

Migrations to and from the Archive

Natural objects (source and cataloged form as well as Xref data), as well as foreign objects can be migrated from the Control status to the Archive status, where they are stored off-line as a sequential file on a medium such as tape, cassette, or DASD.

Archiving allows you to manage your available storage space more effectively by removing redundant or historical objects from the PAC environment in a timely manner leaving only the current and active object versions online in PAC.

Note:

Objects in a production status type are current and cannot be archived because they must be synchronized with PAA.

Archived objects are still under PAC control and can subsequently be restored to a Control, test, or production status using the archive reload facility.

The maintenance of archive events is a PAC administrator function. Refer to the PAC Administration documentation for more information.

Archived objects are not retired objects. A retired object has been physically removed either from a status or from PAC itself. An archived object can subsequently be retired. Refer to the section Migrations to Retire Status for more information.

Note:

It is not yet possible to archive Predict objects.

This chapter covers the following topics:

- Archiving Process
- Phase 1
- Restoring Objects Before Phase 2
- Phase 2 : Removing the Object Versions from the PAC Environment
- Restoring Archived Objects to the PAC Environment

Archiving Process

To preserve the integrity of the objects, the archiving process unloads objects into an archive file in two phases:

Phase Action

- 1 The archive file (dataset) is created. The archive event for Phase 2 is created and placed in an unloaded state. Since the objects are still physically in PAC until Phase 2 is performed, the objects' status is not updated; their status is still what it was before Phase 1.
- 2 Finalizes the archive event: The objects on the archive (output) file are verified and then removed from the PAC environment (ACF system file).

Note:

Phase 2 can only be done by the PAC administrator.

Phase 1

Creating the Archive File

Phase 1 unloads the object version to the archive file as represented in the diagram below. The steps indicated in the diagram are described below:

Step	Action
1	Define the migration event from Control to Archive.
2	PACEX026 user exit archiving User Exit 26 archiving Select objects for the migration event object list. The list can be generated automatically or created manually. User Exit 26 can be set up in your environment to provide the criteria for automatically selecting the objects to be archived.
3	Verify the object version selection. Specified statuses are checked and objects found are checked first against the minimum current version, then against the minimum retention period defined in the application defaults, and finally, against the status.
4	Authorize and submit the event. The archiving process unloads the objects to the archive file and creates the archive event for Phase 2.
5	Unlock the migration event to complete Phase 1.

Setting Up the Job

The archive migration event can be run only in batch with a work file. The JCL (ARCHIVE_UNLOAD) is similar to the JCL for a Natural migration with a work file; however, the processing control statements are different.

Defining the Archive File

You need to define the archive file where the application objects are to be stored.

When the job for the migration is set up, Work File 1 (CMWKF01) should be assigned to the archive file or to an @ variable.

PAC provides the @ARC-DSN and @ARC-VOLSER substitution parameters for this purpose. If you use these substitution parameters, you are prompted for the archive file name and any other variables (for example, DISP) when the migration event is submitted. This information is stored in the archive event (Phase 2).

You may use other substitution parameters; however, the migration event is automatically updated only if you use the @ARC-DSN and @ARC-VOLSER parameters.

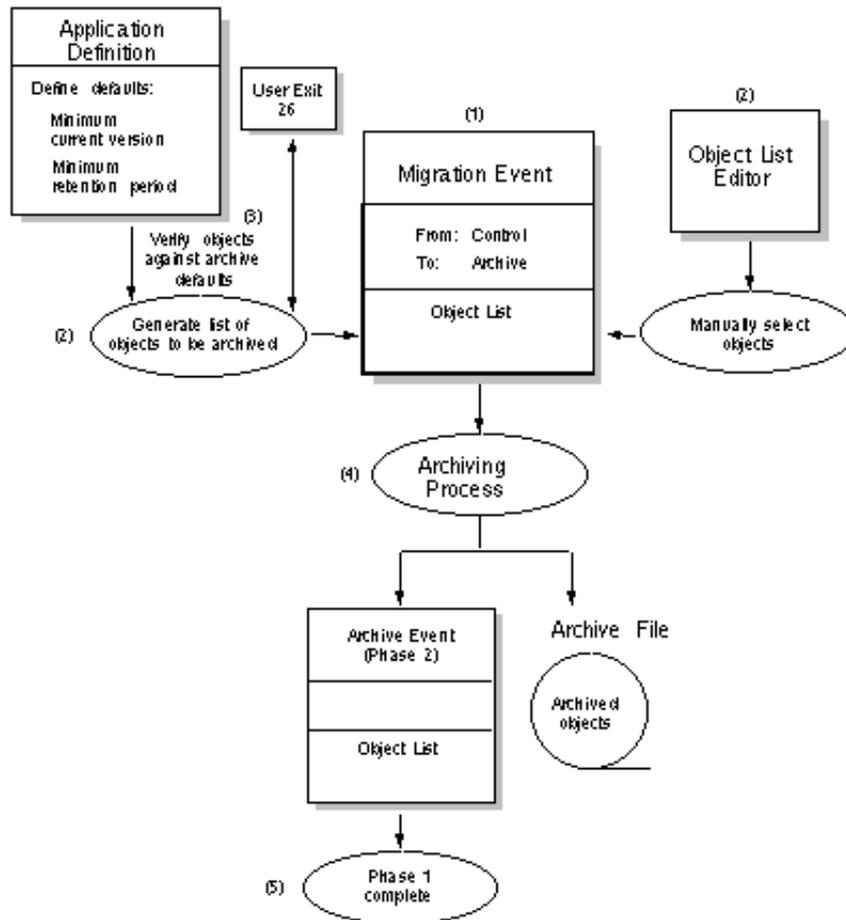
Do not use the MVS JCL or the CMS filedef option DISP=MOD for the disposition of the archive dataset; PAC has no facility for tracking datasets that have been concatenated using this parameter; thus, each archive file should be unique and independent of the others.

Note:

For VM/CMS, the substitution parameters are @ARC-VMUSERID, @ARC-VADDR1, @ARC-FN, and @ARC-FT.

Defining the Unload Migration Event

Once the status Archive has been linked to the application whose objects are to be archived and a migration path between Control and Archive has been defined for the application, the migration event can be defined. The origin status of a migration event is Control; the destination status is Archive.



Creating the Object List

Saved, cataloged, and Xref objects that exist only in Control (that is, have been removed from all other statuses), can normally be migrated to the Archive status. These objects may be selected for unloading in the following combinations:

Saved
Cataloged
Saved/cataloged
Cataloged + Xref
Saved/cataloged + Xref

PAC provides several options for creating the object list: you can generate the object list or create it by "manually" entering the objects in the Object List Editor.

You can use the PAC Object List API (APINOBL5) to create or update an object list without logging on to PAC.

Creating the List Manually

If you manually specify objects to be archived

- a version number reference must be used to specify multiple versions of the same object.
- by default, the object version selected will be the oldest version for the object or the lowest available object version number.

- range notation can be used. To select the oldest version of all the objects, enter *,*. (Note that *,* will not select all of the objects).

The SEL (Select) and EXP (Expand) command options can be used in the normal way to select objects for archiving.

Note:

Because objects are validated after expansion, those objects added as a result of the expansion may be rejected based on the application defaults for minimum current version and minimum retention period, and status roles.

Creating the List Automatically

You can generate an object list automatically for a new migration event by

- specifying a value P for the Generate List option when defining a migration event online; or
- issuing the ADD EVENT command either online or in batch, and specifying the value P in the GENTYPE field.

For an existing migration event, you can generate an object list automatically by issuing the GENLIST command with the value P for the GENTYPE keyword to generate a new or replace an existing object list.

Retention Defaults

When generating a list of objects for archiving, the objects selected can be restricted by using the retention defaults defined in the application:

- The minimum current versions specifies the minimum number of object versions to be retained in PAC.
- The minimum retention period specifies the number of days the object must have existed in PAC before it is eligible to be archived.

The minimum retention period is verified only after PAC establishes the minimum current versions for each object to be archived.

Once the objects have been validated against the minimum current versions and the current retention period, PAC selects the object for archiving. The values are checked again during Phase 2 of the archiving process (finalization).

Example:

```
Current date:.....17 May 1993
Minimum current versions..... 3 (Versions 8, 9, and 10)
Minimum retention period..... 90 days
```

Object in a Status	Version	Date Compiled/Incorp.	Archive Selectability
ORD-MAIN	0001	04 Dec 1992	**already archived**
ORD-MAIN	0002	16 Dec 1992	Objects will be selected for archiving.
ORD-MAIN	0003	27 Dec 1992	
ORD-MAIN	0004	07 Jan 1993	
ORD-MAIN	0005	28 Jan 1993	
ORD-MAIN	0006	26 Feb 1993	Not archived. Objects are within the current retention period.
ORD-MAIN	0007	05 Apr 1993	
ORD-MAIN	0008	11 Apr 1993	Not archived. Objects are the current versions.
ORD-MAIN	0009	01 May 1993	
ORD-MAIN	0010	17 May 1993	

Archiving Objects in a Status Other Than Control

To archive objects in a status other than Control (except production):

- Applmod 23 must be active in your environment to allow objects in any status to be selected for archiving. Objects will be removed from the status and placed in the Archive status.
- User exit 26 can be set up in your environment to select objects in a particular status. This exit, invoked when generating an object list, can be programmed to set up and/or verify the criteria (retention defaults) used to select the objects for archiving.

Note:

The actual physical purging of objects from the physical environments identified by the statuses is optional.

If an object is to be purged, you can issue delete commands using either the SYSMAIN utility or the PAC migration utility (MIGRATE). MIGRATE will write the delete commands to the migrate work file, which is then processed by the MIGRATE LOAD utility.

Unload Migration Event Processing

When the migration event has been defined and the object list has been selected and validated, the event can be authorized and submitted.

If for any reason there is an abnormal termination during the unload, or a new or another copy of the archive file is required, you must run the unload step again. The unload step can be run as many times as necessary before the event is finalized.

Unlocking the Migration Event

Phase 1 of the archiving process is complete when the migration event is unlocked. At this point, all objects have been migrated to the Archive status and unloaded to the archive dataset.

The resulting unloaded archive file can only be used as input to the PAC MIGLOAD utility. The migration must be from Archive to Control, a test, or a production status. The dataset is in a special format that can only be handled by PAC during an archive reload migration.

You can verify the successful completion of the archiving process by:

- Reviewing the audit report for inconsistencies, cautions, or possible errors;
- Verifying the archive event and object list and updating it if necessary.

Archive Event

The archive event for use in Phase 2 is created using the information from the migration event (for example, object list and name of the archive file). The name of the archive event is taken from the migration event that created it.

The information contained in the archive event is used to audit the results of the archive. This information is also used as input during the optional reloading of any or all of the archived objects.

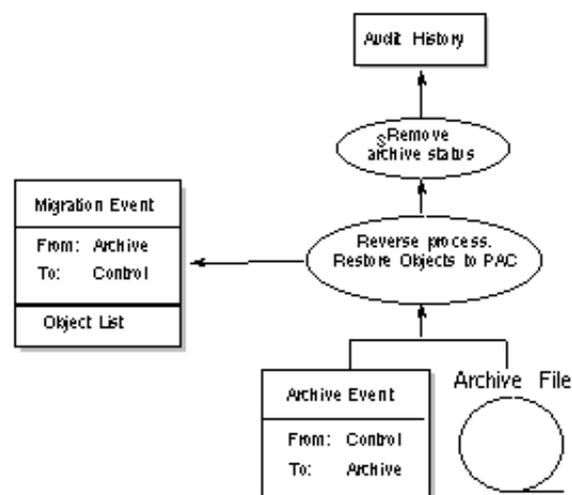
The archive list (that is, an itemized list of every object, including type and version) corresponds to the objects unloaded to the archive file.

Restoring Objects Before Phase 2

It is possible to restore objects to the PAC environment once Phase 1 is completed and before Phase 2 of the archiving process is started; that is, before the cataloged/saved object versions and cross-reference information of the object version(s) are removed from the PAC environment (ACF system file). You may restore the entire archive file or you may restore selected objects from the file.

A migration event with an Archive to Control migration path is set up and submitted. As a result, the objects specified on the object list are removed from the Archive status; that is, the object status link is removed and all activities are written to the audit history.

Any objects restored to PAC after Phase 1 is complete and before Phase 2 begins will not be purged when Phase 2 (the finalize phase) is processed. The following diagram shows this process:



Phase 2 : Removing the Object Versions from the PAC Environment

The second phase of archiving is the actual removal of object versions from PAC. Saved/cataloged object versions and the cross-reference information of each object version are physically removed from the PAC ACF system file.

This is accomplished by the PAC administrator using the finalize option either online or as a batch request. Refer to the PAC Administration documentation for more information about the finalizing procedure.

At this point, the objects are tagged as "archived". You can see the objects' status using the Versioned Objects Reporting sub-function (section Assigning and Maintaining Object Versions).

Selecting the Archive Event

The archive event name corresponds to the migration event that created the archive file in Phase 1.

Finalizing the Archive Event

The finalize step can be run either in batch or online. PAC systematically deletes the saved, cataloged, and cross-reference data for each object version in the archive object list.

The PAC references to the archived objects remain. The reference information retains the object version, date-time saved and cataloged (as applicable), original development directory, and any "Uses" information.

Once the actual deletion has started, deleted objects can be retrieved only by reloading the object from the archive file; however, the archive event may be terminated to prevent any further deletions of the objects.

If the archive event terminates abnormally, it can be restarted. The deletion process will restart at the point where the event terminated.

The archive event is unlocked to complete the archiving of the specified objects.

Restoring Archived Objects to the PAC Environment

Archived objects can be restored (reloaded) to the PAC environment. Once an object has been restored to PAC, it can be used as if it had never been archived.

Reloading is accomplished by running a migration event where the origin (From) status is Archive and the destination (To) status is the Control status.

When objects are reloaded into PAC, the saved and/or cataloged object and cross-reference (Xref) information are restored in the Control status. These objects may then be immediately or subsequently migrated to test or production status types.

Reloading the Migration Event Definition

Only objects from a single archive file may be restored. The archive file created during the actual archiving of the objects must be referenced by the migration event for reloading the archived objects to PAC.

The archive event created when the objects were unloaded to the Archive status supplies the information necessary to identify the archive file. If the archive event for an archived object has been purged from the system, it is not possible to restore any of the objects on the archive file.

The archive file information should be verified. If, for some reason, the file information has changed (for example, the dataset name has been changed or moved to a different volume), the archive file name, the volume serial number of the dataset, and notes may be updated.

The object list will be consistent with the objects unloaded. If the @ARC-DSN and @ARC-VOLSER substitution parameters were used in the job for the migration event that unloaded the objects, the file name and volume serial number will be updated automatically by PAC for the reload migration event.

Note:

For VM/CMS, the substitution parameters are @ARC-VMUSERID, @ARC-VADDR1, @ARC-FN, and @ARC-FT.

Archived objects which have been retired from the Archive status cannot be reloaded.

The object list for the reload migration event should specify only those objects to be restored:

- The *,* notation can be used if all objects on the archive file are to be restored.
- The Generate List option can be used to create a list from the archive event; the list can be modified. Objects can be deleted from the list but not added to it unless they were originally archived by the referenced event.

Reloading Migration Event Processing

The reload migration must be run in batch and must use a work file. The JCL is ARCHIVE_RELOAD.

When the job is submitted, the referenced archive event is determined or verified:

- If the archive event was known (specified): PAC verifies that the archive event name specified on the event is the correct one for reload (that is, that it corresponds with the file specified in the JCL).

Since it is possible to archive and reload an archived object many times, PAC does not require that the user always use the last archive file (with a specific object version) to restore the object.

- If the archive event was not known (not specified), but the archive file is known: PAC determines the correct archive event used to create the archive file during the initiation of the reload.

All objects on the object list for the reload migration event are loaded from the archive file into PAC.

- Any object on the object list that was not archived by the migration event that created the archive file is ignored.
- Any object that is no longer in a status of Archive is ignored, unless Applymod 3 is active.
- If the archive file was created using another PAC system, then all objects are rejected.

As a result of the reload migration, the saved, cataloged and cross-reference data are stored back on to the ACF system file.

- The object status of "ARCHIVE" is removed; and
- The ARCHIVED flag on the object version is removed.