

What is BranchCache

Contributed by David Noel-Davies

Many organizations have distributed environments, with users at remote locations connected via fairly slow, high latency links. It's often a challenge for users at these locations to access data at the hub locations, with end users experiencing slow access. BranchCache aims to improve this experience by allowing data downloaded by one client at a remote location to be accessible to all clients in that location in one of two ways.

- Hosted Cache—With hosted cache, one server at each location has the BranchCache feature installed and enabled. This server will cache data on behalf of the entire remote site. All clients query this server to see if the server has a copy of the required data. When the data isn't available on the server, the client will download directly from the source server and then offer the data to the BranchCache local server to see if it wants to cache it.

- Distributed Cache—With distributed cache, there's no dedicated server to store the data. Each machine in the location has its own local cache for data it downloads. When a client wants to download data from the source server, it sends out a broadcast on the local subnet to see if any machine has the data cached. If any local machine has the data, it will send it to the requesting machine.

In both modes, the source server is still contacted to obtain hash values for the data that's requested. It is through the hash values that clients can request data from its peers or the caching server. Because the hash identifies the data based on content, the use of the hash stops clients from obtaining outdated cached data from the BranchCache. If the data changes on the source, a new hash is generated.

Note that only data that has a hash value can be cached using BranchCache. Currently, only Windows 7 and Windows Server 2008 R2 support BranchCache.