

## Replication Mechanism for a Distributed Version Control System

Disclosed is a mechanism for replication of content repositories of a distributed version control system (DVCS), which is unique in that objects and branch references are transferred between nodes by independent means, with only branch references subject to synchronization among nodes. The object of this disclosure is to simplify deployment of a DVCS in a replicated configuration. Replication can be performed efficiently by transferring objects using a key-value store, such as a distributed hash table (DHT), independently of a repository's branch references, which are transferred using a synchronized data store.

The method begins with a push to a DVCS repository located on one of a plurality of server nodes. For each branch updated by the push, the current branch reference in the DVCS repository is compared to the current branch reference stored in an instance of a synchronized data store.

If the current branch reference in the DVCS repository matches the current branch reference in the synchronized data store, the system attempts to update the branch reference in the synchronized data store to match the branch reference in the push. If this update is successful, the system allows the push to complete normally. Otherwise, the system rejects the push. All other built-in and user-defined restrictions on the push can also be enforced.

The system transfers incoming objects from the push into a key-value store available to the other server nodes. Each node on which a branch reference in the synchronized data store is different from the corresponding branch reference in the corresponding local DVCS repository recursively fetches the required objects from the key-value store into the DVCS repository. Upon successful completion of this recursive fetch on a node, the system updates the branch reference in the DVCS repository to match the corresponding branch reference in the synchronized data store.

Because the key-value store may not have all required objects at the time the branch reference is updated in the synchronized data store, each node periodically checks for branch references in its DVCS repository or repositories that do not match the branch references in its instance of the synchronized data store, and attempts a recursive fetch of the required objects, updating the branch reference in the DVCS repository if successful.

Three operations: transfer of incoming objects into the key-value store from a DVCS repository, fetch of objects from the key-value store into a DVCS repository, and update of a branch reference in a DVCS repository, are not necessarily synchronized among nodes.

It is possible for a push to a node to occur when either (1) the branch reference in the DVCS repository receiving the push does not match the corresponding branch reference in that node's instance of the synchronized data store, or (2) the synchronized data store is unable to establish consensus on the new branch reference; as, for example, when a update is already in progress or there is a failure in an underlying communications medium. In either case (1) or (2), the system refuses the push. This situation is similar to the situation in the ordinary operation of a DVCS, when a previous user's push takes effect before the current push. Upon successful completion of replication, the user will be able to push using ordinary DVCS tools and techniques.