PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2018/044787	01.08.2018	01.08.2017

International Patent Classification (IPC) or both national classification and IPC
INV. G06F16/23

Applicant
SALESFORCE.COM, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☒ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax: +31 70 340 - 3016	see form PCT/ISA/210	Rameseder, Jonathan Telephone No. +31 70 340-0

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed.

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing:

   a. ☐ forming part of the international application as filed:

      ☐ in the form of an Annex C/ST.25 text file.

      ☐ on paper or in the form of an image file.

   b. ☐ furnished together with the international application under PCT Rule 13*ter*.1(a) for the purposes of international search only in the form of an Annex C/ST.25 text file.

   c. ☐ furnished subsequent to the international filing date for the purposes of international search only:

      ☐ in the form of an Annex C/ST.25 text file (Rule 13*ter*.1(a)).

      ☐ on paper or in the form of an image file (Rule 13*ter*.1(b) and Administrative Instructions, Section 713).

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that forming part of the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. IV Lack of unity of invention

1. ☒ In response to the invitation (Form PCT/ISA/206) to pay additional fees, the applicant has, within the applicable time limit:

   ☒ paid additional fees

   ☐ paid additional fees under protest and, where applicable, the protest fee

   ☐ paid additional fees under protest but the applicable protest fee was not paid

   ☐ not paid additional fees

2. ☐ This Authority found that the requirement of unity of invention is not complied with and chose not to invite the applicant to pay additional fees.

3. This Authority considers that the requirement of unity of invention in accordance with Rule 13.1, 13.2 and 13.3 is

   ☐ complied with

   ☒ not complied with for the following reasons:

   see separate sheet

4. Consequently, this report has been established in respect of the following parts of the international application:

   ☒ all parts.

   ☐ the parts relating to claims Nos.

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	2, 10-16, 18-25
	No:	Claims	1, 3-9, 17
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-25
Industrial applicability (IA)	Yes:	Claims	1-25
	No:	Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

Re Item IV

Lack of unity of invention

1 The application does not meet the requirements of unity of invention in that there are two inventions. The reasons for which the inventions are not so linked as to form a single general inventive concept (Rule 13.1 PCT) are as follows.

2 The following inventions are identified:

2.1 invention 1: claims 1-9 and 17;

2.2 invention 2: claims 10-16 and 18-25.

3 With regard to the first invention, D1 (US 2014/324785 A1) discloses all features of the method of independent claim 1:

3.1 A method for a database system synchronizing current state of the database system among a plurality of nodes configured to handle requests for data of the database system stored in a distributed storage

Figure 3, ref. signs 300, 310, 320a-c, 322a-c: shows database system with multiple nodes and a distributed storage service.
[0036]: "the database tier may support the use of synchronous or asynchronous read replicas in the system, e.g., read-only copies of data on different nodes of the database tier to which read requests can be routed."
[0053]: "database engine head node 320a [...] may route read requests [...] to a read replica and/or various storage nodes within distributed database-optimized storage service 310, [...] receive requested data pages from distributed database-optimized storage service 310"

The data of a database system (current state of the database system) is synchronized between multiple nodes that handle read requests for data from (stored in) a distributed storage service. Importantly, the distributed storage service, the primary/head nodes, and the read replica nodes together (not only the distributed storage service alone) constitute the claimed distributed storage.

3.2 with one of the plurality of nodes being currently active and the other nodes of the plurality of nodes being currently standby nodes, the method comprising:

Figure 3, ref. signs 320a-c, 322a-c
[0020]: "the read replica may be configured to convert (e.g., fail over) into a primary node (e.g., after failure of a primary node) without loss of data."

The nodes of the database system are head/primary nodes (active nodes) and read replica nodes. The read replica nodes are on "stand by" from being a head/primary node.

3.3 `receiving, at the active node, a request to perform a first transaction that includes committing data to the distributed storage; and`

[0033]: "the database tier [...] may include a primary node server, which may also be referred to herein as a database engine head node server, that receives read and/or write requests from various client programs, then parses them and develops an execution plan to carry out the associated database operation(s)."
[0088]: "At 610, a write request that specifies a modification to a data record stored by a database service may be received. For example, the write request [...] by a primary node."
[0089]: "the primary node (client side driver) may generate the log record, which may be indicative of the change to the data record"
[0090]: "the log record may be sent (e.g., by the client side driver of the primary node) to a particular server node (or multiple server nodes) of a distributed storage service that stores a version of the data page that includes the given data record. The server node may then apply the modification from the log record to the actual data page stored by the server node."

The primary/head node (active node) receives a write request (transaction with at least one write operation), creates a log record reflecting the data to be written, and sends it (commits it to) to a node of the distributed storage service which in turn stores the data to be written.

3.4 `in response to receiving the request: committing, by the active node, the data to the distributed storage to update the current state of the database system; and`

see §3.3: all the data, including the newly written data, in the distributed storage reflects/is the state of the database system. Thus, causing the new data to be written to the distributed storage constitutes "updating" the database system's state.

3.5 causing storing, by the active node, of first metadata providing an indication of the commitment in a transaction log stored in the distributed storage,

[0054]: "database engine head node 320a may also include transaction log 340"
[0089]: "the primary node (client side driver) may generate the log record, which may be indicative of the change to the data record"

When writing data, the primary node (active node) generates a log record (metadata providing indication of the commitment). Thus, this log record (metadata) must exist in the primary node's (active node's) memory, and is thus, at least temporarily, stored by the primary node (active node). As the primary node (active node) is part of the distributed storage (see §3.1, last sentence), the log record (metadata) that the primary node stores is stored in the distributed storage.

3.6 wherein the transaction log identifies, to the standby nodes, information for the standby nodes to know the current state of the database system.

[0091]: "a cache invalidation indication may be sent [...] to a plurality of read replicas) indicating that a cached version of the given data record stored in the read replica's cache is stale. [...] the cache invalidation [...] may also include the actual log record that was sent to the storage service (e.g., for application by the read replica to its cached version of the data record)."

The primary node (active node) sends the log record (metadata in transaction log) to read replica nodes (standby nodes) in order for them to update their cache (inform them of the data/state of/in the database system).

4 As D1 discloses all features of the method of claim 1, none of these features can define a contribution over the prior art D1 (Rule 13.1 PCT). This applies mutatis mutandis to the medium of claim 17.

5 Moreover, the additional features of the following dependent claims that belong to the first invention are also disclosed by D1 and therefore cannot define a contribution over the prior art:

5.1 Claims 3: D1 discloses that the primary node notifies a read replica node (standby node) of new log records (modifications of transaction log) by sending said new log records (metadata) to the read replica node (standby node). In response, the read replica node (standby node) updates its cache

(update second metadata at the standby node) by reading the new log records (metadata) such that the cache is up to date for future requests for data (client requests).

[0091]: "a cache invalidation indication may be sent to a read replica (or to a plurality of read replicas) indicating that a cached version of the given data record stored in the read replica's cache is stale. [...] the cache invalidation indication may be a simple notification [...] and [...] may also include the actual log record that was sent to the storage service (e.g., for application by the read replica to its cached version of the data record). [...] for a subsequent request for data corresponding to the stale data, the read replica will know that the date is stale and retrieve it from the storage service instead of from its cache. [...] the read replica may apply the modification specified by the log record to its cache."

5.2 Claim 4: D1 discloses that the data that is stored in the distributed storage service (part of the distriuted storage) is cached on the read replica nodes (standby nodes). Thus, the cached data corresponds to the stored data.

Figure 3, ref. signs 630 and 640

5.3 Claim 5: D1 discloses that the primary/head node (active node) has a log (catalog) with log records (metadata). The primary node (active node) uses these log records (metadata) to update read replica nodes (second nodes). Thus, the log (catalog) with the log records (metadata) is, in some way, shared among the nodes (primary node and also multiple read replica nodes). When the primary node (active node) has a log record in its memory (stores metdata), the primary node thereby updates its log (catalog) and notifies the read replica nodes (standby nodes).

Figure 3, ref. signs 320a, 340, 322a 326a
[0091], see §5.1

5.4 Claim 6: D1 discloses that read replica nodes (standby nodes) update their caches based on log records (metadata) received from a primary node (active node) [0091]. To update their cache, they must implicitly access entries in the cache using a memory address. This memory address is the key, and the data in memory is the value.

5.5 Claim 7: D1 discloses that data is stored in not only in the log records (metadata) but also in the distributed storage service (which is different from and therefore extern to the transaction log in the distributed storage).

Figure 3, ref. sign 310
Figure 6, ref. signs 630 and 640

5.6 Claim 8: D1 discloses that a read replica node (standby node) is converted into (becomes) a primary/head node in case of failure (thus being a high availability application). As the read replica node (standby node) takes over all responsibilities of the primary node (active node), when said new primary node (new active node) receives a write request, it writes the data (commits data of second transaction) and performs the logging (indication of second transaction in transaction log).

[0058]: "one of read replicas 322a, 322b, or 322c may be converted into a new database engine head node (e.g., if the head node fails)."
Figure 9

5.7 Claim 9: as elaborated above (§5.4), nodes (part of the distributed storage) have logs (catalogs) with log records (metadata), whereby said log records (metadata) are implicitly accessed using memory addresses (keys). Said references (keys) point to the node's memory location (physical location in memory that is part of the distributed storage) that stores the corresponding log records (metadata providing indications of updates).

6 D1 does not appear to disclose the additional features of the method of dependent claim 2. Thus, the first invention is defined by these features. In view of the originally filed description ([0015], [0016]), the additional features of said method solve the technical problem of providing an alternative manner to inform standby nodes of the active node's log's state.

7 With regard to the second invention, D1 discloses the following features of the database system of claim 10:

7.1 A database system, comprising:

Figure 3, ref. sign 300

7.2 a plurality of nodes configured to implement a database; and

Figure 3, ref. signs 320a-c, 322a-c

7.3 a distributed storage accessible to the plurality of nodes and configured to store data of the database;

Figure 3, ref. sign 310

7.4 wherein a first of the plurality of nodes is configured to: receive a request to perform a first transaction that includes committing data to the distributed storage; for the first transaction, store a first set of data in the distributed storage; and

see §3.3

7.5 store a first record of the first transaction in a transaction log maintained by the distributed storage,

[0054]: "database engine head node 320a may also include transaction log 340"
[0089]: "the primary node (client side driver) may generate the log record, which may be indicative of the change to the data record"

The primary node generates a log record which means that the log record must be at least temporarily stored in the primary node's memory. As the primary node is part of the distributed storage (see §3.1, last sentence), the log record is maintained by the distributed storage.

7.6 ~~wherein the transaction log defines an ordering in which transactions are performed with respect to the database.~~

8 The database system of claim 10 therefore differs from the database system of D1 in the following distinguishing feature:

8.1 the transaction log defines an ordering in which transactions are performed with respect to the database

9 This distinguishing feature therefore defines the second invention. In view of the originally filed description [0025], said feature solves the technical problem of permitting the claimed database system to revert to one of its previous states.

10 Still with regard to the second invention, D1 discloses the following features of the method of claim 18:

10.1 A method, comprising: maintaining, by a first of a plurality of database nodes of a database system, a cache for data stored in a distributed storage shared among the plurality of database nodes,

Figure 3, ref. signs 310, 320a-c, 322a-c, 326a, 335: shows that the distributed database service is accessible to primary/head nodes and read replica nodes, and thus shared among these nodes. The nodes, including read replica nodes (one of which is the first node), have a cache for the data in the distributed storage service.

10.2 wherein the cache includes an entry for a first key-value pair;

[0091]: "after updating its cache, [...] the read replica may remove the stale cache indication (e.g., from the data structure maintaining a list of the stale cache entries) for that particular data record."
[0100]: "At 820, the read replica's cache may be updated with versions of the data pages stored in the primary node's cache."

The primary/head and read replica nodes have a cache that includes stored data records. Data records must be, implicitly, readable into the memory of the node. The data record's memory address is said record's key, and the data record's content is said record's value.

10.3 reading, by the first database node, a transaction log~~, wherein the transaction log identifies an ordering in which transactions of the database system are committed to the distributed storage~~;

[0088]: "a write request that specifies a modification to a data record stored by a database service may be received. [...] The write request may specify a modification to be made to a given data record stored in a database table."
[0089]: "a log record representing the modification to be made to the data record may be generated."
[0091]: "a cache invalidation indication may be sent [...] to a plurality of read replicas) indicating that a cached version of the given data record stored in the read replica's cache is stale. [...] the cache invalidation [...] may also include the actual log record that was sent to the storage service (e.g., for application by the read replica to its cached version of the data record)."

A log record (part of a transaction log) is sent to a read replica node (first node) for updating its cache. Thus, the read replica node (first node) must read the received log record (part of the transaction log).

10.4 based on the reading, the first database node determining that a second of the plurality of database nodes has committed, to the distributed storage, a first transaction that modifies a value of the first key-value pair; and

see §10.3: after writing (committing) data to the distributed storage, the primary node generates a log record and sends it to the read replica node (first node) which determines that its cached data entries (data records/key value pairs reflecting the data in the distributed storage) are stale (value is modified).

10.5 in response to the determining, the first database node updating the entry included in the cache based on the modified value of the first key-value pair.

[0091]: "the read replica will know that the date is stale and retrieve it from the storage service instead of from its cache. In an embodiment in which the cache invalidation indication includes the actual log record, the read replica may apply the modification specified by the log record to its cache. After doing so, that cache entry (and data record) may no longer be indicated as stale."

The read replica node (first node) updates the stale (modified) cache entry (data record/key value pair).

11 Thus, the method of claim 18 differs from the method of D1 in the following distinguishing feature (the same applies to independent claims 24 and 25):

11.1 the transaction log identifies an ordering in which transactions of the database system are committed to the distributed storage

12 This appears to be the same feature as the distinguishing feature of independent claim 10 (§8.1). Thus, claims 18, 24, 25 and 10 appear to share the same special technical feature and belong to the second invention, wherein this special technical feature defines the second invention.

13 The special technical features of the first and second invention are not the same (§6 versus §8.1/11.1). Additionally, these features solve different technical problems (§6, 9) and therefore cannot be corresponding.

14 Thus, the two inventions do not share any of the same or corresponding special technical features (Rule 13.2 PCT). Consequently, these inventions are not linked by a single general inventive concept and the requirement of unity of invention is not met (Rule 13.1 PCT).

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

A. FIRST INVENTION

15 **Summary**

15.1 With regard to Article 33(1) PCT, the claimed subject matter is not new in the sense of Article 33(2) PCT and does not involve an inventive step in the sense of Article 33(3) PCT.

16 **Prior Art**

16.1 Reference is made to the following documents:

D1 US 2014/324785 A1 (GUPTA ANURAG WINDLASS [US] ET AL) 30 October 2014 (2014-10-30)

D2 Anonymous: "Polling (computer science)", Wikipedia, 11 July 2017 (2017-07-11), pages 1-3, XP055507238, Retrieved from the Internet: URL:https://en.wikipedia.org/w/index.php?title=Polling_(computer_science)&oldid=790118174 [retrieved on 2018-09-14]

17 **Novelty (Article 33(2) PCT)**

17.1 With regard to Article 33(1) PCT, the claimed subject matter is not new in the sense of Article 33(2) PCT for the following reasons.

17.2 As elaborated above, D1 anticipates the subject matter of claims 1, 3-9, and 17 (see §3-5.7).

18 **Inventive Step (Article 33(3) PCT)**

18.1 With regard to Article 33(1) PCT, the claimed subject matter does not involve an inventive step within the meaning of Article 33(3) PCT for the following reasons.

18.2 Claim 2 defines the following additional features:

18.2.1 monitoring, by one of the standby nodes, a catalog of the database system that identifies new transactions that have been committed to the transaction log,

18.2.2 wherein the catalog stores a database schema for the database system; and

18.2.3 prior to reading the transaction log, the standby node determining that the catalog identifies a new transaction committed to the distributed storage.

18.3 The method of claim 2 therefore differs from the closest prior art in that the standby nodes (read replica nodes) monitor a catalog for newly committed transactions (new writes). When a new transaction is committed (new write performed), the standby node (read replica node) learns it from the catalog and reads the transaction log (log records). Additionally, the catalog stores a database schema.

18.4 First, it is not apparent what credible technical effect is supposed to be produced by the feature that the catalog stores a database schema (§18.2.2). The claimed database schema is nowhere used and it is not readily apparent, even in light of the description, what its function has. Thus, said feature does not appear, in the context of the method of claim 2, to solve a technical problem and therefore cannot support an inventive step (Guidelines G-VII, 5.2, second paragraph).

18.5 The remaining distinguishing features (§18.2.1, 18.2.3) produce the technical effect that the active node (primary node) does not need to send a notification to the standby nodes (read replica nodes) in case of a newly committed transaction (write operation). Thus, and in view of the originally filed description ([0015], [0016]), these distinguishing features solve the technical problem of how to, in the context of the method of D1, provide an alternative manner to inform standby nodes of the active node's log's state.

18.6 In the method of D1, informing read replica nodes (standby nodes) of updates is implemented using the widely known observer pattern. Multiple read replica nodes "observe" a primary node, meaning that the primary node notifies the

read replica nodes of updates. A well known alternative to the observer pattern is polling. In this commonly used implementation alternative, the read replica nodes (observers) periodically "poll" (monitor) the primary node for changes.

18.7 Polling is common general knowledge of the skilled person and for instance described in Wikipedia article D2.

p.1, first paragraph: "Polling, or polled operation, in computer science, refers to actively sampling the status of an external device by a client program as a synchronous activity."

18.8 For these reasons, when confronted with the aforementioned objective technical problem, the skilled person would adapt the method of D1 and configure the read replica nodes (standby nodes) to poll (monitor) the primary node (active node) for any updates and read its transaction log in case of new write operations (committed transactions, updates). Thus, the skilled person would adapt the method of D1 and arrive at the subject matter of claim 2 without exercising any inventive skill. For these reasons, said subject matter does not involve an inventive step within the meaning of Article 33(3) PCT.

B. SECOND INVENTION

19 **Summary**

19.1 With regard to Article 33(1) PCT, the claimed subject matter does not involve an inventive step in the sense of Article 33(3) PCT.

20 **Prior Art**

20.1 Reference is made to the following documents:

D1 US 2014/324785 A1 (GUPTA ANURAG WINDLASS [US] ET AL) 30 October 2014 (2014-10-30)

D3 Sujoy Paul: "Optimizing Transactional Replication" In: "Pro SQL Server 2008 Replication", 17 June 2009 (2009-06-17), Apress, Berkeley, CA, XP055534102, ISBN: 978-1-4302-1807-4 pages 737-800, DOI: 10.1007/978-1-4302-1808-1

21 **Inventive Step (Article 33(3) PCT)**

21.1 With regard to Article 33(1) PCT, the claimed subject matter does not involve an inventive step within the meaning of Article 33(3) PCT for the following reasons.

21.2 D1 is considered to be the prior art closest to the subject matter of claim 10. As elaborated above (§7), the database system of claim 10 differs from the database system of D1 in the following distinguishing feature:

21.2.1 `the transaction log defines an ordering in which transactions are performed with respect to the database`

21.3 This feature produces the technical effect that the logged records define a temporal sequence. In view of the originally filed description [0025], the distinguishing feature solves the technical problem of how to permit the database system of D1 to revert to one of its previous states.

21.4 Standard textbook chapter D3, which exemplifies the common general knowledge of the skilled person before priority date, discloses the distinguishing feature, namely that records in transaction logs have log sequence numbers that are arranged from small to large and thus define an order (p.738, Figure 18-1, n-1 to n+2; also see p.738, fifth paragraph). D3 also discloses that this feature improves recovery (reverting to a previous state).

21.5 Thus, when confronted with the objective technical problem, the skilled person would use his common general knowledge to adapt the method of D1 and arrive at the subject matter of claim 10 without exercising any inventive skill. For this reason, with regard to Article 33(1) PCT, said subject matter is not inventive in the sense of Article 33(3) PCT. In view of the additional feature mapping provided above (§10), this objection applies mutatis mutandis to claims 18, 24, and 25.

21.6 The subject matter of the following claims is also not inventive in the sense of Article 33(3) PCT:

21.6.1 Claim 11: D1 discloses that a primary node (first node) sends the log record (metadata) to a read replica node (second node) in order for said read replica node (second node) to update its cache (metadata indicative of current database state) (see §3.6). Importantly, the distributed storage service, the primary/head nodes, and the read replica nodes together (not only the distributed storage service alone) constitute the claimed distributed storage

(see §3.1). Thus, when receiving the log record from the primary node (first node), the read replica node (second node) reads the log record from the distributed storage (which includes the primary node).

21.6.2 Claims 12 and 13: D1 discloses that the transaction log includes multiple records ([0054]: "database engine head node 320a may also include transaction log 340 [...] which may [...] track the status of **various** transactions"). Additionally, outsourcing functionality on a separate server as claimed (metadata server) is an obvious design option.

21.6.3 Claim 14: as aforementioned, D1 discloses processing multiple transactions [0054].

21.6.4 Claim 15: D1 discloses a primary node (active node; see §3.2).

21.6.5 Claim 16: D1 discloses a number of read replicas that define a high-availability system (see §3.2).

21.6.6 Claim 19: see §18.3-18.8.

21.6.7 Claim 20: D1 discloses that a database client (client device) accesses the head node (first database node) to write data (second key-value pair) into the distributed storage.
Figure 3, ref. signs 350a-350n, 320
[0052]: "one or more of database clients 350a-350n may access a database head node 320"
[0053]: "execution component 305 may return query responses to database client 350a, which may include write acknowledgements"

21.6.8 Claim 21: updating data in response to a change in data is common general knowledge of the skilled person.

21.6.9 Claim 22: D1 discloses that cache entries are marked as stale (invalidated).
[0091]: "a cache invalidation indication may be sent [...] to a plurality of read replicas) indicating that a cached version of the given data record stored in the read replica's cache is stale."

21.6.10 Claim 23: D1 discloses that nodes serve the read requests of clients.

[0052]: "database clients 350a-350n may access a database head node 320 [...] and/or a read replica (e.g., read replica 322a, 322b, or 322c) via network 360"
[0053]: " each read replica may receive read requests"

Re Item VIII
Certain observations on the international application

22 **Clarity (Article 6 PCT)**

22.1 The application does not meet the requirements of Article 6 PCT because the claimed subject matter is not clear. Claim 11 refers to a database system of claim 8. However, claim 8 defines a method, not a database system. Thus, the database system to which claim 11 refers lacks precedence (Article 6 PCT).

Re Item VII
Certain defects in the international application

23 Contrary to the requirements of Rule 5.1(a)(ii) PCT, the relevant background art disclosed in D1 is not mentioned in the description, nor is this document identified therein.

24 The claims do not contain reference signs to the drawings as required by Rule 6.2(b) PCT.