

An local copy.4. when
the process in the



**ASSIGN
BUSTER**

An Andrew file system (AFS) is a location-independent file system that uses a local cache to reduce the workload and increase the performance of a distributed computing environment. A first request for data to a server from a workstation is satisfied by the server and placed in a local cache. A second request for the same data is satisfied from the local cache.

The Andrew file system was developed at Carnegie-Mellon University.

Question one - B: 1. When a user process in a client computer issues an open system call for a file in the shared file space and there is not a current copy of the file in the local cache, the server holding the file is located and sent a request for a copy of the file. 2. The copy is stored in the local UNIX file system in the client computer.

The copy is then opened, and the resulting UNIX file descriptor is returned to the client. 3. Subsequent read, write and other operations on the file by processes in the client computer is applied to the local copy. 4. When the process in the client issues a close system call, if the local copy has been updated its contents are sent back to the server. The server updates the file contents and the timestamps on the file. The copy on the client's local disk is retained in case it is needed again by a user-level process on the same workstation.

We discuss the observed performance of AFS below, but we can make some general observations and predictions here based on the design characteristics described above:

- For shared files that are infrequently updated (such as those containing the code of UNIX commands and libraries) and for files that are normally accessed by only a single user (such as most of

the files in a user's home directory and its subtree), locally cached copies are likely to remain valid for long periods - in the first case because they are not updated and in the second because if they are updated, the updated copy will be in the cache on the owner's workstation. These classes of file account for the overwhelming majority of file accesses. • The local cache can be allocated a substantial proportion of the disk space on each workstation - say, 100 megabytes. This is normally sufficient for the establishment of a working set of the files used by one user. The provision of sufficient cache storage for the establishment of a working set ensures that files in regular use on a given workstation are normally retained in the cache until they are needed again. - Sequential access is common, and random access is rare.

- Most files are read and written by only one user. When a file is shared, it is usually only one user who modifies it. - Files are referenced in bursts. If a file has been referenced recently, there is a high probability that it will be referenced again soon. These observations were used to guide the design and optimization of AFS, not to restrict the functionality seen by users. • AFS works best with the classes of file identified in the first point above.

There is one important type of file that does not fit into any of these classes - databases are typically shared by many users and are often updated quite frequently. Question two: Routing overlay is: A distributed algorithm known as routing overlay, it locates nodes and objects, it is middleware layer responsible for routing requests from clients to hosts that holds the object to which request is addressed. The Main difference is that routing is implemented in application layer (besides the IP routing at network layer) It is termed an overlay since it implements in the client a routing algorithm that is <https://assignbuster.com/an-local-copy4-when-the-process-in-the/>

quite separate from the routing of individual IP packets, the routing overlay ensures that any node can access any object through a sequence of nodes, by exploiting the knowledge at each of them to locate the destination object. An overlay network decouples network services from the underlying infrastructure by encapsulating one packet inside of another packet. After the encapsulated packet has been forwarded to the endpoint, it is de-encapsulated. Overlay routing is different than IP routing however is strongly influenced by IP routing, Routing tables may be updated synchronously or asynchronously. The authors of Pastry have developed the Squirrel peer-to-peer web caching service for use in local networks of personal computers