

Performance Considerations

This section provides general advice on performance-related issues regarding the buffer pool and the BP cache.

For explanations of the statistics items mentioned in this section, refer to *Buffer Pool Load/Locate Statistics*.

Related Topic:

- *Performance Hints*

This section covers the following topics:

- Internal Fast Locate Table
 - Searching in Steplibs
 - Reusing and Retaining Objects
 - Local versus Global Buffer Pool
-

Internal Fast Locate Table

When a Natural object is referenced for the first time within a Natural session, the buffer pool manager creates a directory entry for this object. A directory entry is used to identify an object and contains information such as the name of the object, the library where it resides (name, database ID and file number) and the address of the object (position) in the buffer pool.

The Natural runtime system remembers names of objects that were recently executed, the libraries (name, database ID and file number) where they reside and the addresses of the corresponding buffer pool directory entries in the internal fast locate table for the duration of a Natural session.

When a user invokes an object that has been used before in the Natural session, the Natural runtime system passes the information from the internal fast locate table to the buffer pool manager, which can then bypass the Normal Locate procedure and perform a time-saving Quick Locate call (see **Quick Locate Calls**). This is the most efficient way to locate an object.

If the position of an object in the buffer pool has changed, the buffer pool manager will automatically schedule a Normal Locate call. The position usually changes, if an object that was deleted from the buffer pool or overwritten by another object is reloaded into the buffer pool.

The internal fast locate table contains a maximum of 128 entries. The fast locate table is reset with the LOGON system command and must be filled again using Normal Locate calls. Therefore, an application that performs a **LOGON** loses performance.

A high ratio of **Normal Locate Calls** to **Quick Locate Calls** indicates the use of LOGON commands in a Natural application in that each LOGON causes the internal fast locate table to be reset.

Searching in Steplibs

The search for a Natural object through a long chain of steplib libraries has a negative impact on performance.

For each steplib, the Natural runtime issues a call to the buffer pool manager until the requested object is found. However, each needless call for a steplib library that does not contain the requested object can usually be avoided.

Depending on the setting of the BPSFI profile parameter (described in the *Parameter Reference* documentation), additional database calls may be required.

A long steplib chain is indicated by a high ratio of **Normal Locate Calls** (including steplib searches) to **Normal Locate Calls** (without steplib searches) that is calculated as follows:

Normal Locate Calls: (Normal Locate Calls - STEPLIB Searches)

Example of a Steplib Search

When searching the default steplib chain (library SYSTEM of FUSER, library SYSTEM of FNAT), each attempt to load an object from SYSTEM (FNAT) will result in the following:

3 Normal Locate Calls and 2 STEPLIB Searches

Explanation:

3 Normal Locate calls result from searches in the current library, the library SYSTEM (FUSER) and the library SYSTEM (FNAT). There are (at least) 2 Normal Locate calls that fail, since the object is stored in neither the current library nor the library SYSTEM (FUSER). Using the above formula, results in a ratio of 3:1.

If the object is located in the current library, the result is as follows:

1 Normal Locate call and 0 (zero) STEPLIB Searches. Using the above formula, results in a ratio of 1:1.

Reusing and Retaining Objects

An application that contains many Natural objects where each object is rarely ever executed has heavy impact on the performance of the buffer pool. None of the objects resides in the buffer pool for a long time and many objects have to be loaded from a system file. For performance reasons, an application should reuse objects as often as possible, for example, by moving identical source code contained in multiple objects into a single object.

You can check the use of objects with the **List Objects** function (see the section *List Objects*). For example: the **Max** column contains information on the maximum number of applications that have executed an object, and the **TotalUC** column contains the total number of Locate calls of an object that was loaded into the buffer pool.

Objects can be made resident individually by using the **List Objects** function or by specifying them in a preload list (see *Preload List Maintenance*).

Local versus Global Buffer Pool

There are no general recommendations for when to use a local or a global buffer pool as different applications have different requirements. However, rules-of-thumb from experience enable us to give general advice on:

This section covers the following topics:

- Using Local Buffer Pools
- Using a Global Buffer Pool

Using Local Buffer Pools

Performance can improve using several small local buffer pools instead of a single global buffer pool if local buffer pools can be assigned to different application environments.

For example, in CICS environments, performance usually improves when using a local buffer pool for each AOR (Application Operating Region).

Using a Global Buffer Pool

For different batch applications that reference the same Natural objects, performance may improve if these applications use a common global buffer pool instead of a local buffer pool for each individual application. If so, there is a higher probability that the objects required by each application have already been loaded into the global buffer pool by one of the other applications.