

PARTIES AND BOARD ONLY

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION and
INTERNATIONAL BUSINESS MACHINES CORPORATION,
Petitioners,

v.

PARALLEL NETWORKS LICENSING, LLC,
Patent Owner.

Case: IPR2015-00483¹
Patent 5,894,554

Before KEVIN F. TURNER, JEREMY M. PLENZLER, and
CHRISTOPHER L. CRUMBLEY, *Administrative Patent Judges*.

CRUMBLEY, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318 and 37 C.F.R. § 42.73

¹ IPR2015-00484 has been consolidated with this proceeding; International Business Machines Corporation was joined as a party to this proceeding via Motions for Joinder in IPR2015-01729 and IPR2015-01731.

I. INTRODUCTION

In this *inter partes* review trial, instituted pursuant to 35 U.S.C. § 314, Petitioners Microsoft Corporation and International Business Machines Corporation (collectively, “Microsoft”) challenge the patentability of claims 12–49 of U.S. Patent No. 5,894,554 (Ex. 1001, “the ’554 patent”), owned by Parallel Networks Licensing, LLC (“Parallel”).

We have jurisdiction under 35 U.S.C. § 6(c). This Final Written Decision, issued pursuant to 35 U.S.C. § 318(a)² and 37 C.F.R. § 42.73, addresses issues and arguments raised during trial. For the reasons discussed below, we determine that Microsoft has not met its burden to prove, by a preponderance of the evidence, that claims 12–49 of the ’554 patent are unpatentable.

A. Procedural History

On December 23, 2014, Petitioner Microsoft Corporation requested an *inter partes* review of claims 12–19, 32, 34, 46 and 48 of the ’554 patent, in a Petition that was assigned case number IPR2015-00483. Paper 1, “Pet.” On the same date, Microsoft Corporation filed a separate, but substantially similar, Petition challenging claims 12, 20–31, 33, 35–45, 47 and 49 of the ’554 patent. IPR2015-00484, Paper 1, “484 Pet.” Parallel filed Patent Owner Preliminary Responses in both cases. Paper 7.³ In two Decisions on

² As permitted by 35 U.S.C. § 316(a)(11) and 37 C.F.R. § 42.100(c) in cases of joinder, the Board adjusted the period for issuing the Final Written Decision beyond one year from the institution of trial.

³ The Preliminary Response was filed as “Parties and Board Only,” accompanied by a Motion to Seal. Paper 8. A redacted version of the Preliminary Response was filed with the Motion. Paper 9. We granted-in-

Institution of *Inter Partes* Review (IPR2015-00483, Paper 10, “Dec.”; IPR2015-00484, Paper 10), we instituted trial. By an Order that same day, we consolidated the grounds of trial in IPR2015-00484 into the instant case, such that trial was instituted on claims 12–49 based on the following grounds:

1. Whether claims 12–17, 20–32, 34–46 and 48 are unpatentable under 35 U.S.C. § 102 as anticipated by SWEB.⁴
2. Whether claims 12–49 are unpatentable under 35 U.S.C. § 103 as having been obvious over SWEB.
3. Whether claims 14, 15, 18, 19, 33, and 47 are unpatentable under 35 U.S.C. § 103 as having been obvious over SWEB and Leaf.⁵
4. Whether claims 17, 29, 43, and 49 are unpatentable under 35 U.S.C. § 103 as having been obvious over SWEB and Bradley.⁶

Paper 11, 2–3.

Following institution of trial, Petitioner International Business Machines Corporation filed two Petitions, substantially identical to those filed by Petitioner Microsoft, accompanied by Motions for Joinder. IPR2015-01729, Papers 1, 3; IPR2015-01731, Papers 1, 3. We granted

part the Motion to Seal (Paper 39), requiring Parallel to file a more narrowly redacted version of the Preliminary Response, which was entered as Paper 41.

⁴ Daniel Andresen et al., *SWEB: Towards a Scalable World Wide Web Server on Multicomputers*, Dept. of Computer Sci. Tech. Report TRCS95-17, U.C. Santa Barbara, 1–25 (Sept. 1995) (Ex. 1009).

⁵ U.S. Patent No. 5,754,772 (May 19, 1998) (Ex. 1060).

⁶ S.W. Bradley et al., *Web-based Access to an Online Atlas of Anatomy: The Digital Anatomist Common Gateway Interface*, 512–16 (1995) (Ex. 1048).

those Motions and joined International Business Machines as a Petitioner to IPR2015-00483. IPR2015-01729, Paper 9; IPR2015-01731, Paper 9.

Parallel filed a Patent Owner Response (Paper 26⁷) and subsequently, with the Board's authorization, filed a Corrected Response (Paper 54, "PO Resp."). Microsoft filed a Reply (Paper 43,⁸ "Reply").

Microsoft supported its Petition with the Declaration of Michael D. Mitzenmacher (Ex. 1007), and Parallel took the cross-examination of Mr. Mitzenmacher via deposition (Ex. 2086). Microsoft submitted a second Declaration of Mr. Mitzenmacher with its Reply (Ex. 1083), and Parallel again took cross-examination by deposition (Ex. 2107). Parallel submitted Observations on Mr. Mitzenmacher's second cross-examination (Paper 61), to which Microsoft replied (Paper 70).

With its Response, Parallel submitted the Declarations of Mark Jones (Ex. 2091), Keith Lowery (Ex. 2092), and Ronald Howell (Ex. 2105), and Microsoft took the cross-examination of Dr. Jones (Ex. 1084), Mr. Lowery (Ex. 1085), and Mr. Howell (Ex. 1086) via deposition.

Microsoft filed a Motion to Exclude various exhibits submitted by Parallel (Paper 64), to which Parallel filed a Response (Paper 68) and Microsoft filed a Reply (Paper 76). Parallel similarly filed a Motion to

⁷ Parallel filed the Patent Owner Response as "Parties and Board Only," accompanied by a Motion to Seal. Paper 25. We denied the Motion to Seal as to the Response (Paper 39), and the Response was subsequently made publicly available.

⁸ Microsoft's Reply was filed as "Parties and Board Only," along with a Motion to Seal (Paper 45) and a redacted version of the Reply (Paper 44). In a Decision issued today, we deny the Motion as to the Reply.

Exclude some of Microsoft's exhibits (Paper 62), to which Microsoft filed a Response (Paper 69) and Parallel filed a Reply (Paper 77).

Oral hearing was requested by both parties, and a hearing was held March 16, 2016. A transcript of the oral hearing is included in the record. Paper 80, "Tr."

B. The '554 Patent

The '554 patent relates to a system and method for managing dynamic Web page generation requests, and providing Web pages in response to those requests. Ex. 1001, 2:20–32. According to the patent, prior art systems were unable to handle large numbers of simultaneous incoming Web page requests, because a single server was responsible for both receiving the requests and providing HTML data in response to the requests. *Id.* at 3:64–4:53 (describing prior art Web server environments). The '554 patent addresses this problem by distributing processing responsibility across a plurality of servers, as illustrated in Figure 4 below:




FIG. 4

Figure 4 depicts a request from Web client 200 received by Web server 201, which is then intercepted by interceptor 400 and routed by dispatcher 402 to one of a plurality of page servers 404.

In contrast to the described prior art systems, instead of Web server executable 201(E) processing the incoming Web page request, in the system of the '554 patent Interceptor 400 intercepts the request and routes it to Dispatcher 402, which may reside on the same machine as the Web server or on a different machine. *Id.* at 4:58–60; 5:8–36. The Dispatcher then dispatches the request to one of a plurality of page servers, which retrieves the data necessary for processing the Web page request from one or more data sources. *Id.* at 5:37–48. The Dispatcher may take into account various criteria in selecting the appropriate page server, including whether the page server is connected to the appropriate data source, as well as the number of

requests already being processed by the server (“load balancing”). *Id.* at 5:60–6:19.

C. Illustrative Claim

Each of the challenged claims is set forth in a Certificate of Correction, which corrects the listing of claims in a Reexamination Certificate issued July 24, 2012. Ex. 1001, Certificate of Correction. Claims 12, 20, 32, 34, 46, and 48 are independent and generally similar in format and substance. Claims 12, 32, and 34 are directed to computer-implemented methods, of which claim 12 is illustrative:

12. A computer-implemented method for managing a dynamic Web page generation request to a Web server, said computer-implemented method comprising the steps of:

routing said request from said Web server to a selected page server, said selected page server receiving said request and releasing said Web server to process other requests, wherein said routing step further includes the steps of intercepting said request at said Web server, routing said request from said Web server to a dispatcher, and dispatching, by said dispatcher, said request to said selected page server;

processing said request, said processing being performed by said selected page server while said Web server concurrently processes said other requests; and

dynamically generating a Web page by said selected page server in response to said request, said Web page including data dynamically retrieved from one or more data sources; and

wherein dispatching includes:

examining said request to make a selection of which page server should process said request from among a plurality of page servers that can each generate said Web page requested by said request;

selecting one of said plurality of page servers to dynamically generate said Web page;

wherein said selection is based on examining dynamic information regarding a load associated with each of said plurality of page servers; and

sending said request to said selected page server based on said examination.

Id.

Independent claims 20, 46, and 48 are directed to machine-readable media having instructions stored thereon, which when executed by a computer cause the computer to perform the steps recited in claims 12, 32, and 34, respectively. *Id.*

II. DISCUSSION

A. Claim Construction

For purposes of our Decision to Institute, we analyzed each claim term in light of its broadest reasonable interpretation, as understood by one of ordinary skill in the art and as consistent with the specification of the '554 patent. 37 C.F.R. § 42.100(b); *Cuozzo Speed Techs., LLC v. Lee*, 136 S. Ct. 2131, 2144–46 (2016) (upholding the use of the broadest reasonable interpretation standard). In the Decision, we construed *Web server* as “software, or a machine having software, that receives Web page requests and may process said requests” and *page server* as “software, or a machine having software, that locates or generates a Web page.” Dec. 10–11.

We also noted in our Decision that the '554 patent was scheduled to expire on April 23, 2016, and that in such situations, the Board has considered it appropriate to reconsider the claim construction standard at the time of a final written decision, in light of the expiration of the claims. *Id.* at

7 n.6. In the case of an expired patent, our claim interpretation is similar to that of a District Court. *See In re Rambus Inc.*, 694 F.3d 42, 46 (Fed. Cir. 2012). Under this standard, we look “principally to the intrinsic evidence of record, examining the claim language itself, the written description, and the prosecution history, if in evidence.” *DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.*, 469 F.3d 1005, 1014 (Fed. Cir. 2006) (citing *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312–17 (Fed. Cir. 2005) (en banc)).

We have reviewed our prior constructions of *Web server* and *page server*, as set forth in the Decision to Institute, and do not consider application of the District Court-like standard to result in a different construction than that resulting from application of the broadest reasonable interpretation standard. Nor have the parties articulated any persuasive reason to depart from these constructions.

Our final decision in this case also requires the interpretation of an additional claim term, namely *request* (or, as set forth in the preamble, *dynamic Web page generation request*). The parties agreed in a copending District Court action that this term should be interpreted as “a message that asks for a Web page.” Ex. 2089, 6. We agree that such a construction is reasonable under the District Court-like standard the Board applies to expired patents, and adopt it herein.

B. Anticipation by SWEB

Microsoft contends that claims 12–17, 20–32, 34–46 and 48 are unpatentable under 35 U.S.C. § 102 as anticipated by SWEB. Pet. 24–42; 484 Pet. 24–41. To establish anticipation, Microsoft must prove that each and every element in a claim is found in SWEB, arranged as recited in the

claim. *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1369 (Fed. Cir. 2008); *Karsten Mfg. Corp. v. Cleveland Golf Co.*, 242 F.3d 1376, 1383 (Fed. Cir. 2001). We first summarize the disclosure of SWEB, followed by analysis why Microsoft has not established by a preponderance of the evidence that SWEB anticipates any of claims 12–17, 20–32, 34–46 or 48.

1. Disclosure of SWEB

SWEB discloses a “Scalable Web server” (SWEB) built on a network of computers, which permits distributed processing of Web page requests. Ex. 1009, 1–2. A scheduler on the network monitors the usage of system resources, and routes requests to computers on the network accordingly. *Id.* at 8–9. The architecture of the SWEB system is shown in Figure 2:




Figure 2: The computing and storage architecture of SWEB.

Figure 2 of SWEB shows client computers on the Internet, which make HTTP requests to SWEB having scheduler and load information.

The scheduler can be centralized on one processor such that all requests go through it, or it can be distributed across all the computers on the

network. *Id.* Figure 6 of SWEB, reproduced below, shows the operation of one system having a distributed scheduler:




Figure 6: The URL redirection.

Figure 6 of SWEB depicts a system comprising client computer C, a DNS server, and servers S0 and S1.

According to Figure 6, client computer C sends a request for a file “myfile” located at URL *r* that is routed by the DNS server to server S0. *Id.* at 10–11. If the distributed scheduler determines that the request is best processed by server S1, S0 returns a rewritten URL *r'* to the client directing the client to obtain “myfile” from server S1. *Id.* Client computer C sends the redirected URL *r'* to S1, which then processes the request and returns file *f*. *Id.* SWEB states that the URL redirection scheme was chosen based on its “simplicity of implementation and universal compatibility,” but also notes that it has the “added overhead of an additional connect/pass request/parse/respond cycle after the redirection occurs.” *Id.* The reference states that “[m]ost Net browsers and clients automatically query the new location, so redirection is virtually transparent to the user.” *Id.* at 10.

Illustrated by Figure 6, and using Microsoft’s analysis of claim 12 as an example, Microsoft’s argument that SWEB discloses all elements of the

claims is as follows. The *dynamic Web page generation request* of the claim is SWEB's *r*, which requests the file "myfile" located on server S0. Tr. 14–15. Microsoft identifies an "SWEB server that receives . . . a user request," such as S0, as the *Web server*. Pet. 25.⁹ Microsoft argues that the *routing* step of the claim, which requires routing the request to a page server, is met by the URL redirection of the request from server S0, back to client computer C, and then to a second server S1 which then fulfills the request. Pet. 25–27. S1 is, therefore, the *page server* of the claims, and its processing of *r'* releases the S0 to process other requests. *Id.* The *dispatcher* limitation of claim 12 is said to be met by the distributed scheduling algorithm running on S0, which intercepts the request *r*, rewrites the URL, and returns *r'* to the client computer. *Id.* at 28–29. The processing server S1 then retrieves a file or invokes a CGI program in order to respond to the request, which Microsoft contends satisfies the step of *dynamically generating a Web page*. *Id.* at 29–30.

2. Analysis of Anticipation

As was illustrated above, the primary distinction between the server system described in SWEB and that disclosed in the specification of the '554 patent is one of architecture. SWEB discloses a distributed structure, in which each of the servers in a network may act as the initial Web server that receives the request, and any of the other servers may be the page server that provides content to the client computer. In addition, SWEB relies on URL

⁹ For simplicity, we cite only to the Petition in this case; similar arguments were made in the Petition in IPR2015-00484, subsequently consolidated with this proceeding.

redirection as the method of routing requests from the Web server to the page server, whereby the client computer is instructed to request the content from the page server directly, rather than communicating through the Web server. As stated in SWEB, “[t]he conceptual model is that of a very short reply going back to the client browser, who then automatically issues another request to the new server address.” Ex. 1009, 14.

By contrast, the ’554 patent discloses a system that is hierarchical in nature, wherein a single Web server receives all incoming requests, and then routes those requests (through the use of an interceptor/dispatcher) directly to a page server for processing. In some embodiments, for example as claimed in dependent claim 33, the content produced by the page server is then routed back through the Web server to the requesting client.

Although the described systems differ in their structures, this alone is not determinative of our inquiry, as we compare the prior art system to the ’554 patent claims, not the disclosure of the specification. The question before us, therefore, is whether the distributed structure of the SWEB system—and its reliance on URL redirection—removes it from the scope of the claims of the ’554 patent. We conclude that it does, for the following reasons.

Parallel argues that the claims of the ’554 patent require a single *request*, and distinguishes SWEB as involving multiple requests. PO Resp. 21–25. As to the former, we agree. Using claim 12 as an example, the preamble recites *a dynamic Web page generation request*, then goes on to only refer to this single request. *Said request* is first routed to a selected page server, wherein the routing involves intercepting *said request*, routing *said request* to a dispatcher, and dispatching *said request* to a selected page

server. The dispatching further includes examining *said request*, selecting a page server, and sending *said request* to the selected page server. The page server then processes *said request*, and dynamically generates a Web page in response to *said request*.

Clearly, the claims of the '554 patent are directed to the routing and processing of a single *request*, the same request as that received by the Web server from the requesting client. Parallel argues that while Figure 6 of SWEB shows a request “r” from client computer C to Web server S0, all that is returned is a rewritten URL in the form of r' that instructs client computer C to obtain the information from server S1. PO Resp. 24–25. The original request itself is not returned. According to Parallel, the client computer C must then make a *new* HTTP request for the information from S0. *Id.* at 25. What is received and processed by the page server S0 in SWEB, therefore, is not the original request (*said request*), Parallel argues. *Id.* at 27.

Microsoft responds that the URL redirection of SWEB is merely a redirection of the original request by the client computer. Reply 3. Regardless of the fact that the URL of the request may have been rewritten, Microsoft argues, the “ask” or “demand” of the request remains the same; in the context of Figure 6, both the original request and the redirection seek the same “myfile,” they merely seek it from different locations. *Id.* at 4. Microsoft emphasizes that the redirection happens automatically at the client computer, and contends that a person of ordinary skill in the art at the time of the invention would have understood a redirected URL to remain the same request as the original. *Id.* at 5.

Both parties cite various documents contemporaneous with the filing of the '554 patent, arguing that they evidence whether a person of ordinary

skill in the art would have understood URL redirection to be a single request or multiple requests. For example, Microsoft cites the February 1996 HTTP specification as referring to a URL redirection as “*the* reformatted request” (Ex. 2010, 6) and an article by Derler, which states that redirection causes the client to “automatically *retransmit*[] its request” (Ex. 1043, 52).

Reply 4–5 (emphases added). Parallel points out, on the other hand, that Derler also states that a redirection using status code 302 “provides the modified URL to be used with *further* requests” (Ex. 1043, 52), and that SWEB refers to the URL redirection concept as a “reply going back to the client browser, who then automatically issues *another* request to the new server request.” (Ex. 1009, 14). Tr. 64–65 (emphases added).

As both parties can find selective quotations from these documents that support their position, we do not find such evidence to be of much help. Nor is the testimony of the parties’ experts, who are in direct opposition on this point, useful. Ex. 1083 ¶ 16 (Mitzenmacher: “[T]he person of ordinary skill in the art at the time of the alleged invention would understand that a redirected request is the same request as the original request, not a different request.”); Ex. 2091 ¶ 74 (Jones: “two different requests . . . is the result of [SWEB’s] specific choice to utilize a distributed Web server and URL redirection.”). In any event, our inquiry is not how the skilled artisan would understand a URL redirection in a vacuum, but rather how she would understand *request* as used in the ’335 patent, and whether she would consider SWEB’s system to involve a single request.

At oral argument, the parties and the Board analogized to a letter that is mailed, returned to sender, and then re-addressed, and whether such modifications to the envelope changed the letter itself. Tr. 56–59, 87–88. In

view of the problem faced by the '554 patent, however, this analogy asks the wrong question. The relevant question is not whether a URL redirection, which in effect re-addresses the same “ask,” creates a new request.¹⁰ Rather, to us, the distinction is that SWEB—by involving the client computer again, and requiring that computer to take further action—invokes a new action by the original requestor that is again sent to a server in the SWEB network. In the context of the '554 patent, this is a new *request*, for the following reasons.

The specification of the '554 patent consistently uses the term *request*, or, in some cases, *URL request*, in the context of the client asking for information. *See, e.g.*, Ex. 1001 at 1:30–32 (“A Web browser allows a Web client to request a particular Web page from a Web site by specifying a Uniform Resource Locator (URL).”); 3:65–66 (“Web client 200 can make URL requests to Web server 201”); 4:13–14 (“the Web client makes a URL request”). It is clear to us, therefore, that it is the act of *asking*—an action taken by the client computer—that defines a request. In URL redirection, the client computer must act a second time by making a second “ask” from the rewritten URL, even if the content being requested is the same as the original.

The parties’ agreed-upon construction of *request*, which we have adopted—“a message that *asks* for a Web page”—also incorporates this link between a request and an act of asking. On cross-examination, Mr. Mitzenmacher’s testimony reflected this link, testifying that the same

¹⁰ Indeed, we note that claim 27 of the '554 patent, which depends from claim 12, refers to “a uniform resource locator *associated with* the request,” implying that the URL is distinguishable from the request itself.

“ask” made at different points in time by the same user (for example, a user refreshing a web page) would constitute different *requests*:

So, again, those might be instances where those requests could have the same format but they’re coming to the server at different points in time. So I would have to consider it. But typically, I would think of that as being *distinct requests* in the sense that they’re distinct – you know, starting from the point that *they’re distinct initiation actions from the user, which would correspond to multiple requests for information*. I haven’t considered these particular scenarios in light of the claims, so I may be missing something.

Ex. 2107, 74–75 (emphasis added). To us, it does not matter in the context of the ’554 patent whether this subsequent “ask” is initiated by the user by hitting “refresh,” or by the client computer automatically through a URL redirection. Both are, in the words of Microsoft’s expert, “distinct initiation actions.”

Not only is this understanding of *request* consistent with its usage in the specification and the parties’ agreed construction, it is also consistent with the purposes of the ’554 patent itself. The ’554 patent is concerned with the “problem of managing numerous Web pages and requests at a Web site,” and states that “[m]anaging . . . hundreds of Web pages in a coherent manner and processing all requests for access to the Web pages is thus a difficult task.” Ex. 1001, 2:2–3. Elsewhere, the specification discusses management of Web sites as a “crucial need” that involves managing “thousands of requests or ‘hits’ in a single day.” *Id.* at 4:36–39. The solution of the ’554 patent is to lessen the load of these thousands of requests on a Web server, by distributing processing of the request to various page servers.

To interpret multiple “asks” for information by a client computer as a single *request* would contravene these purposes, by creating even more load on the servers in the network. As SWEB itself recognizes, “[t]he primary disadvantage of URL redirection in practice is the added overhead of an additional connect/pass request/parse/respond cycle after the redirection occurs.” Ex. 1009, 11. In other words, URL redirection may increase the number of “asks” by client computers up to twofold. This increased load on the servers is contrary to the primary problem the ’554 patent attempts to solve. We are hesitant to conclude that a person of ordinary skill in the art would have understood *request*, as used in the ’554 patent, to encompass multiple “asks” by a client computer, when these “asks” would exacerbate the very problem addressed by the patent.

Finally, we have reviewed the evidence cited by Microsoft in support of its position, and are unpersuaded. For example, Microsoft cites to an article in which one of the inventors of the ’554 patent, Mr. Lowery, is quoted as stating “the patent covers the notion of redirecting that request to other machines.” Reply 8 (quoting Ex. 1089, 1). Microsoft has taken this quote out of context. When read in context, it is clear Mr. Lowey is discussing the redirection of requests from the Web server to page servers, *not* URL redirection of a request back to a client computer. Ex. 1089, 1 (“[T]he patent covers the notion of redirecting that request to other machines [application servers] *within a cluster* and having those other machines dynamically produce a response.”) (emphasis added). We do not find this statement by Mr. Lowery to support the argument that *request* as used in the ’554 patent encompasses URL redirection.

Similarly, we do not find persuasive Microsoft's out-of-context quotation of the cross-examination testimony of Mr. Howell, another named inventor of the '554 patent. While Microsoft quotes the exchange as "Q. 'One way the inventions claimed in the '554 and '335 patents route requests to other servers is redirection; is that correct? A. Yes.'," (Reply 8), the deposition transcript reveals that the question posed to Mr. Howell *expressly* excluded URL redirection. Ex. 1086, 12 ("Q. Let me just drop the URL redirection, so we just deal with that without that. . . . One way the inventions claimed in the '554 and '335 patents route requests to other servers is redirection; is that correct? A. Yes.").

Finally, Microsoft relies on an excerpt of the confidential deposition transcript of Dr. Jones from the copending District Court action between the parties, in which he discusses a Microsoft server module called "URL rewrite." Reply 8–9; Ex. 1097. Dr. Jones testifies that while some of the characters of a request may be rewritten by "URL rewrite," "it is not going to make that a request for something else." Ex. 1097 at 54:1–5; 58:20–24. Microsoft argues that this is an acknowledgement that URL redirection does not result in a new request. Reply 8–9. We disagree. While the nature of "URL rewrite" is not fully explained in the deposition excerpt provided by Microsoft, there is further discussion in the excerpt of Dr. Jones' expert report from the same case. Ex. 1098. In his report, Dr. Jones states that "[t]he URL rewrite module receives the request within the IIS web server 'request pipeline,' . . . and sends the request down the line toward the ARR module, where it can be routed to a specific server for processing." *Id.* ¶ 17. From this explanation, it is clear that "URL rewrite" internally routes a request within a Microsoft server, and is not a URL redirection like that

described in SWEB, which requires a client computer to make another “ask.” Dr. Jones’ statement that the “URL rewrite” module “is not going to make that a request for something else” has no bearing on whether a skilled artisan would understand a URL redirection to result in a new *request* in the context of the ’554 patent.

With the foregoing understanding of how the ’554 patent uses *request*, we must conclude that SWEB does not anticipate claims 12–17, 20–32, 34–46 and 48. All independent claims require that a *request* be intercepted at the Web server, and dispatched by a dispatcher to the page server. The claims also require the same *request* to be processed by the page server. In SWEB, however, a first request is intercepted and dispatched, but it is a second request—the URL redirect—that is processed by the page server.

In addition, we note that some of the independent claims—in particular, claims 12 and 20—require that the dispatching step include “sending said request to said selected page server.” As SWEB was applied to the claims by Microsoft, the dispatcher resides at S0, the Web server. Tr. 89. It is clear, however, that in SWEB the dispatcher at the Web server *sends* nothing to the page server; rather, everything sent by the Web server is back to the client computer C. Even if we were to accept Microsoft’s argument that a URL redirect is the same *request* in the context of the ’554 patent, claims 12 and 20 would not be anticipated, because what *sends* the request *r*’ to the page server S1 in SWEB is the client computer C, not the dispatcher/web server S0.

For these reasons, we find that Microsoft has not proven, by a preponderance of the evidence, that SWEB discloses a single *request* through all steps of the claims, as required by independent claims 12, 20, 32,

34, 46, and 48. Nor do we find that Microsoft has proven, with respect to claims 12 and 20, that SWEB discloses a dispatcher that “send[s] said request to said selected page server.” By extension, Microsoft has not shown that all elements of dependent claims 13–17, 21–31, and 35–45 are disclosed by SWEB. We, therefore, conclude that Microsoft has not proven that SWEB anticipates claims 12–17, 20–32, 34–46 and 48.

C. Obviousness over SWEB

Microsoft asserts that claims 12–49 are unpatentable under 35 U.S.C. § 103 as having been obvious over SWEB. Pet. 42–49; 484 Pet. 41–52. The Petition proposes various modifications that a person of ordinary skill in the art might have made, to address potential arguments that Parallel might make that various elements of the claims are not disclosed by SWEB. *Id.* As our analysis above found SWEB to be lacking disclosure of a single *request* through all steps of the claims, we focus on the portion of the Petition addressing this modification. *See id.* at 42–43.

Microsoft argues that, even if SWEB does not disclose *routing said request from said Web server to a selected page server*, it would have been obvious to include such functionality in SWEB’s scheme. Pet. 42. Citing Mr. Mitzenmacher, Microsoft contends that “[i]t was known in the art to route requests from one machine to another via request forwarding based on a determination of which server should be used to process such a request.” *Id.* (citing Ex. 1007 ¶ 214). Such forwarding had the known advantage of freeing up CPU cycles and memory resources of the sending device. *Id.* According to Microsoft, an “SWEB[] system modified to include request forwarding would indisputably be receiving, intercepting, routing,

dispatching and processing the *same* request, and so satisfy all elements of the claims” even if we accept Parallel’s argument that URL redirection is not a single request. Reply 10.

Parallel argues that a single *request* is not obvious over SWEB, because the architecture of SWEB is fundamentally different, “necessarily requiring the generation of a second dynamic Web page generation request at the Web client or browser in order for a dynamic Web page to be generated.” PO Resp. 41. Parallel notes that while Microsoft argues that routing requests from one machine to another was known, a person of ordinary skill in the art would not have been motivated to apply this knowledge to SWEB given SWEB’s explicit use of URL redirection. *Id.* While Parallel admits that SWEB discusses the possibility of a “UNIX sockets package” alternative to URL redirection, it notes that this alternative was dismissed as “difficult since it requires a substantial modification of the UNIX operating system kernel.” *Id.* at 41–42. According to Parallel, this signifies that the modification proposed by Microsoft was non-trivial and may have been beyond the level of ordinary skill in the art. *Id.* at 42.

Although we believe Parallel’s arguments have merit, our decision rests primarily on the fact that Microsoft has not articulated a sufficient reason why a person of ordinary skill in the art would have modified SWEB to include request forwarding. Even accepting Mr. Mitzenmacher’s assertions that request forwarding was known, and had certain known advantages (Ex. 1007 ¶ 214), neither he nor Microsoft explains *why* a person of ordinary skill in the art would have replaced the URL redirection scheme of SWEB with request forwarding. The advantages highlighted by Mr. Mitzenmacher—that request forwarding frees up CPU cycles and memory

resources of the forwarding computer—are the same advantages provided by URL redirection, which also frees up the resources of the forwarding computer. Neither Microsoft nor Mr. Mitzenmacher provide sufficiently persuasive evidence that request forwarding provides any advantage not already present in SWEB, other than it would meet limitations of the claims of the '554 patent. This type of hindsight-based reasoning is prohibited in an obviousness analysis.

For this reason, we find that Microsoft has not provided sufficient evidence that a person of ordinary skill in the art would have modified the system of SWEB in the manner proposed. We, therefore, cannot conclude that claims 12–49 would have been obvious over the disclosure of SWEB.

D. Remaining Obviousness Grounds

We instituted trial on two additional obviousness grounds: claims 14, 15, 18, 19, 33, and 47 as having been obvious over SWEB and Leaf, and claims 17, 29, 43, and 49 as having been obvious over SWEB and Bradley. Microsoft relies on Leaf¹¹ to disclose various features recited in the dependent claims, such as a connection cache or custom HTML extension templates. Pet. 54–59. Microsoft cites Bradley as teaching maintaining a page cache. Pet. 59–60. Microsoft does not assert that either of these references remedy the deficiency of SWEB identified above; namely, that SWEB involves processing a request that is different than the request

¹¹ Because we conclude that SWEB in combination with Leaf does not disclose all elements of the challenged claims, we need not reach Parallel's argument that Leaf is not prior art to the '554 patent because of antedation. PO Resp. 49–50.

originally received by the Web server. We, therefore, must conclude that Microsoft has not met its burden on either asserted ground of obviousness.

E. Motions to Exclude

As summarized above, both parties filed motions to exclude. Microsoft's Motion seeks to exclude certain evidence submitted by Parallel to antedate the Leaf reference (Exs. 2015–2019, 2021, 2022, 2031, 2032, 2066, 2073, 2097, 2099, 2100, and 2092 ¶¶ 3–43), or to establish objective indicia of nonobviousness (Exs. 2024–2028, 2031, 2032, 2034, 2042–2055, 2059, 2063, 2065, 2068, 2069, 2071–2078, 2080–2085, 2094–2098, and 2092 ¶¶ 44–59). Paper 64, 2–10. Microsoft also seeks to exclude some of these exhibits (Exs. 2015–2016, 2018, 2021–2022, 2031–2032, 2063, 2065–2066, 2069, 2071–2075, 2077–2078, 2080, 2084–2085, and 2095–2100) as lacking authentication. *Id.* at 11–12. Finally, Microsoft objects to a portion of Mr. Lowery's testimony pertaining to objective indicia as improper lay testimony. *Id.* at 12–13. As we find that SWEB did not disclose all elements of the claims as relied upon by Microsoft, we need not reach Parallel's attempted antedation of Leaf or its evidence of objective indicia, and thus our Decision does not rely on these exhibits. We, therefore, dismiss Microsoft's Motion to Exclude as to these exhibits as moot.

Microsoft also moves to exclude Exhibits 2014, 2020, 2023, 2029, 2030, 2033, 2035–2041, 2060–2062, 2056–2058, 2064, 2067, 2070, 2079, and 2088 as irrelevant, as they were not cited in any paper in the proceeding. Paper 64, 11. Parallel concedes that these exhibits may be excluded. Paper 68, 1. We agree, grant Microsoft's Motion to Exclude as to these exhibits, and will expunge them from the record of the case.

Parallel moves to exclude various exhibits (Exs. 1048, 1070–1071, 1087–1089, 1096, and 1101) as lacking authentication and, in some cases, as containing hearsay. Paper 62, 3–8. Exhibit 1048 is the Bradley prior art reference, which we did not rely on in this Decision as we found the disclosure of SWEB lacking. Similarly, Exhibits 1070 and 1071 are documents submitted by Microsoft to support the public availability of Bradley. Parallel’s Motion to these Exhibits is dismissed as moot.

Exhibits 1087–1089, 1096, and 1101 are documents submitted by Microsoft to establish that a *request* in the context of the ’554 patent would include URL redirected requests. We considered these documents, yet still found in favor of Parallel on the relevant issues. Exclusion of the documents, therefore, would not alter the ultimate conclusion of this Decision in favor of Parallel. Parallel’s Motion as to these exhibits is dismissed as moot.

Finally, Parallel objects to portions of Mr. Mitzenmacher’s declarations as “highly prejudicial” as relying on hearsay, and seeks their exclusion under Federal Rule of Evidence 403. Paper 62, 8–9. Proceedings before the Board, however, are not jury trials; in the absence of a jury, the risk of unfair prejudice against which Rule 403 guards is diminished, if not eliminated entirely. *See, e.g., Schultz v. Butcher*, 24 F.3d 626, 632 (4th Cir. 1994) (“in the context of a bench trial, evidence should not be excluded under 403 on the ground that it is unfairly prejudicial”); *see Gulf States Utils. Co. v. Ecodyne Corp.*, 635 F.2d 517, 519 (5th Cir. 1981). The Board is capable of hearing relevant evidence and weighing its probative value, including determining the weight to be given Mr. Mitzenmacher’s

testimony. We, therefore, deny Parallel's Motion as to the Mitzenmacher declarations.

III. CONCLUSION

We conclude that Microsoft has failed to demonstrate, by a preponderance of the evidence, that claims 12–49 of the '554 patent are unpatentable.

IV. ORDER

Accordingly, it is

ORDERED that claims 12–49 of U.S. Patent No. 5,894,554 have *not been proven unpatentable*;

FURTHER ORDERED that, pursuant to 35 U.S.C. § 318(b), upon expiration of the time for appeal of this decision, or the termination of any such appeal, a certificate shall issue confirming claims 12–49 of U.S. Patent No. 5,894,554;

FURTHER ORDERED that Microsoft's Motion to Exclude (Paper 64) is *granted-in-part* and *dismissed-in-part*, as set forth above;

FURTHER ORDERED that Parallel's Motion to Exclude (Paper 62) is *denied-in-part* and *dismissed-in-part*, as set forth above; and

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

IPR2015-00483
Patent 5,894,554

PARTIES AND BOARD ONLY

FOR PETITIONER:

Jason Wolff
wolff@fr.com

Michael Hatcher
iprnotices@sidley.com

Joseph Micallef
iprnotices@sidley.com

Jeffrey Kushan
iprnotices@sidley.com

FOR PATENT OWNER:

Pierre Hubert
phubert@mckoolsmith.com

Christopher Bovenkamp
cbovenkamp@mckoolsmith.com