



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
90/014,552	07/23/2020	7103640		3765

757 7590 09/09/2020
BGL
P.O. BOX 10395
CHICAGO, IL 60610

EXAMINER

SORRELL, ERON J

ART UNIT	PAPER NUMBER
----------	--------------

3992

MAIL DATE	DELIVERY MODE
-----------	---------------

09/09/2020

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

DO NOT USE IN PALM PRINTER

(THIRD PARTY REQUESTER'S CORRESPONDENCE ADDRESS)

ANEETA MISHRA (GENERAL)
50 STEEPLECHASE
IRVINE, CA 92602

***EX PARTE* REEXAMINATION COMMUNICATION TRANSMITTAL FORM**

REEXAMINATION CONTROL NO. 90/014,552 .

PATENT UNDER REEXAMINATION 7,103,640 .

ART UNIT 3992 .

Enclosed is a copy of the latest communication from the United States Patent and Trademark Office in the above identified *ex parte* reexamination proceeding (37 CFR 1.550(f)).

Where this copy is supplied after the reply by requester, 37 CFR 1.535, or the time for filing a reply has passed, no submission on behalf of the *ex parte* reexamination requester will be acknowledged or considered (37 CFR 1.550(g)).

Order Denying Request For Ex Parte Reexamination	Control No. 90/014,552	Patent Under Reexamination	
	Examiner ERON J SORRELL	Art Unit 3992	AIA (FITF) Status No

--The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

The request for *ex parte* reexamination filed 07/23/2020 has been considered and a determination has been made. An identification of the claims, the references relied upon, and the rationale supporting the determination are attached.

Attachments: a) ☐ PTO-892, b) ☒ PTO/SB/08, c) ☐ Other: _____

The request for *ex parte* reexamination is DENIED.

This decision is not appealable (35 U.S.C. 303(c)). Requester may seek review by petition to the Commissioner under 37 CFR 1.181 within ONE MONTH from the mailing date of this communication (37 CFR 1.515(c)). **EXTENSION OF TIME TO FILE SUCH A PETITION UNDER 37 CFR 1.181 ARE AVAILABLE ONLY BY PETITION TO SUSPEND OR WAIVE THE REGULATIONS UNDER 37 CFR 1.183.**

In due course, a refund under 37 CFR 1.26 (c) will be made to requester:

- a) ☒ by Treasury check or,
- b) ☐ by credit to Deposit Account No. _____ , or
- c) ☐ by credit to a credit card account, unless otherwise notified (35 U.S.C.303(c)).

/ERON J SORRELL/
Primary Examiner, Art Unit 3992

cc: Requester (if third party requester)

Decision

A Third Party requested *ex parte* reexamination of claims 17, 18, and 24 of U.S. Patent No. 7,103,640 to Overton et al. ("the '640 Patent"). The Request was granted a filing date of July 23, 2020 and assigned Control No. 90/014,552.

MPEP 2241 recites,

The determination of whether or not to reexamine must be made within 3 months following the filing date of a request. See 35 U.S.C. 303(a) and 37 CFR 1.515(a). If the 3-month period ends on a Saturday, Sunday, or federal holiday within the District of Columbia, then the determination must be mailed by the preceding business day. See the last portion of MPEP § 2240 and also see MPEP § 2283 for multiple copending reexamination proceedings. A determination to reexamine may be made at any time during the period of enforceability of a patent.

No substantial new question (SNQ) of patentability is raised by the request for reexamination and prior art cited therein for the reasons set forth below.

Listing of Prior Art

The following is a list of patents and printed Publications provided in the Request:

a. T. Jaeger and A. Rubin, *Preserving Integrity in Remote File Location and Retrieval*, Proceedings of Internet Society Symposium on Network and Distributed Systems Security, pp. 53-63, Feb 1996, ("Jaeger");

b. Browne et al., *Location-Independent Naming for Virtual Distributed Software Repositories*, Proceedings of the Symposium on Software Reusability, pp. 179-185, Aug 1995 ("Browne I");

c. Browne et al., *The National HPCC Software Exchange*, IEEE Computational Science and Engineering, vol. 2, no. 2, pp. 62-69, 1995 ("Browne II");

d. Browne et al., *Management of the NHSE— a Virtual Distributed Digital Library*, University of Tennessee Technical Report, 1995, ("Browne III");

e. Moore et al., *Resource Cataloging and Distribution System*, University of Tennessee Technical Report, 1997, ("Moore");

f. Fagg et al., *Scalable Networked Information Processing Environment (SNIPE)*, Proceedings of the 1997 ACM/IEEE conference on Supercomputing, 1997, ("Fagg");

g. RFC 1737, *Functional Requirements for Uniform Resource Names*, 1994, ("RFC 1737");

h. RFC 1034, *Domain Names—Concepts and Facilities*, 1987, ("RFC 1034");

i. RFC 1035, *Domain Names—Implementation and Specification*, 1987, ("RFC 1035");

j. D.R. McGregor and J.A. Mariani, *"Fingerprinting"— A Technique for File Identification and Maintenance*, 12 Software Practice & Experience 1165, 1982, ("McGregor");

k. Jack J. Dongarra and Eric Grosse, *Distribution of Mathematical Software via Electronic Mail*, ACM, Vol. 30, pp. 403-407, 1987, ("Dongarra");

l. Andrew Binstock, *Hashing Rehashed*, 1996, ("Binstock");

m. *McGraw-Hill Dictionary of Scientific and Technical Terms* 860 (4th ed. 1989) ("McGraw");

n. Cormen et al., *Introduction to Algorithms*, The MIT Press (1st Ed. 1990) Andrew S. Tanenbaum, *Computer Networks*, 3rd edition (Third Edition) (1996) (excerpted), ("Cormen");

o. Donald E. Knuth, *The Art of Computer Programming, Volume 3: Sorting and Searching*, Addison-Wesley (1973)

p. Brian W. Kernighan and Dennis M. Ritchie, *The C Programming Language*, Second Edition (1988) (excerpted), ("Kernighan");

q. US 6430618 B1 (priority March 13, 1998);

r. David R. Cheriton and Timothy P. Mann, *Decentralizing a Global Naming Service for Improved Performance and Fault Tolerance*, ACM Transactions on Computer Systems 147-183, 1989, "Cheriton").

The eighteen (18) references listed above were neither used in a rejection or even appear to be considered during a prior prosecution of the '640 Patent, thus all are new references.

Third Party Requester's Position

The Requester alleges the following SNQs which have been grouped/organized for discussion purposes:

Group I: Jaeger to raise a SNQ over claims 17, 18, and 24. This group of SNQs is based on proposed rejections of claims 17, 18, and 24 as being obvious over Jaeger.

Group II: Browne I, Browne II, Browne III, Moore, and Fagg (collectively "the Netlib Publications") to raise an SNQ over claims 17, 18, and 24. This group of SNQs is based on proposed

rejections of claims 17, 18, and 24 as being obvious over the Netlib Publications.

Group III: Jaeger in combination with the Netlib Publications to raise an SNQ over claims 17, 18, and 24. This group of SNQs is based on proposed rejections of claims 17, 18, and 24 as being obvious over Jaeger taken in combination with the Netlib Publications.

Review of Prosecution History

The '640 Patent issued on September 5, 2006, having claims 1-25, with claims 1, 10, 17, and 18 being independent claims. The '640 Patent matured from U.S. Pat. Appl. No. 09/661,222 ("the '222 Application") which was filed on September 13, 2000, which claimed priority to U.S. Provisional App. No. 60/153,709, filed September 14, 1999 and U.S. Pat. Appl. 09/111,896 filed July 8, 1998.

Reviewing the prosecution of the history of the '222 Application, it is found that a Notice of Allowance was mailed on November 14, 2005, allowing claims 24-48, with claims 24, 33, 40, and 41 being independent claims. Upon issuance, claims 24-28 were renumbered as claims 1-25 with claims 1, 10, 17, and 18 being independent claims. In the Notice of Allowance, there was no specific reasons for allowance provided by the Examiner. The

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 7 of 39

Notice of Allowance did include an Examiner's Amendment, however the accompanying remarks indicate that the amendment was solely to address an inadvertent claim numbering error.

Looking back further, the previous Office action was a non-final rejection mailed on February 9, 2005 ("the 2/9/05 Office action"), in which claims 1-9 and 12-23 (all of the pending claims at the time) were rejected as being anticipated by: 1) a prior art publication to Beitz titled "Service Location in an Open Distributed Environment" ("Beitz"); 2) U.S. Patent No. U.S. Pat. No. 5,950,173 to Perkowski ("Perkowski"); 3) U.S. Patent No. 5,978,773 to Hudetz et al. ("Hudetz"); and 4) U.S. Pat. No. 5,913,210 to Call.

Given the above, a comparison between the allowed claim set and the claim set that was rejected in 2/9/05 Office action would shed light on why the claims were allowed in the original prosecution.

Issued claims 17, 18, and 24 of the '640 Patent are system claims with claims 17 and 18, being independent claims and claim 24 depending on claim 18. In the claim set that was rejected in the 2/9/05 Office action, only claims 5-9, 19, and 20 were directed toward a system, with claim 5 being independent and claims 6-9, 19, and 20 depending therefrom. Claim 5, that was

rejected in the 2/9/05 Office action, and allowed claims 17 and 18 are reproduced below for discussion purposes:

Rejected claim 5 in 2/9/05 Office action	Issued claim 17 of the '640 Patent
5. A system having a network distributed tracking wire transfer protocol for storing and identifying data with a distributed data collection, comprising: a data repository, the data repository for storing data in a distributed data collection; a client entity, the client entity for manipulating data in the distributed data collection; and a first server entity, the first server entity operative to locate data in the distributed data collection;	17. A system for retrieving data location information for data stored in a distributed network, the system comprising: a plurality of data repositories configured to store data, wherein the data is associated with a respective identifier string in each data repository; a data location sewer [sic] network having a plurality of data location servers, each of the plurality of data location servers containing location strings associated with respective identifier strings

wherein the first server entity is responsive to a client request from the first client entity to provide a redirection message to the first client entity if the first server entity cannot provide at least one location string to the client entity corresponding to an identifier string in the client request.	and each of the plurality of data location servers having computer executable code configured to execute the following steps: in response to receiving a data location request from a client to retrieve a location string associated with an identification string provided in the data location request, transmitting a redirect message to the client if the identification string is not associated with a location string at the data location server, <u>wherein the redirect message contains information for use by the client to calculate a location of a different data location server in the plurality of data location servers, wherein the</u>
--	--

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 10 of 39

	<u>different data location server</u> <u>contains the location string.</u>
--	---

Rejected claim 5 in 2/9/05 Office action	Allowed claim 18
5. A system having a network distributed tracking wire transfer protocol for storing and identifying data with a distributed data collection, comprising: a data repository , the data repository for storing data in a distributed data collection;; a client entity, the client entity for manipulating data in the distributed data collection; and a first server entity , the first server entity operative	18. A system for retrieving data location information for data stored in a distributed network, the system comprising: a data repository configured to store data, wherein the data is associated with an identifier string; a client responsive to a data query to query a data location server for location information associated with the identifier string;

to locate data in the distributed data collection; wherein the first server entity is responsive to a client request from the first client entity to provide a redirection message to the first client entity if the first server entity cannot provide at least one location string to the client entity corresponding to an identifier string in the client request.	a data location server network comprising a plurality of data location servers, at least one of the plurality of data location servers containing location information associated with the identifier string, wherein each of the plurality of data location servers comprises computer executable code configured to execute the following steps in response to receiving a data location request from the client: if the data location server contains the location string associated with the identification string provided in the data location request, the data
---	---

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 12 of 39

	location server transmits location information for use by the client to calculate a location of the data associated with the identification string; if the data location server does not contain the location string associated with the identification string, the location server transmits a redirect message to the client, wherein the redirect message contains redirect information for use by the client to calculate a location of a different data location server in the plurality of data location servers, wherein the different data location
--	---

	server contains the location string.
--	---

As can be seen above, claim 5 that was rejected in the 2/9/05 Office action and allowed claims 17 and 18 all generally recite a system comprising a data repository, a client, and a server entity (or network of servers), wherein the server attempts to locate data requested by the client in the data repository and provide that location information back to the client. Both sets of claims also recite the use of identification strings that identify data and location strings that identify data locations, and the server sending a redirect message to the client if the server cannot locate the requested location string. The distinguishing feature between the rejected claims and the allowed claims appears to be the contents of the redirect message. In rejected claim 5, the contents of the redirect message are not claimed, whereas in the allowed claims, the redirect message is said to, "contain redirect information for use by the client to calculate a location of a different data location server in the plurality of location servers, wherein the different data location server contains the location string." Reviewing all of the claims rejected in the 2/9/05

Office action, the Examiner finds that in addition to claim 5, independent claims 1 and 12, and dependent claims 7-9, 17-23 all include limitations regarding the redirect message, however none of them require the redirect message to comprise information for use by the client to calculate the location of a different location server, wherein the different location server contains the location string.

This limitation was also the focus of the Applicant's remarks filed in response to the 2/9/05 Office action. The Applicant argued,

"Referring to claim 24, when a user enters a request at a client for data at a first client, the data query includes an identifier string (e.g. a file name or URI) and expects to receive information (a location string) on where the data corresponding to that filename resides... During this navigation process, if a data location server does not possess the sought after location string, it transmits a redirect message back along the communication path the data location request traveled to the first client. The redirect message that is sent back to the first server contains information with which the first client is configured to determine a location of a second data location server that does contain the sought after location string. In other words, the redirect message allows the first server to calculate a new data location server address so that the first server "redirects" its query to the right data location server..."

...Similarly, the method of claim 33 also includes includes the feature of redirection, however claim 33 addresses a more simplified scenario such as shown in FIG. 12, where there are not necessarily any intervening devices between the client and the data location servers.

New independent claims 41 and 42 relate to system configurations adapted to incorporating software instructions implementing the redirection process."

See page 11 of the remarks filed 8/9/2005.

"In a hypothetical system that needs to track 100s of billions of objects across 100s of millions of clients and repositories in which any object can "spontaneously and dynamically" change location from any client to any other client in real-time, there are significant challenges. The claimed redirection feature is intended to help address and solve problems like these, using a theoretically one-point-failure-retry redirection method that configurably distributes computation among clients and servers."

See page 12 of the remarks filed 8/9/05.

Regarding the applied prior art the applicant argued,

"In a hypothetical system that needs to track 100s of billions of objects across 100s of millions of clients and repositories in which any object can "spontaneously and dynamically" change location from any client to any other client in real-time, there are significant challenges. The claimed redirection feature is intended to help address and solve problems like these, using a theoretically one-point-failure-retry redirection method that configurably distributes computation among clients and servers."

None of the previously cited references teach or suggest the features of claims 24-49.

Beitz et al. discloses a system that needs some type of improved data management system, for example see p. 32 of Beitz, but fails to address how that data management system would be implemented. Accordingly, Beitz does not teach or suggest a redirection mechanism as recited in the pending claims.

Call discloses mapping UPC codes to IP addresses in a standard relational database using mirrored servers each having a complete copy of a look-up table. (See Col. 3, line 61 - Col. 4, line 9). Call does not disclose or suggest the claimed redirection mechanism.

Perkowski discloses linking consumers to product information using a relational database. Perkowski teaches the use of traditional relational databases using centralized or mirrored deployment. (See Col. 7, lines 38-42; Col. 11, lines 34-61; Col. 12, lines 30-35; and Col. 18, lines 13-35). Perkowski fails to teach or suggest a redirect message containing information for a client or server to determine the location of another data location server, a client or server computing redirection information or of the location of servers changing dynamically.

Similarly, Hudetz et al. discloses mapping UPC numbers to IP addresses that can provide the consumer more information on the product. The disclosure mentions in passing a database that stores the UPC-to-IP data table and does not disclose or suggest a dynamic data management system with a redirection function. (See FIG. 4; Col. 7, lines 1-9)."

See pages 12-13 of the remarks filed 8/9/05.

In light of the foregoing, if a new reference or combination of references teaches or fairly suggests a redirect message containing redirect information, for use by the client that requested a data location string from a data location server, to calculate a location of a different data location server of a plurality of data location servers, wherein the different data location server contains the location string, then that reference or combination of reference would raise an SNQ.

Substantial New Questions of Patentability

Group I: Jaeger

The Requester alleges that Jaeger raises a SNQ over claims 17, 18, and 24 of the '640 Patent. Upon review and for the reasons discussed below, the Examiner disagrees. Jaeger does not raise a SNQ over claims 17, 18, or 24 of the '640 Patent.

Jaeger teaches a service that allows users to share files distributed over the Internet. For example, a technical document can be retrieved using a document identifier from a remote server in a distributed network even if the user does not know the identity of the machine storing the document. See section 1, page 53.

Jaeger's system architecture comprises clients that request files using a file identifier, distribution servers that store and distribute files, and location servers that map file identifiers to distribution servers. This architecture is illustrated in figure 1 of Jaeger reproduced below.




Figure 1: **Service Process:** (1) distribution server publishes a file; (2) client requests a file's location; (3) location server returns a set of distribution servers for a file; (4) client requests a file from a distribution server; and (5) distribution server returns a file and its certificate for authentication.

Jaeger provides a detailed explanation of the file location and file retrieval protocols in figures 3 and 4 and in the related descriptions (reproduced below);




Figure 3: **The file location protocol:** (1) Client sends a file identifier to the client's location server to find the distribution servers that possess a file; (2) Client's location server requests that other location servers find the distribution servers that possess a file (optional); (3) Other location servers return to the client's location server a set of distribution servers that possess a file; (4) Client's location server returns a set of distribution servers that possess the file.




Figure 4: **The file retrieval protocol:** (1) client sends a file identifier to a distribution server; (2) the distribution server returns the corresponding file and its certificate.

Jaeger teaches that a client performing the file location protocol is provided with a list (possibly empty) of

distribution servers that the client's location server claims has the requested file. If the client's location server has no knowledge of the files location, it can ask for help from other location servers and can cache any new information received from those other servers so it will have that information available next time a client requests it. Jaeger concedes that only messages 1 and 4 of the file location protocol have actually been implemented, but does not anticipate that messages 2 and 3 would be problematic. See section 4.3, page 60.

Jaeger's file retrieval protocol makes use of the set of distribution servers returned in the file location protocol to retrieve a file and its certificate. The client uses the same file identifier it sent to the location server to request the file from a selected distribution server. The distribution server returns the corresponding file and its certificate to the client for authentication if it has it. If the distribution server does not have the file it can return a null file and the client and request the file from another distribution server. If the file returned to the client fails authentication, the client can retry retrieval from the same distribution server, try retrieval using another distribution server, obtain a new distribution server set from the same location server, or obtaining a new distribution set from another location server. See section 4.4, page 60.

Jaeger, however, does not appear to teach a redirect message that is sent to a client, wherein the redirect message contains redirect information for use by the client to calculate a location of a different data location server of the plurality of data location servers, wherein the different data location server contains the requested location string. The only transmission from the location server to client in Jaeger is a location structure which comprises an array of distribution servers and the number of distribution servers (see section 4.3, page 59). Jaeger teaches that if a location server does not know the location of a distribution server that has a specific file, it asks for help from other location servers and is returned a location structure as described above (see section 4.3, page 59). In other words, Jaeger teaches that only ***distribution server*** information is transmitted to a client or downstream location server, not information used to calculate to location of another location server that has the location string.

The Requester alleges that the relevant artisan would have discerned that Jaeger's approach would have been better modified by returning a table to the client for the client to use to determine which location server to contact in order to request the file. The Requester also alleges,

"More broadly, the relevant artisan would have recognized that "[h]ashing is useful anywhere a wide range of possible values needs to be stored in a small amount of memory and be retrieved with simple, near-random access." AX0511; see also AX0511 ("As such, hash tables occur frequently in database engines" where "the table is a principal data structure" and "all accesses to the table must be done through the hash function.")".

See page 17 of the Request. The Requester, however does not actually cite any teaching in the prior art of a client that receives a table of other location servers or information that can be used to calculate the location of another location server from a first location server that the client requested a location string from. "Hashing" was known in the prior art, and could be used in place of "for-loop" processing to more efficiently locate the appropriate location server, as argued by the Requester, however there is no teaching or suggestion in Jaeger **for the client** to perform that processing (emphasis added). The location server that receives the request could perform the hashing to locate another location if it has no knowledge of the requested location string.

The Requester also argues that Jaeger refers to the "server-centric" location resolution as "optional" and refers to Jaeger's focus on "trust" to provide motivation to move to a "client-centric" location resolution. See page 17 of the Request.

The Examiner disagrees. While Jaeger does indicate certain steps are optional, Jaeger does not provide any other examples EX2014

or alternatives for location resolution and certainly does not teach or suggest a redirect message as claimed. Jaeger states the steps indicated as optional have not actually been implemented (see section 4.3, page 60), so in context "optional" could simply mean not required rather than selecting one of a plurality of available options. Additionally, Jaeger ensures "trust" by a file authentication protocol. There is no teaching or suggestion in Jaeger to achieve "trust" by the client taking on additional processing (i.e. determining another location server to contact).

During the original concluded prosecution, the Beitz reference was used to reject all of the claims in the 2/9/05 Office action. Beitz teaches a mechanism for name resolution comprising a locator supported by a one or more name servers, wherein each name server is capable of resolving the local name of naming context (See Beitz, page 31-32). Beitz teaches the location function is made up of two levels - a high level that determines a name server required, and low level that accesses the required name server to map a name into an access detail list. See page 32 and figure 3 reproduced below.




Figure 3: The Location Function

Beitz also teaches a locator using another locator to resolve names it is unable to handle and provides suggestions. See Beitz, page 32. This location-resolution mechanism taught by Beitz, including what happens when a location server cannot resolve a location string, is the same type of mechanism disclosed by Jaeger. With that, Jaeger is found to be cumulative to the Beitz reference already considered in the concluded original prosecution.

Given that Jaeger does not teach or suggest the limitations in independent claims 17 and 18 found allowable during the original prosecution, and the finding that Jaeger is cumulative to the reference of Bietz that was considered in the original prosecution of the '640 Patent, it is not likely that a reasonable examiner would consider Jaeger important in

determining the patentability of the '640 Patent. With that, Jaeger is not found to raise an SNQ over claims 17 or 18. Since dependent claims carry with them each of the limitations from the claims from which they depend, it follows that Jaeger does not raise a SNQ over dependent claim 24.

Group II: Netlib Publications

The Requester alleges that the Netlib Publications raise an SNQ over claims 17, 18, and 24 of the '640 Patent. Upon review and for the reasons discussed below, the Examiner disagrees. The Netlib Publications do not raise a SNQ over claims 17, 18, or 24 of the '640 Patent.

Browne I

Browne I is a paper describing a location-independent naming system for network resources that enables easy searching and retrieval of those resources while addressing consistency and integrity issues involved with distributed software repositories. Browne I discloses the design of the naming system, describes a prototype implementation (referred to as "Netlib"), and describes how Netlib fits into the development of the National HPCC Software Exchange (NHSE) - a virtual software repository (see Abstract). The goal of the NHSE is to provide a uniform interface to a virtual software repository built on top

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 25 of 39

of a distributed set of discipline-oriented repositories as shown in figure 1 (reproduced below), while allowing for mirroring of the software packages to increase their overall availability (see Browne I, pages 179-180).




Figure 1: Virtual Repository Architecture

Browne I discloses most problems associated with past repositories can be alleviated by using a location-independent naming service that includes mechanisms for authenticating and checking the integrity of the retrieved resources. Browne I's naming system provides for two levels of naming, a lower level and a higher level. At the lower level there is a binding between a lower level name called a LIFN and file contents which is unchanged and verifiable. At the higher level, a name called

a Universal Resource Name ("URN") which associates with a cataloging record that includes the lower level name and other descriptive information. At any given time, a higher level name can be associated with exactly one lower level name, but that binding can change over time (see page 180). In short, the system works by assigning location-independent names to files and providing a name-to-location service that given a name returns a list of locations for that name. A user can then query the name-to-location service by providing a file name and receiving a list of alternative locations to try (see page 181).

Browne I teaches the following process to resolve a URN so as to access a copy of a file which is illustrated in figure 3 (reproduced below).




Figure 3: File access steps

1. Use DNS to locate an appropriate URN server.
2. Query the URN server to retrieve the URC which contains the currently associated LIFN.
3. Authenticate the URC if desired.
4. Use DNS to locate an appropriate LIFN server.
5. Query the LIFN server to retrieve a list of locations.
6. Choose a location from which to retrieve the file.

In practice, Steps 4 through 6 will often be replaced by using the LIFN to access a local cache server. Because the binding between a LIFN and the byte contents it points to is fixed the cached copy is sure to be correct.

Browne I further discloses that while the naming system returns a list of locations for a given LIFN, it does not guarantee the correctness of the locations. The locations may be wrong if they no longer exist, in which case no file is returned, or if the contents at the location are wrong, which can be detected from an integrity check run at the client (see section 5, page 183).

Browne I teaches that in a prototype implementation, the naming system functioned as part of a Bulk File Distribution (BFD) package, in which the BFD client is a WWW browser which has the capability of returning a file given its URN or LIFN. Browne I discloses that when client program requests a properly formatted URL from a file server, the file server either returns a file that is correct for the given LIFN or an error indicating that the LIFN was not found.

Browne II

Browne II is a publication announcing the availability of the NHSE described in Browne I. Browne II describes the NHSE as in its "infant stages", but explains how the NHSE ("the Exchange") would be beneficial to scientific community by allowing the sharing of resources and cutting down on time invested "reinventing the wheel (See abstract). Browne II provides information on the types of resources provided in the

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 28 of 39

Exchange, on how contributors can add and maintain resources in Exchange; and how the Exchange can be accessed from client web browsers (see pages 63-64).

Browne II provides a brief overview of the name-to-location resolution service, however it does not provide any more details than what is disclosed in Browne I. For example, Browne II provides figure 2 which is essentially the same as figure 1 in Browne I. The associated description at page 65 of Browne II generally describes the two-level naming system, but does not provide any more information than Browne I at page 181.

Browne III

Browne III, another publication regarding the NSHE, describes the authenticated submission and review process for adding resources to the NHSE. At pages 2-3, Browne III gives a general description of the operation of the system, however does not provide any details regarding the name-to-location resolution service beyond figure 1 which, again, is the same as figure 1 in Browne I.

Moore

Moore teaches discloses a system referred to as a Resource Cataloging and Distribution System (RCDS), which is described as architecture for cataloging the characteristics of Internet-

accessible resources, replicating the resources to improve their accessibility, and for registering the current location of the replicated resources. The resources can be named with a URN or an existing URL, and either type of resource name can be resolved to description and ultimately to a set of locations from which the resource can be retrieved (see Abstract).

Moore also describes the following access and retrieval process. First, a user acquires a URN that seems to suit their needs. The URN is resolved by a DNS server to find network addresses of one of more catalog servers. One of those servers is selected and queried for the description of the resource named by the URN. The description may comprise a plurality of sub-descriptions for different versions of the file. Those sub-descriptions also comprise a LIFN for that particular version. The client then resolves the LIFN using DNS to find a list of network addresses of one or more location servers. One of those servers is queried for location of file named by the LIFN. The location server returns a list of URLs at which the file can be located. The client chooses one of those file servers and retrieves the file (see section 2.3, pages 7-10). Moore also teaches efficiency can be improved by having the client avoid fetching the resource if it can determine if it is not needed, or if a locally cached copy exists (see 4.3 at page 12). For scalability, Moore discloses that multiple instances of any

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 30 of 39

particular server and that there is no requirement that all RCDS servers for a particular portion of LIFN-space contain the same set of location so long as the client can query multiple RCDS server until it finds an accessible location for the desired resource (see page 13).

Fagg

Fagg describes a system called SNIPE (Scalable Network Information Processing Environment) which is also a distributed data stores across the internet (see abstract). Fagg discloses the SNIPE uses RCDS as a framework (see section 3). The SNIPE file servers allow for processes to store and retrieve files. SNIPE file servers can also be used to access ordinary data files via URL and LIFNs (see section 4.9). Fagg also discloses improving performance by caching as much information as possible to avoid excess loading on the RCDS servers (see section 5.2).

None of the Netlib Publications are seen as teaching a redirect message. Browne I, which is seen as the most comprehensive of the Netlib Publications, teaches that during the name resolution process if a location is incorrect or no longer exists, then no file is returned (see page 183) or an error is returned indicating the file was not found (see page 184). Moore teaches the client making multiple queries to

different servers until it finds an accessible location (see page 13), but does not teach or suggest a redirect message, or more specifically a redirect message comprising information the client can use to calculate a location server that has the requested location string.

The Requester argues that the Netlib Publications all teach or reasonably suggest decoupling of network resources from their locations in the same manner as taught in claims 17, 18, and 24 (see page 18 of the Request). However this feature is also taught by Beitz which was applied in rejections in the 2/9/05 Office action. In that respect, the Netlib Publications are cumulative to the teachings of Bietz.

Continuing, the Requester also alleges, "The relevant artisan would have readily recognized that the LIFN location server may pass back a resolved location corresponding to a LIFN identifier, or alternatively that the client may calculate from the information returned by the first LIFN server which next LIFN server to contact that could have provided the location from the network resource (see page 18 of the Request)." However, there is no teaching or suggestion in the cited prior art for the alternative action. Browne I teaches during the resolution process if a location is incorrect or no longer

exists, then no file is returned (see Browne I, page 183) or an error indicating the file was not found (see Browne II, page 184). Moore teaches the client making multiple queries to different servers until it finds an accessible location (see Moore, page 13), but does not teach or suggest a redirect message, or more specifically a redirect message comprising information the client can use to calculate a location server that has the requested location string.

Additionally, the Requester alleges that since Moore teaches the client can avoid fetching a resource if it is not needed or locally cached, it suggests a more "client-centric" approach is beneficial over the LIFN-centric approach (see page 18 of the Request). The Examiner disagrees. This statement is in regards to making efficient use of network resources by not requesting duplicative files or files that are already locally cached, and not toward handling a situation in which a location server has no knowledge of the location of a requested file.

Lastly, the Requesters argument that since the LIFN server stores key-value assertions, the relevant artisan would have known to use the information passed back to identify a next LIFN server (see page 18 of the Request) is a conclusion drawn with no factual underpinnings from any of the cited references. The statement also includes a citation from Krager regarding a client using a hash table to determine a server to issue request

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 33 of 39

for a document (see page 18 of the Request); however, Karger's system does not use a name-to-location resolution process so it is not clear how this is applicable to the Netlib Publications and no explanation of its relevance is provided.

During the original concluded prosecution, the Beitz reference was used to reject all of the claims in the 2/9/05 Office action. Beitz teaches a mechanism for name resolution comprising a locator supported by a one or more name servers, wherein each name server is capable of resolving the local name of naming context (See Beitz, page 31-32). Beitz teaches the location function is made up of two levels - a high level that determines a name server required, and low level that accesses the required name server to map a name into an access detail list. See page 32 and figure 3 reproduced below.




Figure 3: The Location Function

This location-resolution mechanism taught by Beitz is the same type of mechanism disclosed by the Netlib Publications. With that, the Netlib Publications is found to be cumulative to the Beitz reference already considered in the concluded original prosecution.

Given that the Netlib Publications do not teach or suggest the limitations in independent claims 17 and 18 found allowable during the original prosecution, and the finding that the Netlib Publications are cumulative to the reference of Bietz that was considered in the original prosecution of the '640 Patent, it is not likely that a reasonable examiner would consider the Netlib Publications important in determining the patentability of the '640 Patent. With that, the Netlib is not found to raise a SNQ over claims 17 or 18. Since dependent claims carry with them each of the limitations from the claims from which they depend, it follows that Jaeger does not raise a SNQ over dependent claim 24.

Group III: Jaeger and the Netlib Publications

The Requester alleges that Jaeger in combination with the Netlib Publications raise a SNQ over claims 17, 18, and 24 of the '640 Patent. For the reasons discussed below, the Examiner disagrees. The combination of Jaeger and the Netlib Publications

EX2014

Kove IO, Inc.

IPR2024-01021

Google, LLC v. Kove IO, Inc.

Page 35 of 39

do not raise a SNQ over claims 17, 18, and 24 of the '640 Patent.

Jaeger taken alone, and the Netlib Publications were not found to raise a SNQ over claims 17, 18, and 24 for the reasons discussed in detail above. The combination of Jaeger and the Netlib Publications is also not found to raise a SNQ. Combining the disclosures of Jaeger and the Netlib Publications still fails to provide a teaching of redirect message comprising information the client can use to calculate a location of a location server that has the requested location string.

The Requester states that Jaeger's file identifiers can be extended Netlib's LIFNs or URNs and that the Netlib architecture could be extended so the LIFN server or catalog server returns an array of similar servers as taught by Jaeger so the architecture would be more client-centric leading to local caching, near-random access, and trust (see pages 18-19 of the Request).

The Examiner disagrees. As mentioned in the discussion of Jaeger above, the array of servers returned is an array of distribution servers and not an array of location servers to try next (see Jaeger, page 60), therefore Jaeger does not teach a "client-centric" approach.

Additionally, while Jaeger and the Netlib Publications both teach using intermediary location servers to decouple a network resource from its location for dynamic lookups for distributed networks such as the Internet (see Request at page 19), this feature appeared in the claims that were rejected as being anticipated by Bietz during the original prosecution. As explained above, Jaeger and the Netlib Publications are cumulative to the Beitz reference.

Given that the combination of Jaeger and the Netlib Publications do not teach or suggest the limitations in independent claims 17 and 18 found allowable during the original prosecution, and the finding that the combination of Jaeger and the Netlib Publications are cumulative to the reference of Bietz that was considered in the original prosecution of the '640 Patent, it is not likely that a reasonable examiner would consider the combination of Jaeger and the Netlib Publications important in determining the patentability of the '640 Patent. With that, the combination of Jaeger and the Netlib Publications are not found to raise a SNQ over claims 17 or 18. Since dependent claims carry with them each of the limitations from the claims from which they depend, it follows that Jaeger does not raise a SNQ over dependent claim 24.

Summary

The Request alleges the following groups of SNQs:

Group I: Jaeger and the Netlib Publications to raise a SNQ over claims 17, 18, and 24;

Group II: the Netlib Publications to raise a SNQ over claims 17, 18, and 24; and

Group III: Jaeger and the Netlib Publications to raise a SNQ over claims 17, 18, and 24.

For the reason discussed above, the Examiner disagrees with each alleged group of SNQs. The references relied upon to establish a SNQ (Jaeger and the Netlib-related Publications) are found to be cumulative to prior art already considered in a prior examination of the '640 Patent.

Conclusion

ALL correspondence relating to this ex parte reexamination proceeding should be directed as follows:

Please mail any communications to:

Attn: Mail Stop "*Ex Parte Reexam*"
Central Reexamination Unit
Commissioner for Patents
P. O. Box 1450 Alexandria VA 22313-1450

Please FAX any communications to:

(571) 273-9900
Central Reexamination Unit

Please hand-deliver any communications to:

Customer Service Window
Attn: Central Reexamination Unit
Randolph Building, Lobby Level
401 Dulany Street
Alexandria, VA 22314

By EFS-Web

Registered users of EFS-Web may alternatively submit such correspondence via electronic filing system EFS-Web, at

<https://efs.uspto.gov/efile/myportal/efs-registered>

EFS-Web offers the benefit of quick submission to the particular area of the Office that needs to act on the correspondence. Also, EFS-Web submissions are "soft scanned" (i.e., electronically uploaded) directly into the official file for the reexamination proceeding, which offers parties the opportunity to review the content of their submissions after the "soft scanning" process is complete.

Any inquiry concerning this communication or earlier communications from the Reexamination Legal Advisor or Examiner, or as to the status of this proceeding, should be directed to the Central Reexamination Unit at telephone number (571) 272-7705.

Signed:

/ERON J SORRELL/
Primary Examiner, Art Unit 3992

Conferees: /JOSEPH R POKRZYWA/
 Primary Examiner, Art Unit 3992

/M.F/
Supervisory Patent Examiner, Art Unit 3992