

United States Court of Appeals for the Federal Circuit

02-1232

STORAGE TECHNOLOGY CORPORATION,

Plaintiff/Third Party Plaintiff-Appellant,

v.

CISCO SYSTEMS, INC. and CISCO TECHNOLOGY, INC.,

Defendants-Appellees,

and

LEVEL ONE COMMUNICATIONS, INC.,

Third Party Defendant.

Ernie L. Brooks, Brooks & Kushman, P.C., of Southfield, Michigan, argued for plaintiff-third party plaintiff-appellant. With him on the brief was Thomas A. Lewry.

Matthew D. Powers, Weil, Gotshal & Manges LLP, of Redwood Shores, California, argued for defendants-appellees. With him on the brief was Edward R. Reines. Of counsel was Sarkis Beudjekian.

Appealed from: United States District Court for the Northern District of California

Judge Susan Illston

United States Court of Appeals for the Federal Circuit

02-1232

STORAGE TECHNOLOGY CORPORATION,

Plaintiff/Third Party Plaintiff- Appellant,

v.

CISCO SYSTEMS, INC. and CISCO TECHNOLOGY, INC.,

Defendants-Appellees,

and

LEVEL ONE COMMUNICATIONS, INC.,

Third Party Defendant.

DECIDED: May 13, 2003

Before SCHALL, Circuit Judge, PLAGER, Senior Circuit Judge, and GAJARSA, Circuit Judge.

PLAGER, Senior Circuit Judge.

Storage Technology Corporation (StorageTek) owns U.S. Patent Nos. 5,842,040 and 5,566,170, both relating to data communication networks. StorageTek filed a patent infringement action in the United States District Court for the Western District of Wisconsin against Cisco Systems, Inc. and Cisco Technology, Inc. (collectively, Cisco) alleging infringement of both patents. After the case was transferred to the Northern District of California, the district court held a Markman hearing and entered a claim construction order. Subsequently, the district court granted Cisco's motions for summary judgment of noninfringement of the '040 and '170 patents. We affirm the district court's judgment of noninfringement of the '170 patent as well as the decision to transfer the case to the Northern District of California. However, because the district court erred in its construction of certain limitations in the '040 patent, we vacate the judgment of noninfringement of the '040 patent and remand for further proceedings.

BACKGROUND

Modern data communication networks consist of network devices, such as switches and routers, interconnected by high-speed transmission media, such as fiber optic cables. Data flowing through a network is divided into packets for transmission over the network. For example, one e-mail message may be broken down into several packets. In addition to user data, each packet includes a header, which contains several pieces of information, including information necessary for routing the packet through the network to its destination.

As each packet passes through a network device, the network device examines the header information to determine where to send the packet next. The network device may also make other decisions regarding the packet based on a set of rules or procedures known as network policies, which are set by system administrators. One such policy may be that packets belonging to certain types of messages get transmission priority over other packets. Another such policy may be that packets from a particular source are not to be forwarded through the network. Network policy information is typically stored in the main memory of a network device.

The '040 patent, assigned to StorageTek, relates to the efficient enforcement of network policies in a high-speed packet data network. When a packet (referred to in the patent as a protocol data unit, or PDU) arrives at a network device, the processor in the network device determines which network policies apply to the packet. To obviate the need for the processor to retrieve all required network policy information from main memory, which may slow down processing of packets, the invention described in the '040 patent provides a cache memory that stores certain network policy information copied from main memory. A cache memory is a small, usually temporary, memory that can be accessed quickly, often used to store a copy of instructions or data obtained from main memory. Thus in the system described in the '040 patent, network policy information that is used frequently by the processor is stored in cache memory for easy and fast accessibility.

A processor in a network device implementing the invention of the '040 patent examines the header of an arriving packet to determine whether, based on selection criteria, it is related to other packets that have been received. For example, packets containing data from the same e-mail message may be related. If the packet is the first of a group of related packets to be received, the network device determines whether any of the network policies in main memory apply to the packet. If so, policy identification information identifying the instance of network policy to be applied to the packet is copied from main memory into an instance classification cache, also referred to as the identification cache. The patent also describes a second cache, which stores a copy of the instance of network policy itself. When a packet related to the first packet arrives, the identification of a policy to be applied to the new packet is retrieved from the instance classification cache, and that identification information is used to retrieve the actual policy from the second cache.

The two independent claims of the '040 patent at issue in this case, method claim 1 and apparatus claim 18, read

(emphasis added to highlight the relevant claim limitations):

1. A policy caching method for use in a communication device, comprising the steps of:
 determining an instance of protocol data unit (PDU) network policy from a plurality of policies to be applied to related-received PDUs based on contents of one of the related-received PDUs, wherein the related-received PDUs are a subset of a stream of PDUs and may be distributed throughout said stream of PDUs; and
 caching policy identification information identifying the instance of PDU policy which is to be applied to other PDUs of the related-received PDUs.

18. A policy cache for use in a communication device, comprising:
 a) exception processing means for determining an instance of protocol data unit (PDU) network policy from a plurality of policies to be applied to related-received PDUs based on contents of one of the related-received PDUs, wherein the related-received PDUs are a subset of a stream of PDUs and may be distributed throughout said stream of PDUs; and
 b) cached instance classification means, operatively coupled to the exception processing means, for caching policy identification information identifying the instance of PDU policy which is to be applied to other PDUs of the related-received PDUs.

The '170 patent, also assigned to StorageTek, relates to the forwarding of packets by a network device. When a packet arrives at a network forwarding device, a processor in the device examines and processes the packet to assign it to an outbound network interface for transmission over the next link in the network. The invention of the '170 patent adds a preprocessor to the network forwarding device. Based upon information in the packet header, the preprocessor adds "next operation information" to the header. The packet is then passed to the processor, which forwards the packet based on the "next operation information" in the header. This approach reduces processing time and increases throughput of the network forwarding device.

The two independent claims of the '170 patent at issue in this case, apparatus claim 14 and method claim 23, read (emphasis added to highlight the relevant claim limitations):

14. A protocol data unit forwarding device for use in a communication network to transfer protocol data units within the communication network, comprising:

(a) protocol data unit processor, comprising:
 (i) identification means for determining media header information of a protocol data unit which is received from over the communication network;
 (ii) validation means, operatively coupled to the identification means, for validating the media header information; and
 (iii) modification means, operatively coupled to the identification means, for adding next operation information to the media header information based upon the determined media header information; and

(b) forwarding processor, operatively coupled to the protocol data unit preprocessor, for forwarding the protocol data unit in the communication network based upon the next operation information.

23. A method of operating a forwarding device within a communication network to forward a protocol data unit received by the forwarding device, the method comprising the device-implemented steps of:

(a) in a first processor, performing the steps of:

- (a)(i) determining media header information of a protocol data unit based upon a portion of the protocol data unit which is received from over a communication network;
- (a)(ii) validating the media header information; and
- (a)(iii) adding next operation information to the media header information based upon the determined media header information; and

(b) in a second processor, forwarding the protocol data unit in the communication network based upon the next operation information.

StorageTek sued Cisco in the United States District Court for the Western District of Wisconsin alleging infringement of the '040 and '170 patents by Cisco's NetFlow Feature Acceleration technology and Cisco's Multi-Protocol Label Switching (MPLS) technology, respectively. Cisco asserted a counterclaim for infringement of its own patent as well as declaratory judgment counterclaims for noninfringement, invalidity, and unenforceability of StorageTek's patents. Cisco later stipulated to the dismissal of its infringement and unenforceability counterclaims. The district court granted Cisco's motion to transfer venue pursuant to 28 U.S.C. § 1404(a) and transferred the case to the Northern District of California.

After a Markman hearing, the United States District Court for the Northern District of California issued a claim construction order. Storage Tech. Corp. v. Cisco Sys., Inc., No. C00-1176 (SI) (N.D. Cal. Nov. 27, 2001) (claim construction order). Relevant to this appeal, the district court construed the term "network policy" in the claims of the '040 patent to mean "a filtering or auditing rule which determines if and why a PDU should be forwarded." In claim 1 of the '040 patent, the district court construed the step of "caching policy identification information" to mean "storing information that identifies and is used to retrieve a separately cached instance of network policy." In other words, the court construed claim 1 to require two caches, one for storing information identifying a policy, and a second for storing a copy of the policy itself. Regarding apparatus claim 18, the court construed "cached instance classification means" under 35 U.S.C. § 112, ¶ 6 to mean "the structure (or equivalent structure) described in the specification of the '040 patent for caching policy identification information identifying the instance of PDU policy which is to be applied to other PDUs of the related-received PDUs."

The district court also construed the term "next operation information" in the claims of the '170 patent to mean "information that specifies the particular type of operation to be performed by the network forwarding device. The information is added to the media header information and must include information to accomplish forwarding of the PDU in the communication network." The court rejected StorageTek's contention that "next operation information" relates to operations to be performed downstream of the preprocessor, holding instead that the preprocessor and the forwarding processor, which uses the next operation information added by the preprocessor, must be part of the same forwarding device.

Cisco moved for summary judgment of noninfringement of the '040 and '170 patents, which the district court granted on February 4, 2002. Storage Tech. Corp. v. Cisco Sys., Inc., No. C00-1176 (SI) (N.D. Cal. Feb. 4, 2002) (order granting motion for summary judgment of noninfringement). Regarding the '040 patent, the district court found that Cisco's NetFlow technology has only one cache, and therefore does not literally infringe the asserted claims, which the district court

construed to require separate caches for policy identification information and the policy itself. Addressing the '170 patent, the district court found that there could be no literal infringement by Cisco's MPLS technology because one network device, a Label Edge Router (LER), applies an MPLS label (alleged to be the claimed "next operation information") to a packet, and a second network device, a Label Switch Router (LSR), forwards the packet based on the MPLS label. Thus, the court found, MPLS does not add next operation information to be performed by the network device that added the label, as required by the asserted claims. The district court further held that arguments made during prosecution of the '170 patent estopped StorageTek from asserting a theory of infringement under the doctrine of equivalents.

StorageTek timely appeals the district court's grant of Cisco's motions for summary judgment of noninfringement with respect to both the '040 and '170 patents. StorageTek also appeals the decision by the United States District Court for the Western District of Wisconsin to transfer the case to the Northern District of California.

DISCUSSION

I.

Before we reach the merits of this case, we first must attend to our jurisdiction. Jim Arnold Corp. v. Hydrotech Sys., Inc., 109 F.3d 1567, 1571, 42 USPQ2d 1119, 1122 (Fed. Cir. 1997). At oral argument, this court informed counsel for both parties that it appeared from the record that Cisco's counterclaims for invalidity of the '040 and '170 patents remained pending in the district court and that the district court had not entered a final judgment or certification pursuant to Rule 54(b) of the Federal Rules of Civil Procedure. Although these circumstances raised a question as to whether we had jurisdiction over the appeal, we nevertheless heard oral argument. Shortly thereafter, this court by letter requested counsel to inform the court as to the status of any counterclaims on the date notice of appeal was filed and whether the district court had entered a final judgment or Rule 54(b) certification on or before that date.

In a joint response to that letter, the parties informed us that on the date StorageTek filed the notice of appeal in this case, Cisco's invalidity counterclaims remained pending in the district court. As of that date, a final judgment had not been entered by the district court, nor had an appropriate judgment or certification been entered by the district court pursuant to Rule 54(b). The parties also informed us that they had jointly submitted a Stipulation and Order to the district court requesting the court to enter a Rule 54(b) certificate nunc pro tunc to February 4, 2002, the date on which the district court granted summary judgment of noninfringement in favor of Cisco. The district court complied with the parties' request, signing the order and entering judgment under Rule 54(b) nunc pro tunc to February 4, 2002. Storage Tech. Corp. v. Cisco Sys., Inc., No. C00-1176 (SI) (N.D. Cal. Dec. 20, 2002) (Judgment Pursuant to Rule 54(b)).

Under Federal Circuit law, which we apply instead of regional circuit law to issues involving Rule 54(b) certification and our appellate jurisdiction, the district court's December 20, 2002 Rule 54(b) certification relates back to February 4, 2002 and ripens the premature notice of appeal. See State Contracting & Eng'g Corp. v. Florida, 258 F.3d 1329, 1334-35, 59 USPQ2d 1498, 1502 (Fed. Cir. 2001). Accordingly, we conclude that we have jurisdiction to decide the appeal in this case.

II.

We review a district court's grant of a motion for summary judgment without deference. Ethicon Endo-Surgery, Inc. v. United States Surgical Corp., 149 F.3d 1309, 1315, 47 USPQ2d 1272, 1275 (Fed. Cir. 1998). Claim construction is a matter of law over which we exercise independent review. Cybor Corp. v. FAS Techs., Inc., 138 F.3d 1448, 1456, 46 USPQ2d 1169, 1174 (Fed. Cir. 1998) (en banc).

A.

StorageTek challenges the district court's construction of two limitations in the claims of the '040 patent—"caching policy identification information" and "protocol data unit (PDU) network policy." We address each in turn.

Claim 1 of the '040 patent includes the step of "caching policy identification information." StorageTek contends that the district court erred when it construed that claim language to require not only caching of policy identification information, but also caching of an instance of network policy. Cisco argues, as it did before the district court, that the written description and prosecution history make it clear that the invention involves utilizing the cached policy identification information to retrieve a cached instance of policy, and therefore the claims must be construed to require the instance of policy to be cached. As discussed below, we agree with StorageTek that the district court's interpretation of claim 1 was overly narrow.

Our analytical focus must begin with and remain centered on the language of the claims themselves. Tex. Digital Sys., Inc. v. Telegenix, Inc., 308 F.3d 1193, 1201, 64 USPQ2d 1812, 1817 (Fed. Cir. 2002). Claim 1 is a method claim with two steps—determining an instance of network policy to be applied to related PDUs based on the contents of one PDU, and caching policy identification information identifying that instance of network policy. The claim does not include any step relating to caching the instance of network policy or any limitation indicating that the instance of network policy is stored in or retrieved from a cache. According to its plain language, the limitation construed by the district court, "caching policy identification information," requires only the identification information, not the instance of network policy, to be cached.

The district court based its interpretation in part on the presence of the phrase "policy caching method" or "policy

cache” in the preamble of every independent claim of the ’040 patent. Whether to treat a preamble as a claim limitation is determined on the facts of each case in light of the claim as a whole and the invention described in the patent. See Catalina Mktg. Int’l, Inc. v. Coolsavings.com, Inc., 289 F.3d 801, 808, 62 USPQ2d 1781, 1784 (Fed. Cir. 2002). In the case of the ’040 patent, the written description consistently uses the terms “policy caching method” and “policy cache” to refer to the invention as a whole, not to the specific step of storing an instance of network policy or to the cache that stores the instance of network policy, which the written description refers to as the “instance policy cache.” Similarly, the term “policy caching method” or “policy cache” in the preamble of each claim serves as a convenient label for the invention as a whole. See IMS Tech., Inc. v. Haas Automation, Inc., 206 F.3d 1422, 1434, 54 USPQ2d 1129, 1137 (Fed. Cir. 2000) (holding that preamble phrase “control apparatus” does not limit claim scope when it merely gives a descriptive name to the claimed invention). In contrast, any claim in the ’040 patent that requires the instance of policy to be cached includes an explicit “instance policy cache” limitation in the body of the claim. See ’040 patent, col. 17, l. 20 (claim 19 includes “instance policy cache means”). Thus, the preamble terms “policy caching method” or “policy cache” do not limit claim scope and simply refer to the invention set forth in the body of the claim, which, depending on the limitations provided in a particular claim, may include caching policy identification information or caching an instance of policy or both. Claim 1 includes only the step of caching policy identification information, and therefore it was improper for the district court to further limit the scope of claim 1 based on language in the preamble.

The district court also relied on the written description, prosecution history, and a declaration by Cisco’s expert witness regarding the purpose of the invention to define claim 1 to require a cache for the instance of network policy. In so doing, the district court disregarded the well-established rule that while proper claim construction requires an examination of the written description and relevant prosecution history to determine the meaning of claim limitations, additional limitations may not be read into the claims. See, e.g., Prima Tek II, L.L.C. v. Polypap, S.A.R.L., 318 F.3d 1143, 1148, 65 USPQ2d 1818, 1821 (Fed. Cir. 2003); Comark Communications, Inc. v. Harris Corp., 156 F.3d 1182, 1186, 48 USPQ2d 1001, 1005 (Fed. Cir. 1998). We have recognized that there is sometimes a fine line between reading a claim in light of the written description and relevant prosecution history, and reading a new limitation into the claim. Comark, 156 F.3d at 1186, 48 USPQ2d at 1005. However, “interpreting what is meant by a word in a claim ‘is not to be confused with adding an extraneous limitation . . . , which is improper.’” Intervet Am., Inc. v. Kee-Vet Labs., Inc., 887 F.2d 1050, 1053, 12 USPQ2d 1474, 1476 (Fed. Cir. 1989) (quoting E.I. du Pont de Nemours & Co. v. Phillips Petroleum Co., 849 F.2d 1430, 1433, 7 USPQ2d 1129, 1131 (Fed. Cir. 1988)).

Cisco relies on the fact that Figure 1 of the ’040 patent shows two caches—a “cached instance classification” cache and a “cached instance policy” cache—to support its theory that claim 1 requires the instance of network policy to be cached. This argument, however, ignores Figure 2, which is a flowchart showing the steps of the inventive method. Figure 2 includes a step for caching policy identification information, but notably does not show a step for caching an instance of

network policy. Given the lack of explicit language in claim 1 requiring the instance of network policy to be cached and the absence of any such step from Figure 2, the written description is not a basis for construing the claim as requiring the step of caching an instance of network policy.

Cisco's use of the prosecution history to narrow the meaning of claim 1 is also misplaced. During prosecution, the patent applicants stated that in the invention as recited in claims 1, 11, and 18, the instance of network policy and the policy identification information are both cached. While on its face this statement appears to limit claim scope, it cannot do so absent some claim language referring to the caching of the instance of network policy. The prosecution history statement describes generally the features of the claimed invention and erroneously suggests that the independent claims include a cache for the instance of network policy. The applicants' inaccurate statement cannot override the claim language itself, which controls the bounds of the claim. See Rambus Inc. v. Infineon Techs. AG, 318 F.3d 1081, 1089, 65 USPQ2d 1705, 1711 (Fed. Cir. 2003) (holding that general statement introducing new limitations does not limit scope of claims not amended to include the new limitations); Intervet, 887 F.2d at 1054, 12 USPQ2d at 1477 (holding that erroneous statement made during prosecution does not limit claim scope because "[t]he claims themselves control").

Finally, the district court improperly relied on extrinsic evidence in the form of a declaration by Cisco's expert to support its construction of claim 1. Resort to extrinsic evidence is appropriate only when an ambiguity remains after consulting the intrinsic evidence of record. Vitronics Corp. v. Conceptoronic, Inc., 90 F.3d 1576, 1583, 39 USPQ2d 1573, 1578 (Fed. Cir. 1996). Moreover, the district court did not use the extrinsic evidence to assist in defining a claim limitation, but rather used it to limit claim scope based on the purpose of the invention, which is impermissible. See Comark, 156 F.3d at 1187, 48 USPQ2d at 1005 (rejecting defendant's attempt to limit claim term to its functional purpose as disclosed in the preferred embodiment).

StorageTek also asserted claim 18 against Cisco. Claim 18 recites a "cached instance classification means . . . for caching policy identification information." The district court correctly interpreted this limitation as "the structure (or equivalent structure) described in the specification of the '040 patent for caching policy identification information identifying the instance of PDU policy which is to be applied to other PDUs of the related-received PDUs." The court also cited to Figure 1 of the '040 patent, indicating that the corresponding structure described in the patent was element 108, the "cached instance classification" cache. Significantly, the district court did not include as part of the corresponding structure the second cache shown in Figure 1—the "cached instance policy" cache (element 110). Instead, the court held that the "cached instance policy" cache corresponds to an additional limitation found in dependent claim 19—the "instance policy cache means." We agree with the district court's claim construction. Thus, although Figure 1 shows two caches, the language of claim 18 requires only one cache and requires only the policy identification information, and not the instance of policy, to be cached.

StorageTek also challenges the district court's construction of "protocol data unit (PDU) network policy," which appears in both claim 1 and claim 18. The parties agree that the term refers to a rule or procedure set by a system administrator. The point of contention is whether the term should be limited to filtering and auditing rules or should be more broadly construed to include rules related to routing, bridging, and switching functions. The district court accepted Cisco's position that the written description and prosecution history limited the meaning to filtering and auditing rules.

We begin with the ordinary meaning of the disputed claim term. Inverness Med. Switz. GmbH v. Princeton Biomeditech Corp., 309 F.3d 1365, 1369, 64 USPQ2d 1926, 1930 (Fed. Cir. 2002). StorageTek asserts, and Cisco does not dispute, that the ordinary meaning of "network policy" as understood by one skilled in the art is a rule or procedure set by a network system administrator. Neither party has proposed a contrary definition from a dictionary or technical treatise. See Tex. Digital, 308 F.3d at 1202, 64 USPQ2d at 1818. Rather, both parties direct our attention to the intrinsic record, i.e., the written description and the prosecution history.

The first sentence of the "Background of the Invention" section states that "routing, bridging, switching, filtering, and auditing functions" of PDUs are handled by network policies. '040 patent, col. 1, ll. 22-26. This information is repeated at column 2, lines 9-13: "These routing, bridging, switching, filtering, and auditing functions are also governed by policy set forth by rules and procedures set by system administrators and the like. These rules and procedures are generally known as communications network policies." Although the "Summary of the Invention" section specifies only filtering and auditing rules as network policies, they are presented as examples. Id., col. 2, l. 65 – col. 3, l. 1 ("Examples of such policies are filtering (e.g., forward or not forward the PDU) and/or auditing (e.g., send a copy of the PDU to a third party)."). When read in context, the statement relied upon by the district court from the next paragraph, that an "instance of network policy is performed by filtering or auditing the related-received PDUs," refers back to the examples described in the previous paragraph and does not narrow the definition provided in the background section. See id., col. 3, ll. 9-11. Based on our independent review of the '040 patent, we conclude that the definition of the term "network policy" provided in the written description is a rule or procedure set by a system administrator governing one or more of a routing, bridging, switching, filtering, or auditing function.

The prosecution history does not warrant a narrower claim construction. In response to a rejection over a prior art reference disclosing a device for bridging local area networks, the patent applicants stated: "Unlike a bridge, which analyzes source or destination addresses to determine where data should be forwarded, the present invention more broadly analyzes the content of PDUs in order to perform filtering and auditing functions which determine if and why data should be forwarded." (Emphasis in original.) When this remark is viewed in context, it is apparent that the applicants were distinguishing their invention from a prior art bridge that, although it performs bridging functions, lacks a mechanism for enforcing any network policies, including policies relating to filtering and auditing functions as well as policies relating to

bridging functions. We therefore do not consider the applicants' statement to be a clear and unambiguous disavowal of claim scope as required to depart from the meaning of the term provided by the written description. See, e.g., Schwing GmbH v. Putzmeister Aktiengesellschaft, 305 F.3d 1318, 1324 -25, 64 USPQ2d 1641, 1645 (Fed. Cir. 2002) (“[P]rosecution history . . . cannot be used to limit the scope of a claim unless the applicant took a position before the PTO that would lead a competitor to believe that the applicant had disavowed coverage of the relevant subject matter.”); IMS Tech., 206 F.3d at 1439, 54 USPQ2d at 1141.

Cisco contends on appeal that even if claims 1 and 18 of the '040 patent do not require the caching of an instance of network policy, there is an independent basis for affirming the district court's grant of Cisco's motion for summary judgment of noninfringement. According to Cisco, StorageTek failed to submit evidence identifying which aspect of Cisco's NetFlow products satisfies the “caching policy identification information” limitation. StorageTek's arguments made in opposition to Cisco's summary judgment motion, however, were premised on an incorrect claim construction that required two caches. “[A] change in the claim construction at the appellate level generally necessitates a remand to the district court to resolve any new factual issues raised by the new claim construction, except, of course, in the rare instance that the record on appeal leaves no genuine issues of material fact and entitles the movant to judgment as a matter of law.” Electro Scientific Indus., Inc. v. Dynamic Details, Inc., 307 F.3d 1343, 1350, 64 USPQ2d 1781, 1785 (Fed. Cir. 2002). Based on the factual record before us, we cannot determine whether there are any genuine issues of material fact regarding infringement. Accordingly, we vacate the district court's grant of summary judgment of noninfringement of the '040 patent and remand for further proceedings consistent with this opinion. We note that although the district court correctly construed claim 18 in its Markman order, the court misapplied its claim construction in deciding Cisco's motion for summary judgment, apparently importing its incorrect construction of “caching policy identification information” into claim 18. Therefore, we vacate the district court's judgment with respect to claim 18 as well as claim 1.

B.

The sole claim construction issue with respect to claims 14 and 23 of the '170 patent is whether the “protocol data unit processor” (or “first processor” in claim 23), which adds “next operation information” to the packet header, and the “forwarding processor” (or “second processor” in claim 23), which forwards the packet based on the added next operation information, must be part of the same forwarding device. StorageTek contends that the district court erred in interpreting the term “forwarding device,” which appears only in the preamble of each claim, to be a limitation. In StorageTek's view there is no “forwarding device,” and the two processors recited in the claims may be anywhere in the communication network.

In this case, we agree with the district court that the preamble term “forwarding device” limits the claims. The written description emphasizes that a key benefit of the invention is that the use of the next operation information within a

forwarding device reduces the processing otherwise performed by the forwarding device. For example, the “Summary of the Invention” states that “a protocol data unit preprocessor is used in a protocol data unit forwarding device . . . such that subsequent processing of the protocol data unit by the protocol data unit forwarding device is reduced.” ’170 patent, col. 4, l. 64 – col. 5, l. 11 (emphasis added). Figure 2 of the patent shows a forwarding device 100 that contains preprocessor 104 and forwarding processor 108. Because the written description makes it abundantly clear that the preprocessor and the forwarding processor are part of a single forwarding device, the term “forwarding device” in the preambles of claims 14 and 23 limits the scope of the claims such that both processors must be contained within one forwarding device. See Catalina Mktg., 289 F.3d at 808, 62 USPQ2d at 1785 (“[W]hen reciting additional structure or steps underscored as important by the specification, the preamble may operate as a claim limitation.”).

The prosecution history supports the district court’s claim construction. Claims 14 and 23 as originally filed contained the phrase “such that subsequent processing of the protocol data unit by the protocol data unit forwarding device is reduced,” referring back to the “forwarding device” recited in the preamble. Had the claims issued in that form, the term “forwarding device” unquestionably would have been a limitation. See id. (“Dependence on a particular disputed preamble phrase for antecedent basis may limit claim scope because it indicates a reliance on both the preamble and claim body to define the claimed invention.”). The patent applicants, however, deleted that phrase in an amendment during the course of prosecution. In some circumstances that deletion might have removed the “forwarding device” as a limitation, because it is the claims as allowed that we must examine. See Intervet, 887 F.2d at 1054, 12 USPQ2d at 1477. Nevertheless, we are not persuaded by StorageTek’s argument that removal of the reference to “forwarding device” eliminated any limiting effect of the preamble. The reason is that in their remarks accompanying the amendment, the applicants stated that, in response to an indefiniteness rejection, the language was deleted as unnecessary “because the use of the added next operation information by the forwarding device is recited as part of each claim.” (Emphasis added.) Because the claims recite the use of the next operation information by the forwarding processor (claim 14) or second processor (claim 23), it follows that the applicants considered the forwarding processor or second processor to be part of the forwarding device recited in the preamble, which therefore operates as a limitation upon the claims.

The applicants also relied on the existence of a “forwarding device” in distinguishing their invention over prior art. In comparing the claimed invention to the Johnston reference, the applicants stated: “Instead of being part of a forwarding device in a communication network as required by the present invention, the ATM/AAL interface of Johnston et al. is more aptly characterized as an interface device.” (Emphasis added.) The passage goes on to allege that there is no teaching or suggestion in Johnston of a preprocessor that “will modify media header information by inserting next operation instructions so as to decrease the amount of processing that would otherwise be required to be performed by the primary processor in the forwarding device.” (Emphasis added.) These statements further support the district court’s conclusion that the term “forwarding device” is a limitation of the claims and that both processors recited in the body of claims 14 and 23 are part of

the forwarding device. See Catalina Mktg., 289 F.3d at 808, 62 USPQ2d at 1785 (“[C]lear reliance on the preamble during prosecution to distinguish the claimed invention from the prior art transforms the preamble into a claim limitation because such reliance indicates use of the preamble to define, in part, the claimed invention.”).

We also reject StorageTek’s argument that even if the claims require both processors to be contained in a forwarding device, the term “forwarding device” can be interpreted broadly to include a data communication network or the Internet itself. The ’170 patent provides as examples of forwarding devices a data packet switch, a bridge, and a router. ’170 patent, col. 1, ll. 31-32; col. 2, ll. 14-15. There is no suggestion in the claim language or the written description that the term could refer to an entire network.

Because we affirm the district court’s claim construction, and it is undisputed that in Cisco’s MPLS technology the alleged “next operation information” is added by one network device and used by a second, downstream network device, we affirm the district court’s grant of Cisco’s motion for summary judgment that it does not literally infringe the ’170 patent.

We also affirm the district court’s judgment regarding infringement under the doctrine of equivalents. In particular, we agree with the district court that prosecution history estoppel bars StorageTek from arguing that Cisco’s MPLS technology contains an equivalent of the forwarding device recited in claims 14 and 23 of the ’170 patent. As discussed above, during prosecution the applicants clearly stated that the next operation information added by the preprocessor is used to reduce the processing time of the processor in the forwarding device. StorageTek cannot now assert that the existence of a forwarding device containing both processors is unnecessary and that any communication network may be the equivalent of a forwarding device. See Allen Eng’g Corp. v. Bartell Indus., Inc., 299 F.3d 1336, 1350, 63 USPQ2d 1769, 1776 (Fed. Cir. 2002) (“[Prosecution history] estoppel also may be found on the basis of arguments made during prosecution of the application to secure the allowance of claims.”).

III.

In reviewing a district court’s decision regarding a motion to transfer under 28 U.S.C. § 1404(a), this court applies the law of the appropriate regional circuit, in this case the Seventh Circuit. See Winner Int’l Royalty Corp. v. Wang, 202 F.3d 1340, 1352, 53 USPQ2d 1580, 1589 (Fed. Cir. 2000). The Seventh Circuit will reverse a district court’s ruling on a motion to transfer only for a clear abuse of discretion. Tice v. Am. Airlines, 162 F.3d 966, 974 (7th Cir. 1998).

The gist of StorageTek’s argument on appeal is that the district court in Wisconsin did not properly consider the District of Minnesota, StorageTek’s alternative choice of forum, when it granted Cisco’s motion to transfer to the Northern District of California. We disagree with StorageTek’s reading of the district court’s order. The district court carefully examined the various factors related to forum convenience and concluded that the factors warranted transfer to the Northern District of

California. Significantly, the district court stated that its analysis did not change when it considered StorageTek's second choice of forum, the District of Minnesota. Contrary to StorageTek's contentions on appeal, the district court did consider that some of StorageTek's witnesses and documents were located in Minnesota. Nevertheless, the district court found that the relevant factors favored the Northern District of California, where almost all of Cisco's relevant witnesses and documents were located. We discern no abuse of discretion in the district court's decision to transfer the case to the Northern District of California.

CONCLUSION

The district court erred in its claim construction of the terms "caching policy identification information" and "protocol data unit (PDU) network policy" in the '040 patent. Accordingly, we vacate the district court's grant of summary judgment of noninfringement of the '040 patent and remand for further proceedings consistent with this opinion. The district court, however, properly granted summary judgment of noninfringement of the '170 patent. Finally, the United States District Court for the Western District of Wisconsin did not abuse its discretion when it transferred the case to the Northern District of California.

AFFIRMED-IN-PART, VACATED -IN-PART, and REMANDED