

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION, CAVIUM, LLC,
WISTRON, INC, and DELL INC.
Petitioner,

v.

ALACRITECH, INC.,
Patent Owner.

Case IPR2017-01392¹
Patent 7,337,241 B2

Before STEPHEN C. SIU, DANIEL N. FISHMAN, and
CHARLES J. BOUDREAU, *Administrative Patent Judges*.

FISHMAN, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318(a) and 37 C.F.R. § 42.108

¹ Cavium, Inc., which filed a Petition in Case IPR2017-01728 (later renamed Cavium, LLC (Paper 76)), Wistron, Inc., which filed a Petition in Case IPR2018-00328, and Dell Inc., which filed a Petition in Case IPR2018-00372, have been joined as petitioners in this proceeding.

I. INTRODUCTION

Intel Corporation (“Petitioner” or “Intel”) filed a Corrected Petition (Paper 4, “Petition” or “Pet.”) requesting *inter partes* review of all claims (1–24) of U.S. Patent No. 7,337,241 B2 (“the ’241 patent,” Ex. 1001) pursuant to 35 U.S.C. §§ 311 *et seq.* Alacritech, Inc. (“Patent Owner”) filed a Preliminary Response. Paper 10 (“Prelim. Resp.”).

On November 30, 2017, based on the record before us at that time, we instituted an *inter partes* review of all claims (1–24) and all grounds. Paper 11 (“Decision” or “Dec.”).

Patent Owner filed a Corrected Response (Paper 34, “PO Resp.”) and Petitioner filed a Reply (Paper 45, “Reply”).

Patent Owner also filed a Contingent Motion to Amend (Paper 25, “Motion” or “Mot.”) to which Petitioner filed an Opposition (Paper 40, “Opposition” or “Opp.”). Patent Owner then filed a Reply in support of its Motion (Paper 46, “PO Reply”) and Petitioner filed a Sur-Reply to Patent Owner’s Reply (Paper 54, “Sur-Reply”). Patent Owner’s Motion is contingent on any of the challenged claims being found to be unpatentable. Mot. 1. We address Patent Owner’s Motion to Amend below.

Each party filed a respective Motion to Exclude certain evidence of the other party. Papers 58, 59. Each party filed a respective Opposition to the other party’s Motion to Exclude (Papers 61, 62) and a respective Reply in support of its Motion to Exclude (Papers 64, 65). We address these motions below.

Upon consideration of the complete record, we are persuaded by a preponderance of the evidence that claims 1–24 are unpatentable. Furthermore, we deny Patent Owner’s Motion to Amend, deny Petitioner’s

Motion to Exclude, deny-in-part and dismiss-in-part Patent Owner's Motion to Exclude, and we grant Patent Owner's Motion to Seal.

A. Related Matters

We are informed that the '241 patent is involved in the following litigations: *Alacritech, Inc. v. CenturyLink, Inc.*, Case No. 2:16-cv-00693-JRG-RSP (E.D. Tex.); *Alacritech, Inc. v. Wistron Corp.*, Case No. 2:16-cv-00692-JRG-RSP (E.D. Tex.); and *Alacritech, Inc. v. Dell Inc.*, Case No. 2:16-cv-00695-RWS-RSP (E.D. Tex.). Pet. 3; Paper 6, 1.

B. The '241 Patent

The '241 patent describes a system and method for accelerating data transfer between a network and storage unit. Ex. 1001, Abstract. An embodiment of the '241 uses the Transmission Control Protocol ("TCP") and Internet Protocol ("IP"). *Id.* at 3:8-10. Another transport protocol is User Datagram Protocol ("UDP"). In particular, the claimed invention of the '241 patent relates to fast-path processing in which processing for headers of a layered network protocol (e.g., TCP/IP or UDP/IP) is offloaded from the host computer to an intelligent network interface. *See id.* at 5:18–38, Fig. 24. Specifically, the intelligent network interface card ("INIC") includes accelerated processing features, "[t]he accelerated processing includes employing representative control instructions for a given message that allow data from the message to be processed via a fast-path which accesses message data directly at its source [in the host computer] or delivers it directly to its intended destination [in the host computer]." *Id.* at 5:18–22.

According to the '241 patent, its invention is based on the “FreeBSD” TCP/IP protocol stack. *Id.* at 37:38–49. The FreeBSD protocol stack is a widely and freely distributed package of software source code that implements the TCP/IP (and other) protocols. “The bulk of the protocol stack is based on the FreeBSD TCP/IP protocol stack.” *Id.* at 48:13–15. “The base for the receive processing done by the INIC . . . is the fast-path or ‘header prediction’ code in the FreeBSD release.” *Id.* at 74:56–59. To simplify the use of that software on the INIC, the '241 patent discloses an embodiment that avoids handling of certain complexities in the TCP/IP protocols. *Id.* at 37:50–38:62. According to the '241 patent, the embodiment results in two modes of operation—a slow path in which the INIC operates as a “typical dumb” network interface and a fast path for processing data that does not fall into one of the exception conditions excluded from its implementation. *Id.* at 38:63–39:2. “In the slow path case, network frames are handed to the system at the MAC layer and passed up through the host protocol stack like any other network frame. In the fast path case, network data is given to the host after the headers have been processed and stripped.” *Id.* at 39:2–7.

C. Illustrative Claims

Claims 1, 9, and 17 are the independent claims of the '241 patent. Claims 1 and 9, reproduced below, are illustrative of the claimed subject matter:

1. A method for network communication, the method comprising:

receiving a plurality of packets from the network, each of the packets including a media access control layer header, a network layer header and a transport layer header;

processing the packets by a first mechanism, so that for each packet the network layer header and the transport layer header are validated without an interrupt dividing the processing of the network layer header and the transport layer header;

sorting the packets, dependent upon the processing, into first and second types of packets, so that the packets of the first type each contain data;

sending, by the first mechanism, the data from each packet of the first type to a destination in memory allocated to an application without sending any of the media access control layer headers, network layer headers or transport layer headers to the destination.

Id. at 98:32–49.

9. A method for communicating information over a network, the method comprising:

obtaining data from a source in memory allocated by a first processor;

dividing the data into multiple segments;

prepending a packet header to each of the segments by a second processor, thereby forming a packet corresponding to each segment, each packet header containing a media access control layer header, a network layer header and a transport layer header, wherein the network layer header is Internet Protocol (IP), the transport layer header is Transmission Control Protocol (TCP) and the media access control layer header, the network layer header and the transport layer header are prepended at one time as a sequence of bits during the prepending of each packet header; and

transmitting the packets to the network.

Id. at 99:19–35.

D. Asserted Grounds Of Unpatentability

Petitioner asserts that claims 1–24 are unpatentable based on the following grounds (Pet. 14–15):

Reference(s)	Basis	Claims challenged
Erickson, ² Tanenbaum, ³ and Alteon ⁴	§ 103	1–8, 18, 22, and 23
Erickson and Tanenbaum	§ 103	9–17, 19–21, and 24

Petitioner additionally contends that the Specification does not provide adequate structural disclosure for certain terms recited in claims 1–5, 7, 8, 17, 20, and 23. *See* Pet. 26, 32–34, 53 n.8. *Inter partes* review proceedings are limited to patentability challenges based on prior art patents and printed publications under 35 U.S.C. §§ 102 and 103. *See* 35 U.S.C. § 311(b); 37 C.F.R. § 42.104(b)(2). Accordingly, we are not authorized in *inter partes* review proceedings to address indefiniteness issues, which arise instead under 35 U.S.C. § 112, second paragraph.

Petitioner relies on the testimony of Robert Horst Ph.D. (Exs. 1003,⁵ 1210, 1223) in support of its assertions. Patent Owner relies on the

² U.S. Patent No. 5,768,618 (“Erickson,” Ex. 1005).

³ Andrew S. Tanenbaum, *Computer Networks*, Third Edition, 1996 (“Tanenbaum,” Ex. 1006).

⁴ Alteon Networks Inc., *Gigabit Ethernet Technical Brief: Achieving End-to-End Performance*, 1996 (“Alteon,” Ex. 1033).

⁵ We note the apparent, harmless, typographic error in Exhibit 1003 in which most pages include a header referring to another patent (7,237,036 rather than the ’241 patent) and appear to refer to the Declaration as the “Lin Decl.” Ex. 1003, 6–103.

testimony of Kevin Almeroth Ph.D. (Exs. 2026, 2305) in support of its assertions.

II. DISCUSSION

A. *Legal Principles*

1. *Obviousness*

A patent claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter 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 to which said subject matter pertains.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) objective evidence of nonobviousness, i.e., secondary considerations. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

2. *Level of Ordinary Skill in the Art*

Petitioner argues a person of ordinary skill in the art related to the ’241 patent would have a Bachelor’s Degree in computer engineering, computer science, or electrical engineering, plus at least five years of experience in computer architecture, network design, network protocols, and software and hardware development. Pet. 38 (citing Ex. 1003 ¶ 18–19).

Patent Owner’s definition of the person of ordinary skill is similar to Petitioner’s definition suggesting “several” years of experience rather than a

specific number of years and argues “[a]ny differences between Petitioners’ proposed level of ordinary skill and that proposed by Patent Owner would not have any bearing on the analysis presented in this Response.” PO Resp. 8.

We are persuaded by Petitioner’s definition of the level of ordinary skill in the art and we find this definition is commensurate with the level of ordinary skill in the art as reflected in the prior art. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001) (“[T]he absence of specific findings on the level of skill in the art does not give rise to reversible error where the prior art itself reflects an appropriate level and a need for testimony is not shown.” (internal quotation marks omitted)); *In re GPAC Inc.*, 57 F.3d 1573, 1579 (Fed. Cir. 1995).

Therefore, we define the level of ordinary skill in the art, at the time of the ’241 patent, to include a Bachelor’s degree in computer engineering or computer science and at least five years of experience in network design or network protocols.

3. Claim Construction

In an *inter partes* review filed before November 13, 2018, as is the case here, we construe claim terms in an unexpired patent according to their broadest reasonable construction in light of the specification of the patent in which they appear. 37 C.F.R. § 42.100(b). Consistent with the broadest reasonable construction, claim terms are presumed to have their ordinary and customary meaning as understood by a person of ordinary skill in the art in the context of the entire patent disclosure. *In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007). Only terms that are in controversy need to be construed and only to the extent necessary to resolve the controversy.

See Wellman, Inc. v. Eastman Chem. Co., 642 F.3d 1355, 1361 (Fed. Cir. 2011); *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999).


In our Decision on Institution, we determined that it was not necessary to provide an express interpretation of any claim terms. Dec. 5. In its Response, Patent Owner agrees that no construction is necessary. PO Resp. 22–23. Petitioner’s Reply does not address any claim construction issues. Therefore, we discern no reasons to depart from our determination in our Decision on Institution that no express claim construction is necessary.

B. Cited Prior Art References

1. Overview Of Erickson (Ex. 1005)

Erickson is directed to a “method of controlling an input/output (I/O) device connected to a computer to facilitate fast I/O data transfers.”
Ex. 1005, Abstract.

Figure 3 of Erickson is reproduced below:




Erickson's Figure 3, reproduced above, depicts data flow in accordance with Erickson's invention. As shown in Figure 3, slow application 306 uses normal stream processing 308 and pass-through driver 310 to send information to I/O device adapter 314 to commodity interface 322. *Id.* at 4:53–61. Alternatively, fast applications 302 and 304 send information directly to I/O adapter 314 via “virtual hardware” 316 and 318 avoiding the overhead of the streams processing and pass-through driver. *Id.* at 4:61–5:3.

In particular, Erickson discloses that I/O device adapter 314 and a user process on a host computer share access to a portion of the user's virtual memory space on the host computer. *Id.* at 1:67–2:7. When applied in the context of network communications, I/O adapter 314 (i.e., a “network interface device”) and the user process pre-negotiate certain fields that will be common to all data transfers to be made—i.e., fields in the various headers used in layered network protocols such as TCP/IP or UDP/IP over an Ethernet medium. *Id.* at 6:42–7:4. To transmit a datagram, the user

process programs registers of the I/O adapter to identify the data to be transmitted, the registers and the identified data accessible through the shared virtual memory. *Id.* at 7:5–32. The I/O adapter then combines the identified data with the pre-negotiated fields of the various headers needed for network transmission, adjusts fields in the pre-negotiated header that vary for each datagram, and transmits the completed packet including the completed headers and the user-supplied data. *Id.* at 7:38–8:26.

Erickson's Figure 4 is reproduced below.

FIG. 4



Erickson's Figure 4, reproduced above, depicts processes 402 and 404 transmit and receive information from interconnect 410 through direct application interface ("DAI") 408. *Id.* at 5:6–11. Information from interconnect 410 is routed directly to a process (402 or 404) through virtual hardware and register rather than through a traditional operating system interface 406. *Id.* at 5:11–15.

2. *Overview Of Tanenbaum (Ex. 1006)*

Tanenbaum describes general principles, as well as detailed aspects, of data transmission in computer networks, including TCP/IP and UDP/IP protocols. *See generally* Ex. 1006.

3. *Alteon Reference (Ex. 1033)*

a. Printed Publication

An *inter partes* review may only request review under 35 U.S.C. §§ 102 and 103 and only based on “prior art consisting of patents or printed publications.” 35 U.S.C. § 311(b); 37 C.F.R. § 42.104(b)(2). Before reaching the merits of Petitioner’s obviousness contentions regarding unpatentability, some of which are based, in part, on Alteon, we must determine as a threshold issue whether Alteon is a prior art printed publication under 35 U.S.C. § 311(b). It is Petitioner’s burden to prove that it is, as Petitioner bears the burden of proving unpatentability by a preponderance of the evidence. *See* 35 U.S.C. § 316(e).

Petitioner argues “Alteon was published on or before January 26, 1997 and is therefore at least § 102(a) prior art.” Pet. 40 (citing Ex. 1087 (the “Butler Declaration”).⁶

Patent Owner argues Petitioner has failed to meet its burden to show Alteon is available as a prior art printed publication in this preliminary proceeding. PO Resp. 27–31. Specifically, Patent Owner argues the

⁶ This assertion appears to be in error because Mr. Butler’s Declaration identifies the archival date in www.archive.org as January 13, 1997. *See* Ex. 1087, 1 (¶ 5), 4.

Petition fails to explain how a “crawler,” referenced in Mr. Butler’s Declaration, automatically located a copy of the Alteon technical brief document. *Id.* at 28–29. Patent Owner further argues the fact that a crawler can locate the Alteon technical brief does not explain how an ordinarily skilled artisan could similarly locate the reference by exercising reasonable diligence. *Id.* at 29 (citing Ex. 2026 ¶ 106). Thus, Patent Owner contends the Petition fails to explain how an interested person of ordinary skill would have located the Alteon document and notes with approval some weaknesses the Board identified in the Petition regarding sufficiency of the evidence relating to printed publication status of Alteon. *Id.* at 29–30 (citing Dec. 8–9). Patent Owner further contends Petitioner fails to even allege that Alteon was ever meaningfully indexed such that interested ordinarily skilled artisans would have been able to locate the reference. *Id.* at 30–31. Still further, Patent Owner contends that a “crawler” has far greater faculties to locate web pages than does a human and, thus, asserts that the fact that a crawler at archive.org can locate a web page does not evidence that an ordinarily skilled artisan would have been able to locate the web page by exercising reasonable diligence at the relevant time. *Id.* at 29 (citing Ex. 2026 ¶ 106).

In its Reply, Petitioner initially points out that Patent Owner had cited Alteon during prosecution as evidence that it was publicly accessible to Patent Owner at the time of the ’241 patent. Reply 2 (citing Ex. 1002, 309–10). Furthermore, Petitioner argues Dr. Horst testifies that Alteon was accessible via a number of hyperlinks at www.alteon.com and testifies that search engines at the time of the ’241 patent, similar to archive.org, used

crawlers to locate web pages as potential search results for users. *Id.* at 3 (citing Ex. 1223 ¶¶ 23–28).

Our Decision on Institution also noted that it was unclear whether “Exhibit A” in Mr. Butler’s Declaration (Ex. 1087) was the same document as the Alteon reference filed as Exhibit 1033. Dec. 8–9. Dr. Horst further testifies that he examined the Alteon reference (Ex. 1033) and “Exhibit A” attached to Mr. Butler’s Declaration (Ex. 1087) and determined they were the same. Ex. 1210 ¶ 16.

Patent Owner’s assertions comparing the capability of a human to search for web pages to that of a web crawler computer program are inapposite. Petitioner asserts,

Moreover, like the Internet Archive, search engines also use crawlers to index web pages for searching. If the Internet Archive could find it, so could a search engine, such as Altavista. Ex. 1223, ¶¶ 24–25. A POSA would certainly have relied on search engines in 1997 to locate relevant art. Ex. 1223, ¶ 25.

Reply 3.

We agree. Search engines, such as Altavista, were operable at the time of the ’241 patent and at the time archive.org recorded a copy of the alteon.com web pages—including the URL that stored the technical writeup that is Exhibit 1033. Thus, we are persuaded by Mr. Butler’s Declaration that his attached “Exhibit A” was locatable by a web crawler for archive.org. We are further persuaded that Alteon was publicly accessible at the time of the ’241 patent’s priority date by Petitioner’s observation that Patent Owner cited Alteon during prosecution.

For at least the above reasons, we are persuaded by a preponderance of the evidence that Alteon was publicly accessible, and hence a printed publication, at the time of the '241 patent.

b. Overview Of Alteon

Alteon describes general principles of operation of a high-speed Ethernet I/O adapter. *See* Ex. 1033, 5–7. In particular, Alteon discloses a problem of prior Ethernet I/O adapters that required multiple interrupts in the processing for each packet, thus, consuming resources of the host computer. *See id.* at 20. Alteon purports to address this problem using an intelligent network I/O adapter that “allows a single interrupt to be issued [(from the adapter to a host system)] for multiple data packets.” *Id.* at 22.

C. Alleged Indefiniteness

Claims 1–8 and 17–24 recite a “first mechanism” and/or a “second mechanism.” Petitioner argues “mechanism” is a nonce word that fails to convey sufficient structure and, thus, these terms should be construed under 35 U.S.C. § 112(6) as means plus function elements. *Pet.* 26. Petitioner further argues the '241 patent Specification fails to disclose any corresponding structure for the functions of these elements and, thus, these claims are invalid as indefinite under 35 U.S.C. § 112(2).⁷ *Pet.* 26–33. Indefiniteness, however, is not an issue for an *inter partes* review.

⁷ Petitioner specifically identifies only claims 1–5, 7, 8, 17, 20, and 23 as indefinite under this reasoning. *Pet.* 26. However, claims 1 and 17 are independent claims from which all of claims 2–8 and 18–24 depend (directly or indirectly). The dependent claims incorporate all limitations of the claims from which they depend. Because Petitioner has not presented argument or

In our Decision on Institution, we determined that the apparent novelty of the claimed invention lies in the performance of the steps to exchange information between two structures—i.e., between any two “mechanisms” capable of performing the recited steps. Dec. 10–11. All claims of the ’241 patent are method claims and the method steps are agnostic regarding the particular type of mechanism involved in each step. Instead, the term “mechanism” bears weight in the claim only to the extent that certain recited processing of the method steps involve one or the other recited mechanism or involve both mechanisms. Patent Owner does not dispute that the prior art discloses “mechanisms.” Therefore, particular structural features/limitations of each mechanism are not dispositive with respect to our determination here of unpatentability.

We are not persuaded claims 1–8 and 17–24 should be interpreted in accordance with 35 U.S.C. § 112(6) and we do not reach Petitioner’s assertions regarding indefiniteness under § 112.

D. Obviousness Over Erickson, Tanenbaum, And Alteon

Petitioner contends claims 1–8, 18, 22, and 23 are unpatentable under 35 U.S.C. § 103(a) as obvious over the combination of Erickson, Tanenbaum, and Alteon and contends that claims 9–17, 19–21, and 24 are unpatentable under 35 U.S.C. § 103(a) as obvious over the combination of Erickson and Tanenbaum. See Pet. 49–92. In general, the combination of Erickson and Tanenbaum is relied upon to disclose fast-path processing for

evidence that claims 6, 18, 19, 21, 22, and 24 recite any additional structure to exclude them from its argument of indefiniteness, Petitioner’s argument regarding indefiniteness applies to all of claims 1–8 and 17–24.

transmitting information to a network (i.e., claims 9–17, 19–21, and 24) and Alteon is added to the combination for those claims that include receiving transmissions from a network using fast-path processing (i.e., claims 1–8, 18, 22, and 23) as well as claims that include limitations relating to interrupts (1–8 and 18).

1. Independent Claim 1

Independent method claim 1 is directed to receiving information from a network using fast-path processing. Petitioner argues Erickson discloses the step of “receiving a plurality of packets . . .” as receiving Ethernet packets having a physical (e.g., media access control or “MAC”) layer header, a network (e.g., “IP”) layer header, and a transport (e.g., “UDP” or “TCP”) layer header. Pet. 50–51 (citing Ex. 1005, Fig. 6, 6:48–56; Ex. 1003 Appendix A-2). Petitioner notes that, although Erickson’s Figure 6 is specific to the UDP transport protocol, Erickson expressly discloses its invention is equally applicable to other protocols including the TCP transport protocol. *Id.* at 52 (Ex. 1005, 5:47–51).

Petitioner argues the step of “processing the packets . . . [so] the network layer header and the transport layer header are validated without an interrupt dividing the processing of [the headers]” is disclosed by Erickson’s I/O adapter executing scripts that validate the network and transport layer headers by computing checksums for each such header. *Id.* at 53–54 (citing Ex. 1005, 4:18–23, 7:50–64, 8:10–25; Ex. 1003 Appendix A-5). Petitioner acknowledges Erickson does not disclose the “without an interrupt” limitation of claim 1 but asserts Alteon in the proposed combination discloses this limitation by teaching a single interrupt of a host system is

generated by an intelligent network interface for the processing of multiple packets. *Id.* at 54–55 (citing Ex. 1033, 15, 22, 23).

Petitioner argues Erickson and Alteon are both concerned with reducing intervention processing by a host computer for each I/O operation and, therefore, contends the ordinarily skilled artisan would have been motivated to combine Alteon with Erickson because Alteon’s “single interrupt processing reduces the need for the host computer to insert itself into the process.” *Id.* at 48.

The Petition next argues Erickson in combination with Tanenbaum discloses the step of “sorting the packets . . . into first and second types of packets, so that the packets of the first type each contain data.” *Id.* at 55–56 (citing Ex. 1006, 584–85). Specifically, the Petition asserts Tanenbaum discloses checking if packets meet certain criteria for fast-path processing and, thus, sorts packets according to claim 1. *Id.*

Petitioner argues Erickson incorporates an earlier version of Tanenbaum (1981 which Petitioner refers to as “Tanenbaum81”) and, thus, provides express motivation to combine Erickson with the teachings of the then-current version of Tanenbaum (1996 which Petitioner refers to as “Tanenbaum96”). *Id.* at 44 (citing Ex. 1005, 4:34–43). Petitioner then argues the ordinarily skilled artisan would have been motivated to seek out the then-current version of Tanenbaum (e.g., Tanenbaum as published in 1996 (Ex. 1006)) at the time of the ’241 patent priority. *Id.* Petitioner contends Tanenbaum and Erickson both relate to fast-path processing of packets and, although Erickson discloses its applicability to the UDP (transport) protocol (Ex. 1005, 5:47–51), Tanenbaum expressly discloses fast-path processing for TCP protocol. *Id.* at 45–47.

Petitioner contends the proposed combination also discloses the step of “sending . . . the data from each packet of the first type to a destination in memory . . . without sending any of the . . . headers to the destination.” *Id.* at 56–58. Specifically, Petitioner argues Erickson discloses sending data to an application by writing the data directly to a memory space of the application. *Id.* at 56–57 (citing Ex. 1005, Figs. 3, 4, 4:53–5:14, 6:1–41, 8:17–37; Ex. 1003 § V.H.1).

Petitioner further argues step 2 of Alteon’s Table 4 expressly discloses that the packet is moved to the application memory without the headers. *Id.* at 57–58 (citing Ex. 1033, 21; Ex. 1003 Appendix A13–A14). Alteon’s Table 4 is reproduced below as annotated by Petitioner.

Table 4. Steps in Data Reception using Second Generation NIC	
Data Steps	Control Steps
1. The NIC moves the first 64 bytes of the packet to the protocol stack through a pre-allocated buffer. The first 64 bytes includes the header information and some data.	A. The NIC notifies the stack that it has moved 64 bytes of data by issuing an interrupt.
	B. The protocol stack analyzes the headers and tells the NIC where in application memory to put the remaining data from the packet.
2. The protocol stack moves any data from the initial 64 bytes (minus the header) to the application memory.	
3. The NIC moves the rest of the data into application memory. If application memory is not accessible, then an intermediate buffer is used and the data is copied to the application memory by the host in Step 4.	C. The NIC tells the stack that it has finished moving the rest of the data packet into application memory by issuing an interrupt.
4. The protocol stack performs a checksum on the packet in the application memory space.	D. The protocol stack informs the application that data has arrived.

Table 4 of Alteon, as highlighted by Petitioner, shows steps in data reception using a second generation network interface card (“NIC”). The Petition highlights step 3 in Alteon’s Table 4, which discloses “[t]he NIC moves the rest of the data into application memory.” Both steps 2 and 3 of Alteon’s Table 4 disclose moving data without headers—the headers being moved in the first 64 byte transfer to the protocol stack. Pet. 57.

a. Motivation To Combine The References

Patent Owner argues there is no motivation to combine Erickson with either Tanenbaum or Alteon. PO Resp. 42–46. Specifically, regarding adding Alteon to the combination, Patent Owner argues Alteon’s disclosure of a single interrupt involves only the receive side of exchanges but complete systems involve both receive and transmit operations. *Id.* at 43 (citing Ex. 2026 ¶ 126). Furthermore, Patent Owner argues the alleged motivation of increased performance “is generic to the point of being inconsequential” and, thus, it is only hindsight that would lead to the proposed combination with Alteon. *Id.* (quoting Ex. 2026 ¶ 127).

Petitioner argues Alteon and Erickson both disclose bidirectional communication devices. Reply 9 (citing Ex. 1005, 1:63–67, 4:53–5:14; Ex. 1033, 22). Furthermore, Petitioner contends its motivation is based on more than a generic increase in performance but, rather, is based on the express teaching of eliminating interrupts during receive processing to improve performance. *Id.* The Petition argues that Erickson partially addresses the problem of avoiding operating system (host CPU) intervention of each I/O operation. Pet. 47 (citing Ex. 1005, 3:9–10; Ex. 1003 ¶ 169). Petitioner further argues Alteon addresses a similar problem and, thus, the ordinarily skilled artisan would have recognized that Alteon’s disclosure of a single interrupt further reduces the need for host CPU intervention for an I/O request. Pet. 47–48 (citing Ex. 1033, 22; Ex. 1003 ¶ 170). Petitioner further contends Erickson already offloads processing of the network protocols from the host CPU to the network interface and Alteon further enhances that same improvement by reducing interrupts to a single interrupt. Pet. 48 (citing Ex. 1003 ¶ 171).

We are persuaded that the ordinarily skilled artisan would have recognized the benefit of further enhancing Erickson's goal of offloading processing from the host CPU by combining it with Alteon's disclosure of reducing interrupts required of the host CPU for processing of an I/O request. Even assuming Patent Owner is correct that a generic motivation for increased performance is insufficient motivation, the motivation articulated in the Petition is more than a mere generic increase but, instead, points to Alteon's specific improvement of reduced interrupts of the host CPU to further enhance Erickson's goal of offloading processing from a host CPU.

Regarding the combination of Erickson with Tanenbaum, Patent Owner argues Tanenbaum teaches away from performing TCP/IP protocol processing by anything other than the host CPU. PO Resp. 43–44 (citing Ex. 2026 ¶ 128). Patent Owner further argues Tanenbaum teaches “[i]t is a myth to think that the main CPU has other work to do while waiting.” *Id.* at 44 (citing Ex. 1006, 588–89).

Petitioner argues that Tanenbaum expressly discloses that transport layer protocol processing may be offloaded to a network interface card and merely expresses its preference for maintaining processing of the network protocol on the host CPU. Reply 10–11 (citing Ex. 1006, 498, 530; Ex. 1003 ¶¶ 144, 146, 163–65). Petitioner further argues Tanenbaum merely expresses a preference for having the main (host) CPU perform the protocol processing but does not suggest that offloading protocol processing to a network interface will not work if the second CPU (network interface CPU) is fast enough. Reply 11. Petitioner contends expressing such a preference

is not criticism, discrediting, or discouragement sufficient to teach away from the proposed combination. *Id.*

We are persuaded by Petitioner’s arguments. A teaching away must clearly discredit or discourage the proposed combination. *See In re Gurley*, 27 F.3d 551, 553 (Fed. Cir. 1994); *see also In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004) (“The prior art’s mere disclosure of more than one alternative does not constitute a teaching away from any of these alternatives because such disclosure does not criticize, discredit, or otherwise discourage the solution claimed . . .”). Patent Owner does not persuade us that Tanenbaum’s preference for protocol processing in the host CPU clearly discredits Erickson’s approach of offloading protocol header processing to a network interface such that Tanenbaum teaches away from the proposed combination.

Lastly, Patent Owner argues Alteon and Erickson are “fundamentally incompatible” because Erickson discloses that its network adapter has limited knowledge of the user process and the memory mapping of the host and, thus, can only access fixed pages of memory in the user address space for transferring data using the fast-path processing. *Id.* at 44–45 (citing Ex. 2029 ¶ 129; Ex. 1005, 7:65–8:2, 8:16–24). Patent Owner contends Petitioner’s Expert, Dr. Horst, admits Erickson does not disclose any feature for data transfers between the host memory and the network adapter that span page boundaries. *Id.* at 45 (citing Ex. 2028, 78:7–13).

Petitioner argues that the ordinarily skilled artisan would understand that a typical page of memory would hold multiple packets and argues that Erickson expressly acknowledges that its scripts would need to be enhanced if data were to span multiple pages of memory. Reply 11–12 (citing

Ex. 1023 ¶¶ 46–47). Furthermore, Petitioner argues the ordinarily skilled artisan “could have implemented TCP without making changes to the Erickson adapter’s ability to cross page boundaries.” Reply 12 (citing Ex. 1223 ¶ 46). Still further, Petitioner argues that any such enhancements to allow a request to span page boundaries in the host system’s memory were “well within the capabilities of a [person of ordinary skill in the art].” *Id.* (citing Ex. 1223 ¶¶ 46–49).

We are persuaded by Petitioner’s arguments supported by Dr. Horst’s declaration. Initially, we observe that the claims do not require that a request be allowed to span multiple pages in memory. Even assuming such a requirement is implied in some aspect of the claims, we are persuaded by Dr. Horst’s testimony that the ordinarily skilled artisan would have recognized how to implement enhancements to allow a request to span page boundaries. Furthermore, although Erickson discloses that the adapter would have a limited range of addresses accessible in the host’s memory, even “as small as a single page,” Erickson does not express one page as an upper limit on the size of shared memory between a process and the adapter. Ex. 1005, 8:16–20. Multiple page of shared memory could be mapped for each process accessing the adapter. *See, e.g., id.* at 3:48–60. Thus, in Erickson, the maximum number of segments/packets that may be generated in the adapter without spanning a page boundary or the maximum size of a segment or packet without spanning a page boundary could be varied for a particular application.

b. “Without An Interrupt Dividing The Processing”

Patent Owner argues the combined references fail to disclose claim 1’s recitation of processing network and transport layer headers without an interrupt dividing the processing. PO Resp. 32–37. Specifically, Patent Owner argues Alteon merely discloses use of an interrupt timer that allows multiple packets to be sent to the host system buffers before a timer expires and discloses altering the number of interrupt allowed per second. *Id.* at 33 (citing Pet. 54–55; Ex. 1033, 23). Patent Owner further argues Alteon’s reference to sending packets to the host “merely refers to ‘copying packets into the operating system buffer space’” but does not disclose validating the headers as claimed. *Id.* (citing Ex. 2026 ¶ 112).

Petitioner argues Patent Owner is improperly attacking the references individually rather than for what the combination would have taught or suggested to the ordinarily skilled artisan. Reply 4 (citing Dec. 14). Petitioner argues Alteon expressly discloses its technology “not only reduces the number of times data is copied . . . it allows a single interrupt to be issued for multiple data packets.” *Id.* at 5 (quoting Ex. 1033, 22). Petitioner contends Alteon expressly discloses validating TCP/IP headers by off-loading TCP/IP checksum computations from the host to the network adapter. *Id.* at 5–6 (citing Ex. 1033, 22). Petitioner further argues Patent Owner’s priority document for the ’241 patent admits that it was known at the time of the patent to require less than one interrupt per packet. *Id.* at 6–7 (citing Ex. 1031, 6; Ex. 1223 ¶¶ 53–61).

We are persuaded by Petitioner’s arguments. As noted in our Decision on Institution, we find Patent Owner’s arguments improperly attack the references individually rather than considering what the

combination would have taught or suggested to the ordinarily skilled artisan. Dec. 14–15. The Petition relies on Erickson and Tanenbaum, not Alteon, in the proposed combination for disclosing checking (i.e., “validating”) headers of packets for potential fast-path processing. See Pet. 50–58. Alteon is relied on in the Petition primarily for disclosing reducing the number of interrupts to a *single interrupt* in the processing of *multiple packets*. Pet. 54. In the proposed combination, Petitioner contends “Alteon discloses the absence of interrupts between validation of the network and transport layer headers by disclosing a single interrupt for multiple data packets,” further contending, “[i]f there was an interrupt after validation of every network layer header, it would not be possible to have only one interrupt for multiple data packets.” *Id.* (citing Ex. 1003 Appendix A-5). We agree. If there were a single interrupt generated for the processing of *multiple* packets, it would have been suggested to the ordinarily skilled artisan that there is at least some packet among the *multiple* packets for which there is no interrupt between the processing of the various headers of that packet.

For the above reasons, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan processing the headers without interrupt as recited in claim 1.

c. Sending Data Without Sending Any Header

Patent Owner argues the Petition relies only on a single screenshot from Alteon to disclose the sending limitation of claim 1 but contends Alteon fails to disclose this feature because Alteon discloses moving a header with the data. PO Resp. 37–39 (citing Pet. 57–58; Ex. 2026 ¶¶ 118–19; Ex. 1033, 21 (Table 4); Ex. 1003 Appendix A-13, A-14). Specifically,

Patent Owner argues the first two steps of Alteon disclosed in its Table 4 expressly teach that the network interface moves a packet *including its header* to the host system memory. *Id.* at 38. Reproduced below is Table 4 from Alteon as annotated by Patent Owner.

Table 4. Steps in Data Reception using Second Generation NIC	
Data Steps	Control Steps
1. The NIC moves the first 64 bytes of the packet to the protocol stack through a pre-allocated buffer. The first 64 bytes includes the header information and some data.	A. The NIC notifies the stack that it has moved 64 bytes of data by issuing an interrupt. B. The protocol stack analyzes the headers and tells the NIC where in application memory to put the remaining data from the packet.
2. The protocol stack moves any data from the initial 64 bytes (minus the header) to the application memory.	
3. The NIC moves the rest of the data into application memory. If application memory is not accessible, then an intermediate buffer is used and the data is copied to the application memory by the host in Step 4.	C. The NIC tells the stack that it has finished moving the rest of the data packet into application memory by issuing an interrupt.
4. The protocol stack performs a checksum on the packet in the application memory space.	D. The protocol stack informs the application that data has arrived.

Table 4 of Alteon is highlighted by Patent Owner showing step 1 that moves the first 64 bytes, including a header, from the network interface to a host system buffer and step 2 that moves remaining data of the first 64 bytes minus the header to the host system buffer. PO Resp. 39 (excerpted from Ex. 1033, 21).

Petitioner argues again that Patent Owner's arguments attack the references individually rather than addressing the teachings of the proposed combination. Reply 7 (citing Dec. 14). Instead, Petitioner argues the Petition identifies features of Erickson and Tanenbaum (fast-path header prediction) in the proposed combination, as disclosing the fast-path transfer of data to a host process without headers and identifies Alteon as also teaching transfer of data without headers to a memory of a host system (steps 2 and 3 of Table 4). *Id.* Therefore, Petitioner contends the combination would have disclosed the fast-path transfer of data without a

header (as in Erickson and Tanenbaum) into a memory buffer of a host (as in Alteon). *Id.* (citing Pet. 56–58; Ex. 1223 ¶ 29).

We are persuaded by Petitioner’s arguments. Again, Patent Owner attacks the teachings of Alteon independent of the Petition’s assertions regarding the combined teachings. The Petition specifically relies on Erickson and Tanenbaum for transfer of data without headers from a network interface to a host system and relies on Alteon in the combination for disclosing in step 2 of Table 4 (as reproduced above) for transferring such data, without a header, to a memory buffer of the host system. Pet. 56–58. We further note step 3 of Alteon’s that discloses moving “the rest of the data” to the application memory after the headers are moved in the first step. Pet. 57. Erickson teaches that a computer may be both a sender and receiver of data. Pet. 50 n.7 (citing Ex. 1005, 3:23–36). Erickson expressly discloses an example of a computer as a sender of data to another device wherein only the data is provided to the interface card. Ex. 1005, 7:5–39. The network interface then appends appropriate headers to transmit the data in accordance with, for example, the UDP/IP network protocol. *Id.* at 7:65–8:9. Erickson also discloses that numerous other protocols may be chosen in an implementation including TCP/IP. *Id.* at 5:47–51. Erickson expressly discloses that the headers appended include a UDP protocol header, an IP protocol header, and a media access control (e.g., Ethernet) protocol header. *Id.* at 6:47–56. Thus, Erickson expressly discloses transferring data without headers from the host to the network adapter, which, in turn, appends the appropriate protocol headers and send the data with headers to a receiving device over the network. *See id.* at 7:44–50. Although the exemplary UDP implementation of Erickson focuses on the transmission of data from the

host process to the network interface without headers, Erickson expressly discloses that its invention is applicable to both sending and receiving of data. *See id.* at 4:66–5:15. Therefore, we are persuaded the ordinarily skilled artisan would have understood Erickson’s approach to be applicable to a network interface receiving packets including the recited protocol headers and transferring the data portions of the received packets to a host process devoid of the protocol headers—i.e., the mirror image of the interface receiving data from a host to packetize and send to another device. Furthermore, Petitioner argues in response to Patent Owner’s Response:

As stated above, in the combination of Erickson in view of Tanenbaum96 and Alteon, Tanenbaum96’s “fast path” header prediction TCP/IP processing is performed by Erickson’s interface. Alteon also teaches the transfer of data without headers from the protocol stack (which, in the combination, is on the network adapter for the “fast path”) to the application (which is on the host). Pet. at 58; Ex. 1223, ¶ 29. Because this transfer is without headers, the combination discloses the transfer of data from each packet to a destination in host memory without the headers. Pet. at 56-58; Ex. 1223, ¶ 29.

Reply 7. Thus, even if Erickson does not expressly teach a step of sending data without headers *from the network interface to the host process*, the combination including Alteon’s express teaching of such a step provides an express disclosure of such a step.

For the above reasons, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan the step of sending data to a host system without headers as recited in claim 1.

d. Conclusion Regarding Obviousness Of Claim 1

We have reviewed the record regarding other recitations of claim 1. Other than the issues discussed above, Patent Owner does not challenge the Petition's arguments regarding claim 1. For the above reasons, we are persuaded the proposed combination taught or suggested all limitations of claim 1.

2. Dependent Claim 2

Claim 2 depends from claim 1 and further recites that the media access control layer header and network layer header of packets are processed without an interrupt dividing the processing. The Petition identifies Alteon in the proposed combination as disclosing the recited interrupt limitation. Pet. 58 (citing Ex. 1033, 22–23). The Petition refers to the same reasoning as presented in its analysis of the similar limitation in claim 1. *Id.* (citing Pet. 53–55).

As previously discussed regarding claim 1, Patent Owner argues “Alteon merely suggests copying the packet data from a buffer to the host without interruption, not performing the complete datagram processing without interruption.” PO Resp. 40.

As above regarding claim 1, Petitioner argues Patent Owner's arguments improperly attack Alteon independent of the proposed combination and further argues “for the same reasons as above that the combination teaches that the TCP and IP headers are validated without an interrupt to the host, the MAC and IP layer headers are validated without an interrupt to the host.” Reply 8 (citing Pet. 58).

For essentially the same reasons as discussed above regarding claim 1, we are persuaded the proposed combination taught or suggested this feature of claim 2. As above, we are persuaded that Erickson and Tanenbaum, not Alteon, in the proposed combination taught or suggested processing of headers. Alteon is relied on primarily for disclosing a single interrupt for processing of multiple packets. In particular, as above, we are persuaded that if there is a single interrupt generated for the processing of *multiple* packets, as expressly disclosed in Alteon, the ordinarily skilled artisan would have recognized that there is at least some packet among the *multiple* packets for which there is no interrupt between the processing of any of the various headers of that packet.

For at least the above reasons, we are persuaded the proposed combination taught or suggested the limitations of dependent claim 2.

3. *Dependent Claim 3*

Claim 1 refers to processing of headers by a “first mechanism.” As discussed *supra*, the Petition identifies the network interface of Erickson (as enhanced in combination with Tanenbaum and Alteon) as the recited “first mechanism.” Claim 3 depends from claim 1 and further recites: “processing an upper layer header of at least one of the packets by a second mechanism, thereby determining the destination, wherein the upper layer header corresponds to a protocol layer above the transport layer.” Though the claim is not so limited, we understand the referenced upper layer to refer to layers of layered networking models such as TCP/IP or OSI.

The Petition argues Tanenbaum, in the proposed combination, discloses this feature. Pet. 59–60. In particular, the Petition argues “when a

packet is slow-pathed, it will be sent to the kernel (a second mechanism) for processing – including application layer headers, which are above the TCP header.” Pet. 59. The Petition further argues that Tanenbaum discloses a number of application layers above the transport layer including, for example, the electronic mail (SMTP) application layer of the TCP/IP layered model. *Id.* Petitioner contends the application header (e.g., an email protocol header), processed by the kernel (a second mechanism), includes an email “To:” field specifying an intended recipient—i.e., a destination—and thus meets the claim limitation. Pet. 60.

Patent Owner argues Petitioner’s reliance on Tanenbaum for processing of upper layer headers (e.g., email application headers) by slow-path processing in the host system’s kernel is contrary to the requirements of claim 3. PO Resp. 41. Specifically, Patent Owner argues:

Claim 3, in contrast, requires the lower layer headers (MAC, network, and transport layer headers) to be processed by the claimed first mechanism (i.e., the network interface device), while an upper layer above the transport layer is processed by the claimed second mechanism (i.e., the host CPU operating the protocol stack). Tanenbaum, alone or in combination with the Erickson and Alteon, does not show this functionality.

Id.

In response to Patent Owner’s Response, Petitioner argues that, in the proposed combination, the combined teachings of Erickson and Tanenbaum disclose processing lower layer headers in the network interface (the claimed first mechanism). Reply 8. Petitioner further argues, “[o]nce the network and transport layer headers are removed and the data is transferred to the host, the host processes the upper layer headers.” *Id.* (citing Pet 59).

We are persuaded by Petitioner’s argument that, in the proposed combination, the transport and lower layer protocol headers are processed by the network adapter and relies on the host for subsequently processing upper layer headers. Reply 9. Petitioner’s proposed combination does not require an all-or-nothing approach—i.e., either all headers are processed in a host by traditional protocol stack operating therein or all headers are processed by an enhanced network adapter.

For the above reasons, we are persuaded the proposed combination taught or suggested the limitations of dependent claim 3.

4. *Dependent Claim 6*

Claim 6 depends from claim 1 and further recites “wherein sorting the packets includes classifying each of the packets of the first type as having an Internet Protocol (IP) header and a Transmission Control Protocol (TCP).” The Petition refers to its earlier arguments regarding the sorting step of claim 1 (*see* Pet. 55–56) and contends that Tanenbaum, in the proposed combination, “discloses that once the protocol is determined, sorting can be completed – i.e., the data is assigned to the fast or slow path.” Pet. 64.

Patent Owner argues the proposed combination only enables a UDP embodiment and does not enable a TCP embodiment. PO Resp. 42.

We are persuaded by Petitioner’s arguments. Initially we observe that “a non-enabling reference may qualify as prior art for the purpose of determining obviousness under § 103.” *Symbol Techs., Inc. v. Opticon, Inc.*, 935 F.2d 1569, 1578 (Fed. Cir. 1991). Furthermore, although Erickson expressly discloses a UDP/IP embodiment, it also expressly discloses that it may be used with other protocols including TCP/IP even if not enabling

therefor. Ex. 1005, 5:47–51. Still further, Tanenbaum, in the proposed combination, expressly discloses that the TCP/IP protocol stack may be implemented on the network interface adapter. Ex. 1006, 530 (the TCP/IP protocol “may also be contained on a coprocessor chip or network board plugged into the host’s backplane”). Thus, even if not fully enabled, the combined references clearly teach or suggest to the ordinarily skilled artisan that Erickson’s adapter could be modified to implement a TCP/IP protocol as well as the exemplary UDP/IP protocol. Still further, Dr. Horst testified in his deposition that modifications to Erickson to implement TCP/IP, rather than the exemplary UDP/IP, would have been obvious to the ordinarily skilled artisan. *See, e.g.*, Ex. 2028, 99:18–101:4. We credit Dr. Horst’s testimony mindful of the fact that the source code for an entire TCP/IP stack was freely accessible to the ordinarily skilled artisan (FreeBSD) and was well documented (*see, e.g.*, Exs. 1006, 1013).

For the above reasons, we are persuaded the proposed combination taught or suggested the limitations of dependent claim 6.

5. Dependent Claims 4, 5, 7, And 8

Dependent claims 4, 5, 7, and 8 each depend from claim 1 and recite further limitations. We have reviewed Petitioner’s arguments regarding each of these dependent claims. Pet. 61–71. Patent Owner does not dispute Petitioner’s assertions regarding claims 4, 5, 7, and 8.

We are persuaded the proposed combination taught or suggested the limitations of dependent claims 4, 5, 7, and 8.

6. Dependent Claims 18, 22, And 23

As noted above, independent claim 1 is generally directed to method steps for receiving packets in a network interface from a network, processing the network layer headers on the network interface, and transmitting the data portions of the packets (stripped of the headers) to a host system. By contrast, independent claims 9 and 17 are generally directed to obtaining data at the network interface from a host, generating network packets from the data, and transmitting generated packets to a network. Dependent claims 18 and 22 each depend from independent claim 17 and add limitations similar to claim 1 relating to receiving packets from a network, processing the headers at the network interface, and transmitting the data portions of the packets to a host system. In other words, although independent claim 17 is directed to transmissions to the network, dependent claims 18 and 22 add to that base claim limitations relating to transmission from the network. Claim 23 depends from claim 22 and adds limitations (similar to those of claim 3) that an upper layer header (higher than the transport layer) is processed by the host.

The Petition identifies features of these dependent claims in the combination of Erickson, Tanenbaum, and Alteon referring to similar argument in the Petition for claims 1 and 3. Pet. 89–91. As discussed below, claim 17 is alleged to be unpatentable over Erickson and Tanenbaum (without Alteon in the combination). Thus, Petitioner adds Alteon to the combination to disclose “without interrupt” features of claims 18 and 22 similar to those features recited in claim 1.

Patent Owner argues these claims “suffer[] from the same deficiencies” as claim 17 from which they depend. PO Resp. 56.

We have reviewed Petitioner’s arguments regarding claims 18, 22, and 23 and find them persuasive assuming base claim 17 is also found to be unpatentable. In other words, we are persuaded the Petition has identified all features of claims 18, 22, and 23 if claim 17, discussed below, is shown to be unpatentable. As discussed below, we find claim 17 unpatentable as obvious over the combination of Erickson and Tanenbaum. Thus, for at least the same reasons as claims 17, 1, and 3, we are persuaded the combination of Erickson, Tanenbaum, and Alteon taught or suggested all features of dependent claims 18, 22, and 23.

E. Obviousness Over Erickson And Tanenbaum

Independent method claims 9 and 17 are directed to transmitting information to the network using fast-path processing—the information received from a host or otherwise provided to a network interface and transmitted from the network interface to a network.

Claim 9 generally recites obtaining data from a first processor (e.g., from a host computer), dividing the received data into segments, prepending a header to each segment to form a packet, and transmitting the packets to the network. Claim 9 further recites that the prepended header includes a media access control layer header, a network layer (Internet Protocol “IP”) header, and a transport layer (Transmission Control Protocol “TCP”) header. Still further, claim 9 requires that the three header portions “are prepended at one time as a sequence of bits.”

1. Independent Claim 9

Regarding claim 9, the Petition identifies Erickson as teaching obtaining data from a first processor. Pet. 72–73 (citing Ex. 1005, 7:44–45, 8:30–35; Ex. 1003 Appendix A-36). The Petition cites Erickson for disclosing that its techniques may be used for TCP/IP and relies on Tanenbaum, in combination with Erickson, for teaching dividing the data into TCP segments. Pet. 73 (citing Ex. 1005, 5:41–51; Ex. 1006, 540, 543). The Petition relies on Erickson for teaching a header template that is prepended to each segment when execution of a script is triggered on the I/O adapter (*id.* (citing Ex. 1005, 7:42–43; Ex. 1003 Appendix A-42)) and further identifies Tanenbaum as teaching that the transport entity may be operable “on a coprocessor chip or network board” (Pet. 74 (citing Ex. 1006, 530, 588)). Petitioner identifies the same disclosures as discussed above regarding the header having three portions (media access control, network, and transport layer headers) and transmitting the constructed packet to the network. Pet. 74–75.

Regarding the recitation of claim 9 that requires the header be prepended “at one time as a sequence of bits,” the Petition refers to its argument for a similar limitation in claim 7. Pet. 74. The related recitation in claim 7 (which depends from claim 1) recites that the headers “are prepended at one time as a packet header,” i.e., no limitation that the headers are prepended “as a sequence of bits.” In the argument regarding this similar recitation of claim 7, Petitioner asserts Erickson discloses: a header template is stored in the network interface device’s memory, the header template includes both a network (IP) header and a transport (UDP) header; the data to be transmitted to the network is first transferred to the network

interface device's memory; the header template is used to create a header that corresponds to the data to be sent; and the datagram (header combined with data) is then transmitted to the network. Pet. 67–68.

Petitioner further argues there are two well-known, obvious ways the header and data can be combined—i.e., either the completed header is prepended to the data to be sent or the data to be sent is appended to the completed header. Petitioner then contends it would have been obvious to the ordinarily skilled artisan to select one of these two options—namely prepend the completed header to the data to be transmitted. Pet. 69 (citing Ex. 1005, 7:38–47, 8:2–9; Ex. 1003 Appendix A-28, § V.B.8).

a. Headers Are “Prepended At One Time As A Sequence Of Bits”

Patent Owner argues the Petition fails to show that the headers are prepended at one time as a sequence of bits and disputes Petitioner's contention that there are only two way to combine the data with the headers. PO Resp. 47–50. Specifically, Patent Owner contends Erickson is silent regarding how (or if) the headers are combined with the data, asserting “Erickson only explains that ‘[t]he I/O device adapter stores the user data provided by the user process in the I/O adapter's memory, and then transmits the completed UDP datagram.’” *Id.* at 47 (quoting Ex. 1005, 7:44–47). Based on Dr. Almeroth's Declaration, Patent Owner posits that the sequential nature of the steps in the exemplary UDP script (Ex. 1005, 7:50–64) would have suggested that Erickson “likely first prepends the Ethernet header 704 to the user data, then prepends the IP header 706 to the Ethernet header, and finally prepends the UDP header 708 to the IP header.” *Id.* at 48 (citing Ex. 2026 ¶ 134).

We are persuaded by Petitioner’s arguments. Initially, we note the term “sequence of bits” is not used, let alone defined, in the Specification of the ’241 patent. Relying on Dr. Almeroth’s testimony, Patent Owner argues, “Rather, prepending these layer headers ‘at one time as a sequence of bits’ further reduces the number of interrupts required to prepare an outbound packet because the packet is not built layer-by-layer as in traditional packet formation, but rather is prepared with fewer data copy operations.” *Id.* at 49 (quoting Ex. 2026 ¶ 135). We understand all data stored in a computer, at the time of the ’241 patent, was stored, retrieved, and manipulated as a sequence of bits. A byte is typically a sequence of eight bits from least significant to most significant. A “word” of data storage is typically understood to be a sequence of multiple bytes and, thus, also a sequence of bits.

Although we agree with Patent Owner that Erickson is silent regarding precisely how the headers and data of a packet are combined for transmission as a completed datagram (packet), we find Dr. Horst’s testimony more credible than Dr. Almeroth’s testimony in this regard. Dr. Horst points out that Dr. Almeroth’s proposed sequential prepending of an Ethernet (MAC) header, then an IP header, then a UDP header would result in the headers being out of order for the requisite transmission of the packet. Ex. 1223 ¶ 30. We agree.

Furthermore, we credit Dr. Horst’s testimony that “[e]ven if one were to assume that Erickson disclosed prepending of headers serially (UDP then IP then Ethernet), it would have been obvious to modify Erickson to prepend the headers at one time given the template structure used.” *Id.* ¶ 31. In other words, even if Erickson’s exemplary script is filling in portions of the

headers in a serial manner, once the header is completed, it would have been obvious to the ordinarily skilled artisan to modify the script to prepend the completed header at one time because the header template is a block of memory that would be moved/copied as a block (i.e., at one time as a sequence of bits).

Moreover, Erickson's exemplary UDP script (Ex. 1005, 7:50–64) discloses steps to construct all three headers based on a header template, but the script makes no reference to moving/copying the data to be associated with the completed header. Thus, Erickson suggests a completed header is constructed by the exemplary script and other processing (not disclosed in detail by Erickson) combines (prepends/appends) the completed header with data to create a completed datagram for transmission to the network. Erickson depicts an exemplary header template for a UDP datagram to be constructed (Ex. 1005, 7:38–49, Fig. 7) and shows an exemplary completed UDP datagram with completed headers and exemplary data (*id.* at 6:42–56, Fig. 6). Lastly, Erickson discloses that the data to be incorporated in the UDP datagram would be copied from the host's memory to a buffer in the network adapter's memory to use it in checksum calculations within the adapter and for subsequent transmission to the network from the network adapter. *Id.* at 8:27–35. These disclosures suggest that, in Erickson, copying of the data, construction of the completed header by the script, and transmission of a completed datagram are separate operations such that the data and a completed header would be combined to generate the completed datagram for transmission.

Patent Owner correctly notes that the Petition is essentially silent regarding the limitation “as a sequence of bits.” PO Resp. 49 n.20

(“Petitioner does not even address the ‘as a sequence of bits’ language recited in claim 9.”). Instead, the Petition’s analysis of this feature of claim 9 refers to earlier analysis of a similar feature in claim 7. Pet. 74 (referring to the analysis at Pet. 67–71 (section 10.7.3)). However, the recitation of claim 7 does not include the limitation “as a sequence of bits.” The Petition argues the header template, shown in Erickson’s Figure 7, does not include the data but the data to be transmitted is transferred from the host to the memory of the network adapter and a completed UDP datagram is transmitted to the network. Pet. 68 (citing Ex. 1005, 7:45–47, Fig. 7).

Accepting, as we have above, that a completed header, generated by the script from the template, is combined with the data to generate a completed UDP datagram (comprising one or more Ethernet packets), then it follows that the header is combined with the data as a block of memory copied or moved to another block of memory—i.e., as a sequence of bytes, words, etc. and, hence, as a sequence of bits.

For the above reasons, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan prepending the headers at one time as a sequence of bits, as recited in claim 9.

b. The References Disclose Dividing Data Into Segments

Claim 9 recites dividing the data obtained from the host into multiple segments (the headers discussed above are then prepended to each segment). The Petition argues Erickson discloses its systems and methods may be applied to TCP protocols as well as the expressly disclosed UDP examples and further argues Tanenbaum, in combination with Erickson, discloses the

function of TCP that divides data to be transmitted into multiple segments. Pet. 73 (citing Ex. 1005, 4:41–51; Ex. 1006, 540, 543).

Patent Owner argues the combined teachings of Erickson and Tanenbaum fail to disclose this feature because Tanenbaum discloses TCP segmentation processing only on a host system—not on a network adapter as required by claim 9. *Id.* at 53. Patent Owner contends Dr. Horst (Petitioner’s expert) acknowledged that the scripts of Erickson would have to be modified to implement TCP segmentation to generate multiple packetized segments (with headers prepended) and further contends Dr. Almeroth testifies that the necessary modifications to Erickson to implement TCP segmentation on the network adapter would be very complex requiring undue experimentation. *Id.* at 52–53 (citing Ex. 2028, 98:16–100:4; Ex. 2026 ¶ 143). Patent Owner particularly identifies the problem of Erickson’s device having limited access to the host memory—e.g., limited to one page of memory when a large request could span multiple pages. *Id.* at 53 (citing Ex. 2026 ¶ 143).

Petitioner argues segmentation is a requirement of the TCP protocol. Reply 15 (citing Pet. 73); *see also* Ex. 1007, 9 (“The TCP is able to transfer a continuous stream of octets in each direction between its users by packaging some number of octets into segments for transmission through the internet system.”). Petitioner further argues that Patent Owner’s assertions that the modifications to implement Tanenbaum’s fast-path TCP (including segmentation) on Erickson’s network adapter would be too complex are incorrect. Reply 16. Petitioner contends the required changes would be straightforward to adapt Erickson’s script to transfer multiple TCP segments

from a single page of host memory. *Id.* (citing Pet. 73; Ex. 1003 Appendix A-39–A-41).

We are persuaded by Petitioner’s argument. Referring to modifications to Erickson’s UDP script to enable processing of the TCP protocol in his deposition, Dr. Horst testified “the combination of Erickson and Tanenbaum would make a script like this obvious.” Ex. 2028, 99:7–9. Dr. Horst’s Declaration proffered three possible alternatives for the modifications required in Erickson’s UDP script and specifically testified “segmentation code is within the skills of a person of ordinary skill in the art in light of the disclosures by Tanenbaum.” Ex. 1210 Appendix A-46. We credit Dr. Horst’s testimony that the modifications would have been within the skills set of the ordinarily skilled artisan.

We emphasize it is what the combined references would have taught or suggested to the ordinarily skilled artisan. The test for obviousness is what the combined teachings of the references would have *suggested* to one of ordinary skill in the art. *See In re Young*, 927 F.2d 588, 591 (Fed. Cir. 1991); *see also In re Keller*, 642 F.2d 413, 425 (CCPA 1981) (a determination of obviousness does not require the claimed invention to be *expressly suggested* by any one or all of the references). Erickson expressly discloses a network interface that performs UDP/IP protocol processing offloaded from the host system CPU—i.e., generating specific UDP/IP headers for a packet based on a template pre-negotiated with the host. Ex. 1005, 6:57–7:38. Tanenbaum in the proposed combination with Erickson discloses details of the TCP protocol and, although expressing a preference for implementing TCP/IP protocols within the host system’s CPU, expressly discloses that the TCP/IP protocol may be offloaded to a network interface.

Pet. 74 (citing Ex. 1006, 530 (“In general, the transport entity may be . . . contained on a coprocessor chip or network board plugged into the host’s backplane.”)). As noted by Petitioner, segmentation of data is a required aspect of the TCP protocol and, thus, Tanenbaum’s disclosure of implementing TCP on a network board, combined with Erickson’s network adapter implementing either UDP or TCP, would have taught or suggested TCP segmentation to the ordinarily skilled artisan.

Still further, the ’241 patent makes clear repeatedly that its implementation is based on the “FreeBSD” software. *See, e.g.*, Ex. 1001, 2:36–39, 32:45–55, 37:38–49, 48:13–16, 48:34–38, 49:67–50:3, 53:10–12. Dr. Horst testifies that computer source code for the “FreeBSD” implementation of a TCP/IP protocol stack was widely and freely available to ordinarily skilled artisans. Ex. 1003 ¶ 27. TCP segmentation is, therefore, implemented in the FreeBSD free software well-known and freely available to ordinarily skilled artisans and actually was the basis for embodiments of the ’241 patent. Ex. 1001, 74:56–59

Therefore, we credit Dr. Horst’s testimony that the modifications required in the exemplary script of Erickson to implement TCP as taught by Tanenbaum on Erickson’s network interface such, including segmentation of data, would have been readily apparent to the ordinarily skilled artisan.

For the above reasons, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan dividing the data into multiple segments as recited in claim 9.

2. *Independent Claim 17*

Independent claim 17, similar to claim 9, is directed to the transmission of user data to the network. Claim 17 differs from claim 9 primarily in that it does not recite the prepending of the headers is at one time as a sequence of bits but, instead, recites “prepending of each outbound packet header occurs without an interrupt dividing the prepending of the outbound media access control layer header, the outbound (IP) header and the outbound TCP header.”

Petitioner identifies the features of claim 17 in the combined disclosures of Erickson and Tanenbaum. Pet. 84–88. Regarding claim 17’s recitation of prepending headers “without an interrupt,” Petitioner refers to the same arguments as for independent claim 9 and dependent claim 7 (dependent from claim 1) as discussed *supra* and contends,

As discussed in claim limitation [7.1.3], the outbound packet has the MAC, IP, and TCP headers appended at one time – *i.e.*, not appended one after another. *See also* Ex.1003, Horst Decl. V.B.8., A-28, A-29. Accordingly, the headers are appended without an interrupt dividing the prepending of the headers.

Pet. 87.

a. Headers Are Prepended Without Interrupt Between Headers

Patent Owner argues the Petition cites to the argument relating to claim 7 but “provides ***no independent analysis.***” PO Resp. 50. We disagree that there is “no independent analysis.” The above quoted portion from page 87 of the Petition cites the arguments relating to claim 7 but then further explains that because the headers are prepended at one time (as in claim 7), they are prepended without interrupt. Pet. 87.

Patent Owner further argues the above-quoted analysis is flawed because: (1) Erickson does not prepend all headers at one time but rather prepends each of the three headers serially, (2) there is no reason to believe prepending multiple headers at one time necessarily implies doing so without an interrupt between the prepending of headers, and (3) Erickson does not disclose use of the TCP protocol that Patent Owner asserts is more complex than UDP. PO Resp. 51–52.

Petitioner argues, as we found in our Decision on Institution, that because processing of headers in Erickson is performed by processing capability of the network adapter, there would be no reason to interrupt the host in between processing of any two headers. Reply 14 (citing Dec. 19; Ex. 1223 ¶ 32).

We are persuaded by Petitioner’s argument. Erickson identifies problems of the prior art that its structure and methods resolve by reducing overhead processing of a host CPU, offloading processing to an intelligent device adapter (e.g., a network interface adapter). *See, e.g.*, Ex. 1005, 1:35–55. Although Erickson is silent regarding the overhead processing associated with interrupts to the host computer, it would be contrary to the purpose of Erickson to presume a need to interrupt the host processor between processing of one protocol header and another protocol header for the network protocol processing by the network adapter. Therefore, Erickson at least would have suggested to the ordinarily skilled artisan that there is no interrupt of the host system CPU between processing of any of the three headers (UDP, IP, or MAC). *See* Reply 14 (citing Ex. 1223 ¶ 32).

For the above reasons, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan prepending the headers without interrupt as recited in claim 17.

b. The References Disclose Dividing Data Into Segments

Similar to claim 9, claim 17 recites dividing the data from the host into multiple segments. Patent Owner presents the same arguments regarding dividing into segments for both claims 9 and 17. PO Resp. 15–16.

For the same reasons discussed above regarding the similar limitation in claim 9, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan dividing the data into multiple segments as recited in claim 17.

3 Dependent Claims 10–16, 19–21, And 24

Petitioner identifies the elements of claims 10–16 (dependent from claim 9) in the combined disclosures of Erickson and Tanenbaum. Pet. 75–84. The Petition identifies the elements of claims 19–21 and 24 (dependent from claim 17) in the same combination of Erickson and Tanenbaum. Pet. 87–88.

Patent Owner does not address these issues in its Response. See PO Resp.

We have reviewed the arguments presented in the Petition regarding these claims and we find them persuasive and un rebutted. Thus, for at least the same reasons as claims 9 and 17 discussed *supra*, we are persuaded that the combined references would have taught or suggested to the ordinarily skilled artisan the elements of dependent claims 10–16, 19–21, and 24.

F. Secondary Considerations

Patent Owner argues that “strong objective indicia of nonobviousness” weighs against obviousness of the challenged subject matter. PO Resp. 57. In particular, Patent Owner contends that the claimed invention addresses a long-felt yet unresolved need in the art for accelerated network communications, that the claimed inventions were commercially successful, that the claimed invention received praise in the industry, that many others tried and failed to develop the claimed technology, and that experts were skeptical of the claimed invention and taught away from it. *Id.* at 57–63. Petitioner responds that Patent Owner has failed to show any nexus between the alleged objective evidence and the features of the challenged claims, as neither Patent Owner nor its expert ties any of Patent Owner’s products or the alleged “claimed network acceleration technologies” to any limitation of any claim. Reply 19.

We address these arguments below in turn.

1. Long-Felt, Unsatisfied Need

Patent Owner alleges that there was significant demand, beginning at least in the early 1990s and recognized in academic papers and prior art publications, “to enhance the efficiency of network protocol processing and network traffic management” and that “[t]he nexus between the long-felt need and the claimed invention is clear and direct” insofar as the accelerated network processing technologies recited in the challenged claims solved recognized “bottlenecks” in data communications caused, for example, by

the processing of protocols. PO Resp. 57–58 (citing Ex. 2026 ¶¶ 147–148; Exs. 2031–2034).

Petitioner responds that “Patent Owner provides no evidence that the ‘accelerated network processing technologies recited in the challenged claims’ actually relate to the ‘challenged claims,’” and that Patent Owner’s “only support is its expert declaration, which is identical to the Response and likewise has no support.” Reply 20.

We agree with Petitioner that Patent Owner’s arguments are insufficient to establish a nexus between the alleged “long-felt, yet unresolved need” and the challenged claims. Although Patent Owner provides citations to four references that afford evidence of networking bottlenecks (Exs. 2031–2034), we agree with Petitioner that Patent Owner has not persuasively established any connection between resolution of those bottlenecks and the claimed invention. To be accorded substantial weight, there must be a nexus between the claimed invention and the evidence of secondary considerations. *In re GPAC Inc.*, 57 F.3d 1573, 1580 (Fed. Cir. 1995). Nexus is a legally and factually sufficient connection between the objective evidence and the claimed invention, such that the objective evidence should be considered in determining nonobviousness. *Demaco Corp. v. F. von Langsdorff Licensing Ltd.*, 851 F.2d 1387, 1392 (Fed. Cir. 1988). The burden of showing that there is a nexus lies with the Patent Owner. *See In re Paulsen*, 30 F.3d 1475, 1482 (Fed. Cir. 1994). In the absence of an established nexus with the claimed invention, secondary consideration factors are not entitled to much, if any, weight and generally have no bearing on the legal issue of obviousness. *See In re Vamco Mach. & Tool, Inc.*, 752 F.2d 1564, 1577 (Fed. Cir. 1985).

Moreover, here, to the extent that offloading protocol processing could be regarded as solving any long-felt need, we note that Erickson previously disclosed offloading protocol processing solves the need reduce overhead processing at the host. Ex. 1005, 4:66–5:3 (“Thus, the overhead of the normal streams processing 308 and pass-through driver 310 are eliminated with the use of the virtual hardware 316 and 318 of the present invention, and fast applications 302 and 304 are able to send receive information more quickly than slow application 306.”). The “long-felt need” must not have been satisfied by another before the patentee. *Newell Co. v. Kenney Mfg. Co.*, 864 F.2d 757, 768 (Fed. Cir. 1988).

2. *Commercial Success And Licensing*

Patent Owner argues that “[t]he features described in the challenged claims also enjoyed great commercial success for over a decade” and that “the offloading and other network acceleration technology described in the challenged claims became ‘the *de facto* standard’ in network acceleration techniques shortly after its introduction and is still the standard today.” PO Resp. 58–59 (citing Ex. 2026 ¶ 149). Patent Owner further contends that its “patent portfolio covering network acceleration techniques was the subject of several successful commercial licenses to many large network and storage players in the industry” and that “[t]his remarkable commercial success was attributed to Patent Owner’s network acceleration technology.” *Id.* at 59 (citing Ex. 2026 ¶ 150).

Petitioner responds that Patent Owner provides no support for its assertions other than paragraphs in its expert declaration that are identical to the paragraphs in its response and are likewise entirely unsupported. Reply

20–21. Petitioner also argues Patent Owner does not demonstrate sufficiently that the alleged licenses were the result of the claimed invention and, therefore, fails to establish a nexus between the claimed invention and the alleged licenses. *Id.* at 21–23. Rather, Petitioner contends, Patent Owner does not attempt to tie the ’241 patent to these licenses, the ’241 patent is not mentioned in any of the licenses, and the licenses “resulted from a lawsuit . . . asserting that Microsoft’s software and Broadcom’s hardware were infringed by different patents on [TCP Offload Engine (“TOE”)] technology,” as a result of which lawsuits, Microsoft and Broadcom took a license and a handful of other manufacturers also took licenses so they could utilize certain Microsoft software that supported the TOE technology. *Id.* at 22–23 (citing Ex. 1227; Ex. 2038).

We are not persuaded by Patent Owner’s argument for at least the reasons set forth by Petitioner. *Id.* at 20–23. We agree with Petitioner that Patent Owner does not provide sufficient information or evidence to establish that the claimed invention, in fact, experienced “commercial success.” In fact, as Petitioner argues, evidence of record indicates that the claimed invention “never went anywhere” and was ultimately “deprecated.” Reply 21–22 (citing Exs. 1224, 1227, 1228, 1230). Further, Patent Owner fails to show that its licensing program was successful because of the merits of the challenged claims of the ’241 patent, as opposed to, for example, other of the patents in Patent Owner’s licensed portfolio, business decisions to avoid litigation, prior business relationships, or for other economic reasons. Although “there is a presumption of nexus for objective considerations when the patentee shows that the asserted objective evidence is tied to a specific product and that product ‘is the invention disclosed and

claimed in the patent’” (*WBIP, LLC v. Kohler Co.*, 829 F.3d 1317, 1339 (quoting *J.T. Eaton & Co. v. Atl. Paste & Glue Co.*, 106 F.3d 1563, 1571 (Fed. Cir. 1997))), Patent Owner carries the burden of demonstrating that the “thing . . . that is commercially successful is the invention disclosed and claimed in the patent” (*Demaco*, 851 F.2d at 1392). Moreover, “[w]hen the thing that is commercially successful is not coextensive with the patented invention . . . the patentee must show prima facie a legally sufficient relationship between that which is patented and that which is sold.” *Id.* Patent Owner has not made a sufficient showing in this case. Additionally, we note that Patent Owner relies on essentially the same evidence and arguments when asserting secondary considerations for at least eight other patents. See Cases IPR2017-01391 (concerning U.S. Patent No. 7,237,036 B2), IPR2017-01393 (concerning U.S. Patent No. 9,055,104 B2), IPR2017-01405 (concerning U.S. Patent No. 7,124,205 B2), IPR2017-01406 (concerning U.S. Patent No. 7,673,072 B2), IPR2017-01409 and IPR2017-01410 (both concerning U.S. Patent No. 8,131,880 B2), IPR2018-00226 (concerning U.S. Patent No. 7,124,205), IPR2018-00234 (concerning U.S. Patent No. 8,805,948), and IPR2018-00401 (concerning U.S. Patent No. 7,945,699). This casts further doubt on the existence of a legally sufficient relationship between the alleged commercial success and the claimed subject matter in *this* case.

Still further, even assuming that the claimed invention experienced “commercial success” or “successful commercial licens[ing]” as Patent Owner alleges, we find that offloading of protocol processing was previously disclosed by Erickson. See discussion above. Thus, to the extent Patent Owner’s alleged commercial success or licensing resulted from that

feature, that success stems from what was known in the prior art so that there can be no nexus. *Tokai Corp. v. Easton Enters., Inc.*, 632 F.3d 1358, 1369 (Fed. Cir. 2011).

3. Industry Praise

Patent Owner alleges that “[t]he industry universally praised commercial embodiments of the features described in the challenged claims.” PO Resp. 60. In particular, Patent Owner contends that HP found that Patent Owner’s NIC was “able to sustain network bandwidth comparable to that of Native NT for large messages, which is close to wire-speed” and “achiev[e] lower processor utilization than native NT’s TCP/IP protocol stack for transmission of large enough messages” and that the test performance of Patent Owner’s INIC was “definitely better than [HP’s] offload.” *Id.* (citing Ex. 2026 ¶ 152; Ex. 2039 ¶ 4). Patent Owner also cites a technology analyst as opining in 2011 that a particular Alacritech product was “an evolutionary advancement of Alacritech’s long standing leadership in protocol acceleration” and that “Alacritech is setting the stage for a next generation of solutions that will accelerate storage from outside the storage array.” *Id.* at 60–61 (quoting Ex. 2040). Patent Owner further contends that the analyst “call[ed] the patented technology ‘game-changing.’” *Id.* (citing Ex. 2040, 3; Ex. 2026 ¶ 153). Petitioner responds that Patent Owner “has provided no evidence that its products practice the challenged claims.”

Reply 23.

Notwithstanding Patent Owner’s arguments and reliance on Exhibits 2039 and 2040, we again agree with Petitioner (Reply 19–20, 23) that Patent Owner has not established a nexus between the challenged claims and the

alleged objective evidence. Whereas Patent Owner argues, for example, that sources stated that Patent Owner’s network interface card “is able to sustain network bandwidth,” “achiev[es] lower processor utilization,” and “is an evolutionary advancement of [Patent Owner’s] . . . protocol acceleration” (PO Resp. 60 (citing Ex. 2039 ¶ 4; Ex. 2040; Ex. 2026 ¶¶ 152–153)), Patent Owner does not demonstrate sufficiently that any of these alleged statements, assuming that any of these statements would have been considered to be “praise” at all, pertain to the claimed invention and in what way. Likewise, whereas Patent Owner contends an analyst “call[ed] the patented technology ‘game changing’” (*id.* at 61), the evidence indicates only that the analyst reported that he had “talked to early-stage customers using the product”—referring in context to one specific product, to which Patent Owner has not persuasively established any connection with the challenged claims—“and they believe it’s game-changing” (Ex. 2040, 3). Again, we note that Patent Owner relies on the same evidence and arguments when asserting secondary considerations for other patents in Cases IPR2017-01391, IPR2017-01393, IPR2017-01405, IPR2017-01406, IPR2017-01409, IPR2017-01410, IPR2018-00226, IPR2018-00234, and IPR2018-00401, casting further doubt on any alleged nexus between the alleged “praise” and the specific subject matter of the challenged claims in this case.

4. *Failure Of Others*

Patent Owner argues that “prior attempts at ‘TCP offload [have] repeatedly failed’” as a “result of the ‘complexities of deploying TCP offload in practice.’” PO Resp. 61 (alteration in original) (citing Ex. 2041,

2). Patent Owner alleges that “[t]he TCP offload described above is a form of network processing offload that is described by the challenged claims, and this failure of others therefore has a direct nexus to the claimed inventions.” *Id.* (citing Ex. 2026 ¶ 154).

Petitioner responds that “Patent Owner provides no evidence of nexus between the single article it cites and the features of the [']241 Patent.”

Reply 23.

We again agree with Petitioner. Even if TCP offload is a form of network processing offload, the Patent Owner provides no evidence linking the failure of others to any limitations of the challenged claims.

5. *Skepticism*

Patent Owner argues “experts and industry were skeptical of offloading processing of complex protocols such as TCP/IP, and expressly taught away from offloading.” PO Resp. 61. Patent Owner points specifically to the same portion of Tanenbaum that it relied upon previously as teaching away from the claimed invention, as well as to a paper published by Dr. Horst that Patent Owner alleges “expressed a high level of skepticism that offloading would result in any beneficial results.” *Id.* at 62–63 (citing Ex. 1006, 588–589; Ex. 2300, 194).

Petitioner responds that, although Tanenbaum “states a preference for an alternative because of expense and complexity, it never suggests a ‘plug-in board with a second CPU and its own program’ will not work with a fast-enough processor, regardless of the offload protocol,” and furthermore, “Dr. Horst’s article in fact confirms the ‘conventional wisdom’ was that special purpose NICs were used for TCP/IP.” Reply 24. Accordingly, Petitioner

contends, Patent Owner's reliance on Tanenbaum and Dr. Horst's article is misplaced. *Id.*

We are not persuaded by Patent Owner's argument for at least the reasons set forth by Petitioner. Further, as previously discussed, Erickson, for example, discloses offloading processing of complex protocols. There can be no nexus if the feature relied upon was previously known in the prior art. *Tokai Corp.*, 632 F.3d at 1369. Nor would one of ordinary skill in the art have been "skeptical" of procedures (e.g., offloading) that a person of ordinary skill in the art would have recognized to have already been disclosed in the prior art (e.g., Erickson).

G. Conclusion Of Obviousness

The secondary consideration evidence in this case is insufficient to persuasively show nonobviousness and does not overcome Petitioner's strong evidence regarding the teachings of Erickson, Tanenbaum, and Alteon with respect to the subject matter of the challenged claims. Accordingly, for the foregoing reasons, we conclude that Petitioner has shown by a preponderance of the evidence that claims 1–8, 18, 22, and 23 would have been obvious to a person of ordinary skill in the art at the time of the '241 patent over Erickson, Tanenbaum, and Alteon and that claims 9–17, 19–21, and 24 would have been obvious to a person of ordinary skill in the art at the time of the '241 patent over Erickson and Tanenbaum.

H. Real Parties-In-Interest

Intel Corporation identified itself as a real party-in-interest in this proceeding and represented that "[n]o other parties exercised or could have

exercised control over this petition; no other parties funded or directed this petition.” Pet. 2.

Patent Owner argues that “[t]he Board should terminate this IPR proceeding because the Petition fails to identify all real parties-in-interest as required by 35 U.S.C. § 312(a)(2) and 37 CFR § 42.8(b)(1).” PO Resp. 63–64. Patent Owner contends, for example, that “Dell is . . . Intel’s . . . customer and indemnitee,” that “Dell, Cavium, and Intel have closely intertwined financial interests and business relationships; express indemnification agreements; shared experts; and common litigation strategy with respect to their defense” and that “the fact that Cavium, Intel, Dell, and Wistron all filed almost verbatim petitions and share the same expert is compelling evidence that they were in privy and cooperating in the drafting of each other’s petitions.” *Id.* at 64.

We note that we previously addressed Patent Owner’s arguments in this regard both in our Decision on Institution and in an Order denying Patent Owner’s request for additional discovery and associated supplemental briefing with respect to real parties in interest in light of the intervening decision of the U.S. Court of Appeals for the Federal Circuit in *Applications in Internet Time, LLC v. RPX Corp.*, 897 F.3d 1336 (Fed. Cir. 2018). *See* Dec. 21–25; Paper 67, 2–3. We additionally note that Cavium, Dell, and Wistron are all acknowledged real parties in interest in this proceeding by virtue of their having joined as parties to this proceeding.

Accordingly, we are therefore not persuaded by Patent Owner’s argument that this proceeding should be terminated for failure to identify all real parties-in-interest.

J. Patent Owner's Motion To Exclude

In *inter partes* review proceedings; documents are admitted into evidence subject to an opposing party asserting objections to the evidence and moving to exclude the evidence. 37 C.F.R. § 42.64.

Patent Owner moves to exclude Tanenbaum (Ex. 1006) as “irrelevant, as Petitioner has failed to establish that [it] is prior art.” Paper 59, 2. Patent Owner also moves to exclude Exhibit 1011, the Declaration of Rice Majors regarding Tanenbaum (“Majors Declaration”), contending that “it is inadmissible hearsay and inadmissible layman opinion.” *Id.*

Petitioner responds that Patent Owner's arguments are directed at to sufficiency of Petitioner's evidence, not the admissibility of Tanenbaum. Paper 62, 3. Moreover, Petitioner contends, Patent Owner never objected to the *relevance* of Tanenbaum and, accordingly, waived the opportunity to move to exclude it based on relevance. *Id.* at 4 (citing Paper 15 (Patent Owner's Objections to Evidence)).

Federal Rule of Evidence 401 provides that evidence is relevant if “it has any tendency to make a fact more or less probable than it would be without the evidence” and “the fact is of consequence in determining the action.” Both the Federal Circuit and the Board have recognized that there is a “low threshold for relevancy.” *See, e.g., OddzOn Prods., Inc. v. Just Toys, Inc.*, 122 F.3d 1396, 1407 (Fed. Cir. 1997); *Laird Techs., Inc. v. GrafTech Int'l Holdings, Inc.*, Case IPR2014-00025, slip op. at 44 (PTAB Mar. 25, 2015) (Paper 45) (“*Laird Techs.*”). There is no question that Tanenbaum is relevant to the patentability of the challenged claims in this case. Tanenbaum is a widely recognized and readily accessible reference explaining network protocols including UDP, TCP, and IP. Further, as

Petitioner points out, the record does not show that Patent Owner timely objected to Tanenbaum based on relevance, and Patent Owner, accordingly, waived that objection. *See* Paper 15, 2–3 (raising objections to Tanenbaum under FRE 901 (authentication) and 801 (hearsay), but not under FRE 401 (relevance)).

Accordingly, we deny Patent Owner’s Motion to Exclude as it relates to Tanenbaum (Ex. 1006).

We also dismiss as moot Patent Owner’s Motion to Exclude as it relates to the Majors Declaration (Ex. 1011), as we do not rely on that Declaration in rendering this Decision.

K. Petitioner’s Motion To Exclude

Petitioner moves to exclude portions of Dr. Almeroth’s Declaration (Exhibit 2026) because, according to Petitioner, portions thereof “are identical to the arguments in the” Patent Owner’s Corrected Response to the Petition and, “[when] counsel for Petitioner asked [Patent Owner’s expert, Dr. Almeroth] why portions of the Patent Owner’s oppositions were identical to the expert’s purported declaration . . . Counsel for Patent Owner instructed Dr. Almeroth not to answer on the basis of privilege.” Paper 58, 1–3.

Although we agree with Petitioner that significant portions of Dr. Almeroth’s declaration indeed are identical to arguments Patent Owner’s Response and lack disclosure of underlying facts or data on which they are based, we nonetheless agree with Patent Owner that “Petitioner’s complaints go to the weight of Dr. Almeroth’s opinions and not their admissibility.” Paper 61, 4. As explained in *Laird Techs.*, “[a] motion to exclude . . . is not

an appropriate mechanism for challenging the sufficiency of evidence or the proper weight that should be afforded an argument.” Case IPR2014-00025, slip op. at 42 (Paper 45). Moreover, “[o]ur general approach for considering challenges to the admissibility of evidence was outlined in *Corning Inc. v. DSM IP Assets B.V.*, Case IPR2013-00053, slip op. at 19 (PTAB May 1, 2014) (Paper 66),” which stated that, “similar to a district court in a bench trial, the Board, sitting as a non-jury tribunal with administrative expertise, is well-positioned to determine and assign appropriate weight to evidence presented.” *Id.* (citing *Donnelly Garment Co. v. NLRB*, 123 F.2d 215, 224 (8th Cir. 1941) (“One who is capable of ruling accurately upon the admissibility of evidence is equally capable of sifting it accurately after it has been received”)).

Accordingly, we deny Petitioner’s motion to exclude.

L. Patent Owner’s Motion To Amend

Patent Owner filed a contingent motion to amend proposing to substitute claims 1–24 with claims 25–48 (a one-for-one substitution). Mot. 1.

Before considering the patentability of any substitute claims, we first must determine whether the motion to amend meets the statutory and regulatory requirements set forth in 35 U.S.C. § 316(d) and 37 C.F.R. § 42.121. A motion to amend may “cancel any challenged patent claim” or, for each challenged claim, “propose a reasonable number of substitute claims.” 35 U.S.C. § 316(d)(1). Our corresponding rule provides that “[t]he presumption is that only one substitute claim would be needed to replace each challenged claim.” 37 C.F.R. § 42.121(a)(3). Furthermore, a motion to

amend “may not enlarge the scope of the claims of the patent or introduce new matter.” 35 U.S.C. § 316(d)(3). Our corresponding rule provides that “[a] motion to amend may be denied where . . . [t]he amendment seeks to enlarge the scope of the claims of the patent or introduce new subject matter.” 37 C.F.R. § 41.121(a)(2)(ii).

In the Federal Circuit’s *en banc* decision in *Aqua Products, Inc. v. Matal*, 872 F.3d 1290 (Fed. Cir. 2017), the lead plurality opinion explains that “the patent owner must satisfy the Board that the statutory criteria in § 316(d)(1)(a)–(b) and § 316(d)(3) are met and that any reasonable procedural obligations imposed by the Director are satisfied.” *Id.* at 1305–06; *see also id.* at 1341 (“There is no disagreement that the patent owner bears a burden of production in accordance [with] 35 U.S.C. § 316(d).” (Reyna, J., writing for a majority)).

On November 21, 2017, the Office provided guidance on motions to amend in view of *Aqua Products*. *See* “Guidance on Motions to Amend in view of *Aqua Products*” (Nov. 21, 2017) (https://www.uspto.gov/sites/default/files/documents/guidance_on_motions_to_amend_11_2017.pdf). In that Guidance, the Office explained:

In light of the *Aqua Products* decision, the Board will not place the burden of persuasion on a patent owner with respect to the patentability of substitute claims presented in a motion to amend. . . . Thus, for example, if the entirety of the evidence of record before the Board is in equipoise as to the unpatentability of one or more substitute claims, the Board will grant the motion to amend with respect to such claims, and the Office will issue a certificate incorporating those claims into the patent at issue.

Beyond that change, generally speaking, practice and procedure before the Board will not change. For example, a patent owner still must meet the requirements for a motion to amend under 37 C.F.R. § 42.121 or § 42.221, as applicable. That

is, a motion to amend must set forth written description support and support for the benefit of a filing date in relation to each substitute claim, and respond to grounds of unpatentability involved in the trial.

Id. at 2.

Still further, on June 1, 2018 (after filing of Patent Owner’s Motion to Amend), the Board designated as “Informative” an order in *Western Digital Corp. v. SPEX Techs., Inc.*, Case IPR2018-00082 (PTAB Apr. 25, 2018) (Paper 13, “*Western Digital Order*”) for its guidance on filing, and analysis, of motions to amend. In the *Western Digital Order*, the Board emphasizes that a motion to amend must “set forth written description support in the originally filed disclosure of the subject patent for each proposed substitute claim, and also set forth the support in an earlier-filed disclosure for each claim for which benefit of the filing date of the earlier filed disclosure is sought.” *Western Digital Order* 7 (citing 37 C.F.R. § 42.121(b)(1)). The *Western Digital Order* further clarifies that “[t]he *motion to amend itself*, not the claim listing . . . , must set forth the written description support.” *Id.* at 8 (emphasis added). Still further, the *Western Digital Order* explains that “the motion must set forth written description support for each proposed substitute claim as a whole, and not just the features added by the amendment. . . . even if the only amendment to the dependent claims is in the identification of the claim from which it depends.” *Id.*

We, therefore, determine whether Patent Owner has met its burden of production of a threshold of evidence sufficient to establish that its Motion to Amend complies with 35 U.S.C. § 316(d)(1)(a)–(b), § 316(d)(3), and 37 C.F.R. § 42.121.

1. Insufficient Written Description For Substitute Claims

With regard to written description support for the proposed substitute claims, Patent Owner's Motion merely argues:

The chart included in Appendix A indicates where support can be found for the substitute claims from the original disclosure of the '241 Patent—from U.S. App. No. 10/260,878 (“the '878 Application”) (Exhibit 2021).

...

The '241 Patent claim priority to U.S. Prov. App. No. 60/061,809, filed on Oct. 14, 1997. “Support in an earlier-filed disclosure” for which “benefit of the filing date of the earlier filed disclosure is sought” is shown in the chart included as Appendix B.

Mot. 3–4.

Initially, we find Patent Owner's argument in the Motion itself is insufficient to satisfy Patent Owner's threshold burden of production for compliance with the requirements of 35 U.S.C. § 316(d) and 37 C.F.R. § 42.121. First, Patent Owner's Motion itself fails to set forth where or how the proposed amended claims are supported, providing instead only a statement that support can be found in the accompanying claim charts. Mot. 3–4. Nevertheless, the information provided in the cited claim charts still fails to produce sufficient evidence of written description support for each substitute claim. For example, the claim chart entry in Appendix A of the Motion for substitute claim 25 (replacing independent claim 1) is reproduced below:

Claims	Exemplary Support in the '878 Application
Proposed Claim 25	
[[1]]25. A method for network communication, the method comprising:	<i>See, e.g.</i> , Ex. 2021 at Abstract, Figs. 3, Figs. 4A, 4B, 4C, and 4D, ¶¶ [0055]-[0064], Cl. 1.
receiving a plurality of packets from the network, each of the packets including a media access control layer header, a network layer header and a transport layer header;	<i>See, e.g.</i> , Ex. 2021 at Abstract, Figs. 3, Figs. 4A, 4B, 4C, and 4D, ¶¶ [0055]-[0064], Cl. 1.
processing the packets by a first mechanism, so that for each packet the network layer header and the transport layer header are validated without an interrupt dividing the processing of the network layer header and the transport layer header;	<i>See, e.g.</i> , Ex. 2021 at Abstract, Figs. 3, Figs. 4A, 4B, 4C, and 4D, ¶¶ [0055]-[0064], Cl. 1.
sorting the packets, dependent upon the processing, into first and second types of packets, so that the packets of the first type each contain data;	<i>See, e.g.</i> , Ex. 2021 at Abstract, Figs. 3, Figs. 4A, 4B, 4C, and 4D, ¶¶ [0055]-[0064], Cl. 1.
sending, by the first mechanism, the data from each packet of the first type to a destination in memory allocated to an application <u>running on a host computer</u> without sending any of the media access control layer headers, network layer headers or transport layer headers to the destination <u>or to a host protocol stack running on the host computer</u> .	<i>See, e.g.</i> , Ex. 2021 at Abstract, Figs. 3, Figs. 4A, 4B, 4C, and 4D, ¶¶ [0017], [0055]-[0064], Cl. 1.

Mot. Appendix A at i. Patent Owner's Appendix A claim chart for substitute claim 25, reproduced above, provides string citations to figures and paragraphs of Exhibit 2021, but provides no explanation as to how the citations disclose the corresponding claim element. The remainder of the claim chart in Appendix A, as well as the claim chart in Appendix B provide similar string citations for each element of each substitute claim, devoid of explanation of how the cited portions disclose the corresponding elements. *Id.* at i—xiv.

In its Opposition to the Motion to Amend, Petitioner argues:

Patent Owner has used string citations to broad swaths of pages and figures in the disclosures (in fact, some of the exact same pages and figures in the original disclosures and priority application are cited for each limitation of the three independent claims) without any attempt to explain how the substitute claims are supported by those disclosures. It is not Petitioner's or the Board's job to sift through these repetitive string citations to piece together Patent Owner's claim of support.

Opp. 1; *see also id.* at 3–5.

In its Reply, Patent Owner argues Petitioner's position is “meritless” and then provides explanation of how the string citations in the claim chart of Appendix A disclose the elements of claims 25–48. PO Reply 3–10. Patent Owner further argues that Petitioner's reliance on *Respironics*⁸ is inapposite because “it is premised upon precedent that has been overturned” by *Aqua Products*. *Id.* at 6. Patent Owner contends *Aqua Products* held that the Petitioner has the burden of persuasion to prove unpatentability and the Patent Owner “need only satisfy its burden of production,” which Patent Owner contends it has done. *Id.* at 6–7.

Patent Owner correctly characterizes the holding of *Aqua Products* that it need only meet the burden of production. However, we disagree with Patent Owner that *Aqua Products* in any manner overturned the Board's decision in *Respironics*. As the Board found in *Respironics*, mere string citations without explanation are insufficient to meet even this lower threshold burden of production. 2014 WL 4715644, at *13; *see also B.E. Tech., L.L.C. v. Google, Inc.*, 2016 WL 6803057, at *7 (Fed. Cir. 2016)

⁸ *Respironics, Inc. v. Zoll Med. Corp.*, No. IPR2013-00322, 2014 WL 4715644, at *13 (PTAB Sept. 17, 2014), *vacated and remanded on other grounds*, 656 F. App'x 531 (Fed. Cir. 2016).

(agreeing with Board that a patent owner did not meet its burden to show written description support for proposed substitute limitations in motion to amend where the patent owner provided only a string citation without explaining how the cited material supported each of the proposed substitute limitations); *Greene's Energy Grp., LLC v. Oil States Energy Servs., LLC*, No. IPR2014-00216, 2015 WL 2089371, at *14 (PTAB May 1, 2015) (“A string citation does not explain how the original disclosure of the application relied upon reasonably conveys to a person the features intended to be encompassed by the proposed substitute claims.”), *aff'd*, 639 F. App'x 639 (Fed. Cir. 2016), *aff'd sub nom. Oil States Energy Servs., LLC v. Greene's Energy Grp., LLC*, 138 S. Ct. 1365 (2018). It is unclear, for example, whether such citations are to be understood as a combination of disclosures that, taken together, disclose the corresponding element, or whether Patent Owner contends that each citation of the string citations is sufficient to disclose the corresponding element. In that regard, Patent Owner provides the same string citations for all of the elements of substitute claim 25 (while adding only a citation to paragraph 17 for the last element). It is not the responsibility of the Board to search through the string citations to find sufficient written description support for each element and we decline an invitation to do so. *See Respironics*, 2014 WL 4715644, at *13 (“Zoll’s string citations amount to little more than an invitation to us (and to Respironics, and to the public) to peruse the cited evidence and piece

together a coherent argument for them. This we will not do; it is the province of advocacy.”).⁹

For the reasons discussed above (Section II.L.1), Patent Owner has failed to meet its burden of production, *in its Motion or in the appended claim charts*, to identify written description support for each substitute claim and we have grounds to deny the Motion to Amend for those reasons. Nevertheless, we address below whether the proposed substitute claims overcome the prior art challenges.

2. *Amendments Do Not Overcome Erickson, Tanenbaum, And Alteon*

Assuming, *arguendo*, that Patent Owner’s Motion to Amend complied with the statutory and regulatory requirements to meet its burden of production, the proposed substitute claims would still not overcome the prior art challenges.

With regard to the Motion to Amend, the parties repeat many of the arguments presented in the Petition and the Patent Owner’s Response. To the extent arguments are repeated, we incorporate our earlier discussion addressing those arguments and, accordingly, we remain persuaded by Petitioner’s repeated arguments and find Patent Owner’s repeated responses unpersuasive. We focus the discussion below on arguments directed specifically to the claim as amended.

⁹ See also *Stampa v. Jackson*, 78 USPQ2d 1567, 1571 (BPAI 2005) (quoting *Ernst Haas Studio, Inc. v. Palm Press, Inc.*, 164 F.3d 110, 111–12 (2d Cir. 1999) (“Appellant’s Brief is at best an invitation to the court to scour the record, research any legal theory that comes to mind, and serve generally as an advocate for appellant. We decline the invitation.”)).

a. Substitute Claim 25 Obvious Over The Art

Proposed substitute claim 25 would amend original, independent claim 1 by modifying the sending step to read,

sending, by the first mechanism, the data from each packet of the first type to a destination in memory allocated to an application running on a host computer without sending any of the media access control layer headers, network layer headers or transport layer headers to the destination or to a host protocol stack running on the host computer.

Mot. Appendix A at i (additions emphasized).

Regarding substitute claim 25, Patent Owner argues that the first step in Alteon's Table 4 (Ex. 1033, 21) specifically recites sending the first 64 received bytes, including the header, to the protocol stack. Mot. 5. According to Patent Owner, this is contrary to the amended requirement that the data be sent without sending any headers to the destination in application memory on the host or to the protocol stack on the host. *Id.* Patent Owner further argues Erickson fails to cure the deficiency of Alteon because "Erickson almost exclusively relates to the transmit side and does not disclose or suggest that that data from received packets are sent to a destination in memory" without sending the headers. *Id.* Patent Owner further argues that Tanenbaum fails to cure the deficiency of Alteon because, in Tanenbaum, all protocol processing is done on the host computer—not on the network adapter. *Id.*

In its Opposition, Petitioner argues that the combination of Erickson, Tanenbaum, and Alteon discloses the sending step, as amended, for the same

reasons presented in the Petition. Opp. 8 (citing Pet. 55–58¹⁰). In particular, Erickson discloses a network interface (direct application interface, “DAI”) that routes received data directly to processes 402 or 404 operable in the host system (i.e., “application[s] running on a host” as claimed). Opp. 8–9 (citing Ex. 1005, 4:53–5:14, Fig. 4; Pet. 55–56). Petitioner further argues that the combination of Erickson, Tanenbaum, and Alteon discloses fast-path processing of packets that avoids processing by the protocol stack on the host system and, thus, discloses sending data to the host applications without sending any headers (TCP, UDP, IP, or MAC layers) to the host as required in the amended claim. Opp. 9–11.

In its Reply, Patent Owner argues Alteon specifically teaches away from the amended recitation of not sending headers to the protocol stack because step 1 of Alteon’s Table 4 specifically recites doing so. PO Reply 11–12. Patent Owner further argues that, because of Alteon’s “teaching away from the amended limitation,” Petitioner pivots to reliance on Erickson and Tanenbaum rather than Alteon for the amended limitation. *Id.* Patent Owner further argues such reliance “is undermined by the fact that its Petition did not argue this combination disclosed or suggested the broader, pre-amendment version of this limitation—let alone the more narrow post-amendment version in the substitute claims.” *Id.* at 12–13 (citing Pet. 56–58). Patent Owner argues Tanenbaum teaches away from the substitute

¹⁰ We note Petitioner (as well as Patent Owner) cites to Paper 2—the original version of the Petition that was supplanted by the Corrected Petition (Paper 4). We discern no pagination difference in the original and corrected versions of the Petition and, therefore, deem the errors harmless. We understand the parties’ references to “Paper 2” to be synonymous with references to the Corrected Petition (Paper 4 (Pet.)).

claims because nowhere does it teach transferring data without headers to the application but instead teaches “traditional header processing by the host.” *Id.* at 13–14 (citing Ex. 1006, 588–89; Ex. 2305 ¶ 32).

Petitioner’ Sur-Reply argues Erickson teaches that protocol processing for fast-path processing resides in the network interface device and, therefore, “Alteon’s teachings when viewed in light of the proposed combination of Erickson and Tanenbaum⁹⁶ are not to transfer headers to the host, but rather to transfer only data without headers to applications on the host.” Sur-Reply 6 (citing Pet. 43–47; Reply 6).

We are persuaded by Petitioner’s arguments. Initially, as discussed *supra*, we are not persuaded by Patent Owner’s repeated arguments that Tanenbaum (or Alteon) teaches away from the claimed features or the proposed combination. Tanenbaum expressly notes that the protocol stack may be implemented on a coprocessor or network card, but articulates a preference for implementing the protocol on the host. Pet. 74; Ex. 1006, 530, 588. Mere expression of a preference, as in Tanenbaum for protocol processing on the host, is insufficient to discredit or discourage the ordinarily skilled artisan from considering alternative embodiments such as implementing the protocol stack on the network interface. *See Gurley*, 27 F.3d at 553; *see also Fulton*, 391 F.3d at 1201.

Moreover, we return to the fundamental conclusion that the combination of references would have taught or suggested the substitute claim features to the ordinarily skilled artisan. Erickson discloses fast-path transmit processing in which data, without headers, is sent from the host application to the network adapter, where headers are prepended (to generate one or more packets) and the generated packets, with prepended headers, are

sent to the network. *See* Pet. 69 (citing Ex. 1005, 7:38–47, 8:2–9; Ex. 1003 Appendix A-28, § V.B.8). Thus, Erickson, in the proposed combination, discloses transferring data between a host and a network interface without transferring any headers. Alteon, in the proposed combination, discloses a fast-path receive process in which a first portion of received data (including a header) is sent to the protocol stack and then the rest of the received data is sent, without headers, to the application bypassing the protocol stack. Pet. 74. In particular, Alteon’s Table 4, cited as support by both parties, discloses “step 1” that expressly transfers headers to the protocol stack, but also discloses “step 2” and “step 3” expressly transferring data without headers to the application. Ex. 1033, 21 (Table 4). Tanenbaum, in the proposed combination, teaches that fast-path processing *may be* implemented on a coprocessor or network board. Ex. 1006, 530. Thus, Tanenbaum suggests to the ordinarily skilled artisan that a protocol stack, including for example the protocol stack disclosed by Alteon or Erickson, may be implemented on a network card. Erickson expressly discloses such an implementation for transmitting data in which header processing for the packet/packets to be transmitted is performed by scripts operable on the network interface card.

Accordingly, we are persuaded that the combination would have taught or suggested to the ordinarily skilled artisan the amended sending step of substitute claim 25 that added the express recitation that data is sent to a destination memory of an application “running on a host computer” without sending any of the specified headers to that destination “or to a host protocol stack running on the host computer.”

b. Substitute Claims 33 And 41 Obvious Over The Art

Independent substitute claim 33 would amend independent claim 9 by adding:

wherein the dividing, prepending, and transmitting occur without the second processor generating an interrupt to the first processor.

Mot. Appendix A at iii (additions emphasized). Proposed substitute claim 41 would amend independent claim 17 by adding:

wherein the dividing, prepending, and transmitting occur without the second mechanism generating an interrupt to the first mechanism.

Id. at v (additions emphasized). Substitute claims 26–32, 34–40, and 42–48 modify dependent claims 2–8, 10–16, and 18–24, respectively, only to change the claim dependency to the corresponding substitute claim

Regarding substitute claims 33 and 41, Patent Owner argues:

In Alteon, the only reference that describes the potential reduction or moderation of interrupts, “[t]he NIC notifies the stack that it has moved 64 bytes of data by issuing an interrupt” in a first control step. Ex. 1033.021. Then, in another control step, “[t]he NIC tells the stack that it has finished moving the rest of the data packet into application memory by issuing an interrupt.” *Id.* There is no disclosure of the dividing, prepending, and transmitting steps all being performed without the generation of an interrupt. Alteon also describes some crude interrupt reduction techniques in connection with data reception in its third generation NICs, but does not disclose any interrupt reduction techniques in connection with data transmission, as described in substitute independent claims 33 and 41. In particular, Alteon does not disclose that the dividing, prepending, and transmitting steps can all be performed without the generation of an interrupt.

Mot. 6.

Initially, we note that the “without . . . interrupt” limitation has been broadly addressed above with respect to related recitations in challenged claims 1–24 (namely without interrupts dividing the prepending of the various headers). Although these amended, substitute claims recite “without . . . interrupt” in relation to other aspects of the recited methods (namely without interrupt between the dividing, prepending, and transmitting steps), we find our analysis above relating to the original challenged claims equally applicable. Thus, for at least the same reasons as discussed *supra* regarding original challenged claims 1–24, we find substitute claims 33 and 41 unpatentable over the combined references.

In addition, we consider below the parties’ arguments regarding Patent Owner’s Motion to Amend and remain persuaded that substitute claims 33 and 41 are unpatentable over the prior art.

Petitioner argues Erickson discloses that all header processing to generate the completed UDP datagram is performed by the network interface device—not the host system—and discloses that its invention is “equally applicable to TCP.” Opp. 15–16 (citing Ex. 1005, 5:47–51, 7:38–47). Therefore, Petitioner argues there would be no reason for the network interface device to interrupt the host computer in between generating (prepending) any of the headers. *Id.* at 16 (citing Dec. 19).

Patent Owner counters that Petitioner has failed to support its conclusions that Erickson could be modified to implement TCP and argues there is no express or inherent disclosure by Erickson that the host is not interrupted. PO Reply 17–19.

Petitioner argues that Patent Owner contradicts itself where it first argues Erickson includes no express or inherent disclosure regarding

interrupts and then argues Erickson discloses reasons why the host “would need to be interrupted to perform the dividing, prepending, or transmitting steps.” Sur-Reply 10 (citing PO Reply 18, 20).

We are persuaded by Petitioner’s arguments. Disregarding the inconsistency of Patent Owner’s argument, in general, Patent Owner’s arguments seek express recitations in the references to match the language of the substitute claims. The test for obviousness is what the combined teachings of the references would have *suggested* to one of ordinary skill in the art. *See Young*, 927 F.2d at 591; *Keller*, 642 F.2d at 425. We are persuaded that the proposed combination including Erickson and Alteon would have suggested to the ordinarily skilled artisan that there is no need to interrupt the host system during Erickson’s processing of headers by the network adapter to prepare a UDP datagram for transmission. As discussed *supra*, Erickson expressly discloses that its invention could be used for TCP protocol processing, Tanenbaum expressly discloses that the protocol stack could be implemented in a coprocessor or network board, and Alteon expressly discloses that a single interrupt may be generated to the host after processing of multiple packets. In view of these express disclosures in the combination, we are persuaded that it would have been suggested to the ordinarily skilled artisan that there is no need in the proposed combination to interrupt the host system while a network adapter prepares one or more TCP segments for transmission as packets.

Patent Owner argues Erickson “teaches and suggests reasons why the host would need to be interrupted” between the recited steps because Erickson’s exemplary script invokes a “nextid() function” that provides a monotonically increasing value (a count) used in the generated headers. PO

Reply 19–20 (citing Ex. 1005, 8:10–13). Patent Owner contends the ordinarily skilled artisan “would understand that the host processor and the network interface processor would want to maintain consistency between the count” by interrupting the host processor. *Id.* at 20 (citing Ex. 2305 ¶ 39).

Petitioner responds that Patent Owner’s assertion that Erickson’s use of the nextid() function would require interrupting the host to assure consistency is mere unsupported speculation. Sur-Reply 10. Petitioner contends that the only description of the nextid() function in Erickson expressly discloses that it is executed on the network adapter and expresses no requirement of storage on the host or a need for consistency with the host. *Id.* at 10–11. Furthermore, Petitioner argues, “Even if consistency was needed, Patent Owner provides no explanation why it would be implemented by undisclosed interrupts rather than the disclosed shared-memory snooping mechanism.” *Id.* at 11 (citing Ex. 1255 ¶¶ 12, 15).

We are persuaded by Petitioner’s arguments. Patent Owner’s contention that interrupts would be required based on Erickson’s invocation of the nextid() function is speculation unsupported by the references. See Ex. 1255 ¶ 12; *see also* Ex. 2305 ¶ 39. Still further, we agree with Petitioner that Erickson expressly discloses a shared memory (“shared registers”) structure between the host system and the network interface card. Ex. 1255 ¶ 15 (citing Ex. 1005, 6:25–31). It would have been obvious to the ordinarily skilled artisan to use such shared memory and Erickson’s “snooping” techniques to assure consistency, if needed, for shared information between the host and the network interface device, such as the monotonically increasing count, thus obviating the alleged need for an interrupt of the host system. *Id.* (citing Ex. 1005, 8:50–52).

Patent Owner further argues Erickson discloses that the user process (on the host) holds certain values that must be set by the user process and “[t]his would require triggering multiple interrupts to the host after each packet is transmitted in order to configure the ‘senduserdatagram’ process for the next packet.” PO Reply 20 (citing Ex. 1005, 7:5–18).

Petitioner asserts that, even if it is true that Erickson would require an interrupt after each TCP packet is transmitted, which Petitioner disputes, the claim limitation would still be met because if an interrupt occurs after each datagram/packet, “the steps of dividing, prepending, and transmitting would have all occurred *prior* to those interrupts being issued.” Sur-Reply 11–12 (citing Ex. 1255 ¶ 18). Petitioner further argues Alteon discloses one interrupt for multiple packets. *Id.* at 12 (citing Ex. 1033, 18, 19, 23; Ex. 1210 Appendix A-5–A-6). Petitioner contends, given Alteon’s express disclosure of one interrupt for multiple packets, “[l]ogic dictates that no interrupt is issued for some of the packets, and thus the amended claim limitation is met.” *Id.*

We are persuaded by Petitioner’s arguments. Initially, we are not persuaded by Patent Owner’s unsupported speculation that modifying Erickson’s UDP scripts to enable TCP processing would necessarily require an interrupt of the host after each packet. Although it is true that no single reference in the proposed combination *expressly* discloses otherwise, we are persuaded by Petitioner’s arguments that the proposed combination at least suggests that Erickson can be modified to provide TCP protocol support (as expressly disclosed in Erickson and at least suggested by Tanenbaum), with reduced interrupts of the host (as expressly disclosed by Alteon), to obviate the alleged need to interrupt the host after each TCP packet transmission (or

reception). We credit Dr. Horst’s testimony that “[i]t would have been within the skills of a [person of ordinary skill in the art] to adapt [Erickson’s] UDP script for TCP.” Ex. 1003 ¶ 139.

Moreover, the “certain values that must be set by the user process” to prepare for a next packet/datagram, as speculated by Patent Owner, could be controlled through Erickson’s disclosed shared memory and snooping structure. Like the monotonically increasing count value generated by the nextid() function, consistency of shared values between the host system and the network interface could be achieved by sharing the values in Erickson’s shared memory, thus obviating the alleged need for interrupting the host system.

c. Obviousness of Substitute Claims 26–32, 34–40, and 42–48

Patent Owner argues these dependent substitute claims, dependent from claims 25, 33, and 41, respectively, overcome the art for the same reasons as their respective base claims.

Petitioner responds that each dependent claim is unpatentable for the same reasons as its corresponding original challenged claim (2–8, 10–16, and 18–24, respectively. Opp. 11–13, 17–19, 20–23.

For at least the reasons discussed above (Sections II.L.2.a–b regarding claims 25, 33, and 41), we are persuaded by Petitioner’s arguments.

3. Conclusions Regarding Substitute Claims

For the reasons discussed *supra* (Section II.L.2), we are persuaded by a preponderance of the evidence that proposed substitute claims are unpatentable as obvious over the combination of Erickson, Tanenbaum, and

Alteon. Therefore, Patent Owner's Motion to Amend is denied because the claims do not overcome the applied references. In addition, Patent Owner's Motion to Amend is denied as failing to meet its burden of production with respect to written description.

M. Joint Motion To Seal

On March 15, 2018, the parties filed a Joint Motion to Enter a Stipulated Protective Order (Paper 38), which was granted on March 27, 2018 (Paper 39).

Patent Owner requests that we seal Exhibit 2038 due to the inclusion of certain licensing terms that Patent Owner regards as confidential. Paper 32. There is a strong public policy in favor of making information filed in an *inter partes* review open to the public, especially because the proceeding determines the patentability of claims in an issued patent and, therefore, affects the rights of the public. See *Garmin Int'l, Inc. v. Cuozzo Speed Techs., LLC*, Case IPR2012-00001 (PTAB Mar. 14, 2013) (Paper 34). Under 35 U.S.C. § 316(a)(1) and 37 C.F.R. § 42.14, the default rule is that all papers filed in an *inter partes* review are open and available for access by the public; a party, however, may file a concurrent motion to seal and the information at issue is sealed pending the outcome of the motion. It is, however, only "confidential information" that is protected from disclosure. 35 U.S.C. § 316(a)(7); see Office Patent Trial Practice Guide, 77 Fed. Reg. 48,756, 48,760 (Aug. 14, 2012). The standard for granting a motion to seal is "for good cause." 37 C.F.R. § 42.54(a). The party moving to seal bears the burden of proof in showing entitlement to the requested

relief, and must explain why the information sought to be sealed constitutes confidential information. 37 C.F.R. § 42.20(c).

In reviewing the Exhibit 2038, we conclude that it may contain confidential information. Accordingly, we are persuaded that good cause exists to have the identified portions remain under seal, and the Motion to Seal is granted.

The Office Patent Trial Practice Guide provides:

Expungement of Confidential Information: Confidential information that is subject to a protective order ordinarily would become public 45 days after denial of a petition to institute a trial or 45 days after final judgment in a trial. There is an expectation that information will be made public where the existence of the information is referred to in a decision to grant or deny a request to institute a review or is identified in a final written decision following a trial. A party seeking to maintain the confidentiality of information, however, may file a motion to expunge the information from the record prior to the information becoming public. § 42.56. The rule balances the needs of the parties to submit confidential information with the public interest in maintaining a complete and understandable file history for public notice purposes. The rule encourages parties to redact sensitive information, where possible, rather than seeking to seal entire documents.

77 Fed. Reg. at 48761.

Consequently, 45 days from entry of this Decision, all information subject to a protective order will be made public by default. In the interim, Patent Owner may file a motion to expunge any such information that is not relied upon in this Decision. See 37 C.F.R. § 42.56.

III. CONCLUSION

For the foregoing reasons, we determine that the information presented establishes, by a preponderance of the evidence, that claims 1–24 of the '241 patent are unpatentable. Furthermore, for the foregoing reasons, we determine that proposed substitute claims 25–48 are unpatentable and, thus, Patent Owner's Motion to Amend is denied.

IV. ORDERS

Accordingly, for the foregoing reasons, it is

ORDERED that claims 1–24 of U.S. Patent No. 7,337,241 B2 are held unpatentable;

FURTHER ORDERED that Patent Owner’s Motion to Exclude Exhibit 1006 is denied;

FURTHER ORDERED that Patent Owner’s Motion to Exclude Exhibit 1011 is dismissed as moot;

FURTHER ORDERED that Petitioner’s Motion to Exclude Exhibit 2026 is denied;

FURTHER ORDERED that Patent Owner’s Contingent Motion to Amend is denied;

FURTHER ORDERED that Patent Owner’s Motion To Seal is granted; and

FURTHER ORDERED that because this is a final written decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2017-01392
Patent 7,337,241 B2

For PETITIONER:

David Xue
Garland T. Stephens
Jeremy Jason Lang
Anne Cappella
Adrian Percer
WEIL, GOTSHAL & MANGES LLP
david.xue@rimonlaw.com
garland.stephens@weil.com
jason.lang@weil.com
anne.cappella@weil.com
adrian.percer@weil.com

For PATENT OWNER:

James M. Glass
Joseph M. Paunovich
Brian E. Mack
Ziyong LiQUINN EMANUEL URQUHART
& SULLIVAN LLP
jimglass@quinnemanuel.com
joepaunovich@quinnemanuel.com
brianmack@quinnemanuel.com