CLAIMS

1. A computer-implemented method, comprising:

receiving, from a user and at a primary database, a data modification request for modifying data in the primary database, wherein the data modification request specifies one or more updated values of one or more respective first database records in first columns;

modifying, in response to the data modification request, one or more previous values of the one or more respective first database records in the first columns to have the one or more updated values;

generating, based on the data modification request, a first log and a second log, wherein the second log comprises a) the one or more updated values and the one or more previous values of the one or more respective first database records in the first columns that are being modified and b) previous values of second database records in second columns that are not being modified, and wherein the first log comprises only the one or more updated values of the one or more respective first database records in the first columns for use in database recovery or synchronization;

performing persistence processing including storing first columns being modified in the first log in persistent storage before returning a modification result to the user and storing second columns not being modified in the second log only after returning the modification result to the user; and

returning the modification result to the user that is indicative of having successfully modified the data in the primary database in response to receiving the data modification request.

2. The computer-implemented method of claim 1, wherein the data modification request comprises at least one of a data update request, a data deletion request, and a data addition request.

3. The computer-implemented method of claim 1, wherein performing the persistence processing on the first log comprises:

recording the first log as a file in a certain format, and

storing the file in a nonvolatile storage device.

4. The computer-implemented method of claim 1, further comprising, prior to returning the modification result:

temporarily withholding the modification result in response to performing the persistence processing on the first log; and

synchronizing the first log to a secondary database so that the secondary database synchronizes the modification based on the first log.

5. The computer-implemented method of claim 4, wherein the secondary database synchronizing the modification based on the first log comprises:

receiving, by the secondary database, the first log synchronized by the primary database;

performing, by the secondary database, the persistence processing on the first log; and

performing a replay on the first log.

6. A non-transitory, computer-readable medium storing one or more instructions executable by a computer system to perform operations comprising:

receiving, from a user and at a primary database, a data modification request for modifying data in the primary database, wherein the data modification request specifies one or more updated values of one or more respective first database records in first columns;

modifying, in response to the data modification request, one or more previous values of the one or more respective first database records in the first columns to have the one or more updated values;

generating, based on the data modification request, a first log and a second log, wherein the second log comprises a) the one or more updated values and the one or more previous values of the one or more respective first database records in the first columns that are being modified and b) previous values of second database records in second columns that are not being modified, and wherein the first log comprises only the one or more updated values of the one or more respective first database records in the first columns for use in database recovery or synchronization;

performing persistence processing including storing first columns being modified in the first log in persistent storage before returning a modification result to the user and storing second

columns not being modified in the second log only after returning the modification result to the user; and

returning the modification result to the user that is indicative of having successfully modified the data in the primary database in response to receiving the data modification request.

7. The non-transitory, computer-readable medium of claim 6, wherein the data modification request comprises at least one of a data update request, a data deletion request, and a data addition request.

8. The non-transitory, computer-readable medium of claim 6, wherein performing the persistence processing on the first log comprises:

recording the first log as a file in a certain format, and

storing the file in a nonvolatile storage device.

9. The non-transitory, computer-readable medium of claim 6, wherein the operations further comprise, prior to returning the modification result:

temporarily withholding the modification result in response to performing the persistence processing on the first log; and

synchronizing the first log to a secondary database so that the secondary database synchronizes the modification based on the first log.

10. The non-transitory, computer-readable medium of claim 9, wherein the secondary database synchronizing the modification based on the first log comprises:

receiving, by the secondary database, the first log synchronized by the primary database;

performing, by the secondary database, the persistence processing on the first log; and

performing a replay on the first log.

11. 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 comprising:

receiving, from a user and at a primary database, a data modification request for modifying data in the primary database, wherein the data modification request specifies one or more updated values of one or more respective first database records in first columns;

modifying, in response to the data modification request, one or more previous values of the one or more respective first database records in the first columns to have the one or more updated values;

generating, based on the data modification request, a first log and a second log, wherein the second log comprises a) the one or more updated values and the one or more previous values of the one or more respective first database records in the first columns that are being modified and b) previous values of second database records in second columns that are not being modified, and wherein the first log comprises only the one or more updated values of the one or more respective first database records in the first columns for use in database recovery or synchronization;

performing persistence processing including storing first columns being modified in the first log in persistent storage before returning a modification result to the user and storing second columns not being modified in the second log only after returning the modification result to the user; and

returning the modification result to the user that is indicative of having successfully modified the data in the primary database in response to receiving the data modification request.

12. The computer-implemented system of claim 11, wherein the data modification request comprises at least one of a data update request, a data deletion request, and a data addition request.

13. The computer-implemented system of claim 11, wherein performing the persistence processing on the first log comprises:

recording the first log as a file in a certain format, and

storing the file in a nonvolatile storage device.

14. The computer-implemented system of claim 11, wherein the operations further comprise, prior to returning the modification result:

temporarily withholding the modification result in response to performing the persistence processing on the first log; and

synchronizing the first log to a secondary database so that the secondary database synchronizes the modification based on the first log, wherein the secondary database synchronizing the modification based on the first log comprises:

receiving, by the secondary database, the first log synchronized by the primary database;

performing, by the secondary database, the persistence processing on the first log; and

performing a replay on the first log.