

Managing and Auditing the Production Environment

This section is organized in the following topics:

- Migrating Objects to the Production Environment
- Loading a Job for Activation in the Future
- Production Applications and Libraries / Datasets
- PAA Administrator Functions
- PAA Reports
- Control Information

The production environment is the culmination and purpose of the application life-cycle. The integrity of this environment is critical.

The nature of development and testing environments requires that they be administered flexibly. A wide range of people may be involved in controlling access to applications at these stages of the life-cycle.

However, the production environment requires a high degree of control, along with special procedures to ensure the timeliness of implementation. For this reason, PAC views each production environment as distinct from other environments and uses a special subsystem, Predict Application Audit (PAA), to control it.

A special problem is posed when an application is implemented in multiple production environments and a change to the application is required in one but not all of them, or if the change is required in all of them simultaneously. Both of these problems are addressed with PAA in the first instance, it is imperative that each be controlled separately. Thus, a separate PAA subsystem controls each environment. In the second case the application has multiple defined locations controlled by a single PAA.

The critical nature of the production environment increases the importance of timely implementation. You may need to introduce objects into production at a specific date and time. PAA allows you to migrate an application or object to the environment ahead of time and specify the date and time that it can become active in production.

Sometimes, when a problem is detected with a production application, it is necessary to restore the previous version. When production objects cause system problems, PAA automates the process of removing them and restoring the previous versions.

Auditing and reporting are particularly important in the production environment. You may need quick answers to questions like the following:

- Which objects were active at a certain date and time (for example, the end of the month, or the point where a problem occurred)?
 - Which objects are currently in production?
 - What has changed in the production environment since last night?
-

Migrating Objects to the Production Environment

A PAA-controlled production environment corresponds to a production status defined in PAC. Thus, a PAC migration to the production status is a migration to the PAA-controlled environment.

A PAA load is the process of storing objects in the PAA-controlled environment. The load is usually executed automatically when a PAC migration event to the production status is processed; when the PAA-controlled environment is remote to PAC, the load is executed manually using the MIGLOAD utility.

PAA jobs are different from PAC jobs. A PAA job (or load job) identifies the destination library and the objects to be loaded. The job is created during the processing of the PAC migration event. The PAA job retains attributes of the PAC migration event (for example, the Schedule Date).

Each PAA job identifies a specific load. Thus, unlike a PAC migration event, which can be run multiple times. Each PAA job must be unique. PAA uses the job number to define a unique job.

Migration Steps

The list below describes the steps in a direct migration (where the PAC and PAA systems communicate) to a PAA-controlled environment:

1. A user submits a PAC migration event to migrate objects from a test status to the production status.
2. PAC locks the object versions specified in the object list.
3. PAC reads the object versions and invokes PAA
4. PAA creates a unique job to identify the load.
5. PAA verifies that the user is authorized to submit the load.
6. PAA verifies the application, destination library, and Predict file and locks the library.
7. PAA backs up each object version that is being superseded and stores the new version in the library. PAA audits these actions and stores the audit data in the PAA system file.
8. When the load is complete, PAA unlocks the library.
9. PAC unlocks the objects that were locked in Step 2.

You can view information about the PAA load job, PAC migration event, application, library, and loaded objects.

The PAA administrator can review the job table at any time to determine the state of the PAA jobs.

Control and Production Versions of Objects

Whenever an object is migrated to a PAA-controlled production environment, PAA assigns the object a production version number. PAA stores both the production version number and the PAC version number (called the Control version) in a control record for the object.

For example, when Control Version 1 of a program is migrated from a PAC test status to production, PAA designates it Production Version 1. Assume that Control Version 2 is never migrated to production; when Control Version 3 is migrated to production, it is designated Production Version 2.

Superseded Object Versions and Backups

An object version is superseded when a load replaces it with another version. PAA automatically backs up object versions that will be superseded during a load and assigns the job number to the backup.

Backouts and Reactivated Versions

When a load abends or a problem is detected with object versions that were loaded, you can back out the load job. Backing out a job reverses all actions completed by the load: object versions that were added are purged, and object versions that were superseded are reactivated.

When a version is reactivated, it is restored from the backup and assigned a new production version number. The assignment of production versions is illustrated in the table below:

Control	Production	Remarks
1	1	Superseded by Production Version 2.
2	?	Never migrated to production.
3	2	Load job subsequently backed out.
1	3	Reactivated from backup of Production Version 1.

Finalizing a job prevents it from being backed out. When a job is finalized, the backups of superseded versions are purged. Backups are removed only when a job is backed out (which restores the backups) or finalized.

Loading a Job for Activation in the Future

Unlike other PAC migrations, a migration event to a PAA-controlled production status can be authorized and executed before the Schedule Date. When it processes a scheduled job, PAA loads the objects into the user system file / dataset for foreign objects that constitutes the production environment. However, it does not load them into the application library / dataset; they are inaccessible to the application. PAA designates the load and each loaded object as a Schedule until the PAA administrator activates the job.

When the job is activated, PAA backs up versions that will be superseded and loads the scheduled objects into the application library / dataset.

The dates and times in the following table are used to describe load jobs and the loaded objects. When a job is loaded as a Schedule, these dates and times might all be different.

Date	Description
Schedule Date	The earliest date and time the job can be activated. All PAA jobs have a Schedule Date; except for jobs loaded as Schedules, the Schedule Date is the same as the Load Date and Effective Date.
Authorize Date	The date and time the PAC migration event is authorized.
Load Date	The date and time the objects are loaded into the PAA-controlled environment.
Effective Date	The date and time the loaded objects become active in production. Except for Schedules, the Effective Date is the same as the Load Date.
Backout Date	The date and time the job is backed out.

For example, a job is scheduled for 1999-11-07 (November 7, 1999) at 18:00:00. On November 5, it is authorized at 12:02:33 and loaded into the production environment at 12:10:05.

On November 7, the PAA administrator activates the job at 18:30:30. Two hours later, after finding that an object loaded by the job is causing system problems, the administrator backs out the job.

PAA reports the following dates and times:

Schedule Date	1999-11-07 at 18:00:00
Authorize Date	1999-11-05 at 12:02:33
Load Date	1999-11-05 at 12:10:05
Effective Date	1999-11-07 at 18:30:30
Backout Date	1999-11-07 at 20:35:22

Production Applications and Libraries / Datasets

When a PAC application is migrated to production, PAA derives most of its information about the application from PAC definitions. However, the PAA application definition controls file assignments for application libraries, the Predict file where views/DDMs are stored and the location of the foreign object dataset.

Multiple production locations can be defined for a deployment (application).

To provide flexibility in mixed environments, applications in a PAA-controlled location can be populated with objects either from PAC or from outside the PAC-controlled environment. However, to protect the integrity of an application, you cannot replace a PAC object with a non-PAC object.

Deployment and Location Definitions

The PAA deployment definition consists of the following elements:

- The name of the application;
- The name of the production status;
- The Xref option (whether to include Predict Xref data in loads for the application);
- The default user system file that PAA assigns to the application’s libraries when the user does not specify a file assignment in the library definition;
- The Predict file where views/DDMs and Xref data are loaded;
- The library definition for each PAA library the application uses.

Its name and the database and file numbers of the user system file in which it is located defines a library. Two libraries can have the same name if they are located in different system files. A foreign location is defined by specifying the ESY node, type, format, dataset name and volume.

You must define a deployment (application and library) and location to PAA before you can load objects for the application into the PAA-controlled environment.

Maintaining the Integrity of Deployments and Locations

PAA protects production deployments by securing file assignments and by locking an application location when objects are loaded into the library.

Once you have loaded objects into a deployment location, the file assignment cannot be changed unless the PAA administrator refreshes the deployment, which removes all objects and audit data and resets the library to a Dormant (never used) state.

Once objects or data for the application have been stored in the Predict file, the database and file numbers cannot be changed. During the load process, PAA locks the location accessed by the load. The lock prevents another migration from concurrently loading objects into the location. When the load is complete, the lock is released.

PAA Administrator Functions

The PAA administrator can perform the online functions listed below. Some functions, such as refreshing a deployment or finalizing a job, can also be performed in batch.

Deployment and Location Functions

- Define a deployment and location to PAA.
- Modify parts of the deployment definitions.
- Refresh or purge deployments and locations. Purging a deployment or location removes all objects, PAA audit data, and the deployment or location definition. Refreshing a deployment or location removes objects and PAA audit data but leaves the definition intact.

Job Functions

- Activate scheduled jobs.
- Back out loaded jobs.
- Finalize loaded jobs.
- Purge backed-out or finalized jobs.
- Purge superseded object versions from the Natural buffer pool.

System Functions

- Activate or deactivate PAA system applymods.
- Activate or deactivate PAA user exits.
- Establish system defaults.

PAA Reports

PAA online reports provide information about deployments, locations, load jobs, object versions, and object backups.

Deployment Reports

Deployment reports include the following information:

- The load state of the deployment (active, pending, etc.);
- The date and time of the last migration for the deployment into the PAA-controlled environment;
- The deployments PAA libraries and their locations;
- The most recent jobs processed for the deployment.

Location Reports

Location reports include the following information:

- The physical location of the library;
- The deployment using it;
- Its load state;
- The date and time of the last migration into it.

Job Reports

Job reports provide the following information about load jobs:

- The application for which the migration loaded objects;
- The processing state of the job (backed out, loaded, pending, etc.);
- The user ID of the user who submitted the job;
- The PAC status from which the objects were migrated;
- The name of the PAC migration event and the date and time it was authorized;
- The user who authorized the migration event;
- The date the job became effective in production;
- The name, version, and type of each object included in the job;
- The name and version of the file translation table used in the job.

Audit Reports About Production Objects

In screens and pop-up windows, audit reports provide a wide range of information about object versions in production:

General Information

PAA reports the following information about each object version:

- The state of the object version (current, backed out, scheduled, etc.);

- The dates the object was unloaded from the PAC-controlled environment, loaded into the PAA-controlled environment, and activated in production;
- The date the object version was superseded or backed out;
- The corresponding production and Control version numbers;
- The application to which the object belongs and the library and status where it resides;
- The PAA load job that included the object version;
- The related PAC migration event.

Directory Information

PAA reports the following information from Natural directories:

- The library in which the object was saved or cataloged;
- The user who saved or cataloged the object;
- The date and time the object was saved or cataloged;
- The sizes of the saved and cataloged object;
- The Natural version, operating system, and TP system used to create, save, or catalog the object.

Summary Reports

On a single screen, summary reports provide the following information about object versions:

- Production and Control version numbers;
- Number, date, and time of the load job;
- The date the object version became effective in production;
- The earliest date the object can be activated (for Schedules only);
- The date the object version was backed out (if applicable).

User Reports

PAA provides APIs for reporting purposes. These APIs enable you to design and customize your own output reports based on PAA held data.

Control Information

The PAA administrator can display a snapshot of the current state of the PAA-controlled production environment:

- The date and time of the last migration into the environment;
- The numbers of active and pending load jobs (these jobs have locked PAA locations);
- The numbers of load jobs that have been backed out;

- The database and file numbers of the PAA system file where audit data is stored;
- The number that will be assigned to the next load job.