

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 JUSTIN T. ARBES, CHARLES J. BOUDREAU, and
PETER D. SHAPIRO, *Administrative Patent Judges*.

SHAPIRO, *Administrative Patent Judge*.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314

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). Patent Owner Portsmouth Network Corp. filed a Preliminary Response (Paper 6, “Prelim. Resp.”) pursuant to 35 U.S.C. § 313. Pursuant to our authorization provided by email, Petitioner filed a Reply (Paper 7, “Reply”) and Patent Owner filed a Sur-Reply (Paper 8, “Sur-Reply”) directed to the issue of discretionary denial under 35 U.S.C. § 314(a).

Pursuant to 35 U.S.C. § 314(a), the Director may not authorize an *inter partes* review unless the information in the petition and any preliminary response “shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” *See* 37 C.F.R. § 42.4(a) (2024) (“The Board institutes the trial on behalf of the Director.”). For the reasons that follow, we institute an *inter partes* review as to claims 1–24 of the ’986 patent on all grounds of unpatentability asserted in the Petition.

B. *Related Matters*

The parties indicate that the ’986 patent is the subject of two district court cases, *Portsmouth Network Corp. v. Juniper Networks, Inc.*, 23-cv-12201-PGL (D. Mass.) (“the Juniper case”), and *Portsmouth Network Corp. v. Cisco Systems, Inc.*, 2:23-cv-00441 (E.D. Tex.) (“the Cisco case”). *See* Pet. 1; Paper 3 (Patent Owner’s Mandatory Notices). Petitioner filed petitions challenging an unrelated patent asserted in the Juniper case. IPR2024-00913, Paper 2; IPR2024-00989, Paper 2. In addition, Cisco Systems, Inc. filed a petition challenging claims 1, 2, 4–7, 9–14, 16–19, and

21–24 of the ’986 patent, and a trial was instituted on August 6, 2024.
IPR2024-00505, Papers 2, 7 (“the ’505 IPR”).

C. The ’986 Patent

The ’986 patent, titled “Fast Failure Protection Using Redundant Network Edge Ports,” issued July 12, 2005, from U.S. Application No. 10/036,518, filed January 7, 2002. Ex. 1001. 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 3:58–63. The goal of the ’986 patent is 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 2:65–3:14, 3:58–69.

More specifically, the background 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

A problem with redundant topologies, however, is that the multiple paths between points in the network can create loops that disrupt the proper functioning. . . . [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.”

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. STP ensures that all data paths in a network

of bridges^[1] are free of loops by disabling forwarding of packets through certain [ports].

Id. at 2:7–36. “The spanning tree algorithm defines one bridge in the network as the root bridge,” where “[e]ach LAN connected to the network has a bridge port that connects it to the root bridge.” *Id.* at 2:48–50. Therefore, “[e]ach bridge has a port defined as its root port, which uniquely connects that bridge to the root bridge [or switch].”² *Id.* at 2:55–56. The other ports on the switch (bridge) that do not connect to the root switch “are put into the blocking state.” *Id.* at 2:56–58. “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 2:58–62. 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 [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.

¹ 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, 1:26–28.

² 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, 2:47–64. The path from a particular switch towards the root switch is called the root path or designated path. *Id.*

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 2:65–3:10.

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, 3:57–59. 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 . . . 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 database so that upstream traffic is immediately directed to the new active link.

Id. at 3:65–4:28.

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



“FIG. 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, 7:53–56. “[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 8:62–64. “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 9:14–18.

“Upon sensing the failure, processor 44 of edge switch 28 blocks its port P3 and activates its port P4 instead.” Ex. 1001, 9:25–27. “Now edge switch 28 and access switch 27 can immediately begin exchanging traffic over link 32.” *Id.* at 9:34–35. “The filtering database maintained by the

switches are no longer accurate [Switch 28] 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 9:35–41. “In order to shorten this delay, edge switch 28 sends dummy frames downstream to access switch 27 via link 32.” *Id.* at 9:42–43. “When switch 27 receives the dummy frame, it also forwards it to switch 26.” *Id.* at 9:56–57. “Once the dummy frame has been sent, switches 26, 27 and 28 update their filtering database to reflect the change in the active link from link 30 to link 32” *Id.* at 9:62–64. 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 4:24–31.

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, 1:31–33. 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 1:33–36. Edge switches are “[c]ore switches that are directly connected to access switches.” *Id.* at 1:36–38.

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, Fig. 1, 7:52–61. 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 1:36–38.

D. Illustrative Claim

Challenged claims 1, 12, 13, and 24 of the '986 patent are independent. Dependent claims 2–11 depend, directly or indirectly, on claim 1. Dependent claims 14–23 depend, directly or indirectly, on claim 13. Claim 1 recites (with letter designations used in the Petition to refer to the various limitations):

1. [p] 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.

E. Evidence

Petitioner relies 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, issued Dec. 11, 2001 (Ex. 1006, “Jain”).

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

F. Prior Art and Asserted Grounds

Petitioner asserts that claims 1–24 would have been unpatentable on the following grounds:

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

II. PROCEDURAL CONSIDERATIONS

A. Discretionary Denial Under 35 U.S.C. § 314(a)

Institution of *inter partes* review is discretionary. *See Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1367 (Fed. Cir. 2016) (“[T]he PTO is permitted, but never compelled, to institute an [*inter partes* review (IPR)] proceeding.”); 35 U.S.C. § 314(a) (2018) (“The Director *may not* authorize an *inter partes* review to be instituted unless the Director determines that the information presented in the petition filed under section 311 and any response filed under section 313 shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” (emphasis added)).

The Board has held that the advanced state of a parallel district court action is a factor that may weigh in favor of denying a petition under § 314(a). *See NHK Spring Co. v. Intri-Plex Techs., Inc.*, IPR2018-00752,

³ The Leahy-Smith America Invents Act (“AIA”), Pub. L. No. 112-29, 125 Stat. 284 (2011), amended 35 U.S.C. § 103, effective March 16, 2013. Because the application from which the ’986 patent issued was filed before this date, the pre-AIA version of § 103 applies.

Paper 8 at 20 (PTAB Sept. 12, 2018) (precedential); Patent Trial and Appeal Board Consolidated Trial Practice Guide (Nov. 2019), 58 & n.2, *available at* <https://www.uspto.gov/TrialPracticeGuideConsolidated>. We consider the following factors to assess “whether efficiency, fairness, and the merits support the exercise of authority to deny institution in view of an earlier trial date in the parallel proceeding”:

1. whether the court granted a stay or evidence exists that one may be granted if a proceeding is instituted;
2. proximity of the court’s trial date to the Board’s projected statutory deadline for a final written decision;
3. investment in the parallel proceeding by the court and the parties;
4. overlap between issues raised in the petition and in the parallel proceeding;
5. whether the petitioner and the defendant in the parallel proceeding are the same party; and
6. other circumstances that impact the Board’s exercise of discretion, including the merits.

Apple Inc. v. Fintiv, Inc., IPR2020-00019, Paper 11 at 5–6 (PTAB Mar. 20, 2020) (precedential) (“*Fintiv*”). In evaluating these factors, we “take[] a holistic view of whether efficiency and integrity of the system are best served by denying or instituting review.” *Id.* at 6.

We also consider “several clarifications” made by the Director of the United States Patent and Trademark Office (“USPTO”). *See* USPTO Memorandum, Interim Procedure for Discretionary Denials in AIA Post-Grant Proceedings with Parallel District Court Litigation, 2 (June 21, 2022) (“Director’s Memo”), *available at* https://www.uspto.gov/sites/default/files/documents/interim_proc_discretionary_denials_aia_parallel_district_court_litigation_memo_20220621_.pdf.

Patent Owner argues that we should exercise our discretion to deny institution based on the Cisco case (which does not involve Petitioner) for *Fintiv* factors 1, 2, and 3 and the Juniper case asserted against Petitioner for factors 4 and 5.⁴ Prelim. Resp. 14–17. According to Patent Owner, the “Cisco trial is expected to take place approximately six months before the projected statutory deadline for a final written decision,” which “weighs against institution.” *Id.* at 14. Petitioner argues that we should decline to use our discretion to deny institution based on the applicability of the *Fintiv* factors in view of both cases. Pet. 66–68; Reply 1–3. Upon consideration of the relevant factors and the clarifications, which are discussed below, we decline to exercise our discretion to deny the Petition.

First, we find no basis for exercising our discretion to deny institution with regard to the Juniper case because, as the parties recognize, that case does not yet have a scheduling order in place, so there is almost no possibility that a trial will occur before our deadline for a final written decision. *See Fintiv* at 5 (“When the patent owner raises an argument for discretionary denial under *NHK* due to an earlier trial date, the Board’s decisions have balanced [*Fintiv*] factors [1-6].” (emphasis added)); Pet. 67 (arguing that “the court has neither entered a schedule nor set a trial date”);

⁴ Patent Owner does not specifically identify the Juniper case in the section of its Preliminary Response arguing for discretionary denial under § 314(a) for *Fintiv* factors 4 and 5, instead referring to a “parallel district court proceeding,” but based on the arguments presented, it is clear that reference is to the Juniper case. Prelim. Resp. 15–16. Patent Owner argues that Petitioner failed to offer a stipulation for factor 4, and that Petitioner “is a defendant in a parallel district court proceeding” for factor 5. *Id.* Patent Owner must be relying on the Juniper case for those factors since neither of those arguments would be accurate with regard to the Cisco case where Petitioner is not a party.

Prelim. Resp. 14 (acknowledging that “the court has not yet entered a schedule or set a trial date for the Juniper case”); Reply 1 (arguing that Patent Owner “effectively concedes that the [Juniper case], in which no trial date has been set, does not raise any *Fintiv* issue”). Accordingly, we analyze *Fintiv* factors 1–6 below in view of the Cisco case.

1. *Whether the Court Granted a Stay or Evidence Exists That One May Be Granted if a Proceeding is Instituted*

According to the parties, the district court has not issued a stay in the Cisco case, so this factor should be neutral. Prelim. Resp. 14; Reply 1; Sur Reply 1. We agree. *See Sand Revolution II LLC v. Cont’l Intermodal Grp. – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (PTAB June 16, 2020) (informative).

2. *Proximity of the Court’s Trial Date to the Board’s Projected Statutory Deadline for a Final Written Decision*

Patent Owner argues that trial in the Cisco case is set for June 2, 2025, approximately six months before our deadline for the final written decision, so this factor favors discretionary denial. Prelim. Resp. 14–15; Sur Reply 1–2. Petitioner does not dispute the timing of the Cisco trial but instead argues that this factor, when viewed collectively with factor 5, weighs against discretionary denial, since the Juniper case has no scheduling order entered yet. Pet. 67–68; Reply 2.⁵

⁵ Petitioner cites the Board’s decision in *Dolby Laboratories, Inc. v. Intertrust Technologies Corp.*, IPR2020-00665, Paper 11 at 15–17 (PTAB Feb. 16, 2021), in which the panel looked at both factors 2 and 5 together when evaluating the impact of two different district court cases on the *Fintiv* analysis. Pet. 67–68. We decline to evaluate the individual *Fintiv* factors in view of both district court cases here given the timing of the Juniper case, as discussed above, and note that the Board’s decision in *Dolby Laboratories* is not precedential.

We agree with Patent Owner that under our precedential decision in *Fintiv*, this factor weighs in favor of discretionary denial in view of the timing of the trial in the Cisco case. *See Fintiv* at 9.

3. *Investment in the Parallel Proceeding by the Court and the Parties*

Patent Owner argues that this factor favors discretionary denial because by the deadline for the institution decision,

Patent Owner will have expended significant resources in the [Cisco case]. Patent Owner and Cisco have already exchanged infringement and invalidity contentions for the '986 patent. Fact discovery has already begun under the district court's schedule and will conclude January 6, 2025, approximately four months after the institution decision is projected to be issued in this case. Furthermore, the claim construction process has already begun in the Cisco case and [the] Joint Claim Construction hearing is scheduled for December 4, 2024, only ten days before the expected date of the institution decision.

Prelim. Resp. 15–16; Sur-Reply 2.

Petitioner responds by pointing out that we already found that factor 3 weighed “strongly against discretionary denial” when analyzing Patent Owner’s discretionary denial arguments in the '505 IPR. Reply 3. Additionally, “[l]itigation of patent validity in the Cisco case will remain at an early stage: fact discovery will still be ongoing, and expert discovery will not even have begun.” *Id.*

The Cisco case is farther along than it was when we analyzed this factor in the '505 IPR, so the result from our analysis there does not apply here. As Patent Owner points out, the parties will have exchanged infringement and invalidity contentions for the '986 patent and the parties will have invested more resources in fact discovery with the looming January 6, 2025 deadline, and all of these investments by the parties are

relevant to this factor. Prelim. Resp. 15; *Fintiv* at 10 (“[The] more work completed by the parties and court in the parallel proceeding tends to support the arguments that the parallel proceeding is more advanced . . . and instituting would lead to duplicative costs.”).

On the other hand, the *Markman* hearing has not yet occurred and no claim construction order has been issued by the district court, fact discovery is still ongoing, and expert discovery has not yet begun. Reply 3; *Fintiv* at 10 (“If, at the time of the institution decision the district court has not issued orders related to the patent at issue in the petition, this fact weighs against exercising discretion to deny institution under *NHK*.”). Additionally, we agree with Petitioner that its diligence in filing also weighs against discretionary denial. Pet. 68. Petitioner filed the Petition on May 3, 2024, about five months before the one-year deadline and prior to receiving Patent Owner’s infringement contentions against Petitioner. *Id.*; see *Fintiv* at 11 (“If the evidence shows that the petitioner filed the petition expeditiously, such as promptly after becoming aware of the claims being asserted, this fact has weighed against exercising the authority to deny institution under *NHK*.”).

Given that some facts weigh in favor of exercising our discretion on this factor, and some weigh against doing so, overall we find this factor to be neutral.

4. *Overlapping Issues with the Parallel Proceedings*

Patent Owner argues that this factor weighs in favor of discretionary denial because Petitioner has not offered a stipulation and “Petitioner does not deny that there are overlapping issues between this IPR and the [Cisco case].” Sur-Reply 3. Petitioner responds that the “IPR’s inclusion of claims not at issue in the Cisco case . . . disfavors discretionary denial” and that

Patent Owner “offers *no evidence* that this IPR’s grounds overlap Cisco’s arguments, which in any event are speculative until trial.” Reply 3.

We agree with Petitioner and therefore find this factor weighs against discretionary denial. Claims 3, 8, 15, and 20 of the ’986 patent are being challenged in this Petition, but are not the subject of the ’505 IPR, which presumably reflects only claims asserted in the Cisco case, and therefore resolution of the validity issues in the Cisco case may not resolve the dispute between Petitioner and Patent Owner here as to those claims. *See Fintiv* at 13. Moreover, Patent Owner has not presented any evidence showing that the same combinations of references being asserted here, Gai and IEEE 1999, and Gai, IEEE 1999, and Jain, are also being considered in the Cisco case, and we note that these combinations of references were not asserted in the ’505 IPR. *See Fintiv* at 12–13 (“[I]f the petition includes materially different grounds, arguments, and/or evidence than those presented in the district court, this fact has tended to weigh against exercising discretion to deny institution under *NHK*.”).

5. *Whether Petitioner and the Defendant in the Parallel Proceeding are the Same Party*

Petitioner is not a defendant in the Cisco case. Pet. 68; Prelim. Resp. 16–17; Sur-Reply 3. Accordingly, we find this factor weighs against discretionary denial. *See Fintiv* at 13–14.

6. *Other Circumstances that Impact the Board’s Exercise of Discretion, Including the Merits*

Regarding factor 6, Petitioner does not make an argument that the merits are compelling. Pet. 68. Patent Owner asserts that “Petitioner’s prior art grounds contain critical deficiencies and do not demonstrate a reasonable

likelihood of—let alone plainly lead to—a conclusion that at least one claim is unpatentable.” Prelim. Resp. 19.

As discussed above, factors 1 and 3 are neutral, factor 2 weighs in favor of discretionary denial, and factors 4 and 5 weigh against discretionary denial. Because we conclude that, on balance, factors 1–5 do not weigh in favor of discretionary denial, it is not necessary for us to make a determination about whether the merits are compelling. *See* Director’s Memo at 3–5; *CommScope Techs. LLC v. Dali Wireless, Inc.*, IPR2022-01242, Paper 23 at 4–5 (Vidal Feb. 27, 2023) (precedential) (“[I]n circumstances where the Board determines that the other *Fintiv* factors 1–5 do not favor discretionary denial, the Board shall decline to discretionarily deny under *Fintiv* without reaching the compelling merits analysis.”). That being said, as discussed below in the Analysis section, we do find that Petitioner has met its burden of demonstrating a reasonable likelihood of showing that at least one of the challenged claims of the ’986 patent is unpatentable.

The parties make no other arguments with respect to factor 6. Therefore, this factor is neutral.

7. *Summary of Fintiv Factors and Conclusion on Section 314(a)*

In view of all six *Fintiv* factors taken together, and taking “a holistic view of whether efficiency and integrity of the system are best served by denying or instituting review,” we conclude that the factors weigh against discretionary denial. *Fintiv* at 6. Only factor 2 favors discretionary denial, while all the other factors are either neutral or weigh against exercising our discretion to deny institution. Therefore, we are not persuaded that the

interests of efficiency and integrity of the system would be best served by invoking our authority under 35 U.S.C. § 314(a) to deny institution.

III. ANALYSIS

A. *Legal Standards*

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) (2006)). 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).⁶

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 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

⁶ Additionally, “[s]uch secondary considerations 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.” *Graham*, 383 U.S. at 17–18. Here, however, Patent Owner has not asserted in its Preliminary Response that any such considerations are relevant in its Preliminary Response, and we have no other evidence of such indicia currently before us.

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). Also, “[t]hough less common, in appropriate circumstances, a patent can be obvious in light of a single prior art reference if it would have been obvious to modify that reference to arrive at the patented invention.” *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1361 (Fed. Cir. 2016).

“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). 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’”); *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.”).

B. Level of Ordinary Skill in the Art

Petitioner argues that at the time of the filing date of the ’986 patent (January 7, 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 (citing Ex. 1003 ¶ 111). Patent Owner does not dispute this assertion or present its own contention regarding the level of ordinary skill in the art. Based on the record presented, including our review of the '986 patent and the types of problems and solutions described in the '986 patent and cited prior art, we agree with Petitioner's proposed definition of the level of ordinary skill for purposes of this Decision, except that we remove the phrase "at least" to avoid unnecessary ambiguity in the level of ordinary skill in the art.⁷

C. Claim Construction

We interpret each challenged claim:

using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim 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 asserts that no claim requires express construction for the grounds presented in the Petition. Pet. 20. Patent Owner applies the ordinary and customary meaning to the '986 patent's claim terms. Prelim. Resp. 3. In view of these positions, we find that no express interpretation of any claim term is necessary to determine whether to

⁷ In the '505 IPR, we preliminarily adopted a similar definition of "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 7 at 22.

institute an *inter partes* review in this proceeding. *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.’” (citation omitted and emphasis added)).

D. Ground 1

Petitioner contends that claims 1–3, 6–15, and 18–24 are unpatentable as obvious over Gai in view of IEEE 1999, citing the testimony of Bruce McNair, M.E., as support. Pet. 21–61; Ex. 1003. Patent Owner filed a Preliminary Response opposing institution. Prelim. Resp. 7–13. We are persuaded that Petitioner has established a reasonable likelihood of prevailing on its asserted ground based on Gai and IEEE 1999 as to all the challenged claims in Ground 1, as discussed below.

1. Gai

Gai is a U.S. patent titled “Method and Apparatus for Rapidly Reconfiguring Computer Networks.” 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.*

Figure 1 of Gai is reproduced below.




FIG. 1

Figure 1 depicts an embodiment of the invention implemented in a computer network 100 with access switches 114–116 and backbone switches 120–125.⁸ Ex. 1005, 7:3–14. “[F]or each access switch 114–116, only one

⁸ Because the ’986 patent claims categorize switches as being access, core, and edge types, it is helpful to apply the definitions for each type from that patent (discussed above) to the switches shown in Figure 1 of Gai. Backbone switches 120–125 in Gai are core switches because they directly connect to routers 130 and 131 (connecting to WANs), or connect switches

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 11:8–13. “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 11:35–41.

To illustrate, Gai provides the following example with reference to switch 114: “[A]ssume that port number three [on switch 114], which is coupled to backbone switch 122, is . . . in the forwarding state.” Ex. 1005, 11:42–44. Because port number three is active, the other ports on switch 114 that connect to the root switch (directly or indirectly) are blocked as explained by the specification:

[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 11:45–52.

The Gai specification then gives an example where back up port number four on switch 114 is chosen to replace port number three, which

120 and 121 that are, in turn, directly connected to the routers. *See* Ex. 1005, Fig. 1, 7:3–52. Backbone switches 122–125 are also edge switches because they are core switches that are directly connected to access switches 114–116. *See id.* Switches 114–116 are access switches because they are directly connected to LANs 102–108 and/or servers 12a and 12c. *See id.* Note that switch 120, in addition to being a core and backbone switch, is also an access switch because it is directly connected to server 112b. *See id.*

has the failed link to switch 122, by transitioning port number four from a blocked state to a forwarding state. Ex. 1005, 12:13–23, 12:32–35. Gai continues explaining that “Switch 114 . . . preferably informs the other devices of its new forwarding port by transmitting dummy multicast packets through the new port.” *Id.* at 12:62–66. 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 13:34–38. Gai continues the example, explaining:

Thereafter, 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 messages[] through all of its forwarding ports so that other devices, as necessary, may learn of the new forwarding ports (i.e., port number four) at access switch 114.

Id. at 13:38–49.⁹

2. IEEE 1999

IEEE 1999 is a published article titled “Automatic Fault Detection and Recovery in Real Time Switched *Ethernet* Networks.” Ex. 1009. According to IEEE 1999, “[t]his paper 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.

⁹ Gai is discussed in the ’986 patent and incorporated by reference into it. *See* Ex. 1001, 3:15–41.

The article describes 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 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. 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 at 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 . . . traffic. 2. Terminate all connections going through the offending link. 3. Rebuild the spanning tree. 4. Reestablish the connections that were torn down.” *Id.*

The article explains that “[t]o 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.” Ex. 1009 at 167.

In describing step 3 of the switch function after detecting a fault, the article states, “[t]he switch[] that detect[s] the failure then initiate[s] a spanning tree reconfiguration, with [itself] as root.” Ex. 1009 at 166. “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 the article, using this method after a link failure will result in the network that “merely sees a small interruption in service.” *Id.*

3. *Analysis of Claim 1*

Petitioner asserts that Gai, in view of IEEE 1999, teaches or suggests all of the limitations of claim 1. Pet. 30–46.

With respect to the preamble, Petitioner asserts Gai teaches the preamble by disclosing its method for rapidly reconfiguring a computer network. Pet. 30. Patent Owner does not dispute this element. Prelim. Resp. 7–13. The preamble here, “[a] method for network communication, comprising,” is not limiting because it is simply stating an intended use. *See Arctic Cat Inc. v. GEP Power Prod., Inc.*, 919 F.3d 1320, 1328 (Fed. Cir. 2019) (“We have long ruled that a preamble is not limiting where a patentee defines a structurally complete invention in the claim body and uses the preamble only to state a purpose or intended use for the invention.”) (citations omitted) (internal quotation marks omitted). In any event, as discussed further below, on this record Gai teaches a method for network communication. *See* Ex. 1005, code (57).

a) Limitation 1a - 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;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation.

Gai discloses edge switches connected to access switches. For example, switch 122 is an edge switch because it is a core switch that is directly connected to access switch 114, and also connects switch 114 to a Wide Area Network (WAN) through router 130. Pet. 31–32; Ex. 1003 ¶¶ 148–150; Ex. 1005, Fig. 1, 6:10–12, 7:26–30. Switch 114 is an access switch because it connects to LANs 102–103. Ex. 1005, 7:4–7 (“The network 100 preferably comprises a plurality of local area networks (LANs) 102–109 each of which is coupled to one or more access switches 114–116.”).

Figure 1 of Gai shows edge switch 122 is coupled to access switch 114 through multiple redundant links, all labeled 128, which are connected through ports 2–4 on switch 114. Ex. 1005, Fig. 1, 11:41–52. Ports 2 and 3 on switch 114 are both directly coupled to switch 122, while port 4 is indirectly coupled to switch 122, first by a link to switch 124 from port 4, then switch 124 is in turn linked to switch 121, which is linked to switch 120, which is linked to switch 122. *Id.* at Fig. 1, 11:36–52. While the first and second ports on edge switch 122 are not labeled in Figure 1, they are clearly there as the figure shows edge switch 122 connected to the three redundant links 128 (two direct and one indirect) going to switch 114. *Id.*

Patent Owner does not dispute Gai teaches this limitation.

b) Limitation 1b - placing the first port in an active state, while placing the second port in a blocking state;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. *See* Pet. 34; Ex. 1003 ¶¶ 153–154. Petitioner contends that Gai teaches placing the first port in an active state (i.e., port 3 on switch 114), and placing the second port in a blocking state (i.e., ports 2 and 4 on switch 114). *Id.* (citing Ex. 1005, 11:42–48).

Gai discloses a preferred embodiment which describes that port 3 on access switch 114 is initially in the active (forwarding) state, connecting directly to edge switch 122. *Id.* at Fig. 1, 11:36–52 (“With reference to switch 114, for example, assume that port number three, which is coupled to backbone switch 122, is . . . in the forwarding state.”). Since port 3 is in the active state, ports 2 and 4 of switch 114, which are redundant links (directly and indirectly, respectively) to edge switch 122, are initially in an inactive (blocked) state. *Id.* (“[P]ort numbers two and four, which also connect switch 114 to backbone switches (i.e., backbone switches 122 and 124), are in the blocked state.”). Therefore, of the two ports on edge switch 122 that are directly connected to access switch 114, one will initially be in the active state (the port directly connected to port 3 of switch 114), and the other will initially be in the inactive state (the one directly connected to port 2 of switch 114).

Patent Owner does not dispute that Gai teaches this limitation. Additionally, the ’986 patent describes Gai as disclosing 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, 3:24–26. As

described by Gai, a “trunk port” is a port that has a link connecting it to another switch. Ex. 1005, 1:43–44 (“Ports used to couple switches to each other are generally referred to as trunk ports . . .”).

c) Limitation 1c - 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;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation by describing that in the preferred embodiment traffic initially flows from port 3 in switch 114 over link 128 to switch 122. *See* Pet. 35–36; Ex. 1005, 7:30–33; Ex. 1003 ¶¶ 156–157. Gai makes clear that when ports are in the forwarding state, data frames, or communication traffic, is sent over the port. Ex. 1005, 2:56–63. 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 11:16–33; Ex. 1003 ¶ 158. “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 1:50–54; Ex. 1003 ¶ 159.

Petitioner’s annotated Figure 2 of Gai, depicting a block diagram of a switch with a database on it, is reproduced below:




FIG. 2

Annotated Figure 2 of Gai

Pet. 35. The annotated version of Figure 2 above shows that each switch will have, among other things, a filtering database 244 (shown on the bottom right inside a purple box). Pet. 36; Ex. 1003 ¶¶ 158–159; Ex. 1005, 11:16–23. The “filtering database 244 [is] configured to store address information corresponding to the entities of network 100. Specifically, filtering database 244 has a plurality of columns or cells, including a destination address cell [DA], a destination port [PORT], and a corresponding timer cell. Each row or entry in the filtering database 244 preferably corresponds to a particular entity.” Ex. 1005, 8:32–39.

Referring back to the preferred embodiment in 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. Ex. 1005, 11:16–33. Once this entry is added, when switch 122 receives a data frame with the destination address for switch 114, it will see 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.*

Patent Owner does not dispute that Gai teaches this limitation.

d) Limitation 1d - 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

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. See Pet. 37–40; Ex. 1003 ¶¶ 160–164. Petitioner contends Gai teaches that in response to “a failure of the active link” (i.e., link 128 from port 3 of switch 114), “access switch 114 places its port 2 in the forwarding state and its port 3 in the blocking state.” Pet. 39 (citing Ex. 1005, 11:48–52, 14:40–51).

Gai teaches that, in the preferred embodiment shown in Figure 1, once the link 128 coming from port 3 of switch 114 fails, port 3 will be placed in a blocked state, and one of the other root ports, either port 2 or port 4, will be placed in the active (forwarding state):

If the link 128 coupled to the port number three at switch 114 fails, either one of these two other ports (i.e., port numbers two

or four) will immediately transition to the forwarding state and begin receiving and sending messages.

Ex. 1005, 11:48–52; Ex. 1003 ¶ 163.

Patent Owner does not dispute that Gai teaches this limitation.

e) Limitation 1e - 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.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai, in view of IEEE 1999, suggests this limitation. See Pet. 40–46; Ex. 1003 ¶¶ 165–175.

Petitioner’s Argument for Limitation 1e

According to Petitioner, the “’986 patent admits that Gai discloses sending dummy multicast messages (i.e., dummy traffic) upstream using a root port after detection of a link failure.” *Id.* at 40 (citing Ex. 1001, 3:26–35); Ex. 1003 ¶ 165. Petitioner also contends that Gai teaches or suggests that its rapid reconfiguration can be implemented on any switch, as shown in Figure 2 (rapid reconfiguration entity 234) and the accompanying description. Pet. 43 (citing Ex. 1005, 8:12–14, 8:66–9:3, 12:59–66); Ex. 1003 ¶ 169. 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.

Pet. 44 (citing Ex. 1005, 13:14–23). “[W]hen other devices within the network receive dummy multicast messages, they enter this new information in their respective filtering databases.” *Id.* (citing Ex. 1005, 13:34–38).

Furthermore, Petitioner argues, a person of ordinary skill in the art (“POSITA”) would understand that IEEE 1999 teaches that it is advantageous to send “self-healing” messages (i.e., dummy traffic) both upstream and downstream when recovering from a link failure. Pet. 46; Ex. 1003 ¶ 135. “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.’” Pet. 25. Petitioner asserts that IEEE 1999’s teardown message is dummy traffic that causes the switches receiving it to update their filtering databases such that the failed connection is terminated. *Id.* at 26–27. According to Petitioner, a POSITA “would have had reason to modify Gai’s technique” in view of IEEE 1999 because it relates to “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.* “[W]hen modified in view of IEEE 1999’s disclosure that it was advantageous to send self-healing messages both upstream and downstream, Gai’s backbone switches send Gai’s dummy traffic both upstream and downstream, regardless of whether a link failure is detected upstream or downstream of the switch.” *Id.* at 46; Ex. 1003 ¶¶ 174–175.

Patent Owner’s Argument for Limitation 1e

Patent Owner does not dispute that Gai discloses sending dummy traffic upstream after a link failure, but argues that the portion of IEEE 1999

relied on by Petitioner does not teach or suggest sending dummy traffic as used in the claims. Prelim. Resp. 7–13. According to Patent Owner, “*IEEE 1999*’s teardown message is not dummy traffic used ‘to cause each of the one or more access switches to modify its respective database responsive to the second port being the active state’ as required by the claims.” Prelim. Resp. 8.

Patent Owner argues that “*IEEE 1999* discloses that its fault recovery mechanism performs three main steps: (1) ‘[t]erminating connections through the faulty link/switch,’ (2) ‘[r]econfiguring the spanning tree,’ and (3) ‘[r]e-establishing the connections that were torn down.’” Prelim. Resp. 8 (citing Ex. 1009 at 167). “As part of the first step, *IEEE 1999*” sends the teardown messages to “upstream/downstream sources of outgoing connections.” *Id.* at 8–9. “[T]he teardown message is used . . . ‘[t]o terminate the connections going through a failed link.’” *Id.* “The teardown message plays no part in the reconfiguration of [*IEEE 1999*’s] spanning tree.” *Id.* Therefore, “*IEEE 1999* clearly teaches that sending the teardown message is used as a preparatory step and that it [is] not used in the actual reconfiguration [process].” *Id.* at 10.

Furthermore, Patent Owner asserts that “[e]ven assuming that *IEEE 1999*’s teardown message is equivalent to the claimed dummy traffic, Petitioner makes no showing that the teardown message is sent 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’ as required by claim limitation 1e.” Prelim. Resp. 11. According to Patent Owner, “*IEEE 1999* does not disclose that the respective databases are modified ‘responsive to the second port being in the active state,’ and Petitioner does not address this point.” *Id.*

Finally, Patent Owner argues that “a POSITA would not be motivated to combine *IEEE 1999*’s teardown message [with] Gai’s switch system.” Prelim. Resp. 11. According to Patent Owner, the “teardown message is sent specifically to tear the connection to the upstream/downstream destinations . . . [and] plays no part in the reconfiguration of *IEEE 1999*’s spanning tree.” *Id.* “As a result, *IEEE 1999*’s teardown message serves a completely different purpose than [] Gai’s dummy multicast messages.” *Id.*

Gai in View of IEEE 1999 Suggests Limitation 1e

Based on the current record, we are persuaded that Gai in view of IEEE 1999 would at least suggest to a POSITA dummy traffic being sent both upstream and downstream. Pet. 26–30, 40–46; Ex. 1003 ¶¶ 165–175; Ex. 1005, 5:35–6:12, 6:39–54, 8:12–14, 8:66–9:3, 11:53–55, 11:64–12:12, 12:59–13:4 13:9–23, 13:34–38, 15:21–25, Fig. 1; Ex. 1009, 161–162, 164–167. This includes sending dummy traffic from an edge switch to an access switch as part of the rapid reconfiguration process disclosed in Gai.

As Petitioner points out, Gai describes that the rapid reconfiguration process can be implemented on both edge and access switches.¹⁰ Pet. 43 (citing Ex. 1005, 8:66–9:3, describing that rapid reconfiguration entity 234 is also implemented on backbone switch 120, which indicates it can be applied to any switch). Gai discusses how to implement fast network recovery

¹⁰ Gai uses the broader term “backbone switch” instead of “edge switch,” describing that “[a]ccess switches are preferably coupled to entities (e.g., LANs, end stations, etc.) whereas backbone switches provide the interconnection between access switches.” Ex. 1005, 6:10–12. So the edge switches in Figure 1, such as switches 122–125, would also be backbone switches because they interconnect access switches 114–116.

generally and says it can be accomplished using commands on both the access switches and edge switches:

[T]he invention relates to a method and apparatus for rapidly reconfiguring a computer network. The network preferably includes a plurality of devices executing the spanning tree algorithm so as to elect a root and place the ports of the devices in either a forwarding or blocked state. In accordance with the method one or more devices are configured and arranged so that one trunk port is in the forwarding state and other trunk ports are in the blocked state. Additionally, one or more of the blocked ports are designated as back-up ports. Upon detection of a failure at the active forwarding port, the state of one of the back-up ports immediately transitions from blocked to forwarding, thereby becoming the new active port for the device. . . . Upon transition to the new forwarding port, the device begins transmitting “dummy” multicast messages through the new port. . . . Upon receipt, the other devices examine the contents of these messages and note the port on which they were received, which may differ from the port on which messages from these entities were previously received (i.e. before the failure and subsequent replacement of the device’s active port) [T]he method and apparatus manifests, in part, as a series of novel commands that may be entered at the devices. *The devices, moreover, may be classified as either access switches or backbone switches.* Access switches are preferably coupled to entities (e.g., LANS, end stations, etc.) whereas backbone switches provide the interconnections between access switches.

Ex. 1005, 5:35–6:12 (emphasis added).

When it is 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. Pet. 45; Ex. 1003 ¶ 173. For example, as Petitioner points out, 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 the dummy traffic would be sent from

switch 122 downstream to access switch 114 in order to reconfigure it.
Pet. 27–28; Ex. 1003 ¶ 139.

A similar mechanism is described in IEEE 1999 where, as part of the reconfiguration process, a teardown message is sent both upstream and downstream from the switch where the failure is detected. Ex. 1009 at 166. According to IEEE 1999, “each switch that detects the fault” takes steps to “[r]ebuild the spanning tree.” *Id.* As part of the spanning tree reconfiguration that takes place after a failure, each switch that detects the failure sends teardown messages to both upstream and downstream switches to “terminate the connection[] through a failed link.” Ex. 1009 at 167. By sending teardown messages in both directions, along with optimizations in rebuilding the spanning tree protocol, the network downtime in response to a link failure is substantially reduced. *Id.* at 164–167. We are persuaded, based on the current record, that a POSITA would apply the teaching from IEEE 1999 that sending messages for reconfiguring the network after a link failure both upstream and downstream can reduce the network downtime to the rapid reconfiguration process in Gai when that process is implemented on an edge switch, such as on switch 122. *See* Pet. 26–30, 40–46; Ex. 1003 ¶¶ 165–175.

We do not understand Petitioner to be arguing that the teardown message itself should be taken from IEEE 1999 and combined with Gai’s rapid reconfiguration process, as Patent Owner seems to contend. Prelim. Resp. 7–12. Rather, Petitioner asserts that in view of IEEE 1999’s teaching that the teardown message is sent both upstream and downstream in order to more rapidly recover from a link failure, a POSITA would understand that the same principle could be applied to the dummy traffic being sent out during the rapid recovery process in Gai. Pet. 26–30. So it is the dummy

traffic described in Gai, rather than the teardown message in IEEE 1999, that Petitioner alleges meets limitation 1e. *Id.*; Ex. 1003 ¶¶ 54, 134–144.

Therefore, at this stage, we find Patent Owner’s arguments trying to distinguish the teardown message from dummy traffic in limitation 1e to be unpersuasive.¹¹

We also are persuaded, based on the current record, that a POSITA would have been motivated to apply the teachings from IEEE 1999 to Gai. Pet. 26–27. Both references pertain to the same subject matter, more rapidly updating the network reconfiguration after a link failure on a network that implements an IEEE 801.1d spanning tree algorithm. Pet. 25; Ex. 1003 ¶¶ 100, 132–134; Ex. 1009 at 161–162; Ex. 1005, 6:42–53, 15:21–25. Additionally, both references use a similar approach to reconfiguring the switches by reconfiguring from failed ports to backup ports and sending messages to neighboring switches so they update their configurations. Ex. 1009 at 166–167; Ex. 1005, 6:42–53, 15:21–25. Therefore, we are persuaded on this record that a POSITA would look to the teachings of IEEE 1999 that lead to time savings when considering how to implement a version of the rapid reconfiguration method described in Gai. *See* Pet. 26–30, 45–46; *Intel Corp. v. PACT XPP Schweiz AG*, 61 F.4th 1373, 1380–81 (Fed.

¹¹ We note that Petitioner asserts that the teardown messages in IEEE 1999 have the same function as Gai’s dummy messages. Pet. 27; Ex. 1003 ¶¶ 134–135, 138. As discussed above, Patent Owner disputes this, pointing out that the teardown messages are part of an earlier step in the IEEE 1999 process that takes place before the filtering databases are updated to reflect the new connections, which is the function of the dummy messages in Gai. Prelim. Resp. 7–12. At this stage, we do not need to resolve this dispute because, as noted above, it is Gai’s dummy message being sent downstream that Petitioner relies on for meeting limitation 1e, not IEEE 1999’s teardown message. Pet. 26–30; Ex. 1003 ¶¶ 54, 134–144.

Cir. 2023) (“That [two prior art references] address the same problem and that [one of the references] was a known way to address that problem is precisely the reason that there’s a motivation to combine under *KSR* and our precedent.”).

Patent Owner’s argument for why a POSITA would not have been motivated to combine the references is that the teardown message from IEEE 1999 is “solely used to terminate faulty connections and is not used in the spanning tree reconfiguration process.” Prelim. Resp. 10. This seems to be based on the same argument Patent Owner makes for why the combination does not teach limitation 1e, that the teardown message does not teach the claimed dummy message. But that is not the motivation relied on by Petitioner; rather, it is that IEEE 1999’s disclosure teaches a POSITA that it would be beneficial for a rapid reconfiguration process (like that of Gai) to send dummy messages both upstream and downstream. Pet. 26–27. Therefore, at this stage we find Patent Owner’s argument unpersuasive.

For the reasons discussed above, we find that Petitioner has demonstrated a reasonable likelihood that claim 1 of the ’986 patent would have been obvious over Gai in view of IEEE 1999.

4. *Analysis of Independent Claim 12*

Independent claim 12 is similar to independent claim 1, with several of the corresponding limitations below having the same wording. At this stage of the proceeding, limitation 12e, corresponding to limitation 1e, is the only limitation of claim 12 that Patent Owner contends the combination of Gai and IEEE 1999 fails to teach or suggest.

Claim 12’s preamble, “[a] method for network communication, comprising:” is the same as claim 1’s preamble, and is similarly non-limiting

because it is simply stating an intended use. *See Arctic Cat Inc.*, 919 F.3d at 1328.

a) Limitation 12a – 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;

Based on the current record, we are persuaded that Petitioner has sufficiently shown that Gai teaches this limitation. *See Pet. 54; Ex. 1003* ¶ 189. Petitioner contends that Gai teaches a “first bridge” (i.e., switch 122) and a “second bridge” (i.e., switch 114). *See Pet. 54.*

The only substantive difference from claim 1 is that instead of requiring an “edge switch” connected by redundant links to an “access switch,” claim 12 uses the broader terms “first bridge” connected by redundant links to a “second bridge.” *Ex. 1001, claim 12.* A switch is a type of bridge as discussed above in footnote 1. Therefore, our discussion above describing why we are persuaded on this record that Gai teaches limitation 1a applies here to explain why we also are persuaded on this record that Gai teaches this limitation.

Patent Owner does not dispute that Gai teaches limitation 12a.

b) Limitation 12b - placing the first port in an active state, while placing the second port in a blocking state;

Based on the current record, we are persuaded that Petitioner has sufficiently shown that Gai teaches this limitation. *See Pet. 54; Ex. 1003* ¶ 190. This limitation uses the same words as limitation 1b and therefore

our finding above for why on this record Gai teaches that limitation also applies here. Patent Owner does not dispute that Gai teaches this limitation.

c) Limitation 12c – 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 the forwarding the traffic;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation for the same reasons presented for limitation 1c. *See* Pet. 54; Ex. 1003 ¶ 191; Ex. 1005, 11:41–44. The only difference from limitation 1c is that the term “bridge” is used instead of “switch,” and since a switch is a type of bridge, our findings above for why on this record Gai teaches limitation 1c also apply here. Patent Owner does not dispute that Gai teaches this limitation.

d) Limitation 12d – 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

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. *See* Pet. 54; Ex. 1003 ¶ 192. This limitation uses the same words as limitation 1d, and therefore our findings above for why on this record Gai teaches limitation 1d also apply here. Patent Owner does not dispute that Gai teaches this limitation.

e) Limitation 12e – 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.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai and IEEE 1999 suggest this limitation. *See* Pet. 55; Ex. 1003 ¶ 193.

This limitation is similar to limitation 1e, except that it uses the terms “first bridge” and “second bridge” instead of “edge switch” and “access switch” and it adds the requirement that the dummy traffic is sent “downstream” over the second link. Since a switch is a type of bridge (*see supra* fn. 1), our finding above with respect to limitation 1e accounts for that difference. Also, as discussed above with respect to limitation 1e, we find, based on the current record, that Gai’s rapid reconfiguration process in view of IEEE 1999’s description of the benefit of sending teardown messages both upstream and downstream would at least suggest sending dummy traffic downstream. *See* Pet. 55; Ex. 1003 ¶ 193.

For the reasons discussed above, at this stage we find that Petitioner has demonstrated a reasonable likelihood that claim 12 of the ’986 patent would have been obvious over Gai in view of IEEE 1999.

5. Analysis of Independent Claim 13

Independent claim 13 is similar to independent claim 1, and limitation 13c, corresponding to limitation 1e, is the only limitation that Patent Owner contends the combination of Gai and IEEE 1999 fails to teach or suggest.

Based on the current record, claim 13’s preamble, “[a] network switching device, comprising:” is taught by Gai’s disclosure of switches that

can implement the invention, e.g., switch 122 as shown in Figure 1. Pet. 55; Ex. 1003 ¶ 194; Ex. 1005, Fig. 1. Patent Owner does not dispute that Gai teaches this preamble.

a) Limitation 13a – a network port, adapted to be coupled over a network trunk link to a core switch in a network so as to exchange communication traffic with the core switch;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. See Pet. 55; Ex. 1003 ¶ 194. Petitioner contends Gai teaches this limitation in the network shown in Figure 1, citing link 128 between switch 122 and switch 120. Pet. 55. Switch 120 is a core switch because it is connected to a (WAN) router 131 and the link 128 is a trunk line because it is between two switches. Ex. 1001, 1:32–35. Patent Owner does not dispute that Gai teaches this limitation.

b) Limitation 13b – first and second access ports, adapted to be coupled via first and second redundant network access links to one or more access switches serving network users; and

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. See Pet. 56; Ex. 1003 ¶ 195. This is similar to limitation 1a, which requires first and second redundant ports coupled to an access switch, and our analysis for that limitation applies here. Patent Owner does not dispute Gai teaches this limitation.

c) Limitation 13c – a protocol processor, adapted to place the first access port in an active state, while placing the second access port in a blocking state; so that the communication traffic is conveyed over the first link while the first port is in the active state, thereby causing each of the access switches to build a respective access switch database for use in forwarding the traffic; the protocol processor being further adapted, responsive to a failure associated with the first link, to place the second access port in the active state and the first port in the blocking state and to send dummy traffic through the second access port over the second link to the one or more access switch switches, so as to cause each of the one or more access switches to modify its respective access switch database responsive to the second access port being in the active state.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai, in view of IEEE 1999, teaches this limitation. *See* Pet. 56–57; Ex. 1003 ¶¶ 196–197. This is similar to limitations 1b–1e with the addition of a “protocol processor” placing the ports into active and blocking states. Petitioner points out that Gai teaches a protocol processor called a protocol entity 232 that includes rapid reconfiguration entity 234, which is software on each switch that will change active ports when it detects a failure. Pet. 56–57 (citing Ex. 1005, 8:12, 8:40–51).

Patent Owner does not dispute that Gai discloses a protocol processor, but instead disputes that Gai, in view of IEEE 1999, teaches or suggests this limitation for the same reasons it relies on with regard to limitation 1e. Prelim. Resp. 12–13. For the same reasons discussed above for limitation

le, based on the current record we do not find Patent Owner's argument persuasive.

For the reasons discussed above, we find that Petitioner has demonstrated a reasonable likelihood that claim 13 of the '986 patent would have been obvious over Gai in view of IEEE 1999.

6. *Analysis of Independent Claim 24*

Independent claim 24 is similar to independent claim 13 but is an apparatus claim instead of a method claim and explicitly mentions the root switch (which was implicitly referenced in claim 13 by requiring that dummy traffic be sent downstream over the second link, which is measured with respect to the root switch). Limitation 24c, corresponding to limitation 13c, is the only limitation that Patent Owner contends Gai and IEEE 1999 fail to teach or suggest. Prelim. Resp. 13.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches claim 24's preamble, "[a] bridge device for use in a communication network having a spanning tree root, the device comprising:" because Gai discloses a communication network having a spanning root, including in Figure 1, where switch 120 is designated as the root for STP purposes. *See* Pet. 60; Ex. 1003 ¶ 207; Ex. 1005, 8:57–59, Fig. 1. Patent Owner does not dispute that Gai teaches the preamble to claim 24.

a) Limitation 24a – a designated port, adapted to be coupled over an upstream link in the network so as to communicate with the root;

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. *See* Pet. 60; Ex. 1003 ¶ 207. The designated port here is similar to the network port in limitation

13a, so our analysis is the same as for that limitation. Patent Owner does not dispute that Gai teaches this limitation.

b) Limitation 24b – 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

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this limitation. *See* Pet. 61; Ex. 1003 ¶ 208. The first and second downstream ports are similar to the first and second access ports in limitation 13b, so our analysis is the same as for that limitation. Patent Owner does not dispute that Gai teaches this limitation.

c) Limitation 24c – a protocol processor, which is adapted to place the first downstream port in an active state, while placing the second downstream port in a blocking state, so that communication traffic is conveyed over the first downstream link while the first port is in the active state, thereby causing the bridge device to build respective databases for use in forwarding the traffic, the protocol processor being further adapted responsive to a failure associated with the first link, to place the second downstream port in the active state and the 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.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 suggests this limitation.

See Pet. 61; Ex. 1003 ¶¶ 209–210. This limitation is nearly identical to limitation 13c, so our analysis is the same as for that limitation. Patent Owner disputes that Gai in view of IEEE 1999 teaches or suggests this limitation for the same reasons as set forth for limitation 1e and, therefore, based on the current record, we do not find that argument persuasive for the same reasons as discussed above for limitation 1e. *See* Prelim. Resp. 12–13.

For the reasons discussed above, we find that Petitioner has demonstrated a reasonable likelihood that claim 24 of the '986 patent would have been obvious over Gai in view of IEEE 1999.

7. *Analysis of Dependent Claim 2*

Dependent claim 2 depends on independent claim 1 and adds “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.” Ex. 1001, claim 2. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches this additional limitation. *See* Pet. 46–47; Ex. 1003 ¶¶ 176–177. Petitioner contends that “Gai discloses switches that execute a spanning tree algorithm, which calculates a spanning tree that is loop-free.” Pet. 46 (citing Ex. 1005, 2:30–41, 7:59–61).

The network disclosed and discussed in Gai is one that implements the STP to eliminate redundant loops, so that would eliminate all loops between the edge switch and one or more access switches as required by this claim. Pet. 46–47; Ex. 1005, 7:59–61 (“To avoid the creation of loops, these switches execute the spanning tree algorithm.”). Patent Owner does not dispute that Gai teaches the additional limitation in claim 2.

8. *Analysis of Dependent Claim 3*

Dependent claim 3 depends on independent claim 1 and adds “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.”

Based on the current record, we are persuaded that Gai teaches, or at least suggests, this additional limitation. Pet. 47–48; Ex. 1003 ¶ 178. As described by Petitioner, Gai accomplishes reducing the radius of a spanning tree network 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, 48; Ex. 1003 ¶¶ 153–155, 178; Ex. 1001, 3:67–4:4, 4:35–37, 4:44–48; Ex. 1005, 5:36–43, 11:42–48.

Patent Owner does not dispute that Gai teaches the additional limitation in claim 3.

9. *Analysis of Dependent Claim 6*

Dependent claim 6 depends on independent claim 1 and adds “wherein sending the dummy traffic comprises sending data frames originating at the edge switch while having source addresses corresponding to network address of other entities in the network that are accessible to the access switches via the edge switch.” Ex. 1001, claim 6. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches the additional limitation in dependent claim 6. *See* Pet. 48; Ex. 1003 ¶ 79.

Gai discloses “[u]pon transition to the new forwarding port, the device begins transmitting ‘dummy’ multicast messages through the new port. These dummy multicast messages carry the source address of each entity

that is directly coupled to the device with the new active port or downstream thereof (relative to the root) and are received by other devices in the computer network.” Ex. 1005, 5:54–60. So dummy messages “carry the source address” of other entities in the network. Patent Owner does not dispute that Gai teaches the additional limitation in dependent claim 6.

10. Analysis of Dependent Claim 7

Dependent claim 7 depends on claim 6 and adds “wherein sending the data frames comprises copying the source addresses for the data frames from the database of the edge switch.” Ex. 1001, claim 7. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai suggests dependent claim 7. *See* Pet. 49; Ex. 1003 ¶ 180. Petitioner contends Gai’s discussion of how the filtering databases get updated on the switch teaches the additional limitation in claim 7. Pet. 49.

Gai’s “filtering database 244 [is] configured to store address information corresponding to the entities of network 100.” Ex. 1005, 8:31–39. Gai also describes loading the address of the corresponding entities into the source address field of the dummy messages. Ex. 1005, 13:14–18 (“[R]apid reconfiguration entity 234 (FIG. 2) loads SA field 412 with the address of an entity directly coupled to switch 114.”), 13:50–55 (“Rapid reconfiguration entity 234 similarly generates one or more separate multicast message 400 for each remaining entity directly coupled to access switch 114 In each message 400 entity 234 loads SA field 414 with the address of the corresponding entity.”). On this record we are persuaded that, in view of these teachings by Gai, a POSITA would have found it obvious for the source address for the data frame for the devices to be copied from the filtering database; indeed, there does not seem to be any other readily apparent place from which it would be obtained to create the dummy

messages. *See* Ex. 1003 ¶ 180. Patent Owner does not dispute that Gai suggests the additional limitation in claim 7.

11. Analysis of Dependent Claim 8

Dependent claim 8 depends on claim 7 and adds “wherein the edge switch has three or more ports, including the first and second ports, and wherein prior to the failure, the database comprises entries associated with a plurality of the ports, and wherein copying the source addresses comprises using the network addresses in all the entries that are not associated with the first port as the source addresses of the data frames.” Ex. 1001, claim 8.

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 suggests dependent claim 8. *See* Pet. 49–52; Ex. 1003 ¶¶ 182–186. As described by Petitioner, Figure 1 of Gai shows that switch 122 has at least three ports. Pet. 49; Ex. 1003 ¶ 181. Gai also describes that switch 122 has a filtering database that includes columns that store an address, a destination port, and a timer. Pet. 50; Ex. 1003 ¶ 182; Ex. 1005, 1:53–154. Gai also describes using the network addresses in the filtering database that were generated by entities coupled to the access switch as the source address for the dummy messages. Pet. 52; Ex. 1005, 13:1–4, 13:14–33, 14:55–58. Patent Owner does not dispute that Gai in view of IEEE 1999 suggests the additional limitation in claim 8.

12. Analysis of Dependent Claim 9

Dependent claim 9 depends on claim 6 and adds “wherein sending the data frames comprises incorporating a Virtual Local Area Network (VLAN) tag in a header of the data frames.” Ex. 1001, claim 9. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 teaches, or at least suggests, dependent claim 9.

See Pet. 52–53; Ex. 1003 ¶ 186. Petitioner contends that Gai’s disclosure of

how VLAN tags are incorporated into dummy frames discloses the additional limitation in claim 9. Pet. 52–53.

Gai discloses that “in formulating dummy multicast messages following the detection of a change, access switches 114–116 may append a VLAN tag 416 to message 400.” Ex. 1005, 17:32–39. “[A]n objective of [the] VLAN arrangement is to allow all ports and entities of the network having the same VLAN designation to interchange messages by associating a VLAN designation with each message. Thus, those entities having the same VLAN designation function as if they are all part of the same LAN.” *Id.* at 16:17–22; *see also id.* at Fig. 4, 15:52–67, 17:8–56; Ex. 1003 ¶ 186. Given this disclosure, we find at this stage that Gai in view of IEEE 1999 teaches, or at least suggests, that the dummy messages originating at the edge switch would also incorporate a VLAN tag into the message. Patent Owner does not dispute Gai teaches this additional limitation in claim 9.

13. Analysis of Dependent Claim 10

Dependent claim 10 depends on claim 6 and adds “wherein sending the data frames comprises multicasting the data frames to all of the one or more access switches.” Ex. 1001, claim 10. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches the additional limitation in dependent claim 10. *See* Pet. 53. Petitioner contends Gai’s discussion of how the dummy multicast messages are transmitted discloses this limitation. Pet. 53; Ex. 1003 ¶ 187.

Gai teaches that sending the dummy message include multicasting the packets so the message will reach all the switches, including any access switches, causing them to update their filtering databases to reflect the change. Ex. 1005, 13:9–13 (“Rapid reconfiguration entity 234 preferably loads DA field 410 with a multicast address that causes message 400 to be

received by all devices (e.g., access switches 114–116 and backbone switches 120–125) within the network 100.”). Patent Owner does not dispute that Gai teaches the additional limitation in dependent claim 10.

14. Analysis of Dependent Claim 11

Dependent claim 11 depends on claim 1 and adds “wherein sending the dummy traffic comprises sending data frames downstream from the edge switch towards the access switches.” Ex. 1001, claim 11. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 teaches, or at least suggests the additional limitation in dependent claim 11. *See* Pet. 53; Ex. 1003 ¶¶ 104, 188. Independent claim 12 also requires that the dummy traffic be sent downstream, so our findings here are for the same reasons we find on this record Gai and IEEE 1999 suggest limitation 12e above.

15. Analysis of Dependent Claim 19

Claim 19 depends on dependent claim 18 and adds: “comprising a memory, wherein the protocol processor is adapted to build an edge switch database in the memory responsive to the communication traffic prior to the failure associated with the first link, and, to copy the source address for the data frames from the edge switch database.” Ex. 1001, claim 19. The protocol processor referred to in this claim is also recited in claim 13, so our analysis for limitation 13c also applies to this claim. Accordingly, based on the current record, we find Petitioner has shown sufficiently that Gai in view of IEEE 1999 suggests claim 19 for the same reasons discussed for claim 13. Patent Owner does not dispute that Gai in view of IEEE 1999 suggests the additional limitation in claim 19. Pet. 59; Ex. 1003 ¶ 202.

16. Analysis of Dependent Claims 14, 15, 18, 20, 21, 22, and 23

Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai teaches dependent claims 14, 15, 18, 20, 21, 22, and 23. *See* Pet. 55–57, 59–60, 62–65. These dependent claims are nearly identical to dependent claims 2, 3, 6, 8, 9, 10, and 11, respectively, but depend on independent claim 13 (or intermediate dependent claims that depend on claim 13) instead of claim 1.

Dependent claim 14 depends on independent claim 13 and adds “wherein the protocol processor is adapted to configure the first and second access ports so that no loop is created in the network between the edge switch and the one or more access switches.” Ex. 1001, claim 14. So, for the same reasons we found that Gai teaches claim 2, we find on this record it also teaches claim 14. Pet. 57; Ex. 1003 ¶¶ 198; Ex. 1005, 7:59–61.

Dependent claim 15 depends on claim 14 and adds “wherein the protocol processor is further adapted to configure the first and second access ports so as to decrease a radius of a spanning tree established in the network.” Ex. 1001, claim 15. So, for the same reasons we found that Gai teaches claim 3, we find on this record it also teaches claim 15.

Dependent claim 18 depends on independent claim 13 and adds “wherein the dummy traffic comprises sending data frames originating at the edge switch while having source addresses corresponding to network addresses of other entities in the network that are accessible to the access switches via the edge switch.” Ex. 1001, claim 18. This limitation is substantively the same as the limitation in claim 6 so our analysis for why Gai teaches it on this record is the same. Pet. 58; Ex. 1003 ¶ 200.

Dependent claim 20 depends on claim 19 and adds “wherein prior to the failure, the edge switch database comprises entries associated with a plurality of the ports, and wherein the protocol processor is adapted to use the network addresses in all the entries that are not associated with the first access port as the source addresses of the data frames.” Ex. 1001, claim 20. This limitation is substantively the same as the limitation in claim 8 so our analysis for why Gai in view of IEEE 1999 teaches it on this record is the same. Pet. 59; Ex. 1003 ¶ 203.

Dependent claim 21 depends on claim 18 and adds “wherein the protocol processor is adapted to incorporate a Virtual Local Area Network (VLAN) tag in a header of the data frames.” Ex. 1001, claim 21. This limitation is substantively the same as the limitation in claim 9, so our analysis for why Gai teaches it on this record is the same. Pet. 59; Ex. 1003 ¶ 204.

Dependent claim 22 depends on claim 18 and adds “wherein the protocol processor is adapted to multicast the data frames to all of the one or more access switches.” Ex. 1001, claim 22. This limitation is substantively the same as the limitation in claim 10, so our analysis for why Gai teaches it on this record is the same. Pet. 60; Ex. 1003 ¶ 205.

Dependent claim 23 depends on claim 18 and adds “wherein the dummy traffic comprises data frames sent downstream from the edge switch toward the access switches.” Ex. 1001, claim 23. This limitation is substantively the same as the limitation in claim 11, so our analysis for why Gai in view of IEEE 1999 teaches it on this record is the same. Pet. 60; Ex. 1003 ¶ 206.

E. Ground 2

Petitioner contends that dependent claims 4, 5, 16 and 17 are unpatentable as obvious over Gai, in view of IEEE 1999 and Jain, citing the testimony of Mr. McNair as support. Pet. 62–66; Ex. 1004. Patent Owner does not dispute that this combination teaches the additional limitations in these four dependent claims. We are persuaded, based on the current record, that Petitioner has established a reasonable likelihood of prevailing on its asserted ground based on the arguments presented as to all the challenged claims in Ground 2, as discussed below.

1. Jain

Jain is a U.S. patent titled “Spanning Tree with Rapid Forwarding Database Updates.” Ex. 1006, code (54). Similar to the ’986 patent, Jain also describes managing a recovery from a link failure in a network that implements a spanning tree protocol/algorithm. *Id.* at code (57) (“Management of forwarding databases in the case of link failures on bridges according to the all improved spanning tree, limits the propagation of notifications of topology change to only those parts of the network which are affected by the link failures . . .”).

Jain discloses detecting a failure and transitioning to a backup port. Ex. 1006, 5:15–18. To reduce recovery time, Jain proposes that the forwarding databases are transferred from the failed port to the backup port which has now become active. *Id.* at 4:28–33, 5:15–20. The updates are made before the failed port is deactivated and the backup port is activated to maintain network availability. *Id.*

2. Analysis of Dependent Claim 4

Dependent claim 4 depends on independent claim 1 and adds “comprising, responsive to placing the second port in the active state,

updating the database of the edge switch so that entries in the database point to the second port instead of the first port.” Ex. 1001, claim 4. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 and Jain at least suggests the additional limitation recited in dependent claim 4. *See* Pet. 62–64; Ex. 1003. As described by Petitioner, when an edge switch, such as switch 122 as described by Gai and discussed above, transitions a backup port to be active after detecting a failure, it would be beneficial to first modify the filtering database for that switch in order to reduce the possibility of misdirecting or losing messages using outdated database information, as taught by Jain. Pet. 62–64; Ex. 1003 ¶ 212–215.

Further, as described by Petitioner, and based on the current record, we are persuaded that a POSITA would have been motivated to modify Gai by incorporating Jain’s technique of immediately updating the filtering database. Pet. 62–63; Ex. 1003 ¶¶ 212–214. Both references relate to rapidly recovering from a link failure in a spanning tree network and this would address the concern raised in Gai about messages being misdirected or lost. Pet. 62–63; Ex. 1005, 6:31–38, 12:59–62, 15:18–21, Fig. 2; Ex. 1006, 4:28–33, 5:15–20; Ex. 1003 ¶ 212.

Patent Owner does not dispute that the combination of Gai, IEEE 1999, and Jain suggests the additional limitation in dependent claim 4 or that a POSITA would have been motivated to combine Jain with Gai.

3. *Analysis of Dependent Claim 5*

Dependent claim 5 depends on claim 4 but adds “wherein updating the database comprises changing the entries in the database to point to the second port before receiving the communication traffic from the one or more access switches on the second port.” Ex. 1001, claim 5. Based on the

current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 and Jain at least suggests the additional limitation recited in dependent claim 4. *See* Pet. 64–65; Ex. 1003 ¶ 216.

This claim requires the switch implementing the method for rapid recovery, such as switch 122 in Petitioner’s example discussed above, to update its filtering database before receiving communication traffic from access switches. Based on the current record, that appears to be what is taught by Jain as discussed above.

Patent Owner does not dispute that the combination of Gai, IEEE 1999, and Jain suggests the additional limitation in dependent claim 5.

4. Analysis of Dependent Claim 16

Dependent claim 16 depends on claim 13 and adds “comprising a memory, wherein the protocol processor is adapted to build an edge switch database in the memory responsive to the communication traffic prior to the failure associated with the first link, and responsive to placing the second access port in the active state, to update the edge switch database so that entries in the edge switch database point to the second access port instead of the first access port.” Ex. 1001, claim 16. Based on the current record, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 and Jain suggests the additional limitation in dependent claim 16. *See* Pet. 65–66; Ex. 1003 ¶¶ 217–219.

Substantively, this claim is similar to dependent claim 4, so our analysis above for that claim applies here for why Gai in view of IEEE 1999 and Jain suggests on this record that the edge switch updates the filtering database responsive to a failure. Patent Owner does not dispute that Gai in view of IEEE 1999 and Jain suggests the additional limitation in claim 16.

5. *Analysis of Dependent Claim 17*

Dependent claim 17 depends on claim 16 (which in turn depends on independent claim 13) and adds “wherein the protocol processor is adapted to update the entries in the database to point to the second access port before receiving communication traffic from the one or more access switches on the second port,” similar to the additional limitation recited in claim 5.

Ex. 1001, claim 17.

Therefore, based on the current record, for the reasons discussed above with respect to claim 5, we are persuaded that Petitioner has shown sufficiently that Gai in view of IEEE 1999 and Jain suggests the additional limitation in dependent claim 17. Pet. 66; Ex. 1003 ¶ 220. Patent Owner does not dispute that Gai in view of IEEE 1999 and Jain suggests the additional limitation in claim 17.

IV. CONCLUSION

Based on the arguments presented in the Petition, we conclude that Petitioner has demonstrated a reasonable likelihood of prevailing with respect to at least one claim of the ’986 patent challenged in the Petition. Accordingly, we institute a trial on all claims and all grounds asserted in the Petition. The Board has not made a final determination under 35 U.S.C. § 318(a) with respect to the patentability of the challenged claims.

V. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that, pursuant to 35 U.S.C. § 314(a), an *inter partes* review of claims 1–24 of the ’986 patent is instituted with respect to all grounds set forth in the Petition; and

FURTHER ORDERED that, pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4(b), *inter partes* review of the ’986 patent shall commence

IPR2024-00893
Patent 6,917,986 B2

on the entry date of this Decision, and notice is hereby given of the institution of a trial.

IPR2024-00893
Patent 6,917,986 B2

FOR PETITIONER:

Michael N. Rader
Scott A. McKeown
WOLF, GREENFIELD & SACKS, P.C.
Michael.Rader@WolfGreenfield.com
Scott.McKeown@WolfGreenfield.com

FOR PATENT OWNER:

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