

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Netskope, Inc.,
Petitioner,

v.

Fortinet, Inc.,
Patent Owner

Case IPR2023-00459
U.S. Patent No. 10,084,825

**SUPPLEMENTAL DECLARATION OF DR. JOHN BLACK JR.
IN SUPPORT OF PATENT OWNER'S RESPONSE**

1. I have been retained on behalf of Fortinet, Inc. (“Patent Owner”) in connection with the above-captioned *inter partes* review (IPR). I have been retained to provide my opinions in support of Fortinet’s Patent Owner Preliminary Response. I am being compensated for my time at the rate of \$625 per hour. I have no interest in the outcome of this proceeding.

2. In preparing this declaration, I have reviewed and am familiar with the Petition for IPR2023-00459, U.S. Patent No. 10,084,825 (“the ’825 patent”) and its file history, and all other materials cited and discussed in the Petition (including the declaration of Petitioner’s expert Michael Franz, Ph.D.), and cited and discussed in this Declaration. I have also reviewed the Board’s Institution Decision.

3. The statements made herein are based upon my own knowledge and opinion. This Declaration represents only the opinions I have formed to date. I may consider additional documents as they become available or other documents that are necessary to form my opinions. I reserve the right to revise, supplement, or amend my opinions based on new information and on my continuing analysis.

I. QUALIFICATIONS

4. My qualifications can be found in my Curriculum Vitae (*see* Ex. 2002), which includes my detailed employment background, professional experience, and list of technical publications and patents.

5. I am an Associate Professor of Computer Science at the University of Colorado, Boulder. I received a B.S. in Mathematics and Computer Science from the California State University at Hayward (now "California State University, East Bay") in 1988. I received an M.S. in Computer Science in 1997, and a Ph.D. in Computer Science in 2000, both from the University of California at Davis.

6. I have taught more than 60 classes in computer science, on subjects including data structures, algorithms, networking, operating systems, software engineering, security, cryptography, discrete mathematics, and quantum computing. I have authored or coauthored more than 20 publications, primarily on issues relating to computer security. I have been involved with computers for over 40 years in both commercial and academic capacities.

7. My earliest interest was in networks and security. My first memories in this regard were around 1975 when a group of friends and I learned about the telephone network and its security. A few years later, personal computers became available and I spent most of my free time studying, programming, and modifying them. I worked extensively with various networking products throughout the 1980s, and developed an interest in cryptography soon thereafter. Although my undergraduate institution had no courses in cryptography or security in the 1980s, I decided to pursue self-study at the time, and opted to double major in Computer

Science and Mathematics because cryptography is a blend of these two subject areas.

8. After earning my B.S. degree in 1988, I worked for six years at Ingres Corp as a software developer, writing and reviewing code written in C. My work primarily was directed at transaction logging, data type support, and security.

9. In 1995, I began my Ph.D. at UC Davis under cryptographer Phillip Rogaway. My area of focus was cryptography and security and my research involved encryption, authentication, hash functions, and network security.

10. After graduation I took a position as an Assistant Professor at the University of Nevada at Reno. In the Fall of 2001, I taught the networking class there, which included coverage of Ethernet, interior gateway protocols, exterior gateway protocols, ARP, DHCP/BOOTP, IP, UDP, TCP, HTTP, SMTP and other protocols. In 2001, a graduate student and I looked at the security of ARP and invented a new protocol "AuthARP" to add security to the protocol.

11. In 2002, I moved to the University of Colorado at Boulder where I am currently employed. In the Fall of 2002, I co-designed and co-taught a new course called "Foundations of Computer and Network Security," which included descriptions of security issues around both wired and wireless security challenges, mostly for public-facing network services including the world-wide web. I have

taught this class seven more times since then, including modern topics such as wireless networking, the Internet of Things, and so forth.

12. In my career at the University of Colorado, I have published several more papers in the area of cryptography and network security. I have taught more than 30 courses in network security and cryptography, and have graduated several PhD students in these areas.

13. I also have worked for a consulting company at times, writing software on contract basis. Although most projects are covered by NDAs, many involved networks and/or aspects of computer security.

14. In 2011, I began technical consulting for a local company called Cardinal Peak, which focuses primarily on video encoding and delivery systems. My work for Cardinal Peak has largely been directed to video encoding, transcoding, compression, encryption, and DRM, but also has included networking projects and embedded device work. For example, I designed the security system for the Pro1 smart thermostat, implemented the DRM for the Yonder Music App, worked on 802.1X code for smart dog collars, and helped design the cryptography used in Fitbit devices for wireless transfer of a Fitbit watch to a phone or laptop.

15. In 2016, I took a two-and-a-half year leave of absence from the University of Colorado to start a company named "SecureSet" in Denver, Colorado. The objective of SecureSet is to take reasonably proficient technical

people and turn them into computer and network security specialists via five months of intensive training. SecureSet was sold to WeWork in 2019 and continues to offer computer security classes today. In 2018 I returned to my position at the University where I remain employed to the present day.

II. LEGAL UNDERSTANDING

A. The Person of Ordinary Skill in the Art

16. I understand that a person of ordinary skill in the relevant art at the time of the invention (also referred to herein as “ordinary artisan”) is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity—not an automaton.

17. I have been asked to consider the level of ordinary skill in the field that someone would have had at the time the claimed invention was made. In deciding the level of ordinary skill, I considered the following:

- the levels of education and experience of persons working in the field;
- the types of problems encountered in the field; and
- the sophistication of the technology.

18. A person of ordinary skill in the art relevant to the '825 patent at the time of the invention would have a Bachelor of Science degree in electrical engineering and/or computer science, and two years of work or research experience

in the fields of network and data security, or a Master's degree in electrical engineering and/or computer science and one year of work or research experience in related fields.

19. I have reviewed the declaration of Dr. Franz, including his opinions regarding the Person of Ordinary Skill in the Art. My description of the level of ordinary skill in the art is essentially the same as that of the Dr. Franz. The opinions set forth in this Declaration would be the same under either my or Dr. Franz's proposal.

20. I am well-qualified to determine the level of ordinary skill in the art and am personally familiar with the technology of the '825 patent. I was a person of at least ordinary skill in the art at the time of the priority date of the '825 patent. Regardless if I do not explicitly state that my statements below are based on this timeframe, all of my statements are to be understood as an ordinary artisan would have understood something as of the priority date of the '825 patent.

B. Legal Principles

21. I am not a lawyer and will not provide any legal opinions. Though I am not a lawyer, I have been advised that certain legal standards are to be applied by technical experts in forming opinions regarding the meaning and validity of patent claims.

1. Obviousness

22. I understand that to obtain a patent, a claimed invention must have, as of the priority date, been nonobvious in view of prior art in the field. I understand that an invention is obvious when the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art.

23. I understand that to prove that prior art, or a combination of prior art, renders a patent obvious, it is necessary to: (1) identify the particular references that singly, or in combination, make the patent obvious; (2) specifically identify which elements of the patent claim appear in each of the asserted references; and (3) explain how the prior art references could have been combined to create the inventions claimed in the asserted claim.

24. I understand that a patent composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art, and that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art to fit the parameters of the patented invention.

25. I also understand that a reference may be said to teach away when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent

from the path that was taken by the applicant. Even if a reference is not found to teach away, I understand its statements regarding preferences are relevant to a finding regarding whether a skilled artisan would be motivated to combine that reference with another reference.

2. *My Understanding of Claim Construction Law*

26. I have been informed that patent claims are construed from the viewpoint of a person of a ordinary artisan and that patent claims generally should be understood consistent with their ordinary and customary meaning at the time of the invention. A review of the patent claim language, the patent specification, and its prosecution history are also necessary to determine the proper meaning and scope of the term at issue.

27. I have further been informed that in the specification and prosecution history the patentee may define a claim term in a way that differs from the ordinary and customary meaning. I understand that during prosecution of the patent before the U.S. Patent and Trademark Office, the Applicant may make representations or provide definitions of terms that may affect the scope of the patent claims. In particular, the Applicant may, during the course of prosecution, limit the scope of the claims to overcome prior art and/or disavow claim coverage by making clear and unambiguous statements to that effect.

28. I have been informed that a ordinary artisan may, among other things, consider dictionaries, publications, other patents, and treatises (*i.e.*, “extrinsic evidence”). I understand that extrinsic evidence may generally not be relied on if it contradicts the meaning of claim language provided by the intrinsic evidence, such as express definitions given for terms in the specification.

III. OVERVIEW OF THE '825 PATENT

A. The '825 Patent Specification


29. The '825 patent, entitled “Reducing Redundant Operations Performed by Members of a Cooperative Security Fabric” was filed June 27, 2017 and claims priority to a provisional application filed on May 8, 2017. Ex. 1001.

30. The '825 patent explains that “[i]n a large network” with “many network security appliances,” each appliance may “work independently and execute operations” on “network traffic transmitted to/from the network” that may be wasteful and redundant. Ex. 1001, 1:33-37, 1:45-52.

31. The '825 patent solves this issue by using a “cooperative security fabric (CSF)” (a product that Patent Owner introduced to the market a few months before the '825 patent’s earliest filing date) with interconnected network security appliances that can “coordinat[e] security operations” using flags carried by incoming/outgoing network traffic to indicate which operations have already been

executed by an appliance, and further indicate whether the packets were transmitted by participating appliances in the CSF.

32. Figure 2 of the '825 patent is an example topology of a CSF. The “root node ... is the outermost network security appliance of the cooperative security fabric [and] acts as a WAN role in topology 200 that connects a private network to other networks or the Internet.” Ex. 1001, 5:52-56. There are also branch nodes that constitute a middle layer and leaf nodes that connect endpoint devices to the network. *Id.*, 5:60-6:5. As is evident, there are several different paths that network traffic could take, thereby requiring different coordination depending on the path. The '825 patent explains that a topology may have multiple root nodes, multiple layers of branch nodes, and paths that do not involve branch nodes (such as the red path below). *Id.*, 6:14-24. The below figure is a modified version of Figure 2 which shows that each leaf node is a network security appliance for a subnetwork of endpoints.



33. Thus, in the '825 patent, when coordinating operations, the network security appliances operate on “downstream network traffic that is transmitted *from the Internet 110 to a subnet*” or “upstream network traffic that is transmitted *from [a] subnet ... to the Internet 110.*” *Id.*, 5:33-38. The network appliances will inspect network packets for flags that, if present, indicate that the flag “was added to the original network packet by a previous member node along the transmission path ... to indicate that the network packet is transmitted from a member node” as well as “information about the distribution of security operations among member nodes of a CSF.” *Id.*, 6:53-7:4. However, if the flag is **not** present, “it means that

the member node is the first member of the CSF to receive the incoming packet” which “happen[s] to root nodes for downstream network traffic and leaf nodes for upstream network traffic.” *Id.*, 7:26-30.

IV. OVERVIEW OF CERTAIN ART

34. Although the petition cites to a number of references, the main references it relies on for its challenge against the independent claims of the ’825 patent are Chandra (for Grounds 1-4) and Keohane (for Grounds 5-8). The petition’s first ground of challenge asserts Chandra alone renders obvious the independent claims (as well as some of the dependent claims). The petition’s fifth ground of challenge asserts that Keohane alone renders obvious the independent claims (as well as some of the dependent claims). Below, I summarize Chandra and Keohane.

A. Chandra (Ex. 1004)

35. Chandra is directed to a network where “data is transferred between network devices,” *i.e.* network traffic only within a network. Ex. 1004, ¶4. Chandra notes that, with respect to a network involving multiple VPN sites, the VPN firewall policies may “be uniform among all the VPN sites.” *Id.*, ¶6. However, “[t]his can result in the same action being performed by more than one VPN site.” *Id.*

36. To resolve this specific issue of redundant actions performed on internal network traffic, Chandra describes a system involving coordination between “a transmitting Customer Edge (CE)” and a “receiving CE,” where the transmitting

CE and receiving CE can be “for example, Intrusion Detection Systems (IDSs), firewalls, etc.” *Id.*, ¶30. In other words, the transmitting CE may be a firewall for the VPN site that receives traffic from the transmitting endpoint, and the receiving CE may be a firewall for another VPN site that sends traffic to the receiving endpoint. See Fig. 2.




FIG. 2

37. Transmitting CE and receiving CE coordinate which actions to perform on the network traffic by using labels that are used to tag the data packets. *Id.*, ¶¶31-32. Prior to that, the transmitting CE and receiving CE will first “authenticate each other” and exchange the “CE labels that can be used for tagging the data packets.” *Id.*, ¶33.

38. Thus, as shown above, Chandra only discloses coordination for traffic between two nodes *within the same network*, not traffic coming into or going out of the network. *Id.*, ¶21. For incoming traffic that originates from outside the network, such as the Internet, only the “receiving CE” receives such traffic and thus only the receiving CE will be able to perform actions on such traffic. *Id.*, ¶53; *see also* ¶77 (“[R]eceiving end can identify a data packet from the Internet as insecure, and perform rigorous actions on it.”).

B. Keohane (Ex. 1008)

39. Keohane is directed to a method for performing iterative security operations on incoming network traffic in a single direction. Ex. 1008, ¶¶33-37. Keohane’s method is depicted in Figure 3 below.




FIG. 3
(Annotated)

40. As shown in the figure, Keohane screens incoming data through a Firewall (300) and a Mail Gateway (302) before ultimately allowing the data to be transmitted to the packet.

41. Keohane's data screen methods involve the attachment of a signature located in the header of a data packet by an upstream network element in order for a downstream network element to be able to confirm the traffic was received from the upstream network element. *Id.*, ¶11. Thus, as depicted in Figure 3, a data packet travels in **one direction** from network element "firewall 300," to "mail gateway 302," and to "client 304," with each network element performing a security action on the data and inserts a "signature ... indicat[ing] that a previous security action has been performed on the data packet." *Id.*, ¶¶32-37.

V. CLAIM CONSTRUCTION

42. In my opinion, all claim terms should be given their ordinary meaning. I understand that Petitioner and Dr. Franz also agree that all the terms should be given their ordinary meaning. Pet. 14; Ex. 1003, ¶53.

43. In my first declaration in support of Patent Owner's Preliminary Response, I opined that the term "fabric" "requires network nodes that process bidirectional communications." Ex. 2001, ¶¶43-51. I understand that while the Board credited my testimony and the evidence I cited to, it construed "fabric" as "a network topology such as the physical structure of a switch or network." Paper 10

at 9. I understand that Petitioner, through its expert Dr. Franz, has broadly interpreted the Board’s construction to cover **any** network configuration, including two computers connected by a single cable or one computer sending unidirectional data to a second computer. *See* Ex. 2005, 14:18-24 (“[T]he network fabric in this case is just these two boxes with the cable.”); 17:13-18 (opining that two computers communicating in one direction is still “two computers networked together”). This is in conjunction with Petitioner’s position that the larger phrase “cooperative security fabric” can be met by just two devices connected directly to each other. Pet. 27 (arguing that “as few as two [appliances] are needed” to “form a CSF”). I note that Dr. Franz admitted he lacked familiarity with the field of “distributed computing,” which is a key technology of the ’825 patent, which describes managing the “**distribution** of security operations among member nodes of a CSF.” Ex. 1001, 7:2-4, 8:50-52; Ex. 2005, 15:4-6, 16:1-7. The CSF can be viewed as distributed computing because the execution of network security operations are distributed within the nodes of the CSF—each node doing one of the requisite security operations.

44. As I discuss below, in view of the context of the ’825 patent and the plain meaning of a “cooperative security fabric,” I believe that Petitioner’s broad interpretation of the Board’s construction is problematic and not supported by the evidence. A “fabric” is analogous to clothing, which is in fact where the term comes

from—it refers to multiple, interwoven threads, or paths. However, Petitioner and Dr. Franz are interpreting it to include a single thread, which is contrary to the plain meaning of fabric and CSF. The '825 patent is directed not just to any fabric, but a specific type of fabric called a “cooperative security fabric,” which, as the patent explains, specifically addresses challenges with coordinating security operations on incoming or outgoing traffic within “a larger network” with “many network security appliances.” Ex. 1001, 1:33-37, 5:38-40. In such a network, there is “a need for a cooperative security fabric (CSF) that may coordinate operations performed on network traffic to avoid or reduce redundant operations among members of the CSF.” *Id.*, 1:64-67. Thus, I discuss below what the plain meaning of a CSF is based on the intrinsic and extrinsic evidence.

45. In my opinion, the plain meaning of a CSF is “a topology of network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.” Petitioner and Dr. Franz’s interpretation, however, is overly broad, removes any meaning of the term within the claims, and is contrary to the very purpose of the '825 patent’s cooperative processing.

A. Intrinsic Evidence

46. I begin with the intrinsic evidence, and specifically the specification of the '825 patent. The specification supports my definition of a CSF.

47. The '825 patent is directed to cooperative security within a “large network” with “many network security appliances ... deployed at different locations within the network.” Ex. 1001, 1:33-37.

48. In its earlier days, computer networks were straightforward. For example, two computers could be connected directly to each other (called a point-to-point topology), multiple computers could have access to the same communication medium (called a bus topology), or they could be connected to each other in a ring (called a ring topology). In these topologies, there were clearly defined paths between individual computers. Over time, however, networks grew more complex as enterprises increased the number of offices, locations, departments, functions, etc., that needed to be connected to each other via a network.

49. By 2017, a “network fabric” was a term of art. As enterprise networks became larger and larger, a complex, interconnected topology known as a “fabric” was developed to handle the management of such networks. In particular, enterprise networks needed to support increasingly numerous subnetworks as a result of, for example, increases in the enterprise’s number of office locations, departments, enterprise functions, etc. Thus, a fabric of network devices (such as switches or access points for each subnetwork) was developed to manage network traffic between subnetworks, and to and from the Internet. In contrast to more traditional network topologies, such as point-to-point, bus, and ring topologies, a “fabric”


topology consists of multiple paths to and from endpoints. The resultant cabling resembles interwoven threads in a piece of cloth—i.e., a “fabric.”

50. The ’825 patent sought to improve such fabrics, by using a specific type of fabric called a “cooperative security fabric.” The patent notes that within a network fabric, “[n]etwork traffic transmitted to/from the network may go through multiple network security appliances along a path within the network.” Ex. 1001, 1:33-37. In the prior art, the network security appliances “work[] independently,” which possibly leads to an appliance performing a “redundant” operation that have “already been done by a previous network security appliance at a previous hop.” *Id.*, 1:46, 1:54-56. The ’825 patent therefore notes a “need for a cooperative security fabric (CSF) that may coordinate operations performed on network traffic,” i.e. traffic “transmitted to/from the network,” “to avoid or reduce redundant operations among members of the CSF.” *Id.*, 1:64-67.

51. In a simple network, where network traffic traverses the same path of network appliances, no coordination is necessary, even if that path traverses multiple network security appliances. In such a simplistic topology, each network security appliance could be configured to always perform the same subset of security operations, thereby eliminating redundancy even without cooperation.

52. A network fabric, however, is more complex. There are many different paths of network appliances through which traffic can flow, and thus, cooperation is

needed to avoid redundant operations. Figure 2 of the '825 patent depicts a topology for a cooperative security fabric, which encapsulates this understanding of a fabric. I have provided below a modified version of Figure 2 where each leaf is further connected to one or more endpoints, such as the endpoints shown in Figure 1 (see Ex. 1001, 6:1-3), and further annotated it to show three different paths through the network.



53. In this example, all three transmission paths traverse Root node 220, but diverge thereafter. In the blue and green transmission path, traffic goes through the same Root node 220 but then through different Branch nodes 230a and 230b, and

different Leaf nodes 240a and 240d. In addition, the blue transmission path is shown as being “downstream,” *i.e.*, from the Internet to an endpoint, whereas the green transmission path is shown as being “upstream,” *i.e.*, from an endpoint to the Internet. See Ex. 1001, 5:32-38 (“[D]ownstream network traffic that is transmitted from the Internet 110 to a subnet 141 goes through three network security appliances, *i.e.* a root, a middle and a subnet network security appliance, while upstream network traffic that is transmitted from subnet 141 to the Internet 110 goes through a subnet, a middle and a root network security appliance.”).

54. The red transmission path also differs from the blue transmission path as it has “no branch nodes.” *Id.*, 6:18-20. Finally, while not pictured, “Topology 200 may have multiple layers of branch nodes, *i.e.*, multiple branch nodes exist between a root node and a leaf node” and nodes that act “as multiple roles in the topology 200,” such as “a network security appliance [that is] connect[ed] to another network security appliance as well as endpoint devices,” thereby acting as “a branch node and a leaf node in the same time.” *Id.*, 6:21-27.

55. It is this complexity and existence of multiple possible paths through a network “fabric” that creates “a need for ... coordinat[ing] operations performed on network traffic to avoid or reduce redundant operations” as traffic traverses any one of multiple potential paths through the network fabric. Ex. 1001, 1:64-67. For example, in the annotated version of Figure 2 above, traffic arriving at a leaf node

may have traversed zero network security appliances (green upstream path), one network security appliance (red downstream path), or two network security appliances (blue downstream path). In this scenario, a mechanism is needed to dynamically determine what security operations have already been performed and what additional security operations still need to be performed. Thus, the '825 patent describes that each “**member node** of the cooperative security fabric may add one or more flags ... to a network packet,” which is then “transmitted **to another member node** at the next hop.” Ex. 1001, 6:28-34. The member node that receives the packet “can be **any node** of the CSF.” *Id.*, 6:50-51.

56. Based on the above, the '825 patent describes that a CSF inspects incoming and outgoing traffic along multiple interconnected paths.

57. There are a few additional points to make. I note that in its Institution Decision, the Board stated that the '825 patent “describes methods of performing security operations on network traffic **passing through** a cooperative security fabric (CSF) of network security appliances that coordinate their security operations to avoid redundant security operations.” Paper 10 at 3. I agree with the Board’s description, because the CSF inspects traffic that passes through the network, *i.e.* traffic transmitted into or out of the network. I also note that Dr. Franz testified that he agrees with this description. Ex. 2005, 19:10-16; *see also* 26:4-7 (“[T]he CSF has the ... objective of coordinating security operations on the network traffic passing

through them.”). Dr. Franz further confirmed that the ’825 patent does not provide any examples of “performing security operations on traffic between two subnets” of the same network. *Id.*, 35:7-17.

58. Dr. Franz was also asked about the ’825 patent’s disclosure that “a subnet 151 and a subnet security appliance 150 [that] are connected to private network 120 through the Internet 110 via VPN” and thus “may be seen as a part of private network 110.” Ex. 1001, 4:65-5:2. He described this disclosure as “subnetworks being able to communicate with each other.” Ex. 2005, 74:14-23. However, the ’825 patent does not describe that such VPN traffic to other endpoints within the network is inspected by the CSF.

59. A POSITA would understand that a VPN is considered to be part of the same network it is connected to. Ex. 2005, 42:12-43:5 (“[O]ne of the purposes of VPN technology is that from the perspective of the two endpoints ... they appear to be connected to the same virtual private network.”); Ex. 1001, 4:65-5:2 (“[A] subnet 151 and a subnet security appliance 150 [that] are connected to private network 120 through the Internet 110 via VPN ... may be seen as a **part of private network 110.**”).

60. But the ’825 patent is concerned with and directed to “downstream network traffic that is transmitted from the Internet 110 to a subnet,” or “upstream network traffic that is transmitted from [a] subnet ... to the Internet.” Ex. 1001,

5:32-36. It is this traffic that the CSF must operate on, and it does not include VPN traffic since that is considered to be internal network traffic.

61. I have also reviewed the provisional application which the '825 patent claims priority to, which is Exhibit 2016. The provisional application provides further support for my definition of a CSF.

62. The provision states that the “dynamic formation of a CSF among multiple network security devices within a private network ***based on hierarchical interconnections among the network security devices.***” Ex. 2016, ¶13. The “hierarchical interconnections” involve different levels of network security devices, and in one embodiment, “a downstream node would not be permitted to query an upstream node (e.g., a parent node or a higher level node in the hierarchy of the CSF) or a node at the same hierarchical level within the CSF.” *Id.*, ¶18.

63. In other words, the provisional describes a hierarchy of network security devices where there are multiple devices on each hierarchical level, and where each device has its own downstream nodes. This means that the CSF constitutes multiple, interconnected paths.

B. Extrinsic Evidence

64. The extrinsic evidence I have reviewed also supports my construction of a CSF as “a topology of network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.”

65. The categories of documents I reviewed were Fortinet documentation as well as evidence regarding the term “fabric.”

1. Fortinet Documentation

66. First, I note that the term “cooperative security fabric” is a term that was introduced by Fortinet in 2016. Because Fortinet is a well-known company that develops and sells network security devices, a POSITA would have been familiar with their products and documentation.

67. In fact, a POSITA would have known that the best (and possibly only) evidence of the plain meaning of “cooperative security fabric” would be Fortinet documentation. I have attempted to look for other uses of the term “cooperative security fabric” but could not find any that did not otherwise come from Fortinet or were a reference to Fortinet products.


68. When performing a search for the phrase “cooperative security fabric” on Google Books, no results were returned. I also performed a search on Google Scholar for any academic papers containing the phrase, and it did not return any results related to network security. When searching for patents or patent applications using Google Patents, the only results that contain “cooperative security” were those assigned to Fortinet, the earliest of which was the ’825 patent. I also used the USPTO and EPO searching sites and likewise only found Fortinet-owned patents and applications.

69. Finally, I performed a search on Google for the term “cooperative security fabric.” I reviewed the first dozen or so results and each result was either Fortinet documentation or describing a Fortinet product. Thus, I conclude that “cooperative security fabric” was introduced to the field of network security by Fortinet via its products, and thus the best (and only) extrinsic evidence for its plain meaning would be found in Fortinet documentation.

70. One document that was published in June 8, 2016 is entitled “Cooperative Security Fabric.” The document describes the steps for “configur[ing] a Cooperative Security Fabric (CSF) ... throughout your network, using a range of Fortinet products.” Ex. 2012, 1. It instructs that the first step is to “install the initial FortiGate which will later be used as the **Internet-facing, or upstream**, FortiGate in the security fabric.” *Id.*, 2. The next step is to “[i]nstall[] internal FortiGates and enabl[e] a security fabric,” where in the example provided in the documentation, “two additional FortiGates are added to the network” and once installed, “a security fabric is set up between them **and the external FortiGate.**” *Id.*

71. This second step includes a link to another document which provides additional detail for setting up the internal FortiGates. Ex. 2013. This document provides a diagram of a simple CSF that depicts the two internal FortiGates (Accounting and Marketing) being routed through the Internet-facing External

FortiGate, thus indicating that network traffic travels through multiple interconnected paths. *Id.*, 1.



72. The document describes how to install the internal FortiGates and enable them as a CSF. Once the internal FortiGates are installed, the CSF can be configured to “allow you to distribute security functions to different FortiGates in the security fabric,” such as “implement[ing] virus scanning on the External FortiGate but add[ing] application control and web filtering to the ISFW [internal] FortiGates.” Ex. 2013, 9. This “results in distributed processing between the FortiGates in the CSF; reducing the load on each one.” *Id.* In other words, the CSF described in these documents uses FortiGates that cooperate to inspect incoming and outgoing network traffic along multiple interconnected paths, as reflected in the plain meaning for “cooperative security fabric.”

73. Exhibits 2012 and 2013 are true and correct copies of web pages that were found on the Internet Archive’s Wayback Machine. I am familiar with the

Wayback Machine. It is a service that is operated by Internet Archive and which is intended to archive web pages. It involves the use of crawlers that surf the Internet and make copies of web pages it encounters. For each archived copy, the Wayback Machine will indicate the date and time that the crawler made the archived copy. This can be found in the URL of the archived page. Thus, for Exhibit 2012, which has as part of its URL the number “20160728,” this indicates the crawler crawled the web page and made an archived copy on July 28, 2016. Similarly, for Exhibit 2013, whose URL has the number “20160828,” this indicates an archival date of August 28, 2016. As a result, Exhibits 2012 and 2013 reflect the authentic copy of those web pages that were publicly available on the Fortinet website as of 2016.

2. Other Extrinsic Evidence

74. I also reviewed a number of documents related to the term “fabric” which also supports the plain meaning of cooperative security fabric. These documents establish that a fabric includes network security appliances arranged in multiple, interconnected paths, and not just a single path.

75. One example, which the Board credited, is the 30th Edition of Newton’s Telecom Dictionary (2016), which defines “fabric” as “[a] descriptive term referring to the physical structure of a switch or network,” where “***[m]uch like a piece of cloth, physical/logical communications channels (threads) [plural] are interwoven from port-to-port (end-to-end).***” Ex. 2003. In other words, a “fabric” consists of

multiple, interwoven communications paths/channels (which are analogous to individual threads within a piece of cloth) that define the network paths from one end to another—much like with clothing, a single thread does not constitute a fabric.


76. Dr. Franz agreed that this dictionary definition “reflects the ordinary and customary meaning of fabric” (Ex. 2005, 29:13-16) and that the term “fabric” was inspired by clothing fabric (Ex. 2005, 7:25-8:8 (“[T]hat’s probably where the analogy comes from.”)).

77. I note that in its Institution Decision, the Board made use of the first sentence of this definition (“A descriptive term referring to the physical structure of a switch or network.”) but did not include the second sentence (“Much like a piece of cloth, physical/logical communications channels (threads) are interwoven from port-to-port (end-to-end).”). I think the entire definition should be considered, especially in view of the additional extrinsic evidence I discuss below. The second sentence provides important context that describes what type of network topology qualifies as a fabric.

78. Another document, a 2002 book titled “Using Storage Area Networks,” identifies three kinds of “Topologies”: point-to-point, loop, and switched fabric. Ex. 2006, 145. Whereas a “point-to-point” topology is “a simple and easily applied topology for connecting two nodes” which is “best used on a small scale,” a “switched fabric topology is appropriate for **large-scale or enterprise networks** in

which *a number of users are interconnected.*” *Id.*, 146. The “fabric is used to connect networks across large distances” and “works in a similar fashion to a switched telephone network in that it can provide a dedicated connection between multiple pairs of nodes.” *Id.*

79. Similarly, a 2013 networking book depicts a “Fabric Based” network as a collection of multiple interconnections between endpoints and the edge nodes of the network. Ex. 2007, 143.



80. The same book also depicts a “Traditional three-tier Network **Fabric** Architecture,” with “access switches” (*i.e.* subnetwork switches) that “are connected

Fortinet Ex. 2004, Page 30

via Ethernet to aggregation switches,” and then “aggregation switches [that] are connected into a set of core switches or routers that forward traffic flows from servers to an intranet and internet.” *Id.*, 145. This fabric is depicted below, and shows multiple paths to and from Access-layer switches.




Figure 3-9 Traditional three-tier Network Fabric Architecture

81. Dr. Wenke Lee, who is Petitioner’s expert in related cases, also agrees that the plain meaning of “fabric” requires multiple paths to and from endpoints. For example, in IPR2023-00030, Dr. Lee testified at his deposition that a fabric is “**multiple** security elements and controls that one would use or one would want to have,” with firewalls (which controls traffic into and out of a network) being “part of the fabric.” Ex. 2009, 109:4-6. And in IPR2023-00175, in describing U.S. Patent No. 8,000,329 (“Fendick”) (which was filed June, 29, 2007), Dr. Lee testified that a

“fabric” is “basically, a communications hub” “for connecting the platform components,” and provided the below annotated figure depicting how a fabric provides interconnected paths between multiple components, including the components that provide ingress/egress. Ex. 2010, ¶60.




Fig. 2
(Annotated)

82. In fact, the Fendick reference explains that a “fabric” “facilitates *multiple parallel communication sessions* of traffic, as well as, permits multiple entities, such as control elements 212 and service cards 202, to communicate with each other in a simultaneous fashion.” Ex. 2011, 7:42-45.

83. In my opinion, the above extrinsic evidence for “fabric” establishes that the plain meaning of “cooperative security fabric” includes network security appliances with multiple interconnected paths.

C. Addressing Dr. Franz’s Testimony

84. During his deposition, Dr. Franz provided testimony that indicates he would likely not agree with certain aspects of the plain meaning of a CSF. Specifically, Dr. Franz testified that (1) a “fabric” is merely “the physical structure of a network,” and “is how ... you physically realize the topology” (Ex. 2005, 7:18-19, 11:3-4), and (2) the term CSF is simply a “compound ... term” whose plain meaning is merely defined by the plain meaning of its constituent parts, that is, it “coordinate[s] security operations on the network traffic passing through” its network appliances (*id.*, 7:6-15, 25:25-26:7).

85. Dr. Franz believes that “every network has a fabric,” because every network has physical structure. *Id.*, 12:3-9. Thus, under his definition, even “two computers [connected] with each other via an ethernet cable,” *i.e.*, “just these two boxes with the cable,” qualify as a “fabric.” *Id.*, 14:14-24; *see also* Pet. 27 (“[A]s few as two [network appliances] are needed” to form a CSF). Under Dr. Franz’s definition, even a first computer sending unidirectional data to a second computer meets the “fabric” claim limitation. Ex. 2005, 17:13-18. I respectfully disagree with Dr. Franz. I do not think a “fabric” can be defined so broadly. What Dr. Franz

describes is a “thread,” whereas a “fabric” (like its clothing analogue) requires multiple threads, *i.e.* multiple paths of network appliances.

86. In my opinion, Petitioner has no intrinsic support for its theory that “as few as two [network appliances] are needed” to form a fabric or CSF. Pet. 27. Petitioner cites ’825 patent at col. 6:18-20 in support of this assertion, but that passage states that one of ***multiple “paths” (plural)*** in topology 200 may pass through just two nodes. Ex. 1001, 6:18-20. Thus, the ’825 patent teaches that a CSF topology consists of multiple paths to and from endpoints, and if anything emphasizes that one of the challenges with a “fabric” is that different paths may have a different number of nodes along the way—the very type of complexity of a fabric compared to a point-to-point connection that motivated the innovation of the ’825 patent. Thus, the cited disclosure of the ’825 patent is consistent with the rest of the specification in describing a network “fabric” as being applicable to “a large network” with “many network security appliances.” Ex. 1001, 1:33.

87. Also, in my opinion, Petitioner’s definition of “fabric” and CSF renders the terms superfluous within the claims of the ’825 patent. I understand that claim construction which renders a term superfluous should generally be avoided. For example, Claim 1 recites a second network security appliance that transmits network traffic to a first network security appliance. According to Dr. Franz, such a transmission is sufficient to meet the “fabric” claim limitation, even without the

additional claim language that requires both appliances to be part of the same “fabric.” Ex. 2005, 14:18-24. Claim 1 also recites coordination between the first and second network security appliances for performing network security operations on traffic. Thus, under Dr. Franz’s definitions, the limitation “cooperative security fabric” would have no effect on claim scope, rendering the term superfluous.

88. Petitioner’s definition is not consistent with the history of networks and the term “fabric.” Computer networks have been around for nearly a century, with the first computer networks being developed in the 1950s, when AT&T released the first commercial modem that allowed two or more computers to communicate with one another. Network fabrics were not developed until many decades later when enterprise networks became significantly larger and administratively complex—the term does not regularly appear in the literature until the 2000s. A POSITA would therefore not equate a fabric with the topology of any network. Instead, a POSITA would consider a “fabric” to be a complex network topology with many paths to and from endpoints, because the term’s emerged in the context of larger, more complex networks.

89. I also disagree with Dr. Franz that the plain meaning of “cooperative security fabric” can be defined by its constituent parts, *i.e.* the plain meaning of “cooperative,” “security,” and “fabric.” As I explained above, the term was introduced by Patent Owner, first found in Patent Owner’s products, and only

appears in Patent Owner's documentation and patents. As a result, a POSITA would have known that the plain meaning of a CSF is found in Fortinet's patents, documents, and products.

90. For these reasons, it is my opinion that the plain meaning of CSF is “a topology of network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.”

VI. GROUND 1: CHANDRA DOES NOT RENDER OBVIOUS CLAIMS 1-5, 12, 15-16, 21-22, AND 27

91. In my opinion, Petitioner has not shown that Chandra renders claims 1-5, 12, 15-16, 21-22, and 27 obvious for a few independent reasons. Chandra does not disclose a “cooperative security fabric,” as claimed. For claims 2, 4, and 15, Chandra also does not disclose or render obvious that actions are taken based on “local network policies.” Finally, for claim 12, Petitioner provides no rationale to include flags in a “first packet of a session.”

A. Independent Claims 1, 15, 21, and 27: Chandra does not disclose a “cooperative security fabric”

92. In my opinion, Petitioner has not shown that Chandra discloses a “cooperative security fabric,” which as I explained above has a plain meaning of “a topology of network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.” Section V.

93. Instead, Chandra is quite clear that to the extent two appliances cooperate to inspect traffic, it is only for traffic between a transmitting and receiving device ***within the same network***, and not for incoming or outgoing traffic. Ex. 2005, 41:13-16 (“[B]oth the transmitting end and the receiving end are connected to the same network 102”). Thus, in Chandra, “network devices 104 [are] connected to network 102,” which “function[s] as a single larger network.” Ex. 1004, ¶21. Chandra describes that, when coordinating network actions, there is a “transmitting end” and “receiving end,” which Dr. Franz described as “endpoints.” *Id.*, ¶24; Ex. 2005, 42:24-43:5. Both the transmitting end and the receiving end have a corresponding “network device” which is one of the network devices within network 102, and which are connected via a “network link 218.” *Id.*, ¶¶26-27; *see also* ¶29 (“transmitting end 202 and receiving end 204 can be any identifiable part or region of network 102”). Each data packet in Chandra originates from “transmitting network device,” then a “transmitting customer edge” or “CE” performs a “first set of actions,” the packet is then transmitted over network 102 to a “receiving customer edge that performs a “second set of actions,” and finally arrives at the receiving network device. But in this context, the transmitting and receiving CE only process internal network 102 communication and thus do not “cooperate to inspect incoming and outgoing traffic,” which is part of the plain meaning of “cooperative security fabric.”

94. Chandra does describe how incoming traffic (such as from the Internet) is inspected. Specifically, only one customer edge (a receiving CE) inspects such traffic and will perform all the necessary actions. For example, “when receiving CE 214 receives a data packet from the Internet, it performs actions corresponding to rules regarding such data packets.” *Id.*, ¶53; *see also* ¶77 (“[R]eceiving end can identify a data packet from the Internet as insecure, and perform rigorous actions on it.”). As Dr. Franz explains it, “[i]f the data packet did not have a recognized label (e.g., ***the traffic originated from the Internet***), the receiving CE identified the packet as non-secure and performed a third set of actions on the packet.” Ex. 1002, ¶96. This is because such a packet “was not transmitted from the transmitting CE.” *Id.* Thus, for incoming traffic, Chandra does not disclose inspection using a “cooperative security fabric,” because only one appliance inspects the traffic and it does not coordinate that inspection with other appliances.

95. Chandra also discloses in connection with Figure 8 that the transmitting CE and receiving CE can be connected via a VPN connection, but that scenario also fails to disclose the claimed “cooperative security fabric.” As addressed above in connection with Patent Owner’s construction of “cooperative security fabric,” VPN sites are considered to be part of the same network. Thus,

any traffic between VPN sites is purely internal network traffic and fails to meet the requirement that a CSF inspects incoming or outgoing traffic.

96. The same reasoning applies here. Chandra’s transmitting and receiving CEs connected via VPN are still part of the same network 102—that is indeed the purpose of a VPN. Ex. 2005, 42:11-43:5. Under that view, even though the transmitting and receiving CEs are physically remote from each other, they are *virtually* in the same network and only cooperate to process traffic within the same network 102. Thus, they do not cooperate to “inspect *incoming and/or outgoing network traffic*,” as required of a “cooperative network security fabric.”

97. Alternatively, if traffic received over a VPN connection is indeed considered “incoming and/or outgoing network traffic,” then that traffic is inspected on each end by just a single transmitting or receiving CE, and thus does not constitute multiple cooperating network security appliances. Regardless of how Chandra’s disclosure is viewed, Chandra lacks a “cooperative security fabric” as defined by the ’825 patent.

B. Dependent Claims 2, 4, and 15: Chandra does not disclose “local network policies”

98. Dependent claims 2 and 4 and independent claim 15 recite that network security operations are executed “based on local network policies.” In my opinion, Chandra does not disclose local policies that define which network actions to take, and Petitioner does not explain or provide evidence for its proposition that

“[a] POSITA would find it obvious that Chandra contained two sets of local policies: one local to receiving CE and another local to transmitting CE.” Pet. 38.

99. First, Chandra does not disclose local policies at the CEs. Chandra discloses that when a receiving CE receives a labeled data packet, it will forego performing a “first set of actions” that were already performed by a transmitting CE. *See, e.g.*, 1004, ¶32. The receiving CE will instead “determine[] a second set of actions” and perform those actions. *Id.* However, Chandra never states that the second set of actions is determined according to a policy that is local to the receiving CE.

100. Instead, Chandra discloses that actions are performed according to a **global** policy, or at least, a policy that is shared by endpoints. Chandra discloses that its system is operating within a network where “policies [are] uniform among all the VPN sites within a VPN.” Ex. 1004, ¶6. Moreover, in the examples provided by Chandra, transmitting CE and receiving CE have the same set of labels and rules. *Id.*, ¶¶51, 53. Chandra, in fact, consistently describes the “second set of actions” performed by receiving CE as a set of actions “that **has to be** performed on the data packet.” *Id.*, ¶¶18, 32, 46, 67. A POSITA would understand this to mean that there is a global policy that defines which actions must be performed on the data packet.

101. Although Petitioner argues that it would be “obvious that Chandra” has policies local to receiving CE and transmitting CE, they do not provide any further explanation or evidence for their statement. I cannot discern why Petitioner (or Dr. Franz) believes this would be obvious, especially when Chandra does not need local policies for its system that shares labels and rules.

C. Dependent Claim 12: Chandra does not disclose or render obvious that flags are carried in “a first packet of a session”

102. Claim 12 recites “one or more packets in which the flag is carried comprises a first packet of a session.” Petitioner acknowledges that Chandra does not disclose that Chandra’s labels are carried in a “first packet of a session.” Pet. 43. Petitioner argues that using the first packet is obvious because “including the flag(s) in the first packet would be sufficient.” *Id.* Petitioner then goes on to explain how in a network with only one route between endpoints, “flags in a given session would be redundant after the initial packet [that contains the flags].” *Id.*

103. I disagree with Petitioner. While including flags may be sufficient, Petitioner does not provide any particular reason or motivation for why a POSITA would include flags in the first packet. Also, Chandra describes that before a labelled data packet is transmitted and received, the transmitting CE and receiving CE must first authenticate each other, then exchange labels. Ex. 1004, ¶34. Thus, within a session, those initial steps will be performed with the first few packets, rather than sending the label first.

VII. GROUND 2: CHANDRA IN COMBINATION WITH BURUGANAHALLI DOES NOT RENDER OBVIOUS 6-7, 17-18, 23-24, 28-29

104. The dependent claims challenged in Ground 2 require that the receiving network security appliance either update or remove flags before sending them to a “next network appliance” in the CSF. Chandra, however, only describes coordination between two nodes, and therefore does not describe a “next network appliance.” See, e.g., Ex. 1004, ¶¶31-32. I note that Dr. Franz acknowledged that Chandra’s receiving CE is an “endpoint,” meaning that at most only two CEs are involved in the coordination. Ex. 2005, 42:24-43:5.

105. I note that Petitioner provides several arguments that Chandra combined with Buruganahalli renders obvious a “next network appliance participating in the CSF.” My response to each one is below.

106. *First*, Petitioner cites Buruganahalli’s disclosure of processing packets by a sequence of three or more network devices. Pet. 53-54. But Buruganahalli only teaches processing packets along a single sequence of network device 54, followed by network device 52, followed by network device 50. There are no multiple paths for packets to traverse Buruganahalli’s network devices, and those devices are thus not part of a “fabric.” This also means that, adding Buruganahalli’s network devices to Chandra fails to add “a next network appliance *participating in the* [cooperative security *fabric*],” as claimed. Petitioner also

provides no reason for adding Buruganahalli's additional network devices to Chandra.


107. *Second*, Petitioner relies on Chandra's disclosure that "a plurality of such CEs [could] be associated with a single PE." Pet. 54 (quoting Ex. 1004, ¶28). Chandra, however, does not disclose that these plurality of CEs ever communicate with each other *in seriatim*. Chandra is clear that a transmitting CE sends a packet to a receiving CE, after which the chain ends. Ex. 2005, 69:15-16 ("Chandra **only** talks about packages going from **one** source to **one** destination."). That multiple CEs may use the same PE to communicate with other CEs does not change the fact that Chandra's communications are only between two CEs.

108. *Third*, Petitioner argues that "receiving CE (214) could transmit the data packet with updated label to receiving network device (212)." Pet. 54. However, in my opinion, "receiving network device (212)" is not a "network appliance participating in the CSF." The receiving network device is the destination endpoint for the transmission of data—by contrast, as discussed above, the nodes of a CSF are the components that make up the communication paths **between** endpoints, such as switches, routers, etc. The '825 patent makes a distinction between "network appliances" of a CSF and "endpoint devices." For example, the '825 patent teaches that "Leaf nodes 250 a-250 d are the network security appliances that are used for connecting endpoint devices of the private

network” and that “LAN ports of leaf node are connected to endpoint devices, such as personal computers, servers, laptops, tablet computers, smart phones, printers and etc.” Ex. 1001, 6:1-9.

109. Similarly, throughout Chandra, there is a distinction between CE’s (which can be “Intrusion Detection Systems (IDSs), firewalls, etc.”) that perform network actions on network data, and “network devices”, which are “personal computers, servers, mobile phones, etc.,” and which send and receive data. Ex. 1004, ¶¶21, 30. Thus, Chandra never discloses that “network devices” perform any network actions on network data. Indeed, Chandra discloses that the receiving CE performs all remaining actions “that ha[ve] to be performed on the data packet.” *Id.*, ¶32. The “receiving network device” therefore does not perform any actions on the data packet and is not an appliance that is part of the CSF, or cooperating with other appliances.

110. *Fourth*, Petitioner argues that “a POSITA would have found it obvious that Figure 8’s VPN could be attached as an extension of Figure 2’s private network,” and provides the below modified Figures.



111. Petitioner, however, does not explain why a POSITA would modify Chandra by combining Figures 2 and 8, or explain why a POSITA would reasonably expect it to succeed. Chandra clearly describes that in its system, there is a transmitting CE and a receiving CE without any indication that it would be beneficial to involve additional CE's and network devices as Petitioner proposes. Moreover, in Chandra's system, the CE's have to authenticate and exchange labels before data is transmitted between them. Ex. 1004, ¶¶33-34. Adding additional CE's would require an exponential increase in the exchange of such data—in Petitioner's proposed combination, Figure 2's transmitting CE would need to exchange labels not only with Figure 2's receiving CE, but also both Figure 8's transmitting and receiving CE; Figure 2's receiving CE would need to exchange labels with Figure 8's transmitting and receiving CE's, *etc.* Instead of one exchange as described in Chandra, Petitioner's proposal would involve six exchanges of labels, which is far less efficient.

112. In my opinion, Chandra discloses that transmitting provider edge 810 already acts as a gateway to the transmitting network device. In other words, if transmitting network device 206 were to transmit data to receiving VPN site 812, it would be able to do so directly through transmitting provider edge 810, rather than first going through transmitting provider edge 210 and receiving provider edge 216. Since Chandra already configures a single transmitting CE and a single receiving CE to perform all necessary actions, there is no need in Chandra for the same traffic to traverse an additional two CEs as proposed by Petitioner (receiving CE 214 and transmitting CE 808). Doing so, as proposed in Petitioner's combination, only leads to inefficiencies.

113. For these reasons, it is my opinion that Chandra does not disclose or render obvious a "next network appliance participating in the CSF."

VIII. GROUNDS 3 AND 4 FAIL FOR THE SAME REASONS AS GROUND 1

114. Grounds 3 and 4 of the petition challenge claims that are dependent on the independent claims challenged in Ground 1. Thus, they rely on the analysis of Ground 1. In my opinion, Grounds 3 and 4 fail for the same reasons I discussed above with respect to Ground 1.

IX. GROUND 5: KEOHANE DOES NOT RENDER OBVIOUS CLAIMS 1-5, 12, 15-16, 21-22, AND 27

115. In my opinion, Petitioner has not shown that Keohane renders claims 1-5, 12, 15-16, 21-22, and 27 obvious because Keohane does not disclose a “cooperative security fabric,” whose plain meaning requires “network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.” Keohane instead discloses a single communication path (i.e., a single thread rather than a fabric) that only handles traffic coming into a network.

A. Independent Claims 1, 15, 21, and 27: Keohane does not disclose a “cooperative security fabric”

116. In my opinion, Petitioner has not shown that Keohane discloses a “cooperative security fabric,” whose plain meaning is “a topology of network security appliances that cooperate to inspect incoming and outgoing traffic along multiple interconnected paths.” Instead, Keohane discloses a single path that only processes communications coming into the network.

117. Petitioner argues that in Keohane, “each participating appliance inserted a signature into a packet to indicate the security operations performed,” and that “[d]ownstream devices ... determined the security operations performed by prior devices., and thus are a “cooperative security fabric.” Pet. 71-72 (citing Ex. 1008, ¶11). But Keohane discloses that firewall (300) and mail gateway (302) are nodes that can only process traffic in one direction: “Data packet 306 is received by firewall 300. ... [F]irewall 300 allows data packet 306 to be passed to

mail gateway 302. ... Then, data packet 306 is passed to client 304.” Ex. 1008,

¶¶33-37.

118. Figure 3 depicts this one-way path, as annotated by Petitioner below.




FIG. 3
(Annotated)

119. As shown in the figure, the data packet 306 only flows in one direction along a single path, with firewall 300 acting as a dedicated data transmitter to mail gateway 302, which acts as a dedicated data transmitter to client 304. The coordination between these nodes only occurs for data that originates outside of the firewall and only along one path.

120. Dr. Franz confirmed that in Keohane, a packet is “first received by firewall 300” and then forwarded to “mail gateway 302,” and that this flow is “illustrative to the specific method that is described in [Keohane].” Ex. 2005,

59:3-12. Dr. Franz also confirmed that Keohane does not disclose “what else is happening in this network” besides this unidirectional flow along a single path. *Id.*

121. Cooperation among network security appliances to inspect traffic along a single path in a single direction is not a “cooperative security fabric” under its plain meaning. Rather than a series of interconnected paths, Keohane only describes a single path. Keohane also only describes inspecting incoming network traffic—there is no disclosure that Keohane’s Mail Gateway and Firewall cooperate to inspect outgoing network traffic.

122. Thus, in my opinion, Petitioner’s Ground 5 should be rejected because Keohane fails to disclose a “cooperative security fabric” under its plain meaning.

B. Dependent Claim 4: Keohane does not disclose or render obvious a first network security appliance that receives traffic “not transmitted from the second network security appliance”

123. Claim 1 recites a “first network security appliance” that receives network traffic, and can determine that the network traffic was transmitted by a “second network security appliance.” Claim 4 further requires that the first network security appliance can determine that the network traffic is “not transmitted from the second network security appliance,” and if it is not, executing security operations specified by a local policy.

124. In my opinion, Keohane does not disclose or render obvious this limitation because none of its alleged network security appliances can both

determine that traffic is from another network security appliance and determine that traffic is not from another appliance. What I mean is that Keohane's network appliances can at best determine that traffic comes from another appliance, **or** that traffic does not come from another appliance, but none of Keohane's network appliances can do both.

125. Petitioner alleges for claim 1 that Keohane's "mail gateway (302)" is a "first network security appliance." Pet. 69. However, Keohane discloses that the mail gateway always receives data from the firewall, which Petitioner alleges to be the "second network security appliance." Ex. 1008, ¶34, Fig. 3; Pet. 74. Thus, Keohane's mail gateway will never determine that the traffic did **not** come from another network security appliance and therefore does not disclose claim 4.




FIG. 3
(Annotated)

126. Petitioner also alleges that “the first network security appliance could also be any network appliance in the appliance chain (e.g., client (304)) that received network traffic.” Pet. 71. However, client 304 has the same issue—it only receives traffic from the mail gateway, an alleged network security appliance. Client 304 can never determine that traffic did not come from another network security appliance.

127. If firewall 300 is the alleged first network security appliance, then the opposite problem occurs. Firewall 300 never receives traffic from an alleged network security appliance. As a result, it cannot perform claim 1’s determination that the network traffic came from a second network security appliance, nor perform the steps that follow such a determination.

128. Although Petitioner relies on step 706 of Figure 7, which determines the presence of a signature header, such a determination does not determine if it was or was not “transmitted by an upstream device” as Petitioner claims. Pet. 85. This is because the signature only indicates “that a previous security action has been performed on the data packet.” Ex. 1008, ¶32. The lack of a signature could simply mean that the upstream device did not perform any security action—not that it did not come from an upstream device.

129. For these reasons, it is my opinion that Keohane does not disclose this limitation.

C. Dependent Claim 12: Keohane does not disclose or render obvious that flags are carried in “a first packet of a session”

130. Petitioner acknowledges that Keohane does not expressly disclose claim 12. However, Petitioner argues that claim 12 is obvious for the same reasons discussed in Ground 1 with respect to Chandra. Petitioner’s arguments therefore fail for the same reasons discussed above in Section VI.C.

X. GROUND 6: KEOHANE IN COMBINATION WITH BURUGANAHALLI DOES NOT RENDER OBVIOUS CLAIMS 6-7, 17-18, 23-24, 28-29

131. The claims challenged in Ground 6 require that the receiving network security appliance either update or remove flags before sending them to the next network appliance in the CSF. Keohane does not render these claims obvious because (1) for claims 6, 17, 23, and 28, a POSITA would not be motivated to use Buruganahalli’s updating of tags in Keohane, (2) for claims 7, 18, 24, and 29, a POSITA would not be motivated to remove Keohane’s signatures, and (3) Keohane does not disclose a “next network appliance participating in the CSF” and Petitioner provides no explanation for incorporating Buruganahalli’s “next network appliance.”

A. Dependent Claims 6, 7, 17, 18, 23, 24, 28, and 29: Keohane in view of Buruganahalli does not render obvious a “next network appliance”

132. These claims recite that the receiving network security appliance sends the network packet to a next network appliance that participates in the CSF.

Petitioner argues that “[i]n the Keohane-Buruganahalli combination, a POSITA would incorporate Buruganahalli’s disclosure of updating a tag into Keohane’s security element updating the signature and transmitting the updated signature to the next security element along the data path.” Pet. 94.

133. However, Keohane’s “next security element” is not a “next network appliance.” Under Petitioner’s invalidity theory, Keohane’s mail gateway in Figure 3 is the claimed “first network security appliance.” Pet 74. As depicted in Figure 3, reproduced below, the downstream element from the mail gateway is “client 304.” However, “client 304” is not a network appliance participating in the CSF. Petitioner also asserted in an offhand statement that any of the appliances could be a first network security appliance, such as the firewall. However, Keohane does not disclose that it is possible the firewall receives network traffic from another appliance—in fact, it does not. Thus, the firewall as a first network security appliance would not meet these limitations.




FIG. 3
(Annotated)

134. Client 304 is the destination endpoint for the transmission of data—by contrast, “network security appliances” are the components that make up the communication paths **between** endpoints, such as switches, routers, etc. In accordance with this plain meaning, the ’825 patent makes a distinction between “network appliances” of a CSF and “endpoint devices.” For example, the ’825 patent teaches that “Leaf nodes 250 a-250 d are the network security appliances that are used for connecting endpoint devices of the private network” and that “LAN ports of leaf node are connected to endpoint devices, such as personal computers, servers, laptops, tablet computers, smart phones, printers and etc.” Ex. 1001, 6:1-9.

135. Similarly, Keohane makes a distinction between “routers, gateways, and firewalls,” “servers,” and “clients,” the latter of which are “personal computers or network computers.” Ex. 1008, ¶23. As a result, Keohane’s clients are not network appliances nor are they part of a fabric under its plain meaning. For these reasons, Keohane does not disclose a “next network appliance.”

136. In addition, Buruganahalli, like Keohane, only discloses forwarding traffic along a single, unidirectional path of network devices. Such a single, unidirectional path is not a “fabric,” and thus Buruganahalli also cannot supply the missing “a next network appliance participating in the [cooperative security *fabric*],” as claimed.

B. Dependent Claims 6, 17, 23, 28: A POSITA would not be motivated to use Buruganahalli’s updating of tags with Keohane’s signatures

137. Recognizing that Keohane does not disclose updating signatures from packets, Petitioner argues that “[a] POSITA would have been motivated to add Buruganahalli’s feature of updating the metadata tag to Keohane,” and specifically combining them so that Keohane’s security elements “updated the signature before passing the data packet to the next element instead of adding a new signature.”

Pet. 93. In other words, Petitioner’s proposal is to use Buruganahalli’s teachings to somehow modify the digital signatures embedded by Keohane’s security elements into a data packet. Petitioner’s proposal is flawed because Buruganahalli does not

teach how to modify digital signatures and making Keohane's digital signatures modifiable would defeat their purpose.

138. Keohane discloses that a "signature from a trusted security element" "is a digital signature." Ex. 1008, ¶¶34-35. However, a POSITA would have understood that a purpose of a digital signature is to prevent tampering. This is because the very purpose of a signature is to authenticate that a message truly came from a certain sender—in Keohane's case, the signature also authenticates which security actions were performed by the sender. *Id.*, ¶34. If a signature could be easily altered, then the message could not be authenticated as coming from the sender.

139. For example, a well-known, conventional method for signing a message involves using a private-public key pair. The sender of the message will, for example, derive a hash of the message (*i.e.* using a one-way function that converts any message into an ideally-unique number), and then encrypt the message hash with the sender's **private key** (a key known only to the sender). The recipient can authenticate the sender and message by decrypting the hash using the sender's **public key** (a key known to the public) and comparing it to the hash of the message calculated by the recipient. The recipient knows that the sender truly sent the message because the sender's public key was able to decrypt the hash.

140. A POSITA would understand that the sender's signature could not be modified. This is because in order for the signature to indicate that it is the sender's signature, the sender's private key must be used. But the sender's private key is not known to anyone but the sender.

141. If the recipient were to instead sign the message using the recipient's private key, downstream elements would be unable to confirm that the packet came from the original sender, which in Keohane must be a "trusted security element," such as an upstream firewall. Ex. 1008, ¶¶35-36. Downstream elements would be unable to verify that the packet was first received by a firewall, or to verify that the firewall performed the security actions, such as ensuring that "the data packet 306 has originated from a trusted source." *Id.*, ¶34.

142. Buruganahalli provides no teachings about modifying a digital signature and Petitioner does not explain how such a modification would be performed, or why a POSITA would expect it to succeed. Indeed, somehow modifying Keohane to render its digital signatures modifiable, as proposed by Petitioner, would defeat the very purpose of those signatures.

143. For these reasons, a POSITA would not be motivated to look to Buruganahalli to modify Keohane's signatures.

C. Dependent Claims 7, 18, 24, 29: A POSITA would not be motivated to remove Keohane's signatures

144. Keohane does not disclose removing signatures from packets.

Petitioner argues a POSITA would be motivated to incorporate Buruganahalli's removal of a tag for two reasons. For the reasons below, I disagree with both reasons.

145. *First*, Petitioner argues that “removing the tag before passing the data packet would keep the network's labeling mechanism private, preventing malicious devices from intercepting the data packet and exploiting the labeling mechanism.” Pet. 95. But Petitioner provides no explanation for how a “malicious device[]” could “exploit[] the labeling mechanism.” I note that because Keohane's labeling mechanism uses digital signatures, such exploitation is difficult—the signature ensures that the label was added by a trusted device.

146. Also, in Keohane, the signatures provide an essential function which a POSITA would not be motivated to remove. For example, the signatures first indicate that the packet came from “trusted security elements.” Ex. 1008, ¶37. Second, the signatures indicate which security actions have already been performed on the network packet. *Id.* Thus, a firewall's signature verifies to the client that “the data packet is from a trusted source,” and the mail gateway's signature verifies that “no viruses were found in the data packet.” *Id.* Not only do these signatures mean the client does not need “to perform a virus scan,” it also

allows the client to trust the data packet—a security action that the client cannot itself perform. *Id.* Thus, without the signatures, the client cannot ensure that data packets are trustworthy.

147. *Second*, Petitioner argues that “removing the tag would prevent compatibility issues in subsequent devices that might not be compatible with the proprietary tags.” Pet. 95. Keohane does not disclose that any downstream devices are incompatible with the “proprietary tags.” Moreover, a POSITA would understand that a downstream device can simply ignore the signatures. However, as explained above, Keohane is clear that the signature is essential for a downstream device to verify that the sender is trusted and that certain security actions were performed. Thus, a POSITA would not be motivated to remove the signature.

XI. GROUND 7 AND 8: KEOHANE IN COMBINATION WITH OTHER REFERENCES DOES NOT RENDER OBVIOUS CLAIMS 10, 13-14, 19-20, 25-26, AND 30-31

148. Grounds 7 and 8 of the petition rely on Keohane for the same limitations discussed in Section VII. Accordingly, these Grounds should be rejected for the same reasons.

XII. CONCLUSION

149. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the

United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

150. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on the information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed: January 11, 2024

A handwritten signature in black ink, appearing to read 'J. Black Jr.', written over a horizontal line.

John Black Jr., Ph.D.