

# **METHOD, SYSTEM AND APPARATUS FOR MANAGING PRIMARY AND SECONDARY DATABASES**

## **TECHNICAL FIELD**

[0001] The present application relates to the field of information security  
5 technologies, and in particular, to a method, system and apparatus for managing  
primary and secondary databases.

## **BACKGROUND**

[0002] Databases are key infrastructures in the fields such as finance, commerce,  
transportation, and even the entire society. Continuous availability of the database can  
10 ensure that services can be continuously provided for a user in the fields such as  
financial and commercial fields. To prevent a single database from being faulty and  
improve database availability, a primary database and a secondary database are  
usually configured. When the primary database is faulty, the primary database is  
switched to a secondary database to continue to serve the user.

15 [0003] In related technologies, switching between primary and secondary  
databases can be automatic. That is, a monitoring system is deployed for the primary  
database, and once finding that the primary database is abnormal, the monitoring  
system sends an alarm and a command for switching between primary and secondary  
databases.

20 [0004] However, in the previous method for switching between primary and  
secondary databases, the following case may occur due to incorrect judgment: The  
primary database is switched when the primary database is not faulty, or the primary  
database is faulty but is not switched. Consequently, normal services provided for the  
user are affected.

25

## **SUMMARY**

[0005] The present application intends to alleviate one of the previous technical

problems at least to some extent.

**[0006]** Therefore, a first objective of the present application is to provide a method for managing primary and secondary databases. According to the method, if a primary database does not send a lock renewal request before a lock held by the primary database expires, it is determined that the lock held by the primary database is invalid and the primary database cannot normally serve a user, so that a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

**[0007]** A second objective of the present application is to provide a lock arbitration server.

**[0008]** A third objective of the present application is to provide a primary database server.

**[0009]** A fourth objective of the present application is to provide a secondary database server.

**[0010]** A fifth objective of the present application is to provide a system for managing primary and secondary databases.

**[0011]** To achieve the previous objective, a method for managing primary and secondary databases in an implementation of a first aspect of the present application includes: determining whether a lock held by a primary database expires, where the primary database and secondary databases share the lock; if determining that the lock held by the primary database has expired, determining whether a lock renewal request of the primary database is received; and if the lock renewal request of the primary database is not received, selecting one of the secondary databases as a new primary database, and controlling the primary database to be switched to a secondary database.

**[0012]** According to the method for managing primary and secondary databases in this implementation of the present application, it is determined whether the lock held by the primary database expires. If it is determined that the lock held by the primary database has expired, it is determined whether the lock renewal request of the primary database is received. If the lock renewal request is not received, one of the secondary databases is selected as a new primary database, and the primary database is controlled to be switched to a secondary database. According to the method, if the primary database does not send the lock renewal request before the lock held by the primary database expires, it is determined that the lock held by the primary database is invalid and the primary database cannot normally serve a user, so that a secondary

database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

[0013] In addition, the method for managing primary and secondary databases in this implementation of the present application further has the following additional technical features:

[0014] In an implementation of the present application, after the determining that the lock held by the primary database has expired, the method further includes: sending a lock expiration notification to the secondary database; and receiving a lock request sent by the secondary database based on the lock expiration notification, and recording a receiving time of the lock request.

[0015] In an implementation of the present application, the selecting one of the secondary databases as a primary database further includes: selecting, as the primary database, a secondary database with the earliest receiving time.

[0016] In an implementation of the present application, a priority of the primary database is higher than a priority of the secondary database.

[0017] In an implementation of the present application, an update cycle of the lock is T1, the primary database sends, at a cycle T2, a lock request for querying a status of the lock, and the secondary database sends, at a cycle T3, a lock request for querying the status of the lock, where T2 is less than T1, and T3 is greater than or equal to T1.

[0018] In an implementation of the present application, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary database and the disaster recovery database are located in different data centers.

[0019] In an implementation of the present application, after the original primary database is restored, the method further includes: receiving a lock request sent by the original primary database, and controlling the lock to be held by the original primary database after the lock expires, so that the original primary database is restored to a primary database, and the current primary database is restored to a secondary database.

[0020] In an implementation of the present application, after the selecting one of the secondary databases as a new primary database, the method further includes: continuing to determine whether a lock held by the current primary database expires; if the lock held by the current primary database has expired, determining whether lock

requests of the original primary database, the current primary database, and other secondary databases are received; and if the lock request of the original primary database is received, restoring the original primary database to a primary database, and restoring the current primary database to a secondary database; or if the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, maintaining the current primary database as a primary database; or if the lock requests of the original primary database and the current primary database are not received and the lock requests of the other secondary databases are received, selecting one of the other secondary databases as a primary database, and restoring the current primary database to a secondary database.

**[0021]** To achieve the previous objective, a lock arbitration server in an implementation of a second aspect of the present application includes: a first determining module, configured to determine whether a lock held by a primary database expires, where the primary database and secondary databases share the lock; a second determining module, configured to: when the lock held by the primary database has expired, determine whether a lock renewal request of the primary database is received; and a first processing module, configured to: when the lock renewal request of the primary database is not received, select one of the secondary databases as a new primary database, and control the primary database to be switched to a secondary database.

**[0022]** The lock arbitration server in this implementation of the present application determines whether the lock held by the primary database expires. If determining that the lock held by the primary database has expired, the server determines whether the lock renewal request of the primary database is received. If the lock renewal request is not received, the server selects one of the secondary databases as a new primary database, and controls the primary database to be switched to a secondary database. If the primary database does not send the lock renewal request before the lock held by the primary database expires, the server determines that the lock held by the primary database is invalid and the primary database cannot normally serve a user, so that a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

**[0023]** In addition, the lock arbitration server in this implementation of the present

application has the following additional technical features:

**[0024]** In an implementation of the present application, the first processing module is configured to select, as the primary database, a secondary database with the earliest receiving time.

5 **[0025]** In an implementation of the present application, a priority of the primary database is higher than a priority of the secondary database.

**[0026]** In an implementation of the present application, an update cycle of the lock is T1, the primary database sends, at a cycle T2, a lock request for querying a status of the lock, and the secondary database sends, at a cycle T3, a lock request for  
10 querying the status of the lock, where T2 is less than T1, and T3 is greater than or equal to T1.

**[0027]** In an implementation of the present application, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary  
15 database and the disaster recovery database are located in different data centers.

**[0028]** In an implementation of the present application, the lock arbitration server further includes a first receiving module, configured to: after the original primary database is restored, receive a lock request sent by the original primary database, where the first processing module is further configured to control the lock to be held  
20 by the original primary database after the lock expires, so that the original primary database is restored to a primary database, and the current primary database is restored to a secondary database.

**[0029]** In an implementation of the present application, the first determining module is further configured to continue to determine whether a lock held by current  
25 primary database expires; the second determining module is further configured to: when the lock held by the current primary database has expired, determine whether lock requests of the original primary database, the current primary database, and other secondary databases are received; a second processing module is configured to: when the lock request of the original primary database is received, restore the original  
30 primary database to a primary database, and restore the current primary database to a secondary database; a third processing module is configured to: when the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, maintain the current primary database as a primary database; and a fourth processing module is configured

to: when the lock requests of the original primary database and the current primary database are not received and the lock requests of the other secondary databases are received, select one of the other secondary databases as a primary database, and restore the current primary database to a secondary database.

5   **[0030]**   To achieve the previous purpose, a primary database server in an implementation of a third aspect of the present application includes: a second sending module, configured to send a lock renewal request to a lock arbitration server when a lock held by a primary database expires, so that the lock arbitration server selects one of secondary databases as a new primary database when the lock arbitration server  
10 does not receives the lock renewal request; and a switching module, configured to control the primary database to be switched to a secondary database.

**[0031]**   When the lock held by the primary database expires, the primary database server in this implementation of the present application sends the lock renewal request to the lock arbitration server, so that the lock arbitration server selects one of the  
15 secondary databases as a new primary database when the lock arbitration server does not receives the lock renewal request; and switches the primary database to a secondary database. When a validity period of the lock held by the primary database cannot be extended, the primary database server switches the primary database to a secondary database, and a secondary database is selected as a new primary database,  
20 thereby improving a speed and accuracy of switching between primary and secondary databases, ensuring that the database normally serves a user, and improving user experience.

**[0032]**   In addition, the primary database server in this implementation of the present application has the following additional technical features:

25 **[0033]**   In an implementation of the present application, a priority of the primary database is higher than a priority of the secondary database.

**[0034]**   In an implementation of the present application, an update cycle of the lock is T1, the primary database server sends, at a cycle T2, a lock request for querying a status of the lock, and the secondary database sends, at a cycle T3, a lock  
30 request for querying the status of the lock, where T2 is less than T1, and T3 is greater than or equal to T1.

**[0035]**   In an implementation of the present application, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary

database and the disaster recovery database are located in different data centers.

**[0036]** In an implementation of the present application, the primary database server further includes: a third sending module, configured to: after the original primary database is restored, send a lock request to the lock arbitration server, so that  
 5 the lock arbitration server receives the lock request sent by the original primary database, and controls the lock to be held by the original primary database after the lock expires, so that the original primary database is restored to a primary database, and the current primary database is restored to a secondary database.

**[0037]** To achieve the previous objective, a secondary database server in an  
 10 implementation of a fourth aspect of the present application includes: a fourth sending module, configured to send a lock request to a lock arbitration server; a second receiving module, configured to receive a lock acknowledgement message sent by the lock arbitration server; and a fifth processing module, configured to switch a secondary database to a new primary database based on the lock acknowledgement  
 15 message.

**[0038]** The secondary database server in this implementation sends the lock request to the lock arbitration server, receives a lock acknowledgement message sent by the lock arbitration server, and switches a secondary database to a new primary database based on the lock acknowledgement message. When a primary database  
 20 cannot extend a validity period of a lock held by the primary database, the secondary database server selects a secondary database as a new primary database, and the primary database is switched to a secondary database, thereby improving a speed and accuracy of switching between primary and secondary databases, ensuring that the database normally serves a user, and improving user experience.

**[0039]** The secondary database server in this implementation of the present application has the following additional technical features:

**[0040]** In an implementation of the present application, the secondary database server includes: a third receiving module, configured to receive a lock expiration notification sent by the lock arbitration server after the lock arbitration server  
 30 determines that a lock held by a primary database has expired, where the fourth sending module sends the lock request to the lock arbitration server based on the lock expiration notification.

**[0041]** In an implementation of the present application, a priority of the primary database is higher than a priority of the secondary database.

[0042] In an implementation of the present application, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary database and the disaster recovery database are located in different data centers.

5 [0043] In an implementation of the present application, the secondary database server further includes: a sixth processing module, configured to control the current primary database to be restored to a secondary database after the lock arbitration server receives a lock request sent by the original primary database, and controls the lock to be held by the original primary database after the lock expires, so that the  
10 original primary database is restored to a primary database.

[0044] To achieve the previous purpose, a system for managing primary and secondary databases in an implementation of a fifth aspect of the present application includes the lock arbitration server in the implementation of the second aspect of the present application, the primary database server in the implementation of the third  
15 aspect of the present application, and the secondary database server in the implementation of the fourth aspect of the present application.

[0045] Some additional aspects and advantages of the present application are provided in the following descriptions, and are apparent in the following descriptions or understandable through practice of the present application.

20

## **BRIEF DESCRIPTION OF DRAWINGS**

[0046] The previous and/or additional aspects and advantages of the present application are apparent and comprehensible in the following descriptions of implementations with reference to the accompanying drawings.

[0047] FIG. 1 is a flowchart illustrating a method for managing primary and  
25 secondary databases, according to an implementation of the present application;

[0048] FIG. 2 (a) to FIG. 2 (c) are diagrams illustrating primary and secondary database deployment, according to a specific implementation of the present application;

[0049] FIG. 3 is a flowchart illustrating a method for managing primary and  
30 secondary databases, according to another implementation of the present application;

[0050] FIG. 4 is a flowchart illustrating a method for managing primary and secondary databases, according to still another implementation of the present

application;

[0051] FIG. 5 is a schematic structural diagram illustrating a lock arbitration server, according to an implementation of the present application;

[0052] FIG. 6 is a schematic structural diagram illustrating a lock arbitration server, according to a specific implementation of the present application;

[0053] FIG. 7 is a schematic structural diagram illustrating a lock arbitration server, according to another implementation of the present application;

[0054] FIG. 8 is a schematic structural diagram illustrating a lock arbitration server, according to still another implementation of the present application;

10 [0055] FIG. 9 is a schematic structural diagram illustrating a primary database server, according to an implementation of the present application;

[0056] FIG. 10 is a schematic structural diagram illustrating a primary database server, according to a specific implementation of the present application;

[0057] FIG. 11 is a schematic structural diagram illustrating a secondary database server, according to an implementation of the present application;

[0058] FIG. 12 is a schematic structural diagram illustrating a secondary database server, according to a specific implementation of the present application;

[0059] FIG. 13 is a schematic structural diagram illustrating a secondary database server, according to another implementation of the present application; and

20 [0060] FIG. 14 is a flowchart illustrating a system for managing primary and secondary databases, according to an implementation of the present application.

## DESCRIPTION OF IMPLEMENTATIONS

[0061] The implementations of the present application are described in detail below, and examples of the implementations are shown in the accompanying drawings. The same or similar reference numerals represent the same or similar elements or elements having the same or similar functions. The implementations described below with reference to the accompanying drawings are examples, are intended to explain the present application, and should not be construed as a limitation on the present application.

30 [0062] With reference to the accompanying drawings, the following describes a method, system and apparatus for managing primary and secondary databases according to an implementation of the present application.

**[0063]** FIG. 1 is a flowchart illustrating a method for managing primary and secondary databases, according to an implementation of the present application. As shown in FIG. 1, the method for managing primary and secondary databases includes the steps below.

5 **[0064]** S110. Determine whether a lock held by a primary database expires, where the primary database and secondary databases share the lock.

**[0065]** S120. If determining that the lock held by the primary database has expired, determine whether a lock renewal request of the primary database is received.

10 **[0066]** It can be understood that, incorrect switching between primary and secondary databases may occur when a monitoring system monitors and controls switching between primary and secondary databases, for example, a primary database that is not faulty is switched or a primary database is faulty but is not switched. To alleviate the problem, in the method for managing primary and secondary databases in  
15 this implementation of the present application, an external arbitration mechanism is introduced and is used to accurately determine whether the primary database is faulty, thereby accurately controlling switching between primary and secondary databases.

**[0067]** There can be a plurality of external arbitration mechanisms based on different specific application scenarios. In this implementation of the present  
20 application, an example that the arbitration mechanism is a lock service is used for detailed description.

**[0068]** Specifically, the primary database and the secondary databases share a lock. Both the primary database and the secondary databases contend for ownership of the lock, and a database that obtains the lock is a primary database. In addition, because  
25 the lock has a mutual exclusion feature, only one database can be a primary database at the same time.

**[0069]** The primary database needs to continuously update the lock to extend a validity period (usually tens of seconds) of the lock before the validity period expires. If the primary database is faulty or is performing an upgrade or maintenance operation  
30 within the validity period, the primary database cannot complete updating the lock to extend the validity period of the lock. Therefore, another secondary database seizes the lock and is upgraded to a primary database, and the primary database is downgraded to a secondary database.

**[0070]** Specifically, in actual applications, to keep the primary database

unchanged, that is, to prevent the primary database from being switched to a secondary database when the primary database is not faulty or performs an upgrade or maintenance operation, a priority of the primary database is higher than a priority of the secondary database, thereby ensuring that a database that holds the lock is the primary database when the primary database is not faulty or performs an upgrade or maintenance operation.

**[0071]** It is worthwhile to note that, there are a plurality of methods for deploying primary and secondary databases, for example, a method for deploying one primary database and one secondary database, or a method for deploying one primary database and a plurality of secondary databases.

**[0072]** For ease of description, this implementation of the present application is described by using an example that the method for deploying primary and secondary databases is a three-data-centers-in-two-locations method. That is, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in the same data center, and the primary database and the disaster recovery database are located in different data centers.

**[0073]** Further, to determine whether the primary database that holds the lock is faulty, it is determined whether the primary database extends the validity period of the lock within the validity period.

**[0074]** Specifically, it is determined whether the lock held by the primary database expires. If it is determined that the lock held by the primary database has expired, it is determined whether the lock renewal request of the primary database is received.

**[0075]** S130. If the lock renewal request of the primary database is not received, select one of the secondary databases as a new primary database, and control the primary database to be switched to a secondary database.

**[0076]** Specifically, if the lock renewal request of the primary database is not received, it indicates that the primary database has a network failure or a power failure and cannot complete updating the lock. Therefore, to normally serve a user, one of the secondary databases is selected as a new primary database, and the primary database is controlled to be switched to a secondary database.

**[0077]** It is worthwhile to note that, one of the secondary databases can be selected as a primary database by using different methods based on different specific

application scenarios. For example, the secondary database can actively query, at a specific cycle, whether the primary database sends the lock renewal request when the lock of the primary database expires, to extend the validity period of the lock. If no, the secondary database sends a lock request, to quickly obtain the lock and switch to a primary database.

**[0078]** To describe the method for managing primary and secondary databases in this implementation of the present application more clearly, with reference to FIG. 2 (a) to FIG. 2 (b), an example that a database deployment method is a three-data-centers-in-two-locations method is used below for description. That is, as shown in FIG. 2 (a), two databases (the primary database and the hot standby database) are deployed in a local primary equipment room and a local hot standby equipment room, and one independent database (the disaster recovery database) is deployed in a remote disaster recovery equipment room.

**[0079]** In this example, the external arbitration mechanism service GOS can be deployed across areas, can tolerate a complete failure of an equipment room and a layout area or a network failure, and can always provide an uninterrupted lock service.

**[0080]** FIG. 2 (b) is a schematic diagram illustrating switching between primary and secondary databases, according to an implementation of the present application.

**[0081]** As shown in FIG. 2 (b), if the primary equipment room is entirely faulty, for example, if the primary equipment room is damaged or a network is faulty, the primary database cannot send the lock renewal request, and therefore cannot extend the validity period of the lock held by the primary database. Consequently, the lock held by the primary database is invalid. This process usually lasts for several tens of seconds. In this period, the primary database cannot normally serve the user.

**[0082]** After the validity period of the lock expires, the hot standby equipment room usually gets a new lock, so that the hot standby database becomes a new primary database. The original primary database is automatically downgraded to a secondary database, and the hot standby database serves the user as a new primary database. The disaster recovery equipment room obtains the latest data from the new primary database and locally synchronizes the data.

**[0083]** In conclusion, according to the method for managing primary and secondary databases in this implementation of the present application, it is determined

whether the lock held by the primary database expires. If it is determined that the lock held by the primary database has expired, it is determined whether the lock renewal request of the primary database is received. If the lock renewal request is not received, one of the secondary databases is selected as a new primary database, and the primary database is controlled to be switched to a secondary database. According to the method, if the primary database does not send the lock renewal request before the lock held by the primary database expires, it is determined that the lock held by the primary database is invalid and the primary database cannot normally serve the user, so that a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

**[0084]** Based on the previous implementation, in the method for managing primary and secondary databases, a method for switching between primary and secondary databases can fall into two types: active switching and passive switching, and detailed descriptions are as follows:

**[0085]** For example, to enable the primary database to update the lock in a timely manner and to enable other secondary databases to quickly obtain the lock when the primary database is faulty, the primary database can send, at a relatively short cycle within an update cycle of the lock, a request for querying a status of the lock, to update the lock in a timely manner when the lock is close to expiring and keep the primary database unchanged.

**[0086]** In addition, the secondary database actively sends, at a relatively long cycle, a lock request for querying the status of the lock. Therefore, when the lock of the primary database is invalid, the secondary database can quickly obtain the lock and actively switch to a new primary database, so that services provided for the user are not affected.

**[0087]** For example, the update cycle of the lock is  $T1$ , the primary database sends, at a cycle  $T2$ , a lock request for querying the status of the lock, and the secondary database sends, at a cycle  $T3$ , a lock request for querying the status of the lock, where  $T2$  is less than  $T1$ , and  $T3$  is greater than or equal to  $T1$ .

**[0088]** For another example, if the secondary database does not always actively send, at a specific cycle, a lock request for querying a status of the lock, when the lock held by the primary database is invalid, a lock expiration notification is actively sent to each of other secondary databases, so that the secondary database sends the lock request based on the lock expiration notification, obtains the lock, and becomes a new

primary database.

[0089] Specifically, FIG. 3 is a flowchart illustrating a method for managing primary and secondary databases, according to another implementation of the present application. As shown in FIG. 3, the method includes the steps below.

5 [0090] S310. Determine whether a lock held by a primary database expires, where the primary database and secondary databases share the lock.

[0091] S320. If determining that the lock held by the primary database has expired, send a lock expiration notification to the secondary database.

[0092] S330. Receive a lock request sent by the secondary database based on the  
10 lock expiration notification, and record a receiving time of the lock request.

[0093] S340. Determine whether a lock renewal request of the primary database is received.

[0094] S350. If the lock renewal request of the primary database is not received, select, as a primary database, a secondary database with the earliest receiving time.

15 [0095] It can be understood that data is synchronized in real time between the primary database and the secondary database. A synchronization speed depends on a backup speed of the database. Generally, a delay of an intra-city equipment room is several milliseconds, and a delay of a remote equipment room is hundreds of milliseconds.

20 [0096] In addition, if the primary database actively initiates switching, for example, performs version upgrade or offline maintenance, data write can be stopped before the primary database gives up an identity of the primary database. A hot standby equipment room obtains the lock, and a hot standby database can synchronize and complete updating all the data. Therefore, generally, data in the hot standby  
25 database is relatively complete, so that the hot standby database is preferably used as a next primary database.

[0097] Specifically, a secondary database closer to the primary database is determined based on the receiving time of the lock request. That is, distances between the primary database and different secondary databases can be different. For example,  
30 a hot standby database is usually closer to the primary database, and a disaster recovery database is usually farther from the primary database. Therefore, receiving times of lock requests of the two databases are different, and the receiving time of the request of the hot standby database is earlier.

[0098] Further, if the lock renewal request of the primary database is not received,

it indicates that the current database is faulty or is performing an upgrade or maintenance operation. Therefore, to normally serve a user as soon as possible, a secondary database with the earliest receiving time is selected as a primary database.

**[0099]** In addition, it should be understood that when the primary database and the hot standby database that are located in the same data center are faulty due to a disaster such as an earthquake, the method for managing primary and secondary databases in this implementation of the present application can still be implemented.

**[0100]** As shown in FIG. 2 (c), when equipment rooms of the primary database and the hot standby database are both faulty, the lock of the primary database becomes invalid, and the hot standby database cannot obtain the lock. Therefore, the disaster recovery database obtains the lock and becomes a new primary database, and continues to serve the user. The entire process does not require manual participation.

**[0101]** In conclusion, according to the method for managing primary and secondary databases in this implementation of the present application, when the lock of the primary database expires and is invalid, the secondary database quickly obtains the lock and becomes a new primary database, thereby ensuring that normal services are still provided for the user in a relatively short time, and improving user experience.

**[0102]** In actual applications, after related operations such as maintenance and upgrade are performed on the faulty original primary database, the original primary database can continue to serve the user. Therefore, the method for managing primary and secondary databases in this implementation of the present application further includes: receiving a lock request from the original primary database after the original primary database is restored, to switch the original primary database to a primary database.

**[0103]** Specifically, in this example, after one of the secondary databases is selected as a new primary database, validity of a lock of the current primary database and the lock request of the original primary database are always detected. When the lock of the current primary database is invalid, one of other secondary databases is selected as a new primary database again. Alternatively, after the lock request of the original primary database is received, because a priority of the original primary database is higher, the original primary database is switched to a primary database, and the current primary database is downgraded to a secondary database.

**[0104]** FIG. 4 is a flowchart illustrating a method for managing primary and secondary databases, according to still another implementation of the present application. As shown in FIG. 4, the method for managing primary and secondary databases includes the steps below.

5 **[0105]** S410. Continue to determine whether a lock held by a current primary database expires.

**[0106]** S420. If the lock held by the current primary database has expired, determine whether lock requests of an original primary database, the current primary database, and other secondary databases are received.

10 **[0107]** S430. If the lock request of the original primary database is received, restore the original primary database to a primary database, and restore the current primary database to a secondary database.

**[0108]** Specifically, after the original primary database is restored, the original primary database sends the lock request. Because a priority of the original primary database is higher, after the lock request is received, the original primary database is restored to a primary database, and the current primary database is restored to a secondary database.

**[0109]** S440. If the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, maintain the current primary database as a primary database.

**[0110]** Specifically, when the lock request of the original primary database is not received and the current primary database normally updates the lock, the current primary database is maintained as a primary database.

25 **[0111]** S450. If the lock requests of the original primary database and the current primary database are not received and the lock requests of the other secondary databases are received, select one of the other secondary databases as a primary database, and restore the current primary database to a secondary database.

**[0112]** Specifically, if the lock requests of the original primary database and the current primary database are not received, it indicates that the original primary database is not restored, and the current primary database is also faulty or needs to perform an upgrade operation. Therefore, to normally serve a user, one of the other secondary databases that send the lock requests needs to be selected as a primary database, and the current primary database is restored to a secondary database.

**[0113]** In conclusion, according to the method for managing primary and

secondary databases in this implementation of the present application, after selecting one of the secondary databases as a new primary database, a lock arbitration server continues to determine whether the current primary database receives the lock request of the original primary database and whether the current primary database updates the lock. When the lock request of the original primary database is received, the server restores the original primary database to a primary database, and restores the current primary database to a secondary database. If the lock request of the original primary database is not received, the server determines whether the current primary database updates the lock within a validity period. If no, the server selects one of the other secondary databases as a primary database, and restores the current primary database to a secondary database. According to the method, it is ensured that the original primary database is restored to a primary database after the original primary database is restored, to better serve the user, thereby improving practicality of the method for managing primary and secondary databases in the present application.

**[0114]** To implement the previous implementation, the present application further provides a lock arbitration server. FIG. 5 is a schematic structural diagram illustrating a lock arbitration server, according to an implementation of the present application.

**[0115]** As shown in FIG. 5, the lock arbitration server includes: a first determining module 510, configured to determine whether a lock held by a primary database expires, where the primary database and secondary databases share the lock; a second determining module 520, configured to: when the lock held by the primary database has expired, determine whether a lock renewal request of the primary database is received; and a first processing module 530, configured to: when the lock renewal request of the primary database is not received, select one of the secondary databases as a new primary database, and control the primary database to be switched to a secondary database.

**[0116]** Specifically, in actual applications, to keep the primary database unchanged, that is, to prevent the primary database from being switched to a secondary database when the primary database is not faulty or performs an upgrade or maintenance operation, a priority of the primary database is higher than a priority of the secondary database, thereby ensuring that a database that holds the lock is the primary database when the primary database is not faulty or performs an upgrade or maintenance operation.

**[0117]** It is worthwhile to note that, there are a plurality of methods for deploying

primary and secondary databases, for example, a method for deploying one primary database and one secondary database, or a method for deploying one primary database and a plurality of secondary databases.

**[0118]** For ease of description, this implementation of the present application is described by using an example that the method for deploying primary and secondary databases is a three-data-centers-in-two-locations method. That is, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in the same data center, and the primary database and the disaster recovery database are located in different data centers.

**[0119]** Further, to determine whether the primary database that holds the lock is faulty, it is determined whether the primary database extends the validity period of the lock within the validity period.

**[0120]** Specifically, the first determining module 510 determines whether the lock held by the primary database expires. If it is determined that the lock held by the primary database has expired, the second determining module 520 determines whether the lock renewal request of the primary database is received.

**[0121]** Further, if the lock renewal request of the primary database is not received, it indicates that the primary database has a network failure or a power failure and cannot complete updating the lock. Therefore, to normally serve a user, the first processing module 530 selects one of the secondary databases as a new primary database, and controls the primary database to be switched to a secondary database.

**[0122]** It is worthwhile to note that, the first processing module 530 can select one of the secondary databases as a primary database by using different methods based on different specific application scenarios. For example, the secondary database can actively query, at a specific cycle, whether the primary database sends the lock renewal request when the lock of the primary database expires, to extend the validity period of the lock. If no, the secondary database sends a lock request, to quickly obtain the lock and switch to a primary database.

**[0123]** In conclusion, the lock arbitration server in this implementation of the present application determines whether the lock held by the primary database expires. If determining that the lock held by the primary database has expired, the server determines whether the lock renewal request of the primary database is received. If

the lock renewal request is not received, the server selects one of the secondary databases as a new primary database, and controls the primary database to be switched to a secondary database. If the primary database does not send the lock renewal request before the lock held by the primary database expires, the lock arbitration server determines that the lock held by the primary database is invalid and the primary database cannot normally serve the user, so that a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

**[0124]** Based on the previous implementation, in the method for managing primary and secondary databases, a method for switching between primary and secondary databases can fall into two types: active switching and passive switching, and detailed descriptions are as follows:

**[0125]** For example, to enable the primary database to update the lock in a timely manner and to enable other secondary databases to quickly obtain the lock when the primary database is faulty, the primary database can send, at a relatively short cycle within an update cycle of the lock, a request for querying a status of the lock, to update the lock in a timely manner when the lock is close to expiring and keep the primary database unchanged.

**[0126]** In addition, the secondary database actively sends, at a relatively long cycle, a lock request for querying the status of the lock. Therefore, when the lock of the primary database is invalid, the secondary database can quickly obtain the lock and actively switch to a new primary database, so that services provided for the user are not affected.

**[0127]** For example, the update cycle of the lock is  $T1$ , the primary database sends, at a cycle  $T2$ , a lock request for querying the status of the lock, and the secondary database sends, at a cycle  $T3$ , a lock request for querying the status of the lock, where  $T2$  is less than  $T1$ , and  $T3$  is greater than or equal to  $T1$ .

**[0128]** For another example, if the secondary database does not always actively send, at a specific cycle, a lock request for querying a status of the lock, when the lock held by the primary database is invalid, a lock expiration notification is actively sent to each of other secondary databases, so that the secondary database sends the lock request based on the lock expiration notification, obtains the lock, and becomes a new primary database.

**[0129]** Specifically, FIG. 6 is a schematic structural diagram illustrating a lock

arbitration server, according to a specific implementation of the present application. As shown in FIG. 6, on the basis of FIG. 5, the lock arbitration server further includes a first sending module 540 and a record module 550.

[0130] Specifically, the first sending module sends a lock expiration notification  
5 to the secondary database after the first determining module 510 determines that the lock held by the primary database has expired. Then, the record module 550 receives a lock request sent by the secondary database based on the lock expiration notification, and records a receiving time of the lock request.

[0131] Further, if the lock renewal request of the primary database is not received,  
10 it indicates that the current database is faulty or is performing an upgrade or maintenance operation. Therefore, to normally serve the user as soon as possible, the first processing module 530 selects, as a primary database, a secondary database with the earliest receiving time.

[0132] In conclusion, according to the lock arbitration server in this  
15 implementation of the present application, when the lock of the primary database expires and is invalid, the secondary database quickly obtains the lock and becomes a new primary database, thereby ensuring that normal services are still provided for the user in a relatively short time, and improving user experience.

[0133] In actual applications, after related operations such as maintenance and  
20 upgrade are performed on the faulty original primary database, the original primary database can continue to serve the user. Therefore, the lock arbitration server in this implementation of the present application is further configured to receive a lock request from the original primary database after the original primary database is restored, to switch the original primary database to a primary database.

[0134] FIG. 7 is a schematic structural diagram illustrating a lock arbitration  
25 server, according to another implementation of the present application. As shown in FIG. 7, on the basis of FIG. 5, the lock arbitration server further includes a first receiving module 560. Specifically, in this example, after the first processing module 530 selects one of the secondary databases as a new primary database, the first  
30 receiving module 560 always receives a lock request of the original primary database. After the first receiving module 560 receives the lock request of the original primary database, because the priority of the original primary database is higher, the first processing module 530 switches the original primary database to a primary database, and downgrades the current primary database to a secondary database.

**[0135]** In an implementation, FIG. 8 is a schematic structural diagram illustrating a lock arbitration server, according to still another implementation of the present application. As shown in FIG. 8, on the basis of FIG. 5, the lock arbitration server further includes a second processing module 570, a third processing module 580, and a fourth processing module 590.

**[0136]** Specifically, after one of the secondary databases is selected as a new primary database, the first determining module 510 continues to determine whether a lock held by the current primary database expires. When the lock held by the current primary database has expired, the second determining module 520 determines whether lock requests of the original primary database, the current primary database, and other secondary databases are received.

**[0137]** When the lock request of the original primary database is received, the second processing module 570 restores the original primary database to a primary database, and restores the current primary database to a secondary database.

**[0138]** Specifically, after the original primary database is restored, the original primary database sends the lock request. Because the priority of the original primary database is higher, after the lock request is received, the second processing module 570 restores the original primary database to a primary database, and restores the current primary database to a secondary database.

**[0139]** If the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, the third processing module 580 maintains the current primary database as a primary database.

**[0140]** Specifically, when the lock request of the original primary database is not received and the current primary database normally updates the lock, the third processing module 580 maintains the current primary database as a primary database.

**[0141]** If the lock requests of the original primary database and the current primary database are not received and the lock requests of the other secondary databases are received, the fourth processing module 590 selects one of the other secondary databases as a primary database, and restores the current primary database to a secondary database.

**[0142]** Specifically, if the lock requests of the original primary database and the current primary database are not received, it indicates that the original primary database is not restored, and the current primary database is also faulty or needs to

perform an upgrade operation. Therefore, to normally serve the user, the fourth processing module 590 needs to select, as a primary database, one of the other secondary databases that send the lock requests, and restores the current primary database to a secondary database.

5   **[0143]**   It is worthwhile to note that, for details that are not disclosed in the implementation of the lock arbitration server in the present application, references can be made to the previous implementation of the method for managing primary and secondary databases described with reference to FIG. 1 to FIG. 4. Details are omitted here for simplicity.

10   **[0144]**   In conclusion, after selecting one of the secondary databases as a new primary database, the lock arbitration server in this implementation of the present application continues to determine whether the current primary database receives the lock request of the original primary database and whether the current primary database updates the lock. When the lock request of the original primary database is received, the server restores the original primary database to a primary database, and restores the current primary database to a secondary database. If the lock request of the original primary database is not received, the server determines whether the current primary database updates the lock within a validity period. If no, the server selects one of the other secondary databases as a primary database, and restores the current primary database to a secondary database. According to the lock arbitration server, it is ensured that the original primary database is restored to a primary database after the original primary database is restored, to better serve the user, thereby improving practicality of the method for managing primary and secondary databases in the present application.

25   **[0145]**   To implement the previous implementation, the present application further provides a primary database server.

30   **[0146]**   FIG. 9 is a schematic structural diagram illustrating a primary database server, according to an implementation of the present application. As shown in FIG. 9, the primary database server includes: a second sending module 910, configured to send a lock renewal request to a lock arbitration server when a lock held by a primary database expires, so that the lock arbitration server selects one of secondary databases as a new primary database when the lock arbitration server does not receives the lock renewal request; and a switching module 920, configured to control the primary database to be switched to a secondary database.

[0147] Specifically, in actual applications, to keep the primary database unchanged, that is, to prevent the primary database from being switched to a secondary database when the primary database is not faulty or performs an upgrade or maintenance operation, a priority of the primary database is higher than a priority of the secondary database, thereby ensuring that a database that holds the lock is the primary database when the primary database is not faulty or performs an upgrade or maintenance operation.

[0148] It is worthwhile to note that, there are a plurality of methods for deploying primary and secondary databases, for example, a method for deploying one primary database and one secondary database, or a method for deploying one primary database and a plurality of secondary databases.

[0149] For ease of description, this implementation of the present application is described by using an example that the method for deploying primary and secondary databases is a three-data-centers-in-two-locations method. That is, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in the same data center, and the primary database and the disaster recovery database are located in different data centers.

[0150] Specifically, when the lock held by the primary database expires, the second sending module 910 sends the lock renewal request to the lock arbitration server. If the lock arbitration server does not receive the lock renewal request, it indicates that the primary database is performing an upgrade operation or is faulty, and therefore cannot normally serve a user. As such, the lock arbitration server selects one of the secondary databases as a new primary database, to normally serve the user as soon as possible. In addition, the switching module 920 controls the primary database to be switched to a secondary database.

[0151] It should be understood that, to enable the primary database to update the lock in a timely manner and to enable other secondary databases to quickly obtain the lock when the primary database is faulty, the primary database can send, at a relatively short cycle within an update cycle of the lock, a request for querying a status of the lock, to update the lock in a timely manner when the lock is close to expiring and keep the primary database unchanged.

[0152] In addition, the secondary database actively sends, at a relatively long

cycle, a lock request for querying the status of the lock. Therefore, when the lock of the primary database is invalid, the secondary database can quickly obtain the lock and actively switch to a new primary database, so that services provided for the user are not affected.

- 5   **[0153]**     For example, the update cycle of the lock is T1, the primary database sends, at a cycle T2, a lock request for querying the status of the lock, and the secondary database sends, at a cycle T3, a lock request for querying the status of the lock, where T2 is less than T1, and T3 is greater than or equal to T1.

- 10   **[0154]**     In conclusion, when the lock held by the primary database expires, the primary database server in this implementation of the present application sends the lock renewal request to the lock arbitration server, so that the lock arbitration server selects one of the secondary databases as a new primary database when the lock arbitration server does not receive the lock renewal request; and switches the primary database to a secondary database. When a validity period of the lock held by the
- 15   primary database cannot be extended, the primary database server switches the primary database to a secondary database, and a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases, ensuring that the database normally serves a user, and improving user experience.

- 20   **[0155]**     In actual applications, after related operations such as maintenance and upgrade are performed on the faulty original primary database, the original primary database can continue to serve the user. Therefore, the primary database server in this implementation of the present application is further configured to receive a lock request from the original primary database after the original primary database is
- 25   restored, to switch the original primary database to a primary database.

**[0156]**     FIG. 10 is a schematic structural diagram illustrating a primary database server, according to a specific implementation of the present application. As shown in FIG. 10, on the basis of FIG. 9, the primary database server further includes a third sending module 930.

- 30   **[0157]**     Specifically, after the original primary database is restored, the third sending module 930 sends a lock request. Because the priority of the original primary database is higher, after the lock request is sent, the lock arbitration server receives the lock request sent by the original primary database, and controls the lock to be held by the original primary database after the lock expires, so that the original primary

database is restored to a primary database, and the current primary database is restored to a secondary database.

[0158] It should be emphasized that, for details that are not disclosed in the implementation of the primary database server in the present application, references  
5 can be made to the method for managing primary and secondary databases described with reference to FIG. 1 to FIG. 4. Details are omitted here for simplicity.

[0159] In conclusion, according to the primary database server in this implementation of the present application, it is ensured that the original primary database is restored to a primary database after the original primary database is  
10 restored, to better serve the user, thereby improving user experience.

[0160] To implement the previous implementation, the present application further provides a secondary database server. FIG. 11 is a schematic structural diagram illustrating a secondary database server, according to an implementation of the present application. As shown in FIG. 11, the secondary database server includes: a fourth  
15 sending module 1010, configured to send a lock request to a lock arbitration server; a second receiving module 1020, configured to receive a lock acknowledgement message sent by the lock arbitration server; and a fifth processing module 1030, configured to switch a secondary database to a new primary database based on the lock acknowledgement message.

[0161] It can be understood that, in actual applications, to keep a primary database unchanged, that is, to prevent the primary database from being switched to a secondary database when the primary database is not faulty or performs an upgrade or maintenance operation, a priority of the primary database is higher than a priority of the secondary database, thereby ensuring that a database that holds a lock is the  
20 primary database when the primary database is not faulty or performs an upgrade or maintenance operation. However, when the primary database cannot effectively extend a validity period of the lock, the lock arbitration server needs to select one of secondary databases as a primary database to continue to serve a user.

[0162] Specifically, the fourth sending module 1010 sends the lock request to the  
30 lock arbitration server, so that when the primary database cannot effectively update the lock, the secondary database quickly seizes the lock. The second receiving module 1020 receives the lock acknowledgement message sent by the lock arbitration server, so that the fifth processing module 1030 switches a secondary database to a primary database based on the lock acknowledgement message.

**[0163]** It is worthwhile to note that, there are a plurality of methods for deploying primary and secondary databases, for example, a method for deploying one primary database and one secondary database, or a method for deploying one primary database and a plurality of secondary databases.

5 **[0164]** For ease of description, this implementation of the present application is described by using an example that the method for deploying primary and secondary databases is a three-data-centers-in-two-locations method. That is, the secondary database includes a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in the same data center, and the  
10 primary database and the disaster recovery database are located in different data centers.

**[0165]** In conclusion, the secondary database server in this implementation of the present application sends the lock request to the lock arbitration server, to switch a secondary database to a new primary database when it is determined that the lock held  
15 by the primary database is invalid and the primary database cannot normally serve the user, thereby improving a speed and accuracy of switching between primary and secondary databases, ensuring that services can be normally provided for the user, and improving user experience.

**[0166]** Based on the previous implementation, if the secondary database server  
20 does not always actively send, at a specific cycle, a lock request for querying a status of the lock, when the lock held by the primary database is invalid, a lock expiration notification is actively sent to each of other secondary databases, so that the secondary database server sends the lock request based on the lock expiration notification, obtains the lock, and becomes a new primary database.

25 **[0167]** Specifically, FIG. 12 is a schematic structural diagram illustrating a secondary database server, according to a specific implementation of the present application. As shown in FIG. 12, on the basis of FIG. 11, the secondary database server further includes a third sending module 1040.

**[0168]** Specifically, the third receiving module 1040 receives a lock expiration  
30 notification sent by the lock arbitration server after the lock arbitration server determines that a lock held by a primary database has expired, where the fourth sending module 1010 sends the lock request to the lock arbitration server based on the lock expiration notification.

[0169] Further, when the lock expires, if the lock arbitration server does not receive a lock renewal request of the primary database, it indicates that the current database is faulty or is performing an upgrade or maintenance operation. Therefore, to normally serve the user as soon as possible, the lock arbitration server selects, as a  
5 primary database, a secondary database with the earliest receiving time.

[0170] In conclusion, based on the secondary database server in this implementation of the present application, when the lock of the primary database expires and is invalid, the secondary database quickly obtains the lock and becomes a new primary database, thereby ensuring that normal services are still provided for the  
10 user in a relatively short time, and improving user experience.

[0171] In actual applications, after related operations such as maintenance and upgrade are performed on the faulty original primary database, the original primary database can continue to serve the user. Therefore, the secondary database server in this implementation of the present application is further configured to: after the  
15 original primary database is restored, switch the current primary database to a secondary database.

[0172] FIG. 13 is a schematic structural diagram illustrating a secondary database server, according to another specific implementation of the present application. As shown in FIG. 13, on the basis of FIG. 11, the secondary database server further  
20 includes a sixth processing module 1050.

[0173] Specifically, after the original primary database is restored, the original primary database sends a lock request. Because the priority of the original primary database is higher, after receiving the lock request, the lock arbitration server restores the original primary database to a primary database, and the sixth processing module  
25 1050 restores the current primary database to a secondary database.

[0174] It should be understood that, for details that are not disclosed in the implementation of the secondary database server in the present application, references can be made to the method for managing primary and secondary databases described with reference to FIG. 1 to FIG. 4. Details are omitted here for simplicity.

[0175] In conclusion, after the original primary database is restored, the secondary database server in this implementation of the present application restores the current primary database to a secondary database, and the arbitration server restores the original primary database to a primary database, to better serve the user, thereby  
30 improving user experience.

**[0176]** To implement the previous implementation, the present application further provides a system for managing primary and secondary databases. FIG. 14 is a schematic structural diagram illustrating a system for managing primary and secondary databases, according to an implementation of the present application. As shown in FIG. 14, the system for managing primary and secondary databases includes a lock arbitration server 1000, a primary database server 2000, and a secondary database server 3000.

**[0177]** It is worthwhile to note that, for descriptions of the lock arbitration server 1000, the primary database server 2000, and the secondary database server 3000 in the present application, references can be made to the previous descriptions of the lock arbitration server, the primary database server, and the secondary database server. Details are omitted here for simplicity.

**[0178]** In conclusion, according to the system for managing primary and secondary databases in this implementation of the present application, it is determined whether the lock held by the primary database expires. If it is determined that the lock held by the primary database has expired, it is determined whether the lock renewal request of the primary database is received. If the lock renewal request is not received, one of the secondary databases is selected as a new primary database, and the primary database is controlled to be switched to a secondary database. According to the system, if the primary database does not send the lock renewal request before the lock held by the primary database expires, it is determined that the lock held by the primary database is invalid and the primary database cannot normally serve a user, so that a secondary database is selected as a new primary database, thereby improving a speed and accuracy of switching between primary and secondary databases.

**[0179]** In addition, the terms "first" and "second" are merely used for description, and shall not be understood as an indication or implication of relative importance or implicit indication of the number of indicated technical features. Therefore, a feature limited by "first" or "second" can explicitly or implicitly include at least one such feature. In the descriptions of the present application, "a plurality of" means "at least two", for example, two or three, unless otherwise specifically limited.

**[0180]** In the descriptions of the present specification, reference terms such as "an implementation", "some implementations", "example", "specific example", and "some examples" mean that specific features, structures, materials, or characteristics described with reference to the implementations or examples are included in at least

one implementation or example in the present application. In the present specification, the previous examples of the terms are not necessarily with respect to the same implementation or example. In addition, the described specific features, structures, materials, or characteristics can be combined in a proper way in any one or more of the implementations or examples. In addition, a person skilled in the art can integrate or combine different implementations or examples and characteristics of different implementations or examples described in the present specification, provided that they do not conflict with each other.

**[0181]** Although the implementations of the present application are shown and described above, it can be understood that the previous implementations are examples and should not be construed as a limitation on the present application. Within the scope of the present application, a person of ordinary skill in the art can make changes, modifications, replacements, and variations to the previous implementations.

## CLAIMS

What is claimed is:

1. A method for managing primary database and secondary databases, comprising:  
determining that a lock held by the primary database has expired, wherein the  
5 primary database and the secondary databases share the lock;

in response to determining that the lock held by the primary database has expired,  
sending a lock expiration notification to the secondary databases, and determining that  
a lock renewal request of the primary database has not been received after the lock has  
expired;

10 receiving a lock request sent by the secondary databases based on the lock  
expiration notification, and recording a receiving time of the lock request; and

in response to determining that the lock renewal request of the primary database  
has not been received, selecting a particular secondary database as a new primary  
database, wherein the particular secondary database has earliest receiving time.

15

2. The method for managing the primary database and the secondary databases  
according to claim 1, wherein a priority of the primary database is higher than a priority  
of the particular secondary database.

20 3. The method for managing the primary database and the secondary databases  
according to claim 1, wherein the particular secondary database comprises a hot standby  
database and a disaster recovery database, the primary database and the hot standby  
database are located in a same data center, and the primary database and the disaster  
recovery database are located in different data centers.

25

4. The method for managing the primary database and the secondary databases  
according to claim 1, further comprising:

performing operations comprising at least one of a maintenance or an upgrade  
on the primary database.

30

5. The method for managing the primary database and the secondary databases  
according to claim 4, further comprising:

determining whether the operations comprising at least one of the maintenance and the upgrade on the primary database are completed; and

in response to determining that the operations comprising at least one of the maintenance and the upgrade on the primary database are completed, restoring the  
5 primary database.

6. The method for managing the primary database and the secondary databases according to claim 5, wherein after the primary database is restored, the method further comprises:

10 receiving a lock request sent by an original primary database, and controlling the lock to be held by the original primary database after the lock expires, wherein the original primary database is restored to the primary database, and a current primary database is restored to one of the secondary databases.

15 7. The method for managing the primary database and the secondary databases according to claim 6, further comprising:

continuing to determine whether the lock held by the current primary database expires;

in response to determining that the lock held by the current primary database has  
20 expired, determining whether lock requests of the original primary database, the current primary database, and other secondary databases are received; and

in response to determining that the lock request of the original primary database is received, restoring the original primary database to the primary database, and restoring the current primary database to one of the secondary databases; or

25 in response to determining that the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, maintaining the current primary database as the primary database; or

in response to determining that the lock requests of the original primary database  
30 and the current primary database are not received and the lock requests of the other secondary databases are received, selecting one of the other secondary databases as the primary database, and restoring the current primary database to one of the secondary databases.

8. A non-transitory, computer-readable medium storing one or more instructions executable by a computer system to perform operations for managing primary database and secondary databases, comprising:

determining that a lock held by the primary database has expired, wherein the primary database and the secondary databases share the lock;

in response to determining that the lock held by the primary database has expired, sending a lock expiration notification to the secondary databases, and determining that a lock renewal request of the primary database has not been received after the lock has expired;

receiving a lock request sent by the secondary databases based on the lock expiration notification, and recording a receiving time of the lock request; and

in response to determining that the lock renewal request of the primary database has not been received, selecting a particular secondary database as a new primary database, wherein the particular secondary database has earliest receiving time.

9. The non-transitory, computer-readable medium of claim 8, wherein a priority of the primary database is higher than a priority of the particular secondary database.

10. The non-transitory, computer-readable medium of claim 8, wherein the particular secondary database comprises a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary database and the disaster recovery database are located in different data centers.

11. The non-transitory, computer-readable medium of claim 8, the operations further comprise:

performing at least one of a maintenance or an upgrade on the primary database.

12. The non-transitory, computer-readable medium of claim 11, the operations further comprise:

determining whether the operations comprising at least one of the maintenance and the upgrade on the primary database are completed; and

in response to determining that the operations comprising at least one of the maintenance and the upgrade on the primary database are completed, restoring the

primary database.

13. The non-transitory, computer-readable medium of claim 12, wherein after the primary database is restored, the operations further comprise:

5       receiving a lock request sent by an original primary database, and controlling the lock to be held by the original primary database after the lock expires, wherein the original primary database is restored to the primary database, and a current primary database is restored to one of the secondary databases.

10       14. The non-transitory, computer-readable medium of claim 13, the operations further comprise:

continuing to determine whether the lock held by the current primary database expires;

15       in response to determining that the lock held by the current primary database has expired, determining whether lock requests of the original primary database, the current primary database, and other secondary databases are received; and

in response to determining that the lock request of the original primary database is received, restoring the original primary database to the primary database, and restoring the current primary database to one of the secondary databases; or

20       in response to determining that the lock request of the original primary database is not received and the lock requests of the current primary database and the other secondary databases are received, maintaining the current primary database as the primary database; or

25       in response to determining that the lock requests of the original primary database and the current primary database are not received and the lock requests of the other secondary databases are received, selecting one of the other secondary databases as the primary database, and restoring the current primary database to one of the secondary databases.

15. A computer-implemented system, comprising:

one or more computers; and

one or more computer memory devices interoperably coupled with the one or more computers and having tangible, non-transitory, machine-readable media storing one or more instructions that, when executed by the one or more computers, perform operations for managing primary database and secondary databases, comprising:

determining that a lock held by the primary database has expired, wherein the primary database and the secondary databases share the lock;

in response to determining that the lock held by the primary database has expired, sending a lock expiration notification to the secondary databases, and determining that a lock renewal request of the primary database has not been received after the lock has expired;

receiving a lock request sent by the secondary databases based on the lock expiration notification, and recording a receiving time of the lock request; and

in response to determining that the lock renewal request of the primary database has not been received, selecting a particular secondary database as a new primary database, wherein the particular secondary database has earliest receiving time.

16. The computer-implemented system of claim 15, wherein a priority of the primary database is higher than a priority of the particular secondary database.

17. The computer-implemented system of claim 15, wherein the particular secondary database comprises a hot standby database and a disaster recovery database, the primary database and the hot standby database are located in a same data center, and the primary database and the disaster recovery database are located in different data centers.

18. The computer-implemented system of claim 15, the operations further comprise:

performing at least one of a maintenance or an upgrade on the primary database.

19. The computer-implemented system of claim 18, the operations further comprise:

determining whether the operations comprising at least one of the maintenance and the upgrade on the primary database are completed; and

in response to determining that the operations comprising at least one of the maintenance and the upgrade on the primary database are completed, restoring the  
5 primary database.

20. The computer-implemented system of claim 19, wherein after the primary database is restored, the operations further comprise:

receiving a lock request sent by an original primary database, and controlling the  
10 lock to be held by the original primary database after the lock expires, wherein the original primary database is restored to the primary database, and a current primary database is restored to one of the secondary databases.