
In the
UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NetNut Ltd.
Petitioner,

v.

Bright Data Ltd.
Patent Owner.

Case IPR2021-00465

U.S. Patent No. 9,742,866

**PETITIONER’S OPPOSITION TO PATENT OWNER’S REVISED
NON-CONTINGENT MOTION TO AMEND**

Contents

1	INTRODUCTION	1
2	UNFULFILLED STATUTORY AND REGULATORY REQUIREMENTS	2
2.1	Respond to a ground of unpatentability involved in the trial/“Good Cause”	2
2.2	New Matter/Lack of Adequate Written Description	4
2.2.1	“a first request . . . indicating that the content slice is to be fetched directly from the first server”	4
2.2.2	“(c) in response to receiving the first request, sending the content slice identifier from the selected tunnel device to the first server using the second identifier;”	6
2.2.3	“the selected tunnel device does not cache a copy of the content slice”	6
2.3	Duty of Candor	8
3	UNPATENTABILITY	8
3.1	Lack of Enablement (§ 112(a)) and/or Utility (§ 101)	8
3.2	Claim Construction and Indefiniteness	9
3.2.1	“so long as”	9
3.2.2	“the first request . . . indicating that the content slice is to be fetched directly from the second server”	10
3.2.3	“idle resource”	10

3.3	Obviousness	10
3.3.1	Obviousness of revised claim 29 over Shribman in view of Durso (or any one or more of Chauhan, Liu, MorphMix, Sigurdsson, Mithyantha, or Jackowski) . .	11
3.3.2	Claims 30-38	19
4	CONCLUSION	20

EXHIBIT LIST	
1001	United States Patent No. 9,742,866 to Shribman <i>et al.</i>
1002	Excerpts from Prosecution History of the '866 patent (s/n 14/930,894)
1003	Expert Declaration of Keith J. Teruya
1004	CV of Keith J. Teruya
1005	Appendix of Challenged Claims
1006	Postel, RFC 791, Internet Protocol, DARPA Internet Program Protocol Specification (Information Sciences Institute, University of Southern California, September 1981)
1007	Fielding, <i>et al.</i> , RFC 2616, Hypertext Transfer Protocol – HTTP/1.1, (Internet Engineering Task Force, Network Working Group, June 1999)
1008	Reserved
1009	Reserved
1010	Exhibits considered by Keith J. Teruya
1011	Excerpts from Webster's Third New International Dictionary, ISBN: 978-0-87779-789-0 (2017) (Cover, Copyright Page, Definitions of "include" and "website")
1012	Wessels, "Squid: The Definitive Guide," ISBN-10: 9780596001629, ISBN-13: 978-0596001629, O'Reilly Media; 1st Ed. (January 1, 2004)
1013	Reserved
1014	Luotonen, "Web Proxy Servers," ISBN-10: 0136806120, ISBN-13: 978-0136806127, Prentice Hall; 1st Ed. (December 30, 1997)
1015	Reserved
1016	Reserved
1017	U.S. Patent Publication No. 2011/0087733 to Shribman <i>et al.</i> ("Shribman"), Published April 14, 2011
1018	European Patent Publication EP 2 597 869 A1 by Sharp KK, Published May 29, 2013
1019	Reserved

Exhibit List (Continued)	
1020	Cooper <i>et al.</i> , RFC 3040, Internet Web Replication and Caching Taxonomy (January 2001)
1021	Reserved
1022	Reserved
1023	Claim construction order, Luminati Networks Ltd. v. UAB Tesonet, No. 2:18-CV-299 (E.D. Tex.), slip op. (D.I. 121, Aug. 20, 2019)
1024	Reserved
1025	Reserved
1026	Reserved
1027	ISO/IEC 23009-1:2012(E), MPEG-DASH standard, Jan. 5, 2012
1028	Reserved
1029	Microsoft Computer Dictionary (5th Ed. 2002), definitions of Web site and Web page
1030	Reserved
1031	Reserved
1032	Claim construction order, Luminati Networks Ltd. v. BIScience Inc., No. 2:18-CV-483 (E.D. Tex.), slip op. (D.I. 130, Dec. 6, 2019)
1101	Expert Declaration of Keith J. Teruya in Support of Petitioner's Opposition to Patent Owner's Non-Contingent Motion to Amend
1102	CV of Keith J. Teruya
1103	Exhibits considered by Keith J. Teruya
1104	U.S. Patent No. 8,560,604 to Shribman <i>et al.</i>
1105	U.S. Patent No. 6,266,704 to Reed <i>et al.</i>
1201	U.S. Patent No. 10,469,614 to Shribman <i>et al.</i>
1202	Reexamination request in reexamination application ser. no. 90/014,880
1203	Reexamination grant in '880 reexamination
1204	First office action in '880 reexamination
1205	Grant of extension request in '880 reexamination
1206	Redline comparison of revised amended claims against first amended claims

Exhibit List (Continued)	
1207	Expert Declaration of Keith J. Teruya in Support of Petitioner's Opposition to Patent Owner's Revised Non-Contingent Motion to Amend
1208	U.S. Patent Pub. No. 2010/0262650 A1 by Chauhan <i>et al.</i>
1209	U.S. Patent Pub. No. 2010/0082513 A1 by Liu
1210	U.S. Patent Pub. No. 2007/0174246 A1 by Sigurdsson <i>et al.</i>
1211	U.S. Patent Pub. No. 2011/0066924 A1 by Dorso <i>et al.</i>
1212	U.S. Patent No. 8,972,602 to Mithyantha
1213	U.S. Patent Pub. No. 2013/0304796 A1 by Jackowski <i>et al.</i>
1214	MorphMix—A Peer-to-Peer-based System for Anonymous Internet Access, by Marc Renhard (2004)
1215	Berz declaration, printed publication support for Morphmix
1216	U.S. Patent Publication No. 2013/0080575 to Prince <i>et al.</i> ("Prince"), Published March 28, 2013
1217	Printout of VIP72 Youtube web page at https://www.youtube.com/watch?v=L0Hct2kSnn4 , retrieved Nov. 21, 2019
1218	VIP72 Scene Images extracted from VIP72.com/nvpnnet , MPEG-4 video recording of "nVPN.net Double your Safety and use Socks5 +nVpn", accessed from https://www.youtube.com/watch?v=L0Hct2kSnn4 , published September 11, 2011
1219	Declaration of Ronald Abramson authenticating Ex. 1218 Scene Images
1220	Certification dated Nov. 8, 2019 of Anjali Shresta of Google, Proof of Date for VIP72 Youtube web page and video
1221	Rennhard declaration, printed publication support for Morphmix
1222	U.S. Pat. No. 4,403,286 to Fry <i>et al.</i>
2020	Appendix A-1: Written Description Support
2021	Appendix A-2: Proposed Substitute Claims
2022	61870815 Provisional - specification
2023	61870815 Provisional - drawings

Exhibit List (Continued)	
2024	14468836 Nonprovisional - abstract
2025	14468836 Nonprovisional - specification
2026	14468836 Nonprovisional - drawings
2027	14468836 Nonprovisional - claims
2028	Revised Appendix A-1 Written Description Support
2029	Revised Appendix A-2 Revised Proposed Substitute Claims

1. INTRODUCTION

Patent Owner's ("PO's") revised motion to amend ("RMTA") would limit the client device's receipt of an identifier for a (first) tunnel device, "so long as the first tunnel device has an idle resource that satisfies a criterion," thus bringing in the further factor of basing tunnel device selection on a measurement of the device's availability.

What the RMTA fails to mention, however, is that adding such a resource-based or "idling" limitation would transform the claims to closely track the claims of a related patent, U.S. Pat. No. 10,469,614 (Ex. 1201 hereto), *e.g.*, claims 1, 20, and particularly with regard to the '866 patent, claim 29, which also recite brokering of content retrieval proxies (including piece-wise retrieval in claim 29), where the proxies are likewise selected for service based on their state of resource utilization or idling. Significantly, though not noted in the RMTA, the '614 patent is currently under reexamination (no. 90/014,880), in which all challenged claims stand rejected, under prior art going to the resource utilization/idling limitations. *See* Ex. 1202 (reexamination request), Ex. 1203 (grant), and Ex. 1204 (office action).

The proposed amendment raises a host of procedural and disclosure issues as a result of the foregoing, in addition to obviousness over the same references on which the '614 patent stands rejected.

In addition to the foregoing, which reflects those problems in common with the current state of the '044 patent, the proposed amendments to the '866 patent also suffer from PO's effort to force the same edits that it made for the

'044 patent onto the claim set of the '866 patent, which was premised on a different data flow, *i.e.*, content partitioning and peer-to-peer distribution. This has resulted in a clash of claim terms raising additional issues under § 112 as well as § 101 utility.

The following addresses the most significant issues with the RMTA.

2. UNFULFILLED STATUTORY AND REGULATORY REQUIREMENTS

2.1. Respond to a ground of unpatentability involved in the trial/“Good Cause”

Since *Lectrosonics, Inc. v. Zaxcomm, Inc.*, IPR2018-1129, Paper 15 (Precedential) (2019), and the Pilot Program Notice at 84 FR 9497 (“PPN”), PTAB amendment practice has often taken as a given that any amendment that goes to a ground of the IPR challenge thereby “respond[s] to a ground of unpatentability involved in the trial.” Further, to the extent a revised MTA might be considered more than the “1 motion to amend” provided for under 35 U.S.C. §316(d), statements at page 9497 of the PPN suggest that “good cause” under 37 CFR 42.121(c) may automatically be satisfied where a revised MTA follows a petitioner’s opposition (“Opposition”) and/or the Board’s preliminary guidance (“PG”). The circumstances of the present case, however, put into question such easy answers to this threshold issue under *Lectrosonics*.

PO makes only the same cursory statement in support of a revised MTA that it would make for an initial amendment, as required by 37 CFR 42.121(a)(2)(i), *i.e.*, that “Patent Owner . . . submits the RPSC to further narrow the claimed

methods,” to overcome issues of novelty and obviousness that have been raised. RMTA at 5.

Petitioner submits that, where an RMTA introduces new features beyond the scope of what was addressed in the Opposition or PG (here, the new “so long as the tunnel device has an idle resource” limitation), a bare recitation that the amendment is narrowing is insufficient to support the motion.

The PPN also states: “A revised MTA may not include amendments, arguments, and/or evidence that are *unrelated to issues raised in the preliminary guidance or the petitioner’s opposition to the MTA.*” PPN at 9501 (emphasis added). The PPN referenced public comments expressing concern that the proposed RMTA procedure might “read out” the requirement for good cause. In response, the PPN noted: “the revised motion to amend is not a second motion to amend per se, but rather *a revised version* of the initially-filed motion to amend.” *Id.* at 9504 (emphasis added).

The non-contingent revised MTA here seeks to introduce an entirely new technical feature, at a late stage, in an attempt to inject a basis for patentability where the prior effort to do so by a lesser amendment failed. PO seeks to justify this on the same basis used to satisfy 37 CFR 42.121(a)(2)(i), which also applies (*i.e.*, that the revised MTA responds to a ground of unpatentability involved in the trial). PO does not address why the current subject matter could not have been included in the initial MTA.

If the Board accepted as “good cause” the limited rationale offered here, it would be relying on the *same* showing that would have been necessary for *any* amendment (“revised” or not), under 37 CFR 42.121(a)(2)(i)—thereby

effectively eliminating the *further* good cause requirement under the 37 CFR 42.121(c). The PPN cannot override a formal regulation, which still binds the agency. *See Accardi v. Shaughnessy*, 347 U.S. 260 (1954). Thus, the PPN’s blanket “good cause” provision should only be applied where the RMTA is limited to a scope reasonably commensurate with the issues addressed in the Opposition and/or the PG. Where (as here) the RMTA inserts a new limitation that goes beyond that baseline, a further showing of good cause (beyond the blanket authorization of the PPN) should be required, and it is lacking here.

2.2. New Matter/Lack of Adequate Written Description

Once again, added claim features must find written support in the original disclosure of the application.¹ (Some new matter issues herein depend on claim construction, as to which, see page 9.)

2.2.1. “a first request . . . indicating that the content slice is to be fetched directly from the first server”

This limitation is not supported by the written description (which follows from the fact that it also makes no sense).

¹Regarding the PG at 6, Petitioner respectfully again submits that the provisional may not serve as a source of written description, as opposed to establishing an earlier priority date. The provisional here fails to fill any gaps left by the non-provisional specification. Should the Board have reason to look to the provisional for written description support, Petitioner asks that it reconsider this issue.

The preamble (as amended) recites a method for fetching “fresh content” from a first server identified in the Internet, and then partitioning “the” fresh content, *i.e.*, the fresh content thus fetched from the first server, into a plurality of content slices.

But it then turns around and recites requests that the content slices are to be “fetched *directly* from the first server.” This makes no sense, and even if some logic could be twisted out of it, has no written description support.

PO points to nothing that supports a content request from a requesting client (first device) to a tunnel device, in which the request itself “indicates” (whatever that may mean) a *content slice* is supposed to be fetched “directly” from an origin server (first server). All that PO points to in Ex. 2025 are requests identifying content without indicating partitioning. Moreover, nothing about the disclosed requests come with any kind of instruction or other indication to the tunnel to fetch the content directly.

The Opposition at 5 previously pointed out (in connection with “complete URL” in the prior amendment on the ’044 patent), that the ’866 specification “does not disclose the format of the claimed first request itself, *i.e.*, the content request from the client device to the tunnel device.” There is no disclosure, then, of how the “first request” could also be somehow “indicating” that the tunnel is supposed to do a direct fetch, as opposed to some indirect retrieval, *e.g.*, from a peer cache. Indeed, retrieval from peer cache appeared to be the original thrust of these claims, but clearly that has all changed, with this latest “direct fetch” wording. The problem is that this total redirection of the claims has zero support in the written description.

PO also refers to direct fetches by the client to the server (without a tunnel), but this is irrelevant.

PO also cites to the summary of the invention, which recites the original claims, but since this does not include the current amendment, this is to no avail (also not enabling in any case).

2.2.2. “(c) in response to receiving the first request, sending the content slice identifier from the selected tunnel device to the first server using the second identifier;”

This makes sense if the “first server” is a peer device that has previously retrieved content slices cached. It make no sense in juxtaposition with step (b), which as currently sought to be amended would now require a direct fetch from the origin server. PO points to no disclosure of a content slice retrieval directly from an origin server.

2.2.3. “the selected tunnel device does not cache a copy of the content slice”

PO’s amendment is again seeking to turn this claim on its head, without written description support. There is repeated disclosure of the tunnel device caching content (*e.g.*, Ex. 2015, 140)—which could make sense if the tunnel device partitioned the content into slices—but there is no disclosure of a tunnel device as used in the disclosure, which does *not* cache.

PO points to disclosure that content “may be stored” – for example if it was previously retrieved. *E.g.*, Ex. 2025, pp. 94, 117. But the disclosure that the tunnel device may not have cached particular requested content does not support a claim to a tunnel device that “does not cache” content. Declaration

of Keith J. Teruya, Ex. 1207 at ¶ 5. Disclosure of a mere “cache miss” does not support a claim for a non-caching device.

Indeed, the claims of the ’866 patent specifically address piece-wise retrieval of content through intermediary devices, and make no sense if content is not cached by the intermediate devices. There is no disclosure of piece-wise retrieval from origin servers. Rather, per the disclosure, content is fragmented because it is broken up into fragments by the devices acting as “tunnels” (and simultaneously, peers) and cached by them to support later piece-wise retrieval. Ex. 2015 at 140. Without this, the claims of the ’866 patent would not be operable. Ex. 1207 ¶ 6. Nor is this a matter of alternate embodiments. There is no suggestion in the disclosure that would support a conclusion that the claims of the ’866 patent are based on different embodiments than those that support the claims of the ’044 patent. No non-caching embodiment is evident in the disclosure.²

² A device running a tunneling protocol generally will not cache what goes through the tunnel, because without additional facilities it generally cannot see the content tunneled through it. Ex. 1207 ¶ 6’. However, the claims here, like the disclosure, only recite a tunneling protocol for the *initial leg* of passing the second request to the tunnel device. While a tunneling proxy could perhaps have been used in a non-caching embodiment, the specification does not describe this. *Id.*

2.3. Duty of Candor

Lectrosonics instructs with regard to the duty of candor that “[i]nformation about an added limitation may be material even if it does not include the rest of the claim limitations.” *Lectrosonics* at 10. The rejection of claims 1, 20, and 29 of the ’614 patent in the ’880 reexamination, also asserting resource utilization/idling limitations, which are highly material to the RMTA, and the prior art cited with respect thereto (discussed beginning on page 11), fall squarely under this requirement. Yet, while inserting substantially the same limitations herein, PO failed to disclose these developments or art, though in the reexamination itself, PO claims to be “diligently working towards drafting a response.” *See* Ex. 1205 at 6.

3. UNPATENTABILITY

3.1. Lack of Enablement (§ 112(a)) and/or Utility (§ 101)

There is no reasonable explanation that would untangle the chicken-and-egg circularity of how the same content slices partitioned from fresh content obtained from the first server can then be fetched directly from the first server. Nor does the disclosure or the RMTA address how the devices can partition fresh content and then serve the resulting content slices on later request, without caching them. Such claims lack enablement. *See Newman v. Quigg*, 877 F.2d 1575, 1581-82 (Fed. Cir. 1989). Furthermore, there is no explanation provided or apparent so to how a method in accordance with these claims (assuming they could be construed) would be operable (and thus have utility under § 101).

Although utility can often be upheld even based on improbable operations, it is difficult to see a coherent set of actual steps with actual devices that could operate in accordance with these claims. *See id.*; *In re Swartz*, 232 F.3d 862, 863-63 (Fed. Cir. 2000).

3.2. Claim Construction and Indefiniteness

3.2.1. “so long as”

The specification does not use the exact phrasing “so long as.” It does use the words “as long as” eleven times, in each case to mean that a specified state or ongoing operation continues “as long as” a specified condition is met. *See, e.g.*, 119:26-27 (tunnel device remains in an ONLINE state “as long as no congestion is detected”). In the proposed claim, however, the subject of the “so long as” wording is not a state or ongoing operation, but rather a discrete action (“selecting”). “Selecting” a tunnel device would happen (if and when it did happen) in a discrete operation. This fact eliminates “so long as” as referring to a characteristic concerning the *continuation* of the state or operation. The words “so long as” in these claims must rather be in the nature of an “if,” “when,” or “only if” type of condition.

The “if” and “when” constructions, however, are not supported by the specification. The tunnel device autonomously may report itself (*e.g.*, to the first server) as not congested (or idle). Ex. 2025, 161-62. But the first server is not constrained under the disclosure to hand out that tunnel device’s ID immediately. There are numerous other considerations involved that may result in other tunnel devices being selected first. *E.g.*, 106:46-50.

Petitioner submits that this term is indefinite. However, solely for purposes of addressing a patentability analysis, Petitioner will assume that “so long as” means “only if”. *I.e.*, the actor performs the step of selecting the tunnel device only if the the condition (however determined) has been met.

3.2.2. “the first request . . . indicating that the content slice is to be fetched directly from the second server”

The word “indicating” is indefinite. The word itself does not suggest any bounds to what “the first request . . . indicating” might encompass, or how it is carried out. Nor does PO point to any disclosure that might shed light on what this term means.

3.2.3. “idle resource”

The ’866 specification at 121:25-30 addresses “idling” as being inversely related to CPU utilization (“CPU utilization relates to the relative time that the CPU is not idling”). Disclosure at 158:62-159:9 (though not directly related to determining a “congestion” state as discussed at col. 121) likewise relates an “idleness resource” inversely to a “CPU busy” condition, or (non-inversely) to “idleness of communication,” among others. The claim term “idle resource” should be understood accordingly, as the inverse of CPU utilization, or another idle condition of a system, such as idleness in communications.

3.3. Obviousness

The primary focus here will be on differences introduced by the current revisions. Ex. 1206 is a redline showing the incremental changes of the revised amendment. However, the claims are now sufficiently restructured

(and scrambled in the respective meanings of the limitations), that in some respects they can only be addressed as a whole.

3.3.1. Obviousness of revised claim 29 over Shribman in view of Durso (or any one or more of Chauhan, Liu, MorphMix, Sigurdsson, Mithyantha, or Jackowski)

Preamble – “A method for fetching a fresh content over the Internet, requested by a first device, identified in the Internet by a first identifier, from a first server identified in the Internet by a second identifier via a group of multiple tunnel devices, each identified in the Internet by an associated group device identifier”

No new changes. *See* Opposition at 12. *See generally* PG at 10.

“the method comprising the step of partitioning the fresh content into a plurality of content slices, each content slice containing at least part of the fresh content, and identified using a content slice identifier, and for each of the content slices, comprising the steps of:”

No new changes. *See* Opposition at 12; Petition at 38, 49; Ex. 1017 ¶¶ 0089; Ex. 1018 ¶¶ 0245-0248.

“(a) selecting a tunnel device from the group so long as the tunnel device has an idle resource that satisfies a criterion”

As addressed in the Opposition (at 14), Shribman discloses its corresponding first device receiving an identifier (second identifier) for an intermediate retrieval device (first tunnel device) in response to its prior request (first request). Furthermore, this “receiving” in Shribman occurs only in the event the first server has sent the second identifier to the first device, and the first server does so only when a resource utilization condition of the tunnel device

is satisfied, *i.e.*, where first server has recorded that the tunnel device is in an ONLINE state. Shribman, cl. 32, ¶¶ 38, 66, 69, 80, and 99.

As previously discussed (Opposition at 22), Shribman was also utilizing consumer clients as tunnel devices, and as such, device load and potential interruption of these clients would naturally have been a concern, and thus a motivation for the incremental improvement of checking that the devices were sufficiently free (in addition to just being online), before assigning them for retrieval tasks. Ex. 1207 ¶ 8.

As cited and applied in the '880 reexamination, ample examples existed in the art as of mid-2013 as to how to introduce such load controlling features in allocating tasks to a plurality of distributed agents.

It would have been obvious to modify Shribman to further require that the tunnel device whose ID would be sent to the requesting client (first device) for it to use for content retrieval not only be determined to be online (as expressly disclosed in Shribman), but also in a sufficiently idle state so as not to be overly burdened by assigning it such a task.

For example, Ex. 2011, Dorso *et al.*, US 2011/066924 A1, cited in the '880 reexamination, teaches fetching content over the Internet from other devices via “relay peers,” managed by a “tracker peer” (itself another peer). The disclosed architecture, shown in Fig. 1 of Dorso, also involving the fetching of Internet content through a plurality of distributed agents, closely resembles that of both Shribman and the '044/'866 patents. Dorso discloses “a peer-to-peer computer environment”—“a distributed network architecture that is composed of participants that make a portion of their resources (such as processing power,

disk storage, and network bandwidth) available directly to their peers.” Dorso, ¶ 0025. Peers are implemented on systems schematically illustrated in Fig. 3 of Dorso, which may be consumer devices. *See* ¶ 0069. Dorso’s relay peers correspond to the “tunnel devices” of the ’044/’866 patents. *See* Ex. 1207 ¶ 10. Similar to the acceleration server of the ’044/’866 disclosure, Dorso’s tracker peer sends a list of peers to the user device. ¶ 0049. Corresponding to step (c) here, Dorso further discloses that the selection of relay peers may be “optimized,” based on “factors related to the capabilities and resources, *such as CPU and bandwidth.*” ¶¶ 0050, 0059 (where the steps are performed by the tracker peer, corresponding to the “active” construction of step (c)). In another embodiment (¶ 0032), the optimizing selection is performed by the requesting user device, corresponding to the “passive” construction of step (c).

Given the close similarities in its objects and operation to both Shribman and the ’044/’866 patents, Dorso would have been an obvious source to look to, to further develop Shribman’s tunnel device selection methods to include an assessment of the tunnel’s resource utilization. Step (c) expresses this as an assessment of the tunnel device’s “idle resource,” but the ’044/’866 disclosure itself makes clear (at 148:25-29) that idling is merely the inverse of CPU utilization as disclosed in Dorso (*i.e.*, the one is obvious and/or inherent based on the other). Ex. 1207 ¶ 9.

Numerous other references, all cited in the reexamination, are to the same effect as Dorso with regard to selecting retrieval devices based on resource utilization and/or idling. For example, Ex. 1208, Chauhan, US 2010/0262650 A1, cited and applied in the ’880 reexamination (*see* ’880 Office Action Ex. 1204 at

21), relates to “using an intermediate device to manage network connections” (§ 0002), disclosing “appliances” that “work in cooperation or conjunction with each other to accelerate network traffic” (§ 0056), similar to the objects addressed in Shribman and the ’044, ’866, and ’614 patents.

Chauhan’s appliances comprise load-balancing virtual servers (§ 0239), which can perform operations such as “expir[ing] client connections by tracking the amount of time a connection has been idle” (§ 0240), or “terminating the execution of a virtual machine” (§ 0155). Back-end servers may be selected by a load balancing virtual server “that manages the load placed on each server.” § 0294. Similarly, § 0251 describes virtual servers that may distribute load “across multiple appliances.” Chauhan further describes a “monitoring agent” that “measures and monitors the memory, CPU and disk usage and performance of the appliance (§ 0129), and also has the capability to “monitor and measure performance of of any application” on a client (§§ 0148, 0071).

Chauhan thus discloses both the advantages of managing the loads of a plurality of devices assigned to perform network task, as well as how to measure device loads to achieve such control. This is what a POSITA would need to know to incorporate a criterion based on resource utilization of idling into Shribman to arrive at the present claim. The techniques taught in Chauhan would have been readily applicable to Shribman, with predictable results. Ex. 1207 § 10.

Ex. 1210, Liu, US 2010/0082513 A1, similarly addresses management of network operations by distributed agents that monitor the state of distributed servers. At § 0037, Liu discloses agents that collect data on “resource usage,

performance, workload, etc.,” and to “transform or classify collected data . . . as discrete values (e.g., CPU usage equals ‘high,’ ‘medium,’ of ‘low’), according to programmable boundaries.” This fully corresponds to the claimed comparison with a “set threshold.” Liu tracks device state with a “state machine” (*see* Fig. 5), where state transitions reflect a measured condition crossing such a set threshold. Ex. 1207 ¶ 11.

Ex. 1212, Mithyantha, US 8972601, and Ex. 1213, Jackowski, US 2013/0304791 A1 (both prior art under 35 U.S.C. § 102(a)(2) (AIA)) are from the same company as Chauhan, have overlapping disclosures, and are described with reference to the same types of network appliances. Mithyantha discloses “status monitor 704 may determine that a monitored status or value exceeds a threshold, such as a CPU utilization rate past a threshold for a predetermined time,” causing it to be removed from routing (so that it will no longer receive requests). 57:11-14, 2:15-19, 3:4-8, claim 4.

Jackowski discloses at ¶ 0220 a connection manager that can manage a list of peers and detect congestion (corresponding to idleness of its connection *below* a threshold), and in ¶ 0221 sending a request to a peer via a tunnel to establish a new connection after the connection manager determines that the new peer connection has not been flagged for congestion (has an idling resource *above* a threshold).

Ex. 1214, MorphMix (authenticated as to its publication date by Exs. 1215 and 1221) addresses conditions under which a content relaying node may or may accept a role in retrieving content, including whether it has “spare resources” (p. 134), such as “computing power” (p. 146).

Ex. 1210, Sigurdsson, US 2007/0174246 A1, at ¶ 0061, teaches Sigurdsson wherein determining whether a peer is available is performed by the client device “monitoring polling by the client.”

Any of the foregoing analogous references would provide sufficient suggestion and guidance for a POSITA to easily have modified Shribman by also tracking the resource utilization and/or idling of the devices acting as tunnels and determining to utilize them only if their performance fell above a set threshold of idling (or below a set threshold for CPU or bandwidth utilization.

“(b) sending over the Internet a first request from the first device to the selected tunnel device using a tunneling protocol and the group device identifier of the selected tunnel device, the first request including the content slice identifier and the second identifier and indicating that the content slice is to be fetched directly from the first server”

The substantive changes in this language are the tunneling protocol (previously addressed at length in the Opposition), and the further limitation that the first request now is “indicating” (contrary to the preamble and partitioning step) “that the content slice is to be fetched directly from the first server.”

To the extent this limitation by itself could make any sense, it might be understood as the “first server” being a peer that has the content. Granted, this interpretation may be seen as conflicting with the preamble, but no worse than the interpretation therein of the first server being the origin server.

Putting aside § 112 issues, Shribman definitely does disclose direct fetches of content slices from servers distinct from the device serving as the tunnel and distinct from the origin server, *i.e.*, peers. Ex. 1017 ¶ 0089-90. This construction is also happens to reflect as close as the ’866 specification gets to

supporting this limitation. If the Board is inclined to construe the limitation at all, and construe it that way, then Shribman meets the limitation.

“(c) in response to receiving the first request, sending the content slice identifier from the selected tunnel device to the first server using the second identifier;”

Subject to the same § 112 concerns as apply to clause (b), Petitioner asserts that Shribman discloses this limitation, insofar as it discloses in response to (*i.e.*, after) the selected tunnel device receives the first request, it sends a content slice identifier to a peer device (first server), using its identifier (second identifier). Ex. 1017 ¶ 0085.

“(d) following step (c), receiving over the Internet the content slice from the first server through the selected tunnel device”

Consistent with the construction assumed for (b) and (c), Shribman discloses that the requesting client, following the request to the first server per step (c), receives the content slice it requested over the Internet, which came from the first server (a peer), through the selected tunnel device. Ex. 1017 ¶ 0087.

“wherein the method further comprises the step of constructing the fresh content from the received plurality of content slices, and”

Previously addressed, *see* Ex. 1017 ¶ 0096; Ex. 1018 ¶¶ 0032, 0364, 0559-61, Fig. 32; Opposition at 22.

“wherein each of the tunnel devices in the group is a client device comprising a consumer electronic device with a client operating system, and”

No new changes. *See* Opposition at 22. As noted therein, Shribman is based on consumer devices. *See also* Dorso at ¶ 0069; Morphmix at 116 (“any computer with a dial-up connection capable of running a modern graphical web browser should be sufficient”).

“wherein the criterion is satisfied when an idleness of the idle resource is above or below a set threshold value and”

The references cited either expressly disclose this limitation or render it obvious.

Mithyantha: “status monitor 704 may determine that a monitored status or value exceeds a threshold, such as a CPU utilization rate past a threshold for a predetermined time.” Ex. 1212 at 57:11-14, 2:15-19, 3:4-8, claim 4.

Liu: ¶ 0058: “As illustrated in FIG. 4, if the data collection agent detects a change of state (e.g., a transition from ‘medium high’ CPU usage to ‘high’ CPU usage), a threshold value that is reached, or a threshold rate of change that is reached . . . it may send a message to the information layer agent indicating this event”—substantially identical to the corresponding disclosure of this feature in the ’044 and ’866 patents.

Dorso: ¶¶ 0031-32, 0050. Given that Dorso teaches to use an algorithm to optimize relay peer selection based on CPU resources, it would have been obvious to compare observed levels of CPU activity (or conversely, idling) to a set threshold, in order to make a selection based thereon.

Furthermore, without limiting the foregoing, using the utilization information obtained in accordance with any of Chauhan, MorphMix (determining “spare resources”), Sigurdsson, or Jackowski implies comparing it to a thresh-

old to make it actionable as disclosed. *See* Ex. 1222, *Fry et al.*, 6:21-48; Ex. 1207 ¶ 20.

“wherein the selected tunnel device does not cache a copy of the content slice”

The Opposition (at 17-21) cited a number of additional references (including Shribman itself) that disclosed tunneling protocols for content retrieval, further including Ex. 1216, Prince and Ex. 1218, VIP72 (supplemented by Ex. 1217 and authenticated by Exs. 1219 and 1220). A tunneling protocol generally does not “see” what is being tunneled through it and cannot cache. Ex. 1207 ¶ 6. Thus, the “tunneling” references already cited also meet this further limitation.

“wherein the first device remains anonymous as to the first second server.”

Previously addressed under Shribman. Opposition at 24; Ex. 1017 ¶ 0089.

3.3.2. Claims 30-38

Changing “content” to “fresh content” does not change the applicability of Shribman at all, because Shribman’s agents likewise partitions fresh content they receive directly from the origin server. Ex. 1017 ¶ 0089; Fig. 10, steps 400, 402. In all other respects, as seen in Ex. 1206, these claims present no new changes, and have been previously addressed.

4. CONCLUSION

The RMTA should be denied and all challenged claims cancelled.

Respectfully submitted,

/Ronald Abramson/

Ronald Abramson

Dated: April 21, 2022

LISTON ABRAMSON LLP
The Chrysler Building
405 Lexington Avenue, 46 FL
New York, NY 10174
Telephone: (212) 822-0163

CERTIFICATE OF COMPLIANCE

The undersigned hereby certifies that Petitioner's within OPPOSITION TO REVISED MOTION TO AMEND is within the 25 page-limit, in compliance with 37 C.F.R. §§ 42.24(a)(1)(vi).

/Ronald Abramson/

Ronald Abramson

Dated: April 21, 2022

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. § 42.6(e), the undersigned hereby certifies that Petitioner's within OPPOSITION TO REVISED MOTION TO AMEND and all supporting exhibits were served on April 21, 2022 via email, as authorized by Patent Owner, on the following:

elizabetho@ruyakcherian.com
tomd@ruyakcherian.com

/Ronald Abramson/

Ronald Abramson

Dated: April 21, 2022