

What is claimed and desired to be secured by patent is:

1. A method for handling clashes during the synchronization of operations performed on a disconnected first computer with operations performed on a second computer, the synchronization being performed after the first computer and the second computer are reconnected by a network link, the second computer and the first computer each containing a replica of a distributed database, said method comprising the computer-implemented steps of:

merging out a representation of the operations performed on the first computer by applying at least a portion of the operations to the second replica;

merging in a representation of the operations performed on the second computer by applying at least a portion of the operations to the first computer replica;

detecting persistent clashes during at least one of said merging steps; and

recovering from at least a portion of the detected persistent clashes.

2. The method of claim 1, wherein the representation of the operations performed on each computer includes an update log.

3. The method of claim 2, wherein said method further comprises the steps of detecting transient clashes and eliminating at least a portion of the detected transient clashes.

4. The method of claim 3, wherein said method further comprises the step of compressing at least one of the update logs and said steps of detecting and eliminating transient clashes both occur during said compressing step.

5. The method of claim 3, wherein said step of detecting transient clashes comprises identifying a clashing update in an update log, and said eliminating step comprises scanning forward in the log from the clashing update to determine whether a subsequent repairing update will remove a clash condition associated with the clashing update.

6. The method of claim 5, wherein temporally consistent sections of the update log are merged atomically.

7. The method of claim 3, wherein said eliminating step comprises providing a shadow database for at least a portion of one of the replicas, and one of said merging steps applies operations atomically to the shadow database and subsequently alters the replica to correspond with the shadow database.

8. The method of claim 1, wherein one of said merging steps is preceded by the step of locking the replica used during merging to prevent access to the replica during merging by any process other than a process that is performing the merging step.

9. The method of claim 1, wherein said steps of detecting persistent clashes and recovering from persistent clashes are performed primarily on the second computer.

10. The method of claim 1, wherein said merging out step is preceded by the step of compressing an update log on the first computer to create the representation of operations performed on the first computer.

11. The method of claim 1, wherein said step of detecting persistent clashes comprises detecting unique key clashes.

12. The method of claim 1, wherein said step of detecting persistent clashes comprises detecting incompatible manipulation clashes.

13. The method of claim 1, wherein said step of detecting persistent clashes comprises detecting file content clashes.

14. The method of claim 1, wherein said step of detecting persistent clashes comprises detecting permission clashes.

15. The method of claim 1, wherein said step of detecting persistent clashes comprises detecting a clash between at least a portion of the distributed database and a structure external to the distributed database.

16. The method of claim 1, wherein said step of recovering from persistent clashes comprises comparing a value in the first computer replica with a corresponding value in the second replica and modifying the value stored in one replica to make it equal to the corresponding value in the other replica.

17. The method of claim 1, wherein the representation of the operations performed on the first computer includes an update log, and said step of recovering from persistent clashes comprises regressing persistently clashing updates from the first computer replica.

18. The method of claim 17, wherein said detecting, recovering, and regressing steps are applied recursively after an update regression to handle clashes caused by the update regression.

19. The method of claim 1, wherein said step of recovering from persistent clashes comprises inserting an update into the update stream before a clashing update.

20. The method of claim 1, wherein said step of recovering from persistent clashes comprises creating a recovery item in a recovery storage.

21. The method of claim 20, wherein the recovery storage comprises a directory hierarchy for holding recovery items in a file system.

22. The method of claim 21, wherein all recovery items created in response to clashes detected during a single instance of said detecting step reside in a single directory hierarchy whose directories were created during said recovering step.

23. The method of claim 20, wherein the recovery storage comprises a container hierarchy for holding recovery items in a directory services tree.

24. The method of claim 20, wherein said step of creating a recovery item comprises moving an object in the distributed database into the recovery storage.

25. The method of claim 20, wherein said step of creating a recovery item comprises duplicating an object in the distributed database and placing the duplicate object into the recovery storage.

26. The method of claim 20, wherein said step of creating a recovery item comprises creating in the recovery storage an object that is not present in the distributed database.

27. The method of claim 1, wherein at least one of said merging steps includes applying an operation to the replica by modifying a grouped attribute.

28. The method of claim 1, further comprising the steps of:

associating a sequence number with each operation performed on the first computer while the first computer is disconnected from the second computer; and

maintaining a record on the second computer of which sequence numbers have been applied to the second replica during said merging out step.

29. The method of claim 1, wherein the first computer is a mobile computer and the second computer is a central server computer.

30. The method of claim 1, wherein the first computer and the second computer are each server computers in a computer network.

31. A computer-readable storage medium having a configuration that represents data and instructions which cause a first computer and a second computer connected by a network link to perform method steps for handling clashes, the first computer and the second computer each containing a replica of a distributed database, the method comprising the steps of:

merging out a representation of operations performed on the first computer by applying at least a portion of the operations to the second replica;