precisely

Assure iOptimize

Assure iOptimize User Guide

Version 10.0

Notices

Assure iOptimize User Guide
October 2021
Version: 10.0.01.00

Precisely
1700 District Ave Ste 300
Burlington MA 01803-5231
USA

Trademarks

See www.precisely.com for information about our valuable trademarks.

The following are trademarks or registered trademarks of their respective organizations or companies:

- AIX, AIX 5L, AS/400, DB2, eServer, FlashCopy, IBM, Informix, i5/OS, iSeries, MQSeries, OS/400, Power, PowerHA, System i, System i5, System p, System x, System z, and WebSphere—International Business Machines Corporation.
- Adobe and Acrobat Reader—Adobe Systems, Inc.
- HP-UX—Hewlett-Packard Company.
- Teradata—Teradata Corporation.
- Intel—Intel Corporation.
- Linux—Linus Torvalds.
- Excel, Internet Explorer, Microsoft, Windows, and Windows Server—Microsoft Corporation.
- Mozilla and Firefox—Mozilla Foundation.
- Java, Solaris, Oracle—Oracle Corporation.
- Red Hat—Red Hat, Inc.
- Sybase—Sybase, Inc.
- UNIX and UNIXWare—the Open Group.

All other brands and product names are trademarks or registered trademarks of their respective owners.

Support

Getting technical support: Customers with a valid maintenance contract can get technical assistance via Support. There you will find product downloads and documentation for the products to which you are entitled, as well as an extensive knowledge base.

Contents

Introduction

1

About this Manual

This manual introduces Assure iOptimize and helps you begin to configure basic optimization functions as well as providing detailed instructions for using the product.

This manual is for use with iOptimize version 10.0.01.00 and later.

It is recommended that you interact with iOptimize through the Assure iOptimize portal application which is automatically installed as part of the Assure iOptimize product installation when the Assure iOptimize installation wizard is used. You can configure the iOptimize portal application as an instance in the Assure Unified Interface using the *Assure Unified Interface User's Guide*. Also see the *Assure iOptimize Portal Application User Guide* for more information about using iOptimize through its portal application.

The Assure Unified Interface (the Assure UI portal) is a browser-based, graphical interface that presents Precisely product portal applications for management and monitoring in a single, unified view. See the *Assure Unified Interface User's Guide* manual for more information about the Assure UI portal.

About Assure iOptimize

Assure iOptimize is a complete management tool; it is not just a monitor or a guide to problem areas. iOptimize automatically covers all the fundamental areas of object management and optimization where performance can be improved or disk space can be released. The range of functionality for monitoring, managing, and optimizing an IBM Power Systems includes object management and optimization, automated job and user tracking, and complete cross-reference including IFS objects, with full forecast action control and action history. iOptimize has relatively low system overhead and requires only specific time slots outside of normal working hours during which it manages the execution of tasks such as reorganization, optimization, or housekeeping.

iOptimize gathers information, performs housekeeping actions, monitors performance, tracks jobs, and reports its findings and results. Additionally, it:

- gathers information on disk, object, and job throughput
- highlights areas where disk savings can be made and where performance improvements can be realized automatically
- performs the necessary jobs, unattended, to make the performance gains and space savings it previously identified
- identifies change (for example, day-to-day or week-to-week and highlights where user intervention could result in additional savings
- tracks and reports on object growth and job throughput including temporary storage and response times
- monitors performance and summarizes system health as a fundamental aid to capacity management

NOTE

See "Security Policy" on page 205, for information on iOptimize's security policy.

Support for independent ASPs

In addition to supporting the functionality described in the previous sections for *SYSBAS and user ASPs, iOptimize supports the same functionality on IASPs. A few limitations or special behaviors are of note for the IASP:

- iOptimize cannot be installed on an IASP. It can be installed only on the *SYSBAS to which the IASP is configured.
- There are some limitation to object types that are supported on the IASP. For a list of objects supported by the OS installed on your systems, refer to the IBM Power Systems Independent Disk Pools manual for your OS version.
- The IASP must be in an available and varied on state for iOptimize to have access to manage objects. (Other IASP states include varied off and unavailable.)
- An IFS cross-reference is not run on an IASP.

For managing objects on IASPs:

- When v10.0 is installed, iOptimize automatically creates a library on the IASP (IOPTIASP) that facilitates management of objects.

NOTE

A full cross-reference is automatically processed the first time iOptimize is started. If you want to configure the cross-reference, you must do it before starting the product (see "Real-Time Cross-Reference" on page 21).

- When a cross-reference is run, iOptimize processes all libraries on all available IASPs by default.
- When working with objects on an IASP from within the iOptimize product, you do not have to use the command SETASPGRP to access the IASP — iOptimize automatically does that for you. When you are no longer working with objects on the IASP, iOptimize returns you to the last known settings.
- Although an IASP may be varied off or not available, libraries and other objects will be maintained by the cross-reference tables. Even though the device is temporarily not available, the objects may still exist. If the objects no longer exist on the IASP when it becomes available or is varied on, an automatic cross-reference is requested to remove these libraries.

Support for Notifications in the Assure iOptimize Portal Application

The Notifications portlet in the Assure iOptimize portal application provides a method to provide notification within the portal application for iOptimize subsystem and job status issues that are predefined for iOptimize. A notification is an automatic report associated with an event that already occurred.

You can add your own notifications using the iOptimize command IDADDNFY. The severity of any new notifications is reflected in the overall status of the instance. See "Add a Notification" on page 208, for more information on adding your notifications using the IDADDNFY command.

Accessing and Using the Assure iOptimize Interface

2

Accessing iOptimize

You must have a user profile set up in the IBM i environment and your System Administrator must have previously given you access to the areas of iOptimize for which you have responsibility.

Use the following command to start the Assure iOptimize product:

```
installation-library/STRID *ALL
```

See "Starting and Ending Product Subsystems" on page 200 for information on the command parameters.

```
IDMAIN                    iOptimize V10.0 Main Menu
                                                              System: ABCDE
Select one of the following:

     1. Object analysis
     2. Optimization
     3. Operations
     4. Configuration
     5. Accounting code analysis

    15. Data Manager

    90. Sign off


                                                                      Bottom
Selection or command
===>
F3=Exit   F4=Prompt     F7=Task Manager   F9=Retrieve   F12=Cancel
F16=System i Main menu   F23=Set initial menu
```

How to Lock/Unlock a Session

Assure iOptimize has an Unattended Terminal Security policy that allows you to lock a user session. This allows you to secure access to your login if you have to interrupt your work activities briefly.

To lock a session, press F2 (Lock). The iOptimize Unattended Terminal Security screen is displayed.

To unlock a session, type your password (the IBM Power Systems password for the user profile with which you logged on and locked the session) and press Enter.

How to Use the iOptimize Interface

iOptimize is designed to easily access its features. The features are described in the following table.

Menu Item	Description of the Features
Object Analysis	Track and maintain information on the usage, occupancy, and growth of objects on the system. At any time, you can easily locate the answer to questions you may have about an object, including its size, usage, ownership, contents, relationships, condition, and who created it. You can also manage collections and multi-dimensional domains. Object Analysis has many preconfigured inquiries (see Chapter 7).
Optimization	Optimize physical files and tables, logical files and tables, programs and modules, temporary storage, spool members, oversized objects, device files, menus, and panel groups (see Chapter 9).
Operations	Work with housekeeping, spool file archiving, managing spool files, object archiving, subsystem analysis, and history of work with active jobs.
Configuration	Work with cross-references, calendars, schedules, and licensing (see Chapter 5). Configuration, controls, and policies are managed through the iOptimize portal application.
Accounting code analysis	Enables you to collect and analyze job accounting computer usage data (see the *Assure iOptimize Job Accounting User Guide* for more information.
Data Manager	Accesses Data Manager, if installed.

Using the Position to Field

Screens displaying lists may include the “Position to” field. This field refers to the primary column on the screen. Entering a full or partial name in the “Position to” field will adjust the position so that the list displayed for the primary column starts with the name entered. In the following example, the Position to field represents the Library column:

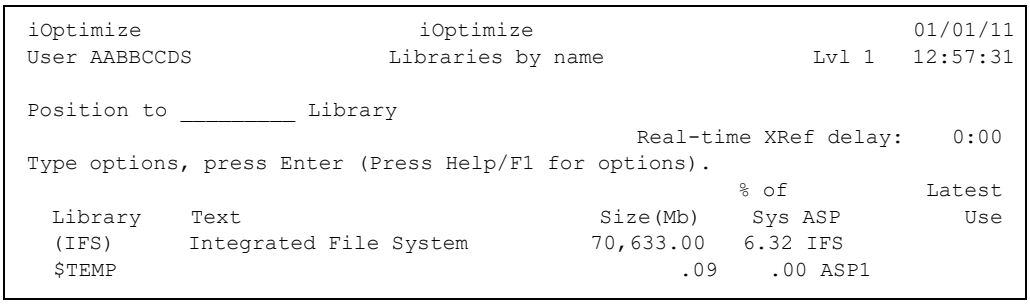

Accessing Screens and Menus

The following graphic is an example of the method used in this manual to help you identify the path to locate screens and menus in iOptimize.

Operations —> Option 1 (Housekeeping policies)

In the example, the path indicates that you will access the Housekeeping policies menu by using option 1 from the Operations menu.

Getting Started

3

Jobs Started During Installation

The options listed below are for Assure iOptimize jobs that start during installation. These job are shipped with standard options that can be changed at any time to effectively use iOptimize.

You can run all of the iOptimize jobs from any job scheduler (including the IBM job scheduler) since iOptimize is almost entirely command driven. It is recommended that you use only the standard shipped options initially.

Parameter	
Job queue (Operations)	iOptimize contains a real-time cross-reference that is object level. Other processes are required to keep the cross-reference up to date (for example Job, System and File statistics). Users can select any job queue and/or any job description on the system to run cross-referencing jobs. This configuration is defined in the Cross-Reference Configuration window in the iOptimize portal application (see the *Assure iOptimize Portal Application User Guide*). Note**:** When starting the iOptimize subsystem after a new installation, please ensure it has sufficient memory to complete the initial object collection, which can take some time on a large system.
Job Queue (IFS)	iOptimize periodically submits a job to keep the IFS cross-reference tables up to date. This job will be released during the install procedure. It will provide extremely valuable information about your machine, will not affect any other jobs and can be terminated at any time without adverse effect.
The Parameters	The parameters which have been provided as defaults will ensure that iOptimize is set in forecast mode (*GENFCS) for functions such as Optimization Physical Files, Logical Files, Programs and Outqueues, just some of the areas where you can start to save space. You can view the forecast but no objects will be changed.

iOptimize is preconfigured to perform housekeeping functions during Unused run times using the default job queue (see “Understanding Run Times” on page 27). You can change these default settings , however when first setting up configuration it is recommended that you leave all tasks set to run during Unused time. This way there will be minimal impact to your users. After becoming more familiar with it, reconfigure the autorun times to make best use of Prime, Off-Peak, and Unused time periods.

Performing Object Analysis

One of the most powerful iOptimize features is the ability to get information on the usage, occupancy, and growth of objects on the system. Object Analysis tracks and maintains this information automatically and online. At any time, you can easily locate the answer to questions you may have about an object, including its size, usage, ownership, contents, relationships, condition, and who created it. Object Analysis has many preconfigured enquiries.

Exploring Object Analysis Enquiries

Object Analysis contains a variety of enquiries that provide you with valuable information about the objects on your system. You can use this information to save disk space by identifying objects that could be reduced in size or removed from disk. In addition, programmers and systems personnel will benefit from a wealth of easily accessible and useful information that is produced through the enquiries.

It is important to understand that these enquiries (and the corresponding reports) are simply different ways to isolate information. The standard enquiries are those that provide the greatest immediate benefit to a wide range of organizations. The real power of the enquiry is the ability to find out everything you need to know about an object (by drilling down using action codes) and then using the tools (such as archiving) to act on the information obtained.

All of the enquiries are similar in layout and have the same features available for examining objects (for example, displaying file layouts, program calls, usage, dependencies, and more).

On each enquiry screen, you can access a help feature that displays a list of available action codes, which you can use to perform various functions on that screen. See “Action Codes (F1/HELP)” on page 89 for an explanation of each code.

Real-Time Cross-Reference

To provide online access to detailed information about your IBM Power Systems status, iOptimize maintains cross-reference files automatically, using a real-time cross-reference feature. A *FILE cross-reference is scheduled to run every day at 4:00:30 a.m. by default. A*FILE cross-reference can also be run or scheduled through task manager. For more information, see “Cross-Reference Jobs in Task Manager” on page 67.

The real-time cross-reference feature maintains iOptimize’s database of IBM i object information on a nearly real-time basis. iOptimize detects object changes by reading journal entries in the security journal, QAUDJRN. A new object, an object change, or an object deletion results in an immediate updating of the iOptimize database.

The real-time cross-reference is made up of two processes. The first monitor has a job name of IDDYNMON. This process reads the QAUDJRN journal, picking out relevant journal entries and passing the appropriate information to the second monitor which has a job name of IDDYNPRC. IDDYNPRC then applies the update to the iOptimize cross-reference database. These processes are designed with autonomics such that if the jobs are cancelled or fail for any reason, the iOptimize subsystem will automatically restart them.

NOTE

The real-time cross-reference is automatically configured when the subsystem starts for the first time. The iOptimize subsystem should be active at all times. Cross-reference activity will change an object’s auditing value from *NONE to *CHANGE for all objects on the system except for files and the following object types: *DTAARA, *DTAQ, *FILE, and *USRSPC (the cross-reference will leave their values at *NONE). However, if the object audit value was set to *CHANGE or *ALL prior to running the cross-reference, that object auditing value remains for all object types. The value of *CHANGE (or *ALL) is subsequently used by the real-time cross-reference to pick up object changes.

Real-time cross-reference parameters can also be specified from the iOptimize Configuration menu. In addition, libraries, files, and spool files can be manually added to the cross-reference list using IDXREF. For instructions on configuring cross-reference parameters or manually adding to the cross-reference list, see “Working with Audit Monitors: Real-Time Cross-Reference” on page 51).

Setting System Values

System values are automatically set up by iOptimize. If they need to be changed or re-set, it should be done when the iOptimize subsystem is inactive. The table below includes the system values required for cross-referencing, job accounting, and spool file data. This allows iOptimize to collect the information required by the real-time cross-reference feature. Once the subsystem is restarted all processes should start or adjust automatically. If the Audit Monitor is not active you may need to configure the current receiver and re-start the monitor (see "Working with Audit Monitors: Real-Time Cross-Reference" on page 51).

SYSVAL	Value	Description
QAUDCTL	*NOQTEMP	Suggested, but not required
QAUDCTL	*OBJAUD	Required for real-time cross-reference
QAUDCTL	*AUDLVL	Required for real-time cross-reference
QAUDLVL	*CREATE	Required for real-time cross-reference
QAUDLVL	*DELETE	Required for real-time cross-reference
QAUDLVL	*OBJMGT	Required for real-time cross-reference
QAUDLVL	*PRTDTA	Required for Job Accounting
QAUDLVL	*SAVRST	Required for real-time cross-reference
QAUDLVL	*SECCFG	Required for real-time cross-reference (Profiles)
QAUDLVL	*SECRUN	Required for real-time cross-reference (Profiles)
QAUDLVL	*SECVLDL	Required for real-time cross-reference (Profiles)
QAUDLVL	*SPLFDTA	Required for real-time Spool File cross-reference

Excluding Libraries from Cross-Referencing and Optimization

You can exclude libraries from real-time cross-reference processing and optimization from the Cross-Reference Configuration window in the Assure iOptimize portal application (see the *Assure iOptimize Portal Application User Guide*). From the 5250 emulator, you can exclude libraries using the IDEXCLIB command. The excluded libraries are added to the IDAUXMP file.

1. Prompt the following command:

   ```
   IDEXCLIB
   ```

   The Exclude from Processing screen is displayed.

2. Enter the following parameters:

Parameters	Description
Library	The library to exclude.
Object type	The object type to exclude.
Exclude from auditing	Whether to exclude the library from cross-referencing. • *YES–The library is excluded. • *NO–The library is not excluded.
Exclude from optimization	Whether to exclude the library from optimization. • *YES–The library is excluded. • *NO–The library is not excluded.
Reset object auditing value	Press F9 (All parameters) to display this parameter. Resets the object auditing value for all objects to their values before being cross-referenced. • *YES–the value is reset. • *NO–the value is not reset.

3. Press Enter to run the command. The cross-reference for excluded libraries can also be run from task manager. For more information, see "Cross-Reference Jobs in Task Manager" on page 67.

Managing System Files to Increase Disk Space

The Optimization menus contain a variety of options that enable you to increase available disk space. These include:

- clearing unused spool members
- optimizing oversized objects
- clearing replaced objects

Managing Queues

iOptimize saves disk space by regularly managing the IBM Power Systems queues (Output queues, Job queues, Message queues). A number of system queues are supplied by IBM, including such standard defaults as QPRINT (output), QBATCH (job) and QSYSOPR (messages).

In addition to these, every time that you configure a new workstation or add a new user to your system, the operating system (IBM i) automatically creates an associated message queue and, for every printer that you attach, it creates an associated output queue. In addition, if you use packaged application software, it is typical that the authors created their own set of queues for use within the application. If you develop your own software, it is extremely likely that you have set up special job queues and output queues for use in the development and test environments.

Since the number of queues on your system can grow rapidly and it is difficult to keep track of them, iOptimize can help by automatically keeping track of all the queues on your system and clearing and/or archiving entries that have exceeded your configured time limits. Once you understand the basics, you can run iOptimize housekeeping on some of your own queues, resulting in a real saving on disk space.

Managing Logs

The IBM i logs serve as a record of almost everything that happens on your system, which is useful for understanding the circumstances surrounding system events.

Logs record details such as normal system activity (History log), system problems (Problem log—used in IBM Power Systems Problem Analysis), network communications (Distribution log), journaled data (IBM journals—used for recovery purposes in the event of system failure), and copies of messages (Alerts), including internal IBM i messages and any user-specified messages.

The amount of space taken up by these logs can be quite considerable if they are not managed regularly. This is because the operating system continues adding to them, doing very little to clean them up. The result is a continually increasing volume of historical data, which decreases in usefulness with the passage of time.

As with the proliferation of queues, you may be unaware of many of the internal logs and, consequently, take no action to manage them. Once again, iOptimize can help by automatically keeping track of system logs and clearing entries that have exceeded your configured time limits.

Reclaiming Temporary Storage

Temporary storage is taken up by temporarily expanded objects. This space can, and should, be released shortly after it has been temporarily used. Where it hasn't, iOptimize seeks out and releases this space for general use.

Spool members are physical file members which hold printed output data. They are cleared by the system after printing is complete and are available for re-use, but are not removed from the file. Consequently, the number of

such members always corresponds to the maximum number of prints that have ever been processed on your IBM Power Systems. Each one, though clear of data, continues to hold the amount of disk space that was allocated when it was first created.

Many IBM Power Systems objects, such as libraries and message queues, grow dynamically to meet processing requirements. Even if you clear them regularly, the space allocated for these objects remains at the largest size the object has ever been.

When creating a new object, IBM i frequently saves the object being replaced, giving it a new name and description. The library containing these increases in size unless it is managed.

See "Working with Temporary Storage" on page 182 for instructions on reclaiming storage.

Optimizing User Data to Increase Disk Space

You can increase disk space by managing your application files and programs and by accomplishing a reorganize of files. Performing all of these tasks manually takes extensive time and knowledge, as well as consistency. But iOptimize performs all of these tasks automatically, according to your configured specifications.

Optimizing Physical Files

iOptimize automatically reorganizes physical files to remove deleted records. As part of this function, you can configure it to reorganize physical files so that the data is held in a particular sequence. This is very important, and is especially useful for applications that use unkeyed physical files, relying on logical files for keyed access.

As an example, consider an orders file with a most frequent access route by customer. Regularly reorganizing your orders physical file so that data is held in customer sequence ensures considerably better access performance, resulting in your application running faster.

See "Working with Physical Files and Tables" on page 160 for instructions on optimizing physical files.

Deleting Unused Records

There may be real (undeleted) records in a file that have no use, but your application has not deleted them (for example, old transaction history records, records for a company or product that no longer exists). In this case, the file needs purging which can release significant amounts of space and improve performance.

Data Manager, iOptimize's companion product, removes and optionally archives these redundant records from within the physical data file. These purge routines can be set up, by non-technical staff, for any physical file in approximately ten minutes. Once created, these file purges and archives can be scheduled to run on a regular basis to keep files in check. Naturally, any data that is archived can be immediately retrieved back online at any time. Contact Precisely for additional information on Data Manager.

Optimizing Logical Files

The order in which logical files are created can have a significantly adverse effect on the efficiency with which IBM i shares access paths. The importance of this issue to the efficient use of disk space and to the optimum performance of your IBM Power Systems should not be underestimated.

iOptimize checks all of your logical files (even those in packaged applications) and recreates them as necessary, making the maximum use of shared access paths. iOptimize also analyzes all of your logical file access paths currently being maintained on an immediate (*IMMED) basis and switches those files that are clearly not being used regularly to rebuild (*REBLD). This saves considerable space by reducing the overall number of access paths and significantly improves update times by ensuring that only essential access paths need updating.

See "Working with Logical Files and Indexes" on page 169 for instructions on optimizing logical files.

Optimizing Programs

Programs can be made to run more efficiently by organizing the code in a particular way. This fact is not widely understood, and the means of achieving the desired result is even less so. iOptimize can optimize all of your programs automatically, allowing you to have the best possible performance from them.

In terms of saving space, iOptimize can perform tasks on programs such as compressing, or even removing, the observable information.

NOTE

Compressing the observable information will save disk space and has no adverse impact on your applications. Removing this information will reduce the ability to debug program failures and you will want to discuss this with Support staff before taking this option.

Optimizing Other User Data Types

For instructions on optimizing additional types of user data, see "Working with Device Files, Menus & Panel Groups" on page 188.

Setting Up Automatic Processing

In previous sections, you saw how easy it is to configure iOptimize to process the various internal (IBM i) and external (user) elements on your system to suit your needs. Processing options include daily, on selected days only, and according to time limits (for example, only clear save status prints after 30 days). Automatic processing takes you a step further, to put the configured processing into action.

Understanding Run Times

When scheduling jobs to run in normal (*NORMAL) mode, it is possible to schedule all the automatic optimization and management tasks to run at any time. However, it is important to carefully select the times at which certain jobs run (for example, Reorganize physical files), to ensure that other users are not running applications that use the relevant files or resources at the same time. This section describes how to restrict iOptimize to run only at specified times; it will not run during any other times.

NOTE

In this section, the focus is on ensuring that iOptimize tasks run only when you define them to. To avoid potential conflict with your users, set up the run times before releasing the clean-up jobs to run in normal (*NORMAL) mode.

Within the specified run time windows, iOptimize automatically organizes its tasks, enabling it to complete the necessary optimization and housekeeping tasks on a cyclical basis. For example, prior to initiating a task, iOptimize estimates how long it will take to complete the task, to ensure that it does not start a job that can't be completed by the specified end time. Where an action (for example, a reorganization of a very large file with many logical files) cannot be completed within the run time window, iOptimize bypasses the action on the object, records that it was unable to perform the task due to time restrictions, then carries on processing other objects (there may well be many smaller files that could be reorganized in the available time) until the end of the run time window. When iOptimize runs out of time before completing all of the specified tasks, the next time it runs it starts where it left off, giving preferential treatment to objects that were bypassed due to lack of time in the previous run time window, thereby ensuring that all tasks will eventually be implemented, no matter how small the run time window.

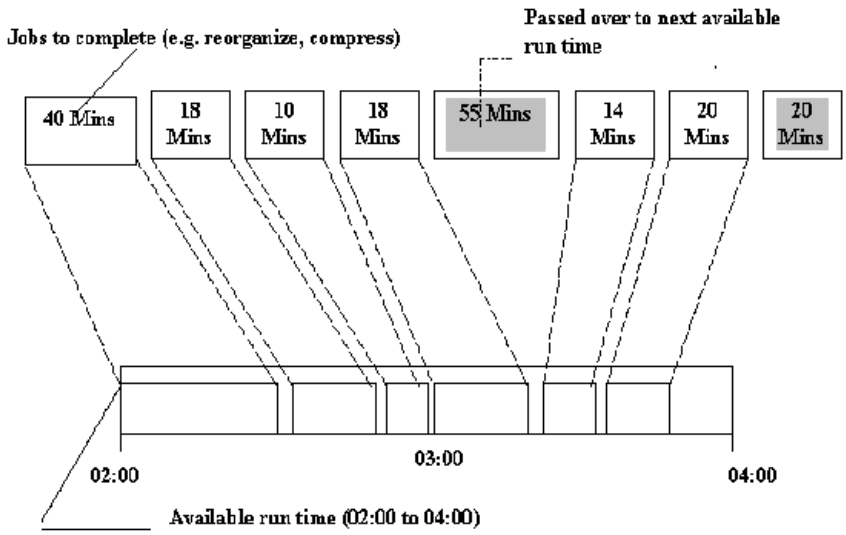

Setting Run Time Windows

You configure the different run time classifications (Peak, Off-peak, and Unused time periods) from the Policies portlet of the iOptimize portal application. It is not necessary to categorize all of the possible timeslots; categorize only the timeslots during which you want to run iOptimize. If, for example, you only want to let iOptimize run in 'U' (unused) time, just define the unused) time periods, and leave the rest empty. (See “Setting Processing Windows” on page 198 for more information.)

Selecting Run Time Windows

iOptimize is preconfigured to perform housekeeping functions during Unused run times using the job queue that you specified during installation. You can change these default settings to suit your needs. However, when first setting up configuration, it is recommended that you leave all tasks set to run during Unused time, so that there will be no impact on your users. After you become more familiar with iOptimize and its functions, you will want to reconfigure the autorun times to make best use of Peak, Off-Peak, and Unused time periods (see “Maintain Autorun Parameters” on page 192 for instructions).

Changing the Default Job Queue

You can change the job queue that a particular function uses, from the iOptimize Maintenance screen. A single-streaming job queue that does not conflict with any task that you want to run in that time frame is recommended. Typically, the same queue is used for all functions.

Using this approach, iOptimize reorganizes one library at a time, addressing a file only after the previous file has either been reorganized or rejected from the task, either because no work is required or because there is insufficient time left to complete the task.

You can also select a number of Job Queues, which may be used, for example, to reorganize the files in a number of libraries simultaneously within that time frame.

See "Configuring Multi-Threading Parameters" on page 194 for instructions.

You can change all of the actions currently on a queue to a new queue. See "Making Global Changes" on page 195 for instructions.

Working With Processing Modes

You can view and change current processing mode configuration from the Maintain Processing Parameters screen. See "Maintaining Processing Parameters" on page 197 for instructions.

It is important to review the action parameters set for each job before releasing it to run in *NORMAL mode, to ensure that you are aware of all the objects the clean-up jobs are set to action. This is good practice and has the added benefit of enabling you to change the processing mode from the same location. See "Configuring Action Parameters" on page 194 for instructions.

See "Understanding Processing Modes" on page 44 for a review of the three processing modes.

Reviewing System Logs

You can review system logs, such as run time history, iOptimize action history, iOptimize errors. See "Viewing System Logs" on page 195 for instructions.

Ending Submitted Tasks

Occasionally you may have the need to use the IBM Power Systems during a time when you have scheduled iOptimize to run. For example, you may have allocated Sunday afternoon/evening as Unused time, and scheduled iOptimize to run during this period, but then you elect to perform an operating system upgrade on Sunday afternoon.

See "Ending Submitted Jobs" on page 198 for instructions.

Managing Archiving

iOptimize makes it easy to manage your near-line and off-line storage. iOptimize maintains a list of all your archive tapes, together with details of what they contain and when they were used for archiving. It also allows you to import your current archive tapes. When importing tapes, iOptimize analyzes and catalogs each tape, thus giving you greater control over these important items.

iOptimize archiving is a two-stage process. iOptimize maintains a library on disk (/VisionSolutions/IOPTIMIZE/Archive) that holds objects you want to save. Objects are transferred to this library when you select the Archive option, which is available from all of the iOptimize inquiries. When you decide that you want to move these objects to tape, you use the Archiver portion of iOptimize. This is what we refer to as “near-line” storage—it's off-line (on tape, optical, etc.) but cataloged and very easy to find and restore. The beauty of this approach is that it simplifies archiving by making it easy for you to select what you want to save. You can then transfer your selections to tape when it is most convenient for you. And, it means that you only have to save one library.

Restoring objects to your IBM Power Systems is simple. iOptimize logs everything that you save, so it is easy to determine which tape an object is on, load the tape, and let iOptimize restore it.

See “Working with Archiving” on page 145 for instructions.

Exploring the Integrated File System

The Integrated File System (IFS) is the term applied as a generic to the various filing systems now supported on IBM Power Systems processors. The IBM Power Systems can now support filing systems from the OEM environments, such as the UNIX filing systems, which can be managed by the IBM Power Systems alongside the traditional QSYS filing systems. The earliest of these alternate filing systems, and the one most users will have been aware of, was the QDOC library. This is now the QDLS filing system within the IFS.

iOptimize’s Cross-Reference job runs IFS cross-references which iOptimize catalogues and reports on all of the QSYS objects in great detail. In addition, it reads the rest of the IFS and catalogs it into a new set of files for user enquiry (without updating the ‘Date Last Accessed’ information). This file is then manipulated by the various enquiry/report programs to highlight areas that may be of concern (for example, the largest object in the IFS is 300MB and it is owned by Peter, or file XYZ takes up 70MB and has not been used since January 27, 1997).

Object analysis for the IFS is configured and managed using iOptimize portal application. See the *Assure iOptimize Portal Application User Guide* for more information on iOptimize analysis for the IFS.

Managing Jobs

iOptimize is a comprehensive product and most of the system management jobs that you would like to get around to, if you only had time, are performed automatically by iOptimize.

From time to time you will need to know which jobs iOptimize is running, when it intends to run cleanup jobs, when it intends to run collection jobs, and, of course, what jobs it has run. Incidentally, many of these job-tracking functions are just as useful for looking at your own applications.You can run all of the iOptimize jobs from any job scheduler (including the IBM job scheduler) since iOptimize is almost entirely command driven.

You can check the status of scheduled jobs, add new jobs, view history, amend scheduling, configure the Task Manager environment and start/stop the job scheduler from the Task Manager screen. See "Task Manager" on page 66 for instructions.

iOptimize needs to collect file level information and this part will need to be scheduled as frequently as possible. This part of the cross-reference is referred to as a *FILE cross-reference. To ensure that all files on your system get cross-referenced, a *FILE cross-reference is scheduled to run every day at 4:00 by default. A*FILE cross-reference can also be run or scheduled through task manager. The file can be rebuilt either in its entirety or on a selective basis at intervals you specify. For more information about using task manger to change parameters for the *FILE cross-reference job, see "Changing Job Parameters" on page 70.

Submitted Housekeeping Jobs

Housekeeping jobs are the jobs that manipulate objects on the system to optimize, compress, and reorganize. Typically these jobs would be run at times when the system does not require the objects being worked upon but, for the vast majority of the tasks, there are ways around this if free time is at a premium. You can limit, for example, the jobs to run only during a fixed time such as 02:00 hours to 07:00 hours over the weekend.

iOptimize can be set to ensure that, regardless of where the jobs are actually submitted and at what time, the jobs are only authorized to run in a specific time window. This means that if jobs get inadvertently held up, and are then started outside of their appropriate time slot, iOptimize ensures that the jobs immediately close down harmlessly. They will automatically reset to run during the next authorized run time. This makes iOptimize entirely safe.

See “Maintain Autorun Parameters” on page 192 for instructions.

Working with Job Tracking

Job tracking is an independent task that records information about the jobs that are moving through the system (without manipulating them). It can therefore be ignored as far as its potential to affect other jobs. Any job in the system can be tracked, whether it has been submitted by Task Manager or otherwise. As a job is tracked its identity is kept for a user-specified period, say 30 days, and details of its execution can be viewed for this period.

Job tracking is controlled using the Scheduler function in the iOptimize portal application. You can view more about the frequency of this job and the length of time it is keeping information, and of course change these values. Job tracking is a low-to-no overhead function, and the greater the frequency the more up-to-date the available information.

For example, if we decided to keep a track of information on jobs for 20 days or so (and we have not cleared output queues and joblogs within this period using iOptimize) we could look back at a specific job to see how it executed, whether it completed normally, look at its Job log, view its spool file entries, etc. If the job was currently running we could actually view the job from this display. iOptimize’s powerful Task Manager enables you to view this job information easily.

By using this data intelligently, we can provide excellent information on response times, job run and completion data, reconstruct a pseudo WRKACTJOB for any historical point in time (within the tracked data time limit you have specified), etc.

Job Tracking is managed through the iOptimize portal application.

Reviewing Job History

Once job tracking has been initiated, you can review the history for jobs on the Task Manager schedule. Using the iOptimize Job Tracking feature, you can see exactly which jobs completed and when they started and stopped. This is useful for confirming that jobs ran, and, more importantly, if they did not run, why they failed.

Response Time Monitoring

Job response monitoring is one of the functions of the Job Monitor. The Job Monitor is automatically started when the iOptimize Sub-System is started.

Response times can be monitored by date, user or job. As jobs are tracked, one of the important pieces of information gathered is the response time of interactive jobs. This enquiry groups, summarizes, and details actual

response times in a way that helps your operations staff to highlight areas of poor response time, and identify the causes (perhaps with the help of other Job tracking functions).

Response time monitoring is managed through the iOptimize portal application.

ASP Monitoring

The ASP Monitor is automatically started when the iOptimize subsystem is started. The ASP Monitor records ASP usage statistics for Event Manager.

Job Accounting

Job accounting enables you to collect and analyze job accounting computer usage data including the following:

- Number of print files, lines, and pages created by programs
- Number of database read, write, and update operations
- Number of communications read and write operations
- Actual lines and pages printed
- Time the job was active and suspended
- Actual number of bytes of control information and print data sent to the printer

The data collected by job accounting is used to cross charge computer usage, including one-off and recurring charges, to user divisions within the organization. This charging process can be used to manage and control costs ensuring the best use of time and resources.

iOptimize records job resource data at the job level, however job accounting can provide detailed information about resources such as database, printer, communications, or segments within a job. Job accounting can also provide information for printed output as well as active and suspended time information.

Instructions to set up and use job accounting are contained in the *Assure iOptimize Job Accounting User Guide*.

Planning for Optimization and Management

4

This chapter provides an overview of the many benefits of Assure iOptimize and the resulting Direct Access Storage Device (DASD) space savings you can attain by configuring iOptimize to process the various internal (IBM i) and external (user) elements on your system.

NOTE

Throughout this manual you will find suggestions for setting parameters to obtain meaningful and useful results. You can change any of the parameters to suit your own requirements, either initially or after a forecast. When using parameters other than those suggested in this manual, please keep a record of them.

The initial objective of tasks provided in this manual is to set up global initial parameters and to pick up valuable information about your system, which will both assist you in assessing your system and help you to manage your system more effectively.

About Performance Optimization

Before you start full-scale performance management by administering the pools and the timeslices of jobs, and before you even set the IBM Power Systems automatic pool management, you need to address the fundamentals.

Assure iOptimize has many features where performance, as well as storage optimization, is considered; good optimization of the storage itself can also bring huge benefits in performance.

There are many areas in which iOptimize can aid or boost performance directly and automatically. Once iOptimize is in general use, there is a cycle which runs, constantly monitoring, forecasting, actioning, and reporting on object management.

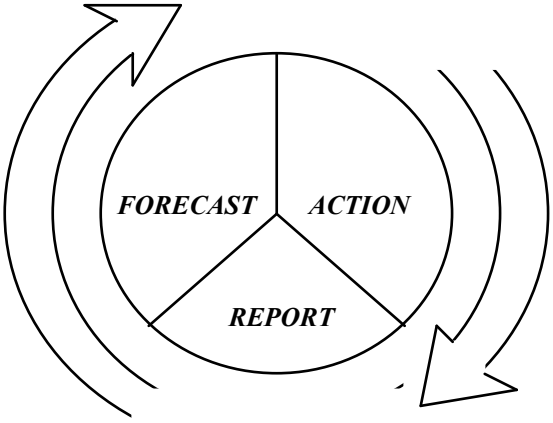

The iOptimize configuration in this manual will help you begin this cycle.

Object Management

The IBM i keeps a vast number of facts about each object. Much of this can be accessed by various IBM i commands. iOptimize selects the most relevant meta data and holds this for further analysis. By grouping, filtering, and logically presenting this information, you can more easily focus on objects which are in some way exceptional. Once you have identified an object and decided the action to take, you need the tools to effect this action.

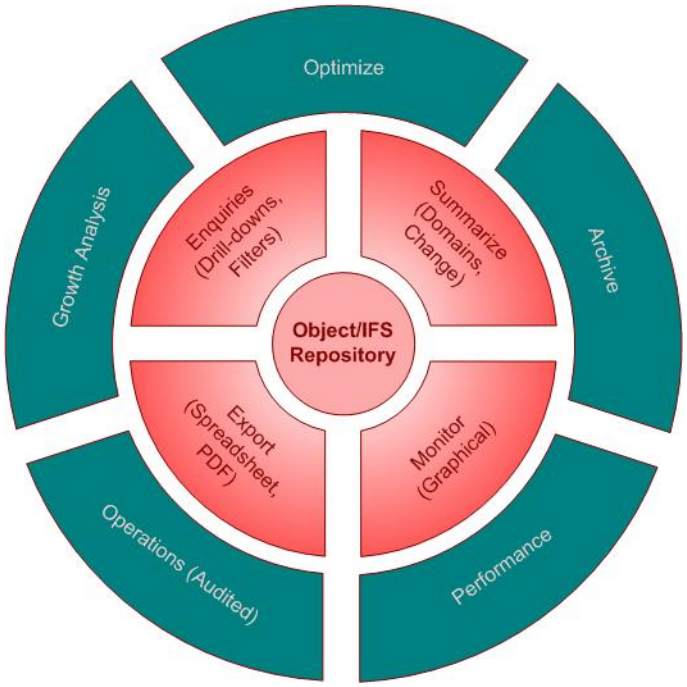

Object Enquiries

The majority of the iOptimize enquiries are backed by a report hard copy feature, have drill down capability of the sub enquiries (for example, file layout or references) and action parameters (clear, archive, or reorganize) and can be accessed independently of the iOptimize menu system. For example, you could get the file layout or the logical file keys while editing a program (SEU) or from any command line.

The Object Analysis Object Enquiries menu offers the ability to run enquiries for the following types of information (example of only some enquiries):

- libraries by name (there are many of options you can select in this screen to get additional information on a library)
- library usage summary
- files by number of records
- further file information
- file keys
- file layout
- file dependencies
- logical access path keys over a physical file
- object growth (presented as a bar chart)
- file-to-program
- program-to-file
- objects by name, size, or date last used
- files with deleted records

Methods of Increasing Available Disk Space

You can increase available disk space using some of the features available in iOptimize, including:

- automating the Reclaim Disk Storage command (an important IBM i utility)
- managing queues
- managing logs
- reclaiming temporary storage

- clearing unused spool members
- optimizing oversized objects
- clearing replaced objects

Queue Management

Queue (spool file) management is just one element of the automated space saving routines. Keeping on top of output and message queues is an important job and one which IBM has recognized with the introduction of routines in Operational Assistant.

With iOptimize, you can specify which day of the week you want a queue to be cleared. You can also view a print member, message queue or job queue directly and effect a clean up operation immediately for all or just one queue at a time. In addition, you can save some printed output to tape. You can forecast what iOptimize clears and iOptimize can record the actions that it takes.

You can change any of the default information on a queue to new specific values by typing 2 (change) next to any queue displays a window.

For example, you could change the number of days before clearing a HELD item to, say, 10 days for just this one queue. You could prevent this queue from ever being cleared by changing all of the day entries to 0 or set just the Saturday entry to 1 so that the entries are only cleared on that day.

Further processing

You can execute the clearing process for one queue by typing 10 (Action) beside the queue. Alternatively you can elect to run the Queue management routines against all of the queues right now by either running interactively (F17) or submitting to batch (F18).

NOTE

You will normally have iOptimize clear output queues using the automatic routines and not via the manual method. However this method is always available, it can override the automatic routines and has the benefit of being able to be run immediately.

Forecasting

The system is initially in forecast mode (*GENFCS) for automatic operation. iOptimize can be set to perform the clearing action on those items it identifies (*NORMAL) or, alternatively, to only clear items on its forecast (*PRCFCS). This gives you the chance to modify the forecast before entries are cleared.

Archiving Print Data

Individual prints can be archived. You can build a link between the iOptimize output queue cleanup routines and the print archive routines. In this way you can specify that certain prints (for example all INVOICES or all prints belonging to STEVE etc.) are archived just before they become eligible for deletion.

For more information on this feature, call Precisely.

User Data Optimization

You can perform optimization on system objects, including:

- physical file (see "Physical File Optimization" on page 38)
- logical files (see "Logical File Optimization" on page 39)
- program objects (see "Program Object Optimization" on page 41)
- compression (see "Compression" on page 42)

Physical File Optimization

The IBM Power Systems does not scatter individual records randomly over the available storage (DASD). Instead, it scatters blocks of records (8 pages long). A whole block is read when the disk arm reacts to a GET operation, but only the record data for the actual request is passed to the program.The difference is significant and, if managed properly, yields savings in access times and aids performance, particularly in batch.

Reorganizing your physical files is important. As well as reducing the overall size of the file by dropping deleted records, it can also improve sequential processing speeds. We can also resequence the physical file data at the same time that we remove the deleted records.

If another record is required, the current block is checked to see if the record is already there. If it is, it passes this record to the program; if not, the system reads in another block.

The following diagram illustrates the problem. It is not difficult to see that if you get a number of contiguous requests in the same block (a block could hold 10, 20, 30 records), then the number of SEEK and GET operations drops considerably.

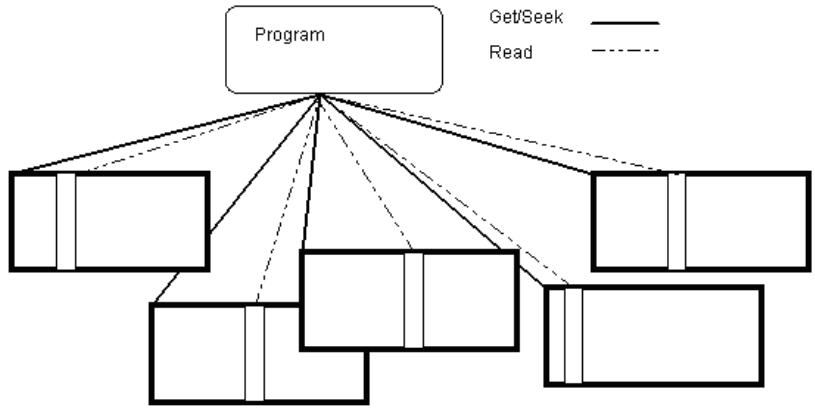

Records in blocks (8 pages) on disk

Physical file reorganization makes a considerable improvement in performance on any busy machine and is essential if you are using client/server technology. Maintaining data sequencing in this way will result in a reduction in database paging (which can be viewed on the WRKSYSSTS screen).

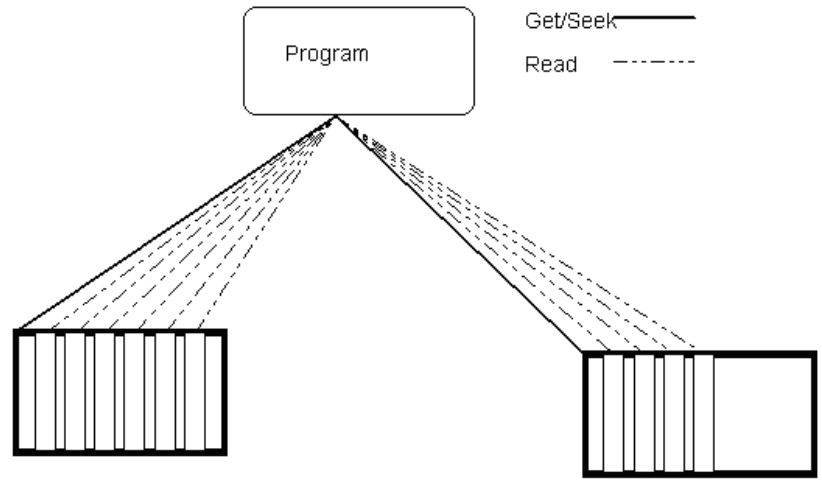

Records in blocks (8 pages) on disk

Logical File Optimization

Logical files are a feature of the IBM Power Systems and, in many ways, data manipulation via logical files is what sets the IBM Power Systems apart. However, they do affect both storage space and performance.

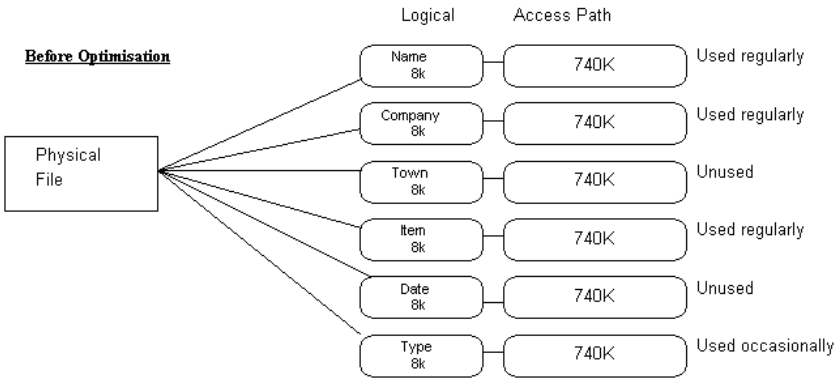

Logical files and access paths are not quite the same thing. Many logical files can use the same access path if you manage them. The IBM i provides implicit sharing of access paths, but it can only do this if the logical access paths are created in a specific sequence which can neither be dictated by IBM i or, in most cases, by you. Also, access paths can be large. Each logical file access path dictates another update which must be performed by IBM i as a physical file record is added or altered.

iOptimize will optimize all logical file usage and:

- force sharing of access paths where IBM i has been unable to effect an implicit share action
- analyze all logical files where the access path is being maintained immediately and, if eligible, reassign this to a rebuild status based on non-usage over a given, user-defined, period.

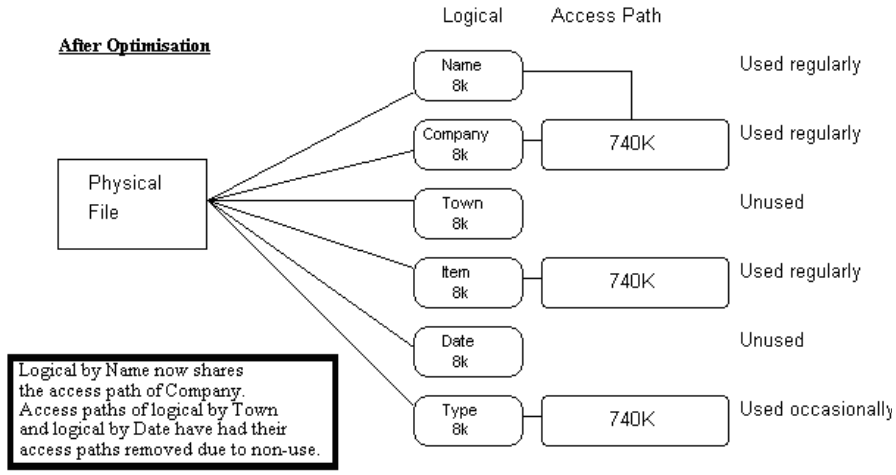

As a result, you will quickly see a useful savings in storage and an improvement in performance.

NOTE

You will never lose a logical file. However, sometimes you will share access paths and sometimes an access path will be removed (*REBLD). iOptimize automatically switches a file which it has changed to *REBLD back to *IMMED if it detects that the file was subsequently opened.

Program Object Optimization

The IBM i provides excellent program optimization routines that deliver performance improvements to both individual programs and overall machine performance.

NOTE

Before beginning optimization, validate with the program vendor that the product can be optimized. Precisely's program objects should not be optimized.
When you perform user data optimization on program objects, the IBM command CHGPGM OPTIMIZE(*YES) is used to optimize objects. This will affect the ability to debug and trace in ILE and OPM program objects. For more information, see help for the CHGPGM command.

You do not need source code to perform program optimization. When you compile a program on the IBM Power Systems, it will not optimize the program code unless specifically directed to do so. IBM also provides the Off-line Optimizer to optimize the code.

The Optimizer carries out two major tasks: replaces poor code, (for example, replacing repetitive indicator checks with a structured DO group) and resequences the physical program make-up so that the definition of each of the variables is placed near the code where it is first needed. This reduces the amount of Non Database Paging (seen on the WRKSYSSTS display).

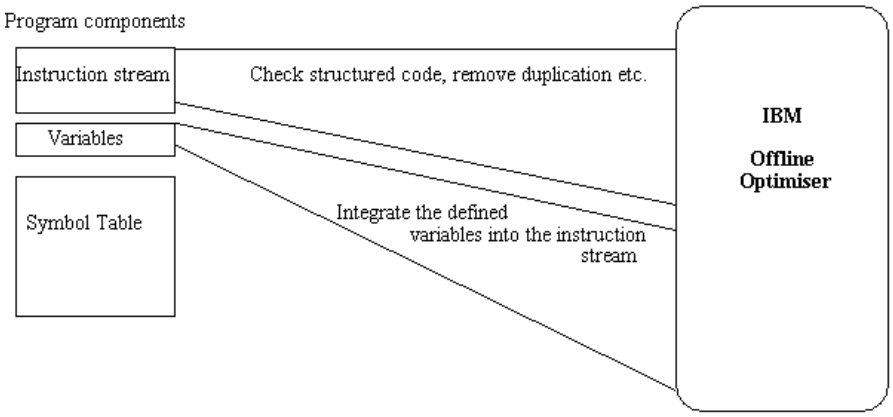

Compression

NOTE

Database files are an object type that cannot be compressed by iOptimize. The compression explained in this section will not affect the normal running of regularly used programs.

iOptimize can:

- compress programs and several other object types (files are not included) which are not being used (can be up to 50%-plus saving on a normal, uncompressed object size)
- compress observable information attached to programs (can be up to 30% saving on normal program size)
- remove the observable information completely (60% saving on normal program size), if comfortable with removing observability

Generally, the IBM IBM i compression feature, which iOptimize uses, compresses an object to approximately 50% of its original size.

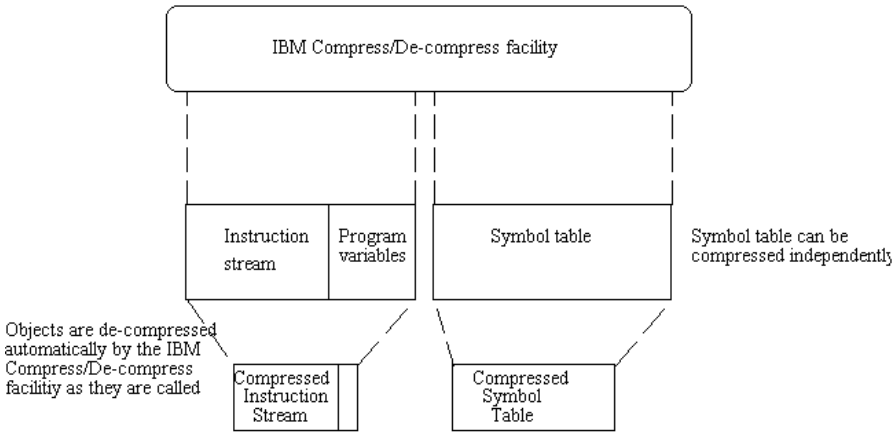

If a compressed program is called, IBM i automatically issues a message - “Operation in progress please wait” - to the user and then automatically expands (de-compress) the program, making it ready for use. This operation occurs quickly, taking seconds at most. The program can then be used as normal. iOptimize re-compresses the program after a further specified period of non-usage.

About Processing Options

There are several ways of processing actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. It is important that you are familiar with these options.

Processing Option	Description
Automatic processing	Once you have set up the action parameters, processing takes place automatically according to the rules you specified when configuring iOptimize run times. No further action is required.
Batch processing	To initiate a one-time processing run of an action (for example, your whole output queue list) in batch mode, press F18. This submits the job to the job queue that has been specified as the default for your user profile.
Interactive running - entire list	To initiate a one-time processing run of an action (for example, your entire list of libraries containing physical files to reorganize) in interactive mode, press F17. The keyboard is locked until the job is completed for all items in the list. This processing method is not recommended, due to the potentially lengthy amount of time the system may be unavailable for use.
Interactive running -single queue	To initiate a one-time processing run of selected entries from your list in interactive mode, enter 10 in the Opt field next to the item(s) you want to process. There will be a short delay while the action is processed.

NOTE

This feature is only available for programs with lists of objects or libraries, such as message queues. It is not available for generic jobs such as History logs or Problem logs.

Understanding Processing Modes

There are three main methods of processing jobs during automatic system processing.

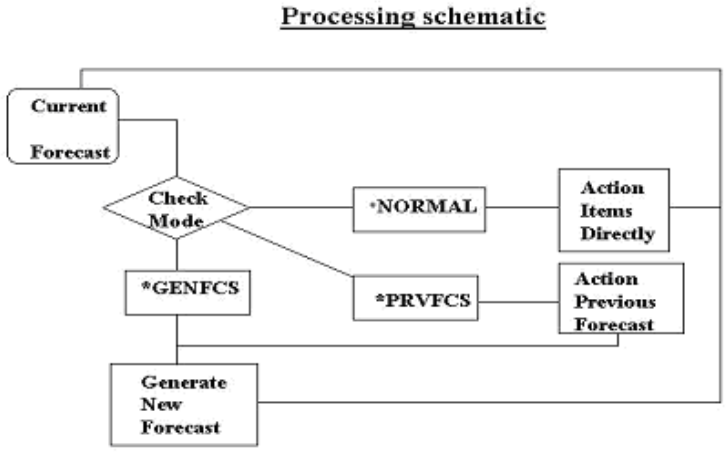

Processing Mode	Description
NORMAL (1)	When the current mode field is set to 1 (Normal), the system will action each item it finds that falls within the parameters that are set in the actual function. For example, if the Physical file maintenance function finds that a file has deleted records and it confirms that there is sufficient processing time available, it will action the file to be reorganized.

Processing Mode	Description
GENFCS (2)	Setting the mode to 2 (Generate forecast only) forces iOptimize to assess each item and forecast if and when it will action the item. Using the previous example, when iOptimize finds a file that it determines to need reorganization, it uses the parameters to estimate when the reorganization should take place. It does not, however, implement the reorganization. In this mode, each time iOptimize runs it updates the forecasts, but never performs the actions.
PRCFCS (3)	Setting the mode to 3 (Process previous forecasts), results in the system actually performing the actions. However, it will only do so based on the forecast generated in the previous run, not on the current object conditions. Again, using our example of the file reorganization, and using a simple single-day forecast, a forecast is generated on the first run and the system might schedule a file for reorganization on the next run. It does not perform any reorganization during the first run. On the following run, iOptimize generates a new forecast, but then reads the old forecast and performs the actions it predicted on that previous run, regardless of the current object conditions. The advantage of this processing method is that you have the opportunity to review the generated forecast and make modifications to it before the following run, during which the actions are implemented (so you can prevent an undesirable forecasted reorganization from taking place).

Setting Up iOptimize to Run Automatically

You can set up iOptimize to perform actions automatically at prescribed times. Although you may want to wait until you have performed some of the following optimization and management task to configure automatic run times, it is important to understand the concepts at this point.

You can configure iOptimize to process the various internal (IBM i) and external (user) elements on your system to suit your needs. Processing options include daily, on selected days only, and according to time limits (for example, only clear save status prints after 30 days).

Archiving/ Nearline

About iOptimize Archival Management

When a library or object has been identified as redundant, you can clear or delete the object or library. However, sometimes the information needs to be taken off-line for occasional retrieval.

iOptimize provides a feature to archive and store off-line almost any object. It keeps a comprehensive catalog of all archiving activity that can be scanned and used to initiate an instant recall of any archived item.

Thus, iOptimize provides NEARLINE management (off-line but available for immediate recovery) for your IBM Power Systems.

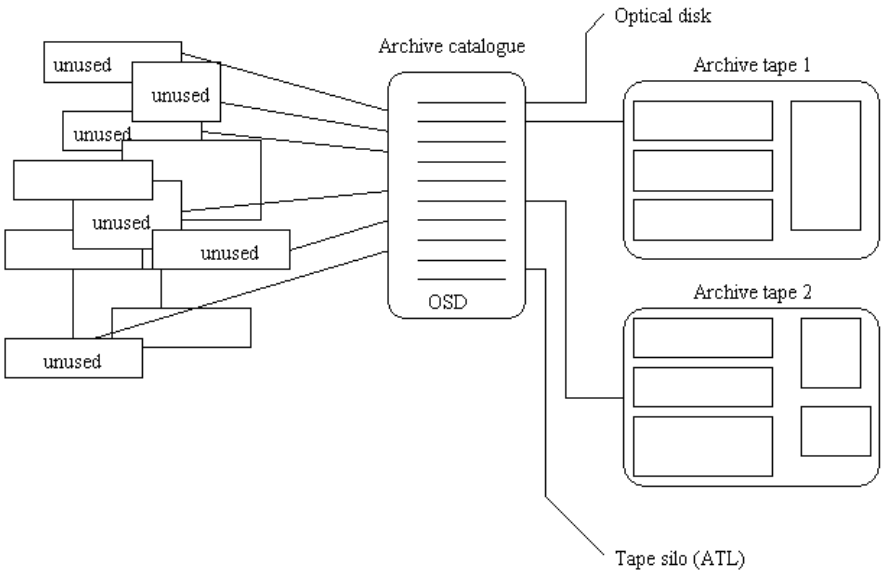

Integrated File System Management

The Integrated File System (IFS) is the term generically applied to the various filing systems now supported on the IBM Power Systems processors. The IBM i can now support filing systems from the OEM environments, like the UNIX filing systems, which can be managed by IBM i alongside the traditional QSYS filing systems. The earliest of these alternate filing systems was the QDOC library. This is now the QDLS filing system within the IFS.

Object analysis for the IFS is configured and managed using iOptimize portal application.

Managing Jobs and Tracking Performance

As iOptimize manages itself, it picks up information about the jobs that are running and general machine performance.

iOptimize can track all jobs on the system and provide information such as failed job summary, grouping of CPU usage for particular users, and analysis of sub-system usage. It presents this information in a number of readily available and easily understood panels and maintains the links to vital Goblets and Spool entries related to any job. In addition to

performance information on jobs, iOptimize also tracks the success and failure of jobs.

About Task Manager

Task Manager controls automatic processing of jobs on your IBM Power Systems. In addition to handling iOptimize tasks, it can also manage batch jobs. Task Manager can track all jobs, whether they were submitted by Task Manager or not. Tracking and logging these jobs provides a unique method of accessing the history of a job, and simplifies tasks such as finding the joblog for a particular failed job. See "Task Manager" on page 66, for more information.

Configuration

5

Configuration tasks for Assure iOptimize are performed from the iOptimize Configuration Menu. It is important to note that all security configuration for iOptimize is completed in the portal application on the Administration page (see the *Assure iOptimize Portal Application User Guide*).

iOptimize Main Menu —> Option 4 (Configuration)

Configuration options on this menu are as follows:

Option	Description
1 (iOptimize Controls)	• 1 - "Working with Exit Points" on page 49 • 3 - "Working with Audit Monitors: Real-Time Cross-Reference" on page 51 • 4 - "Working with Calendars" on page 55 • 5 - "Working with Schedules" on page 58 • 6 - "Working with iOptimize Subsystem Jobs" on page 62
2 (License Maintenance)	• 1 - "Displaying System Information" on page 63 • 2 - "Update License Key via AutoValidate" on page 64 • 3 - "Build AutoValidate Up-Load File" on page 64 • 4 - "Apply AutoValidate License Key" on page 65
Related options	These options include the following: – 21 - "Object Analysis" on page 73 – 22 - "Optimization" on page 156 – 23 - "Operations" on page 114 – 25 - Displays the iOptimize Accounting Code Analysis Menu.

Working with Exit Points

iOptimize Configuration Menu—>Option 1 (iOptimize Controls)—>Option 1 (Work with Exit Points)

The Work with Exit Points screen shows the available Exit Points in iOptimize and provides the option to configure whether exit point programs are called (by default, exit point programs are not called unless they are enabled). The following information shows when exit points are called:

Exit Program	Exit Point
COU020CL	Called during execution of the End iOptimize Process (ENDID) command, the program is called twice BEFORE and AFTER the requested process is ended.
COU030CL	Called during execution of the Start iOptimize Process (STRID) command, the program is called twice BEFORE and AFTER the requested process is started.
COU040CL	Called within the SYSSTATS scheduled job for each job that is found at status MSGW, the frequency of this job cannot be changed. Monitors can also be used trigger actions related to user jobs at MSGW status.
COU050CL	Called when a new library is detected and being added to Library Registration
IDUGRP1CL	Called within Library Registration to enable a value to be provided for the User defined Group field 1 when initially loading libraries and when new libraries are created and loaded into Library Registration.
IDUGRP2CL	Called within Library Registration to enable a value for the User defined Group field 2 when initially loading libraries and when new libraries are created and loaded into Library Registration.
IDUGRP3CL	Called within Library Registration to enable a value for the User defined Group field 3 when initially loading libraries and when new libraries are created and loaded into Library Registration.
IDUGRP4CL	Called within Library Registration to enable a value for the User defined Group field 4 when initially loading libraries and when new libraries are created and loaded into Library Registration.

Exit Program	Exit Point
IDUGRP5CL	Called within Library Registration to enable a value for the User defined Group field 5 when initially loading libraries and when new libraries are created and loaded into Library Registration.
IDU080CL	Called when any iOptimize subsystem monitor job ends abnormally.

NOTE

Exit point templates can be found in: *installation-library*/IOEXTCLE.

The screen contains the following:

Column	Description
Program	The program that will be called.
Library	The library for the exit point.
Status	The current status of the exit point. • *INACTIVE–The program will not be called. This is the default. • *ACTIVE–The program will be called at the appropriate point.
Description	A description of where the exit point is called from.
Obj	If the object currently exists in the specified library or in the current library list. Note that L may be different at run-time depending on where the exit point is called.

The following actions are available:

Option	Description
2 (Change)	Modify the library associated with or the status of the exit point (see "Modifying Exit Points" on page 51).
5 (Display)	Display the exit point attributes (see "Displaying Exit Point Attributes" on page 51).

Modifying Exit Points

You can modify the library associated with the exit point and the status of the exit point.

1. In the Work with Exit Points screen, enter Option 2 (Change) next to the Exit Point you want to change.
2. Modify the parameters, as required.

Displaying Exit Point Attributes

In the Work with Exit Points screen, enter Option 5 (Display) to display detailed information about the selected exit point.

Working with Audit Monitors: Real-Time Cross-Reference

On the Work with Audit Monitors screen, you can change real-time cross-reference monitor parameters, display current parameters, view real-time cross-reference monitor logging messages, display journal attributes and objects, and Start or End the real-time cross-referencing.

NOTE

The QACGJRN journal can also be configured from this screen. For more information, see *Assure iOptimize Job Accounting User Guide*.

iOptimize Configuration Menu—> Option 1 (iOptimize Controls)—> Option 3 (Work with Audit Monitors)

This section includes the following information:

- "About Real-Time Cross-Referencing" on page 52 provides an overview of real-time cross-referencing.
- "Changing Real-Time Cross-Reference Monitor Parameters" on page 52 provides instructions for specifying real-time cross-reference parameters.
- "Adding to Cross-Reference" on page 53 provides instructions for manually adding a library, file, or spool file for cross-reference processing.
- "Excluding Libraries from Cross-Reference" on page 54 provides information for excluding libraries from cross-reference processing.

About Real-Time Cross-Referencing

To provide online access to detailed information about IBM Power Systems status, iOptimize maintains cross-reference files automatically, using a real-time cross-reference feature. To ensure that all files on your system get cross-referenced, a *FILE cross-reference, FILEBUILD, is scheduled to run every day at 4:00:30 a.m. by default. A FILEBUILD task can also be run or scheduled through task manager. The file can be rebuilt either in its entirety or on a selective basis at intervals you specify. For more information, see "Cross-Reference Jobs in Task Manager" on page 67.

This automatic real-time cross-reference feature maintains iOptimize's database of IBM Power Systems object information on a nearly real-time basis. iOptimize detects object changes by reading journal entries in the security journal, QAUDJRN. New spool files or objects, spool file or object changes, or spool file or object deletion results in an immediate update of the iOptimize database.

For more information about real-time cross-reference processing, see "Real-Time Cross-Reference" on page 21.

Note the following:

- For spool file data to be cross-referenced, QAUDLVL must include the parameter value *SPLFDTA, which is automatically added when iOptimize is started, provided the value of IDDYNSPL in IDCTLMP is set to *YES (default). Once the subsystem is started, configure and start the Audit Monitor.
- If you do not want the spool database IDSPLTP to be updated dynamically, set IDDYNSPL in IDCTLMP to *NO. *YES is the default.

For more information about cross referencing spool file data, see "Setting System Values" on page 22.

Changing Real-Time Cross-Reference Monitor Parameters

1. To change real-time cross-reference monitor parameters, from Work with Audit Monitors, enter option 2 (Change) and press Enter.
2. Press Page Down to see more options.
3. Configure the following parameters and press Enter:

Parameter	Value	Description
Apply Priority	*LOW *MEDIUM *HIGH	Priority level for applying object changes.
Receiver delay	*NEXTENT Seconds	*NEXTENT - Next journal entry will be sent. Seconds - Delay time, in seconds, in sending the journal entry to the receiver.
Start up Sequence	1 - 999	The start up sequence number.
Sequence control	*NEXT *SEQN *TIME	The sequence number to be selected. *NEXT - Next sequence number *SEQN - A specific sequence number *TIME - The sequence number of a specific date/time.
Apply from Sequence	Number	Apply object changes from a specific journal sequence number.
Apply from Receiver	Name	Apply object changes from a specific journal receiver.
Apply from Date	Date	Apply object changes from a specific date.
Apply from Time	Time	Apply object changes from a specific time.

Adding to Cross-Reference

System Command —> IDXREF

With iOptimize's real-time cross-referencing, new libraries, files, and spool files are automatically detected and cross-referenced. You can also manually add one into the cross-reference list using the IDXREF command. Manually adding to the cross-reference list is only required when real-time cross-referencing is not being used or when a file has been manually excluded from cross-referencing.

Running this command selects the list of libraries to be processed based on the selections made. The selected libraries are processed by the IDXRFPRC job running in the iOptimize subsystem. Based on the system set-up, the IDXRFPRC job submits the cross-reference jobs that processes each selected library (by default these jobs run in the iOptimize subsystem). Progress can be monitored from the iOptimize portal application (see the *Assure iOptimize Portal Application User Guide* for more information).

1. Enter the following command to add a library, file, or spool file for cross-referencing:

   ```
   IDXREF
   ```

   The Request Cross-Reference screen is displayed.

2. In the Scope field, enter one of the following to add to the cross-reference list:

   - *EXCL - adds libraries defined in IDAUXMP (libraries excluded from real-time cross referencing).
   - *FILE – updates the cross-reference file level attributes (for example, number of records)
   - *IFS – updates the IFS information
   - *LIB – updates the database for an entire library.
   - *SPL–updates spool information.

3. If you are including a library, in the Library field, enter one of the following:

   - Name-the name of the library.
   - *ALL-adds all libraries to the cross-reference.

4. Enter the following:

Parameter	Description
ASP Name	The name of the ASP on which the library resides. Specify *ALLAVL for all available ASPs.
Objects to Process	The objects to process: • *CHG processes only changed objects. • *ALL process all objects (note this is a full cross-reference and can take considerably longer to process since it is processing all objects).

Excluding Libraries from Cross-Reference

System Command —> IDEXCLIB

You can exclude libraries from real-time cross-reference processing using the IDEXCLIB command. For instructions, see "Excluding Libraries from Cross-Referencing and Optimization" on page 22.

Working with Calendars

iOptimize Configuration Menu—>Option 1 (iOptimize Controls)—>Option 4 (Work with Calendars)

The calendar and schedule work together to support event processing in iOptimize. The calendar provides the dates for the schedule. In the schedule, you configure the days on which the schedule operates and the user to which messages are sent to alert users of certain event outcomes. See "Working with Schedules" on page 58 for more information on schedules.

The screen contains the following:

Column	Description
Name	The name of the calendar.
Description	A description of the calendar.

The following actions are available:

Option	Description
1 Create	Adds a new calendar (see "Creating a New Calendar" on page 55).
2 Change	Modifies an existing calendar (see "Modifying a Calendar Description" on page 56).
4 Delete	Deletes an existing calendar (see "Deleting a Calendar" on page 57).
5 Display	Displays attributes for the selected calendar (see "Displaying Calendar Attributes" on page 56).
7 Dates	Sets the holidays and working days (see "Setting Holidays and Working Days" on page 56).

Creating a New Calendar

New calendars are created for the current year and don't have any holidays or working days configured (unless the calendar is copied from an existing calendar). Once a calendar is created, you need to configure the holidays and working days for the calendar (see "Setting Holidays and Working Days" on page 56).

1. In the Work with Calendars screen, enter Option 1 (Create) on the first line to create a calendar, then enter the calendar name.
2. Enter the following:

Field	Description
Description	A description of the calendar.
Copy Existing	The name of an existing calendar on which to base this new calendar. This copies all calendar dates, including non-working time and holidays onto the new calendar.

3. Press Enter to create the new calendar.

Modifying a Calendar Description

1. In the Work with Calendars screen, enter Option 2 (Change) next to the calendar you want to modify.
2. Modify the description as required.

Displaying Calendar Attributes

In the Work with Calendars screen, enter Option 5 (Display) next to the calendar for which you want to display attributes (see "Creating a New Calendar" on page 55 for parameter information).

Setting Holidays and Working Days

You set holidays and working days in the Work with Calendar Dates screen.

1. From the Work with Calendar screen, enter Option 7 (Dates) next to the calendar for which you want to set holidays and working days.

   The Work with Calendar Dates screen displays.

The screen contains the following:

Column	Description
Position to Date	Does the following: • Enter a date to move to a calendar date. • If a date is entered that is not in the current calendar, a calendar is automatically generated for the year entered. The new year is automatically added to any schedules that use this calendar. **Note:** New calendar years are also generated automatically as they occur.
Date	The calendar date.

Column	Description
Wrk	If it is a working day. • Y–It is a working day. • N–It is not a working day.
Exl	If the date is excluded. Excluded dates cannot be used in a schedule. • Y–The date is excluded. • N–The date is included.
UPH	If the date is a User Public Holiday. A User Public Holiday is a holiday that falls on a different date each year. • If blank, the date is not a user public holiday. • If set to Y, the date is a user public holiday.
SPH	If the date is a Static Public Holiday. A Static Public Holiday is a holiday that falls on the same date each year. • If blank, the date is not a static public holiday. • If set to Y, the date is a static public holiday.
Day	The day of the week.
Week	The week number for the date.

2. Navigate to a date.
   - Roll up or down to the calendar date.
   - Use the Position to Date field to locate the date.
3. Enter any of the following options:

Option	Description
1 (Include/Exclude)	Excludes the date. Excluded dates cannot be used in a schedule. Use the option again to reverse the status.
2 (UPH)	Sets the date as a holiday that falls on a different date each year. Use the option again to reverse the status.
3 (SPH)	Sets the date as a holiday that falls on the same date each year. Use the option again to reverse the status.

Deleting a Calendar

You delete a calendar from the Work with Calendars screen.

1. On the Work with Calendars screen, enter Option 4 (Delete) next to a calendar to delete it.
2. Press Enter to confirm the deletion.

Working with Schedules

iOptimize Configuration Menu—>Option 1 (iOptimize Controls)—>Option 5 (Work with Schedules)

The calendar and schedule work together to support event processing in iOptimize. The calendar provides the dates for the schedule and must be configured before configuring the schedule (see "Working with Calendars" on page 55 for more information). In the schedule, you configure the days on which the schedule operates and the user to which messages are sent to alert users of certain event outcomes.

The screen contains the following:

Column	Description
Name	The name of the schedule.
Description	A description of the schedule.

You can do the following in this screen:

Option	Description
1 (Create)	Create a new schedule (see "Creating a New Schedule" on page 58).
2 (Change)	Modify an existing schedule (see "Modifying a Schedule" on page 60).
4 (Delete)	Deleting a schedule (see "Deleting a Schedule" on page 62).
5 (Display)	Displaying schedule attributes (see "Displaying Schedule Attributes" on page 60).
7 (Dates)	Set the user contacts for a schedule by date (see "Setting User Contacts by Date" on page 60).

Creating a New Schedule

Before creating a new schedule, make sure that you have created the calendar that will be used by the schedule (see "Creating a New Calendar" on page 55) and have configured the user contact information.

1. In the Work with Schedules screen, enter Option 1 (Create) on the first line to create a schedule, then enter the schedule name.
2. Enter the following:

Field	Description
Description	A description of the schedule.
Calendar	The calendar associated with the schedule. The calendar provides the dates for the schedule. This field is required.
Build Parameters	
Days	The default working and non-working days for the working week. (The working week begins with Monday and ends with Sunday.) Y designates a working day; N designates a non-working day. The default is YYYYYNN.
Working	The days on which the schedule operates. • Y–The schedule operates on working days. • N–The schedule operates on non-working days. • B–The schedule operates on both working and non-working days. If B is selected, Days is automatically set to YYYYYYY.
User Holidays	Determines if the schedule operates on user holidays (defined in "Setting Holidays and Working Days" on page 56). • Y–The schedule operates on user holidays. • N–The schedule does not operate on user holidays.
Static Holidays	Determines if the schedule operates on static holidays (defined in the associated calendar, see "Setting Holidays and Working Days" on page 56). • Y–The schedule operates on static holidays. • N–The schedule does not operate on static holidays.
Every nth day	The schedule operates every nth day, starting with the day specified in from day.
Every nth week	The schedule operates every nth week, starting with the week specified in from week.
Week in month	The schedule operates on a given week number within a month.
User	

Field	Description
User Name	The name of the user. **Note:** You must add at least one user with configured contact information in order to create a new schedule. (See the *Assure iOptimize Portal Application User Guide* for information on setting the contact information.)
Time	The time range during which the user can be contacted. Enter the time using a 24 hour clock (for example, enter a range of 4:00 pm to 2am as 16:00 to 02:00). If no time is entered, the user can be contacted at all times.
Device Type	The means by which the user is contacted. • *EMAIL–Device capable of receiving email.

3. Press Enter to create the new schedule.

Modifying a Schedule

1. In the Work with Schedules screen, enter Option 2 (Change) next to the schedule you want to modify.
2. Modify the schedule as required (see "Creating a New Schedule" on page 58 for field definitions).

Displaying Schedule Attributes

In the Work with Schedules screen, enter Option 5 (Display) next to the schedule for which you want to display attributes (see "Creating a New Schedule" on page 58 for parameter information).

Setting User Contacts by Date

The user contact is the user to which messages are sent when specified events occur (see the *Assure iOptimize Portal Application User Guide* to configure the events). You set user contacts by date for a schedule in the Work with Schedule Dates screen.

1. From the Work with Schedule screen, enter Option 7 (Dates) next to the schedule for which you want to set holidays and working days.

   The Work with Schedule Dates screen displays.

2. Navigate to a date.
   - Roll up or down to the calendar date.
   - Use the Position to Date field to locate the date.

3. Modify any of the following:

Column	Description
User	The user contact to which messages are sent.
From	The start time in the time range during which the user can be contacted. Enter the time using a 24 hour clock (for example, enter 4:00 pm as 16:00). If no time is entered, the user can be contacted at all times.
To	The end time in the time range during which the user can be contacted. Enter the time using a 24 hour clock (for example, enter 2am as 02:00). If no time is entered, the user can be contacted at all times.
Device	The means by which the user is contacted. • *EMAIL–Device capable of receiving email.

4. Optionally, enter Option 5 (Display Contact Details) to display the contact information for the selected user. (See the *Assure iOptimize Portal Application User Guide* for information on setting the contact information.)

   The screen contains the following:

Field	Description
User	The user name.
E-mail	The email address for the user.
Mobile	The mobile phone number for the user.

Testing a Schedule

Once a schedule is configured, you can verify that the schedule is running on a specific date.

1. On a command line, enter:

   ```
   IDRTVRSC
   ```

2. Press F4 to prompt.

   The iOptimize Retrieve Resource (IDRTVRSC) screen is displayed.

3. Enter the following:

Field	Description
Schedule	Enter the schedule to verify.
Date (job format)	Enter the date in the schedule.
Time	Optionally, enter a time on the specified date.

4. Press Enter.

   One of the following occurs:

   - A message confirming who is scheduled to work on that date (and time, if specified) is displayed.
   - If the schedule is not running on the specified date, a notification message is displayed.

Deleting a Schedule

You delete a Schedule from the Work with Schedule screen.

1. On the Work with Schedule screen, enter Option 4 (Delete) next to a schedule to delete it.
2. Press Enter to confirm the deletion.

Working with iOptimize Subsystem Jobs

iOptimize Configuration Menu—>Option 1 (iOptimize Controls)—>Option 6 (iOptimize Subsystem Jobs)

Option 6 (iOptimize Subsystem Jobs) opens the Work with Active Jobs screen. From this screen you can Change, Hold, End, Work with, Release, Display messages, Work with spooled files, or Disconnect active jobs.

iOptimize contains the following subsystem jobs:

Subsystem Job	Description
IDALTMON	Alert monitor–Monitors new requests from the Event Monitor and process actions related to Events.
IDASPMON	ASP monitor–Monitors disk utilization across all ASP's

Subsystem Job	Description
IDDIRCTL	Main iOptimize–Schedules tasks and occasionally purges records and reorganizes files.
IDDIRPRC	iOptimize Process–Controls yearly, daily, and hourly processing.
IDDYNACG	Accounting Monitor–Reads QACGJRN and populates files IDACGMP and IDACGTP with detail and summary accounting information.
IDDYNMON	QAUDJRN Monitor–Reads QAUDJRN and populates files IDDXDTP with changed object information.
IDDYNPRC	Real-time updates–Processes IDDXDTP and updates iOptimize cross-reference information. During periods of significant change to objects this will show higher usage.
IDETSCTL	Email Transaction Server–The controlling job for the ETS JVM.
IDETSJVM	Email Transaction–This is the Email Transaction Server which sends the email alerts.
IDEVTMON	Event Monitor–Checks events and adds to Event Transaction file when they are triggered.
IDJOBLST	Job list–Maintains a list of active jobs.
IDJOBMON	Job collection–Updates statistics on active jobs.
IDXRFPRC	Cross-reference processor–cross-reference processing requested with IDXREF is processed here. This may show higher usage when cross-references are requested.

Displaying System Information

iOptimize Configuration Menu—> Option 2 (License Maintenance)—>Option 1 (Display System Information)

To display the system information needed to obtain a license, take option 1 (Display System Information) on the License Maintenance menu, then press Enter.

The following information is included:

Parameter	Description
System serial number	The serial number of the system.
Model	The model number of the system.
Processor feature code	The processor feature code of the system.
Processor group	The processor group number of the system.
LPAR number	The partition number of the system.
Product versions	The version number of the installed Precisely products. License Manager and iOptimize product versions will always be displayed. Other installed products that use License manager will also be displayed.

Update License Key via AutoValidate

iOptimize Configuration Menu—> Option 2 (License Maintenance)—>Option 2 (Update License Key via AutoValidate)

Read the screen text for a description of how to use this screen. Complete instructions for generating and applying a license key for iOptimize are included in the *Using License Manager* book.

Build AutoValidate Up-Load File

iOptimize Configuration Menu—> Option 2 (License Maintenance)—>Option 3 (Build AutoValidate Up-Load File)

Instructions for building the AutoValidate upload file for iOptimize licensing are included in the *Using License Manager* book. This menu option can be used to generate the License Request Package that is saved in the IFS, which must then be uploaded to the Assure Unified Interface.. The name of the file generated is displayed in the Update License Keys screen (see "Update License Key via AutoValidate" on page 64). It contains the system information (see "Displaying System Information" on page 63) that is used to generate the license key.

Apply AutoValidate License Key

iOptimize Configuration Menu—> Option 2 (License Maintenance)—>Option 4 (Apply AutoValidate License Key)

Instructions for applying the license key for iOptimize are included in the *Using License Manager* book. This menu option can be used to apply the key after the License Key Package has been saved to the IFS.

Task Manager

6

Task Manager controls automatic IBM Power Systems job processing. It processes tasks for Assure iOptimize and, if you wish, manages batch jobs on your behalf. Task Manager can track all jobs, whether or not the job was submitted by Task Manager. Tracking and logging these jobs provides a unique ability to access the history of a job, and simplifies tasks such as finding the joblog for a failed job.

When active, Task Manager generates an hourly job called SYSSTATS, which gathers data about your system status. This information is used by Job Analysis to provide information about processor usage including which jobs are using the most processor time. It also tracks the submission of housekeeping tasks to the appropriate batch job queue.

You can use Task Manager to submit jobs. Simply define the job, including when and how often to run it. You can configure a job to be run after another task has been completed, or even require another task to be completed successfully before running the job.

NOTE

Task Manager is accessed from *any* iOptimize Menu screen by pressing the F7 (Task Manager) function key.

Task Manager Column Descriptions

Task Manager is accessed from *any* iOptimize Menu screen by pressing F7.

<Any iOptimize Menu> —> F7 (Task Manager).

The current Task Manager status is displayed at the top of the screen (Active or Inactive). The following is an explanation of each column of data displayed on the screen:

Column	Description
Opt	Used to enter an option to perform.
St	Displays whether the task is in held (HLD) Status.
Date	The date on which the task is scheduled to run.
Time	The time at which the task is scheduled to begin.
Task name	The name of the task.
Seq	Used along with the Task name to ensure that each task is unique.
Description	Describes the task's function.

Cross-Reference Jobs in Task Manager

Task manager contains the following jobs that support real-time cross-reference processing:

Task	Description
IFSBUILD	IFS cross-reference build. Object analysis for the IFS is configured and managed using iOptimize portal application. See the *Assure iOptimize Portal Application User Guide* for more information on iOptimize analysis for the IFS.
EXCLBUILD	Cross-reference of excluded libraries. For more information, see "Excluding Libraries from Cross-Reference" on page 54.
FILEBUILD	Complete system cross-reference file build. This file can be rebuilt, either in its entirety or on a selective basis. FILEBUILD should be run shortly after the EXCLBUILD. For more information, see "About Real-Time Cross-Referencing" on page 52.

By default the IFSBUILD and EXCLBUILD jobs are scheduled to run at 4:00 a.m. each day and the FILEBUILD job is scheduled to run at 4:00:30 a.m. These times can be rescheduled at regular intervals specified by selecting option 2 from Task Manager.

If you choose to change these times from their defaults, the FILEBUILD job should be scheduled to run shortly after the EXCLBUILD job to allow time for the EXCLBUILD job to build its list of libraries. The FILEBUILD job will then select any remaining libraries for processing.

NOTE

These jobs are dependent on the settings for machine processing times which controls when cross-reference processor jobs can run. For more information, see "Setting Processing Windows" on page 198.

For information about running or rescheduling these jobs, see "Changing Job Parameters" on page 70.

Adding User Jobs

Task Manager—> F6 (Add Task)

1. To add a job to the schedule jobs list, from the Task Manager screen, press F6 (Add Task) to display a job configuration screen.
2. Use the following information to complete each field:

Field	Description
Task name	The name of the job to schedule (this name should be easily recognizable and describe the task it performs).
Sequence	Used in conjunction with the task name to identify a job. You can schedule multiple jobs that have the same name as long as their sequence numbers are different. (Range: 001 - 999)
Description	A description of the job.
Next date	The date on which the job will run next. The date format is the same as the system date format (for example, in the UK: *DD/MM/YY,* in the USA: *MM/DD/YY*).
Next time	The time of day at which the job will run next (in 24-hour format).
No. of times	The number of times that the job is to be run. An entry of 99999 means it will run indefinitely.
Frequency	The frequency at which the job is to be run, specified in minutes, hours, days, weeks, months, years, or a combination thereof. Entering numbers in more than one frequency field results in a cumulative frequency (for example, 1 week, 3 days would result in the job running every 1 week + 3 days (or 10 days)).

Field	Description
Dependent on job	If the task is to run only after another task is completed successfully, specify the Task Name and Sequence Number of the other task. Note that a task may be configured to be "dependent on" one job and "wait for" a different job.
Wait for job	If the task is to run only after another task is completed (not necessarily successfully), specify the Task Name and Sequence Number of the other task. Note that a task may be configured to be "dependent on" one job and "wait for" a different job.
Day Mask	Enter 1 to indicate the days of the week the job can run. Enter 0 to indicate the days the job must not run.
Holidays	Enter 1 if the job is to run regardless of company shutdown or holidays. Enter 0 if the job must not run on a company holiday.
Run if overdue when Job Scheduler starts	Enter 1 to run the job immediately if it is found to be overdue when Job Scheduler becomes ACTIVE after a period of INACTIVITY. Enter 0 to reschedule the overdue job to its next regular run time.

NOTE

The Dependent on job and Wait for job parameters can be combined for a single job to ensure that a previously submitted job has both completed and completed successfully before releasing the current job (for example, Invoice selection/generation to have completed and completed successfully before releasing the invoice print/update).

3. After entering information to add a task, press F8 (Update) to update the configuration.

   The Submit Job (SBMJOB) screen displays showing the parameters for submitting a job.

4. Enter the job details as you would typically do when using this command and press Enter to save the job details and return to the Task Manager screen.

NOTE

For more details about the SBMJOB command, please refer to the appropriate IBM manual.

Changing Job Parameters

Task Manager—> Option 2 (Change)

1. To change job parameters, enter Option 2 (Change) in the Opt field next to the job you want to change. You can select multiple jobs at one time; when you exit the first job configuration screen Task Manager displays the next selected job configuration screen.
2. When you are finished changing parameters, press F8 (Update).
3. Press F20 to toggle between displaying the job list in chronological or alphabetical order.

Holding Scheduled Jobs

Task Manager—> Option 3 (Hold)

To put a job on hold, enter Option 3 (Hold) in the Opt field next to the job name.

Deleting User Jobs

Task Manager—> Option 4 (Delete)

1. To delete a job from the schedule list, enter Option 4 (Delete) in the Opt field next to the job to be deleted. A confirmation window is displayed.
2. Press F8 to confirm the delete, or press F12 to return to the Task Manager screen without deleting the job.

Viewing Job History

Task Manager—> Option 5 (History)

1. To obtain the history of a particular job, enter Option 5 (History) in the Opt field next to the job. The history for the selected job displays.

NOTE

Option 5 (History) is not available with the jobs SYSSTATS and TRACKJOBS.

2. Select any of the following options to obtain more historical information: Option 1 (More Information), Option 5 (Display job), Option 6 (Spool files), and Option 8 (Joblog).
3. To display a history of all submitted jobs, on the Task Manager screen, press F15 (Jobs Submitted).

Releasing a Held Job

Task Manager—> Option 6 (Release)

To release a job that has been put on hold (status HLD), enter Option 6 (Release) in the Opt field next to the job.

Displaying Submitted Job Enquiries

Task Manager—> F15 (Submitted Jobs)

Press F15 (Jobs submitted) to display the Submitted Jobs Enquiry screen.

Displaying Scheduler Jobs

Task Manager—> F24 (More Keys)—> F13 (Display Scheduler Job)

Press F13 (Display scheduler job) to display the Task Manager job (DSPJOB).

Sorting Task List

Task Manager—> F24 (More Keys)—> F20 (Sort by task/Sort by date)

Press F20 (Sort by task/Sort by date) to toggle between the two sorting methods; chronological or alphabetical order.

Maintaining Holidays

Task Manager—> F24 (More Keys) —> F22 (Forthcoming Holidays)

You maintain your holiday list on the Forthcoming holidays screen.

Enter the holiday dates, using the same format as your system date. Press F8 (Update) to save your changes.

Object Analysis

7

Object Analysis tasks are performed from the iOptimize Object Analysis Menu.

iOptimize Main Menu —> Option 1 (Object Analysis)

Options on this menu are as follows:

- "Object Enquiries" on page 73
- "File Enquiries" on page 84
- "Domains, Snapshots and Growth" on page 94

NOTE

All metrics are available only in the iOptimize portal application (see the *Assure iOptimize Portal Application User Guide* for more information).

Object Enquiries

Object Analysis—>Option 1 (Object Enquiries)

From the Object Enquiries menu, you can access the following types of enquiries:

- "Library Enquiries" on page 73
- "Object Enquiries" on page 77
- "ASP Enquiry" on page 79

Library Enquiries

You can run the following library enquiries:

- Libraries by Name–Displays a list of all the libraries currently on your system by name (see "Libraries by Name" on page 74).
- Libraries by Size–Displays a list of all the libraries currently on your system by name (see "Libraries by Size" on page 74).
- Library Usage Analysis–Scans the system for objects not used within a specified time frame (see "Library Usage Analysis" on page 75).
- Library Dates Analysis–Displays significant dates relating to objects in a library but presented at a library level (see "Library Dates Analysis" on page 75).

Libraries by Name

Object Analysis—>Option 1 (Object Enquiries)—>Option 1 (Libraries By Name)

This inquiry is often the starting point in identifying what is occupying the DASD. It displays a list of all the libraries currently on your system, by name. The inquiry shows information on the total size (the library object and all of the objects within the library) as well as the percentage of the system that the library takes up.

The following actions are available:

Option	Description
F1 (Help)	Displays a list of the currently enabled user options.
F14 (Reverse)	Toggles the order in which the library list is displayed.
F17 (Filter)	Filters the displayed data (see "Filter the Displayed Data" on page 74).
F18 (% of ASP)	Toggles between displaying % of ASP or % of system space used.

Filter the Displayed Data

F17 (Filter) enables you to filter the displayed data, reducing the time it takes to find particular information.

Libraries by Size

Object Analysis—>Option 1 (Object Enquiries)—>Option 2 (Libraries By Size)

Libraries by size displays a list of all the libraries currently on your system, by size. The inquiry shows information on the total size (the library object

and all of the objects within the library) as well as the percentage of the system that the library takes up.

Library Usage Analysis

Object Analysis—>Option 1 (Object Enquiries)—>Option 3 (Library Usage Analysis)

This inquiry scans the system for objects not used within a specified time frame. By grouping these within their respective libraries, then sequencing the list, those libraries with the most DASD occupied by unused objects can be immediately highlighted. This then identifies the areas where the highest DASD gains can be made in the shortest possible time, with some manual housekeeping,

The time period ranges from No information (objects not used at all) to Over 1 Year. If No information is selected, the inquiry shows a list of libraries containing objects that have never been used since they were loaded.

If a timeframe is used, in particular the Over 1 year option, the inquiry shows only objects which have been used but not within the selected timeframe. From here you can drill down to further analyze whether there is an opportunity to perform an action such as delete, clear or archive on a particular library.

Library Dates Analysis

Object Analysis—>Option 1 (Object Enquiries)—> Option 4 (Library Dates Analysis)

This inquiry displays significant dates relating to objects in a library but presented at a library level. It provides you the ability to check on the management of the objects in a library by tracking such items as the timing of the latest backup against the changes made to an object. The dates show object usage within the library as follows:

Date	Description
Earliest save	This field shows the date of the object that has the earliest save date within the library. Generally objects are saved by library, but this is obviously not mandatory. Usually this date will indicate the last time that the library was saved, although there may be individual objects that have been saved more recently.
Latest used	This field shows the date of the object in the library that has been used most recently (meaning actual use, such as a file being opened for processing or a program being run). If this date is long past, there is potential for removing the whole library, as all other objects in the library will have been used either on or before this date.
Latest Change	This field shows the latest date on which any object has been changed. This could mean the last time the object was compiled, compressed, or restored. These changes are to the aspect or position of an object, not the last time it was used (a program run or a file opened for reading). It is quite normal, although undesirable, for this date to be quite recent (usually restore dates), even though none of the objects have actually been used for some time. If a file is opened for read only, the Latest use date is updated but the Latest change date is not. If the file is opened and changed, the Latest change date is also updated.
Latest Create	This field shows the latest creation date of any object in the library (for example, the date on which a file was created or an object was originally compiled). Create dates that are later than use dates could indicate that you are running processes that have no use for the created objects.
Latest Restore	This field shows the date of the last object to be restored to the library. As with earliest save dates, this usually refers to the whole library, but is not necessarily the case.

The following actions are available:

Option	Description
F1 (Help)	Displays a list of the currently enabled user options.

Option	Description
F14 (Reverse)	Toggles the order in which the library list is displayed.
F17 (Filter)	Filters the displayed data (see "Filter the Displayed Data" on page 77).

Filter the Displayed Data

F17 (Filter) enables you to filter the displayed data, reducing the time it takes to find particular information.

Object Enquiries

You can run the following object enquiries:

- Objects by Name–Lists all objects on your system, regardless of library, by name (see "Objects by Name" on page 77).
- Objects by Size–Lists all objects on your system, regardless of library, by size (see "Objects by Size" on page 77).
- Objects by Date Used–Lists all objects on your system, regardless of library, by date last used (see "Objects by Date Used" on page 78).
- Object Compression–Specify what to display from various options related to object compression (see "Object Compression" on page 78).
- Damaged Objects–Lists all objects on the system that the operating system has determined are in a damaged state (see "Damaged Objects" on page 78).
- Objects by Owner–Lists all objects owned by a user profile (see "Objects by Owner" on page 79).
- Analysis by Object Type–Displays objects by their object type and object attribute, regardless of library or domain (see "Analysis by Object Type" on page 79).

Objects by Name

Object Analysis—>Option 1 (Object Enquiries)—>Option 5 (Objects By Name)

This inquiry lists all objects on your system, regardless of library, by name.

Objects by Size

Object Analysis—>Option 1 (Object Enquiries)—>Option 6 (Objects By Size)

This inquiry lists all objects on your system, regardless of library, by size. It is helpful to inquire on the objects by size, then scroll down looking for objects that have either no usage or usage in the distant past. These objects are candidates for archiving.

Objects by Date Used

Object Analysis—>Option 1 (Object Enquiries)—>Option 7 (Objects By Date Used)

This inquiry lists all objects on your system, regardless of library, by date last used.

Object Compression

Object Analysis—>Option 1 (Object Enquiries)—>Option 8 (Object Compression)

When you select Object Compression, you can choose what to display from various options related to object compression.

The following options are available:

Option	Description
Compressed Objects	Displays a list of all objects that have been compressed.
Objects Eligible For Compression	Displays a list of all objects eligible for compression after selecting a period of time that objects have not been used.
Objects Not Eligible For Compression	Displays a list of all objects on the system not eligible for compression. Typically, these are objects with type *FILE, *DTAARA, *CMD, etc.
Objects Temporarily Decompressed	Displays a list of all objects temporarily decompressed.
No Benefit From Compression	Display objects for which the operating system has determined that there is no tangible benefit from compression.

Damaged Objects

Object Analysis—>Option 1 (Object Enquiries)—>Option 9 (Damaged Objects)

The Damaged Objects inquiry lists all objects on the system that the operating system has determined are in a damaged state. It gives information about the object name, size, type and whether the object has full or partial damage, as well as that last date changed and used.

NOTE

The only way to remove or repair damaged objects is to use the IBM Reclaim Storage capability.

Objects by Owner

Object Analysis—>Option 1 (Object Enquiries)—>Option 10 (Objects By Owner)

The Objects by Owner inquiry lists all objects owned by a user profile. It is important to be aware of these objects, as they may override security routines.

Analysis by Object Type

Object Analysis—>Option 1 (Object Enquiries)—>Option 11 (Analysis by Object Type)

The Analysis by Object Type inquiry groups objects by their object type and object attribute, regardless of library or domain. In this way, all objects in one group (for example, System 38 Physical files (*FILE, *PF38)) can be grouped together, enabling you to drill down (using option 1) into more detailed information for further analysis. Understanding the spread of object types can help you reclaim space rapidly and also helps Application Administrators look for changes. For example, you could look for Save Files and Journal Receivers to help you control disk growth and reduce DASD usage.

ASP Enquiry

Object Analysis—>Option 1 (Object Enquiries)—>Option 12 (ASP Enquiry)

ASP usage, content and overflows are available using the ASP inquiry, which displays the percentage of disk/ASP used.

The following describes each column on this screen:

Column	Description
Opt	You can perform various functions on the ASPs from the list at the top of the screen.
ASP	The ASP number.
Name	The name of the ASP.

Column	Description
Size In GB	The total size of the ASP in gigabytes (GB).
% Used	The percentage of the total ASP that has been used.
Available GB	The available size, in gigabytes (GB), that remains in the ASP. User ASPs are allowed to overflow into the system ASP.
Threshold	The threshold percent of storage used that, when met, causes the system to send a message to QSYSOPR.

The following actions are available:

Option	Description
5 (Display)	Display ASP attributes (see "Displaying ASP Attributes" on page 80).
9 (Libraries)	Display a list of libraries in the ASP (see "Displaying ASP Libraries" on page 83).

Displaying ASP Attributes

Enter 5 (Display) to display information about the ASP.

The following information is displayed:

Parameter	Description
Auxiliary Storage Pool	The number of the ASP.
Type	The type of ASP.
Short description	The description of the ASP.
Number of disk units	Number of disk units assigned to the ASP.
Storage percent used	The percentage of the total ASP that is being used.
Threshold percentage	The threshold percent of storage used that, when met, causes the system to send a message to QSYSOPR.
Total capacity in MB	The total size of the ASP in megabytes (MB).
Available capacity in Mb	The available size, in megabytes (MB), that remains in the ASP. User ASPs are allowed to overflow into the system ASP.

Parameter	Description
Protected capacity in Mb	The amount of the ASP that is protected in megabytes (MB).
Overflow storage	The number of megabytes of storage that has overflowed from the user ASP into the system ASP.
Overflow indication	An indicator of the result of the ASP overflow recovery operation, which is performed during IPL at the user's request. When this operation is requested, the machine attempts to recover the user ASP from an overflow condition by moving overflowed auxiliary storage from the system ASP back to the user ASP during the storage management recovery step of an IPL. • 0–All overflowed storage was successfully moved. • 1–ASP overflow recovery failed due to insufficient space in the user ASP. • 2–ASP overflow recovery was cancelled prior to completion.
Compression recovery policy	The compression recovery policy for the ASP. If the ASP has compressed drives as part of its configuration, this value controls how overflow situations are handled for this ASP. The following policies allow the user to control what is done when the ASP appears full. The following values are returned: • 1–When the system detects that the ASP capacity is about to be exceeded, the system posts system reference code (SRC) A6xx 0277 in the system control panel and waits for a limited time for space to become available. If space becomes available before the limited time ends, the SRC is removed from the system control panel and normal operations resume. If space does not become available before the limited time ends, data overflows into the system ASP. • 2–When the system detects that the ASP capacity is about to be exceeded, the data immediately overflows into the system ASP. • 3–When the system detects that the ASP capacity is about to be exceeded, the system posts SRC A6xx 0277 in the system control panel and waits indefinitely for space to become available. The user must take action before normal operation resumes. Possible actions include deleting objects from the ASP or changing the compression recovery policy to a value that allows the ASP to overflow.

Parameter	Description
Compressed disk units	Whether there are compressed disk units in the ASP. The following special values are returned: • 0–No compressed disk units in this ASP. • 1–Compressed and uncompressed disk units in this ASP. • 2–All disk units in this ASP are compressed.
Balance type	The type of balance activity that is currently running or was done last. The possible types of ASP balancing are: • 0–No ASP balance activity was requested for the ASP. • 1–Capacity balancing to redistribute data so that the percentage of disk space used is the same on all disk units within the ASP. • 2–Usage balancing to redistribute data so that the percentage of disk activity is the same on all disk units within the ASP. High-use and low-use data is identified by running the Trace ASP Balance (TRCASPBAL) command. Usage balancing moves data among the disk units, guided by the trace results, in an attempt to equalize the utilizations. • 3– Hierarchical Storage Management (HSM) balancing can be run only on an ASP that contains a mixture of high-performance and low-performance disk units. HSM balance moves high-use data to high-performance units and moves low-use data to low-performance units. The high-use and low-use data is identified by running the Trace ASP Balance (TRCASPBAL) command.
Balance status	The current status of the balance function for this ASP. The following special values are returned: • 0–No balance activity has occurred for this ASP. • 1–The ASP balance function is currently running for this ASP. • 2–The ASP balance function is currently in the process of ending. Either the time limit has run out or the End ASP Balance (ENDASPBAL) command was issued for this ASP. • 3–The ASP balance function has run, but was ended before the ASP was completely balanced. The Start ASP Balance (STRASPBAL) command can be used to restart the balance function. • 4–The ASP balance function has completed running. The ASP is completely balanced.

Parameter	Description
Balance date and time	The date and time of the last status change for the balance function. The 13 characters are: • 1–Century, where 0 indicates years 19xx and 1 indicates years 20xx. • 2-7–The date, in YYMMDD (year, month, day) format. • 8-13–The time of day, in HHMMSS (hours, minutes, seconds) format.
Balance data moved	The number of megabytes that have been moved by the balance function. A special value of -2 is returned if the size of this field is exceeded.
Balance data remaining	The number of megabytes that remain to be moved by the balance function before the move is considered complete. A special value of -2 is returned if the size of this field is exceeded.
QSTGLOWLMT system value	The system value of QSTGLOWLMT.
QSTGLOWACN System Value	The system value of QSTGLOWACN.
ASP system storage	The amount of storage in megabytes currently allocated in the ASP for system use.
Allocated to error log	The number of megabytes of auxiliary storage allocated to the error log.
Allocated to machine log	The number of megabytes of auxiliary storage allocated to the machine log.
Allocated to machine trace	The number of megabytes of auxiliary storage allocated to the machine trace.
Allocated to main stg dump	The number of megabytes of auxiliary storage allocated to the main storage dump space.
Allocated to microcode	The number of megabytes of auxiliary storage allocated to the microcode and space used by the microcode.

Displaying ASP Libraries

Enter 9 (Libraries) to display the ASP libraries.

The following information is displayed.

Parameter	Description
ASP	The ASP number.
Library	The library on the ASP.
Text	A description of the library. .
Size	The size of the library.
% of Sys	The percentage of the ASP the library occupies.
Latest Use	The date of the last cross-reference.
Ovr Flw	If the user ASP/IASP has exceeded its capacity limit and overflowed into the system ASP. Y–It has overflowed. N or blank–It has not overflowed.

File Enquiries

Object Analysis—>Option 2 (File Enquiries)

You can run the following file enquiries:

- Deleted Records–Lists all physical files containing deleted records (see"Files With Deleted Records" on page 85).
- Logical File Analysis–Displays all logical files on the system with access path *IMMED, *REBLD or *DELAY, regardless of library (see"Logical File Analysis" on page 85).
- Analysis by Object Type–Displays objects by object type and object attribute, regardless of library or domain (see "Analysis by Object Type" on page 86)
- Index Analysis–Combines the total size of the physical data and logical access paths, then sequences the enquiry on total size with the largest at the top (see "Index Analysis" on page 87).
- Files Nearing Capacity–Lists files getting close to their size limitation (see "Files Nearing Capacity" on page 87).
- Files LVLCHK (*NO)–Displays the files on the system with the level checking set to (*NO) (see "Files Without Level Checking" on page 87).

- Files being journaled–Lists the files on the system that are being journaled, including both user and system files (see"Files Being Journaled" on page 87).
- Files without Members–Displays physical and logical files that cannot hold data or information (see "Files Without Members" on page 87).
- Members by Name/Members by Size–Lists all members across the entire disk, regardless of library, in either name or size sequence (see "Members by Name/Size" on page 88).
- Objects by Name–Lists all objects on your system, regardless of library, by name (see "Objects by Name" on page 88).
- ASP Enquiry–Displays the percentage of disk/ASP used (see "ASP Enquiry" on page 88).

Files With Deleted Records

Object Analysis—>Option 2 (File Enquiries)—>Option 1 (Deleted Records)

This Files With Deleted Records inquiry lists all physical files containing deleted records. This is typically a list of space that can be reclaimed through iOptimize automatic reorganization routines (if you have configured them) or reorganized directly from this screen, using Option 9 (Reorganize physical file member).

Logical File Analysis

Object Analysis—>Option 2 (File Enquiries)—>Option 2 (Logical File Analysis)

The Logical File Analysis inquiry displays all logical files on the system with access path *IMMED, *REBLD or *DELAY, regardless of library.

The following options are available:

Option	Description
Maintenance *IMMED	Selecting Maintenance *IMMED displays a list of logical files with access path *IMMED, along with information on object, library and member name, member size, important dates (changed and last used) as well as the amount of separate days on which this file has been used, up to and including the last used day. These logical files provide fast, random access to your physical data and are invaluable. However, they do take up both a great deal of DASD and a significant amount of processing power – even when they are not being used. iOptimize can optimize and manage these logical files.
Maintenance *REBLD	Displays all logical files with the Rebuild parameter set to *REBLD. These logical files only build the access path at the point they are required (opened) by a user or job. Once opened, they, too, provide fast, random access to your physical data. They do take up a very small amount of DASD, but take no processing power unless actually in use. However, the building of the access path at the point of use can take some time.
Maintenance *DELAY	Displays all logical files with the Rebuild parameter set to *DELAY. These logical files only build the access path at the point they are required (opened) by a user or job. Once opened, they too provide fast, random access to your physical data. They do take up large amounts of DASD, but take no processing power unless actually in use. The building of the access path at the point of use does take time, but not as long as a *REBLD object.

Analysis by Object Type

Object Analysis—>Option 2 (File Enquiries)—>Option 3 (Analysis by Object Type)

The Analysis by Object Type inquiry, groups objects by their object type and object attribute, regardless of library or domain. In this way, all objects in one group (for example, System 38 Physical files (*FILE, *PF38)) can be grouped together, enabling you to drill down (using option 1) into more detailed information for further analysis.

Index Analysis

Object Analysis—>Option 2 (File Enquiries)—>Option 4 (Index)

The Index Analysis inquiry combines the total size, in megabytes, of the physical data and logical access paths (object). The data is sequenced by total size (Mb) with the largest at the top.When considering the footprint taken by a data file, it can be vitally important to consider both the DASD taken by the physical data and the DASD taken by the dependent logical access paths.

Files Nearing Capacity

Object Analysis—>Option 2 (File Enquiries)—>Option 5 (Files Nearing Capacity)

The Files Nearing Capacity inquiry lists files that are getting close to the limitations set on their size. By reviewing these files and taking appropriate action, you can avoid requiring operator intervention to extend the file DASD allowance.

Files Without Level Checking

Object Analysis—>Option 2 (File Enquiries)—>Option 6 (Files LVLCHK (*NO)

The Files LVLCHK(*NO) inquiry displays a list of all the files on your system with the level checking set to (*NO). It might be completely legitimate to have the files set to LVLCHK(*NO), however it is worthwhile to browse through the list for files that shouldn't be without level checking, as these objects can be the source of significant system failures.

Files Being Journaled

Object Analysis—>Option 2 (File Enquiries)—>Option 7 (Files Being Journaled)

The Files Being Journaled inquiry lists all files on your system that are being journaled, including both user and system files.

Files Without Members

Object Analysis—>Option 2 (File Enquiries)—>Option 8 (Files Without Members)

The Files Without Members inquiry lists files that do not contain members, but are taking up DASD. Physical and logical files can be created without

the means to actually hold data or information. Occasionally this is legitimate but, for the most part, files without members are of little or no use and merely take up valuable DASD.

Members by Name/Size

Object Analysis—>Option 2 (File Enquiries)—>Option 9/10 (Members By Name/Members By Size)

The Members by Name and Members by Size inquiries list all members across the entire disk, regardless of library, in either name or size sequence. This information can be invaluable in locating a named member, checking for duplicates or highlighting individual large members.

Objects by Name

Object Analysis—>Option 2 (File Enquiries)—>Option 11 (Objects By Name)

The Objects by Name inquiry lists all objects on your system, regardless of library, by name.

ASP Enquiry

Object Analysis—>Option 2 (File Enquiries)—>Option 12 (ASP Enquiry)

Displays the percentage of disk/ASP used including ASP usage, content, and overflows.

The following columns are included:

Column	Description
ASP	The ASP number.
Size In GB	The total size of the ASP in gigabytes (GB).
% Used	The percentage of the total ASP that has been used.
Available GB	The available size in gigabytes (GB) that remains in the ASP. User ASPs are allowed to overflow into the system ASP.
Threshold	The threshold percent of storage used that, when met, causes the system to send a message to QSYSOPR.

The following actions are available:

Option	Description
5 (Display)	Display ASP attributes.
9 (Libraries)	Display a list of libraries in the ASP.

Action Codes (F1/HELP)

There are many options from which to select. While all options are available at all time, some options do not apply to the current inquiry type. For example, you cannot compress a physical file (IBM i will not allow this). Some allowable options may be difficult to understand at first glance, but the value becomes clear once you are more familiar with using the inquiries. For example, if you were inquiring on the objects in a library and would like to see a list of all of the files in this library (without losing your current position) you could enter F next to any of the objects on your current screen. This takes you to another level (Lvl 3), from which you could inquire on the file then return to the previous screen, exactly where you left it.

At all stages the same functions exist (in other words, when at level 4 you can still see the layout of a particular file displayed on an overlay window). There are a maximum of five levels of enquiry using the select fields. Not all selections, however, result in moving down a level. The available options can always be viewed by pressing the F1 (Help) key.

NOTE

This screen is user-configurable, enabling you to decide which character performs which function.

The following table explains several of the available options:

Option	Key	Description
Cross-reference object	**X**	Performs a one-time rebuild of the cross- reference file for the selected library. Normally, you do this only if you believe that the library or its contents have changed since the last automatic cross-reference.
Clear	**C**	Removes all objects from the selected library. Exercise caution when using this option, make sure there is nothing in the library you want to keep. You may consider first archiving the library. Selecting this displays a confirm action screen.

Option	Key	Description
Delete object	**4**	Deletes the selected library from the disk. As with clearing a library, you should be very careful when using this option. Selecting this function displays a confirm action screen.

Option	Key	Description

Growth	**G**	Displays a histogram showing the growth pattern of the selected library so that you can see which libraries are growing fastest and to forecast growth.
		You can also generate a forecast based on parameters that are displayed when you press F8 (Forecast) from the Object Size Growth screen. These parameters are configurable. The system uses the most appropriate of three smoothing algorithms, based on the amount of historical data with which it has to work.
		1. To use Library Growth Control, use SQL Select to verify that the LIBGTHDAY parameter exists in file IDCTLMP:
		`SELECT * FROM`
		`*Installation-library*/IDCTLMP`
		`WHERE OCPARM = 'LIBGTHDAY'`
		2. Do one of the following:
		– If the parameter does exist, change the day Library Growth runs using the following SQL Update statement:
		UPDATE *Installation-library*/IDCTLMP SET OCVAL = *'xxx'* WHERE OCPARM ='LIBGTHDAY'
		– If the parameter does not exist, add it to IDCTLMP using SQL Insert:
		`INSERT INTO`
		`*Installation-library*/IDCTLMP`
		`(OCPARM, OCVAL, OCDESC,`
		`OCMOD, OCSCOPE,OCDFT,`
		`OCTYPE,`
		`OCLEN)VALUES('LIBGTHDAY',`
		`'xxx', 'Library Growth - day`
		`to run, 'ID', 'U', 'SUN',`
		`'CHAR', 3)`
		xxx = Valid values for this parameter are MON, TUE, WED, THR, FRI, SAT, and SUN. If an invalid value is entered then it defaults to SUN.
		Installation-library = The default installation library name is IOPT71.
		If the Library Growth value is updated the iOptimize subsystem will need to be ended and restarted to register the change.

Option	Key	Description
Change object description	**2**	Enables you to change the name of the selected library.
Archive	**A**	Archives the contents of the selected library to the default staging location `/VisionSolutions/IOPTIMIZE/Archive`, and catalogs its details so it is ready for immediate retrieval. Selecting this function displays a confirm action screen. A separate operation moves all those objects awaiting archiving to an offline storage device (for example, tape), and updates the catalog with the storage media information.
Files by number of records	**F**	Displays all the files in the selected library in descending order based on the number of records each contains. For example, locate the installation library (usually IOPT71), then enter an F next to the library name.
Objects by library and size	**N**	Displays all the files in the selected library in descending order based on the file size.
Reorganize physical file member	**9**	Allows you to reorganize individual files and remove all the deleted records, outside your scheduled housekeeping runs. Selecting this function displays a confirm action screen. This option is not applicable to libraries.
File information	**I**	Displays information about the current status of the selected file and its status the last time that it was analyzed. You can see, at a glance, information such as record length, number of active and deleted records and the amount of space taken up by the deleted records in the file (a very useful bit of information!). There is also information about when it was created, the last time it was used, the last time it was saved and so on. This is an extremely convenient way of displaying this type of file information— and much simpler than it would be using IBM commands!
Display keys	**K**	Displays the file key.
Display file layout	**L**	Provides a quick way of looking at the layout of a particular file. Press F8 to display the key to the selected file.
Display physical file member	**5**	Enables you to review the data in a file. Enter 5 next to the Enquiries file to display the IBM display physical file member (DSPPFM) screen. This option is only valid with physical files.

Option	Key	Description
View dependencies	**D**	Lists all of the logical files that are based on the selected physical file.
Library usage window	**S**	Displays a breakdown of information about the selected library. You can see usage information, which tells you how many objects in the library have never been used (No usage information), how many have not been used for a year, six months and so on. Armed with this information, you can investigate why certain objects have not been used for a long time by drilling down using some of the other features described in this section. You will almost certainly find that there are objects that you can safely remove from your disk, thereby releasing more space for essential use.
Objects by library and name	**O**	Displays all the objects in the selected library in alphabetical order. The information displayed includes the object type (*PGM, *FILE, etc.), its size in Kilobytes and when it was last used. You could also display this information using the Objects by library and date used (J) or Objects by library and size (N) functions. It is important to note that the actions and enquiries you can perform on an object are dictated by the object type itself, not the screen from which you are executing the routine (for example, viewing file dependencies can be done from any screen on which the file is listed).
Begin size monitoring	**B**	Starts monitoring the size of an object. This will then allow you to look at the growth pattern of monitored objects in histogram format. You would not begin size monitoring on all objects, only those that are most indicative of the growth of your applications (Stock transaction history file, Customer master file, etc.). See "Object Size Monitoring" on page 111 for more information. All libraries are automatically configured to monitor size.

Option	Key	Description
References	**R**	Displays all the objects that are referenced by that particular program. Locate an object then enter R next to the object name to display a screen similar to the following: Because you selected a program, the screen displays all the objects that are referenced by that program. From this screen, you can enter Option 1 (Information) next to an object to view the IBM Display Object Description—Full or Option 1 (Service) to view the Display Object Description—Services. In the case of files, you can also view the keys by entering Option 3 (Keys) or view the layout by entering option 5 (Layout). Return to the Object enquiry screen (by pressing F12), then enter R next to the object ACCPTH to display a similar screen, but this time, because you selected a file, the screen displays a list of the programs that reference the selected object.
Members by library and date used	**Y**	Displays each member in each file in the selected library in order of date last used (oldest first).
Move	**M**	Enables you to move objects between libraries, either on the same ASP/IASP or on different ASPs/IASPs. For a list of objects that can be moved using this command, see "Valid Object Types for Move Command" on page 202.

Domains, Snapshots and Growth

From the iOptimize Domains, Snapshots and Growth menu, you can do the following:

- Group libraries by one or more categories for easier analysis (see "Library Registration" on page 95).
- Group profiles by one or more categories for easier analysis (see "Profile Registration" on page 97).
- Display information about each of the configured domains (see "Library Domain Summaries" on page 99).
- Display summary information about each of the configured domains (see "Profile Domain Summaries" on page 100).
- Display information about the configured domain collections (see "Domain Changes" on page 101).

- Display information about the configured library change collections (see “Library Changes” on page 105).
- Display information about configured object change collections, by library (see “Object Changes” on page 108).
- Display information about changes to configured profiles (see “Profile Changes” on page 108).
- Start a collection for a library, domain, object, or profile (see “New Snapshot” on page 110).
- Display all objects on your IBM Power Systems for which the size is currently being monitored (see “Object Size Monitoring” on page 111).

Library Registration

Object Analysis—>Option 3 (Domains, Snapshots, and Growth)—>Option 1 (Library Registration)

Library registration allows you to group libraries by one or more categories for easier analysis.

The Work with Library Registration screen contains the following parameters:

Parameter	Description
Lib/Mask	The library or mask.
Group	The group for the library.
Domain	The domain for the library.
App	The application for the library.
Owner	The library owner.
Size	The size of the library, in megabytes.
Objs	The objects in the library, in megabytes.

The following actions are available:

- Create a mask used to group libraries (see “Creating a Mask” on page 96)
- Modify a library registration (see“Modifying a Mask” on page 97)
- Apply a mask created in “Creating a Mask” on page 96 to the list of libraries/profiles (see “Applying a Mask” on page 97)

- Delete a mask (see "Deleting a Mask" on page 97)
- Display a library/mask registration (see "Displaying a Library/Mask" on page 97)
- Display the objects in the library/mask (see "Displaying Objects in the Library/Mask" on page 97)

Creating a Mask

1. On the Work with Library Registration screen, enter option 1(Create) on the first row.
2. Enter the name of the mask and press Enter. When naming a mask, note the following:
   - The characters can be defined using the System Parameters in IDCTLMP.
   - ? can be used for a single Alphanumeric character.
   - % represents a numerical digit.
   - Asterisks (*) can be a prefix, middle or suffix.
3. In the screen that displays, enter the following:

Parameter	Description
Group	The group for the mask.
Domain	The domain for the mask.
Application	The application for the mask.
Journal	The journal for the mask. The journal is populated from QDFTJRN if it is configured in the corresponding library.
Library	The library for the mask, if required
Accounting Code	The accounting code for the mask. The Accounting code is loaded from the settings on the profile owning the library.
Group 1, 2, 3, 4, 5	Additional groups for the mask, if required.

4. Press Enter when done to create the mask.

Modifying a Mask

1. On the Work with Library Registration screen, enter option 2 (Change) next to a mask to modify its parameters.
2. In the screen that displays, modify the parameters as required (see "Creating a Mask" on page 96 for parameter descriptions).
3. Press Enter when you are finished making changes.

Applying a Mask

On the Work with Library Registration screen, enter option 3 (Apply) next to the mask that you want to apply. A message displays that describes the number of entries that were updated.

Deleting a Mask

1. On the Work with Library Registration screen, enter option 4 (Delete) next to the mask registration you want to delete.
2. In the screen that displays, press Enter to confirm the deletion.

Displaying a Library/Mask

On the Work with Library Registration screen, enter option 5 (Display) next to a library/mask to display the associated library registration information (see "Creating a Mask" on page 96 for parameter information).

Displaying Objects in the Library/Mask

The Displaying Objects in the Library/Mask inquiry lists all objects in the library/mask by name. On the Work with Library Registration screen, enter 8 (Objects) next to the library for which you want to display objects.

Profile Registration

Object Analysis—>Option 3 (Domains, Snapshots, and Growth—>Option 2 (Profile Registration)

Profile registration allows you to group profiles by one or more categories for easier analysis.

The Work with Profile Registration screen contains the following parameters:

Parameter	Description
Prf/Mask	The profile or mask.
Group	The group for the profile.

Parameter	Description
Domain	The domain for the profile.
App	The application for the profile.
Class	The profile class.
Storage	The amount of storage, in gigabytes.
The following additional fields are available on the screen when F11=Alt. View is pressed.	
Owner	The owner of the profile.
Accounting Code	The accounting code for the profile.
Group 1, Group2, Group3, Group4, Group5	Additional groups for the profile.

The following actions are available:

- Create a mask for profiles (see "Creating a Profile Mask" on page 98)
- Modify a profile/mask (see"Modifying a Profile/Mask" on page 99)
- Apply a mask to a list of profiles (see "Applying a Mask" on page 99)
- Delete a profile/mask (see "Deleting a Mask" on page 99)
- Display a profile/mask (see "Displaying a Profile/Mask" on page 99)

Creating a Profile Mask

1. On the Work with Profile Registration screen, enter option 1(Create) on the first row.
2. Enter the name of the mask.
3. In the screen that displays, enter the following:

Parameter	Description
Group	The group for the mask.
Domain	The domain for the mask.
Application	The application for the mask.
Accounting Code	The accounting code for the mask.
Group 1, 2, 3, 4, 5	Additional groups for the mask.

4. Press Enter when done to create the mask.

Modifying a Profile/Mask

1. On the Work with Profile Registration screen, enter option 2 (Change) next to a profile/mask to modify its parameters.
2. In the screen that displays, modify the parameters as required (see "Creating a Profile Mask" on page 98 for parameter descriptions).
3. Press Enter when you are finished making changes.

Applying a Mask

On the Work with Profile Registration screen, enter option 3 (Apply) next to the mask you want to apply. A message displays that describes the number of entries that were updated.

Deleting a Mask

1. On the Work with Profile Registration screen, enter option 4 (Delete) next to the mask you want to delete.
2. In the screen that displays, press Enter to confirm the deletion.

Displaying a Profile/Mask

On the Work with Profile Registration screen, enter option 5 (Display) next to a profile/mask to display the associated information (see "Creating a Profile Mask" on page 98 for parameter information).

Library Domain Summaries

Object Analysis—>Option 3 (Domains, Snapshots, and Growth—>Option 3 (Library Domain Summaries)

The Library Domain Summaries displays information about each of the configured domains.

The Work with Domain Summaries contains the following information for each domain.

Parameter	Description
Domain	The name of the domain.
Size	The size of the domain.
Objects	The number of objects in the domain.
Libs	The number of libraries in the domain.

If required, use F16 (Sequence) to toggle between the different groups configured in Library and Profile Registration (see "Library Registration" on page 95 and "Profile Registration" on page 97).

Profile Domain Summaries

Object Analysis—>Option 3 (Domains, Snapshots, and Growth—>Option 4 (Profile Domain Summaries)

The Profile Domain Summaries displays summary information about each of the configured domains.

The Work with Domain Summaries contains the following information for each domain.

Parameter	Description
Domain	The name of the domain.
Size	The size of the domain.
Libs	The number of libraries in the domain.

If required, use F16 (Sequence) to toggle between the different groups configured in Library and Profile Registration (see "Library Registration" on page 95 and "Profile Registration" on page 97).

Change Analysis

You can collect information about changes to domains, libraries, objects, and profiles. This is useful to reference before and after significant operation events (for example, before and after an upgrade or application installation). See the following sections for more information:

- "Domain Changes" on page 101 to work with domain change collections.
- "Library Changes" on page 105 to work with library change collections.
- "Object Changes" on page 108 to work with object change collections.
- "Profile Changes" on page 108 to work with profile change collections.
- "New Snapshot" on page 110 to configure a collection for library, domain, object or profile.

Domain Changes

Object Analysis—>Option 3 (Domains, Snapshots, and Growth—>Option 6 (Domain Changes)

The Work with Domain Changes screen displays information about the configured domain collections.

The Work with Domain Changes contains the following information:

Parameter	Description
Date	The date on which the domain collection was configured.
Time	The time at which the domain collection was configured.
Reason	A description of the domain collection.

The following actions are available:

- Create a domain collection (see "Creating a Domain Change Collection" on page 101)
- Change a configured domain collection (see "Modifying a Domain Change Collection" on page 102)
- Delete a configured domain collection (see "Deleting a Domain Change Collection" on page 105)
- Display the domain collection information (see"Displaying Domain Change Collection Information" on page 102)
- Print the domain collection information (see "Printing a Domain Change Collection" on page 103)
- Display the changes since the start of the domain collection (see "Displaying Domain Changes for a Collection" on page 102)

Creating a Domain Change Collection

1. On the Work with Domain Changes screen, enter 1 (Create) on the first line to create a collection.
2. In the Start Collection screen, enter a reason for the collection and the retention period for the collection.
3. Press Enter to create the collection.

Displaying Domain Changes for a Collection

You can display the changes that have occurred since a collection was created.

1. On the Work with Domain Changes screen, enter 7 (Changes) next to a collection to display the changes that have occurred since the collection was created.
2. Optionally, enter 5 (Display) next to a category to display additional information about that category.

Parameter	Description
Reason	A description of the collection, entered when the collection was created.
Date	The date on which the collection was started.
Time	The time at which the collection was started.
Requested By	The user that requested the collection.
Detail Entries	The total number of detail records that have been recorded against this collection. Note that when the changes are analyzed, the number of differences determines the number of records displayed in the list. For example, if only one library has changed then you will see only one item listed.
Retention	The retention period. Optional values include: *PERM, *WEEK, *MONTH, *YEAR, *MONTHS2, *MONTHS3, *MONTHS4, *MONTHS5, *MONTHS6, AND *DAY.

Modifying a Domain Change Collection

You can modify the description of a domain change collection.

1. On the Work with Domain Changes screen, enter 2 (Change) next to the domain change collection you want to modify.
2. In Reason, modify the description as required.
3. Press Enter to save your changes.

Displaying Domain Change Collection Information

1. On the Work with Domain Changes screen, enter 5 (Display) next to a domain change collection to display information about the collection.

The following information is displayed:

Parameter	Description
Date	The date on which the collection was created.
Time	The time at which the collection was created.
Requested By	The user that created the collection.
Detail Entries	The number of changes in the collection (see "Displaying Domain Changes for a Collection" on page 102 for additional information).
Retention	The retention period. Optional values include: *PERM, *WEEK, *MONTH, *YEAR, *MONTHS2, *MONTHS3, *MONTHS4, *MONTHS5, *MONTHS6, AND *DAY.

Printing a Domain Change Collection

1. On the Work with Domain Changes screen, enter 6 (Print) next to a domain change collection to print information about the collection. The General Report Request (IDGENREP) screen appears.
2. Refer to the following table to configure the parameters on the General Report Request screen.

Parameter	Value	Description
Report identifier	Character value	The report or enquiry ID of the web object you want to generate the report from, It will look something like ZZOSDE02 (the ID for Libraries by size).
Records to process	0001-9999 *ALL	• 0001-9999—The number of records you want the report to contain. The default is 25. • *ALL—processes all records.
Output	*PRINT *EMAIL	• *PRINT—generates a System i spooled file containing the enquiry output. • *EMAIL—allows you to email the report (after selecting additional options (F24) relating to the format of the email and other settings such as destination, subject etc.)

Parameter	Value	Description
Additional EMAIL Options (only available if *EMAIL is entered above, and F24 or Enter is pressed)		
Email type	*ATTACH (the default) *EMBED	This option controls the format of the email and how it treats the data it gets from iOptimize. *ATTACH—(the default) converts the report to a file and attaches it to the email which can then be detached and viewed externally from the mail client (e.g., saved to the desktop). *EMBED—creates a formatted report inside the body of the mail message itself so it is viewable from inside the mail software immediately upon reading the email.
Attachment type (only available if Email type is *ATTACH)	*PDF *XLS	The output format of the attached file. *PDF—Adobe PDF *XLS–Excel spreadsheet
Compress attachment (only available if Email type is *ATTACH)	*NO *YES	Compresses the email attachment (in ZIP format). WinZIP, or another product, is needed to unzip and view the files.
Address	SMTP address	The email address (or a number of addresses separated by semi-colons) of the person(s) you want to send the report to.
Subject	*REPORT (the default) Text	*REPORT—(the default) uses the title of the report as generated by the enquiry (e.g., ZZOSDE02 will give you the title Libraries by size). Text—Any text can be entered here for the email subject.
Message	Text	Any additional text you want to include in the email outside of the information included on the report itself. For example, "Hi John, here are our current top ten most used files". You can leave this blank if you do not want to include any additional text.

Parameter	Value	Description
F10 (Additional Parameters)		
Extra parameters	*NONE Text	Pressing F10, displays additional parameters. *NONE—(the default) generates the standard report with no filters applied to it. However, some reports come pre-filtered on the GUI and those filters will also need to be applied here. Text—It is possible to add parameters here to re-define a report. You can use this to produce customized versions of standard iOptimize reports.

3. Press Enter to Print or Email the report.

Deleting a Domain Change Collection

1. On the Work with Domain Changes screen, enter 4 (Delete) next to a domain change to delete it.
2. Press Enter to confirm the deletion.

Library Changes

Object Analysis—>Option 3 (Domains, Snapshots, and Growth—>Option 7 (Library Changes)

The Work with Library Changes screen displays information about the configured library change collections.

The following information is included:

Parameter	Description
Date	The date of the library collection was configured.
Time	The time at which the library collection was configured.
Reason	A description of the library collection.

The following actions are available:

- Create a library change collection (see "Creating a Library Change" on page 106)

- Modify a configured library change collection description (see "Modifying a Library Change Collection" on page 107)
- Delete a configured library change collection (see "Deleting a Library Change Collection" on page 108)
- Display a library change collection (see "Displaying Library Change Collection Information" on page 107)
- Print a library change collection (see "Printing a Library Change Collection" on page 107.
- Display the changes since the start of the library collection (see "Object Changes" on page 108)

Creating a Library Change

1. On the Work with Library Changes screen, enter 1 (Create) on the first line to create a library change.
2. In the Start Collection screen, enter a description and the retention period for the library change collection.
3. Press Enter to create the collection.

Creating a Library Change Collection

You can display the changes that have occurred since a collection was created.

1. On the Work with Library Changes screen, enter 7 (Changes) next to a collection to display the changes that have occurred since the collection was created.

   The following information is included:

Parameter	Description
Library	The library for which changes are displayed.
Diff Size	The size of the change.
Diff Object	The number of objects that changed.
Information	Additional information about the library (for example, new libraries, deleted libraries, or ownership changes).
F11=Change view (Additional parameters)	
Cur Size (Mb)	The current object size.

Parameter	Description
Cur Objects	The current number of objects.
Percent Growth	The percentage growth between the current size and the size when the collections was created.

2. The following actions are available:
   - Enter 5 (Display) to display the changes for a specific library.
   - Enter 7 (Changes) to display library changes.

Modifying a Library Change Collection

You can modify the description of a domain change collection.

1. On the Work with Library Changes screen, enter 2 (Change) next to the library change collection you want to modify.
2. Modify the parameters as required (see "Creating a Library Change Collection" on page 106 for parameter information).
3. Press Enter to save your changes.

Displaying Library Change Collection Information

On the Work with Library Changes screen, enter 5 (Display) next to a library change collection to display information about the collection.

The following information is displayed:

Parameter	Description
Date	The date on which the collection was created.
Time	The time at which the collection was created.
Requested by	The user that created the collection.
Detail Entries	The number of changes in the collection.

Printing a Library Change Collection

On the Work with Library Changes screen, enter 6 (Print) next to a library change collection to print information about the collection, and press Enter. The General Report Request (IDGENREP) screen appears.

This is the same General Report Request (IDGENREP) screen that appears when entering 6 (Print) on the Work with Domain Changes screen. For information on the fields, options, and parameters on this screen, see "Printing a Domain Change Collection" on page 103.

Deleting a Library Change Collection

Once you no longer need a collection, you can delete it.

1. On the Work with Library Changes screen, enter 4 (Delete) next to a domain change to delete it.
2. Press Enter to confirm the deletion.

Object Changes

Object Analysis—>Option 3 (Domains, Snapshots, and Growth)—>Option 8 (Object Changes)

The Work with Object Changes screen displays information about configured object change collections, by library.

Enter the library containing the object changes you want to view. The Work with Library Changes screen displays (see "Library Changes" on page 105 for additional information).

Profile Changes

Object Analysis—>Option 3 (Domains, Snapshots, and Growth)—>Option 9 (Profile Changes)

The Work with Profile Changes screen displays information about changes to configured profiles.

The following information is included:

Parameter	Description
Date	The date on which the profile change collection was configured.
Time	The time at which the profile change collection was configured.
Reason	A description of the profile collection.

The following actions are available:

- Create a profile change (see "Creating a Profile Change Collection" on page 109)
- Modify the description of a profile change collection (see "Modifying a Profile Change" on page 109)
- Display profile change collection information (see "Displaying Profile Change Collection Information" on page 109)

- Print a profile change collection (see "Printing a Profile Change Collection" on page 110.
- Delete a profile change (see "Deleting a Profile Change Collection" on page 110)

Creating a Profile Change Collection

1. On the Work with Profile Changes screen, enter 1 (Create) on the first line to create a profile change collection.
2. In Reason, enter a description for the collection.
3. In Retention period, enter the amount of time to keep the collection.
4. Press Enter to create the collection.

Displaying Profile Changes for a Collection

You can display the changes that have occurred since a collection was created.

1. On the Work with Profile Changes screen, enter 7 (Changes) next to a collection to display the changes that have occurred since the collection was created.
2. Optionally, do any of the following:
   - Enter 5 (Display) to display the changes for this profile .
   - Enter 7 (Changes) to display the profile changes.

Modifying a Profile Change

You can modify the description of a profile change collection.

1. On the Work with Profile Changes screen, enter 2 (Change) next to the profile change collection you want to modify.
2. In Reason and/or Retention, modify the description, as required.
3. Press Enter to save your changes.

Displaying Profile Change Collection Information

On the Work with Profile Changes screen, enter 5 (Display) next to a profile change collection to display information about the collection.

The following information is displayed:

Parameter	Description
Date	The date on which the collection was created.
Time	The time at which the collection was created.
Requested by	The user that created the collection.
Detail Entries	The number of changes in the collection (see "Displaying Profile Changes for a Collection" on page 109)
Retention	The retention period. Optional values include: *PERM, *WEEK, *MONTH, *YEAR, *MONTHS2, *MONTHS3, *MONTHS4, *MONTHS5, *MONTHS6, AND *DAY.

Printing a Profile Change Collection

On the Work with Profile Changes screen, enter 6 (Print) next to a profile change collection to print information about the collection, and press Enter. The General Report Request (IDGENREP) screen appears.

This is the same General Report Request (IDGENREP) screen that appears when entering 6 (Print) on the Work with Domain Changes screen and Work with Library Changes. For information on the fields, options, and parameters on this screen, see "Printing a Domain Change Collection" on page 103.

Deleting a Profile Change Collection

1. On the Work with Profile Changes screen, enter 4 (Delete) next to a profile change to delete it.
2. Press Enter to confirm the deletion.

New Snapshot

Object Analysis—>Option 3 (Domains, Snapshots, and Growth)—>Option 20 (New Snapshot)

New snapshot enables you to start a collection for a library, domain, objects, profile, or iOptimize configuration.

1. Enter the following:

Parameter	Description
Type	The type of collection: S–Library *DOMAINS–Domain *OBJS–Object *PROFILES–Profile *CFG–iOptimize configuration
Reason	A description of the collection.
Retention Period	The amount of time to keep the collection.

2. Press Enter.

Object Size Monitoring

Object Analysis—>Option 3 (Domains, Snapshots, and Growth)—>Option 21 (Object Size Monitoring)

This inquiry lists all objects on your IBM Power Systems for which the size is currently being monitored. Monitoring the size of volatile objects gives you valuable information, which can be used in forward planning.

The following parameters are included:

Parameter	Description
Opt	You can perform various functions on the objects from the options listed at the top of the screen. If more than one object is selected, iOptimize displays the information for each subsequent object upon exiting the screen for the previous object.
Library	The name of the library in which each object resides.
ASP	The name of the ASP on which the library resides.
Object Name	The object name.
Object Type	The IBM Power Systems object type classification for each object in the list (for example, if it is a file (*FILE), a program (*PGM)).

Parameter	Description
Date Last Monitored	The date on which iOptimize last checked each object's size.
Size	The size, in bytes, at the time of the last monitor of each monitored object.

Add Objects for Monitoring (F6)

Press F6 (Add) to add new objects to the list of objects being monitored. The Add new object for size monitoring screen is displayed.

Enter the name, library and object type of the object for which you wish to start size monitoring, then press F8 (Confirm).

NOTE

You can also add objects for monitoring from several of the other iOptimize inquiries, which is explained later in this manual.

Specify Gathering Parameters (F18)

Press F18 to specify the parameters that iOptimize needs when monitoring the size of objects.

The following parameters are available:

Parameter	Description
Next Date	The date on which you want the next size monitoring run to occur.
Next Time	The time at which you want the next run to start.
Minutes/Hours/Days/ Weeks/Months/Years	The number of minutes/hours/days, weeks, months and/or years between runs. These fields provide a lot of flexibility in how often you gather size data on the objects being monitored. You can run it as often as you like. This feature could be very useful for checking on something that appears to be growing quickly. In such a case, you could set the monitor to run every half-hour. When you have investigated the results, you could reset it to run less often. Running it every week is recommended.

Press F8 (Confirm) to update the size gathering parameters with the data you have just entered, start the monitor and return to the Object size monitoring screen.

Starting and Stopping Size Monitoring

Size monitoring starts automatically when you first enter the size gathering parameters. To stop size monitoring, press F18 on the Object size monitoring screen to display the Size gathering parameters window, then press F11. To restart size monitoring, press F11 again.

Deleting From Monitor (4)

Enter Option 4 (Delete from Monitor) next to an object to remove it from the list of objects being monitored.

Displaying a Growth Analysis Graph (5)

Enter Option 5 (Display Growth Analysis) next to an object to display a growth analysis histogram for that object.

On the graph, size, in megabytes, is displayed on the vertical axis; the time span is displayed along the horizontal axis. iOptimize automatically constructs the scale, according to the amount of disk space taken up by the object and the amount of historical data it has at its disposal.

To view the graph from a specific date, press F7 (Position).

Press F12 to return to the Object size monitoring screen.

Operations

8

Setting Housekeeping Policies

Operations—> Option 1 (Housekeeping Policies)

From the Housekeeping Policies screen, you can set the processing parameters to clean up the following:

- Output Queues (see "Output Queues" on page 114)
- Job Queues (see "Setting Job Queue Housekeeping" on page 119)
- Message Queues (see "Setting Message Queue Housekeeping" on page 123)
- History Logs (see "History Logs" on page 127)
- Problem logs (see "Setting Problem Log Housekeeping" on page 129)
- Distribution logs (see "Setting Distribution Log Housekeeping" on page 130)
- IBM Journals (see "Setting IBM Journal Housekeeping" on page 132)
- Alerts (see "Setting Alert Housekeeping" on page 133)
- Replaced Objects (see "Setting Replaced Objects Housekeeping" on page 135)
- Optimization (see "Managing Optimization" on page 192).

Output Queues

Operations—> Option 1 (Housekeeping Policies)—>Option 1 (Output) Queues

All IBM Power Systems computers are delivered with a number of system output queues, (for example, QPRINT in library QGPL). In

addition, an output queue is automatically created for every printer attached to a system.

Most companies define their own specific output queues. These are frequently linked to user profiles so that each user has only one place to look for printed output. In some cases, further queues are created for individual departments or to hold particular kinds of output such as job logs.

There are many more output queues on a system. Output queues can fill up quickly for many reasons. For example, a user may run a job that produces a long report and instead of printing the report, they view it on a screen, extract only the information required, and forget to delete the report.

As output queues fill up, more disk space is used. The disk space could be released by performing regular clean up of the output queues and removing unwanted items.

iOptimize can be used to perform scheduled clean up of output queues on a regular basis resulting in freed up disk space each time it runs. The processing parameters are set to specify which output queues to be cleaned up, how frequently to perform the clean up, and how long the output queue should be kept on the system.

The following describes each field:

Field	Description
Queue	The name of the output queue.
Library	The library in which the output queue resides.
ASP	The ASP in which the output queue resides.
S M T...	The day(s) on which the actions are to be performed. The default is daily.
Last Action Date	The date on which the last action was performed on a particular queue.
Last Action Time	The time at which the last action was performed on a particular queue
Days Old	The number of days that old normal entries (RDY) in an output queue must be for them to be actioned by iOptimize. The default is 7.
Save Yes	The number of days jobs that have been printed and saved (SAV) will be kept. The default is 28.

Field	Description
Save No	The number of days jobs that have not yet printed but that will be saved once printed will be kept. The default is 28.
Hld Days	The number of days old held entries (HLD) in an output queue must be for them to be actioned by iOptimize. The default is 14.

Setting Default Parameters

You can set default action values to be applied to all output queues in your list of output queues to action. Doing this saves you having to key in the same values over and over again for all your queues. If there are some queues that you want processed differently from the majority, you can set their individual values by following the steps explained later in this section. Note that if you intend to import your existing output queues, you must have set your default action values.

NOTE

When you change default values, those defaults are applied to all queues. Therefore, if you change the default values after customizing values for individual files, your customized values will be replaced by the new default values.

To set default values for your list of folders, press F22. The Defaults for output queues screen is displayed.

The following describes each field:

Field	Description
Action On...	Type 1 next to each day on which you want queues to be actioned. Type 0 if you do not want them to be actioned.
Days Old	RDY (ready) entries are printed and then deleted from the queue. Enter the number of days old that normal (RDY) entries must be for them to be actioned. The maximum number of days that can be specified is 998. Entering 999 in this field results in no action being performed.
Save Days Printed	SAV (saved) entries are released and printed. Once the print is finished, these entries are returned to a status of SAV. Enter the number of days to keep jobs that have been printed and saved.

Field	Description
Save Days Not Printed	Enter the number of days to keep jobs that have not yet printed but will be saved when they are printed. The maximum number of days that can be specified is 998. Entering 999 in this field results in no action being performed.
Held Days	HLD (held) entries are released and printed. However, once printing is complete these entries are deleted from the queue. Enter the number of days old that held (HLD) entries must be for them to be actioned. The maximum number of days that can be specified is 998. Entering 999 in this field results in no action being performed.

Adding List Items (F6 Add)

To add a new queue to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new queue in the Queue field and the name of the library in which it exists in the Library field. For information on the remaining fields, see "Setting Default Parameters" on page 116. Press F8 (Update) to save your changes and return to the previous screen.

Modifying List Items (Option 2 Change)

You can modify the action values of queues in the list by entering Option 2 (Change) next to the library to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a queue are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of queues, so that they are no longer included in processing runs. To do this enter Option 4 (Remove) next to the queue(s) to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the queue from the list. Press F8 (Update) to confirm.

Viewing Queue Contents (Option 5 View)

You can view the contents of the queues in the list. To do this, enter Option 5 (View) next to the queue. The IBM Work with output queue (WRKOUTQ) screen is displayed. If you selected more than one queue, the contents of each subsequent queue are displayed upon existing the previous queue.

Displaying Summary History

iOptimize maintains a history of all actions taken on output queues at both a detail and summary level. The length of time that history information is held is determined on the Maintain Processing parameters. Enter Option 6 (History) to display a detail history of management on this queue, including printer file name, user name, job number and name, how many pages were in the print member, the date and time it was removed from the queue and whether it was archived prior to deletion. This list displays the details of prints removed in the last run but can be positioned to a specific date and time.

Entering Option 7 (Summary) next to a queue displays a summary history.

Forecasting

iOptimize forecasts the prints to remove, up to a given future point. The forecast can be either created as part of the normal processing for output libraries or can be created on demand.

- Option 8 (View forecast) —shows the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing queue actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the process (action) for the selected queue only (regardless of the current processing mode setting).
- F17 (Run)—actions all of the queues in the list interactively.
- F18 (Submit)—actions all of the queues in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the queues from the managed list, leaving it ready for you to rampart the details. This is particularly useful if you wish to reset the defaults for the company.
- F22 (Defaults)—used to reset the default values for the list of output queues. See "Setting Default Parameters" on page 116 for more information.
- F15 (Import)—imports all the libraries on your system into the list, without the need for you to type all the names in individually. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function. There is a short delay while iOptimize collects all the libraries.

NOTE

If you are using the dynamic maintain function, queues not already in the list (including those that you may have deleted from the list) are automatically found and included into the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Job Queue Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 2 (Job Queues)

All System i computers are delivered with a number of system job queues (for example, QBATCH in library QGPL). If you develop your own application software, it is likely that you will have set up your own job queues for particular uses (for example, compiling programs, test running, live running). If your application software has been written for you by a third party or purchased in package form, the authors will probably have created a number of job queues to handle the processing of batch jobs (for example invoicing, accounting month-end).

iOptimize performs this function for you on a regular basis, freeing-up valuable disk space each time it runs. You set the processing parameters

according to which job queues you want cleaned up regularly, how frequently you want it done and how long you want to keep jobs on the system before iOptimize removes them.

The following is a description of each field:

Field	Description
Queue	The name of the job queue.
Library	The library in which the job queue resides.
ASP	The auxiliary storage pool in which the job queue resides.
S M T...	The day(s) on which the actions are to be performed.
Last Action Date	The date on which the last action was performed on a particular queue.
Last Action Time	The time at which the last action was performed on a particular queue
Days Old	The number of days old entries in a queue must be for them to be actioned by iOptimize.
Save Yes	The number of days jobs that have been printed and saved (SAV) will be kept. The default is 28.
Save No	The number of days jobs that have not yet printed but that will be saved once printed will be kept. The default is 28.
Hld day	The number of days old held entries (HLD) in an output queue must be for them to be actioned by iOptimize. The default is 14.

Setting Default Parameters

You can set default action values to be applied to all queues in your list of queues to action. Doing this saves you having to key in the same values over and over again for each queue. If there are some queues that you want processed differently from the majority, you can set their individual values by following the steps explained later in this section. Note that if you intend to import your existing queues, you must have set your default action values.

To set default values for your list of queues, press F22 (Defaults). The Defaults for job queues screen is displayed.

Following is a description of each field:

Field	Description
Action On...	Type 1 next to each day on which you want queues to be actioned. Type 0 if you do not want them to be actioned.
Days Old	The number of days old jobs that have not run must be for them to be actioned.

Adding List Items (F6 Add)

To add a new queue to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new queue in the Queue field and the name of the library in which it exists in the Library field. For information on the remaining fields, see "Setting Default Parameters" on page 116.

Modifying List Items (Option 2 Change)

You can modify the action values of queues in the list very simply by entering Option 2 (Change) next to the queue to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of queues, so that they are no longer included in processing runs. To do this enter Option 4 (Remove) next to the queue(s) to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the queue from the list. Press F8 to confirm.

Viewing Queue Contents (Option 5 View)

You can view the contents of the queues in the list. To do this, enter Option 5 (View) next to the queue. The IBM Work with job queue (WRKJOBQ) screen is displayed.

Displaying History

iOptimize maintains a history of all actions taken on job queues at a summary level. Enter Option 6 (History) to display a detailed history or Option 7 (Summary) to displays a summary history.

Forecasting

iOptimize forecasts the jobs to remove, up to a given future point. The forecast can be either created as part of the normal processing for output libraries or can be created on demand.

- Option 8 (View Forecast)—shows the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing job queue actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the process (action) for the selected queue only.
- F17 (Run)—actions all of the queues in the list interactively.
- F18 (Submit)—actions all of the queues in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the queues from the managed list, leaving it ready for you to reimport the details. This is particularly useful if you wish to reset the defaults for the company.
- F22 (Defaults)—used to reset the default values for the list of output queues. See "Setting Default Parameters" on page 120 for more information.
- F15 (Import)—imports all the new default values for the list of output queues. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function.

NOTE

If you are using the dynamic maintain function, queues not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Message Queue Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 3 (Message Queues)

All System i computers are delivered with a number of system message queues (for example QSYSOPR in library QGPL). In addition to this, a message queue is created for every user profile that you add and every VDU that you attach to your system. If you develop your own application software, it is likely that you could have set up your own message queues for particular uses within your application. If your application software has been written for you by a third party or purchased in package form, the authors may have set up specific message queues.

Messages are added to queues very frequently. For example, every time you submit a job to be run in batch mode, a message is written when it starts and when it ends. i5/OS itself writes a very large number of messages to internal message queues and, unless these are regularly reviewed by the system operator, they can remain on the system, unnoticed, for long periods.

Message queues can easily be cleared, but many users fail to do this. You will have noticed that whenever you sign on to the System i, any messages on your user profile message queue are displayed before you arrive at your first menu. Many users merely press Enter to progress to the menu, rather than checking the queue and deleting unwanted messages.

The result of neglecting to manage the message queues on your system is that ever increasing volumes of space are taken up on your disk, a great deal of which could be freed up for use by performing regular housekeeping. iOptimize performs this function for you on a regular basis, freeing up valuable disk space each time it runs. You set the processing parameters according to which queues you want cleaned up regularly, how frequently

you want it done and how long you want to keep entries on the queue before iOptimize removes them.

The following is a description of each field:

Field	Description
Queue	The name of the queue.
Library	The library in which the queue resides.
ASP	The ASP in which the queue resides.
S M T...	The day(s) on which the actions are to be performed.
Last Action Date	The date on which the last action was performed on a particular queue.
Last Action Time	The time at which the last action was performed on a particular queue
Days Old	The number of days old entries in a queue must be for them to be actioned by iOptimize.
Save Yes	The number of days jobs that have been printed and saved (SAV) will be kept. The default is 28.
Save No	The number of days jobs that have not yet printed but that will be saved once printed will be kept. The default is 28.
Hld day	The number of days old held entries (HLD) in an output queue must be for them to be actioned by iOptimize. The default is 14.

Setting Default Parameters

You can set default action values to be applied to all queues in the list of queues to action. Doing this saves you having to key in the same values over and over again for each queue. If there are some queues that you want processed differently from the majority, you can set their individual values by following the steps explained later in this section. Note that if you intend to import your existing queues, you must have set your default action values.

To set default values for your list of queues, press F22 (Defaults). The Defaults for message queues screen is displayed.

The following describes each field:

Field	Description
Action On...	Type 1 next to each day that you want queues to be actioned and 0 if you do not want them to be actioned.
Days Old	The age of a message, in days, when it is actioned.

Adding List Items (F6 Add)

IMPORTANT

To add a new queue to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new queue in the Queue field and the name of the library in which it exists in the Library field. For information on the remaining fields, see "Setting Default Parameters" on page 120.

Modifying List Items (Option 2 Change)

You can modify the action values of queues in the list by entering Option 2 (Change) next to the queue to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of queues, so that they are no longer included in processing runs. To do this enter Option 4 (Remove) next to the queue(s) to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the queue from the list. Press F8 (Opiate) to confirm.

Viewing Queue Contents (Option 5 View)

You can view the contents of the queues in the list. To do this, enter Option 5 (View) next to the queue. The IBM Work with message queue (WRKMSGQ) screen is displayed.

Displaying History

iOptimize maintains a history of all actions taken on job queues at a summary level. Enter Option 6 (History) to display a detailed history or Option 7 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the messages it intends to remove, up to a given future point. The forecast can be either created as part of the normal processing for output libraries or can be created on demand.

- Option 8 (View Forecast)—shows the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

More detail on forecasting and other additional options are explained in "Maintaining Processing Parameters" on page 197.

Additional Processing

There are several ways of processing message queue actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the process (action) for the selected queue only.
- F17 (Run)—actions all of the queues in the list interactively.
- F18 (Submit)—actions all of the queues in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the queues from the managed list, leaving it ready for you to reimport the details. This is particularly useful if you wish to reset the defaults for the company.
- F22 (Defaults)—used to reset the default values for the list of message queues. See "Setting Default Parameters" on page 124 for more information.
- F15 (Import)—imports all the new default values for the list of message queues. This overrides any customized settings. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function.

NOTE

If you are using the dynamic maintain function, queues not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in"Understanding Processing Modes" on page 44.

History Logs

Operations—> Option 1 (Housekeeping policies)—>Option 5 (History Log)

The History Log provides a basic record of all system activity on a System i for every job and every user. It contains details about such things as:

- Job start and end times
- Subsystem start and end times
- System operator messages
- Certain error messages.

There is a default system value that determines the size of the History Log. However, when the log gets full, i5/OS freezes the old log and creates a new one, which, in turn, starts to grow, until it, too, gets full and a further one is built. Left to its own devices, i5/OS will continue to create History Logs, using up more and more disk space as it does so. iOptimize enables you to avoid unnecessary disk usage by managing your History Log on a regular basis. You can specify how long you want your History Log kept and then let iOptimize clean them up as often as you like.

Following is a description of each field:

Field	Description
Days To Keep	The number of days old entries in the log must be for them to be actioned and cleared by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter 1 next to each day on which you want the log to be actioned; enter 0 next to each day on which you do not want the log to be actioned.

Press F8 (Update) to update the parameters. The Save History Logs screen is displayed.

If you do not wish to save the history logs (the default) leave the fields on this screen empty and press F8 (Update) to continue. To save the logs, enter a 1 in the Save history logs before deletion field, enter a current library to which to save the logs, then press F8 (Update) to continue.

It is recommended that the library used is backed up to tape and then cleared. Alternatively, you could divert these logs to the archive library and control the archiving using the archiving routines (see "Working with Archiving" on page 145 for more information about archiving).

Displaying History

iOptimize maintains a history of all actions taken on the log at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the log entries it intends to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. Press F22 (Forecasts) to display the current forecast, if one exists. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing job queue actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the log clearing interactively.
- F18 (Submit)—performs the log clearing by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Problem Log Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 6 (Problem Log)

The Problem Log keeps comprehensive records about system problems in order to build up a database for use with the i5/OS problem analysis function.

Certain messages that are sent by i5/OS to the system operator message queue (QSYSOPR) have an asterisk (*) beside them. This is accompanied by the message "Problem analysis allowed for message." Each time a message of this type is sent to QSYSOPR, a substantial amount of associated information is written to the problem log database. This information is then available for use in the future with the problem analysis function.

In most System i installations, problem analysis is rarely, if ever, used and even if it is, the information that it works with is not cleared after an analysis run. Consequently, the Problem Log database continues to grow and grow, unless it is managed. iOptimize helps you manage your Problem Log by enabling you to specify how long you want to keep the information and how often you want the housekeeping procedure to be run.

Following is a description of each field.

Field	Description
Days To Keep	The number of days old entries in the log must be for them to be actioned and cleared by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter 1 next to each day on which you want the log to be actioned; enter 0 next to each day on which you do not want the log to be actioned.

Displaying History

iOptimize maintains a history of all actions taken on the log at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the log entries it intends to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. Press F22 (Forecasts) to display the current forecast, if one exists. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing logs in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the log clearing interactively.
- F18 (Submit)—performs the log clearing by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Distribution Log Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 7 (Distribution Log)

The Distribution Log is used in much the same way as the History Log, but in the area of communications. Whenever there is any activity on your communications network, an entry is written to this log. The size of this log depends on to what extent you use your System i communications function. However, even if you do not have a communications network, entries are written to the Distribution Log for any IBM Electronic Customer Support

(ECS) activity, such as PTFs received electronically or access to DIAL-IBM.

i5/OS does little to manage the Distribution Log, so any use, however small, means that disk space is, after a while, being used unnecessarily. iOptimize manages your Distribution Log on a regular basis, by clearing it according to parameters set by you.

The following is a description of each field.

Field	Description
Days To Keep	The number of days old entries in the log must be for them to be actioned and cleared by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter `1` next to each day on which you want the log to be actioned; enter `0` next to each day on which you do not want the log to be actioned.

Displaying History

iOptimize maintains a history of all actions taken on the log at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the log entries it intends to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing logs in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the log clearing interactively.
- F18 (Submit)—performs the log clearing by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:

– *GENFCS—Forecast only

– *PRCFCS—Action previous forecasts only

– *NORMAL—Action items directly

More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Setting IBM Journal Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 8 (IBM Journals)

Journals are copies of data, written to allow recovery, should your system go down in the middle of processing. A number of IBM products (for example, Office/400, SNADS, PC Support) use journaling for recovery purposes. In addition, journal entries are written whenever any problem analysis information is collected (see "Setting Problem Log Housekeeping" on page 129). IBM Journals contain all the data necessary to allow recovery from processing problems. These journals can take up a great deal of space.

i5/OS does not manage these journals very effectively, therefore journal receivers continue to grow unless managed. iOptimize enables you to clear journals according to your requirements.

The following is a description of each field.

Field	Description
Days To Keep	The number of days old journal receivers must be for them to be actioned and cleared by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter 1 next to each day on which you want the journal to be actioned; enter 0 next to each day on which you do not want the journal to be actioned.

Displaying History

iOptimize maintains a history of all actions taken on the journals at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the journals to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. See the "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing journals in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the clearing interactively.
- F18 (Submit)—performs the clearing by submitting a job to the default batch queue.

  While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.
- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Alert Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 9 (Alerts)

Alerts enable you to specify that certain messages are to be sent to other message queues in addition to the queues to which they would normally be sent. They are primarily used in a network environment, so that, for example, copies of certain messages about the network will always be sent to the system operator.

In addition to copies of messages being sent to other message queues, each alert writes comprehensive data to the alert database, including information about the circumstances prevailing at the time the message was generated. This data can then be referenced at a later date, although, in reality, it rarely is.

The amount of disk space used up by the alert database depends largely on whether you have many user-defined alerts. However, even without defining any alerts yourself, the database continues to grow. This is because there are a number of internal i5/OS messages that have been defined as alerts by IBM.

On its own, i5/OS will not do a great deal towards the management of your System i alert database. But, with iOptimize, you can keep this aspect of your system cleaned up, specifying the age of alerts to be removed and the frequency with which you want the job performed.

The following is a description of each field.

Field	Description
Days To Keep	The number of days old alert messages must be for them to be actioned and cleared by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter `1` next to each day on which you want the Alerts to be actioned; enter `0` next to each day on which you do not want them to be actioned.

Displaying History

iOptimize maintains a history of all actions taken on alerts at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the entries in the alert database to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing alerts in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the clearing interactively.
- F18 (Submit)—performs the clearing by submitting a job to the default batch queue.

  While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.
- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only

– *PRCFCS—Action previous forecasts only

– *NORMAL—Action items directly

More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Setting Replaced Objects Housekeeping

Operations—> Option 1 (Housekeeping Policies)—>Option 10 (Replaced Objects)

Whenever a new object is created to replace an old one in recent use, i5/OS moves the old object to a system library (QRPLOBJ), giving it a new name and text that describes the original library and name. This means that the old object can be retrieved if necessary, although this rarely happens in practice. The replaced objects library continues to grow, as more and more objects are recreated, unless it is managed. iOptimize enables you to automatically manage replaced objects, by checking them as regularly as you wish and cleaning up the library, according to the object age that you specify.

The following is a description of each field.

Field	Description
Days To Keep	The number of days since an object was replaced that it must be for it to be assessed and actioned by iOptimize.
S M T...	The day(s) on which the actions are to be performed. Enter 1 next to each day on which you want the objects to be actioned; enter 0 next to each day on which you do not want them to be actioned.

Displaying History

iOptimize maintains a history of all actions taken on objects at a summary level. Pressing F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts objects to remove, up to a given future point. The forecast can be either created as part of the normal processing for the job. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing objects for removal in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the clearing interactively.
- F18 (Submit)—performs the clearing by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs are run frequently and regularly.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Managing Spool File

Operations —> Option 2 (Manage Spool)

iOptimize supports functions that archive and restore printed output. Prints can be selected manually or can be linked to the iOptimize output queue management in the base system. Here, prints that have become eligible to be removed from the queue due to the age of the print, can be selectively archived for possible retrieval later (for example, invoices). At all times iOptimize maintains a catalog of the items that have been archived (prints are kept separate from general archived objects) and allows the prints to be restored to any queue at any time.

On the Spool file manager menu there is a subwindow that indicates the prints that have been manually or automatically selected for archive—removed from the relative queue but not yet moved to off-line media.

Manually Archiving

Operations —> Option 2 (Manage Spool) —> Option 1 (Manual Archive)

At all times you can select this option to archive a particular print. However, it is typically easier to use the Archive from list option, as this option requires that you have all of the spool member available.

The following is an explanation of each field:

Field	Description
File name	The spool file name of the print to be archived.
Job	The job name of the print to be archived.
User	The username of the user who created the print.
Number	The print job number of the print to be archived.
File Number	The spool file number of the print to be archived.
Archive and Delete	Enter *YES if the print being archived is also to be deleted from the print queue once the print is archived (default). Enter *NO if the archived print is to remain on the print queue after archiving (a copy has been made).

Once all parameters have been entered, press Enter to archive the print. The display returns to the menu and the Archiver status window is updated to reflect the new archived print.

NOTE

Remember that this print is in the staging area on disk and has yet to be taken off-line (see "Archiving to Tape" on page 139). The archive print entry can be viewed using any of the Archived by options.

Archiving from List

Operations —> Option 2 (Manage Spool) —>Option 2 (Archive From List)

This option enables you to group all current prints across all queues by any linking feature, such as Forms type. You may find it much easier to start your archiving from this point, because as soon as a print or group of prints are identified for archiving they are simply flagged and all of the spool member information is automatically collected and used in the archiving.

The following is an explanation of each field:

Field	Description
List Parameters	iOptimize creates a list of those prints of interest, based on the search criteria (list parameters). At this stage the list is empty, so you need to either enter selection criteria or press Enter to view all prints for the current user. The list parameters are cumulative, so you could see, for example, all prints belonging to user FRED, where the Forms type is INVOICE and the user data is ACC-CO1.
List spooled files for User	Enter the name of the user whose prints you wish to see in the list. The parameter defaults to *CURRENT, which means the currently signed on user. The only other valid generic entry is *ALL which will list prints for all users.
Output Queue	Enter the name of the output queue containing the prints you wish to see, or enter *ALL.
Library	Enter the library name in which a specific output queue resides (only necessary if you have a duplicate queue name), or enter *ALL.
User Data	Enter the specific user data that must match prints for selection, or enter *ALL.
Forms Type	Enter the specific forms type of prints to be selected, or enter *ALL.
Delete Spooled Files After Archiving	Enter *YES if the spool files are to be removed from the respective output queues after they have been put into the archive subsystem (default). Enter *NO if the original spool file entries are to remain even after an archive copy has been taken. A list is generated. You may start with the parameters set to view quite a wide variety of prints (for example, all prints for user ACCOUNTS) and then fine-tune this by changing the parameters to a more specific search list (for example, all prints with forms type of INVOICE for user ACCOUNTS).

The following is an explanation of each option:

Field	Description
Option 1 (Archive)	Typing 1 next to a print (or a number of prints) in the list starts the archive procedure. A confirmation message is displayed advising that the selected print(s) is being archived. When you rebuild the list, you will see that the prints are gone.
Option 5 (Display)	Entering 5 next to the print invokes the IBM Display Spool File (DSPSPLF) process, just as when reviewing any output queue currently.
Option 8 (Attributes)	Entering 8 next to a print invokes the IBM Work with Spooled file attributes (WRKSPLFA) command, in the same way as you would when reviewing output queue items currently. Selecting F16 starts the archive for all of the prints in the currently created list. A confirmation message indicates that the prints have been archived. Using this method, the list is rebuilt and should now be empty. These prints can now be viewed by any of the Archived by options and are reflected in the Archiver status window until they are moved to off-line media. They will, however, remain in the catalog but they will now be marked with tape details.

Archiving to Tape

Operations —> Option 2 (Manage Spool) —>Option 3 (Archive To Tape)

As prints are archived using either the manual archive or archive from list functions, they are placed into a staging area and are reflected in the Archiver status window. These prints need to be moved to off-line media (for example, tape, optical, silo). This option accomplishes that task by moving the prints from the available DASD to the off-line media.

Following is an explanation of each field:

Field	Description
Nbr Days In Archiving	Enter the number of days old the entries must be for them to be actioned.
User	Enter the name of the user whose prints you wish to archive to tape. The parameter defaults to *CURRENT (the user currently signed on). *ALL archives prints for all users.

Field	Description
Output Queue	Enter the name of the output queue containing the prints you wish to archive to tape. *ALL archives prints from all output queues.
User Data	Enter the specific user data that selected prints must match. Enter *ALL to select all prints.
File Name	Enter a specific spool file name. Enter *ALL to select all files.
Library	If you have duplicate queue names, enter the library name in which the specified output queue resides. Otherwise enter *ALL.

Press F16 (Archive list), then type the device name to be used for archiving the prints. If necessary, load a tape, cassette, cartridge or disk into the appropriate device, then press F8 (Continue). The screen displays the task progress and resets the Archiver status when completed.

Once archiving is complete you are prompted to archive onto a second tape, if desired. Press F8 to rearchive or press F12 to ignore. To verify the archive, view the print catalog, which now displays the tape/device details on the print archive record.

Automating Spool File Archiving

You can invoke or schedule the following command to save the spool file staging area to tape:

OSDSPL2TAP

Auto Archive/Override Parameters

Operations —> Option 2 (Manage Spool) —>Option 4 (Auto Archive Parameters)

iOptimize enables you to maintain output queues. Using this function, you can set parameters to ensure that old prints are cleared from the queue after a certain length of time.

There may be the requirement, either occasionally or permanently, to override the rules for a particular output queue, library, user, forms type, user data or printer file. The override function described in this section enables you to change the rules by which the old prints are cleared based on a selection of some or all of the aforementioned criteria, overriding the previously set default rules for an output queue.

There may also be occasions when you wish to archive selected prints before clearing them, either using either the criteria established in the iOptimize or the override function. An example might be to remove the

monthly sales analysis print after 45 days, despite its SAVE status on the queue. These prints are usually large and the space saving resulting from their removal worthwhile. Additionally, removing such a report causes less confusion when the new months figures are produced. However, there may be a requirement to occasionally view the print after the 45 day period has expired, and rerunning the print job may be either lengthy or impossible. It makes sense, then, to archive this print rather than just deleting it. You could then build up a catalog of these reports, which could be restored and browsed or reprinted at any time without having to overload the output queue, the disk or the printer.

These options forge the link between the iOptimize output queue cleanup routines and the archiving function.

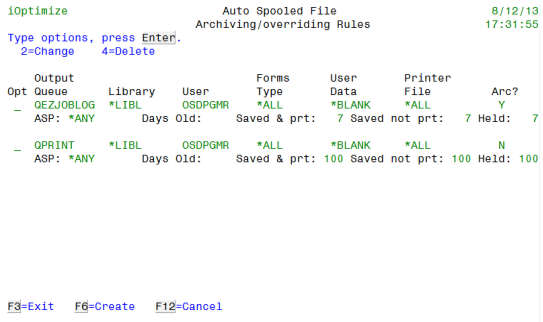

The column headings correspond to the IBM use of these fields in the general IBM Power Systems output queue management routines. The following information further explains the fields in this example.

Setting Archive Link Rules

Using the previous screen as an example, note the following:

QPRINT / *ALL / *ANY / ANDY/ *ALL / *BLANK / *ALL / Y

Days old : 100

Saved & prt: 100

Saved not prt: 100

Held: 100

This rule indicates that any print from the output queue QPRINT in library *ALL, in any ASP, belonging to user ANDY, regardless of forms type and printer file, but with the user data set to *BLANK, must be archived at the point it becomes eligible for clearing, provided that it fulfills the age and status settings rules attached to this override entry. These rules override the defaults set up in iOptimize.

iOptimize detects this rule, along with the archive before deletion indicator, and archives the print, placing it in the staging area ready for moving off-line and then clearing the print from the queue.

As you can see, you can be extremely selective or quite generic and still, with only a handful of rules, ensure that all important items are archived before being automatically cleared.

You can set up as many of these rules as you like, or you can set up a rule to archive ALL of your prints (*ALL/*ALL/*ALL/*ALL/*ALL/*ALL/ Y) that comply with the default rules set up in iOptimize (leave the override Days Old fields blank), in addition to whatever ad hoc archiving you may perform.

Creating, Modifying, Deleting a Rule (F6)

Press F6 (Create) to create a rule.

Enter the values that bracket or pinpoint the print(s) you would like to archive before deletion. Press F8 to confirm and update. The rule, if accepted, is displayed in sequence on the main screen (once updated).

Modifying a Rule (Option 2 Change)

Enter Option 2 (Change) next to a rule you wish to modify. A window similar to the Add window is displayed. The fields are filled in with the values of the selected rule. Change the rules as desired then update by pressing F8.

Deleting a rule (Option 4 Delete)

Enter Option 4 (Delete) next to a rule you wish to delete. A confirmation window is displayed. Press F8 to confirm the deletion or F12 to cancel.

NOTE

The output queue clearing can be set to forecast the actual work that would be done on clearing prints (*GENFCS), even to the point of naming which print will be cleared and when. As part of this function, the forecast identifies whether the print will archive before any clearing would be done. This is a simple way of making sure that the rules are working in the way you intended.

Displaying the Print Archive Catalog

Operations —> Option 2 (Manage Spool) —>Options 5–9 (Archive By Archive Date/User/Print Date/Print File/User Data)

This series of enquiries allows you to view the print archive catalog in a variety of ways, all of which are designed to help you find the print you would like to restore to DASD and to a live output queue for printing. This section uses the Archived by archive date option as an example, however the same rules apply for all other archive enquiries.

The columns describe the archived print in the same manner as the IBM functions describe prints in a queue. This list is in reverse date and time sequence.

Restoring an Object

If the object is to be restored back to the queue from which it was saved, simply enter Option 1 (Restore) next to the print(s) to be restored. If the print was in the staging area (it has not yet been taken to off-line media), the print is restored to the relevant output queue ready for reprinting. If the object was on off-line media you will be prompted to enter the relevant tape (iOptimize specifies which one) and from which device to restore. This can be overridden.

Restoring to Another Queue (F4)

When you restore as described in the previous section but press F4 instead of Enter, you are prompted for the output queue to which you would like to restore. So, even though the print was originally in QPRINT and was archived from there, it can be restored to PRT04 just as easily.

NOTE

When a print is restored from archive, it is copied (not moved) from archive.

Displaying Tape Details

Entering Option 7 (Tape details) next to a print displays details about the media on which it is currently stored, as well as details about any other areas where it might be stored off-line (dual backups, etc.). The details are displayed in an overlay window. If the print is still in the staging area, no tape details are displayed.

Duplicating a Spool File

Operations —> Option 2 (Manage Spool) —>Option 10 (Duplicate Spool File)

From time to time you may need to duplicate a particular print, perhaps to get a number of prints going at the same time or to distribute the same print to a number of different printers in different locations.

The following is an explanation of each field:

Field	Description
Spooled File, Job Name, User, Number & Spooled File Number	These fields relate to the spooled file you wish to duplicate. The field meanings are the same as the equivalent IBM meanings. For further information on the field meanings see the equivalent IBM Manual.
Output Queue, Library	The destination queue in which the duplicate is to be placed.
User To Own New Splf	Enables you to change ownership of duplicated spool file. Default is *SAME.
User Data	Enter the new user data to override the current setting and identify the print. Default is *SAME.
Hold Spooled File	Enter *YES or *NO to dictate whether the print is to be held before printing on the new queue. Default is *SAME.
Save Spooled File	Enter *YES or *NO to dictate whether the print is to be saved after printing on the new queue. Default is *SAME.
Form type	Enter any value to override the current Form type value. The default is *SAME.
Copies	Enter any value to override the current number of copies value. The default is *SAME.

Saving all Spool Files

Operations —> Option 2 (Manage Spool) —>Option 11 (Save All Spool Files)

It is important to save printed output in output queues at the time that you perform your regular backups. Otherwise, if you have to reload from backup, you will have lost all your prints. iOptimize enables you to save prints using the OSDSAVSPLF command via option 11. The command can be integrated into your regular backup routines.

Although the option indicates only a mass save, you can be completely selective, saving just one output queue for instance. Allowable entries are unlimited, but it is recommended that you include at least one save of all output to coincide with full backups.

Restoring all spool files

Operations —> Option 2 (Manage Spool) —> Option 12 (Restore All Spool Files)

If you need to reload from backup following a disaster, and have been prudent enough to save your prints using the iOptimize spool file save function OSDSAVSPLF, you can perform a recovery using the Restore all spool files function.

Using the OSDRSTSPFL command via option 12, you either can restore all of the backed-up prints or be specific and restore only those prints that are critical for your operation to continue. The fields are used in the same manner as all IBM output queue management functions.

Working with Archiving

Operations —> Option 3 (Archiving)

The Archiving menu option on the Operations window opens the Object Archiving screen.

iOptimize enables you to select objects for archive at any time. As objects are selected, they are moved out of the library they are in (all objects are moved if a library is selected for archive) and placed into a staging area on disk. When you have selected a number of items for archiving you will move them off-line together, although they are always cataloged and referenced as separate objects. While the items are in the staging area they are reflected in the Archive Status window on the Archiving menu. Press F14 (Archive Status) to display the Archive Status window.

The Archive Status window reflects the number of objects and the total size of the objects that have been selected for archiving but not yet moved to off-line media.

Managing Media

Operations —> Option 3 (Archiving)—>Option 1 (Manage Media)

This function gives you control over the tapes that you use for archiving. It helps you track all tapes currently containing archive material, without having to insert it in the tape drive and then wait until the information is displayed to you by IBM's DSPTAP command.

Column	Description
Tape ID	The archive tape name.
Libs	The number of libraries that have been archived to each tape.
Objs	The total number of objects that have been archived to each tape.
Earliest	The date of the earliest archive operation to each tape (for the objects and libraries currently resident).
Latest	The date of the most recent archive operation to each tape.

Initialize Tape (Option 6)

You can initialize tapes that are in the list from this screen. When you do this, iOptimize runs the IBM command INZTAP (Initialize tape). This clears the tape(s). For this reason, be sure that there is nothing of importance on tapes before initializing them.

In addition to clearing the tape, this function clears all iOptimize references to the contents of the tape, so that they will no longer be displayed on the Archived objects screen (explained later in this section).

To initialize a tape, put the tape into the IBM Power Systems tape drive. Enter Option 6 (Initialize tape) next to the tape you wish to initialize. Your keyboard is locked while the operation is carried out.

Delete Details (Option 4)

This option will remove a tape from the archive list. You can remove tapes from the list when they become faulty or are no longer used for archiving. To do this, enter Option 4 (Delete details) next to the tape(s) you want to remove. The selected tape(s) are removed from the list.

Archiving Now

Operations —> Option 3 (Archiving)—> Option 2 (Archive Now)

iOptimize enables you to select objects to archive at any time. Archive Now provides a window, Specify Archive Device, that includes a prompt to select the device. An option is included for *BRMS which prompts for Media Policy and Control Group. (These fields require that you specify valid values. They may not be left blank and *NONE is not permitted. If unsure about what to enter, refer to the appropriate IBM documentation for Backup and Recovery Media Services.)

As objects are selected, they are moved out of the library they are in (all objects are moved if a library is selected for archive) and placed into a staging area on disk. When you have selected a number of items for archiving you will move them off-line together, although they are always cataloged and referenced as separate objects. While the items are in the staging area they are reflected in the Archive Status window on the Archiving menu. Press F14 (Archive Status) to display the Archive Status window.

The Archive Status window reflects the number of objects and the total size of the objects that have been selected for archiving but not yet moved to off-line media.

Multiple Media Archive

When the first archive operation has completed, a screen is displayed prompting you as to whether you would like to make additional copies of the archive library on tape. This is often a wise thing to do so that, for reasons of security, you can keep one copy on site and one off-site. Press F8 to make another copy. The Which device window is displayed.

Automating Object Archiving

You can invoke or schedule the following command, to save the near-line archiving data to tape:

OSDOBJ2TAP

Displaying Archived Objects on Media

Operations —> Option 3 (Archiving)—> Option 3 (Objects Archived)

This function enables you to look at a list of all your archived objects and displays what they are, which tape, when archived, and so on. It can be useful if you need to restore something quickly, as you can determine very quickly where the object can be found.

The following describes each column:

Column	Description
Name	The name of the archived object.
Library	The name of the library in which each object in the list resided prior to being archived.
Member	The name of the file member.

Column	Description
Type	The type classification of each object in the list (*PGM, *FILE etc.).
Internal Name	The name assigned to each object in the list when it was archived.
Type	The type of archive done to create this object entry (by library (L) or by object (O)).
Vol ID	The name of the off-line tape on which each object can be found.
Seq	The sequence number on the tape, of each object in the list (exactly where it is on the tape).
Date	The date that each object was first archived. There may be other copies of each object, on the same or different tapes, which were made for security purposes. How these can be displayed is explained in the following section, “Viewing multiple copies.”

Viewing Multiple Copies (Option 1)

If you have created more than one copy of an object on tape, you can determine the location of the other copies. Enter Option 1 (More Information) next to the relevant object. A window is displayed, describing all of the occurrences of this object on different media.

Searching for Objects

If you are not sure of the name of an archived object with which you wish to work, but you know that, for example, the word orders appears in the associated descriptive text for it, you can use this function to search the text. Press F22 (Search) to display a window from which you can enter search criteria. Type the text that you want to use to search for your object. You can specify up to ten characters. A subset of the whole archive list is displayed, in the same format as above.

Restoring Objects to DASD (Option 2)

You can restore an archived object to your IBM Power Systems from either the main display or from the Information window. Where items are restored from the main display, if multiple occurrences are found then the latest object will be restored. If you know there are multiple occurrences of the object and you do not want the latest occurrence, use the Information window (1) to select the particular occurrence to restore (2).

Insert the tape containing the object that you wish to restore (iOptimize displays which tape it is), then enter Option 2 (Restore) on the appropriate object/tape line.

From the Restore menu, enter the name of the device which holds the archive tape in the field. The default for this field is always TAP01, but you may change this if another device is being used. There is a delay while the archived object is transferred back to your IBM Power Systems, during which time your keyboard will be locked. When the object has been successfully restored, the archive copies window is displayed.

To abandon a restore, press F12 from either window. To return directly to iOptimize's main menu without restoring the object press F3.

Displaying Library Archives (F14)

This function displays a panel summarizing the list of archived objects by library. Press F14 (Libraries) to display Archived Libraries.

The options available to you on this screen work in the same way as those for the full Archived objects display. Press F14 (Show objects) to toggle between the two types of display.

Displaying Objects Awaiting Archive

Operations —> Option 3 (Archiving)—> Option 4 (Objects Waiting to be Archived)

This enables you to view a list of all objects that have been selected for archive that have not yet been sent to off-line media, but have been saved to a temporary staging area. It displays what they are, when they were archived, and so on. It can be very useful if you need to restore something quickly, as you can identify very quickly where a particular object can be found. Once you have found the desired object(s), you can use this function to restore them to your IBM Power Systems.

The following is a description of each column:

Column	Description
Opt	You can select objects to restore by entering 2 next to the object. If more than one object is selected, iOptimize displays the information for each upon exiting the previous one.
Name	The name of the archived object.
Library	The name of the library in which each object in the list resided prior to being archived.
Member	The name of the file member.
Type	The type classification of each object in the list (*PGM, *FILE etc.).

Column	Description
Internal Name	The name assigned to each object in the list when it was archived.
[Type]	The type of archive done to create this object entry (by library (L) or by object (O)).
Vol ID	No details are displayed.
Seq.	No details are displayed.
Date	The date that each object was first archived.
[Status]	The status of the archive object at this time (I for in progress, empty for awaiting archive).

Searching for Objects

If you are not sure of the name of an archived object with which you wish to work, but you know that, for example, the word orders appears in the associated descriptive text for it, you can use this function to search the text. Press F22 (Search) to display a window from which you can enter search criteria. Type the text that you want to use to search for your object. You can specify up to ten characters. A subset of the whole archive list is displayed, in the same format as above.

Restoring Objects to Regular Libraries

You can restore an archived object to your IBM Power Systems without having to take the object off-line first. Enter Option 2 (Restore) next to the object/library to restore it. There is a delay while the archived object is transferred, during which time your keyboard is locked. When the object has been successfully restored, the archive screen is displayed.

Display Library Archives

This function displays a screen summarizing the list of archived libraries. Press F14 (Libraries) to display the Archived Libraries screen.

The options available on this screen work in the same way as those for the full Archived objects display. Press F14 (Show objects) to toggle between the two types of display.

Displaying Failed Archive Attempts

Operations —> Option 3 (Archiving)—> Option 5 (Failed Archive Attempts)

This function enables you to view a list of all objects that have been selected for archive but not yet taken to off-line media. It tells you what they are, when they were archived, and so on.

The Failed Archive Attempts screen can be useful if you need to restore something quickly, since you can quickly identify where a particular object can be found. Once you have found the desired object(s), you can use this function to restore them to your IBM Power Systems.

Importing Tape

Operations —> Option 3 (Archiving)—>Option 6 (Import Tape)

The import tape feature enables iOptimize to record the relevant details of all your current archived material. It works on one tape at a time, so you can import all of your tapes at your own pace. Once the details have been recorded, you can use iOptimize to simplify and standardize tape management.

Make sure that the tape you want to import is in the tape drive. If the device you are using is not TAP01, enter the correct device name. There is a delay while iOptimize collects and logs the relevant details from the tape. During this time, your keyboard is locked and a series of messages are displayed regarding the import status.

Archiving by Command

Operations —> Option 3 (Archiving)—> Option 9 (Archive by Command)

Archive by Command enables you to archive specific objects.

1. Enter the following information:

Parameter	Value	Description
Object	Name generic* *ALL	Name—Specify the name of the object to be archived. generic—Specify the generic name of the object to be archived. A generic name is a character string that contains one or more characters followed by an asterisk (*). *ALL—All objects
Library	Name *LIBL *CURLIB	Name—The name of the library where the object(s) to be archived are located. *LIBL—All libraries in the library list for the current thread are searched until the first match is found. *CURLIB—Only the current library is searched. If no current library entry exists in the library list, QGPL is used.
Type	*ALL *FILE *PGM *AUTL and more.	Specify *ALL, or the object type for objects to be archived. Types include: *FILE (file), *PGM (program), *AUTL (authorization list), *CMD (command), and more. Press F4 for this field to see entire list.

2. Press F10 to see the following additional parameters:

Parameter	Value	Description
Attribute	*ALL DSPF PRTF DKTF and more.	Specify *All, or the attribute file type. Types include: *ALL , DSPF (Display files), PRTF (Printer files), DKTF (Diskette files) and more. Press F4 for this field to see entire list.
Omit types	*NONE *FILE *PGM *AUTL and more.	Specify *NONE, or the object type to omit. Types include: *FILE (file), *PGM (program), *AUTL (authorization) list, and more. Press F4 for this field to see entire list.

Displaying Subsystem Analysis

Operations —> Option 4 (Subsystem Analysis)

Reviewing the occurrence of IPLs on the system and the activity of starting and stopping of subsystems is a very powerful way to view historical jobs. For example, you may start with an IPL, then see the subsystems started that day, perhaps view the average response time for the subsystem, then drill down to the actual jobs run in that subsystem to see the actual job response times, ultimately weeding out those jobs which may be either causing or suffering particular problems. Using this method you could, for example, single out batch jobs that were running in ostensibly exclusive interactive subsystems.

The following is a description of the inquiry screen columns:

Column	Description
Subsystem	The name of the subsystem started. An entry of *IPL indicates an IPL occurrence.
Start Date/Time	The date and time at which the subsystem was started or the IPL occurred. The list is displayed in reverse date and time sequence.
End Date/Time	The date and time that the subsystem was ended. If this is blank then it indicates that the subsystem is currently active. IPL occurrences have no end date/time.
No. Of Jobs	The number of jobs that were active in the subsystem while it was active.
CPU Time Used	The CPU time used by all of the jobs run in this subsystem while it was active on this occurrence (*HH:MM:SS)*.
Average Response	The average terminal response time for all of the interactive jobs run in this subsystem for this occurrence.

NOTE

To move through the list, use the roll/page keys or specify a date and time in the Position to date field near the top of the screen.

Displaying Subsystem/IPL Analysis History

Operations —> Option 4 (Subsystem Analysis) —> Option 8 (History)

As iOptimize gathers job information, it summarizes and filters out the most useful information . One such item is the occurrence of IPLs on the system and the activity of starting and stopping subsystems (for as far back as you have configured Job Tracking Options to save information). This can form a simple and very powerful hierarchy of viewing historical jobs. For example, you could start with an IPL, then see the subsystems started that day, perhaps view the average response time for the subsystem, then drill down to the actual jobs run in that subsystem to see the actual job response times, thereby weeding out those jobs which may be either causing or suffering particular problems. An example might be to single out batch jobs that were running in ostensibly interactive subsystems.

Entering Option 8 (History) on the Subsystem/IPL Analysis screen displays a Subsystem Activity Enquiry screen.

Option 1 (More Information) displays more subsystem information, which makes it easy to see classes of jobs that should not be running in the subsystem.

Option 5 (Display Job) displays actual job information, including the jobs that ran during the selected occurrence of the subsystem, broken down by the job class. If this option was called from the Subsystem Analysis screen, then only the selected class is displayed. It shows the start and end date and time, the CPU time used (*HH:MM:SS*) and the completion code if the job has completed. Jobs for which a start time is displayed but there is no completion time are still active jobs.

The following is a description of the options available on this screen:

Option	Description
More Information (Option 1)	Displays more detailed job information, such as average response times (interactive jobs), I/Os, and the job source.
Display Job (Option 5)	Displays a job description for the selected job, assuming the job is still on the system.
Spool Files (Option 6)	Displays the spool file output, if any, for the selected job.
Job Log (Option 8)	Displays the job log for the selected job, if it exists.

Displaying Historical WRKACTJOB

Operations —> Option 5 (Historical WRKACTJOB)

When you are looking to see if a particular job is running you would probably use the WRKACTJOB command, an invaluable IBM command. However, if the job has finished you have little or no easily usable information about the job or what other jobs were running at the same time. This Historical WRKACTJOB allows you to pick a point in time and view the main information for a WRKACTJOB command retroactively (for as far back as you have configured Job Tracking Options to save information).

On the Work with Historical WRKACTJOB screen, enter the date and time (Position to date) for which you want to display job information.

A reconstruction of the WRKACTJOB information for the historical date and time entered is displayed. You may be able to see users who signed on multiple times, batch jobs running alongside interactive jobs (perhaps at times when they shouldn't be), whether SNADS was active at that time, and more. Further information can be obtained by displaying the job if it is still active or looking at the joblog for the job if it is completed.

Optimization

9

Optimization Action Codes

On the left side of each list there is an Opt column (underlined) that allows you to enter action codes. Action codes are used to perform an operation or invoke a function for a single item in the list. The valid action codes are:

Action Code	Description
Option 10 (Action)	This code invokes the action for a single item in a list (for example, clear old messages). It is an override function, therefore the days of the week parameters configured for automatic routines and mass action functions (F17 or F18) are ignored. The remainder of the parameters (for example, the number of days old a message must be before being cleared) are adhered to. Note: This action overrides the mode, therefore items are cleared interactively regardless of whether the mode is *NORMAL, *GENFCS or *PRCFCS.
Option 11 (Forecast only)	Creates a forecast for the selected item (for example, forecasts all file reorganization activity to occur over the forecast period). See "Option 8 (View Forecast)" on page 158 for more information.
Option 12 (Process existing forecast)	Similar to Option 10 (Action), this code invokes the action for a single item in a list (for example, clear old messages), but it only works from the previously created forecast, not from the current object status. See "Option 8 (View Forecast)" on page 158 for more information.
Option 1 (Overrides)	Where a list item can have overrides to its action parameters, and these overrides can be themselves subject to a list (for example, a list of individual files within a library, as in Physical file maintenance), this action code displays a sub-list screen.

Action Code	Description
Option 2 (Change)	This action code displays a screen from which you can make changes or overrides to the existing action parameters for the selected item in the list.
Option 4 (Delete)	This action code displays an action window requesting confirmation that you wish to remove the item from the managed list.
Option 5 (View)	Where the items in the list have a view capability, this calls the necessary viewing procedures (for example, Work with output queues (WRKOUTQ)), for the selected item.
Option 6 (History)	This option displays the detail history of management on this item. For example, on output queues this consists of the printer file name, user name, job number and name, how many pages were in the print member, the date and time it was removed from the queue and whether it was archived prior to being deleted. The list begins with details of prints removed in the last run, but can be positioned to a specific date and time or browsed using the Roll/Page keys. The length of time that information is kept is specified using the Maintain Processing Parameters (see "Maintaining Processing Parameters" on page 197).

Action Code	Description
Option 7 (Summary)	This option displays a summary history containing the date and time of the run and, for example, information such as how many print members were removed and how much space (Kb) was saved. If the message is truncated, enter Option 5 (View) to display the full message. If the message shows an error, enter Option 8 (Attributes) to display the error details. Errors usually occur because the print has been printed on paper and removed between a forecast and an action to remove the print from the queue.
Option 8 (View Forecast)	iOptimize can be set to forecast the actions to carry out, up to a given point in the future. The forecast can either be created as part of the normal processing for the action, or can be created on demand. If no forecast is created automatically, or if you believe that the forecast would benefit from an immediate update (for example, to take into consideration newly created prints or newly added logical files), enter Option 11 (Forecast only) next to the list item. iOptimize can be configured to carry out processing against this forecast rather than implementing the actions in real-time. When reviewing forecasts using this option, you can change the forecast before processing. In this way you can view a forecast print using Option 5 (View), review or change its attributes using Option 8 (View Forecasts), which evokes the **WRKSPLFA** command, or remove the print from the current forecast using Option 4 (Delete). Note: Option 5 (View) and Option 8 (Attributes) are only available when working with output queues.

Build Lists

Once you have set up the relevant defaults for an option using F22 (Defaults), you can use iOptimize to directly import into the list all of the system queues and libraries that are necessary for a particular function, without having to type the names individually. Importing is a one-shot build of the list. Maintenance of the list after importing (either when iOptimize is first installed or after F20 (Clear) has been used to clear the list) is carried out for each library or queue either by real-time cross-reference function or manually.

To import the system queues and libraries, press F15 (Import). There is a short wait while iOptimize scans the disk and collects all of the items required for the particular function (for example, all libraries or all job

queues). Until the operation is complete, your keyboard is locked and the following message is displayed at the bottom of the screen:

Building list of all xxxxxxx xxxxxxx on system

As items are added to the list, the Default parameters entered for this function are applied to each item. These can of course be overridden after the list has been completed.

When the task has completed, the first part of the list is displayed on the screen. Use the Page Up or Page Down keys to review the list entries. The lists are usually displayed in alphabetical order.

If you attempt to import system queues and libraries into a list, but have not previously set up the defaults for the list, the following error message is displayed:

You cannot perform an import operation as defaults do not exist.

Press F12 (Cancel) to clear the error message, then follow the procedure for setting up defaults for your list. Once you have set the defaults, Press F15 (Import) to import the system queues and libraries into the list.

Processing Options

There are several ways of processing actions in iOptimize, some requiring manual operations, others needing no manual intervention at all.

Automatic processing

Once you have set up the action parameters, processing takes place automatically according to the rules you specified when configuring iOptimize run times. No further action is required.

Batch processing

To initiate a one-time processing run of an action (for example, your entire output queue list) in batch mode, press F18 (Submit). This submits the job to the job queue that has been specified as the default for your user profile.

Interactive running—entire list

To initiate a one-time processing run of an action (for example, your entire list of libraries containing physical files to reorganize) in interactive mode, press F17 (Run). The keyboard is locked until the job is completed for all items in the list. This processing method is not recommended, as the keyboard is locked until the job is completed.

Interactive running—single queue

To initiate a one-time processing run of selected entries from your list in interactive mode, enter option 10 (Action) in the Opt field next to the item(s) you want to process. There is a short delay while the action is processed.

NOTE

This feature is only available for programs with lists of objects or libraries, such as message queues. It is not available for generic jobs such as history logs or problem logs.

Working with Physical Files and Tables

Optimization —> Option 1 (Physical files and tables)

This section describes the functions available to you when working in the Libraries to Reorganise Files screen (Option 1 Physical files and tables).

When you delete a record from a physical file it becomes inaccessible but the space it took up when active remains allocated, thereby using up disk space that could be used more effectively.

Imagine a file of your company's orders. Some orders are deleted by the Orders Department as part of normal processing operations and you probably have routines to delete orders that have been complete for a pre-defined period of time. Unfortunately, it is unlikely that the order deletion routines do anything more than that. Thus, the longer your order-processing system runs, the more space is taken up by deleted records. Apply this concept to other files in other applications and you will see that the space allocated to deleted records can grow considerably.

iOptimize reorganizes all the files you specify, at the times that you specify, removing deleted records altogether and freeing up the disk space that had been allocated to them. You can also specify a number of other parameters. For example, you can reorganize a file only if a minimum amount of space will be saved by doing so (you specify the amount). You can also reorganize a file into a particular sequence. This can lead to performance improvements by changing the physical sequence of records to the one that is used most often.

The following describes each column available:

Parameter	Description
Library	The name of each physical file library in the list.
ASP	The name of the ASP on which the library resides.
Ovr	An asterisk denotes that there are file level overrides specified within this library.
Auto Gen	Allows iOptimize to automatically generate file level overrides within a library.
S M T ...	The day(s) (Sunday, Monday etc.) on which the actions are to be performed.
Min. Del. Space	The minimum amount of space that should be possible to save in a physical file before it is actioned by iOptimize.
Jrn Files	Whether journaled files should be included in physical file housekeeping.
Use Key	Whether the data records in a file should be reorganized into its physical file key sequence (if any).
Last Action	The date and time at which the last action was performed on a particular library (this is not the date on which the library was checked to see if action was required).

When using the Position to library field, you do not necessarily need to know the exact name of the library. You may type a letter or group of letters to reposition the list beginning with the letter(s) you typed. You can also use the Roll or Page up/Page down keys to move through the library list.

Default Parameters

You can set default action values to be applied to all libraries listed on the Libraries to Reorganise Files screen. Doing this saves you having to enter the same values over and over again for each library. If there are physical file libraries you want processed differently from the majority, you can set their individual values by following the steps in "Adding, Modifying, and Deleting List Items" on page 163. Note that you must set default action values prior to importing existing libraries.

To set default values for your list of physical file libraries press F22 (Defaults). The Defaults for PF Libraries screen is displayed.

The following describes each field on this screen:

Field	Description
Action On	Type 1 next to each day on which you want physical files to be actioned. Type 0 next to each day on which you do not want files to be actioned. The default is daily.
Minimum Space	The minimum amount of space to save in each processed file (in megabytes). This avoids working on files for which only a small saving can be made. It is recommended that you set this field to 0 in the initial stages, which allows iOptimize to schedule the reorganize processes in the most efficient manner. The default is 0.
Do Journaled Files	Specifies whether or not journaled files should be included in physical file housekeeping. You need to be somewhat careful with these because, once a journaled file has been reorganized, it is only possible to use the journals back to the point of reorganization. If you want to include your journaled files and increase the space savings, you should save them immediately after the reorganization. This ensures that, if recovery is necessary, it can be done from the backup copy. Type 0 in this field if you do not want iOptimize to action these files. If you are using journaling solely for audit purposes or are not using journaling, type 1. The default is 0.

Field	Description
Use Physical File Key	You can specify whether, during file reorganization, iOptimize should resequence the data records according to each physical file's key. Reorganizing the records in a file into the sequence of the most frequently used key improves file access performance. Type 0 in this field if you do not want to use physical file keys, otherwise type 1. The default is 1. If you have unkeyed physical files and rely on your logical files for access purposes, you can still have the physical data organized into the key sequence of one of the associated logicals. Instructions on how to do this can be found under "Physical File Overrides" on page 165.
Allow Auto-Overrides	Entering 1 allows iOptimize to automatically generate file level overrides within a library. Entering 0 does not allow auto-overrides. The default is 1. Instructions on working with overrides can be found under "Physical File Overrides" on page 165.

Press F8 (Update) to save your changes and return to the previous screen.

Adding, Modifying, and Deleting List Items

To add a new library to the list press F6 (Add). The Add Library to List screen is displayed.

The displayed values are those that have been set up as defaults. Type the name of the library you want to add in the Library field. You can override the default parameters (see "Default Parameters" on page 161 for an explanation of each parameter).

To modify a list item, enter Option 2 (Change) next to the library name. A screen similar to the Add Library to List screen is displayed. Modify the parameters as required (see "Default Parameters" on page 161 for an explanation of each parameter). Note that, when modifying, there is no reference to Dynamic Maintenance.

To remove individual entries from the list of physical file libraries, so that they are no longer included in processing runs, enter Option 4 (Delete) next to the library. A confirmation prompt is displayed. Press F8 to confirm that you wish to remove the library from the list.

History

iOptimize maintains a history of all actions taken on physical files at both a summary and a detail level. The length of time that history information is held is determined by the Maintain Processing parameters. This is described in detail in "Maintain Autorun Parameters" on page 192.

Entering Option 6 (History) next to any queue provides a detail history of management on this queue. Entering Option 7 (Summary) next to the queue displays a summary history.

Forecasting

iOptimize forecasts the files it intends to reorganize, up to a given future point. The forecast can be either created as part of the normal processing for output queues or can be created on demand.

- Option 8 (View Forecast)— displays the current forecast, if one exists.
- Option 11 (Forecast only)— generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast) —works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing physical file reorganization, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the reorganize process (action) for the files in the selected library only.
- F17 (Run)—actions all of the files in all of the libraries in the list interactively.
- F18 (Submit)—actions all of the files in all of the libraries in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines, as this is where the best time savings are made and the jobs are run frequently and regularly. In addition, when iOptimize is working in automatic mode, it ensures that the process does not run during your busy times. Any manual initiation of the reorganization routines will run until they are completed.

Other Functions

- F20 (Clear) —removes all of the libraries from the managed list, leaving it ready for you to re-import the details. This is particularly useful if you wish to reset the defaults for the whole system.
- F22 (Defaults)—used to reset the default values for the list of physical file libraries. See "Default Parameters" on page 161 for more information.
- F15 (Import)—imports all the libraries on your system into the list, without the need for you to type all the names in individually. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function. There is a short delay while iOptimize collects all the libraries.

  NOTE

  If you are using the dynamic maintain function, libraries not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  See "Understanding Processing Modes" on page 44 for more information.

Physical File Overrides

The overrides feature allows you to work at individual file, rather than library, level. From the Libraries to Reorganize Files screen, press F23 (All Overrides). Enter Option 1 (Overrides) next to the library containing the file(s) you wish to override. The File Overrides for Library screen is displayed.

The first time you select this option for a library, the screen is blank. This is because iOptimize normally works across an entire library, and only maintains a list of files that are being overridden.

The following describes each column:

Column	Description
Opt	You can perform various functions on the physical file libraries from the table at the top of the screen. If more than one file is selected, iOptimize displays each selection upon exiting the previous one.
File	The name of each overridden file in the list.
Lck	Displays a 1 if the override is locked.
S M T ...	The day(s) on which the actions are to be performed.
Min. Del. Space	The minimum amount of space that must be possible to save in a physical file before it will be actioned by iOptimize.
PF Key	Whether the physical file key should be used in the file override or the access path.
Jrn File	Whether this journaled file should be included in physical file housekeeping.
Logical File	The name of a logical file, based on the physical file, from which the key should be used by iOptimize to sequence the physical data records.
Logical Library	If a logical file is being used to sequence the records in a physical, the library in which that logical file resides.
Force Days	The maximum number of days, after the last time that a physical file was actioned, that can pass before it must be processed again (for example, an entry of 1 will force this file to be reorganized every day).

Adding Physical File Overrides

To add a new file to the list, from the File Overrides for All Libraries screen, press F6 (Add). The Add File Override screen is displayed.

The following is a description of each field:

Field	Description
File	The name of the file to be overridden.
ASP Name	The name of the ASP on which the file resides.

Field	Description
Lock Override	Enter 1 in this field to prevent library-level overrides from affecting this file.
Action On	See "Default Parameters" on page 161 for information.
Minimum Space	See "Default Parameters" on page 161 for information.
Journaled Files	See "Default Parameters" on page 161 for information.
Force Days	The frequency with which a file must be actioned (in number of days), even if it would not qualify under normal circumstances. The benefit of this is that the file will get the performance benefit from being resequenced even though it does not have any deleted records to remove (for example, transaction history files, sales analysis files). For example, if you type 2 in this field, iOptimize actions the file every two days. **NOTE**: To restrict a file from being reorganized, set the action days to 0. Force days of 999 does not stop files from being reorganized.
Physical File Key	Specify whether, during file reorganization, iOptimize should resequence the data records according to each Physical File's key. Enter 1 if you want this to be actioned. Enter 0 if you do not want this to be actioned.
Use Access Path Library	As outlined earlier, file access times can be dramatically improved by having the data records kept in the sequence of the most commonly used access path. You can name a logical file for iOptimize to use during file reorganization, for resequencing the data according to its key. In general you would set iOptimize to automatically monitor for the best access path and gain the benefits directly. However, this is based on I/O seniority of logicals over a physical and you may wish to override this for particular files. Type the name of the logical file that you want iOptimize to use in the Access path field. Type the name of the library in which the logical file is located in the Library field. **NOTE:** You can not configure iOptimize to action both physical file key and have an access path.

Modifying Physical File Overrides (Option 2 Change)

To modify the action values for an overridden physical file enter Option 2 (Change) next to the file you want to modify. A window similar to the Add window is displayed. Rules for changing the parameters are the same as for adding a new file override.

Removing Physical File Overrides (Option 4 Delete)

You may remove individual entries from the list of overridden files by entering Option 4 (Delete) next to the file to be removed from the list. A window similar to the Add window is displayed, prompting you to confirm that you wish to remove the file from the list. Press F8 to confirm.

Viewing Physical File Data (Option 5 View)

You can select files from the list in to look at the data records in them. To do this, enter Option 5 (View) next to file. The IBM Display Physical File Member (DSPPFM) screen is displayed. For more information about this panel, refer to the appropriate IBM reference manual.

Viewing Logical Files (Option 7 Logicals)

iOptimize supports a wealth of system information that can be used to present difficult-to-access information when needed. You can gain information about the logical files that are associated with each of the physical files, to assist in selecting the logical file to override. To do this, enter Option 7 (Logicals) next to the selected file. The Display Dependencies screen is displayed.

Under Dependencies all the logical files are listed, based on the selected physical, together with the library in which each resides. Following is a description of the options on this screen.

Option	Description
K = Keys	Displays the key to the selected file.
L = Layout	Displays a screen containing the field layout of the selected file.
S = Select	You can use this window to select a logical file whose key iOptimize will use to sequence the data when reorganizing the physical file. This merely attaches the selected logical file to the instruction to reorganize the file.

Monitoring Access Paths

iOptimize can automatically monitor, on an hourly basis, which access paths are being used over which physical file and assess the I/O usage of each. iOptimize then switches the physical sequence of the records into the sequence of the most used logical file, thereby improving access times and responses.

To allow iOptimize to change the physical file data to the optimum sequence, on the Libraries to Reorganise Files screen enter 1 in the Monitor access paths? field. Otherwise, enter 0. The default value is 1.

NOTE

The Task Manager job, which normally runs for a few seconds each hour, could now be run for a few minutes, depending on the size and the loading of your system. If this impact becomes noticeable, return to this screen and set the Monitor access paths? field to 0 to stop iOptimize from gathering usage information.

Jobs to Monitor

Enter the number of most commonly used access paths monitored to qualify for file-level overrides.

An asterisk (*) in the Ovr column indicates that there are file-level overrides specified within that library.

Working with Logical Files and Indexes

Optimization —> Option 2 (Logical files and indexes)

iOptimize has as much to do with optimizing system performance as with space saving. It is often the case that logical files are created at different times during the development of software applications. This usually means that optimum use of shared access paths is not made, which, in turn, has a detrimental effect on both system performance and disk usage. In addition, there are many instances where a logical access path is being maintained interactively (*IMMED) where the logical file is either never used or only used occasionally.

The product optimizes logical file access paths by looking at all logical files, determining where access paths could be shared for greater efficiency, and regenerating them to do so. It also releases access paths from immediate update where not required by the system at this time. With fewer access paths being maintained, your system has less work to do when performing

file updates, and benefits from considerable disk space savings and performance improvement.

The following describes each column:

Column	Description
Library	The library name.
ASP	The name of the ASP on which the library resides.
Ovr	An asterisk denotes that there are file level overrides specified within this library.
S M T ...	The day(s) on which the actions are to be performed.
Share Acc	This parameter should always be set to 1 for all libraries that can receive some dedicated processing time. It indicates that iOptimize will optimize access paths (for example, when one logical file is subordinate to another).
Days *REBLD	This is the number of days that a logical file can be maintained at full *IMMED maintenance status without even being opened. Set to 180 it will normally save many access paths. **NOTE**: If you start to use an access path using a logical file that has been changed to *REBLD by iOptimize, it will be reset to *IMMED next time it runs.
Last Action Date and Last Action Time	The date and time at which the last action was performed on a particular library.

Default Parameters

You can set default action values to be applied to all libraries in your list to action. Doing this saves you having to key in the same values over and over for each of your libraries. If there are some libraries that you want processed differently from the majority, you can set their individual values by following the steps explained in this section. Note that if you intend to import your existing libraries, you must have set your default action values.

To set default values for your list of libraries, press F22. The Defaults for LF Libraries screen is displayed.

The following is a description of each field:

Field	Description
Action On ...	Type 1 next to each day on which you want logical file libraries to be actioned. Type 0 if you do not want them to be actioned. The default is daily.
Force Share Of Access Paths	Set this entry to 1 (the default).
Change Maintenance To *REBLD After Days Unused	Enter the number of days of non-use before *IMMED maintained logicals are set to *REBLD. **NOTE**: This is dependent on your own company philosophy, however remember that access paths can be even bigger than the physical file on which they are based. Additionally, they take a great deal of space and they can impact your system performance.

Adding List Items (F6 Add)

To add a new library to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new library in the Library name field. For information on the remaining fields, see "Default Parameters" on page 170.

Modifying List Items (Option 2 Change)

You can modify the action values of libraries in the list very simply by entering Option 2 (Change) next to the library to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of libraries, so that they are no longer included in processing runs. To do this enter Option 4 (Delete) next to the libraries) to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the library from the list. Press F8 to confirm.

History

iOptimize maintains a history of all actions taken on logical file libraries at both a summary and a detail level. The length of time that history information is held is determined in the Maintain Processing parameters (see "Maintaining Processing Parameters" on page 197). Enter Option 6

(History) next to any library to display a detail history of management on this library. Enter Option 7 (Summary) next to a library to display a summary history.

File Overrides

The overrides feature allows you to work at individual file, rather than library, level. From the Libraries to Action Logical File screen, do one of the following:

- Enter Option 1 (Overrides) next to the library containing the file(s) you wish to override. The File Overrides for Library screen is displayed.
- Press F23 (All Overrides) to work with the files in all libraries. The File Overrides for Library screen is displayed.

The first time you select this option for a library, the screen is blank because iOptimize normally works across an entire library, and only maintains a list of files that are being overridden.

The following describes each column.

Column	Description
Lib	The library in which the file resides. • An asterisk (*) denotes no, one, or many characters up to the maximum available. • A percent sign (%) denotes a single numeric character. • A question mark (?) denotes a single alphanumeric character. NOTE: Only available if displaying overrides for all libraries.
ASP	The name of the ASP on which the file resides. NOTE: Only available if displaying overrides for all libraries.
File	The name of each overridden file in the list. • An asterisk (*) denotes no, one, or many characters up to the maximum available. • A percent sign (%) denotes a single numeric character. • A question mark (?) denotes a single alphanumeric character.
SMTWTFS	The day(s) on which the actions are to be performed.

Column	Description
Share	If set to 1, the library can receive some dedicated processing time. It indicates the iOptimize will optimize access paths (for example, when one logical file is subordinate to another). If set to 180 it will normally save many access paths.
***REBLD**	The number of days that a logical file can be maintained at full *IMMED maintenance status without being opened.

Adding Logical File Override

To add a new file to the list, press F6 (Add). The Add File Override screen is displayed.

The following is a description of each field:

Field	Description
File	The name of the file to be overridden.
Library	The library in which the file resides.
ASP	The name of the ASP on which the file resides.
Action on	Do one of the following: • Type 1 next to each day on which you want physical files to be actioned. • Type 0 next to each day on which you do not want files to be actioned. The default is daily.
Share access paths	Set to 1 for all libraries that can receive some dedicated processing time. It indicates the iOptimize will optimize access paths (for example, when one logical file is subordinate to another).
***REBLD days**	The number of days that a logical file can be maintained at full *IMMED maintenance status without even being opened. Set to 180 it will normally save many access paths. NOTE: If you do start to use an access path using a logical file that has been changed to *REBLD by iOptimize, it will be reset it to *IMMED next time it runs.

Forecasting

iOptimize forecasts the logical files it intends to change, up to a given future point. The forecast can be either created as part of the normal processing for

output libraries or can be created on demand. When a forecast is rerun, all logging records are removed.

- Option 8 (View forecast)—displays the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for more information.

Additional Processing

There are several ways of processing logical file library actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 9 (Logging)—displays the results from running Option 11 (Forecast). If running in normal mode, this screen is updated with the next forecast. When a forecast is rerun, all logging records are removed. Use Option 1 (Full Message) on this screen to display the complete message.
- Option 10 (Action)—runs the process (action) for the selected library only.
- F17 (Run)—actions all of the libraries in the list interactively.
- F18 (Submit)—actions all of the libraries in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the libraries from the managed list, leaving it ready for you to re-import the details. This is particularly useful if you wish to reset the defaults for the whole system.
- F22 (Defaults)—used to reset the default values for the list of Physical File libraries. See "Default Parameters" on page 170 for more information.
- F15 (Import)—imports all the libraries on your system into the list, without the need for you to type all the names in individually. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using

the F15 (Import) function. There is a short delay while iOptimize collects all the libraries.

NOTE

If you are using the dynamic maintain function, libraries not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  See "Understanding Processing Modes" on page 44 for more information.

Working with Programs and Modules

Optimization—> Option 3 (Programs and Modules)

Two-thirds of the space taken up by a compiled program is dedicated to the symbol table. This is only used when debugging programs (usually during the development phase). Also, when a program crashes, you get more information if you have the symbol table than if you do not. Once development is over, it is highly likely that this will not be used again and so that space, or part of it, could be released (the symbol table can either be removed or just compressed).

There are also many programs on your system that are used infrequently or not at all. You could identify and remove these objects, but they can be reduced in size by 50 percent by automatically compressing them here. In addition, certain optimization measures can be taken to improve the performance of compiled programs without the need to reference the source.

By specifying particular libraries and a number of days since compilation, you can configure iOptimize to work on your application programs, gradually freeing disk space and improving their efficiency.

NOTE

Before beginning optimization, validate with the program vendor that the product can be optimized. Precisely's program objects should not be optimized.

iOptimize can save considerable space and result in more efficient application programs. Working in a strict sequence, iOptimize could:

- Optimize your programs
- Compress the symbol table (this could be done even during development to save disk space)
- Remove the symbol table altogether (this could be done a number of days after program acceptance, saving even more disk space)
- Compress your programs to save yet more space (you can specify the programs to be compressed, as compression of often-used programs means that, each time they are run, they need to be expanded, a process that may take up to a few seconds but is generally unnoticeable).

NOTE

iOptimize uses only the IBM compression routines. Therefore it is actually IBM i that is forced to compress an object. This means that, without any intervention from iOptimize or even the need for iOptimize to be on the system, compressed objects that are called are automatically and instantaneously expanded for immediate use by IBM i. This process is extremely quick and is usually unnoticeable.

Following is a description of each column:

Column	Description
Library	The library name.
ASP	The name of the ASP on which the library resides.
S M T ...	The day(s) on which the actions are to be performed.
Optim	The number of days after their creation that programs will be optimized. A value of 1 is recommended.
Compress Symbol Table	The number of days after their creation that program symbol tables will be compressed. A value of 1 is recommended.

Column	Description
Cmp	The number of days that must have elapsed since a program was last used for it to be compressed. A value of 37 is recommended.
Remove Symbol Table	The number of days after their creation that programs should have their symbol table removed. A value of 999 is recommended.
Last Action Date	The date on which the last action was performed on a particular program library.
Last Action Time	The time at which the last action was performed on a particular program library.

Setting Default Parameters

You can set default action values to be applied to all libraries on the Libraries to action programs screen. Doing this saves you having to key in the same values over and over for each of your libraries. If there are some program libraries that you want processed differently from the majority, you can set their individual values by following the steps explained in this section. Note that if you intend to import your existing libraries, you must have set your default action values.

To set default values for your list of libraries, press F22 (Defaults). The Defaults for Program Libraries screen is displayed.

Following is a description of each field:

Field	Description
Action On ...	Type 1 next to each day on which you want programs to be actioned. Type 0 if you do not want them to be actioned. The default is daily.
Optimize Days	You can configure iOptimize to optimize your programs a number of days after their creation. Optimization reorganizes the program code to gain maximum efficiency. The time period could be quite short (the default is 1 day). All programs will benefit from optimization, though the benefits can only be seen when all or most of the programs have been optimized. The amount of non-database paging will reduce (WRKSYSSTS).

Field	Description
Compress Symbol Table	As with optimization, you can specify a number of days after creation for program symbol tables to be compressed. This could be done immediately, even in the case of programs in the development phase, as the overhead in expanding the symbol table when you want to work on a program (DEBUG) is negligible and the space savings are considerable. The default is 1.

Field	Description
Compress Program	Again, you can specify a time period for the compression of programs. However, in this case, it is expressed in terms of the number of days since a program was last used. This allows you to avoid compressing production programs that are in frequent use. Typically you would set this to a number between 37 and 180, to compress programs which are clearly not being used or not used regularly. The time taken by i5/OS to expand a program is negligible. You will compress many hundreds of programs to half their original size and expand only a handful.
Remove Symbol Table	The complete removal of program symbol tables can save you a considerable amount of disk space (remember, a symbol table accounts for about two-thirds of the space taken up by a compiled program). Importantly, the symbol table is NOT used when the program runs normally. However, the table is required when debugging or dumping the program. Many users know obliquely of the symbol table as the “observability” of the program. Removing the symbol table (Observability) from a program does not mean that it is lost forever. If you should need to debug a program in the future, simply regenerate the table is to recompile the program. So, if you have the source (or a third party maintains your source and can make it available for debug etc.) you can safely remove the symbol table. Significantly, IBM required the symbol table to be present during conversion from CISC to RISC. For this reason, and as we have no advance warning of what IBM may need of this information in the future, it is recommended that you compress the symbol table rather than delete it, unless you have easy access to recreating the table. If not then just compress the symbol table, a saving of half the space, by setting this parameter to 999 (the default). Once again, you can specify how many days after program creation you want this to be done. Proven programs could be actioned immediately, while those being developed could wait until all the bugs have been ironed out.

Adding List Items (F6 Add)

To add a new library to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those

that have been set up as defaults. Type the name of the new library in the Library name field. For information on the remaining fields, see "Setting Default Parameters" on page 177.

Modifying List Items (Option 2 Change)

You can modify the action values of libraries in the list very simply by entering Option 2 (Change) next to the library to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section except that, when modifying, there is no reference to dynamic maintenance.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of program libraries, so that they are no longer included in processing runs. To do this enter Option 4 (Delete) next to the libraries) to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the library from the list. Press F8 to confirm.

Building Lists

iOptimize enables you to directly import all the libraries on your system into the list, so you do not need to type all the names in individually. To do this press F15 (Import). There is a short wait while iOptimize collects all the libraries. Until the import operation is complete, your keyboard is locked.

Displaying History

iOptimize maintains a history of all actions taken on logical file libraries at both a summary and a detail level. The length of time that history information is held is determined in the Maintain Processing parameters (see "Maintaining Processing Parameters" on page 197). Enter Option 6 (History) next to any library to display a detail history of management on this library. Enter Option 7 (Summary) next to a library to display a summary history.

Forecasting

iOptimize forecasts the programs it intends to change, up to a given future point. The forecast can be either created as part of the normal processing for output libraries or can be created on demand.

- Option 8 (View forecast)—displays the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.

- Option 12 (Process existing forecast)—Works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing program libraries in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the process (action) for the selected library only.
- F17 (Run)—actions all of the libraries in the list interactively.
- F18 (Submit)—actions all of the libraries in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the libraries from the managed list, leaving it ready for you to re-import the details. This is particularly useful if you wish to reset the defaults for the whole system.
- F22 (Defaults)—used to reset the default values for the list of physical file libraries. See "Setting Default Parameters" on page 177 for more information.
- F15 (Import)—imports all the libraries on your system into the list, without the need for you to type all the names in individually. Both F20 (Clear) and F22 (Defaults) options need to be processed before using F15 (Import).

NOTE

If you are using the dynamic maintain function, libraries not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only

– *NORMAL—Action items directly

More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Working with Temporary Storage

Optimization—>Option 4 (Temporary Storage)

Temporary storage is space allocated by the operating system during general processing and when creating temporary objects (for example, expanding a previously compressed object). This space is only reclaimed when managed or when running a full Reclaim Storage command.

The temporary storage used depends on the processing in any particular environment, but can quite easily take up well over one percent of your disk. iOptimize enables you to automatically manage temporary storage, by checking it as regularly as you wish and releasing the storage to available free space.

The following describes each field:

Field	Description
Library	The library name.
ASP	The name of the ASP on which the library resides.
S M T...	The day(s) on which the actions are to be performed. The default is daily.
Days Unused	The maximum amount of time you want the space held before iOptimize releases it to free space. There is no detriment to releasing this space. A value of 1 to 3 is recommended.
Last Action Date	The date on which the last action was performed on a particular library.
Last Action Time	The time at which the last action was performed on a particular library.

Default Parameters

You can set default action values to be applied to all libraries in your list of libraries to action. Doing this saves you having to key in the same values over and over again for each of your libraries. If there are some libraries that you want processed differently from the majority, you can set their individual values by following the steps explained later in this section. Note

that if you intend to import your existing queues, you must have set your default action values.

To set default values for your list of libraries, press F22 (Defaults). The Temporary storage library defaults screen is displayed.

The following describes each field:

Field	Description
Action On ...	Type 1 next to each day on which you want libraries to be actioned. Type 0 if you do not want them to be actioned.
Days Unused	The maximum number of days you want temporarily allocated space to remain in the library.

Adding List Items (F6 Add)

To add a new library to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new library in the Library field. For information on the remaining fields, see "Default Parameters" on page 182.

Modifying List Items (Option 2 Change)

You can modify the action values of libraries in the list very simply by entering Option 2 (Change) next to the library to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of libraries, so that they are no longer included in processing runs. To do this enter Option 4 (Remove) next to the libraries to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the queue from the list. Press F8 (Update) to confirm.

Displaying History

iOptimize maintains a history of all actions taken on libraries at a summary level. Enter Option 6 (History) to display a detailed history or Option 7 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the space to retrieve, up to a given future point. The forecast can be either created as part of the normal processing for libraries or can be created on demand.

- Option 8 (View Forecast)—shows the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected library interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Additional Processing

There are several ways of processing library actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- Option 10 (Action)—runs the process (action) for the selected library only.
- F17 (Run)—actions all of the libraries in the list interactively.
- F18 (Submit)—actions all of the libraries in the list by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs are run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the libraries from the managed list, leaving it ready for you to re-import the details. This is particularly useful if you wish to reset the defaults for the company.
- F22 (Defaults)—used to reset the default values for the list of libraries to reclaim storage. See "Default Parameters" on page 182 for more information.
- F15 (Import)—imports all the libraries on your system into the list, so you do not have to type all the names individually. Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function.

NOTE

If you are using the dynamic maintain function, libraries not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Working with Spool Members

Optimization—>Option 5 (Spool Members)

Whenever any print output is generated, the data to be printed is held in a physical file member. Once the data has been printed (or the print queue entry deleted), the member is cleared by the system and is then available for re-use.

If a large number of prints build up on an output queue, then the number of members increases according to the total number of prints on the queue. As each queue entry is printed or deleted, the corresponding file member is cleared. However, the members, once cleared of data, are not removed from the file in library QSPL.

The file holding spool members contains the number of members corresponding to the maximum number of prints that have ever been resident, at one time, on your IBM Power Systems. Each remaining member, although containing no physical data, will still be allocated the amount of disk space necessary for it when it was created. Thus, depending on the maximum number of prints ever held on your system, a considerable amount of disk space could be allocated unnecessarily.

iOptimize keeps the number of spool members to a minimum by checking the file and getting rid of empty members as regularly as you specify.

The following describes each field.

Field	Description
Days To Keep	The number of days the unused members must be for them to be actioned and cleared by iOptimize. The default is 5.
S M T ...	The day(s) on which the actions are to be performed. Enter 1 next to each day on which you want the members to be actioned; enter 0 next to each day on which you do not want them to be actioned. The default is daily.

Displaying History

iOptimize maintains a history of all actions taken on spool members at a summary level. Press F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts the spool members it intends to remove, up to a given future point. The forecast can be either created as part of the normal processing for the log. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing spool members in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the clearing interactively.
- F18 (Submit)—performs the clearing by submitting a job to the default batch queue.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs are run frequently and regularly.

Other Functions

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Working with Oversized Objects

Optimization—>Option 6 (Oversized Objects)

Over time, many different types of objects (for example, libraries, message queues, job queues) get bigger than they really need to be, as they grow to accommodate current requirements. Using libraries as an example, as the number of objects in them increases, so their size has to be extended. Even if objects are then removed (as would be the case, say, of a library used for the development of application programs) the library size is not automatically reduced by the system. Thus, a library remains as large as the largest size it has ever been.

By specifying how and when you want oversized objects to be cleaned up, iOptimize checks on all relevant objects regularly and reduces any that are too large back down to the minimum size necessary, thereby releasing disk space for more effective use.

Following is a description of each field.

Field	Description
Days To Keep	The number of days since last updated that an object must be for it to be assessed and actioned by iOptimize. The default is 3.
S M T ...	The day(s) on which the actions are to be performed. Enter `1` next to each day on which you want the objects to be actioned; enter `0` next to each day on which you do not want them to be actioned. The default is daily.

Displaying History

iOptimize maintains a history of all actions taken on objects at a summary level. Pressing F20 (History) to display a detailed history or F21 (Summary) to display a summary history.

Forecasting

iOptimize forecasts objects to remove, up to a given future point. The forecast can be either created as part of the normal processing for the job. See "Maintaining Processing Parameters" on page 197 for further detail on creating and processing from forecasts.

Press F22 (Forecasts) to display the current forecast, if one exists.

Additional Processing

There are several ways of processing objects for resize in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:

- F17 (Run)—performs the resizing interactively.
- F18 (Submit)—performs the resizing submitting a job to the default batch queue.

  While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.
- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on processing modes can be found in "Understanding Processing Modes" on page 44.

Working with Device Files, Menus & Panel Groups

Optimization—> Options 7-9 (Device Files/Menus/Panel Groups)

Similar to programs, most systems have many device files, menus and panel groups that are seldom or never used. These should, at a minimum, be compressed.

iOptimize compresses all device files, menus and panel groups that haven't been used for a period of time that you specify, thereby saving disk space by reducing their actual size. Of course, as they are used, they expand again, but running iOptimize regularly keeps them in compressed format for most of the time that they are not in use. Expansion times are extremely fast and, in most cases, expansion by i5/OS is unnoticeable. However, you will find that very few of these objects even get expanded once.

The following describes each field:

Field	Description
Library	The name of the library in which the objects reside.
ASP	The name of the ASP on which the library resides.
S M T...	The day(s) on which the actions are to be performed.
Days Unused	The number of days that entries in a library have not been used before they require action by iOptimize.
Last Action Date & Time	The date and time on which the last action was performed on a particular library.

Setting Default Parameters

You can set default action values to be applied to all libraries in your list of libraries to action. Doing this saves you having to key in the same values over and over again for all your libraries. If there are some libraries that you want processed differently from the majority, you can set their individual values by following the steps explained later in this section. Note that if you intend to import your existing libraries, you must have set your default action values.

To set default values for your list of folders, press F22 (Defaults). The Device file library defaults screen is displayed.

The following describes each field:

Field	Description
Action On ...	Type 1 next to each day on which you want libraries to be actioned. Type 0 if you do not want them to be actioned. The default is daily.
Days Unused	iOptimize lets you specify how long to wait before compressing each entry in the library. In the Days unused field, type the maximum number of days that objects can remain unused in their expanded form. Set this value between 37 and 180 days.

Adding List Items (F6 Add)

To add a new library to the list press F6 (Add). A window similar to the Default parameters window is displayed. The values displayed are those that have been set up as defaults. Type the name of the new library in the

Library name field. For information on the remaining fields, see "Setting Default Parameters" on page 189.

Modifying List Items (Option 2 Change)

You can modify the action values of libraries in the list very simply by entering Option 2 (Change) next to the library to be modified. A window similar to the Add List Items window is displayed. The steps for modifying the action values for a library are similar to those explained in the previous section.

Removing List Items (Option 4 Remove)

You may remove individual entries from the list of libraries, so that they are no longer included in processing runs. To do this enter Option 4 (Delete) next to the libraries to be removed from the list. A window similar to the Add List Items window is displayed, asking you to confirm that you wish to remove the library from the list. Press F8 to confirm.

Displaying History

iOptimize maintains a history of all actions taken on libraries at a summary level. The length of time that history information is held is determined in the Maintain Processing parameters (see "Maintaining Processing Parameters" on page 197). Enter Option 6 (History) or Option 7 (Summary) next to a library to display the history.

Forecasting

iOptimize forecasts the objects it intends to action, up to a given future point. The forecast can be either created as part of the normal processing for libraries or can be created on demand.

- Option 8 (View forecasts)—shows the current forecast, if one exists.
- Option 11 (Forecast only)—generates a forecast for the selected item interactively.
- Option 12 (Process existing forecast)—works from the forecast instead of the actual current position.

More detail on forecasting and other additional options are explained in "Maintaining Processing Parameters" on page 197.

Additional Processing

- There are several ways of processing library actions in iOptimize, some requiring manual operations, others needing no manual intervention at all. The manual options are:
  - Option 10 (Action)—runs the process (action) for the selected library only.
  - F17 (Run)—actions all of the libraries in the list interactively.
  - F18 (Submit)—actions all of the libraries in the list by submitting a job to the default batch library.

While these functions are always available, it is suggested that you set the parameters to use the automatic routines as this is where the best time savings are made and the jobs will be run frequently and regularly.

Other Functions

- F20 (Clear)—removes all of the libraries from the managed list, leaving it ready for you to re-import the details. This is particularly useful if you wish to reset the defaults for the company.
- F22 (Defaults)—used to reset the default values for the list of libraries. See "Setting Default Parameters" on page 189 for more information.
- F15 (Import)—imports all the libraries on your system into the list, without the need for you to type all the names in individually. (Libraries that have been excluded from optimization are not imported. See "Excluding Libraries from Cross-Referencing and Optimization" on page 22 for more information.) Both F20 (Clear) and F22 (Defaults) functions need to be processed before using the F15 (Import) function. There is a short delay while iOptimize collects all the libraries.

  NOTE

  If you are using the dynamic maintain function, libraries not already in the list (including those that you may have deleted from the list) are automatically found and included in the managed list.

- F24 (Mode)—allows you to switch the processing method for this option. The processing methods are:
  - *GENFCS—Forecast only
  - *PRCFCS—Action previous forecasts only
  - *NORMAL—Action items directly

  More detail on the processing modes can be found in "Understanding Processing Modes" on page 44.

Managing Optimization

Optimization —> Option 10 (Manage Optimization)

From the Managing Optimization window, you can do the following:

- Specify how often you wish various tasks to be performed on certain objects (see "Maintain Autorun Parameters" on page 192).
- End submitted jobs (see "Ending Submitted Jobs" on page 198).
- Set up automatic processing (see "Setting Processing Windows" on page 198).

Maintain Autorun Parameters

Manage Optimization —> Option 1 (Maintain Autorun Parameters)

iOptimize enables you to specify how often you wish various tasks to be performed on certain objects. iOptimize then submits the relevant jobs at the specified times. Once you have specified when and where to submit the jobs, iOptimize manages the housekeeping tasks automatically.

iOptimize breaks down daily processing time into three time zones, based on system usage: Prime time, Off-peak time and Unused time. These time zones enable you to manage a shortage of dedicated processing times, by scheduling jobs to run during periods of lower system usage (such as at night or on weekends). The time zone names can be modified, using F19 (Control Parameters). The following is an explanation of each time zone:

- Prime time (P)—Users are signed on and performance is at a premium. You would not want to run any housekeeping tasks that might cause conflicts (for example, Physical file reorganize) during this period. Some jobs that require minimal processing time could be run (for example, clearing the alert database).
- Off-peak time (O)—Most users are off the system, and perhaps the system is being backed up. Clearing output queues could be scheduled during this period, but, again, reorganizing files might result in certain files not being saved.
- Unused time (U)—Users are off the system and all batch jobs have finished. The system can run dedicated jobs such as file reorganizations.

NOTE

iOptimize ships with all housekeeping tasks pre-set to be run in unused time, using the job queue that you specified during product installation.

On some lines an asterisk is displayed to the left of the Task column, indicating that caution is recommended. This means that the function can take some time to complete, therefore running it during Prime or Off-peak periods is not recommended. If you still choose to run it during those time periods, do so with extreme caution.

In addition to scheduling jobs to run in different time zones, you can schedule jobs to run only on specific days for specific libraries (for more information, see the instructions for the actual function). By combining these parameters you could, for example, run logical file optimization for all libraries except the DESPATCH library (in the warehouse) on weekdays in Off-peak time but run the function for the DESPATCH library in unused time on Sunday only.

Use the Roll or Page up and Page down keys to move through the list of housekeeping tasks. The following is a description of each column on this screen.

Option	Description
Task	The name of the housekeeping task to be run.
Prime Queue/Lib	This is the Prime time zone column. If there is an entry in this column (a valid job queue name and library), the job in the task column is eligible to run during the Prime time zone.
Off Peak Queue/Lib	This is the Off peak time zone column. If there is an entry in this column (a valid job queue name and library), the job in the task column is eligible to run during the Off-peak time zone.
Unused Queue/Lib	This is the Unused time zone column. If there is an entry in this column (a valid job queue name and library), the job in the task column is eligible to run during the Unused time zone (default).

Switching Time Zones

If you can allocate some unused time to iOptimize during evenings, overnight, or on weekends, then just accept the defaults and leave all jobs to run in the Unused time zone.

To enable a job to be run in a particular time zone, move the cursor to the job queue field in the desired processing period column. Type the name of the job queue you wish to use on the first line. Type the name of the library in which the job queue resides on the second line. Press F8 (Update) to save your changes.

Viewing and Changing Command Details

Enter Option 1 (Command details) next to a task to review the command details associated with that task. The iOptimize Submit Job (OSDSBM) screen is displayed. The OSDSBM command includes the following configurable parameters:

JOB(*JOBD)
JOBQ(*JOBD)
JOBD(*USRPRF)
MSGQ(*USRPRF)
INLLIBL(*CURRENT)
PRTTXT(*CURRENT)
JOBPTY(*JOBD)
USER(*CURRENT)
INQMSGRPY(*JOBD)
SWS(00000000)
OUTPTY(*JOBD)
OUTQ(*CURRENT)
SCDDATE(*CURRENT)
SCDTIME(*CURRENT)

Configuring Action Parameters

Enter Option 2 (ACTxx Parameters) next to a task to review or configure action parameters related to that task.

Configuring Multi-Threading Parameters

Enter a P (Prime), O (Off-peak) or U (Unused) next to a task to initiate a multithreading procedure. This procedure enables you to specify a number of job queues that may be used, for example, to reorganize the files in a number of libraries simultaneously within that time frame. The Multi-Thread Processing Parameters screen is displayed:

The increment factor is used to calculate how much time is required to complete that task. The default is 10%.

To select a subsystem and job queue, press F4 (JobQs). A window is displayed, from which you can select a subsystem. After selecting the subsystem, press Enter. The Select Subsystem screen is displayed, from which you can select a job queue:

The Max Act column indicates the maximum number of jobs that job queue can run simultaneously.

After selecting a job queue, press Enter to return to the Multi-Thread Processing Parameters screen. The Job Description, Job Queue and Library fields are now completed. Press F8 (Update) to save your changes and return to the System Management Maintenance screen.

Once a multi-threading procedure has been initiated, the > symbol (MT Overrides) is displayed in the time zones selected for multi-threading.

If you choose not to use the multi-threading feature, iOptimize reorganizes one library at a time, addressing a file only after the previous file has either been reorganized or rejected from the task (either because no work is required or because there is insufficient time left to complete the task).

Making Global Changes

To change all of the actions on a queue to a new queue, on the System Management Maintenance screen press F16 (Global change). A screen is displayed on which you can enter the names of the old queues/libraries and the new queues/libraries. Press Enter to save your changes.

Viewing System Logs

You can view Optimization action history, run time history, and error logs.

Displaying Action History

Optimization—> Option 10 (Manage Optimization)—> Option 1 (Maintain Autorun Parameters)—>F20 (Actions History)

iOptimize logs actions at a summary level for all functions and at a detail level for selected items. Summary level history for all functions can be viewed globally within this option.

The following describes each option available from this screen.

Option	Description
Show Details (Option 1)	Expands the summary line to show additional details, where detailed history is kept for a function.
Whole Message (Option 5)	For truncated messages, displays the entire message.
Show Errors (Option 8)	Where the history record shows an error condition, this option displays the error. The error is often because a change in status has occurred between the forecast and the current processing (for example, a spool file has been printed already so it cannot be cleared now).

Displaying Run Time History

Optimization—>Option 10 (Manage Optimization)—> Option 1 (Maintain Autorun Parameters)—>F21 (Run Times History)

iOptimize maintains a high-level summary of all the times the various Optimization tasks ran.

This screen shows the last start and end dates that the particular task has run and the time it took to complete. Information is displayed on this screen only after some iOptimize processing has been completed (at least one overnight run).

Press F9 (Date/time sequence) to toggle the display format between date/time and alphabetical sequence.

Displaying History (Option 5)

Entering Option 5 (History) next to a job displays the history of that job (the start and end times of previous executions).This latter figure includes work areas used in the job, which are dropped just after the job ends.

Displaying Error Summary

Optimization—> Option 10 (Manage Optimization)—> Option 1 (Maintain Autorun Parameters)—>F22 (Error Summary)

Enter Option 5 (Whole message) next to a truncated message to display the entire message.

NOTE

iOptimize ignores any actions for an object for which another user has an exclusive lock (for example, it will not reorganize a file which is being used by another user). However, if iOptimize obtains exclusive use of the object and starts processing (for example, the reorganization of a file), users may experience application difficulties as they try to access the file. For this reason we recommend that jobs such as Reorganization are run at times when other users or applications do not require access to the objects. To resolve this error, review the scheduled run times.

Maintaining Processing Parameters

Optimization—> Option 10 (Manage Optimization) —> Option 1 (Maintain Autorun Parameters)—>F24 (Processing Parameters)

iOptimize can forecast the actions to perform up to a given future point and records history on actions at both detail (where applicable) and summary level. The forecast can be either created as part of the normal processing for the function or can, in most cases, be created on demand. History data limits and forecast windows are set using this option.

iOptimize can carry out processing against previously created forecasts rather than performing actions in real time (applying the parameters to the actual objects as it encounters them). When reviewing forecasts you have the opportunity to change the forecast before processing. In this way, you have the opportunity to stop iOptimize from, say, clearing a particular print from an output queue.

NOTE

It is recommended that you don't make manual changes to forecasts before processing them unless you are using the NORMAL processing mode. If you are using PRCFCS mode, the next time the forecast is processed automatically it will not reflect your manual changes (see "Understanding Processing Modes" on page 44 for more information on processing modes).

The following explains each column:

Column	Explanation
Current mode	The current processing mode. Options: *NORMAL— Action as found *GENFCS—Generate forecast only *PRCFCS—Process previous forecast
Summary days	The length of time, in days, that summary history information is kept for the corresponding function.
Detail days	The length of time, in days, that detail history information is kept for the corresponding function.
Forecast days	The number of days into the future to forecast.

For a detailed explanation of the three processing modes, see "Understanding Processing Modes" on page 44.

Ending Submitted Jobs

Optimization—>Option 10 (Manage Optimization) —> Option 2 (End Submitted Jobs)

Occasionally, you may need to use the IBM Power Systems during a time you have allocated for iOptimize to run. For example, you may have allocated Sunday afternoon/evenings as Unused time (iOptimize can run any task during this period) but then one Sunday afternoon you elect to perform an operating system upgrade.

This option assesses the situation, closes down all of the active iOptimize jobs (submitted by automatic routines) in a controlled manner, and removes pending jobs from the job queues. So, you don't have to worry that terminating the jobs will cause a problem.

When this option is selected, a confirmation screen is displayed.

Press F8 (Continue) to confirm the action. Ended jobs may need to be resubmitted, as there is not an automatic restart procedure. However, it is probably easiest to let iOptimize continue during the next available run time period. iOptimize notes that a job was terminated and, broadly speaking, pick up where it left off.

Setting Processing Windows

Machine processing times called, processing windows, are configured from the Policies portlet of the iOptimize portal application. Processing windows define the busyness of a system during a week. Processes such as

cross-referencing, optimization, and forecasting use processing windows to determine when to run and how many resources to use.

Processing windows can be defined in 20 minute increments and are categorized by Peak, Off-Peak, and Unused. Processing windows typically represent the following:

- Peak - The system's busiest times. Users are signed on and performance is high.
- Off-Peak - Most users are off the system. The system may be being backed up, but performance is less than peak.
- Unused - All users are off the system and all batch jobs have finished.

For more information about configuring processing windows, see the *Assure iOptimize Portal Application User Guide*.

Starting and Ending Product Subsystems

A

This appendix describes generally how to start and end subsystems.

Ending Subsystems

An end command syntax is as follows, where the values are defined in the table below:

```
ENDID value
```

Value	Definition
***ALL**	Ends the subsystem associated with the product.
***ID**	Ends main task scheduler

Starting Subsystems

A start command syntax is as follows, where the values are defined in the table below:

```
STRID value
```

Value	Definition
***ALL**	Starts the subsystem associated with the product.
***ID**	Starts main task scheduler
***INT**	Starts the interactive menus.
***VSP**	Runs the STRVSISVR command to start the Assure Unified Interface server.

Valid Object Types for Move Command

B

Object Types

The following is a list of the object types that can currently be moved using the Move Object command (Option M on the F1/Help menu):

Object Type	Description	Can Move to IASP?
*AUTHLR	Authority holder	N
*AUTL	Authorization list	N
*CFGL	Configuration list	N
*CSI	Communication Side Information	N
*CNNL	Connection list	N
*COSD	Class of service description	N
*CSPMAP	Cross system product map	N
*CSPTBL	Cross system product table	N
*CTLD	Controller description	N
*DDIR	Distributed file directory	N
*DEVD	Device description	N
*DOC	Document	N
*DSTMF	Distributed stream file	N
*DTAARA	Data area	N
*EDTD	Edit description	N
*EXITRG	Exit routine	N
*FILE	File	N
*FLR	Folder	N

Object Type	Description	Can Move to IASP?
*FNTRSC	Font resource	N
*FNTTBL	Font table	N
*FORMDF	Form definition	N
*FTR	Filter	N
*GSS	Graphics Symbol Set	N
*IGCTBL	DBCS font table	N
*IMGCLG	Image	N
*IPXD	IPX description	N
*JOBD	Job description	N
*JOBQ	Job queue	N
*JOBSCD	Job schedule	N
*LIND	Line description	N
*MENU	Menu	N
*MGTCOL	Management collection objects	N
*MODD	Mode description	N
*MODULE	Module	N
*MSGF	Message file	N
*M36	AS/400 Advanced 36 machine	N
*M36CFG	AS/400 Advanced 36 configuration	N
*NODL	Node list	N
*NTBD	NetBIOS description	N
*NWID	Network interface description	N
*NWSD	Network server description	N
*OUTQ	Outqueue	N
*PAGDFN	Page definition	N
*PAGSEG	Page segment	N
*PGM	Program	N
*PNLGRP	Panel group	N

Object Type	Description	Can Move to IASP?
*RCT	Reference code translate table	N
*SOCKET	Local socket	N
*SSND	Session description	N
*S36	System/36 machine description	N
*USRPRF	User profile	N
*ALRTBL	Alert table	Y
*BNDDIR	Binding directory	Y
*CLD	C locale description	Y
*CLS	Class	Y
*CMD	Command	Y
*CRQD	Change request description	Y
*PSFCFG	PSF/400 configuration object	Y
*QMFORM	Query management form	Y
*QMQRY	Query management query	Y
*QRYDFN	Query definition	Y
*SBSD	Subsystem description	Y
*SCHIDX	Information search index	Y
*SRVPGM	Service program	Y
*TBL	Table	Y
*USRIDX	User index	Y
*USRSPC	User space	Y
*VLDL	Validation list	Y
*WSCST	Work station customizing object	Y

Security Policy

C

Assure iOptimize Security Policy

The following describes the iOptimize Security Policy.

- Libraries and Objects
  - All base product objects are contained in the product's installation library.
  - Object archive library default is /VisionSolutions/IOPTIMIZE/Archive.
  - Spool archive library default is /VisionSolutions/IOPTIMIZE/Archive.
  - The library VSI*nnn*LIB contains objects used to manage the Assure Unified Interface server under which the Assure iOptimize Portal application runs.
  - Non QSYS objects for iOptimize core are stored in /VisionSolutions/iOptimize and non QSYS objects for the Assure Unified Interface and the Assure iOptimize Portal Application are stored in /VisionSolutions/HTTP.
- All objects are owned by IOPTOWNER, except as noted below:
  - IFS objects in /VISIONSOLUTIONS are owned by QPGMR.
  - The VSI* Library is owned by QPGMR.
- iOptimize does not adopt authority in programs that display a command line.
- Library level authority is controlled via Authorization List that is named the same as the installation library. At install *PUBLIC is set to *EXCLUDE, this can be manually changed and will not be reset by iOptimize.
- Program adopted authority is used to allow iOptimize to process actions that otherwise might be restricted due to the security level or access permissions of the operator.

    - The programs in iOptimize that adopt authority can be cross-referenced against file IDADPMP.
- Authority is not adopted to give the operator more authority to perform a function that they are trying to perform themselves; for example, viewing spooled output, deleting objects, reorganizing physical files.
    - Additional authority message will be displayed when options are taken that perform specific user requested actions and that user does not have authority. For example, if a user requests that iOptimize reorganize a file, then the subsequent execution of the RGZPFM command will cause security checks of *OBJALTER, *OBJMGT, data authorities of *UPD, *READ, *DLT and *ADD and library authority of *EXECUTE.
    - These additional checks allow us to display a message that says “Not authorized to object” but authorization failures can also cause error screens to be displayed as in the V7.0 product
- iOptimize users must be authorized to the product using the Authorized Users portlet in the iOptimize portal application (see the *Assure iOptimize Portal Application User Guide*). When users are authorized they will be granted access to legitimate authority adopting programs specifically and also placed into the iOptimize Authorization List.

Archiving, Purging, and Internal Controls

iOptimize is a fully autonomic product including self-protecting, self-optimizing and self-managing, more specifically:

- Records are purged in all files based on days to keep parameters
- MSGQ objects that are created to support interactively signed on iOptimize users are removed if not active when the subsystem restarts
- If Monitor jobs are ended or fail they will be automatically restarted
- All physical files within iOptimize have re-use deleted records switched on
- Most physical files are set to *NOMAX but some control files have specific size allocations
- Space is not pre-allocated in physical files
- Jobs will not have resources increased unless specifically stated

Add a Notification

D

About Notifications

The Notifications portlet in the Assure iOptimize portal application displays notifications that have been configured through the iOptimize IDADDNFY command. The Notifications portlet is updated every minute to provide new information. The severity of any new notifications is reflected in the overall status of the instance. The status for any iOptimize event monitor defined for a system propagates up to the Instance Status portlet for each instance in the Assure Unified Interface. For more information, see the *Assure iOptimize Portal Application User Guide*.

Add a Notification

You define notifications using the IDADDNFY command in iOptimize.

NOTE

If you prompt the command IDADDNFY outside of the iOptimize product, qualify the command with the iOptimize product library.

1. From a command line in iOptimize, prompt the following command:

   ```
   IDADDNFY
   ```

   The iOptimize Add Notification (IDADDNFY) screen is displayed.

2. Configure the following parameters:

Parameter	Values
Description	In the text field, enter a brief description of this notification enclosed in apostrophes. The field allows entry of up to 132 characters in length.
Severity	Select the severity that is assigned to this notification. Values available for severity include: • *ACTREQ - Action required. • *ACTREQUNK - Action required. Status is unknown • *CAUTION - Caution. • *STOPPED - Stopped. • *ATTN - Attention. • *ATTNUNK - Attention. Status is unknown. • *INFO - Informational. • *OK - OK. No action is required.
Details	Enter or select a detail description associated with this notification. Possible values are: • Enter a detailed description. The field allows you to specify up to 1024 characters of text, enclosed in apostrophes. A link to a web page can be included using the following format: [http://mylink] [My Link Alias|http://mylink] • *NONE - No description is associated with this notification.
Optional Parameter (F10)	
Related event	Specify an event to be linked to this notification. Possible values are: • *NONE - No event is to be associated with this notification.

Error Messages

The following are error messages that can be issued for the IDADDNFY command:

- IDE0173 - iOptimize Add Notification command ended abnormally.
- IDE0174 - iOptimize Add Notification command completed successfully.