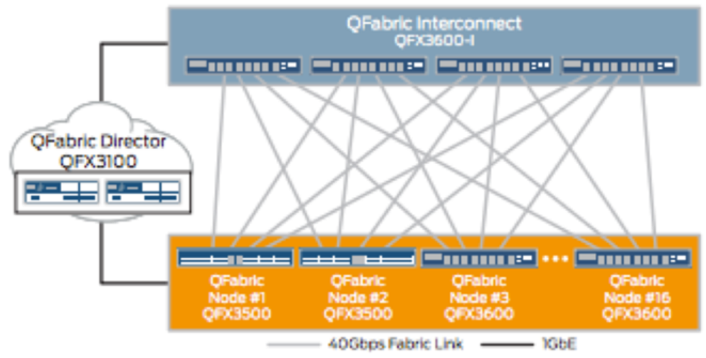

Figure 1: The QFX3000-M system topological view

Data Sheet, *Juniper QFabric System*, JUNIPER NETWORKS TECHNICAL DOCUMENTATION (2015), at 2.

66. On information and belief, a Juniper ’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., edge node) and an exit end module (i.e., edge node) at the edges of a QFabric network. In a QFabric network, there can be multiple intermediate modules (i.e., interconnects).

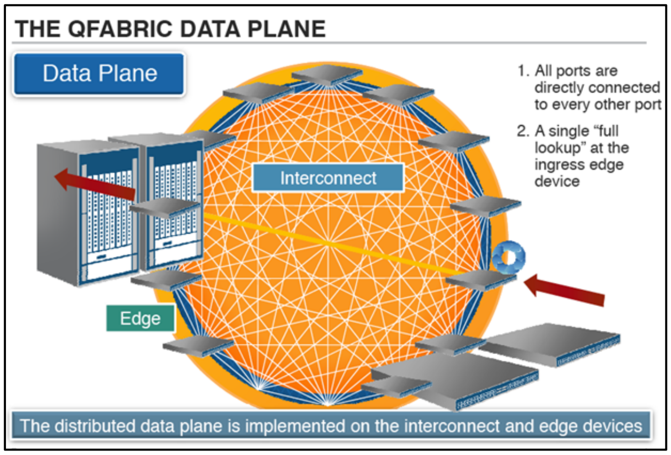

THE QFABRIC DATA PLANE

1. All ports are directly connected to every other port
2. A single "full lookup" at the ingress edge device

The distributed data plane is implemented on the interconnect and edge devices

Suptapa Bansal, *Data Center Fabric Architecture: Competitive Differentiators*, JUNIPER NETWORKS PRESENTATION (2011), at 12.

67. On information and belief, a Juniper '462 Accused Product Network comprises a first addressing domain for identifying each of the end modules and the intermediate module. For example, each switch (i.e., each Node and Interconnect) within a QFabric network is assigned a unique PFE-ID.

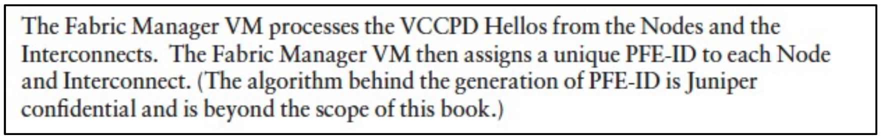

The Fabric Manager VM processes the VCCPD Hellos from the Nodes and the Interconnects. The Fabric Manager VM then assigns a unique PFE-ID to each Node and Interconnect. (The algorithm behind the generation of PFE-ID is Juniper confidential and is beyond the scope of this book.)

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 17.

68. On information and belief, a Juniper '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 nodes in a Juniper '462 Accused Product Network can be addressed using IP addresses, but IP addresses are not used to address intermediate modules (i.e., interconnects) when forwarding packets within a QFabric network.

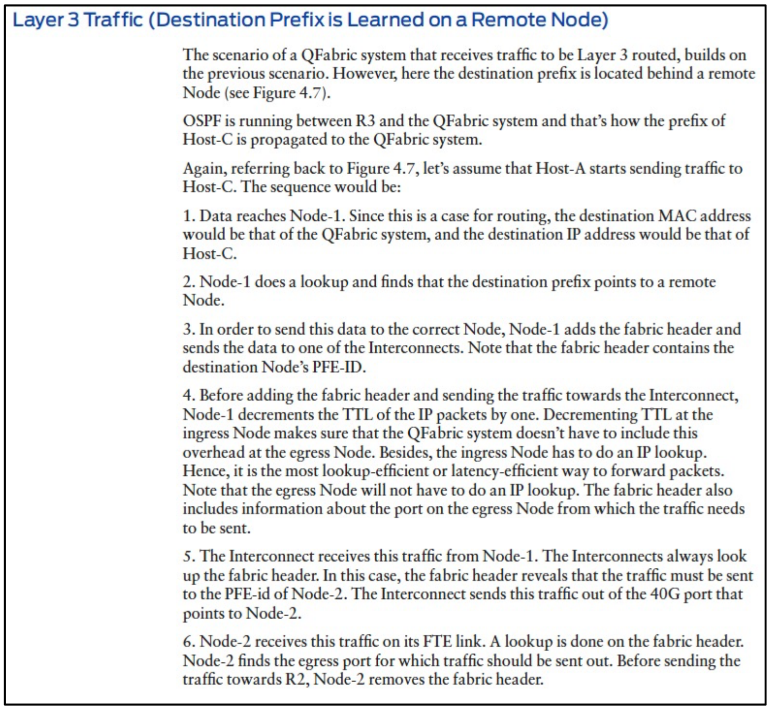

Layer 3 Traffic (Destination Prefix is Learned on a Remote Node)

The scenario of a QFabric system that receives traffic to be Layer 3 routed, builds on the previous scenario. However, here the destination prefix is located behind a remote Node (see Figure 4.7).

OSPF is running between R3 and the QFabric system and that's how the prefix of Host-C is propagated to the QFabric system.

Again, referring back to Figure 4.7, let's assume that Host-A starts sending traffic to Host-C. The sequence would be:

1. Data reaches Node-1. Since this is a case for routing, the destination MAC address would be that of the QFabric system, and the destination IP address would be that of Host-C.

2. Node-1 does a lookup and finds that the destination prefix points to a remote Node.

3. In order to send this data to the correct Node, Node-1 adds the fabric header and sends the data to one of the Interconnects. Note that the fabric header contains the destination Node's PFE-ID.

4. Before adding the fabric header and sending the traffic towards the Interconnect, Node-1 decrements the TTL of the IP packets by one. Decrementing TTL at the ingress Node makes sure that the QFabric system doesn't have to include this overhead at the egress Node. Besides, the ingress Node has to do an IP lookup. Hence, it is the most lookup-efficient or latency-efficient way to forward packets. Note that the egress Node will not have to do an IP lookup. The fabric header also includes information about the port on the egress Node from which the traffic needs to be sent.

5. The Interconnect receives this traffic from Node-1. The Interconnects always look up the fabric header. In this case, the fabric header reveals that the traffic must be sent to the PFE-id of Node-2. The Interconnect sends this traffic out of the 40G port that points to Node-2.

6. Node-2 receives this traffic on its FTE link. A lookup is done on the fabric header. Node-2 finds the egress port for which traffic should be sent out. Before sending the traffic towards R2, Node-2 removes the fabric header.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 73.

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

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

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

72. On information and belief, Juniper intended to induce patent infringement by third-party customers and users of the Juniper ’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. Juniper specifically intended and was aware that the normal and customary use of the accused products would infringe the ’462 patent. Juniper

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, Juniper provides the Juniper ’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 Juniper further provides documentation and training materials that cause customers of the Juniper ’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 Juniper ’462 Accused Products, Juniper specifically intended to induce infringement of the ’462 patent, including at least claim 1. On information and belief, Juniper engaged in such inducement to promote the sales of the Juniper ’462 Accused Products and to actively induce its customers to infringe the ’462 patent. Accordingly, Juniper 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.

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

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

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

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

76. 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.

77. 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.

78. 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).

79. 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.

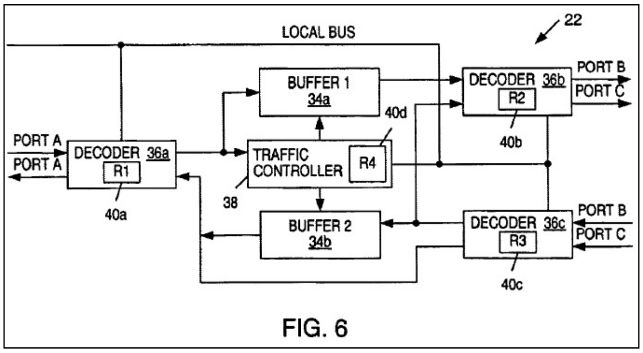

FIG. 6

’701 Patent, Fig. 6.

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

81. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX series data center switches, including but not limited to the QFX3008-I switches, the QFX3100 switches, the QFX3500 switches, the QFX3600 switches, the QFX3600-I switches, and the QFX5100 switches (collectively, “the Juniper QFX Switches”).

82. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper EX Series switches, including but not limited to the EX2200 switches, the EX2300 switches, the EX3300 switches, the EX3400 switches, the EX4200 switches, the EX4300 switches, the EX4550 switches, the EX4600 switches, the EX8200 switches, and the EX9200 switches (collectively, “the Juniper EX Switches”).

83. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper Junos Network Operating System.

84. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX Switches, the Juniper EX Switches, and the Juniper Junos Network Operating System (collectively, “the Juniper ’701 Accused Products”).

85. Juniper makes, uses, sells, and/or offers to sell networks comprised of the Juniper ’701 Accused Products (“a Juniper ’701 Accused Product Network”).

86. On information and belief, a Juniper ’701 Accused Product Network implements at least Juniper’s QFabric technology.

87. On information and belief, the Juniper ’701 Accused Products comprise a switch.

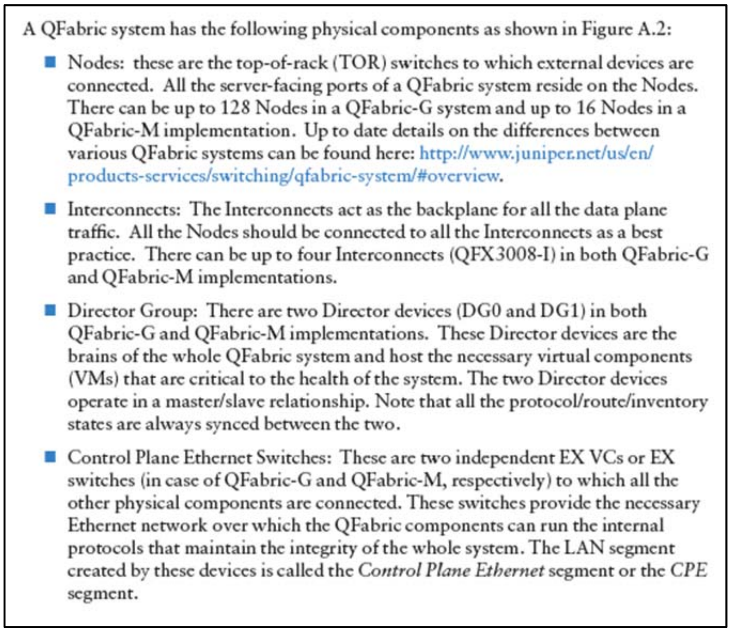

A QFabric system has the following physical components as shown in Figure A.2:

- Nodes: these are the top-of-rack (TOR) switches to which external devices are connected. All the server-facing ports of a QFabric system reside on the Nodes. There can be up to 128 Nodes in a QFabric-G system and up to 16 Nodes in a QFabric-M implementation. Up to date details on the differences between various QFabric systems can be found here: http://www.juniper.net/us/en/products-services/switching/qfabric-system/#overview.
- Interconnects: The Interconnects act as the backplane for all the data plane traffic. All the Nodes should be connected to all the Interconnects as a best practice. There can be up to four Interconnects (QFX3008-I) in both QFabric-G and QFabric-M implementations.
- Director Group: There are two Director devices (DG0 and DG1) in both QFabric-G and QFabric-M implementations. These Director devices are the brains of the whole QFabric system and host the necessary virtual components (VMs) that are critical to the health of the system. The two Director devices operate in a master/slave relationship. Note that all the protocol/route/inventory states are always synced between the two.
- Control Plane Ethernet Switches: These are two independent EX VCs or EX switches (in case of QFabric-G and QFabric-M, respectively) to which all the other physical components are connected. These switches provide the necessary Ethernet network over which the QFabric components can run the internal protocols that maintain the integrity of the whole system. The LAN segment created by these devices is called the *Control Plane Ethernet* segment or the *CPE* segment.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at vi.

88. On information and belief, the Juniper ’701 Accused Products within a QFabric network comprise a traffic manager which dispatches a series of read operations to a memory coupled within a data flow path. For example, a Juniper ’701 Accused Product Network includes a Director Group that includes a Network Node Group virtual machine, a Fabric Manager virtual machine, and a Fabric Control virtual machine. Each of the Juniper ’701 Accused Products includes memory and at least one processor (i.e., a Routing Engine).

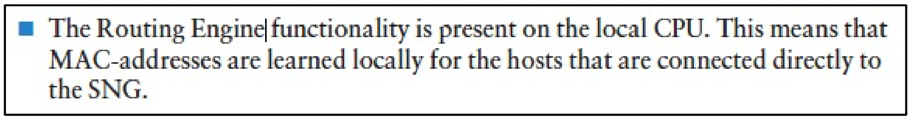

- The Routing Engine functionality is present on the local CPU. This means that MAC-addresses are learned locally for the hosts that are connected directly to the SNG.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 40.

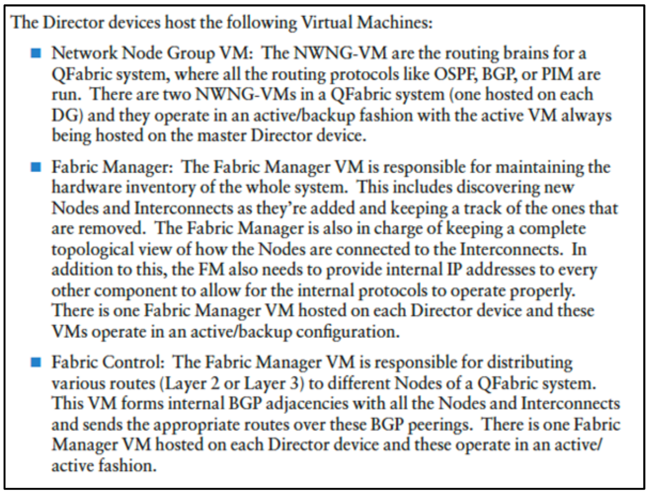
The Director devices host the following Virtual Machines:

- Network Node Group VM: The NWNG-VM are the routing brains for a QFabric system, where all the routing protocols like OSPF, BGP, or PIM are run. There are two NWNG-VMs in a QFabric system (one hosted on each DG) and they operate in an active/backup fashion with the active VM always being hosted on the master Director device.
- Fabric Manager: The Fabric Manager VM is responsible for maintaining the hardware inventory of the whole system. This includes discovering new Nodes and Interconnects as they're added and keeping a track of the ones that are removed. The Fabric Manager is also in charge of keeping a complete topological view of how the Nodes are connected to the Interconnects. In addition to this, the FM also needs to provide internal IP addresses to every other component to allow for the internal protocols to operate properly. There is one Fabric Manager VM hosted on each Director device and these VMs operate in an active/backup configuration.
- Fabric Control: The Fabric Manager VM is responsible for distributing various routes (Layer 2 or Layer 3) to different Nodes of a QFabric system. This VM forms internal BGP adjacencies with all the Nodes and Interconnects and sends the appropriate routes over these BGP peerings. There is one Fabric Manager VM hosted on each Director device and these operate in an active/active fashion.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at vii.

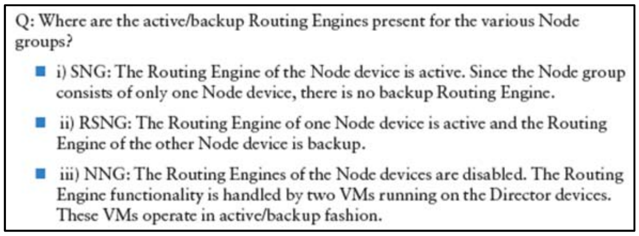
Q: Where are the active/backup Routing Engines present for the various Node groups?

- i) SNG: The Routing Engine of the Node device is active. Since the Node group consists of only one Node device, there is no backup Routing Engine.
- ii) RSNG: The Routing Engine of one Node device is active and the Routing Engine of the other Node device is backup.
- iii) NNG: The Routing Engines of the Node devices are disabled. The Routing Engine functionality is handled by two VMs running on the Director devices. These VMs operate in active/backup fashion.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 62.

89. On information and belief, the Juniper '701 Accused Products within a QFabric network include forwarding tables and routes (i.e., Layer 2 routes, Layer 3 routes, and multicast routes) 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.

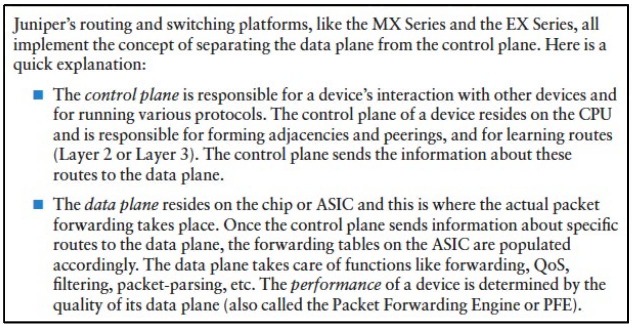
Juniper's routing and switching platforms, like the MX Series and the EX Series, all implement the concept of separating the data plane from the control plane. Here is a quick explanation:

- The *control plane* is responsible for a device's interaction with other devices and for running various protocols. The control plane of a device resides on the CPU and is responsible for forming adjacencies and peerings, and for learning routes (Layer 2 or Layer 3). The control plane sends the information about these routes to the data plane.
- The *data plane* resides on the chip or ASIC and this is where the actual packet forwarding takes place. Once the control plane sends information about specific routes to the data plane, the forwarding tables on the ASIC are populated accordingly. The data plane takes care of functions like forwarding, QoS, filtering, packet-parsing, etc. The *performance* of a device is determined by the quality of its data plane (also called the Packet Forwarding Engine or PFE).

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 40.

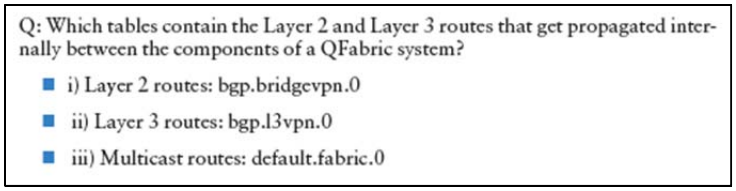

Q: Which tables contain the Layer 2 and Layer 3 routes that get propagated internally between the components of a QFabric system?

- i) Layer 2 routes: bgp.bridgevpn.0
- ii) Layer 3 routes: bgp.l3vpn.0
- iii) Multicast routes: default.fabric.0

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 62.

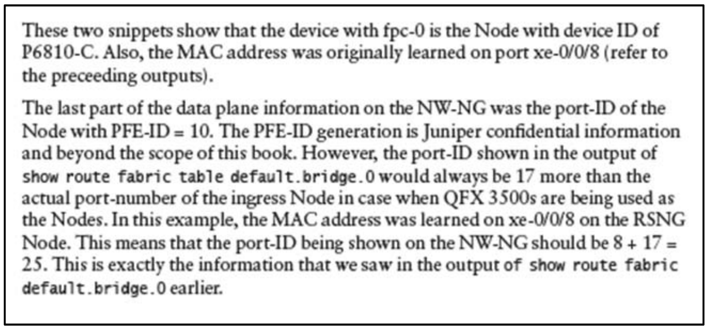

These two snippets show that the device with fpc-0 is the Node with device ID of P6810-C. Also, the MAC address was originally learned on port xe-0/0/8 (refer to the preceeding outputs).

The last part of the data plane information on the NW-NG was the port-ID of the Node with PFE-ID = 10. The PFE-ID generation is Juniper confidential information and beyond the scope of this book. However, the port-ID shown in the output of `show route fabric table default.bridge.0` would always be 17 more than the actual port-number of the ingress Node in case when QFX 3500s are being used as the Nodes. In this example, the MAC address was learned on xe-0/0/8 on the RSNG Node. This means that the port-ID being shown on the NW-NG should be 8 + 17 = 25. This is exactly the information that we saw in the output of `show route fabric default.bridge.0` earlier.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 52.

90. On information and belief, the Juniper ’701 Accused Products comprise an input port and an output port.

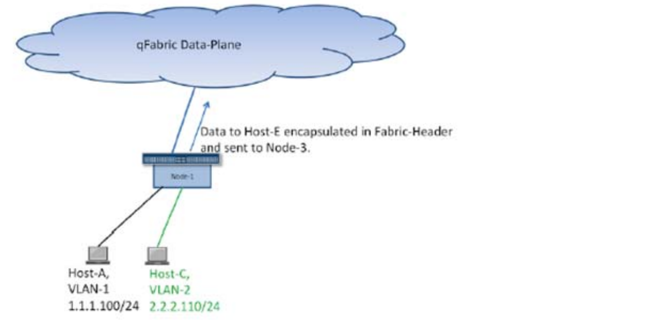
qFabric Data-Plane
Data to Host-E encapsulated in Fabric-Header and sent to Node-3.
Node-1
Host-A, VLAN-1 1.1.1.100/24
Host-C, VLAN-2 2.2.2.110/24

Figure 4.5 Node-1 Encapsulates the Data with Fabric Header and Sends It to the Interconnects

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 68.

91. On information and belief, the memory in the Juniper ’701 Accused Products comprises packets of data dispatched from the input port. For example, the Juniper ’701 Accused Products encapsulate incoming data packets within a Fabric Header. Incoming data

packets are comprised in memory within an edge Node as they are encapsulated within a Fabric Header as forwarding decisions are made.

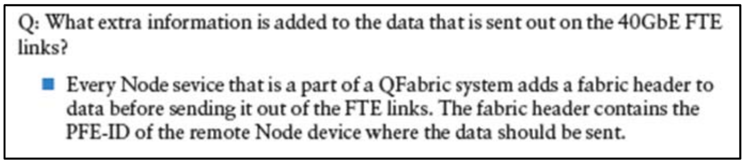

Q: What extra information is added to the data that is sent out on the 40GbE FTE links?

- Every Node sevice that is a part of a QFabric system adds a fabric header to data before sending it out of the FTE links. The fabric header contains the PFE-ID of the remote Node device where the data should be sent.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 78.

92. On information and belief, the Juniper ’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 Fabric Header includes a PFE-ID field, along with other information. The PFE-ID comprises the destination address.

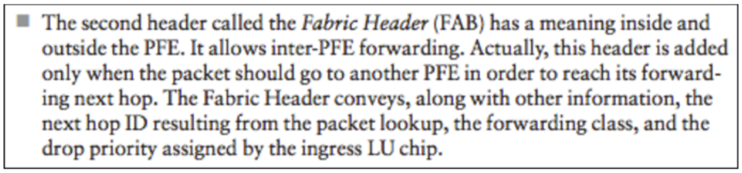

- The second header called the *Fabric Header* (FAB) has a meaning inside and outside the PFE. It allows inter-PFE forwarding. Actually, this header is added only when the packet should go to another PFE in order to reach its forwarding next hop. The Fabric Header conveys, along with other information, the next hop ID resulting from the packet lookup, the forwarding class, and the drop priority assigned by the ingress LU chip.

David Roy, *This Week: An Expert Packet Walkthrough on the MX Series 3D*, Juniper Networks Whitepaper (2015), at 31.

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

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

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

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

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

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

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

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

100. 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.

101. 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.

102. 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.

103. 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.

104. 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.

105. 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.

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

107. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX series data center switches, including but not limited to the QFX3008-I switches, the QFX3100 switches, the QFX3500 switches, the QFX3600 switches, the QFX3600-I switches, and the QFX5100 switches (collectively, “the Juniper QFX Switches”).

108. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper EX Series switches, including but not limited to the EX2200 switches, the EX2300 switches, the EX3300 switches, the EX3400 switches, the EX4200 switches, the EX4300 switches, the EX4550 switches, the EX4600 switches, the EX8200 switches, and the EX9200 switches (collectively, “the Juniper EX Switches”).

109. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper Junos Network Operating System.

110. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX Switches, the Juniper EX Switches, and the Juniper Junos Network Operating System (collectively, "the Juniper '235 Accused Products").

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

112. On information and belief, a Juniper '235 Accused Product Network implements at least Juniper's QFabric technology.

113. On information and belief, a Juniper '235 Accused Product Network comprises a communication network.

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

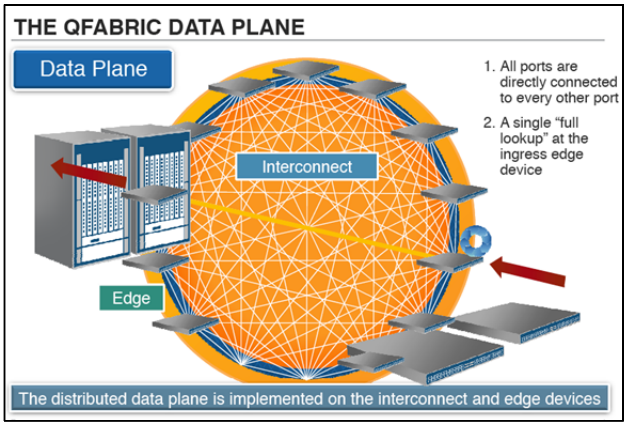

THE QFABRIC DATA PLANE

1. All ports are directly connected to every other port
2. A single "full lookup" at the ingress edge device

The distributed data plane is implemented on the interconnect and edge devices

Suptapa Bansal, *Data Center Fabric Architecture: Competitive Differentiators*, JUNIPER NETWORKS PRESENTATION (2011), at 12.

115. On information and belief, modules within a Juniper '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 (i.e., edge node or interconnect) within a QFabric network is assigned a unique PFE-ID, which is a unique identification number that is independent of and separate from the MAC address.

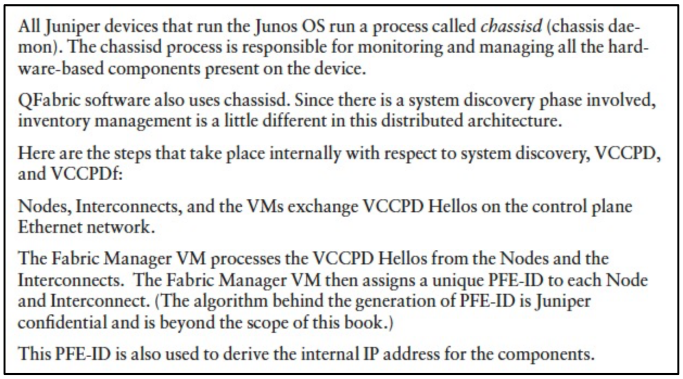

All Juniper devices that run the Junos OS run a process called *chassisd* (chassis daemon). The chassisd process is responsible for monitoring and managing all the hardware-based components present on the device.

QFabric software also uses chassisd. Since there is a system discovery phase involved, inventory management is a little different in this distributed architecture.

Here are the steps that take place internally with respect to system discovery, VCCPD, and VCCPDf:

Nodes, Interconnects, and the VMs exchange VCCPD Hellos on the control plane Ethernet network.

The Fabric Manager VM processes the VCCPD Hellos from the Nodes and the Interconnects. The Fabric Manager VM then assigns a unique PFE-ID to each Node and Interconnect. (The algorithm behind the generation of PFE-ID is Juniper confidential and is beyond the scope of this book.)

This PFE-ID is also used to derive the internal IP address for the components.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 17.

116. On information and belief, the first model layer used in a Juniper ’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 QFabric network, which already include headers from higher layers, are further wrapped/encapsulated within a Fabric Header that includes the PFE-ID of the egress edge node.

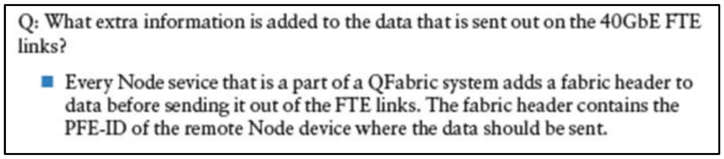

Q: What extra information is added to the data that is sent out on the 40GbE FTE links?

- Every Node sevice that is a part of a QFabric system adds a fabric header to data before sending it out of the FTE links. The fabric header contains the PFE-ID of the remote Node device where the data should be sent.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, Juniper Networks Whitepaper (2014), at 78.

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

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

Juniper 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).

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

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

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

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

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

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

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

125. 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.

126. 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.

127. 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.

128. 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.

129. 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.

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

131. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX series data center switches, including but not limited to the QFX3008-I switches, the QFX3100 switches, the QFX3500 switches, the QFX3600 switches, the QFX3600-I switches, and the QFX5100 switches (collectively, “the Juniper QFX Switches”).

132. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper EX Series switches, including but not limited to the EX2200 switches, the EX2300 switches, the EX3300 switches, the EX3400 switches, the EX4200 switches, the EX4300 switches, the EX4550

switches, the EX4600 switches, the EX8200 switches, and the EX9200 switches (collectively, "the Juniper EX Switches").

133. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper Junos Network Operating System.

134. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX Switches, the Juniper EX Switches, and the Juniper Junos Network Operating System (collectively, "the Juniper '286 Accused Products").

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

136. On information and belief, a Juniper '286 Accused Product Network implements at least Juniper's QFabric technology.

137. On information and belief, a Juniper '286 Accused Product Network comprises a communication network.

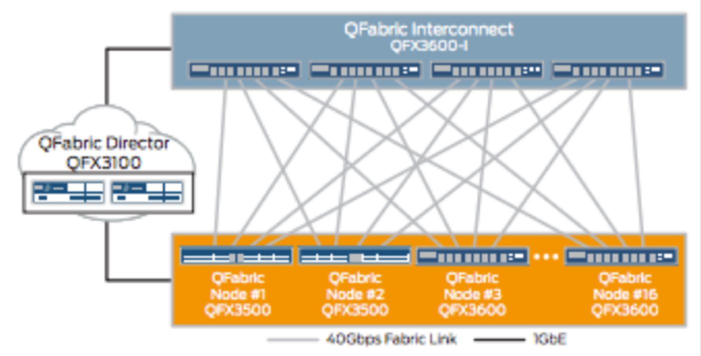

Figure 1: The QFX3000-M system topological view

Data Sheet, *Juniper QFabric System*, JUNIPER NETWORKS TECHNICAL DOCUMENTATION (2015), at 2.

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

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

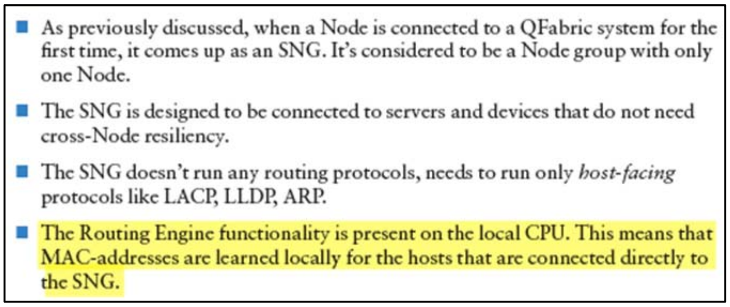

- As previously discussed, when a Node is connected to a QFabric system for the first time, it comes up as an SNG. It's considered to be a Node group with only one Node.
- The SNG is designed to be connected to servers and devices that do not need cross-Node resiliency.
- The SNG doesn't run any routing protocols, needs to run only *host-facing* protocols like LACP, LLDP, ARP.
- The Routing Engine functionality is present on the local CPU. This means that MAC-addresses are learned locally for the hosts that are connected directly to the SNG.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 40 (highlighting added).

140. On information and belief, a Juniper '286 Accused Product Network comprises a plurality of 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., edge Node), an exit end switch (i.e., edge Node), and multiple intermediate switches (i.e., Interconnects) in between them.

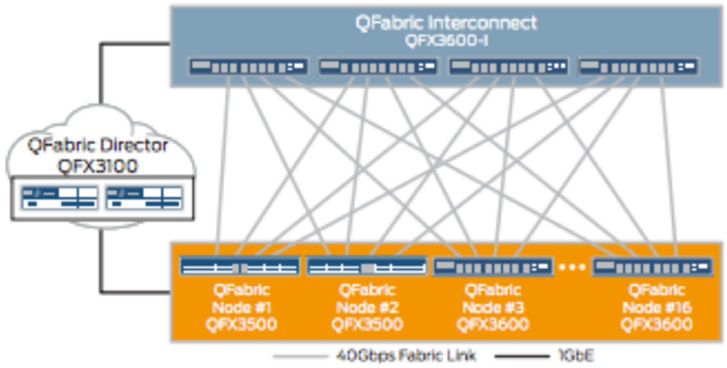

Figure 1: The QFX3000-M system topological view

Data Sheet, *Juniper QFabric System*, JUNIPER NETWORKS TECHNICAL DOCUMENTATION (2015), at 2.

141. On information and belief, an entry end switch in a Juniper '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., QFabric edge Node) encapsulates an incoming data packet within a Fabric Header. The Fabric Header includes the unique PFE-ID of the exit end switch (i.e., QFabric edge Node).

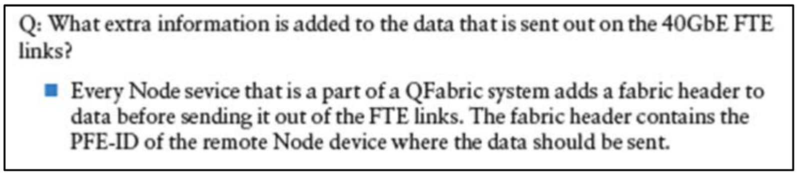

Q: What extra information is added to the data that is sent out on the 40GbE FTE links?

- Every Node sevice that is a part of a QFabric system adds a fabric header to data before sending it out of the FTE links. The fabric header contains the PFE-ID of the remote Node device where the data should be sent.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 78.

142. On information and belief, in a Juniper ’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., QFabric Interconnect) uses the PFE-ID within the Fabric Header to point the packet to the next switch.

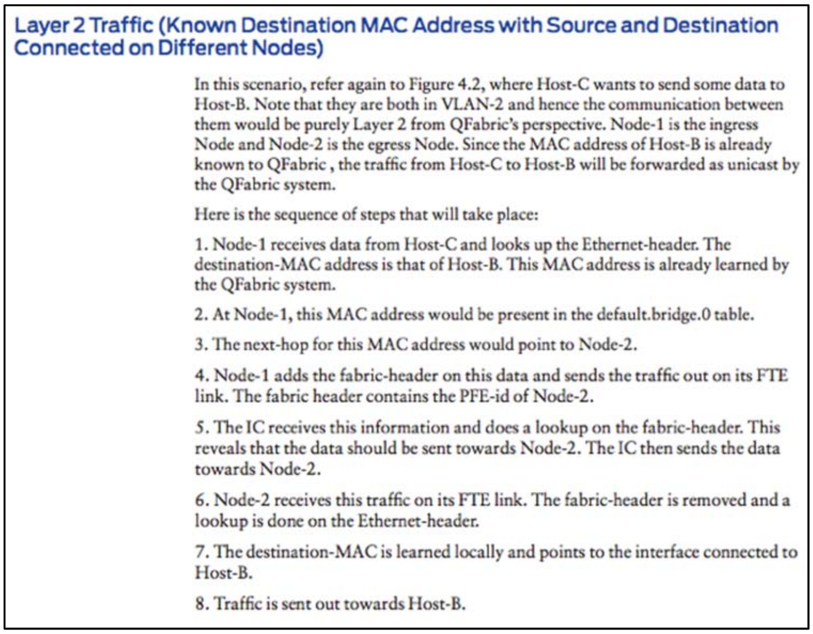

Layer 2 Traffic (Known Destination MAC Address with Source and Destination Connected on Different Nodes)

In this scenario, refer again to Figure 4.2, where Host-C wants to send some data to Host-B. Note that they are both in VLAN-2 and hence the communication between them would be purely Layer 2 from QFabric's perspective. Node-1 is the ingress Node and Node-2 is the egress Node. Since the MAC address of Host-B is already known to QFabric , the traffic from Host-C to Host-B will be forwarded as unicast by the QFabric system.

Here is the sequence of steps that will take place:

1. Node-1 receives data from Host-C and looks up the Ethernet-header. The destination-MAC address is that of Host-B. This MAC address is already learned by the QFabric system.

2. At Node-1, this MAC address would be present in the default.bridge.0 table.

3. The next-hop for this MAC address would point to Node-2.

4. Node-1 adds the fabric-header on this data and sends the traffic out on its FTE link. The fabric header contains the PFE-id of Node-2.

5. The IC receives this information and does a lookup on the fabric-header. This reveals that the data should be sent towards Node-2. The IC then sends the data towards Node-2.

6. Node-2 receives this traffic on its FTE link. The fabric-header is removed and a lookup is done on the Ethernet-header.

7. The destination-MAC is learned locally and points to the interface connected to Host-B.

8. Traffic is sent out towards Host-B.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 70-71.

143. In addition, on information and belief, the exit end switch (i.e., QFabric egress edge Node) performs the entirety of all translation needed by the QFabric network to route the

packet from the QFabric egress edge Node to the exit termination device (i.e., the packet's final destination).

144. By making, using, testing, offering for sale, and/or selling communication network products and services, including but not limited to the Juniper '286 Accused Products, Juniper has injured Dunti and is liable to Dunti for directly infringing one or more claims of the '286 patent, including at least claim 1, pursuant to 35 U.S.C. § 271(a).

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

146. On information and belief, Juniper 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, Juniper knew of the '286 patent and knew of its infringement, including by way of this lawsuit.

147. On information and belief, Juniper intended to induce patent infringement by third-party customers and users of the Juniper '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. Juniper specifically intended and was aware that the normal and customary use of the accused products would infringe the '286 patent. Juniper 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, Juniper provides the Juniper '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 1, and Juniper further provides documentation and training materials that cause customers of the Juniper '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 Juniper '286 Accused Products, Juniper specifically intended to induce infringement of the '286 patent, including at least claim 1. On information and belief, Juniper engaged in such inducement to promote the sales of the Juniper '286 Accused Products and to actively induce its customers to infringe the '286 patent. Accordingly, Juniper

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.

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

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

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

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

151. 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.

152. 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.

153. 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.

154. 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.

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

156. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX series data center switches, including but not limited to the QFX3008-I switches, the QFX3100 switches, the QFX3500 switches, the QFX3600 switches, the QFX3600-I switches, and the QFX5100 switches (collectively, “the Juniper QFX Switches”).

157. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper EX Series switches, including but not limited to the EX2200 switches, the EX2300 switches, the EX3300 switches, the EX3400 switches, the EX4200 switches, the EX4300 switches, the EX4550 switches, the EX4600 switches, the EX8200 switches, and the EX9200 switches (collectively, “the Juniper EX Switches”).

158. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper Junos Network Operating System.

159. Juniper makes, uses, sells, offers to sell, and/or imports the Juniper QFX Switches, the Juniper EX Switches, and the Juniper Junos Network Operating System (collectively, “the Juniper ’259 Accused Products”).

160. Juniper makes, uses, sells, and/or offers to sell networks comprised of the Juniper ’259 Accused Products (“a Juniper ’259 Accused Product Network”).

161. On information and belief, a Juniper ’259 Accused Product Network implements at least Juniper’s QFabric technology.

162. On information and belief, the Juniper ’259 Accused Products perform a method of transporting packets across a network.

163. On information and belief, the Juniper ’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 QFabric network are encapsulated within a Fabric Header, which contains the PFE-ID of the exit end switch.

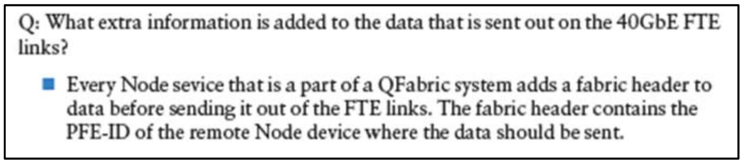

Q: What extra information is added to the data that is sent out on the 40GbE FTE links?

- Every Node sevice that is a part of a QFabric system adds a fabric header to data before sending it out of the FTE links. The fabric header contains the PFE-ID of the remote Node device where the data should be sent.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 78.

164. On information and belief, the Juniper ’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 QFabric network, the Juniper ’259 Accused Products are grouped into ingress edge Nodes, egress edge Nodes, and Interconnects based on whether they are entry switches, intermediate switches, or exit switches.

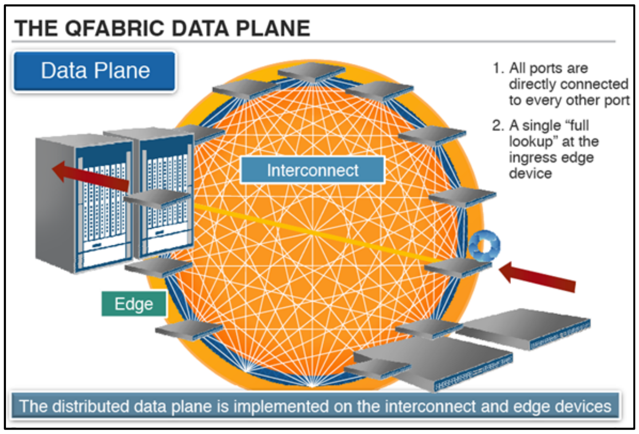

THE QFABRIC DATA PLANE

1. All ports are directly connected to every other port
2. A single “full lookup” at the ingress edge device

The distributed data plane is implemented on the interconnect and edge devices

Suptapa Bansal, *Data Center Fabric Architecture: Competitive Differentiators*, JUNIPER NETWORKS PRESENTATION (2011), at 12.

165. On information and belief, the Juniper ’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 QFabric Node is assigned a unique PFE-ID, which is part of a network topology database that is built and maintained by the Director Group using VCCPDf.

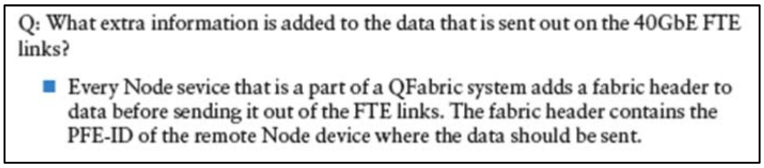

Q: What extra information is added to the data that is sent out on the 40GbE FTE links?

- Every Node sevice that is a part of a QFabric system adds a fabric header to data before sending it out of the FTE links. The fabric header contains the PFE-ID of the remote Node device where the data should be sent.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 78.

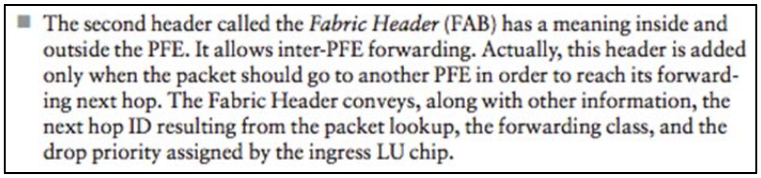

- The second header called the *Fabric Header* (FAB) has a meaning inside and outside the PFE. It allows inter-PFE forwarding. Actually, this header is added only when the packet should go to another PFE in order to reach its forwarding next hop. The Fabric Header conveys, along with other information, the next hop ID resulting from the packet lookup, the forwarding class, and the drop priority assigned by the ingress LU chip.

David Roy, *This Week: An Expert Packet Walkthrough on the MX Series 3D*, JUNIPER NETWORKS WHITEPAPER (2015), at 31.

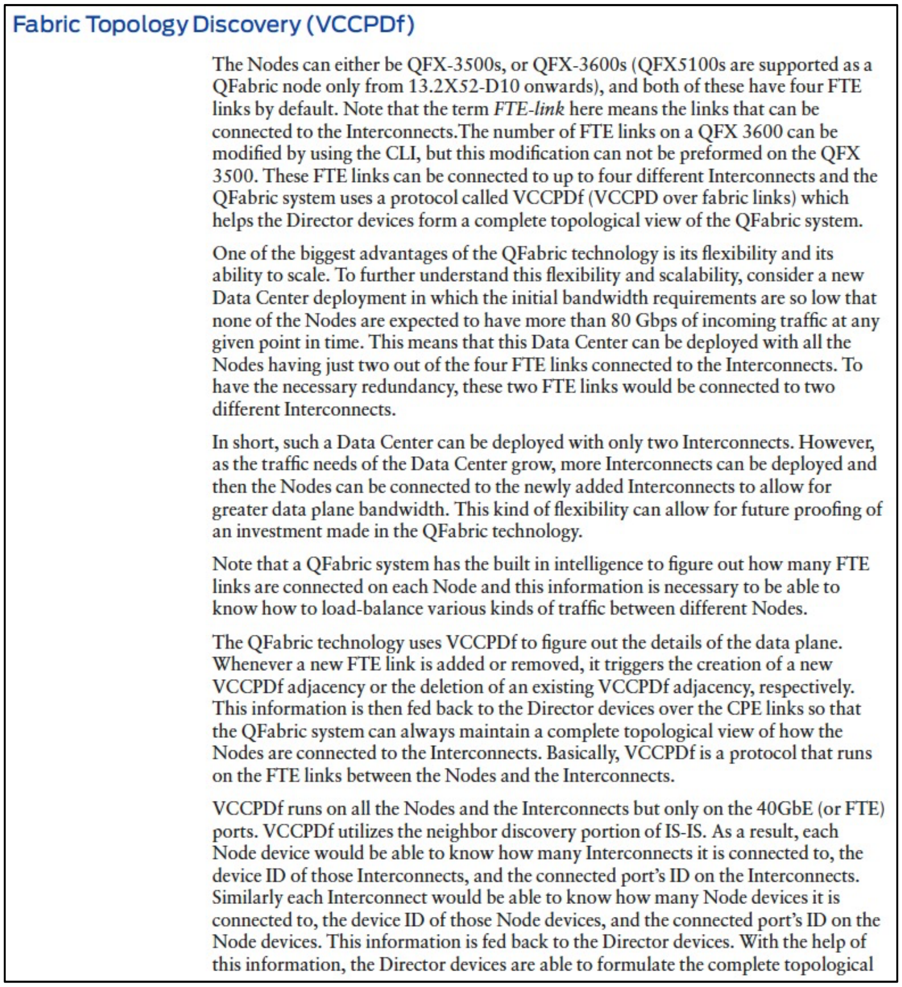

Fabric Topology Discovery (VCCPDf)

The Nodes can either be QFX-3500s, or QFX-3600s (QFX5100s are supported as a QFabric node only from 13.2X52-D10 onwards), and both of these have four FTE links by default. Note that the term *FTE-link* here means the links that can be connected to the Interconnects.The number of FTE links on a QFX 3600 can be modified by using the CLI, but this modification can not be preformed on the QFX 3500. These FTE links can be connected to up to four different Interconnects and the QFabric system uses a protocol called VCCPDf (VCCPD over fabric links) which helps the Director devices form a complete topological view of the QFabric system.

One of the biggest advantages of the QFabric technology is its flexibility and its ability to scale. To further understand this flexibility and scalability, consider a new Data Center deployment in which the initial bandwidth requirements are so low that none of the Nodes are expected to have more than 80 Gbps of incoming traffic at any given point in time. This means that this Data Center can be deployed with all the Nodes having just two out of the four FTE links connected to the Interconnects. To have the necessary redundancy, these two FTE links would be connected to two different Interconnects.

In short, such a Data Center can be deployed with only two Interconnects. However, as the traffic needs of the Data Center grow, more Interconnects can be deployed and then the Nodes can be connected to the newly added Interconnects to allow for greater data plane bandwidth. This kind of flexibility can allow for future proofing of an investment made in the QFabric technology.

Note that a QFabric system has the built in intelligence to figure out how many FTE links are connected on each Node and this information is necessary to be able to know how to load-balance various kinds of traffic between different Nodes.

The QFabric technology uses VCCPDf to figure out the details of the data plane. Whenever a new FTE link is added or removed, it triggers the creation of a new VCCPDf adjacency or the deletion of an existing VCCPDf adjacency, respectively. This information is then fed back to the Director devices over the CPE links so that the QFabric system can always maintain a complete topological view of how the Nodes are connected to the Interconnects. Basically, VCCPDf is a protocol that runs on the FTE links between the Nodes and the Interconnects.

VCCPDf runs on all the Nodes and the Interconnects but only on the 40GbE (or FTE) ports. VCCPDf utilizes the neighbor discovery portion of IS-IS. As a result, each Node device would be able to know how many Interconnects it is connected to, the device ID of those Interconnects, and the connected port's ID on the Interconnects. Similarly each Interconnect would be able to know how many Node devices it is connected to, the device ID of those Node devices, and the connected port's ID on the Node devices. This information is fed back to the Director devices. With the help of this information, the Director devices are able to formulate the complete topological

(continued)

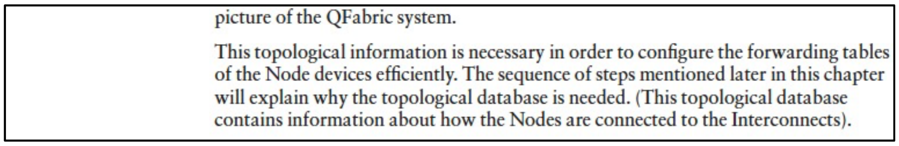

picture of the QFabric system.

This topological information is necessary in order to configure the forwarding tables of the Node devices efficiently. The sequence of steps mentioned later in this chapter will explain why the topological database is needed. (This topological database contains information about how the Nodes are connected to the Interconnects).

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 16-17.

166. On information and belief, the Juniper ’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 QFabric network are encapsulated within a Fabric Header, which includes the PFE-ID of the egress edge Node. QFabric switches, such as the Juniper ’259 Accused Products, compare the value in the PFE-ID field to values in the switches’ memory.

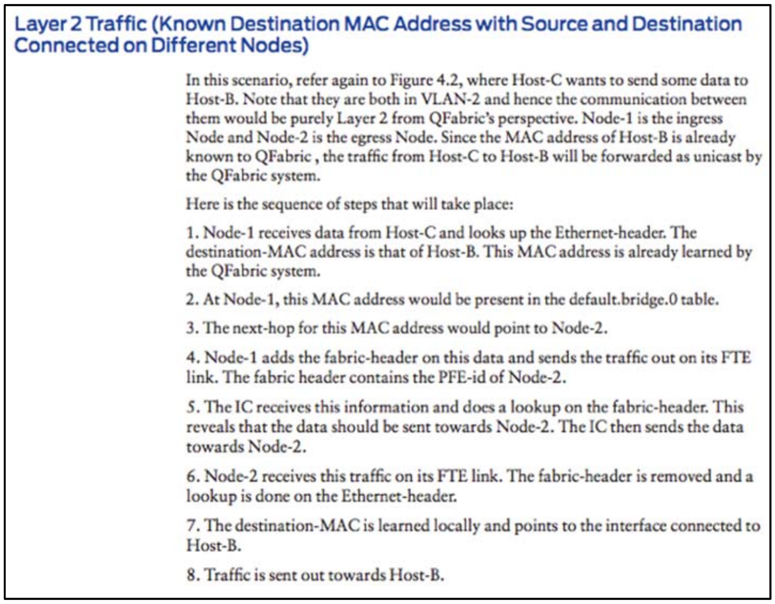

Layer 2 Traffic (Known Destination MAC Address with Source and Destination Connected on Different Nodes)

In this scenario, refer again to Figure 4.2, where Host-C wants to send some data to Host-B. Note that they are both in VLAN-2 and hence the communication between them would be purely Layer 2 from QFabric's perspective. Node-1 is the ingress Node and Node-2 is the egress Node. Since the MAC address of Host-B is already known to QFabric , the traffic from Host-C to Host-B will be forwarded as unicast by the QFabric system.

Here is the sequence of steps that will take place:

1. Node-1 receives data from Host-C and looks up the Ethernet-header. The destination-MAC address is that of Host-B. This MAC address is already learned by the QFabric system.

2. At Node-1, this MAC address would be present in the default.bridge.0 table.

3. The next-hop for this MAC address would point to Node-2.

4. Node-1 adds the fabric-header on this data and sends the traffic out on its FTE link. The fabric header contains the PFE-id of Node-2.

5. The IC receives this information and does a lookup on the fabric-header. This reveals that the data should be sent towards Node-2. The IC then sends the data towards Node-2.

6. Node-2 receives this traffic on its FTE link. The fabric-header is removed and a lookup is done on the Ethernet-header.

7. The destination-MAC is learned locally and points to the interface connected to Host-B.

8. Traffic is sent out towards Host-B.

Ankit Chadha, *This Week: QFabric System Traffic Flows and Troubleshooting*, JUNIPER NETWORKS WHITEPAPER (2014), at 70-71.

167. On information and belief, the Juniper ’259 Accused Products forward data packets through a network based on the comparison of destination transport identification. For example, QFabric switches, such as the Juniper ’259 Accused Products, forward encapsulated data packets using the PFE-ID values.

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

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

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

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

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

173. As a result of Juniper’s infringement of the ’259 patent, Dunti has suffered monetary damages, and seeks recovery in an amount adequate to compensate for Juniper’s infringement, but in no event less than a reasonable royalty for the use made of the invention by Juniper 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 Juniper 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 Juniper’s acts of infringement in accordance with 35 U.S.C. § 284;

C. A judgment and order requiring Juniper 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