

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,
Petitioner,

v.

PORTSMOUTH NETWORK CORP.,
Patent Owner.

IPR2024-00893
Patent 6,917,986 B2

Before ST. JOHN COURTENAY III, JUSTIN T. ARBES, and
CHARLES J. BOUDREAU, *Administrative Patent Judges*.

ARBES, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining Some Challenged Claims Unpatentable
35 U.S.C. § 318(a)

I. INTRODUCTION

A. Background and Summary

Petitioner Juniper Networks, Inc. filed a Petition (Paper 2, “Pet.”) requesting *inter partes* review of claims 1–24 of U.S. Patent No. 6,917,986 B2 (Ex. 1001, “the ’986 patent”) pursuant to 35 U.S.C. § 311(a).

On December 11, 2024, the Board instituted an *inter partes* review as to all challenged claims on all grounds of unpatentability asserted in the Petition. Paper 12 (“Decision on Institution” or “Dec. on Inst.”). Patent Owner Portsmouth Network Corp. filed a Patent Owner Response (Paper 14, “PO Resp.”), Petitioner filed a Reply (Paper 18, “Reply”), and Patent Owner filed a Sur-Reply (Paper 23, “Sur-Reply”). An oral hearing was held on September 12, 2025, and a transcript of the hearing is included in the record (Paper 30, “Tr.”).

We have jurisdiction under 35 U.S.C. § 6. This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a). For the reasons that follow, we determine that Petitioner has shown, by a preponderance of the evidence, that claims 3, 8, 11, 12, 15, 20, 23, and 24 of the ’986 patent are unpatentable.¹

¹ As explained below, subsequent to institution in this proceeding, the Office issued an *inter partes* review certificate (U.S. Patent No. 6,917,986 K1, Trial No. IPR2024-00505, Certificate Issued Nov. 6, 2025) cancelling claims 1, 2, 4–7, 9, 10, 13, 14, 16–19, 21, and 22 of the ’986 patent. *See* Ex. 3001; *infra* Section I.B. Accordingly, Petitioner’s request for relief as to those claims (i.e., that they be cancelled) is moot, and only claims 3, 8, 11, 12, 15, 20, 23, and 24 remain at issue. *See* Pet. 6, 73; 35 U.S.C. § 311(b); *United Therapeutics Corp. v. Liquidia Techs., Inc.*, 74 F.4th 1360, 1372 (Fed. Cir. 2023) (“claims are cancelled when the Director issues a certificate confirming unpatentability, which occurs only after ‘the time for appeal has expired or any appeal has terminated’” (citing 35 U.S.C. § 318(b))); *Nokia of Am. Corp. v. Gen. Access Sols., Ltd.*, IPR2023-01416, Paper 35 at 2–3, 76–77 (PTAB Apr. 14, 2025) (final written decision not addressing certain challenged claims that were statutorily disclaimed after institution, and addressing only the challenged claims that had not been disclaimed); *Markforged Inc. v. Continuous Composites, Inc.*, IPR2022-00679, Paper 24 at 4, 41–42 (PTAB Sept. 26, 2023) (same).

B. Related Matters

The parties indicate that the '986 patent was the subject of *Portsmouth Network Corp. v. Cisco Systems, Inc.*, No. 23-cv-00441 (E.D. Tex.), which has been dismissed, and is the subject of *Portsmouth Network Corp. v. Juniper Networks, Inc.*, 23-cv-12201 (D. Mass), which has been stayed. *See* Pet. 1; Paper 21, 2. Petitioner filed petitions challenging another patent asserted in the district court cases in IPR2024-00913 (instituted) and IPR2024-00989 (instituted and consolidated with IPR2024-00913). A different petitioner, Cisco Systems, Inc. ("Cisco"), filed petitions challenging other patents asserted in the district court cases in IPR2024-00503 (instituted), IPR2024-00506 (instituted), IPR2024-00954 (denied), and IPR2024-01450 (instituted).

Cisco also filed a petition challenging claims 1, 2, 4–7, 9–14, 16–19, and 21–24 of the '986 patent in IPR2024-00505 ("the Cisco IPR"). On August 4, 2025, the Board issued a final written decision in the Cisco IPR determining that claims 1, 2, 4–7, 9, 10, 13, 14, 16–19, 21, and 22 of the '986 patent are unpatentable but determining that Cisco had not shown that claims 11, 12, 23, and 24 are unpatentable by a preponderance of the evidence on the particular ground presented in that proceeding. IPR2024-00505, Paper 29. Neither party appealed the Board's decision. On November 6, 2025, the Office issued an *inter partes* review certificate stating that "[c]laims 1, 2, 4–7, 9, 10, 13, 14, 16–19, 21, and 22 are cancelled." Ex. 3001.

C. The '986 Patent

The '986 patent, titled “Fast Failure Protection Using Redundant Network Edge Ports,” issued July 12, 2005, from U.S. Patent Application No. 10/036,518, filed on January 7, 2002. Ex. 1001, codes (21), (22), (45), (54). The '986 patent relates to “methods and devices [for] enabling rapid reconfiguration in response to network topology changes” or “rapid recovery from network failures.” *Id.* at col. 3, ll. 58–64. The goal of the '986 patent was to find a way to recover rapidly from a failure in a network that uses a spanning tree protocol to avoid loops in redundant network paths. *Id.* at col. 2, l. 65–col. 3, l. 14, col. 3, ll. 58–64.

More specifically, the “Background of the Invention” section of the '986 patent describes:

To protect against failures of network equipment and communication media, it is common to provide multiple, redundant paths between switches in the network. For example, [] switches are typically connected . . . by dual links. These links may be in the form of connections between each of a pair of ports on one [] switch with a corresponding pair of ports on [another] switch. . . . [T]hese configurations are meant to ensure the connection between the [first switch] and [a second] switch will be maintained even if there is a failure in one of the ports

Ex. 1001, col. 2, ll. 7–22. “A problem with redundant topologies, however, is that the multiple paths between points in the network can create loops that disrupt the proper functioning of the learning process” *Id.* at col. 2, ll. 23–26. “[B]roadcast packets may be forwarded repeatedly by the switches in the loop, causing a serious problem of network congestion, known as a ‘broadcast storm.’” *Id.* at col. 2, ll. 26–29. “In order to avoid these problems, the 802.1D standard defines a communication protocol, known as the spanning tree protocol (STP), for managing the creation and

updating of the network topology.” *Id.* at col. 2, ll. 30–33. “STP ensures that all data paths in a network of bridges^[2] are free of loops by disabling forwarding of packets through certain [ports].” *Id.* at col. 2, ll. 33–36.

“The spanning tree algorithm defines one bridge in the network as the root bridge. Each LAN connected to the network has a bridge port that connects it to the root bridge.” *Id.* at col. 2, ll. 48–50. Therefore, “[e]ach bridge has a port defined as its root port, which uniquely connects that bridge to the root bridge [or switch].”³ *Id.* at col. 2, ll. 55–56. The other ports on the switch (bridge) that do not connect to the root switch “are put into the blocking state.” *Id.* at col. 2, ll. 56–58. “Frames destined to cross the network will thus travel from the designated bridge of the originating LAN along a root path toward the root bridge. If the destination bridge does not lie along the root path, the frame will be routed through the root bridge” *Id.* at col. 2, ll. 58–64. The ’986 patent Specification goes on to describe that networks implementing an STP can take a long time to recover from a link failure:

Although STP is effective in removing network loops, it is very slow in adapting to topology changes that may occur when a network link fails or is subsequently restored. Responding to a failure typically necessitates recalculating the spanning tree to accommodate the topology change and accordingly transferring

² A network switch that connects various components of networks together is a type of bridge, one that has multiple ports on it which, in turn, can each be connected to a port on another switch, local area networks (LANs), wide area networks (WANs), or other devices. Ex. 1001, col. 1, ll. 27–29.

³ Note that the connection from the root port on a switch to the root switch (bridge) does not necessarily have to be a direct connection but can be a path going through other switches. Ex. 1001, col. 2, ll. 48–64. The path from a particular switch toward the root switch is called the root path or designated path. *Id.*

[the] bridge ports from the blocking to the forwarding state. The [] 802.1D standard specifies a procedure for responding to topology changes, which include reducing the aging time for filtering database entries to 15 sec. Using this procedure with default STP timing parameters, . . . it still takes 30 sec to transfer ports from the blocking state to the forwarding state.

Id. at col. 2, l. 65–col. 3, l. 10.

Accordingly, the '986 patent describes that it “is an object of some aspects of the present invention to provide methods and devices enabling rapid reconfiguration in response to network topology changes.” Ex. 1001, col. 3, ll. 57–59. In order to do this, the '986 patent describes a preferred embodiment as follows:

[A]n edge switch . . . is attached by redundant links to one or more access switches. The edge switch selects one of its access ports (i.e., the ports that are connected to the downstream links to the access switches) to be its active access port, while placing its remaining access port(s) in the blocking state. . . .

When the edge switch detects a failure in its active link to the access switches, it immediately chooses one of the blocked links to serve as the new active link and blocks the failed port. In order to inform the access switches and other downstream equipment that they should now transmit their upstream traffic through the new active link, the edge switch sends dummy frames through the new active link to the access switches. . . . Upon receiving these dummy frames, the access switches . . . rebuild their own filtering databases so that upstream traffic is immediately directed to the new active link.

Id. at col. 3, l. 65–col. 4, l. 28.

Figure 2A, illustrating a preferred embodiment, is reproduced below.




Figure 2A is “a block diagram that schematically shows details of the interconnection of switches 26, 27 and 28, in accordance with a preferred embodiment of the present invention.” Ex. 1001, col. 7, ll. 53–56. “[W]hen users 40 and 42 attempt to communicate over network 20, all the traffic between them will pass through switch 26, via links 30 and 34.” *Id.* at col. 8, ll. 62–64. “The method [of the invention] is invoked when processor 44 of switch 28 senses a failure at port P3, . . . referred to [] as the primary port, i.e., the port that is used originally at switch 28.” *Id.* at col. 9, ll. 14–18.

“Upon sensing the failure, processor 44 of edge switch 28 blocks its port P3 and activates its port P4 instead.” Ex. 1001, col. 9, ll. 26–28. “Now edge switch 28 and access switch 27 can immediately begin exchanging traffic over link 32.” *Id.* at col. 9, ll. 34–35. “The filtering databases maintained by the switches are no longer accurate. . . . [Switch 26] will not

begin sending traffic over links 34 and 32 . . . until its filtering database has been updated. Communication is delayed for as long as the database entries are incorrect.” *Id.* at col. 9, ll. 35–41. “In order to shorten this delay, edge switch 28 sends dummy frames downstream to access switch 27 via link 32.” *Id.* at col. 9, ll. 42–44. “When switch 27 receives the dummy frame, it also forwards it to switch 26.” *Id.* at col. 9, ll. 56–57. “Once the dummy frame has been sent, switches 26, 27 and 28 update their filtering databases to reflect the change in the active link from link 30 to link 32” *Id.* at col. 9, ll. 62–65. Sending the dummy frames allows all of the switches to rapidly update their filtering databases after a port failure and switchover so the network can recover more quickly. *Id.* at col. 4, ll. 24–31.

The ’986 patent describes what constitutes access switches, core switches, and edge switches. Access switches have “ports that connect to LANs and/or to other network endpoints, such as user computer systems.” Ex. 1001, col. 1, ll. 31–33. A core switch “interconnect[s] the access switches into a Wide Area Network (WAN), by direct connection between the core switch and access switches and/or by connection to other core switches.” *Id.* at col. 1, ll. 33–36. Edge switches are “[c]ore switches that are directly connected to access switches.” *Id.* at col. 1, ll. 36–38.

Applying these definitions to Figure 2A, switches 26 and 27 are “access switches” because they are connected to network endpoints or LANs as shown in Figure 2A’s depiction of switch 27 being connected to end user/computer system 40. Ex. 1001, col. 7, ll. 52–61, Fig. 1. Switch 28 is an “edge switch” because it is a “core switch” that is “directly connected to access switches”—i.e., it is directly connected to access switches 26 and 27, and also connected to WAN 24, so it connects the access switches to the WAN. *Id.* at col. 1, ll. 36–38.

D. Illustrative Claims

Claims 1, 12, 13, and 24 of the '986 patent are independent. Claims 3, 8, and 11 depend from claim 1, and claims 15, 20, and 23 depend from claim 13. Claims 1, 2, 3, and 12 recite (with letter designations used in the Petition to refer to the various limitations):

1. A method for network communication, comprising:

[A] coupling an edge switch in a network to one or more access switches via at least first and second redundant links, respectively connected to first and second ports of the edge switch;

[B] placing the first port in an active state, while placing the second port in a blocking state;

[C] conveying communication traffic over the first link while the first port is in the active state, thereby causing each of the switches to build a respective database for use in forwarding the traffic;

[D] responsive to a failure associated with the first link, placing the second port in the active state and the first port in the blocking state; and

[E] sending dummy traffic from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state.

2. A method according to claim 1, wherein placing the second port in the blocking state comprises configuring the first and second ports so that no loop is created in the network between the edge switch and the one or more access switches.

3. A method according to claim 2, wherein configuring the first and second ports comprises configuring the edge switch so as to decrease a radius of a spanning tree established in the network.

12. A method for network communication, comprising:

[A] coupling a first bridge in a network to a second bridge located downstream of the first bridge, by connecting at least first

and second redundant links to respective first and second ports of the first bridge so as to communicate with the second bridge;

[B] placing the first port in an active state, while placing the second port in a blocking state;

[C] conveying communication traffic over the first link while the first port is in the active state, thereby causing the bridges to build respective databases for use in forwarding the traffic;

[D] responsive to a failure associated with the first link, placing the second port in the active state and the first port in the blocking state; and

[E] sending dummy traffic from the first bridge over the second link downstream to the second bridge, so as to cause the second bridge to modify its database responsive to the second port being in the active state.

E. Evidence

The grounds of unpatentability in the instant *inter partes* review are based on the following prior art:

U.S. Patent No. 6,032,194, issued Feb. 29, 2000 (Ex. 1005, “Gai”);

U.S. Patent No. 6,330,229 B1, issued Dec. 11, 2001 (Ex. 1006, “Jain”); and

Srinidhi Varadarajan & Tzi-cker Chiueh, *Automatic Fault Detection and Recovery in Real Time Switched Ethernet Networks*, IEEE (1999) (Ex. 1009, “IEEE 1999”).

Petitioner filed a declaration from Bruce McNair, M.E. (Ex. 1003) with its Petition. Patent Owner filed a declaration from James Olivier, Ph.D. (Ex. 2002) with its Response. Also submitted as evidence is the transcript of the deposition of Mr. McNair (Ex. 2004).

F. Asserted Grounds

Petitioner asserts in the Petition that claims 1–24 of the ’86 patent would have been obvious on the basis of prior art patents and printed publications under 35 U.S.C. § 103(a):⁴

Claims Challenged	35 U.S.C. §	References/Basis
1–3, 6–15, 18–24	103(a)	Gai, IEEE 1999
4, 5, 16, 17	103(a)	Gai, IEEE 1999, Jain

Claims 1, 2, 4–7, 9, 10, 13, 14, 16–19, 21, and 22 have been cancelled. Ex. 3001. Accordingly, only Petitioner’s first asserted ground as to claims 3, 8, 11, 12, 15, 20, 23, and 24 remains at issue.⁵

II. ANALYSIS

A. Level of Ordinary Skill in the Art

In determining the level of ordinary skill in the art for a challenged patent, we look to “1) the types of problems encountered in the art; 2) the

⁴ The Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (“AIA”), amended 35 U.S.C. §§ 102 and 103. Because the challenged claims of the ’86 patent have an effective filing date before the effective date of the applicable AIA amendments, we refer to the pre-AIA versions of 35 U.S.C. §§ 102 and 103. *See* Ex. 1001, code (22).

⁵ On July 29, 2025, the Chief Administrative Patent Judge issued a Memorandum titled “Final Written Decision Procedures for AIA Trial Proceedings,” *available at* <https://www.uspto.gov/sites/default/files/documents/20250729095727582.pdf>, stating that, “absent extraordinary circumstances, all Board panels assigned to an inter partes review . . . proceeding . . . will address all grounds raised in the petition in their final written decisions.” To the extent applicable, we view the *inter partes* review certificate’s cancellation of claims 1, 2, 4–7, 9, 10, 13, 14, 16–19, 21, and 22 as an extraordinary circumstance justifying not reaching Petitioner’s asserted ground based on Gai, IEEE 1999, and Jain. *See* Ex. 3001.

prior art solutions to those problems; 3) the rapidity with which innovations are made; 4) the sophistication of the technology; and 5) the educational level of active workers in the field.” *Ruiz v. A.B. Chance Co.*, 234 F.3d 654, 666–67 (Fed. Cir. 2000). “Not all such factors may be present in every case, and one or more of them may predominate.” *Id.* at 667.

Petitioner argues that at the time of the application for the ’986 patent (January 2002), a person of ordinary skill in the art would have had “a bachelor’s degree in electrical engineering or a related field, plus at least two years of industry experience or graduate-level academic experience with computer networking.” Pet. 20. In the Decision on Institution, the Board agreed with and applied Petitioner’s proposed definition, removing the phrase “at least” to avoid unnecessary ambiguity. Dec. on Inst. 18–19. Patent Owner uses Petitioner’s proposed definition for its Response. PO Resp. 5. Neither party made any further argument regarding the level of ordinary skill in the art in its papers during trial.

In the Cisco IPR, we adopted a similar definition, determining that a person of ordinary skill in the art would have had “a bachelor’s degree in computer science, computer engineering, electrical engineering, or in an equivalent field, or approximately two years of experience working in the field of information technology and networking,” where “[l]ack of professional experience can be substituted by additional education, and vice versa.” IPR2024-00505, Paper 29 at 10–11 (alteration in original); *see* Dec. on Inst. 19 n.7 (noting the proposed definition in the Cisco IPR). Based on the full record developed during trial, including our review of the ’986 patent and the types of problems and solutions described in the ’986 patent and cited prior art, we adopt the same definition of the level of ordinary skill in the art as the Cisco IPR and apply it for purposes of this

Decision, but note that our decision would be the same even applying Petitioner's similar proposed definition in this proceeding.

B. Claim Interpretation

We interpret the claims of the challenged patent

using the same claim construction standard that would be used to construe the [claims] in a civil action under 35 U.S.C. 282(b), including construing the [claims] in accordance with the ordinary and customary meaning of such [claims] as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.

37 C.F.R. § 42.100(b). Petitioner states that no claim terms require express interpretation. Pet. 20. Patent Owner also does not propose any express interpretations and states that it “appl[ies] the ordinary and customary meaning to the ’986 patent’s claim terms.” PO Resp. 5. The Board did not interpret any terms in the Decision on Institution. Dec. on Inst. 19–20.

In the Cisco IPR, we addressed four claim terms in the final written decision: “access switch” and “edge switch,” which appear in independent claims 1 and 13 and various dependent claims, and “downstream” and “upstream,” which appear in claims 11, 12, 23, and 24. IPR2024-00505, Paper 29 at 11–14. We interpreted “access switch” to mean “a switch with ports that connect to LANs and/or to other network endpoints, such as user computer systems,” interpreted “edge switch” to mean “a core switch that is directly connected to an access switch,” interpreted “downstream” to mean “away from the root switch,” and interpreted “upstream” to mean “toward the root switch.” *Id.* As noted above, Patent Owner did not appeal that decision. We incorporate herein our previous analysis in the Cisco IPR, but note that none of the analysis of the disputed issues below depends on such

interpretations. *See id.* No other terms require express interpretation to determine whether Petitioner has met its burden to prove unpatentability of the challenged claims. *See Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (“Because we need only construe terms ‘that are in controversy, and *only to the extent necessary* to resolve the controversy,’ we need not construe [a particular claim limitation] where the construction is not ‘material to the . . . dispute.’” (emphasis added) (citation omitted)).

C. Legal Standards

To prevail in its challenge to the patentability of claims 3, 8, 11, 12, 15, 20, 23, and 24 of the '986 patent, Petitioner must demonstrate that the claims are unpatentable by a preponderance of the evidence. 35 U.S.C. § 316(e). “In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.” *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016). This burden of persuasion never shifts to Patent Owner. *Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015); *see also In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1376 (Fed. Cir. 2016) (“Where, as here, the only question presented is whether due consideration of the four *Graham* factors renders a claim or claims obvious, no burden shifts from the patent challenger to the patentee.”).

A claim is unpatentable for obviousness if, to one of ordinary skill in the pertinent art, “the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made.” *KSR Int’l Co. v.*

Teleflex Inc., 550 U.S. 398, 406 (2007) (quoting 35 U.S.C. § 103(a)). The question of obviousness is resolved on the basis of underlying factual determinations, including “the scope and content of the prior art”; “differences between the prior art and the claims at issue”; and “the level of ordinary skill in the pertinent art.” *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Additionally, objective indicia of nonobviousness, such as “commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.”⁶ *Id.* When conducting an obviousness analysis, we consider a prior art reference “not only for what it expressly teaches, but also for what it fairly suggests.” *Bradium Techs. LLC v. Iancu*, 923 F.3d 1032, 1049 (Fed. Cir. 2019).

A patent claim “is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.” *KSR*, 550 U.S. at 418. An obviousness determination based on a combination of references requires finding “both ‘that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.’” *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367–68 (Fed. Cir. 2016) (quoting *Kinetic Concepts, Inc. v. Smith & Nephew, Inc.*, 688 F.3d 1342, 1360 (Fed. Cir. 2012)); see *KSR*, 550 U.S. at 418 (for an obviousness analysis, “it can be important to identify a reason that would have prompted a person of

⁶ Patent Owner does not present any evidence of objective indicia of nonobviousness in this proceeding.

ordinary skill in the relevant field to combine the elements in the way the claimed new invention does”).

“Although the *KSR* test is flexible, the Board ‘must still be careful not to allow hindsight reconstruction of references . . . without any explanation as to *how* or *why* the references would be combined to produce the claimed invention.’” *TriVascular, Inc. v. Samuels*, 812 F.3d 1056, 1066 (Fed. Cir. 2016) (quoting *Kinetic Concepts*, 688 F.3d at 1368) (alteration in original). Further, an assertion of obviousness “cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness.” *KSR*, 550 U.S. at 418 (quoting *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)); accord *In re NuVasive, Inc.*, 842 F.3d 1376, 1383 (Fed. Cir. 2016) (stating that “conclusory statements” amount to an “insufficient articulation[] of motivation to combine”; “instead, the finding must be supported by a ‘reasoned explanation’” (quoting *In re Lee*, 277 F.3d 1338, 1342 (Fed. Cir. 2002))); *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (“To satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements. The petitioner must instead articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness.”).

D. Obviousness Ground Based on Gai and IEEE 1999

1. Gai

Gai is a U.S. patent titled “Method and Apparatus for Rapidly Reconfiguring Computer Networks.”⁷ Ex. 1005, code (54). Similar to the ’986 patent, Gai also describes a method for rapidly recovering from a switch port or link failure in a network that implements a spanning tree protocol/algorithm. *Id.* at code (57) (“A method that rapidly reconfigures a computer network having a plurality of devices executing the spanning tree algorithm.”).

Gai describes that “one or more devices are configured and arranged so that one port, providing connectivity to the root, is in the forwarding state and the remaining ports, providing connectivity to the root, are in the blocked state.” Ex. 1005, code (57). “Upon detection of a failure at the active forwarding port, one of the back-up ports immediately transitions from blocked to forwarding, thereby becoming the new active port for the device.” *Id.* “[D]ummy multicast messages are transmitted, each containing the source address of an entity directly coupled to the affected device or downstream thereof. By examining the dummy multicast messages, other devices in the network learn to use the new forwarding port of the affected device.” *Id.*

⁷ The ’986 patent refers to Gai by number and states that Gai’s “disclosure is incorporated herein by reference.” Ex. 1001, col. 3, ll. 15–41, col. 4, ll. 32–43.

Figure 1 of Gai is reproduced below.




FIG. 1

Figure 1 depicts an embodiment of the invention implemented in computer network 100 with access switches 114–116 and backbone switches 120–125. Ex. 1005, col. 7, ll. 3–14. “[F]or each access switch 114–116, only one port . . . that represents a path from the access switch to the root . . . will be forwarding. All other ports . . . that represent paths from the access switch to the root will be blocked.” *Id.* at col. 11, ll. 8–13. “Should a . . . failure disabling the link coupled to a root port of an access switch [occur], the affected access switch will be able to rapidly reconfigure the network 100 without the significant interruption . . . experienced through conventional operation of the spanning tree algorithm.” *Id.* at col. 11, ll. 36–41.

“[A]ssume that port number three [on switch 114], which is coupled to backbone switch 122, is the root port for switch 114 and thus in the forwarding state.” Ex. 1005, col. 11, ll. 41–44. Because port 3 is active, the

other ports on switch 114 that connect to the root switch (directly or indirectly) are blocked as explained by Gai:

[P]ort numbers two and four, which also connect switch 114 to the backbone switches [122 and 124], are in the blocked state. If the link 128 coupled to the port number three at switch 114 fails, either one of these other two ports (i.e., port numbers two or four) will immediately transition to the forwarding state and begin receiving and sending messages.

Id. at col. 11, ll. 44–52.

Gai then gives an example where backup port number 4 on switch 114 is chosen to replace port 3, which has the failed link to switch 122, by transitioning port 4 from a blocked state to a forwarding state. Ex. 1005, col. 12, ll. 13–23, 32–35. Gai continues explaining that “[s]witch 114 . . . preferably informs the other devices of its new forwarding port by transmitting dummy multicast packets through the new port.” *Id.* at col. 12, ll. 62–66. When the dummy multicast messages are received at a new source port on backbone switch 124 that has a different route to network elements, such as to server 112a, switch 124 will update its filtering database to reflect the new information about the network. *See id.* at col. 13, ll. 34–38. Gai further explains that

[t]hereafter, if backbone switch 124 receives a message intended for server 112a, it will use this new destination port and the message will be received at port number four of switch 114, which is now capable of receiving and forwarding messages. Accordingly, the dummy multicast message [] effectively apprises backbone switch 124 of the change in forwarding ports that occurred at access switch 114. Backbone switch 124 distributes the multicast message [] through all of its forwarding ports so that other devices, as necessary, may learn of the new forwarding port (i.e., port number four) at access switch 114.

Id. at col. 13, ll. 38–49.

2. IEEE 1999

IEEE 1999 is a published article titled “Automatic Fault Detection and Recovery in Real Time Switched *Ethernet* Networks.”⁸ Ex. 1009. IEEE 1999 “describes the fault detection and recovery mechanism supported by the *EtheReal* architecture The heart of *EtheReal*’s fault detection and recovery mechanism is a fast spanning tree reconfiguration algorithm to reduce the total fault recovery time.” *Id.* at 161. The article states that “[c]urrent transparent switched *Ethernet* networks use the IEEE 802.1d distributed spanning tree discovery algorithm to detect loops and logically block redundant paths, so that only one route exists between any pair of end hosts.” *Id.* at 164. “While the IEEE 802.1d [algorithm] works well, . . . it has several shortcomings in a real time connection oriented environment.” *Id.* “Since the size of the network is unknown, conservative global timers are used The recommended timer values result in a spanning tree discovery phase of 30 seconds, during which the entire network is unusable.” *Id.* “[F]aults in any part of the network invalidate all the connections through the network. Coupled with a large spanning tree reconfiguration (fault recovery) time, network faults can result in large interruptions in service in IEEE 802.1d based networks.” *Id.*

To address this problem with IEEE 802.1d-based networks, the article proposes a three-phase spanning tree algorithm as follows:

1. A startup phase (SP), which goes from the root of the tree to the leaves.

⁸ Petitioner provides evidence sufficient to show that IEEE 1999 is a prior art printed publication under 35 U.S.C. § 102(b). *See* Pet. 25; Ex. 1010; Ex. 1011; Ex. 1003 ¶ 132. Patent Owner does not argue otherwise in its Response.

2. A startup completion phase (SCP), which goes from the leaves of the tree to the root.

3. A topology discovery complete (TDC) phase which goes from the root of the tree to the leaves.

Ex. 1009, 165. Using this approach, after a fault occurs, “[i]n the interior of the network, the fault would be independently detected by switches both upstream as well as downstream of the fault.” *Id.* at 166. “The fault recovery mechanism on each switch that detects the fault, performs the following functions: 1. Discard all best effort traffic. 2. Terminate all connections going through the offending link. 3. Rebuild the spanning tree. 4. Reestablish the connections that were torn down.” *Id.* Regarding step 2, “[t]he fault recovery mechanism . . . terminates all real time connections going through the offending link and informs the source of the connection that the network is unstable.” *Id.* As to step 3, “[t]he switch[] that detect[s] the failure then initiate[s] a spanning tree reconfiguration, with [itself] as root.” *Id.* “At the end of the spanning tree rebuild phase, all the hosts connected to the network are informed that the network is stable. Hosts that had their connections torn down due to network failure now initiate a new connection to their destination.” *Id.* According to IEEE 1999, using this method after a link failure will result in a network that “merely sees a small interruption in service.” *Id.*

The article further addresses the speed at which the phases of the reconfiguration process take place. Ex. 1009, 167. “From the perspective of the affected connections, the fault recovery mechanism spends time in: 1. Terminating connections going through the faulty link/switch. 2. Reconfiguring the spanning tree[.] 3. Re-establishing the connections that were torn down.” *Id.* “To terminate the connections through a failed link we need to (a) identify the connection identifiers of the failed connections

(b) send a teardown message to the upstream/downstream sources of outgoing connections and (c) tear the connection to the upstream/downstream destinations [of] incoming connections.” *Id.*

3. *Claim 3*

We begin with the first remaining challenged claim, claim 3. Claim 3 depends from claim 2, which depends from claim 1. Patent Owner does not dispute claim 3 specifically, but challenges Petitioner’s asserted combination of Gai and IEEE 1999 with respect to one limitation of parent claim 1, limitation [1E]. *See* PO Resp. 7–11, 19; Sur-Reply 1–12. Thus, the majority of our discussion will be in the context of parent claim 1.⁹

a) Undisputed Limitations

The preamble of claim 1 recites a “method for network communication.” Petitioner persuasively shows that Gai teaches such a method, regardless of whether the preamble is limiting. *See* Pet. 30.


⁹ In the Cisco IPR, we determined claims 1 and 2 to be unpatentable under 35 U.S.C. § 103(a) over Gai alone. IPR2024-00505, Paper 29 at 60–61. Cisco’s arguments in that proceeding, however, were substantially different from those advanced by Petitioner here. For example, Cisco’s first theory of unpatentability (designated “1B”), which we found persuasive, relied on Gai’s switch 114 as an “edge switch,” switch 120 as an “access switch,” and switch 117 added to the embodiment shown in Figure 1. *Id.* at 22–23, 29, 34–35. Cisco’s second theory of unpatentability (designated “1A”), which we did not find persuasive, relied on Gai’s switch 124 as an “edge switch” and switches 115 and 116 as “access switches.” *Id.* at 48–50, 58–59. Among other differences, Petitioner in this proceeding instead relies on switch 122 as an “edge switch” and switch 114 as an “access switch,” and asserts a combination with IEEE 1999 for purposes of limitation [1E] (rather than relying on Gai alone). Pet. 31–33, 40–46. We make our determination herein based on the arguments made by the parties in this proceeding.

Limitation [1A] recites “coupling an edge switch in a network to one or more access switches via at least first and second redundant links, respectively connected to first and second ports of the edge switch.”

Petitioner contends that Gai teaches an “edge switch” (i.e., switch 122), “one or more access switches” (i.e., switch 114), a “first redundant link” (i.e., link 128 connecting port 3 of switch 114 to switch 122), and a “second redundant link” (i.e., link 128 connecting port 2 of switch 114 to switch 122).

Pet. 31–33. Petitioner asserts that switch 114 is an “access switch” because it is described as such in Gai and the ’986 patent, and it has multiple ports for communication paths to other network devices. *Id.* at 31 (citing Ex. 1005, col. 7, ll. 3–9, 34–37, 53–57, Fig. 1; Ex. 1001, col. 3, ll. 15–21 (acknowledging that Gai teaches “access switches”)). As shown in Figure 1 of Gai, switch 114 is connected to LANs 102 and 103. *See id.; supra* Section II.B (interpreting “access switch” to mean “a switch with ports that connect to LANs and/or to other network endpoints, such as user computer systems”). Petitioner further argues that switch 122 is an “edge switch” because it is directly connected to an access switch (i.e., switch 114) and interconnects switch 114 into a WAN. Pet. 28, 31–32 (citing Ex. 1001, col. 1, ll. 33–38); *see supra* Section II.B (interpreting “edge switch” to mean “a core switch that is directly connected to an access switch”).

Petitioner provides the following color version of a portion of Figure 1 of Gai to illustrate what occurs in Gai “before [a] link failure” (Pet. 32).



The figure above shows the alleged “edge switch” (switch 122 in orange), “access switch” (switch 114 in blue) connected to server 112a (in purple), “first redundant link” (in green) connected to a “first port” of the “edge switch” (green dot), and “second redundant link” (in red) connected to a “second port” of the “edge switch” (red dot). *Id.* at 32–33.

Limitation [1B] recites “placing the first port in an active state, while placing the second port in a blocking state.” Petitioner contends that Gai teaches placing a “first port” (shown with a green dot in the figure above) in an active state and placing a “second port” (shown with a red dot in the figure above) in a blocking state. Pet. 34. As Petitioner points out, the ’986 patent describes Gai as a method where “[o]ne of the trunk ports of each switch is placed in the forwarding state, while the others are designated

as back-up ports.” Ex. 1001, col. 3, ll. 24–26; *see* Pet. 34. A “trunk port” is a port that has a link connecting it to another switch. Ex. 1005, col. 1, ll. 43–46 (“Ports used to couple switches to each other are generally referred to as . . . trunk ports . . .”).


Limitation [1C] recites “conveying communication traffic over the first link while the first port is in the active state, thereby causing each of the switches to build a respective database for use in forwarding the traffic.” Petitioner argues that Gai describes traffic initially flowing between port 3 of switch 114 and a corresponding port on switch 122 over link 128 (shown in green in the figure above) “while those ports are in the active state.” Pet. 35 (citing Ex. 1005, col. 7, ll. 30–33 (“Links 128 represent communications paths between various components of the network 100, that carry messages, such as data frames, between switches 114–116, 120–125 and routers 130, 131.”)). Gai makes clear that when ports are in the forwarding state, data frames, or communication traffic, is sent over the port. Ex. 1005, col. 2, ll. 56–63. Communication traffic received at a switch for the first time with an unknown source or destination address on the network will cause a database to be built on that switch. *Id.* at col. 11, ll. 16–33. “Switches typically learn which destination port to use in order to reach a particular entity by noting on which source port the last message originating from that entity was received. This information is then stored by each switch in a block of memory referred to as a filtering database.” *Id.* at col. 1, ll. 50–54.

Petitioner also points to Gai’s description of filtering database 244 in switch 214. Pet. 35–36. “[F]iltering database 244 [is] configured to store address information corresponding to the entities of network 100.” Ex. 1005, col. 8, ll. 32–34. “Specifically, filtering database 244 has a plurality of columns or cells, including a destination address cell [DA],

a destination port cell [PORT] and a corresponding timer cell. Each row or entry in the filtering database 244 preferably corresponds to a particular entity.” *Id.* at col. 8, ll. 34–39. As described in reference to Figure 1, the first time a data frame is transmitted from port 3 of switch 114 to switch 122, switch 122 will check its filtering database and look to see if it has switch 114 in the DA (destination address) field, so the first time a data frame is received from switch 114 there will be no entry, causing a new entry to be added with the corresponding active port added to the PORT field and network address for switch 114 to the DA field. *Id.* at col. 11, ll. 16–33. Once this entry is added, when switch 122 receives a data frame with the destination address for switch 114, it will see that it has that switch in the DA field in its filtering database, and therefore transmit the data frame out from the port specified by the PORT (which is the port that it last received a data frame from switch 114 on). *Id.*

Limitation [1D] recites “responsive to a failure associated with the first link, placing the second port in the active state and the first port in the blocking state.” Petitioner argues that Gai teaches, in response to a failure associated with the “first link” (i.e., link 128 connecting port 3 of switch 114 to switch 122), placing the “second port” of switch 122 in the active state and the “first port” of switch 122 in the blocking state. Pet. 37–40.

Petitioner provides the following color version of a portion of Figure 1 of Gai to illustrate what occurs in Gai “after [a] link failure” (Pet. 39).



The figure above shows a failure associated with the “first link” (link 128 connecting port 3 of switch 114 to switch 122), in response to which the second port of the “edge switch” transitions to a forwarding state (now shown with a green dot) and the “first port” transitions to a blocking state (now shown with a red dot). *Id.*

As Petitioner points out, the ’986 patent itself acknowledges that Gai teaches that, “[u]pon detection of a failure at the active forwarding port, the switch immediately places one of the back-up trunk ports in the forwarding state, so that it becomes the new active port.” Ex. 1001, col. 3, ll. 26–29; *see* Pet. 37. Petitioner further argues that in Gai, both “access switches 114–116 and backbone switches 120–125 detect failures.” Pet. 37 (quoting Ex. 1005, col. 11, ll. 64–67). And the switches detect failures in two different ways:

(1) when “a switch stops receiving [bridge protocol data units (BPDUs)] on its root port,” or (2) “by performing a link integrity test that ‘typically exchanges test messages across the respective link 128 at a relatively high rate (e.g., every ten milliseconds).’” *Id.* (quoting Ex. 1005, col. 4, ll. 1–3, col. 12, ll. 4–9). Thus, according to Petitioner, a person of ordinary skill in the art would have understood that in Gai, “a failed link between a backbone switch (e.g., backbone switch 122) and an access switch (e.g., access switch 114) would be recognized both by the backbone switch (as a downstream link failure) and by the access switch (as an upstream link failure).” *Id.* (citing Ex. 1003 ¶ 161).

Claim 2 adds the additional limitation that “placing the second port in the blocking state comprises configuring the first and second ports so that no loop is created in the network between the edge switch and the one or more access switches.” Petitioner contends that “Gai discloses switches that execute a spanning tree algorithm, which calculates a spanning tree that is loop-free,” and “[f]or each switch, only one root port is placed in the forwarding state while the remaining root ports on that switch are blocked.” Pet. 46 (citing Ex. 1005, col. 2, ll. 30–41, col. 6, ll. 33–38, col. 7, ll. 59–61 (“To avoid the creation of loops, these switches execute the spanning tree algorithm . . .”), col. 11, ll. 8–13). Petitioner argues that, in the context of Figure 1 discussed above,

[b]y placing only one designated port of backbone switch 122 that is connected to access switch 114 in the forwarding state while placing other designated ports that are connected to access switch 114 in the blocking state, no loop is created in the network between backbone switch 122 and access switch 114.

Id. at 47.

Finally, claim 3 adds the additional limitation that “configuring the first and second ports comprises configuring the edge switch so as to decrease a radius of a spanning tree established in the network.” Petitioner contends that Gai accomplishes reducing the radius of a spanning tree network in the same way as described in the ’986 patent, by placing one of its ports connecting to an access switch in an active state while blocking the remaining ports that connect to that switch. Pet. 34, 47–48 (citing Ex. 1001, col. 3, l. 67–col. 4, l. 4, col. 4, ll. 35–37, 44–48; Ex. 1005, col. 5, ll. 36–43, col. 11, ll. 42–48).

Patent Owner in its Response does not dispute that Gai teaches the above limitations of claims 1–3. Petitioner’s analysis for each of the limitations, supported by the testimony of Mr. McNair, which we credit, is persuasive. See Pet. 30–40, 46–48; Ex. 1003 ¶¶ 147–164, 176–178.

b) Disputed Limitation [1E]: Combination with IEEE 1999

Limitation [1E] recites “sending dummy traffic from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state.” Petitioner argues that the ’986 patent, in describing Gai, “admits that Gai discloses sending dummy multicast messages (*i.e.*, dummy traffic) *upstream* using a root port after detection of a link failure.” Pet. 40 (citing Ex. 1001, col. 3, ll. 26–35) (second emphasis added). Petitioner further contends that Gai teaches that its rapid reconfiguration can be implemented on “any switch,” as shown in Figure 2 (rapid reconfiguration entity 234) and the accompanying description. *Id.* at 42–43 (citing Ex. 1005, col. 7, ll. 62–64, col. 8, ll. 12–14,

col. 8, l. 66–col. 9, l. 3, col. 12, ll. 59–66). Therefore, according to Petitioner,

Gai discloses sending dummy multicast message to upstream devices after a link failure. As an example, Gai discloses access switch 114 selecting port 4 as a new root port after a link failure Access switch 114 then sends a dummy multicast message with a source address of server 112a through a new active link to backbone switch 124. . . . Backbone switch 124 updates its filtering database by replacing previous information for server 112a with information referencing new port 4 of access switch 114.

Id. at 44 (citing Ex. 1005, col. 13, ll. 14–23, 34–38).

Petitioner argues that a person of ordinary skill in the art would have been motivated, based on the teachings of IEEE 1999, to modify Gai’s operation to “send dummy multicast messages both upstream, as explained [in Gai], as well as downstream,” “regardless of whether a link failure is detected upstream or downstream of the switch.” Pet. 26, 29, 45–46. Thus, as modified, Gai’s “edge switch” (i.e., switch 122) would send dummy traffic upstream toward the root switch and also downstream to an “access switch” (i.e., switch 114). *Id.* at 45–46. Doing so would be performed using the same Enable_Uplinkfast command already described in Gai to block a failed port and use a new active port, and would “caus[e] the [downstream] access switch to update its filtering database.” *Id.* at 27–28, 46.

Petitioner explains why and how a person of ordinary skill in the art would have been motivated to combine the teachings of Gai and IEEE 1999. *See* Pet. 25–30, 45. Petitioner argues that “IEEE 1999 addresses the same challenges as Gai and the ’986 patent with respect to the same IEEE 802.1D standard.” *Id.* at 25 (citing Ex. 1009, 161–162, 164–165). According to Petitioner, IEEE 1999 teaches that it is advantageous to send network

reconfiguration messages both upstream and downstream when recovering from a link failure. *Id.* at 27, 46. “As part of the spanning tree reconfiguration, to improve recovery time, each switch that detected a link failure sends ‘teardown messages’ to both upstream and downstream switches to ‘terminate the connections through a failed link.’” *Id.* at 25 (quoting Ex. 1009, 167). Petitioner asserts that IEEE 1999’s teardown messages cause the switches receiving them to update their filtering databases, terminating the failed connection. *Id.* at 26.

Petitioner argues that a person of ordinary skill in the art would have had reason to modify Gai to send dummy traffic both upstream and downstream “for the benefit of reducing the total link failure recovery time because the link failure can be addressed by upstream and downstream switches in parallel.” Pet. 45. Petitioner contends that IEEE 1999 has “the same purpose that Gai describes [for] sending dummy traffic: increasing the speed of self-healing following [a] link failure in an 802.1D network.” *Id.* at 27. According to Petitioner, sending dummy traffic both upstream and downstream would “achieve[] Gai’s goal (shared by IEEE 1999) of reducing the time taken to rectify network failures.” *Id.* at 26, 29.

Finally, Petitioner argues that a person of ordinary skill in the art would have had a reasonable expectation of success in modifying Gai’s operation to send dummy traffic upstream and downstream “because signaling protocols are routinely modified to provide for different control signaling, such that there would have been no additional technical challenge in sending those messages both upstream and downstream.” Pet. 30. Petitioner notes that “Gai’s mechanism for sending dummy traffic (*i.e.*, rapid reconfiguration entity 234) is implemented by both its backbone switches and its access switches,” and the modified technique for sending dummy

traffic would be performed at backbone switch 122 (i.e., the alleged “edge switch,” as required by limitation [1E]). *Id.* Petitioner’s contentions regarding a motivation to combine and reasonable expectation of success are supported by testimony from Mr. McNair. *See* Ex. 1003 ¶¶ 132–145, 165–174.

Patent Owner argues in response that Gai and IEEE 1999 do not teach or suggest limitation [1E]. PO Resp. 9–11 (citing Ex. 2002 ¶¶ 54–57). Patent Owner argues that IEEE 1999’s teardown messages are “transmitted to tear down specific connections at the source of these connections, before the spanning tree is reconfigured and before each switch begins forwarding traffic along new routes.” *Id.* at 10–11. Therefore, according to Patent Owner, a person of ordinary skill in the art would understand that the teardown messages “play no role in modifying or updating filtering databases in the switches.” *Id.* at 11. Patent Owner further argues that Petitioner fails to explain how a person of ordinary skill in the art would “incorporate IEEE 1999’s disclosure of sending teardown messages to host computers to terminate connections going through a faulty link to modify Gai’s dummy messages that are sent upstream so that ‘other devices within the network learn to utilize the new forwarding port, rather than the failed port.’” Sur-Reply 12 (quoting Ex. 1005, col. 5, l. 65–col. 6, l. 1). Similarly, Patent Owner contends that even if Gai were modified to send dummy messages downstream, Petitioner does not show that they would cause the receiving switch to “modify its respective database responsive to the second port being in the active state,” as recited in limitation [1E]. *Id.* at 10–12.

To the extent Patent Owner’s position is that IEEE 1999’s teardown messages are not “dummy traffic,” or that IEEE 1999’s sending of those messages does not cause an access switch to “modify its respective

database” as recited in limitation [1E], such arguments are unavailing, given Petitioner’s asserted combination. *See* PO Resp. 9–11; Sur-Reply 10–12; *Palo Alto Networks, Inc. v. Centripetal Networks, LLC*, 122 F.4th 1378, 1386 (Fed. Cir. 2024) (“Where, as here, the grounds for obviousness are based on a specific combination of references, arguments that ‘attack the disclosures of the two references individually’ lack merit. . . . [References asserted in an obviousness combination] must be read together, not in isolation.”). Petitioner relies on *Gai’s dummy messages*, which are sent upstream and cause a receiving switch to modify its database,¹⁰ and asserts that an ordinarily skilled artisan would have been motivated, based on IEEE 1999, to modify *Gai* to also send them downstream. *See* Pet. 26 (referring to “*Gai’s dummy traffic*” being sent both upstream and downstream), 29 (same), 46 (same); Ex. 1003 ¶¶ 137, 144, 174. Petitioner does not argue that the teardown message itself should be taken from IEEE 1999 and incorporated into *Gai’s* rapid reconfiguration process. In other words, Petitioner relies on IEEE 1999 solely for its “teaching [of] the known benefits of upstream and downstream message transmission for reconfiguration—not for the content and/or format of the messages themselves.” *See* Reply 3.

Petitioner also explains sufficiently how *Gai* would operate if modified based on IEEE 1999. As Petitioner points out, *Gai* teaches that the rapid reconfiguration process can be implemented on both edge and access switches. Pet. 43 (citing Ex. 1005, col. 8, l. 66–col. 9, l. 3, disclosing rapid

¹⁰ The parties agree that *Gai* transmits its dummy messages upstream; the only dispute is whether there is sufficient basis for Petitioner’s asserted modification to also send them downstream. *See* Pet. 26; PO Resp. 7; Sur-Reply 8.

reconfiguration entity 234 implemented on backbone switch 120, which indicates that it can be applied to any other switch, such as backbone switch 122); *see* Ex. 1005, col. 6, ll. 6–9 (“[T]he method and apparatus manifests, in part, as a series of novel commands that may be entered at the devices. The devices . . . may be classified as either access switches or backbone switches.”), col. 12, ll. 62–66 (describing rapid reconfiguration entity 234 in the context of access switch 114). When implemented on an edge switch, that dummy traffic could be sent both upstream and downstream in order to reconfigure the network to reflect the updated configuration of active ports. *See* Pet. 44–45 (citing Ex. 1003 ¶¶ 171–173). For example, the rapid reconfiguration process could be implemented on edge switch 122 reacting to a failed port connected to access switch 114, and in that case Gai’s configuration command would cause dummy traffic to be sent from switch 122 downstream to access switch 114 in order to reconfigure it. *See id.* at 27–28. We credit the following explanation provided by Mr. McNair, which is consistent with the disclosure of Gai:

Gai discloses entering an Enable_Uplinkfast command at an access switch (*e.g.*, access switch 114) to block a failed port (*e.g.*, the root port) and configure the access switch to send dummy traffic upstream through a new active port. . . . When modified in view of IEEE 1999’s disclosure that self-healing network traffic can advantageously be sent both upstream and downstream, the command would be entered at a backbone switch (*e.g.*, backbone switch 122) to block a failed port (*e.g.*, a designated port connected to access switch 114) and configure the backbone switch to send dummy traffic downstream to an access switch (*e.g.* access switch 114) through a new active port. While Gai uses its protocol signaling command, Enable_Uplinkfast, to send dummy traffic through a new root port, [a person of ordinary skill in the art] would deem it an obvious implementation detail to modify the parameters of the

command to be applicable to other ports, such as a designated port on a backbone switch.

See Ex. 1003 ¶ 139 (citing Ex. 1005, col. 6, ll. 39–46).

Petitioner has made a sufficient showing that Gai teaches “dummy traffic” sent to a switch so as to cause it to “modify its respective database responsive to [a] second port being in the active state,” and that Gai, as modified based on IEEE 1999, teaches the full limitation of “sending dummy traffic from the edge switch over the second link [downstream] to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state.” We find, based on the full trial record, that the combination of Gai and IEEE 1999 teaches limitation [1E].

Next, Patent Owner argues that a person of ordinary skill in the art would not have been motivated to combine the teachings of Gai and IEEE 1999 because the “purpose and timing” of the two references’ messages are different. PO Resp. 7–11 (citing Ex. 2002 ¶¶ 46–48); Sur-Reply 1–7. As to the purposes of the messages, Patent Owner contends that “IEEE 1999’s teardown messages are transmitted to host computers to tear down connections going across the faulty link and are not received by switches in the network,” and “are not used to learn a MAC address” but rather “to inform host computers that the network is unstable and that their particular connections have been torn down.” PO Resp. 8 (citing Ex. 1009, 166; Ex. 2004, 39:22–40:2). By contrast, according to Patent Owner, Gai’s dummy messages are sent “as part of [a] reconfiguration process” “so that ‘other devices within the network learn to utilize the new forwarding port, rather than the failed port, when directing messages to these entities.’” *Id.* (quoting Ex. 1005, col. 5, l. 54–col. 6, l. 5). As to timing, Patent Owner

contends that IEEE 1999's teardown messages are sent "before the spanning tree is reconfigured and before each switch begins forwarding traffic along new routes," whereas Gai's dummy messages are "only transmitted once a device has already successfully transitioned to a new forwarding port." *Id.* (citing Ex. 1009, 166–167). Patent Owner argues that these "significant differences" show that a person of ordinary skill in the art "would not have been motivated to modify Gai based on [IEEE 1999's] generalized statement about bidirectionally sending messages," as Petitioner contends. Sur-Reply 3–5.

We disagree with Patent Owner that the differences between Gai and IEEE 1999 are so substantial that a person of ordinary skill in the art would not have been motivated to incorporate IEEE 1999's teaching regarding bidirectional communication into the system of Gai. As an initial matter, despite the differences pointed out by Patent Owner, Gai and IEEE 1999 are similar in a number of respects. The references pertain to the same subject matter—updating the network reconfiguration after a link failure on a network that implements an IEEE 802.1d spanning tree algorithm—and both are concerned with the speed at which reconfiguration takes place after a link failure. *See* Pet. 25; Ex. 1003 ¶¶ 100, 132–134; Ex. 1005, col. 4, l. 54–col. 5, l. 19 (explaining that reconfiguration is "a time consuming process" where "message delivery is often delayed as ports transition between states," and stating that it was "an object of the present invention" to "reduc[e] the time necessary to reconfigure the network following a change, such as a link failure or recovery"), col. 6, ll. 42–53 ("Reconfiguration of the network may . . . be accomplished substantially sooner than the time required by the conventional spanning tree algorithm while still avoiding the formation of loops."), col. 15, ll. 21–25; Ex. 1009,

161–162 (describing a “fast spanning tree reconfiguration algorithm to reduce the total fault recovery time,” which was “a significant improvement over the standard IEEE 802.1d implementation”), 165 (“The fault recovery mechanism has to be fast, so that the affected connections can be reestablished with minimal interruption in service.”). And they employ similar approaches to reconfiguring the switches by reconfiguring from failed ports to backup ports and sending messages to neighboring switches so they update their configurations. *See* Ex. 1005, col. 6, ll. 42–53, col. 15, ll. 21–25; Ex. 1009, 166–167.

Further, as Patent Owner acknowledges, the references relate to “individual parts of” the same “overall process” for reconfiguration after a link failure. *See* Tr. 11:8–17, 30:15–19. To illustrate, IEEE 1999 assesses the speed at which three phases of the process take place: “1. Terminating connections going through the faulty link/switch. 2. Reconfiguring the spanning tree[.] 3. Re-establishing the connections that were torn down.” Ex. 1009, 167. The teardown messages sent upstream and downstream are part of the first phase, whereas Gai’s dummy messages are sent as part of what would occur in a second phase (reconfiguring the spanning tree). *See id.*; Ex. 1005, col. 5, l. 35–col. 6, l. 5. IEEE 1999 discusses the phases together as part of the same overall fault recovery process, providing the following table (Ex. 1009, 167).

Fault Recovery Component		Time Spent
Connection Termination	Connection Queue Processing	95 ns/connection
	Message transmission	2us/message
Spanning Tree Reconfiguration	SP Reply Timer	30ms
	Message Processing (SP, SCP, TDC)	35ns, 20ns, 15ns per message
	Message transmission	2us
Connection Reestablishment	Total Cost	8-14ms (on the host) 0.1-0.6 ms/link

Table 1: *Timing breakdown of the various components of fault recovery.*

Table 1 above provides a “[t]iming breakdown” of the speed at which the three phases of the process take place, showing, for example, the time for “[m]essage transmission” in both the first phase where the teardown messages are sent and the second phase of spanning tree reconfiguration. *Id.* Although IEEE 1999’s teardown messages and Gai’s dummy messages serve different purposes in each phase, they are both “network reconfiguration messages” intended to inform other devices on the network of something. *See* Reply 1; Ex. 1005, col. 5, l. 35–col. 6, l. 5; Ex. 1009, 166–167.

Patent Owner contends that the two phases are “discrete parts of [the] process” and a person of ordinary skill in the art would not have been motivated to take a teaching from one (connection termination) and apply it to the other (spanning tree reconfiguration). *See* Tr. 30:15–31:24. Given the similarities between the references discussed above, as well as the fact that IEEE 1999 discusses the phases together, though, we are persuaded that a person of ordinary skill in the art, desiring to make the overall reconfiguration process faster, would have considered doing so. *See KSR*, 550 U.S. at 418–420 (holding that it is often necessary to “look to interrelated teachings of multiple patents” to determine whether there was a reason to combine known elements in the manner claimed, as “in many cases a person of ordinary skill will be able to fit the teachings of multiple patents together like pieces of a puzzle”); *Intel Corp. v. PACT XPP Schweiz AG*, 61 F.4th 1373, 1379 (Fed. Cir. 2023) (“The motivation-to-combine analysis is a flexible one. ‘[A]ny need or problem known in the field of endeavor at the time of invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.’” (quoting *KSR*, 550 U.S. at 420) (alteration in original)).

Importantly, there does not appear to be any factual dispute that, if Gai were modified as Petitioner asserts to send dummy messages both upstream and downstream, the reconfiguration process would be faster. For example, in the context of Figure 1 of Gai, switch 122 sending dummy messages upstream to switch 120 and downstream to switch 114 would be faster (i.e., less “hops” between switches) than Gai’s original operation of sending upstream to the root switch¹¹ and the root switch propagating information downstream to switch 114. *See* Tr. 9:26–10:10, 16:3–25, 39:7–24, 41:11–16; Ex. 1005, col. 12, ll. 56–66; Ex. 1003 ¶¶ 56 (explaining that in the standard spanning tree algorithm, “if a node interior to the network detects a failure, . . . the node will transmit a message upstream to the Root Bridge alerting it of this failure” and “[t]he Root Bridge, in turn, rebroadcasts a message to all the downstream nodes in the network, alerting them to the failure”), 173–174 (describing “the known benefit of reducing the total link recovery time since the failure notification is propagated in all directions substantially simultaneously” such that “the link failure can be addressed by upstream and downstream switches in parallel”). Patent Owner does not argue otherwise or point to any testimony from Dr. Olivier to the contrary. *See* Tr. 35:11–25 (stating that Patent Owner does not “necessarily disagree with the fact that if Gai’s system were capable and designed to work that way, you would decrease the number of hops if you were at switch 122 and send both upwards and downstream”). Based on the full trial record, we are persuaded that a person of ordinary skill in the art

¹¹ Gai describes either switch 120 or switch 121 as the root switch of the network shown in Figure 1. *See* Ex. 1005, col. 10, ll. 37–41, col. 11, ll. 60–64, col. 14, ll. 16–20, col. 15, ll. 27–29, 38–51, col. 16, ll. 42–45.

would have readily appreciated the time savings that would occur by modifying Gai to send upstream and downstream.

Finally, we note that the cases cited by Patent Owner in its Sur-Reply are distinguishable. *See* Sur-Reply 1–2, 6–7, 11–12. First, in *Apple Inc. v. INVT SPE LLC*, 851 F. App’x 200 (Fed. Cir. 2021), the Court found substantial evidence for the Board’s finding of no motivation to combine two references where, although the references both “involve[d] measuring transmission power, they [did] so for different reasons, on different time scales, and compare[d] the results to different thresholds.” *Id.* at 202–203; *see* Sur-Reply 1–2. As explained above, Gai and IEEE 1999, in contrast, both relate to an overall network reconfiguration process in the event of a link failure, they both involve messages sent from one device to others in the network as part of that process, and they both are concerned with increasing the speed at which reconfiguration takes place. Although the two references largely pertain to different phases within that process, we are not persuaded that their functions are so different that a person of ordinary skill in the art would not have considered combining them, as was the case in *INVT*.

Second, Patent Owner argues that “general motivations” and the fact that references are “directed to the same art or same techniques” are insufficient to show a motivation to combine, citing *Microsoft Corp. v. Enfish, LLC*, 662 F. App’x 981, 990 (Fed. Cir. 2016), and other cases. Sur-Reply 6–7 & n.3. That is not the sum of Petitioner’s argument, though. Petitioner points to a specific benefit that would be achieved by modifying Gai to also send dummy messages downstream—faster reconfiguration after a link failure (which Gai already expressed a desire for)—and points to support for that benefit in the cited references and the detailed testimony of Mr. McNair. *See* Pet. 25–30 (citing Ex. 1005, col. 6, ll. 42–53, col. 11,

ll. 53–55, col. 11, l. 64–col. 12, l. 12, col. 15, ll. 21–25; Ex. 1009, 161–162, 164–166; Ex. 1003 ¶¶ 132–145); *Intel Corp. v. Qualcomm Inc.*, 21 F.4th 784, 797 (Fed. Cir. 2021) (stating that “a rationale [to combine] is not inherently suspect merely because [it is] generic in the sense of having broad applicability or appeal,” and citing *DyStar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006), as “referencing an ‘implicit motivation to combine’ to make a device ‘more desirable, for example because it is . . . faster’”).

Third, Patent Owner cites *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1362–63 (Fed. Cir. 2016), for the proposition that arguments premised on “common sense . . . must still be supported by evidence and a reasoned explanation.” Sur-Reply 6, 11–12. Petitioner, however, does not rely on common sense as a reason to combine with IEEE 1999, but rather on express disclosures in IEEE 1999 of communicating network reconfiguration messages upstream and downstream and on statements in the reference indicating that doing so allows for faster reconfiguration after a link failure. *See* Pet. 25–30, 45–46. Indeed, the words “common sense” do not appear in the Petition or Reply.

Petitioner’s explanation regarding the combination of Gai and IEEE 1999, supported by the testimony of Mr. McNair, which we credit, is persuasive. *See* Pet. 25–30, 40–46; Ex. 1003 ¶¶ 132–145, 165–174. Based on the full trial record, we find that Gai and IEEE 1999 teach limitation [1E] and that a person of ordinary skill in the art would have been motivated to combine their teachings in the manner described by Petitioner.

c) Conclusion

Based on all of the evidence of record, we determine that Petitioner has proven, by a preponderance of the evidence, that claim 3 would have been obvious over Gai and IEEE 1999 under 35 U.S.C. § 103(a).

4. Claims 8, 15, and 20

Claim 8 depends indirectly from independent claim 1. Claims 15 and 20 depend indirectly from independent claim 13, which is similar to claim 1 but recites a network switching device rather than a method. As to independent claim 13, Patent Owner refers to its earlier arguments regarding claim 1, asserting that limitation [13C] is “substantially similar” to limitation [1E]. PO Resp. 14–15. We disagree with Patent Owner’s arguments for the reasons explained above. *See supra* Section II.D.3.b. Patent Owner does not argue dependent claims 8, 15, and 20 separately from their respective parent claims 1 and 13. *See* PO Resp. 19. We have reviewed Petitioner’s contentions and supporting evidence, including the testimony of Mr. McNair, and are persuaded that Petitioner has proven, by a preponderance of the evidence, that claims 8, 15, and 20 would have been obvious based on Gai and IEEE 1999 under 35 U.S.C. § 103(a), for the reasons stated by Petitioner. *See* Pet. 21–46, 48–52, 55–59; Ex. 1003 ¶¶ 113–175, 179–185, 194–203.

5. “Downstream” Claims

a) Claim 12

Claim 12 is independent and recites “sending dummy traffic from the first bridge over the second link *downstream* to the second bridge, so as to cause the second bridge to modify its database responsive to the second port

being in the active state” (limitation [12E]) (emphasis added). In the context of the ’986 patent, the term “downstream” means away from the root switch and the term “upstream” means toward the root switch. *See supra* Section II.B.

Petitioner’s obviousness theory is premised on modifying Gai’s operation to send dummy traffic both upstream and downstream. Pet. 45–46 (addressing “downstream directionality” with respect to claim 1); *see* Tr. 5:6–12 (stating that Petitioner is “essentially assuming” that all of the independent claims require downstream communication). Accordingly, Petitioner refers to its earlier arguments regarding claim 1, arguing that, as modified based on IEEE 1999, a “first bridge” in Gai (i.e., switch 122) would send dummy traffic both upstream and downstream to a “second bridge” (i.e., switch 114). *See* Pet. 29 (explaining that “[b]ridge’ is a general term encompassing both edge and access switches” and, “for the same reasons described [for claim 1], the modified Gai technique involves a ‘bridge’ that sends dummy traffic downstream to another ‘bridge’ as required, for example, by claim 12”), 45–46, 54–55 (referencing Section VII.A.4.f of the Petition); *CRFD Research, Inc. v. Matal*, 876 F.3d 1330, 1345–46 (Fed. Cir. 2017) (agreeing with the petitioner that it “incorporated [an] argument into other grounds of unpatentability . . . by direct citation to [the] argument in the petition”).

Patent Owner argues that the ’986 patent “references and distinguishes Gai” because Gai transmits dummy messages upstream toward the root switch, rather than downstream away from the root switch. PO Resp. 12–13. We agree that the ’986 patent distinguishes Gai on that basis. *See* Ex. 1001, col. 3, ll. 29–33 (“[Gai’s] switch . . . transmits dummy multicast messages toward the root (upstream) through the newly-active

trunk port”), col. 4, ll. 36–43 (“The [’986 patent’s disclosed] bridge is modified . . . so as to manage its redundant downstream ports in a novel way. This modification is *in contrast to* the methods described in [Gai] and the 802.1w draft, which require network entities to be aware of and manage redundancies in their upstream ports, and use special commands to cause them to do so.” (emphasis added)). That difference, however, does not account for Petitioner’s asserted combination. Petitioner does not rely on Gai alone, but rather asserts that a person of ordinary skill in the art would have been motivated—based on the teachings of IEEE 1999—to *modify* Gai’s operation to send dummy traffic downstream as well as upstream.¹² See Pet. 45–46, 54–55. As to the combination with IEEE 1999, Patent Owner refers to its earlier arguments regarding claim 1, which we disagree with for the reasons explained above. See PO Resp. 13–14; *supra*

¹² In the Cisco IPR, we determined that Cisco had not proven the “downstream” claims (claims 11, 12, 23, and 24) to be unpatentable under 35 U.S.C. § 103(a) over Gai alone. IPR2024-00505, Paper 29 at 60–61. Again, Cisco’s arguments in that proceeding were substantially different from those advanced by Petitioner here. For example, Cisco’s first theory of unpatentability (designated “1B”), which we found to be unsupported by the disclosure of Gai, relied on designating added switch 117 as the root switch for the network. *Id.* at 41–48. Cisco’s second theory of unpatentability (designated “1A”), which we also found unpersuasive, relied on Gai’s switch 124 as a “first bridge” and switch 115 or switch 116 as a “second bridge.” *Id.* at 48–50, 58–59. Among other differences, Petitioner in this proceeding instead relies on switch 122 as a “first bridge” and switch 114 as a “second bridge,” and asserts a modification to Gai’s operation based on the teachings of IEEE 1999. Pet. 29, 31–33, 54–55. Given the differences between this proceeding and the Cisco IPR—in particular the introduction of a new secondary reference, IEEE 1999, for limitation [12E]—we conclude that a different outcome for claims 11, 12, 23, and 24 is warranted.

Section II.D.3.b. The parties argue all of the independent claims together in their Reply and Sur-Reply. *See* Reply 1–6; Sur-Reply 1–12.

We have reviewed Petitioner’s contentions and supporting evidence, including the testimony of Mr. McNair, and are persuaded that Petitioner has proven, by a preponderance of the evidence, that claim 12 would have been obvious based on Gai and IEEE 1999 under 35 U.S.C. § 103(a), for the reasons stated by Petitioner. *See* Pet. 21–46, 54–55; Ex. 1003 ¶¶ 113–175, 189–193.

b) Claims 11, 23, and 24

Claim 11 depends from claim 1 and recites that “sending the dummy traffic comprises sending data frames *downstream* from the edge switch toward the access switches” (emphasis added). Claim 23 depends from claim 13 and recites that “the dummy traffic comprises data frames sent *downstream* from the edge switch toward the access switches” (emphasis added). Patent Owner argues that Gai and IEEE 1999 fail to teach sending dummy traffic downstream “from the edge switch toward the access switches” “for the same reasons as discussed” for claim 1. PO Resp. 21. We disagree with Patent Owner’s arguments for the reasons explained above. *See supra* Section II.D.3.b.

Claim 24 is independent and recites a bridge device comprising “first and second *downstream* ports, adapted to be coupled via first and second redundant links to a further bridge device located *downstream* in the network relative to the spanning tree root,” and a protocol processor adapted to “send dummy traffic over the second link *downstream* to the further bridge device, so as to cause the further bridge device to modify its database responsive to the second *downstream* port being in the active state” (emphasis added).

Patent Owner refers to its earlier arguments regarding claim 12, asserting that limitation [24C] is “substantially similar” to limitation [12E].

PO Resp. 15–16. We disagree with Patent Owner’s arguments for the reasons explained above. *See supra* Sections II.D.3.b, II.D.5.a.

We have reviewed Petitioner’s contentions and supporting evidence, including the testimony of Mr. McNair, and are persuaded that Petitioner has proven, by a preponderance of the evidence, that claims 11, 23, and 24 would have been obvious based on Gai and IEEE 1999 under 35 U.S.C. § 103(a), for the reasons stated by Petitioner. *See* Pet. 21–46, 53, 55–57, 60–61; Ex. 1003 ¶¶ 113–175, 188, 194–197, 206–210.

III. CONCLUSION

Petitioner has demonstrated, by a preponderance of the evidence, that claims 3, 8, 11, 12, 15, 20, 23, and 24 of the '986 patent are unpatentable on the basis set forth in the following table.

In summary:

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claim(s) Shown Unpatentable	Claim(s) Not Shown Unpatentable
1–3, 6–15, 18–24	103(a)	Gai, IEEE 1999 ¹³	3, 8, 11, 12, 15, 20, 23, 24	
4, 5, 16, 17	103(a)	Gai, IEEE 1999, Jain ¹⁴		
Overall Outcome			3, 8, 11, 12, 15, 20, 23, 24	

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that claims 3, 8, 11, 12, 15, 20, 23, and 24 of the '986 patent are determined to be unpatentable; 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.

¹³ We do not reach this ground as to claims 1, 2, 6, 7, 9, 10, 13, 14, 18, 19, 21, and 22 because those claims have been cancelled. *See* Ex. 3001.

¹⁴ We do not reach this ground because claims 4, 5, 16, and 17 have been cancelled. *See* Ex. 3001.

IPR2024-00893
Patent 6,917,986 B2

FOR PETITIONER:

Michael N. Rader
Scott A. McKeown
WOLF, GREENFIELD & SACKS, P.C.
mrader-ptab@wolfgreenfield.com
smckeown-ptab@wolfgreenfield.com

FOR PATENT OWNER:

James R. Nuttall
Katherine D. Cappaert
Sadaf Misbah
STEPTOE LLP
jnuttall@steptoe.com
kcappaert@steptoe.com
sadaf.misbah2@gmail.com