FORTRΔ



Copyright Terms and Conditions

Copyright © Fortra, LLC and its group of companies. All trademarks and registered trademarks are the property of their respective owners.

The content in this document is protected by the Copyright Laws of the United States of America and other countries worldwide. The unauthorized use and/or duplication of this material without express and written permission from Fortra is strictly prohibited. Excerpts and links may be used, provided that full and clear credit is given to Fortra with appropriate and specific direction to the original content.

202504230907

Table of Contents

Welcome to Robot HA	
Replication Basics	10
Replication through journaling	
Replication through replication/sync jobs	10
Ending sync/apply jobs	
Journaling Basics	
Remote Journaling	
Journaling	
Synchronization Attributes	
*LIBRARY	
*IFS	14
*GROUP	
*SYSTEM	
About Libraries	
Initial Library Copy	
Let Robot HA Copy the Library	
Copy the Library Manually	
Cross-library Dependencies	
Multiple Journals Per Library	
Library Sets Overview	
Considerations for Library Sets	

Using Library Sets	
Use a different journal for each set	23
Use Include/Omit to Ensure Each Set is Unique	
Take Care when Starting Replication from Scratch	23
Target Triggers	
Row/Column Access Control (RCAC)	
Checking your Synchronization Status	
On the production system	
On the backup system	
Replication Groups	
Creating Replication Groups	
Assigning Libraries to a Group	
Example: Group Precedence	
Swap Audit Overview	
Starting and Ending Audit Monitoring	
Creating an Audit	
To create a new Audit:	
To create an Audit from a copy:	
Audit Notification Email Addresses	
Audit Notification in Robot Network	
Swap Audit Steps	
Swap Audit Stop/Recover Matrix:	
Maintain and Running an Audit	

To run an unscheduled audit:	
To edit an audit:	46
To delete an audit:	
Audit History	
To view run results of an audit:	47
To rerun an audit job:	47
To delete an audit job:	
Role Swaps	
Additional Tools	
Compare	
Clean	
Journal Management	
Archive or Refresh	
Reference	
Robot HA Menus	
Robot HA Main Menu	
Options	
Control Menu	
Setup Menu	
Reports Menu	
Status Reports Menu	
Integrity Reports Menu	
History Menu	64

Audit Menu	
Role Swap Menu	67
Robot HA Panels	
Maintain Audit	
Audit History	
Audit History Detail	71
Audit History Report Panel	
Add Audit Attributes	
Audit Steps	
Add Group Sync Attributes	
Add IFS Sync Attributes	
Add Library Sync Attributes	
Add System Sync Attributes	
Change Library Sync Attributes	
Change IFS Sync Attributes	
Change Group Sync Attributes	
Change System Sync Attributes	
Change RSF Defaults (CHGRSFDFT)	
Change RSF User Options (CHGRSFUO)	
Options	
Function Keys	
Check Defined Items (CHKATTRSF)	
Check IFS Directory (CHKDIRRSF)	

Check IFS Object (CHKIFSRSF)	
Check Libraries (CHKLIBRSF)	
Check Objects (CHKOBJRSF)	
Check Replication Group (CHKGRPRSF)	
Check System Items (CHKSYSRSF)	
Clean Defined Items (CLNATTRSF)	
Clean Directory (CLNDIRRSF)	
Clean Library (CLNLIBRSF)	
Clean Replication Group (CLNGRPRSF)	
Clean System Items (CLNSYSRSF)	
Display Defined Item Status (DSPRSFDSS)	
Display Group Sync Status (DSPRSFGSS)	
Private Authority Status	
Display IFS Sync Status (DSPRSFISS)	
Display Library Sync Status (DSPRSFSS)	
Display Synchronization Log	
Display System Sync Status (DSPRSFSSS)	
Email Address List	
End Synchronization Job (ENDSYNCRSF)	
Initialize Robot HA Setup (INZRSFHA)	
Purge History	
Sync Authorities	
Select Email Address	

Synchronization Attributes Report Panel	
Start Synchronizing	
Work with Active Jobs	
Work with All Spooled Files	
Work with User Jobs	
Work with RSF Servers	
Work with Synchronization Attributes	
Display Synchronization Attributes	
Confirm Delete of Synchronization Attributes	
Subset Window	
Appendixes	
Appendix A: Terminology	
Appendix B: Defining System Connections	
Overview	
Defining TCP/IP for LAN/WAN	
Defining TCP/IP Over the Internet	
Appendix C: Running a System 36 Environment	

Welcome to Robot HA

Advanced Configuration and Product Considerations

Robot HA (built on the successful RSF-HA software) is an IBM i high availability solution that reduces downtime and provides data replication and business continuity. It allows you to replicate your libraries, IFS directories, and system information from production to backup, which helps limit planned and unplanned downtime.

With Robot HA, you can use remote journaling for instant replication across any distance the moment changes are applied to your data. You can also test five different role swap scenarios to verify this replication is done correctly and ensure your systems are secure in the event of unplanned downtime. Role swaps can be tested while users are using production to minimize the disruption of working hours.

Robot HA monitors all aspects of replication – including replication tasks in error, excessive transport or apply lags, user-defined conditions, and more – and sends notification to your phone or email for any issues that need attention.

Robot HA is self-healing. It continuously audits and repairs objects and IFS directories to ensure they match the production system. This synchronization process ensures that your data is up-to-date on the backup without the need for manual monitoring.

You can replicate the following data with Robot HA:

Library data	User profiles
IFS directories	Object authorities
Authorization lists	Spooled files
System values	Message queues
Program objects	

NOTE: These instructions are intended as a guide for advanced configuration and indepth product considerations. You can supplement this information, where needed, with the *Robot HA Implementation Guide*.

Replication Basics

In Robot HA, data is replicated in two different ways: through journaling and through submitted replication/sync jobs.

Replication through journaling

Journaling handles database file changes, such as when a record is added, changed, or deleted from a file. It sees that there is a journal entry for the change made on the production system and sends it to the backup system, where the changes are then applied to the same file.

Replication done via journaling happens in real time. So changes made to the production system are immediately made to the backup system.

The recovery point for Robot HA is the point in which you have ended the journal apply job on the target system.

Replication through replication/sync jobs

Data can be replicated through these jobs via the <u>Work With Synchronization Attributes</u> panel.

A replication or synchronization job looks at a library or IFS directory to see what's changed for objects that cannot be journaled. It saves the objects and sends them from the production to the backup, where the changes are then applied. When the job finishes running, it resubmits itself to run based on the time interval set for each synchronization attribute.

This process continues until you stop synchronization.

Ending sync/apply jobs

Run the following command(s) on the Prod system to end sync jobs, apply jobs, or both. These jobs should only be ended using these commands.

Ending sync jobs only

ENDSYNCRSF JOB(*ALL) OPTION(xxxxx) RMTJRNJOB(*NONE)

Ending remote apply jobs only

ENDSYNCRSF JOB(*NONE) OPTION(xxxxx) RMTJRNJOB(*ALL)

Ending sync and remote apply jobs

ENDSYNCRSF JOB(*ALL) OPTION(xxxxx) RMTJRNJOB(*ALL)

See End Synchronization Job (ENDSYNCRSF) Panel for information on parameters.

Consider why are you ending the sync jobs or remote apply jobs.

Journaling Basics

Robot HA has journaling and remote journaling abilities, which are used to replicate your libraries, IFS directories, IFS objects, data areas, and data queues.

Remote Journaling

We recommend you use remote journaling. This manages all objects within your library and IFS directories that can be journaled by the operating system. Journal changes are sent quickly and efficiently, and they are applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply. The remote journaling process happens at a low level and improves performance when getting transactions to the backup system.

At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries. The remote journaling process happens at a low level and improves performance when getting transactions to the backup system.

NOTE: If you use remote journaling, be aware that filtering is not supported. *NO is assumed for the **Filter Journal Entries** parameter.

In certain situations, you might want to consider omitting journaling. This could be necessary when a library contains no database files, data areas, or data queues, or when a library only contains source files and other non-file objects.

TIP: Remote journaling uses TCP ports **446** (DRDA), and **447** (DDM) and **3777** (remote journaling proper). If you intend to use remote journaling and your source and target machines are separated by a firewall, be sure to open ports 446, 447 and 3777. In addition, if you are using the value *USRIDPWD for **lowest authentication method** on the CHGDDMTCPA command, refer to <u>Using a Secure DDM TCP/IP Connection with Robot</u> HA

Journaling

Journaling should be used if you want to filter the journal entries instead of sending them all over for replication.

Synchronization Attributes

Synchronization attributes can be set for each library, IFS directory, system object type (*USRPRF and *SYSVAL, for example), and group to tailor how Robot HA synchronizes these to the backup system.

*LIBRARY

You can <u>add synchronization attributes</u> for a given source library, target library, server ID, and set name combination.

You can specify:

- Whether files are replicated at the object level or the record level, using the Local Journal parameter. Journaling is recommend for most libraries being replicated.
- Whether to use remote journaling. Remote journaling is also recommended for most libraries being replicated.

NOTE: When using remote journaling, be sure to specify different libraries for the local and remote journals in order to be ready for a role swap. The recommended name for the remote journal library is RMTJRNLIB.

- Whether to synchronize changes to object authorities as well as changes to the objects themselves.
- Which objects to include/omit.
- Robot HA has restricted the ability to set up synchronization attributes for IBMowned libraries. Synchronizing some IBM libraries could result in many issues. The libraries in the table below can be set up as sync attributes but contain other restrictions. All other IBM-owned libraries cannot be set up as synchronization attributes.

Library	QGPL	QMQ*	QS36F	QUSRSYS	SYSIBM	SYSIBMADM	SYSTOOLS
Setup allowed	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Allow to be cleared*	No	No	Yes	No	No	No	No

Library	QGPL	QMQ*	QS36F	QUSRSYS	SYSIBM	SYSIBMADM	SYSTOOLS
Omit required	No	No	No	No	No	No	No
Include required	Yes	No	No	Yes	No	No	No
*JRN allowed	Yes	No	Yes	Yes	No	No	No
Allow multiple sets	Yes	No	Yes	Yes	No	No	No
Allow fix objects/wrong JRN	Yes	No	Yes	Yes	No	No	No
Allow in a group	No	No	No	No	No	No	No
Disable Trigger/Constraint s	Yes	No	No	Yes	No	No	No
Allow to be cleaned*	Yes	No	Yes	Yes	No	No	No

NOTE:

*- Allow to be cleared is for the first time action. We normally clear the library on the target side and refresh the complete library to ensure everything is the same. In some cases with IBM libraries we cannot do this.

Allow to be cleaned (CLNxxxRSF) command is used to reconcile differences for nonjournaled objects between the source and the target.

*IFS

You can <u>add synchronization attributes</u> for a source directory, target directory, and server combination.

We recommend using local and remote journaling when syncing IFS directories.

NOTE: When using remote journaling, be sure to specify different libraries for the local and remote journals in order to be ready for a role swap. The recommended name for the remote journal library is RMTJRNLIB.)

*GROUP

You can add synchronization attributes for a replication group.

Groups are only allowed with libraries. Library synchronization attributes must be included in the group for the *GROUP synchronization attribute to be created.

Everything included in a group uses the same journal, so be cautious when adding groups. Apply performance can be affected if you have multiple libraries with many transactions.

You can specify:

- Whether files are replicated at the object level or the record level, using the Local Journal parameter.
- How to manage items associated with any library synchronization attributes added to the group.
- Whether to use remote journaling. Remote journaling is also recommended.

NOTE: When using remote journaling, be sure to specify different libraries for the local and remote journals in order to be ready for a role swap. The recommended name for the remote journal library is RMTJRNLIB.

• The libraries to include in your group.

For more information on *GROUP synchronization attributes, see <u>Replication Groups</u>.

***SYSTEM**

You can <u>add synchronization attributes</u> for system information such as user profiles and system values.

You can specify:

- Whether to set target user profiles to *DISABLED status when they are synchronized.
- The name of the library in which to store configuration information when synchronizing *CFG information.
- Which objects to include or omit. Include and omit parameters are not valid for *NETA, *PVTAUT, and *CFG.

About Libraries

The following section contains introductory information about replicating libraries in Robot HA. We suggest you read through it carefully before proceeding with replication groups or library sets.

Initial Library Copy

There are two ways a library can be synchronized for the first time.

Let Robot HA Copy the Library

Robot HA can send the entire library to the backup system in order to establish a clean synchronization boundary. All production and backup objects must have identical ownership, attributes, and data. In the process of synchronizing the library, Robot HA clears it from the backup system and restores the copy saved on the production system.

If Robot HA copies the library over to the backup system and the library already exists on the backup, it is possible to receive an error if the objects cannot be cleared. In this case, an inquiry message is sent to the system operators message queue on the production system. When this happens, there are a few responses you can send:

- I Ignore the error and continue. If you choose this reply, the objects in the library will not be fully synced with the library on the production system.
- **C** Cancel replication.
- **R** Retry replication. This should be done after the problem objects have been manually deleted from the library. If all issues are resolved, replication of the library will continue.

The maximum size of a library for this process is 2TB. If this process is attempted with a library larger than 2TB, the process fails with error CPF3708 indicating the SAVF is too small. To perform the initial library copy in these environments, follow the steps in <u>Using</u> Tape to Manually Replicate a Library or IFS directory in Robot HA.

Copy the Library Manually

If you do not want Robot HA to copy the library over, you can choose to do it manually using a tape, save file, optical, or other method. This is done outside of Robot HA. For more

information on this process, see <u>Using Tape to Manually Replicate a Library or IFS directory</u> in Robot HA.

Cross-library Dependencies

If you store logical files in a different library from the physical files they are dependent on, you have a cross-library dependency. We recommend you avoid organizing your data this way if possible. It can be complicated to manage and in some instances is not possible due to IBM restrictions.

If you are using library redirects and have cross-library dependencies, IBM has limitations for how objects can be restored. Because of these limitations, you cannot use library redirects with cross-library dependencies.

Examples:

Cross-library dependencies with no library redirection

There is a physical file that resides in one library and its dependent logical file resides in a different library. The libraries are the same on both the source and target system. Robot HA can handle this situation as long as the physical file library is restored first. In this example, if LibraryB restores before LibraryA, it will fail, but we will retry later. Robot HA will self heal and LibraryB will successfully restore after the next sync cycle because LibraryA will have been restored.

Object	Source Library Target Library	
PF1	LibraryA	LibraryA
LF1	LibraryB	LibraryB

Cross-library dependencies with library redirection

Below are two scenarios that use cross-library dependencies with library redirection,.

1. Source library names do not reside on the target

There is a physical file that resides in one library and its dependent logical file resides in a different library. You are using library redirects so the libraries are different on the source and target. Robot HA cannot handle this situation because of the limitations on the IBM restore command.

Object	Source Library	Target Library
--------	----------------	----------------

PF1	LibraryA	LibraryZA	
LF1	LibraryB	LibraryZB	

During replication, we know you are redirecting the library because of your definition, so we format the RSTOBJ command to reflect this. *RSTOBJ savlib(LibraryB) RSTLIB (LibraryZB)*, this would be okay if all the dependent objects were in LibraryB. However, in this example, there is a dependent object (PF1) in LibraryA. IBM doesn't know about the LibraryA to LibraryZA redirection. So, when the RSTOBJ is executed, it is looking for LibraryA/PF1 to exist on the target side. In this case, it will not exist, because it was restored to LibraryZA instead. Therefore you will get an error on the restore and it will fail.

If you are in this situation, you have a few options.

- Ensure that all dependent objects are in the same library. Then you can use library redirects.
- Dependent objects can be in different libraries, but library redirects cannot be used.

2. Source library names also reside on the target

There are two physical files that reside in two separate libraries. There is a join logical file that joins these two files. The join logical file resides in the library of only one of the physical files. You are using library redirects so the libraries have different names on the source and target systems. The source library resides on the target system, but it is not where you are replicating to. Robot HA cannot handle this situation because of the limitations on the IBM restore command.

Object	Source Library	Target Library
PF1	LibraryA	LibraryZA
PF2	LibraryB	LibraryZB
JLF1	LibraryB	LibraryZB

In this scenario, the join logical file will be restored, but it will not have the correct relationship to the physical file in Library A. There is no way for the join logical file JLF1 to know that PF1 is being redirected to LibraryZA. It expects it to be in LibraryA.

If you are in this situation, you have a few options.

• Ensure that all dependent objects are in the same library. Then you can use library redirects.

• Dependent objects can be in different libraries, but library redirects cannot be used.

Multiple Journals Per Library

While it is often simpler if all objects in a library are journaled to a single journal, Robot HA can also replicate changes for objects journaled to multiple journals in a library. This is useful if you have a library that represents the majority of your data and hinders the process of replication because of its size. By defining multiple journals for this library, replication of the library is spread out between journals, making it faster and more efficient. On the backup system, multiple apply jobs run for this library which keeps it up-to-date.

NOTE: it is important to ensure that the set of remote journals is complete and all apply jobs are done before attempting a role swap, or the library may be out of sync.

There are a few things to consider before setting up multiple journals per library:

- Think carefully before you define multiple journals for a library. Once in use, it becomes difficult to determine if the entire library is being replicated and will require you to verify that it is.
- Setting up multiple journals for a library requires you to use the **Objects to include** and **Objects to omit** parameters. Make sure your includes and omits are such that you do not have duplicates or missing items after replicating the library a few times.
- Make sure the synchronization attributes for each set you define use different journals but sync to the same library.

Library Sets Overview

Library replication entries are created with a unique combination of from-library, to-library, server-ID, and set-name. Multiple replication sets can be defined for a library by specifying different set names with the same from-library, to-library, and server-ID.

Generally, you should define a single set for each library. However, there are cases where you might want to define multiple sets:

- When performance is critical and you want to have multiple target journal apply jobs run for the library.
- When different objects in the library must be journaled to different journals due to user requirements or some other reason.
- When different objects in the library must have different sync attribute settings for some other reason.

Library sets are only needed when a) some of the library's objects must be treated differently than others and b) you are replicating the objects to the same target system and library.

In the following cases, you'd need to define multiple replication entries but wouldn't need to define multiple sets. The sync entries in each case below could all be defined with the *DFT set.

Case 1: If you replicate a library to two or more target systems, the sync entries would have the same from-library and to-library values but different server-IDs.

Case 2: If you replicate a library to two or more target libraries on the same target system, the sync entries would have the same from-library and server-ID values but different to-library values.

Considerations for Library Sets

Before you define a library set, consider the following points:

- It is important to use different set names. The sync entries for this library must also be unique.
- The library can be split up in multiple ways. Think about how you want this done.
- Each set should contain include/omit specifications that uniquely select a set of objects in the library, and everything in the library should be covered. If the entire library is not replicated, things can get complicated.

• There are multiple ways you can move the library to the backup system for the first time. Do you want Robot HA to move it, or do you want to move it manually? This is important to consider, as there will be multiple sets synchronizing to the same library. If you use sets and your set definitions do not include everything, you may not have a complete library on your backup. The first time you run a set, Robot HA only sends the things included in the set itself, *not* the entire library, so take special care to make sure your library is complete.

Using Library Sets

Use these instructions to set up your library sets.

Use a different journal for each set

When defining multiple sets for a single library, you must define a different journal (and associated remote journal) for each set. This is needed to ensure that each target journal apply job sees only the entries related to the appropriate set.

For example, if replicating library ABC with sets *DFT, SET2, and SET3, you might assign journal JRNLIB/ABC and remote journal RMTJRNLIB/ABC to the *DFT set, following the usual standard; then assign journal JRNLIB/ABC2 and remote journal RMTJRNLIB/ABC2 to SET2; and journal JRNLIB/ABC3 and remote journal RMTJRNLIB/ABC3 to SET3.

Use Include/Omit to Ensure Each Set is Unique

When defining multiple sets for a library, define include/omit specifications so as to ensure that no object is included in more than one set. Objects included in more than one set will not replicate properly.

Take Care when Starting Replication from Scratch

Starting replication for a library from scratch means starting replication for the very first time or refreshing the whole library after replication has been running for a while. The whole library is sent from the source to the target. After the library is successfully restored, incremental changes can be sent.

When starting from scratch for a library that has multiple sets defined, special considerations apply.

 You must manually clear the target library. Synchronization jobs for libraries with multiple sets defined ignore the Clear target library first parameter on the Synchronize Libraries command and assume a value of *NO.

NOTE: Normally, when replication is started from scratch, a value of CLEAR(*YES) is used and Robot HA automatically clears the target library to ensure the initial restore is done cleanly. With multiple sets, however, each entry sends a unique set of objects to the same target library. Since the order in which sync jobs run is unpredictable, any set that clears the target library could delete objects belonging to other sets.

- If multiple sets are defined for a library and you restart replication for one of the sets from scratch, you must restart all of the sets for the library from scratch. This follows from the fact that the target library will be manually cleared at the beginning of the refresh process.
- Due to the difficulty refreshing a whole library for which multiple sets are defined, you should specify *NOMAX for the **Max items to refresh** parameter when defining sync attributes for libraries with multiple sets.

Once an initial sync point is established for all sets for the library, only incremental changes are sent. From that point forward, the sync jobs for the various sets behave like all other sync jobs.

Target Triggers

To avoid duplicate updates on the backup system, we recommend you disable triggers on the target. If you decide to leave target triggers enabled, make sure your trigger programs exist on the target system. In addition, the job that applies journal changes to the target system must be able to find any existing trigger programs.

If you choose to disable your target triggers, you can do so by using the default value *YES for the Disable triggers on target parameter when setting up the sync attributes for a library.

To change this parameter for an existing library's sync attributes:

- 1. From the Robot HA Main Menu, select option **1** Synchronization attributes.
- 2. Review the parameters and make any changes, then press Enter to search.
- 3. Enter option **2**, Attributes, next to the library you want to change and press Enter. The Change Library Sync Attributes panel displays.
- 4. On the second page, make sure ***YES** is specified for Disable triggers on target.
- 5. Press Enter to save your changes.

Row/Column Access Control (RCAC)

Robot HA does not support Row/Column Access Control. Any rules you have applied for RCAC need to be manually applied to the backup as we do not sync them.

Checking your Synchronization Status

After your attributes have been synchronized, you can check the synchronization process on the production and backup systems to make sure everything is running correctly.

On the production system

Use any of the following methods on the production system to check the status of your synchronization jobs.

Method #1

Select option **1** from the Robot HA Main Menu to display the Work with Synchronization Attributes panel. Press Enter to build the list. Then, use the following options to check the synchronization status:

- Check the synchronization dates and error dates listed for each item.
- Enter option **5** next to a sync attribute, then press Enter to display its job status.

Method #2

Select option **2** from the Robot HA Main Menu to display the Control Menu. Then, use the following options to check your job logs:

- Select option **13** to display a synchronization log. Press Enter on the Display Synchronization Log screen to display the Synchronization Error Log.
- Select option **11** to display synchronization jobs.

On the backup system

Use the following method on the backup system to check the status of your synchronization jobs.

Select option **2** from the Robot HA Main Menu to display the Control Menu. Then, use the following options to check your job logs:

- Select option **13** to display a synchronization log. Press Enter on the Display Synchronization Log screen to display the Synchronization Error Log.
- Select option **11** to display synchronization jobs.

NOTE: When an error occurs while applying journal entries to the backup system, the save file containing transmitted journal entries is renamed to Jdddddhhmm, where *ddddd* is the Julian date and *hhmm* is the time in hours and minutes. A save file renamed in this way is not deleted by Robot HA; you must delete them manually.

Replication Groups

Using replication groups allows you to collect replication entries for several libraries and control them as a single unit. The libraries and their journal entries are submitted together, ensuring that database integrity is maintained and that no records are out of sync between the production and backup systems.

Replication groups are especially helpful if you do not want to replicate libraries individually. However, they do use a single replication journal for all the libraries in a group, so keep that in mind as you consider performance.

We recommend you use replication groups in the following situations:

- When you want to ensure that changes across multiple libraries are processed on the target system in the same order they were processed on the source system.
- When an application uses multiple libraries and you want to start/stop replication for the application easily.
- When an application uses commitment control in a way that spans multiple libraries (where objects in those libraries can be part of the same commit transaction).
- When an application has related objects in multiple libraries, like physical files in one library and logical files in another.

Creating Replication Groups

Before you can start using replication groups, you first need to create a group and assign libraries to it.

NOTE: If you already have a group you want to use, skip to the section on <u>assigning</u> <u>libraries to a group</u>.

To create a group:

- 1. From the Robot HA Main Menu, select option **1** Synchronization attributes.
- 2. Review the filter options and make any changes, then press Enter to search.
- 3. Press F6 to create a new synchronization attribute.
- 4. On the Add Sync Attributes dialog box, enter option **1** next to *GROUP and press Enter. The Add Group Sync Attributes panel displays.

NOTE: For more information on any of the parameters in this panel, press **F1** or see <u>Add Group Sync Attributes</u>.

- 5. Enter a name for the replication group in the Group field.
- 6. Enter the name of a server directory entry in the Server ID field, or press **F4** to select a name from a list.
- 7. Review the rest of the parameters on this page and make changes as necessary.
- 8. Page down to add information about each library you want to assign to the group. You'll need to fill in the following parameters:

NOTE: *YES or *NO appear in the Exists column depending on whether the library sync attribute you entered already exists. If it does not exist, it will be created.

- From lib: Enter a valid library name.
- To lib: Enter * FROMLIB to specify that this library is the same as the From Library, or enter a valid library name.
- Server id: Enter *GRPSERVER to specify that the server ID in this library entry is the same as the group server ID, or enter a valid server ID.

NOTE: While a different server ID may be needed to identify the desired library sync entry, the group server ID overrides the library server ID at run time.

- Set name: Enter * DFT to specify the default set name that is used to identify the library sync entry, or enter a valid set name.
- 9. Page down to review the parameters on the next page and make changes as necessary.
- 10. When you are finished, press Enter to create the replication group. This adds the new group to the Work with Synchronization Attributes panel.

Assigning Libraries to a Group

To set up a replication group, you need to assign one or more libraries to it. There are two ways you can do this: by assigning libraries to a group or by assigning a group to each library.

NOTE: If you have not created a group for replication yet, see the previous section on <u>creating replication groups</u>.

To assign one or more libraries to a group:

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- Toggle (F11) until you get to the Work with Sync Attributes-Groups panel. Enter option 2, Change Attributes, next to the group you want to use and press Enter. The <u>Change Group Sync Attributes</u> panel displays.

Change Group Sy	nc Attributes	(CHGRSFGSA)
Type choices, press Enter.		
Group	GRPTEST	Name
Group server ID	<u>PROD</u>	Name, *SAME
Library information:		
From library	RBTHALIB	Name, *NONE, *SAME
To library	*FROMLIB	Name, *FROMLIB
Library server ID	*GRPSERVER	Name, *GRPSERVER
Set name	<u>*DFT</u>	Name, *DFT
+ for more values _		
Create library entries	<u>*NO</u>	*YES, *NO
Library management option	<u>*ENDJOB</u>	*ENDJOB, *ENDRMTJRN
Journal	*NONE	Name, *SAME, *NONE, *GROUP
Library		Name, JRNLIB
Remote journal	*NONE	Name, *SAME, *NONE, *JRN
Library		Name, RMTJRNLIB
RSF manages the journal	*YES	*YES, *NO, *SAME
Change receiver every	*DAY	*INTERVAL, *DAY, *SYS, *SAME
		More
F3=Exit F4=Prompt F5=Refresh	F10=Additiona	al parameters F12=Cancel
F13=How to use this display	F24=More keys	5
	-	

3. Under Library information, add information about each library you want to assign to the group. You'll need to fill in the following parameters:

NOTE: Type a + and press Enter to display additional library entries.

- From library: Enter a valid library name.
- To library: Enter *FROMLIB to specify that this library is the same as the From Library, or enter a valid library name.
- Library server ID: Enter *GRPSERVER to specify that the server ID in this library entry is the same as the group server ID, or enter a valid server ID.

NOTE: While a different server ID may be needed to identify the desired library sync entry, the group server ID overrides the library server ID at run time.

• Set name: Enter *DFT to specify the default set name that is used to identify the library sync entry, or enter a valid set name.

4. When you are finished assigning libraries to the group, press Enter. This saves your changes and returns you to the Work With Sync Attributes panel.

NOTE: When using this method to assign libraries to a group, the sync attribute entries for the libraries must already exist <u>or</u> you must specify ***YES** for the Create Library Entries parameter.

To assign a group to a library:

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. Type option **2**, Change Attributes, next to the library you want to use and press Enter. The Change Library Sync Attributes panel displays.
- 3. Press **F10** to see all parameters, then page down to the last screen.

Change Librar	y Sync Attributes (CHGRSFSA)
Type choices, press Enter.	
Replication group Ignore remote error	<u>*NONE</u> Name, *NONE, *SAME <u>*NO</u> *YES, *NO
F3=Exit F4=Prompt F5=Refresh F24=More keys	Bottom n F12=Cancel F13=How to use this display

- 4. For the **Replication group** parameter, enter the name of the group you want to assign this library to.
- 5. When you are finished changing the group name, press Enter to save your changes.

NOTE: A library can only be assigned to a group once. If there is more than one sync attribute entry for a from-lib, only one of the sync attributes can be assigned to a group. This is because a) there is only one journal and one remote journal apply job for the group, and b) while journal entries tell you the library for a changed object is on the production system, they do not contain additional information that distinguishes between multiple sync attribute entries for the library.

Example: Group Precedence

Sometimes, sync attributes are set at both library and group levels. When this happens, the group sync attributes (such as the journal and server ID) override any similar sync attributes that are specified at the library level.

Consider the following library and group sync entries:

Library entry

- From library: TESTLIB1
- To library: TESTLIB2
- Server ID: TEST
- Set name: *DFT
- Journal: JRNLIB
- Journal library: TESTLIB1
- Remote journal: RMTJRNLIB
- Remote journal library: TESTLIB1

Group sync entry

- Group: GROUP1
- Server ID: BACKUP
- Journal: JRNLIB
- Journal library: GROUP1
- Remote journal: RMTJRNLIB
- Remote journal library: GROUP1

In this example, when the entry for library TESTLIB1 is added to GROUP1, the server ID and journal information specified in the library sync entry are overridden at run time by the values specified for the group. However, the To library, include/omit specifications, and any attributes that can only be specified at the library level are still honored.

Swap Audit Overview

Because of the business critical nature of a role swap, it is imperative that you ensure that both the production system and the backup system are in a swap ready state. Creating and running a Swap Audit is a great way to constantly monitor the state of both systems.

When you create a Swap Audit, you are defining an interval to run the audit, as well as the audit steps that will be run to verify readiness. The steps within an audit are the checklist that Robot HA uses to verify that you are ready to perform a successful role swap. After the audit is run, you can view the results of each run step to determine the best course of action to correct the issue.

The following sections and steps prepare you to run and view audit results.

Starting and Ending Audit Monitoring

You can start or end audit monitoring from the Audit Menu and the Control Menu. Starting audit monitoring submits a job called RHAAUDMON in the QSYSWRK subsystem that initiates the monitoring of scheduled audit job intervals and will start the audit job when the interval is satisfied. Scheduled audit jobs will not run when audit monitoring is off. Audit monitoring should only be ended in the event of a backup or other type of maintenance.

To start audit monitoring from the Control Menu:

- 1. From the Robot HA Main Menu, select option **2** and press **Enter**.
- 2. On the Control Menu, select option **5** and press **Enter** to start the audit monitor. This starts the monitoring of scheduled audit jobs.

To stop audit monitoring from the Control Menu:

- 1. From the Robot HA Main Menu, select option **2** and press **Enter**.
- 2. On the Control Menu, select option **6** and press **Enter** to end audit monitoring.

To start audit monitoring from the Audit Menu:

- 1. From the Robot HA Main Menu, select option 6 and press Enter.
- 2. On the Audit Menu, select option **2** and press **Enter** to start the audit monitor. This starts the monitoring of scheduled audit jobs.

To stop audit monitoring from the Audit Menu:

- 1. From the Robot HA Main Menu, select option **6** and press **Enter**.
- 2. On the Audit Menu, select option **3** and press **Enter** to end audit monitoring.

Creating an Audit

Before you can realize the benefits of running an audit, you must first create them. Audits can be created from the Maintain Audits panel.

Note: You must have Audit Monitor started to run a scheduled audit.

To create a new Audit:

- 1. From the Robot HA Main Menu, select option **6** Audit Menu. Then, select option **1** Maintain Audits.
- 2. Press **F6** Create. The Add Audit Attributes panel displays.
- 3. Add the required information about the Audit. You'll need to fill in the following fields:

NOTE: Press F4 Prompt within fields requiring a selection from a predefined list.

- Audit name: Enter a valid Audit name.
- Audit type: Currently, the only selectable value is "SWAP".
- Audit description: Enter a description of the Audit.
- Target Server: Select the server that will be the target for the swap.
- 4. When you are finished filling out the required fields, edit the rest of the fields on the panel as desired.

NOTE: Audits are scheduled to run via the parameters defined in the Interval, Units, Start, and End fields.

- Interval: Enter the repeat interval.
- Units: Choose the repeat interval unit.
- Start: Enter the repeat interval start time.

- End: Enter the repeat interval end time.
- Robot HA email address: Enter an email address or choose from a list of email addresses to send audit emails via the F4 prompt. See Audit Notification Email Addresses for more information. If you choose just one, that email address will be displayed in the field. If you enter more than one email address, *LIST will be displayed in the field.

NOTE: Ensure that the necessary email notification fields have been entered in System Defaults (See "Email Notification Setup" in the Robot HA Implementation Guide for details).

- Robot Alert device: Choose a Robot Alert device to use for notifications. Only
 available if Robot Alert is installed and at least one device and one vendor are
 defined within Robot Alert. See the Robot Alert User Guide for additional device
 setup instructions.
- **Robot Network:** Choose the message type(s) to send to Robot Network. Only available if Robot Network is installed.

NOTE: For notifications to be sent to Robot Network, you must first create a user application in Robot Network. See Audit Notification in Robot Network for more information.

- Send to message queue: Choose the message queue where notifications are sent.
- Message Queue Library: Choose the library where the message queue exists.
- **Message Queue Library ASP group:** The ASP group of the message queue library.
- Hold audit: Choose whether you want to prevent the audit from running.
- History runs to retain: Enter the number of runs to retain in history.
- 5. Press **F10** Audit Steps. The Audit Steps panel displays.
- 6. Edit the Run, Stop, and Recover fields for selected steps as needed.

NOTE: Not all step fields are editable. See <u>Swap Audit Steps</u> for more information.

- 7. Press **Enter** to save the steps, or press **F3** if no changes were made.
- 8. On the Add Audit Attributes panel, Press **Enter** to save the audit. The audit now displays on the Maintain Audits panel for running or editing. See <u>Maintaining an Audit</u> for more information.

To create an Audit from a copy:
- 1. From the Robot HA Main Menu, select option **6** Audit Menu. Then, select option **1** Maintain Audits.
- 2. Enter **3** in the option field of the audit you wish to copy. Press Enter.
- 3. In the Confirm copy of audits panel, enter a name for the new audit.
- 4. Press Enter to accept the copy confirmation
- 5. If necessary, enter **2** in the option field of the newly created audit and press Enter to edit the audit's parameters. See <u>Maintaining an Audit</u> for more information.

Audit Notification Email Addresses

When creating or editing an audit, you can choose whether to have email notifications of the audit status sent to one or more email addresses.

NOTE: The email address list that displays after an **F4** Prompt on the **Email Address List** panel displays email addresses that have been added to previous audits. This is where you can add those existing email addresses to your current audit notification list.

To add one or more new email addresses:

1. From the Add or Change Audit Attributes panel, tab down to the **Robot HA email** address field.

2. Press **F4** Prompt. The Email Address List displays. Any email addresses that you have added to this audit will display.

3. Tab to the first available line and enter the new email address. Tab and add more if necessary.

RHA8450		E	mail Address L	ist		11:04:47
Enter v	alid Email	Addresses				
sample1	@sample.co					
sample2	@sample.com	h				
						More
F3=Exit	F4=Prompt	F5=Refresh	F17=Clear All	F21=System	Command	F22=Expand

4. Press Enter to save the list. If you have just one email address, that email address will be displayed in the **Robot HA email address** field. If you enter more than one email address, *LIST will be displayed in the field.

To select a previously used email address:

1. From the Add or Change Audit Attributes panel, tab down to the **Robot HA email** address field.

- 2. Press **F4** Prompt. The Email Address List displays. Any email addresses that you have added to this audit will display.
- 3. Press **F4** Prompt again. The Select Email Addresses panel displays. Any email addresses that you have added to previously defined audits will display.
- 4. Enter a **1** before the email addresses you would like added to the current audit, or press **F15** to select all.

RHA8000	Select Email Addresses	11:23:59
1=Multi-Select Select Email Addresses <u>bob@bob.com</u>	Start:	
F3=Exit F15=Select All	F17=Clear All F22=Expand	Bottom

- 5. Press Enter. The selected email addresses are added to the Email Address List.
- 6. Press Enter again to save the list. If you have just one email address, that email address will be displayed in the **Robot HA email address** field. If you enter more than one email address, *LIST will be displayed in the field

To clear all or some email addresses from an audit:

1. From the Add or Change Audit Attributes panel, tab down to the **Robot HA email** address field.

- 2. Press **F4** Prompt. The Email Address List displays. Any email addresses that you have previously added to this audit will display.
- 3. Delete any email address(s) you do not want included in the list OR you can delete all email addresses by pressing **F17** Clear All.
- 4. Press Enter to save the list.

Audit Notification in Robot Network

If desired, and if you have Robot Network installed on your system, you can enable the Robot Network notification on your audit. For the notification to make its way into Robot Network, you must first setup a Robot HA user application in Robot Network.

To create a user application in Robot Network:

- 1. Connect to and expand the Host system in Robot Network.
- 2. Right-click on User Applications and select New.

3. Enter ROBOT HA for the Name and add a description. The Library you should use is QGPL.

NOTE: There must be space between ROBOT and HA in the **Name** field.

User Application Properties	×			
Name: ROBOT HA				
Description: Robot HA Messages				
Library: QGPL	۵			
Escalation Options				
Send a Robot ALERT Message	Device Name:			
	Wait Time: (HH:MM:SS) 00:00:05			
Send an SNMP Trap	Q			
	Wait Time: (HH:MM:SS) 00:00:05			
Call a User Defined Program	Program Name:			
	Library:			
	Wait Time: (HH:MM:SS) 00:00:05			
Statuses to Escalate from the Hos	t			
Attention Warning Informational				
Days to retain acknowledged statuse	is: 70			
	OK Cancel			

- 4. Click **OK** to save.
- 5. Expand Nodes.

- 6. Right-click on the name of the Robot HA *PROD system that you will run audits from. Select **Properties**.
- 7. Select the **Product Statuses** tab.
- 8. Choose the statuses that you want sent to Robot Network from the Robot HA application. (Attention, Warning, Informational)

lode Information Ho	at Information Product Statuses Parfor	mance Metrics Product Metrics	Diagnostics	
elect the types of statuse	s you wish to escalate to the Host from each	product.		
Product	Description	Attention	Warning	Informational
AUTOTUNE	Robot AUTOTUNE	M	N	
BGBTEST	TEST			
역 CLIENT	Robot CLIENT	M	2	
	Robot CONSOLE			
≅⇔ LPAR	Robot LPAR	M	S	
THE MONITOR	Robot MONITOR	M	2	2
MONITOR	Monitor user app	M	1	
H NETNODE	Robat NETWORK Node	M	2	
REPORTS	Robot REPORTS	M	~	
😫 ковот	Robot SCHEDULE	N	P	
ROBOT HA	Robot HA Messages	2	2	
ROBOTSAV	Robot SAVE	2	2	
ROBOTUPS	Robot UPS	M	2	
⇔ SCHENT	Robot SCHEDULE ENTERPRISE		2	
SPACE .	Robot SPACE			
TRANSFORM	Robet TRANSFORM	M	2	
18 TRAPPER	Robot TRAPPER	M	2	
USERAPPL	USER DEFINED	M	1	

9. Click OK to save.

Swap Audit Steps

The steps defined to run in the Swap Audit are the actions that Robot HA performs to determine your readiness to run a role swap.

Step Name	Description	Can Edit Run?	Can Edit Stop?	Can Edit Rcvr?
Roles are *PROD and *BACKUP	The current system needs to be set to role *PROD and the target needs to be set to *BACKUP. Also, the TCP/IP servers must be active on both systems.			
RELMOD on target and source are equal	RELMODS must match on source and target. Both the source and target need to be at the same release.			
User defined before audit exit point	The user defined program added to this step is run. The program can pass back error information that will be written to audit history. For an example, see member BEFOREAPI in RBTHALIB/RHASAMPLE.	\checkmark	\checkmark	
*CFG synced to *BACKUP	Sync attribute for *CFG from *PROD to *BACKUP needs to exist on the production system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.	\checkmark	\checkmark	\checkmark
*CFG synced from *BACKUP	Sync attribute for *CFG from *BACKUP to *PROD needs to exist on the backup system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.	\checkmark	\checkmark	\checkmark
Product lib synced to *BACKUP	Sync attribute for library RBTHALIB from *PROD to *BACKUP needs to exist on the production system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.		\checkmark	

Step Name	Description	Can Edit Run?	Can Edit Stop?	Can Edit Rcvr?
Product lib synced from *BACKUP	Sync attribute for library RBTHALIB from *BACKUP to *PROD needs to exist on the backup system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.		\checkmark	
*USRPRF synced to *BACKUP	Sync attribute for *USRPRF from *PROD to *BACKUP needs to exist on the production system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.	\checkmark	\checkmark	\checkmark
*AUTL synced to *BACKUP	Sync attribute for *AUTL from *PROD to *BACKUP needs to exist on the production system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.	\checkmark	\checkmark	\checkmark
*PVTAUT synced to *BACKUP	Verifies that your *PVTAUT attribute exists if you have library attributes syncing object authorities. If you are syncing object authorities and the *PVTAUT attribute doesn't exist we will fail. If it is there, but not active, we will restart it.			
RSFUSER synced to *BACKUP	Sync attribute for library RSFUSER from *PROD to *BACKUP needs to exist on the production system. If not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.		\checkmark	

Step Name	Description	Can Edit Run?	Can Edit Stop?	Can Edit Rcvr?
QSTRUPPGM's are different	Retrieves the Primary and Secondary startup programs as defined in the INZRSFHA command and checks that the names and/or libraries of those programs are different. Also, verifies that the programs exist on the *PROD and *BACKUP systems.	\checkmark	\checkmark	
System startup pgm synced to *BACKUP	Retrieves the Primary and Secondary startup programs as defined in the INZRSFHA command and checks that the names and/or libraries of those programs are different. Also, verifies that the programs exist on the *PROD and *BACKUPsystems.		\checkmark	
	Sync attribute for the startup program from *PROD to *BACKUP needs to exist on the production system. If the sync entry is not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.			
System startup pgm synced from *BACKUP	Retrieves the Primary and Secondary startup programs as defined in the INZRSFHA command and checks that the names and/or libraries of those programs are different. Also, verifies that the programs exist on the *PROD and *BACKUPsystems.		\checkmark	
	from *BACKUP to *PROD needs to exist on the backup system. If the sync entry is not found, one is created and synchronization will be attempted. If found, but it is late or lagging, synchronization will be attempted.			
Groups are synced to *BACKUP	Verifies that all Group sync attribute entries from your *PROD to your *BACKUP system are not late or lagging. If they are late or lagging, synchronization will be attempted.		\checkmark	

Step Name	Description	Can Edit Run?	Can Edit Stop?	Can Edit Rcvr?
Libraries are synced to *BACKUP	Verifies that all Library sync attribute entries from your *PROD to your *BACKUP system are not late or lagging. If they are late or lagging, synchronization will be attempted.		\checkmark	
IFS directories are synced to *BACKUP	Verifies that all IFS sync attribute entries from your *PROD to your *BACKUP system are not late or lagging. If they are late or lagging, synchronization will be attempted.		\checkmark	
Identify libraries not synced	Finds libraries that do not have sync attributes set up that match the pattern defined in the step. A record of those libraries found is written to history.	\checkmark	\checkmark	
User defined after audit exit point	The user defined program added to this step is run. The program can pass back error information that will be written to audit history. For an example, see member AFTERAPI in RBTHALIB/RHASAMPLE.	\checkmark	\checkmark	

Swap Audit Stop/Recover Matrix:

See the table below for the run definition of Stop and Recover within the Audit Steps

Stop	Recover	Scenerio Description
Ν	Ν	A step failure is neither recovered nor does it stop processing. The audit process will continue to the next step and flag the failure.
Ν	Y	An attempt to recover will be made after a step failure. The audit process will continue regardless of the outcome of the recovery.
Y	Ν	A failure in the step will stop the audit process. No attempt at recovery of the step will be attempted.
Y	Y	A step failure will stop the audit process if the step cannot be recovered. However, if the step recovers successfully, the audit process will continue.

Maintain and Running an Audit

Once an audit has been created, you can submit it to run as an unscheduled audit job, edit it, or delete it from the Maintain Audits panel.

Note: You must have Audit Monitor started for a scheduled audit job to run.

To run an unscheduled audit:

- From the Robot HA Main Menu, select option 6 Audit Menu. Then, select option 1 Maintain Audits.
- 2. Enter **8** in the option field of the audit you wish to run. Press Enter.
- 3. In the resulting Confirm run of audits panel, press Enter to submit the audit to run.
- 4. After the audit job completes, Enter **5** in the option field of the audit you just ran. Press Enter.
- 5. Press Enter on the Display Audit History panel.
- 6. Enter **5** in the option field of the most recent audit run. Press Enter to view run results. See <u>Audit History</u> for details of the run results.

To edit an audit:

- From the Robot HA Main Menu, select option 6 Audit Menu. Then, select option 1 Maintain Audits.
- 2. Enter **2** in the option field of the audit you wish to edit. Press Enter.
- 3. In the resulting Change Audit Attributes panel, edit as necessary.

NOTE: The Interval fields (Interval, Units, Start, and End) control when scheduled audit jobs are run.

- 4. If desired, press **F10** Audit Steps to edit the run parameters of the audit steps.
- 5. Press Enter to save.

To delete an audit:

- From the Robot HA Main Menu, select option 6 Audit Menu. Then, select option 1 Maintain Audits.
- 2. Enter **4** in the option field of the audit you wish to delete. Press Enter.

- 3. In the resulting Confirm delete of audits panel, verify the audit you wish to delete is displayed.
- 4. Press Enter to delete.

Audit History

You can view current and historical results of audit job runs in Audit History. You can also submit an audit job to run again and delete jobs that are no longer necessary to keep in history.

To view run results of an audit:

- 1. From the Robot HA Main Menu, select option **6** Audit Menu. Then, select option **4** Audit History. The Audit History panel displays.
- 2. For the audit job which you wish to see run results, enter **5** in the option field. Press Enter.
- 3. In the resulting Audit History Detail panel, review the Audit Status field in the topright corner of the panel. Possible values are:
 - PASSED
 - FAILED
 - WARNING
- 4. Page down to view all the steps and the results of each.
- 5. Press **F22** to view the full details of each step. The full text description of the step result is displayed.
- 6. Page down to view all results. If the results of your job shows failures or warnings, use this information to correct the failure or warning and rerun the audit job.

Note: An audit failure or warning will not prevent a planned or unplanned role swap from occurring.

To rerun an audit job:

- 1. From the Robot HA Main Menu, select option **6** Audit Menu. Then, select option **4** Audit History.
- 2. Edit the parameters in the Display Audit History panel as necessary and press Enter. The Audit History panel displays.
- 3. Enter 8 in the option field of the audit you wish to run. Press Enter.

- 4. In the resulting Confirm run of audits panel, press Enter to run the audit.
- 5. Press **F5** to refresh the panel to see the new run results.

To delete an audit job:

- 1. From the Robot HA Main Menu, select option **6** Audit Menu. Then, select option **4** Audit History.
- 2. Edit the parameters in the Display Audit History panel as necessary and press Enter. The Audit History panel displays.
- 3. Enter **4** in the option field of the audit you wish to delete. Press Enter.
- 4. In the resulting Confirm delete of audits panel, press Enter to delete.

Role Swaps

In a two-system HA environment, one system acts as the production system and the other acts as the backup system. Users only interact with the production system, and replication keeps the backup system synchronized with the production system. Any changes made to the production system are automatically replicated on the backup, provided you have defined rules for replication and the data has been frequently synchronized.

A role swap switches the perspective and flow of data. For example, let's say A is the production system and B is the backup system. If a role swap is triggered, B then acts as the production system and A (when available) acts as the backup system.

For detailed information on role swaps, including considerations for your role swap, how to prepare for a role swap, and which ways you can execute a role swap, contact Technical Support.

Additional Tools

The following sections provide an overview on the additional tools Robot HA offers to help you tune your HA environment.

Many of these tools are used internally in the course of replicating your system; others are used in the Synchronization Start and Role Swap programs provided with the product. Every tool has a corresponding panel and command with extensive help text.

Compare

The following options allow you to compare replication groups, system items, and objects in your libraries or directories.

Check Libraries (CHKLIBRSF)

This panel can be viewed by taking option **3** from the <u>Integrity Reports Menu</u> or by prompting (F4) option **16** in front of a library on the <u>Work with Sync Attributes</u> panel.

The Check Libraries panel is used to determine if the objects in two libraries match. After running, a report or output file is generated, listing any unmatched objects. In a replication environment, you can use the Refresh unmatched objects parameter to cause any unmatched objects to be refreshed automatically.

<u>Check Library Objects</u> (CHKOBJRSF)

This panel can be viewed by running the command RBTHALIB/CHKOBJRSF. Library objects can also be checked by entering a specific name for the Object Name parameter on the <u>Check Libraries</u> panel.

The Check Objects panel is used to determine if two objects match.

Check Replication Group (CHKGRPRSF)

This panel can be viewed by taking option **2** from the <u>Integrity Reports Menu</u> or by prompting (F4) option **16** in front of a group on the <u>Work with Sync Attributes</u> panel.

The Check Replication Group panel is used to determine if the objects in a replication group match. After running, a report or output file is generated, listing an unmatched objects.

Check IFS Directory (CHKDIRRSF)

This panel can be viewed by taking option **1** from the <u>Integrity Reports Menu</u> or by prompting (F4) option **16** in front of a directory on the <u>Work with Sync Attributes</u> panel.

The Check IFS Directory panel is used to determine if two IFS directories, including their contents and subdirectories, match. After running, a report or output file is generated, listing any unmatched objects. In a replication environment, you can use the Refresh unmatched objects parameter to cause any unmatched objects to be refreshed automatically.

Check IFS Object (CHKIFSRSF)

This panel can be viewed by running the command RBTHALIB/CHKIFSRSF.

The Check IFS panel is used to determine if two IFS objects match.

Check System Items (CHKSYSRSF)

This panel can be viewed by taking option **4** from the <u>Integrity Reports Menu</u> or by prompting (F4) option 16 in front of *AUTL or *USRPRF on the <u>Work with Sync Attributes</u> panel.

The Check System Items panel is used to determine if system objects, such as user profiles or authorization lists, match. After running, a report or output file is generated, listing an unmatched objects. In a replication environment, you can use the Refresh unmatched objects parameter to cause any unmatched objects to be refreshed automatically.

Check Defined Items (CHKATTRSF)

This panel can be viewed by running the command RBTHALIB/CHKATTRSF.

The Check Defined Items panel is used to check the objects in some or all of the libraries and IFS directories defined on the <u>Work With Sync Attributes panel</u>. The libraries and directories are compared to their counterparts on the backup system to determine if they match.

Clean

The following panels allow you to delete orphan objects from a DR target system. Objects that are replicated without journaling are not always deleted from the backup system when

the associated object is deleted from the production system. These panels provide a way to periodically clean up those orphan objects.

Clean Library (CLNLIBRSF)

This panel can be viewed by prompting (F4) option **20** in front of a library on the <u>Work with</u> <u>Sync Attributes</u> panel.

The Clean Library panel is used on a backup system to delete objects that do not exist in the associated library on the production system

<u>Clean Replication Group</u> (CLNGRPRSF)

This panel can be viewed by prompting (F4) option **20** in front of a group on the <u>Work with</u> <u>Sync Attributes</u> panel.

The Clean Replication Group panel is used to clean the libraries in a replication group. Libraries on the server are cleaned, using libraries and directories on the production system for reference.

Clean Directory (CLNDIRRSF)

This panel can be viewed by prompting (F4) option **20** in front of a directory on the <u>Work</u> with Sync Attributes panel.

The Clean Directory panel is used on a backup system to delete objects that do not exist in the associated directory on the production system.

Clean System Items (CLNSYSRSF)

This panel can be viewed by prompting (F4) option **20** in front of *AUTL or *USRPRF on the <u>Work with Sync Attributes</u> panel.

The Clean System Items panel is used to delete authorization lists or user profiles on the backup system that do not exist on the production system.

Clean Defined Items (CLNATTRSF)

This panel can be viewed by running the command RBTHALIB/CLNATTRSF.

The Clean Defined Items panel is used to check the objects in some or all of the libraries, IFS directories, and system items defined on the <u>Work With Sync Attributes panel</u>. The libraries, directories, and system items on the server are cleaned, using libraries and directories on the production system for reference.

When you are finished, the Clean Library command will run for selected libraries, the Clean Directory command will run for selected IFS directories, and the Clean System Items command will run for selected system items.

Journal Management

Use these panels to manage your journals and remote journals.

A couple things to note before you continue:

- Robot HA can create replication journals and remote journals for you when you
 define synchronization attributes using the Change Library Sync Attributes (Option 2
 from the Work with Sync Attributes Libraries panel), Change IFS Sync Attributes
 (Option 2 from the Work with Sync Attributes IFS panel), and Change System Sync
 Attributes (Option 2 from the Work with Sync Attributes Systems panel) panels.
- Robot HA can manage replication journal receivers for you (creating new ones and deleting old ones per your specifications) by properly setting the "Robot HA manages the journal," "Change receiver every," and "Days to keep receivers" parameters on the three synchronization attribute commands mentioned above.

The following panels are currently only accessible by using one of the commands. To open a command line in Robot HA, press **F21**.

Analyze Journal Entries	The Analyze Journal Entries panel allows you to analyze the entries in a journal to determine the number of entries attributable to objects with names that begin with each letter of the alphabet. This can be useful in deciding how to divide up journaling for a library when the transaction rate is such that journaling all objects in the library to a single journal wouldn't be efficient or convenient.
Check Journal Usage	The Check Journal Usage panel is used to determine the oldest journal receiver still in use by Robot HA and the synchronization entry that is using it
	Analyze Journal Entries Check Journal Usage

Command	Panel	Description
CLRJRNRSF	Clear Journal	The Clear Journal panel is used to clear all current journal entries from a journal in a given library.
DLTJRNRSF	Delete Journal	The Delete Journal panel is used to completely delete a journal and all of its associated journal receivers.
DSPJRN	Display Journal	The IBM Display Journal panel.
ENDRJRSF	End Remote Journaling	The End Remote Journaling panel is used to suspend or remove any remote journals associated with a local journal.
RTVJRBRSF	Retrieve JRNRCV Boundary	The Retrieve JRNRCV Boundary panel is used within a CL program to determine the first journal receiver containing entries generated on or after a specific date.
WRKJRNA	Work with Journal Attributes	The IBM Work With Journal Attributes panel.

Archive or Refresh

Use these panels to save a copy of a library or directory, experiment with changes, and then restore the library or directory to its previous state.

The following panels are currently only accessible by using one of the commands. To open a command line in Robot HA, press **F21**.

Command	Panel	Description
ARCLIBRSF	Archive Library	The Archive Library panel is used to take a snapshot of a library and store it in another library. The archive can be used in the future to restore the library to the state it had at the time the archive was created.
ARCDIRRSF	Archive IFS Directory	The Archive IFS Directory panel is used to take a snapshot of an IFS directory and store it in a library. The archive can be used in the future to restore the directory to the state it had at the time the archive was created.
ARCATTRSF	Archive from Sync Attributes	The Archive from Sync Attributes panel is used to archive the libraries and IFS directories defined on your Work With Synchronization Attributes panel.

Command	Panel	Description
REFLIBRSF	Refresh Library	The Refresh Library panel is used to refresh a library from an archive that was created previously with the Archive Library panel.
REFDIRRSF	Refresh IFS Directory	The Refresh IFS Directory panel is used to refresh an IFS directory from an archive that was created previously with the Archive IFS Directory panel.
REFATTRSF	Refresh from Sync Attributes	The Refresh from Sync Attributes panel is used to refresh the libraries and IFS directories defined on your Work With Synchronization Attributes panel.

Reference

Robot HA Menus

This section contains a reference of all the menus in Robot HA.

Robot HA Main Menu

The Robot HA Main Menu provides access to the most commonly used functions in the product..

RHA1000	Robot HA Main Menu	12:47:14
R134001200-	410	*PROD
	1. Synchronization Attributes	
	2. Control Menu	
	3. Setup Menu	
	4. Reports Menu	
	5. History Menu	
	6. Audit Menu	
	7. Role Swap Menu	
	Select Option:	
F3=Exit	F21=System Command HelpSystems (C)	Copyright 🖌

Options

1. Synchronization attributes

The Synchronization attributes menu option is a quick link to the <u>Synchronization attributes</u> panel.

2. Control Menu

The <u>Control Menu</u> allows you to start and end your servers, synchronizing, audit monitor, and view different jobs and job logs.

3. Setup Menu

The <u>Setup Menu</u> offers access to server definitions, replication groups, synchronization attributes, and system defaults.

4. Reports Menu

The <u>Reports Menu</u> offers access to status reports and integrity reports.

5. History Menu

The <u>History Menu</u> allows you to view audit history, synchronization history, and the purge synchronization log.

6. Audit Menu

The <u>Audit Menu</u> allows you to maintain audits, view audit history and jobs, and allows you view the Audit History Report.

7. Role Swap Menu

The <u>Role Swap Menu</u> allows you to start a role swap.

Control Menu

The Control Menu gives you options for monitor and job control.

How to Get There

From the Robot HA Main Menu, select option **2**, Control Menu.

RHA1020 R13M001200416	Control Menu	12:47:39
K101001200410		*PROD
	TCP/IP Server is: ACTI Auditing is: INAC	VE TIVE
Monitor Control 1. Start TCP/IP Server 2. End TCP/IP Server 3. Start Synchronizing 4. End Synchronizing 5. Start Audit Monitor 6. End Audit Monitor	Job Control 10. Audit Jobs 11. Synchronization Jobs 12. RSFSRV Jobs 13. Display Synchronization Lo 14. Display Synchronization Jo	og ob Logs
	Select Option:	
F3=Exit F21=System Comm	and HelpSustems (C) Copuright

Monitor Control Options

1. Start TCP/IP server

Starts the TCP/IP server on your system. The server must be started before remote systems can access your IBM i to transfer objects, pass-through, or call remote programs.

2. End TCP/IP server

Ends the TCP/IP server on your system. Ending the server prevents clients for accessing your system.

3. Start synchronizing

Starts the Robot HA synchronization on your system. This starts replication for your environment and guarantees consistency for your replication tasks.

4. End synchronizing

Ends the Robot HA synchronization on your system.

5. Start audit monitor

Starts the Robot HA audit monitor on your system.

6. End audit monitor

Ends the Robot HA audit monitor on your system.

Job Control Options

10. Audit jobs

Displays the <u>Work with Active Jobs panel</u>, which allows you to view the performance and status information for Audit jobs that are currently active on the system. All information is gathered on a job basis.

11. Synchronization jobs

Displays the <u>Work with Active Jobs panel</u>, which allows you to view the performance and status information for Synchronization jobs that are currently active on the system. All information is gathered on a job basis.

12. RSFSRV jobs

Displays the <u>Work with User Jobs panel</u>, which allows you to view the performance and status information for RSFSRV jobs that are on the system. All information is gathered on a job basis.

13. Display Synchronization log

Opens the <u>Display Synchronization Log panel</u>, which allows you to display or print selected synchronization errors. These errors are logged on the source system whenever a problem occurs synchronizing to the target system.

14. Display Synchronization job logs

Displays the <u>Work with All Spooled Files panel</u>, which allows you to view some or all of the spooled files that are currently on the system.

Setup Menu

The System Setup Menu offers access to server definitions, maintain audits, synchronization attributes, and system defaults.

How to Get There

From the Robot HA Main Menu, select option **3**, Setup Menu.

RHA1030		Setup Menu		12:47:53
R13M001200	416			*PROD
		1. Server Definition	S	
		2. Maintain Audits		
		3. System Defaults		
		Select Option:		
F3=Exit	F21=System	Command	HelpSystems	(C) Copyright

Options

1. Server definitions

Displays the <u>Work With RSF Servers panel</u>, which allows you to view, add, change, copy, rename, or remove RSF server directory entries on your system.

2. Maintain Audits

Displays the Maintain Audits selection panel.

3. System defaults

Displays the <u>Change Product Defaults panel</u>. You can use this panel to tailor product default values to reflect the preferences of your installation.

Reports Menu

The Reports Menu provides access the Status Reports and Integrity Reports menus.

How to Get There

From the Robot HA Main Menu, select option **4**, Reports Menu.



Options

1. Status Reports

Displays the <u>Status Reports Menu</u>. From this menu, you can display statuses for groups, libraries, defined items, and more, as well as work with attributes.

2. Integrity Reports

Displays the <u>Integrity Reports Menu</u>. From this menu, you can check directories, libraries, and system items for equality, as well as compare items in replication groups.

3. Audit History Report

Displays the Audit History Report panel.

4. Synchronization Attributes Report

Displays the Synchronization Attributes Report panel.

Status Reports Menu

The Status Reports Menu allows you to display statuses for groups, libraries, defined items, and more, as well as work with attributes.

How to Get There

From the Robot HA Main Menu, select option **4**, Reports Menu, then option **1**, Status Reports.

RHA1060	Status Reports Menu	13:45:36
R13M0312103	330	* DBOD
		*PRUD
	1. Display Status for Defined Items	
	2. Display Group Status	
	3. Display IFS Status	
	4. Display Library Status	
	5. Display System Status	
	6. Private Authority Status	
	Select Option:	
F3=Exit	F21=System Command	
	HelpSystems (C)	Copyright

Options

1. Display status for defined items

Opens the <u>Display Defined Item Status panel</u>, which displays the status for some, or all, of the libraries, replication groups, IFS directories, and system information defined on the Work With Sync Attributes panel.

2. Display group status

Opens the <u>Display Group Sync Status panel</u>. Use this panel to display the status of a replication group.

3. Display IFS status

Opens the <u>Display IFS Sync Status panel</u>. Use this panel to display information about an IFS synchronization task.

4. Display library status

Opens the <u>Display Library Sync Status panel</u>. Use this panel to display information about a library synchronization task.

5. Display system status

Opens the <u>Display System Sync Status panel</u>. Use this panel to display information about a system synchronization task.

6. Private Authority Status

Opens the <u>Private Authority Status panel</u>. Use this panel to display information about authority status.

Integrity Reports Menu

The Integrity Reports Menu allows you to check directories, libraries, replication groups, and system items for equality.

How to Get There

From the Robot HA Main Menu, select option **4**, Reports Menu, then option **2**, Integrity Reports.

RHA1070	Integrity Reports Menu	12:52:36
R13M0012004	16	*PROD
	1. Check Directories for Equality	
	2. Compare Items in Replication Group	
	3. Check Libraries for Equality	
	4. Check System Items for Equality	
	Select Option:	
F3=E×it	F21=System Command	

Options

1. Check directories for equality

Displays the <u>Check IFS Directory panel</u>. Use this panel to determine if two IFS directories, their contents, and their sub-directories match.

2. Compare items in replication group

Displays the <u>Check Replication Group panel</u>. Use this panel to determine if the objects in a replication group match.

3. Check libraries for equality

Displays the <u>Check Libraries panel</u>. Use this panel to determine if the objects in two libraries match.

4. Check system items for equality

Displays the <u>Check System Items panel</u>. Use this panel to determine if system objects, such as user profiles or authorization lists, match.

History Menu

The History Menu allows you to view audit history, synchronization history, and the purge synchronization log.

How to Get There

RHA1200	History Menu	12:48:16
K15H0012004		*PROD
	1. Audit History	
	2. Synchronization History	
	3. Purge Synchronization Log	
	Select Option:	
F3=Exit	F21=System Command	
	HelpSystems (C)	Copyright 🦼

From the Robot HA Main Menu, select option **5**, History Menu.

Options

1. Audit History

Displays the <u>Audit History Panel</u>, which allows you to select the audit history you want to display.

2. Synchronization History

Opens the <u>Display Synchronization Log panel</u>, which allows you to display or print selected synchronization errors. These errors are logged on the source system whenever a problem occurs synchronizing to the target system.

3. Purge synchronization log

Opens the <u>Purge History panel</u>, which is used to purge synchronization logs. You can choose to run the purge report immediately or schedule it to run using Robot Schedule.

Audit Menu

The Audit Menu provides access to maintain audits, view audit history and jobs, and allows you access to the Audit History Report.

How to Get There

From the Robot HA Main Menu, select option **6**, Audit Menu.

RHA1090	Audit Menu	12:49:46
R130012004		*PROD
	Auditing is: INACTI	VE
	1. Maintain Audits	
	2. Start Audit Monitor	
	3. End Audit Monitor	
	4. Audit History	
	5. Audit Jobs	
	6. Audit History Report	
	Select Option:	
F3=Exit	F21=System Command HelpSystems (C)	Copyright

Options

1. Maintain Audits

Displays the <u>Maintain Audits panel</u>, which shows the Audits that have already been setup, and allows you to create new audits.

2. Start Audit Monitor

Starts the Robot HA audit monitor on your system.

3. End Audit Monitor

Ends the Robot HA audit monitor on your system.

4. Audit History

Displays the <u>Audit History Panel</u>, which allows you to select the audit history you want to display.

5. Audit Jobs

Displays the <u>Work with Active Jobs panel</u>, which allows you to view the performance and status information for Audit jobs that are currently active on the system. All information is gathered on a job basis.

6. Audit History Report

Displays the Audit History Report panel.



The Role Swap Menu allows you to start a roles swap.

How to Get There

From the Robot HA Main Menu, select option **7**, Role Swap Menu.

RHA1110 R13M0012004	Role Swap Menu	12:50:00
K10110012004		*PROD
	10. Role Swap	
	Select Option:	
F3=Exit	F21=System Command	
l	HelpSystems (C)	Copyright 🌙

Options

10. Role Swap

Use the Role Swap function to switch the perspective and role of your systems. For example, if A is the production system and B is the backup system, then a role swap will cause B to act as the production system and A (when available) to act as the backup system. See the section on role swapping for more information.

Robot HA Panels

This section contains a reference of all the panels in Robot HA.

Maintain Audit

The Maintain Audit panel is used to show details about the Audits that have already been setup, and allows you to create new audits.

How to Get There

From the Robot HA Main Menu, select option **6**, then select option **1**, Maintain Audits.

RHA1433			Maintain A	ludits		07:44:19
						*PR0D
Start at	Audit Nam	e:				
Type opt 2=Chang Audi	ions, pres e 3=Copy t	s Enter. 4=Delete	5=History 7 Aud	'=Hold∕Releas Iit	e 8=Submit Target	Now 12=Steps Status/
Opt Name AUDI	De T1 te	scription st audit	Typ S₩A	e Hel IP	d Server PROD	Recovered
F3=Exit	F5=Refres	h F6=Creat	e F9=Sort O	ptions F21=	System Comm	Bottom and

Options

You can type an option number next to one or more files. When you then press the Enter key, the function associated with the selected option is performed for each of the selected files.

OPT

The possible values are:

- 2 Change: Allows you to change the attributes of the selected audit.
- **3 Copy:** Copies the selected audit.

4 - Delete: Deletes the selected audit. When you choose this option, the Confirm Delete panel displays.

- 5 History: Displays the Audit History panel.
- 7 Hold/RIs: Allows you to hold or release the selected audit.
- 8 Submit Now: Submits the selected audit for immediate processing.

12 - Steps: Displays the <u>Audit Steps panel</u>, which shows the step descriptions and results of the selected audit.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F6 (Create): Displays the Add Audit Attributes panel, which allows you to create a new audit.

F9 (Sort Options): Displays a window to allow changing the order and content of the list shown.

F21 (System Command): Displays a system command line window.

Audit History

The Audit History panel is used to view, delete, and submit audit history records.

How to Get There

1. From the Robot HA Main Menu, select option **6**. The Audit Menu displays.

2. Select option **4** to display the Audit History panel.

You can also access the Audit History panel from the History Menu (option 1).

When you press Enter, the Audit History panel displays.

RHA14371		Audit	History		11:04:16			
Start at Co	Start at Completed Date: <u>*END</u> *END, *CURRENT, Date Time: <u>*AVAIL</u> *AVAIL, Time							
Type options	s, press Ent	ter.						
4=Delete	5=Display	8=Submit Now						
Audit	Status/		Audit	Target				
Opt Name	Recovered	Description	Туре	Server	Comp Date/Time			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:51:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:41:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:31:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:21:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:11:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 11:01:56			
AUDIT6	FAIL NO	swap	SWAP	BACKUP	09/22/17 10:51:56			
AUDIT4	FAIL NO	Swap ready Audi	SWAP	BACKUP	09/21/17 14:42:16			
AUDIT2	FAIL NO	Audit 2	SWAP	BACKUP	09/21/17 14:41:20			
AUDIT1	FAIL NO	test audit	SWAP	PROD	09/21/17 14:41:20			
AUDIT2	FAIL NO	Audit 2	SWAP	BACKUP	08/24/17 12:03:51			
AUDIT4	FAIL NO	Swap ready Audi	SWAP	BACKUP	08/24/17 12:03:04			
AUDIT4	WARN NO	Swap ready Audi	SWAP	BACKUP	08/24/17 07:42:13			
					More			
F3=Exit F5	Refresh F9	=Options F12=Ca	ancel F	21=System Com	mand			

Options

4 - Delete

The Delete Audit History record confirmation window displays.

5 - Display

Displays the <u>Audit History Detail panel</u> for the record.

8 - Submit Now

Submits immediately to batch.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (Sort Options): Displays a window to allow changing the order and content of the list shown.

F21 (System Command): Displays a system command line window.

Audit History Detail

The Audit History Detail panel lists the individual step results for the selected audit run.

How to Get There

- 1. From the Robot HA Main Menu, select option 6. The Audit Menu displays.
- 2. Select option **4** to display the Audit History panel.
- 3. Enter the option **5** for any Audit listed. Press Enter.

RHA1	438	flue	dit Hist	ory Detail		12:1	1:47
Audi	t Name:	AUDIT4	Audit	Status:	WARNING	Recovered:	NO
Huui	it Type: not Sorvor:	BACKUP	Audit	Starteu:	08/23/17	11:40:04	
rarg	let server:	BHCKUP	Huart	compteted:	00/23/1/	11:40:10	
Sev	Step Results	5 4 is starting		Reco	very/Stat	us	
	Roles and D	a is starting. Coperly set to w	DDD and	* BOC			
	RELMODS mat	ob on source and	target	*DHC			
ы	Simulating (step 0300000500	taiget.				
ω.	Simulating :	step 0300000600.					
ŵ	Simulating 9	step 0300000700.					
Ŵ.	Simulating a	step 0300000800.					
W	Simulating :	step 0300000900.					
W	Simulating s	step 0300001100.					
W	The Primary	and Secondary Si	tartup P	rogra			
W	Primary star	rtup program *SYS	SBAS/QSY	S/QST			
W	Secondary s	tartup program **	SYSBAS/Q	SYS/Q			
W	Simulating s	step 0300001500.					
W	Simulating s	step 0300001600.					
						M	lore
E3=E	it F21=Sys	stem Command F23	2=View F	ull Text			

Columns

Severity

The severity code text. Possible values are:

- ''= Info
- 'R' = Recovered
- 'W' = Warning
- 'H' = Held
- 'F' = Failed

Step Results

The text description of the step.

Recovery/Status

The Recovery action message.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F21 (System Command): Displays a system command line window.

F22 (View Full Text): Allows you to view the full text message within the Step Results column for each step.

Audit History Report Panel

The Audit History Report panel is used to view the history of an audit. A report or an output file is generated.

How to Get There

- 1. From the Robot HA Main Menu, select option **4** to display the Reports Menu.
- 2. Select option 3
- 1. From the Robot HA Main Menu, select option **6** to display the Audit Menu.
- 2. Select option 6.

RHA1410 Auc	dit History Report	11:22:07
Starting date <u>*CURRENT</u> Completion date <u>*CURRENT</u> Starting time <u>*AVAIL</u> Completion time <u>*AVAIL</u>	Date, *CURRENT, *BEGIN Date, *CURRENT Time, *AVAIL Time, *AVAIL	
Audit Type . . *ALL Audit Name . . *ALL Target Server . . *ALL Audit Status . . *ALL Recovered . . *ALL	*ALL, name, *generic*, F4=Prom *ALL, name, *generic*, F4=Prom *ALL, name, *generic*, F4=Prom *ALL, P=Passed, F=Failed, W=Wa *ALL, Y=Yes, N=No	npt npt npt arning
Output queue <u>*</u> Library ASP Group *SYSBAS	*, Name, F4=Prompt Name, F4=Prompt	
F3=Exit F4=Prompt F5=Refresh	F12=Cancel F21=System Command	

Options

Starting Date

Indicates which dates to include in the entries that are printed.

The possible values are:

*CURRENT

Entries from the current date are included.

*BEGIN

All entries from the earliest available date are included.

Date

Enter a date in your system date format. Entries added on or after the date specified are included.

Completion Date

Indicates which dates to include in the entries that are printed.

The possible values are:

*CURRENT

Entries from the current date are included.

Date

Enter a date in your system date format. Entries added on or before the date specified are included.

Starting Time

Indicates which times to include in the entries that are printed.

The possible values are:

*AVAIL

All entries from the earliest available time are included.

Time

Enter a time in hhmmss format. Entries added on or after the time specified are included.

Completion Time

Indicates which times to include in the entries that are printed.

The possible values are:

*AVAIL

All entries up to the latest available time are included.

Time

Enter a time in hhmmss format. Entries added on or before the time specified are included.

Audit Type

This is the type of audit. It can be selected from a pre-set list by using **F4** Prompt.

Audit Name

This is the name of the audit. It must be a valid name comprising of alphabetic and numeric characters plus allowed special characters "@", ".", "_", "#', and "\$". It must start with an alphabetic character. It can be selected from a pre-set list by using **F4** Prompt.

Target Server

This is the target server name. It can be selected from a pre-set list by using **F4** Prompt.

Audit Status

This is the status of the audit.

The possible values are:

*ALL

All statuses are included.

Ρ

Only those audits with a status of Passed are included.

W

Only those audits with a status of Warning are included.

F

Only those audits with a status of Failed are included.

Recovered

This is the recovery state of the audit.

The possible values are:

*ALL

All recovery states are included.

Υ

Only those audits with a recovery state of Y are included.

Ν

Only those audits with a recovery state of N are included.

Output Queue and Library

This is the output queue and library to be used for the report. Use **F4** prompt to select from a list. If you enter a library name before pressing **F4** prompt, only output queues from that library will be

listed.

The possible values are:

*

The output queue will be the default queue for the job running this request.

Name

The output will be sent to this named queue and library.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exits the panel without processing any pending changes.

F21 (System Command): Displays a system command line window.

Add Audit Attributes

The Add Audit Attributes panel allows you to create new Audits.

How to Get There

- 1. From the Robot HA Main Menu, select option **6** Audit Menu.
- 2. From the Audit Menu, select option **1** Maintain Audits panel.

3. From the Maintain Audits panel press **F6**, Create. The Add Audit Attributes screen appears.

RHA1433	Add Audit Attributes	08:26:31
		*PR0D
Audit name Audit type Audit description	(F4=Prompt)	
Target server	(F4=Prompt)	
Repeat every:	Interval <u>10</u> Units D D=Daus. H=Hours.	M=Minutes
Restrict to window:	Start <u>*NONE</u> Time, *NONE End Time	
Notification options	:	
Robot HA email add	ress (F4)	
Robot Alert device Robot Network Send to message qu	(F4) Informational _ Warning eue (F4)	_ Attention
ASP group	*SYSBAS	
Hold audit	N Y=Yes, N=No	
History runs to reta	in <u>7</u> 1-999	
E3=Exit E4=Prompt	F5=Refresh F10=Audit Steps F2 <mark>1=System C</mark>	ommand

Options

Audit Name

Enter the name of the Audit to be created. Allowed characters are alphabetic and numeric characters plus special characters "@", ".", "_", "#", and "\$". The name must start with an alphabetic character.

Audit Type

Use the **F4** Prompt command to select a type of audit. It is not possible to type a value into this field.

Audit Description

Enter a text description of the audit.

Target Server

Use the **F4** Prompt command to select the target server name. It is not possible to type a value into this field.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Audit Steps): Displays the <u>Audit Steps panel</u>, which allows you to view and edit the audit steps.

F21 (System Command): Displays a command line.

Audit Steps

The Audit Steps panel allows you to view and edit the predetermined steps that each audit will perform.

How to Get There

1. From the Robot HA Main Menu, select option **6** Audit Menu.

2. From the Audit Menu, select option **1** Maintain Audits.

3. From the Maintain Audits panel, press **F6**, Create. The Add Audit Attributes screen appears.

3. Fill in the required fields and press **F10**, Audit Steps.

RHA1435	Audit Steps			13:14:27
				*PROD
Audit name: Audit type: Target server:	AUDIT SWAP BACKUP	Audit st	ate:	released
Run Step Descri Y Roles are * Y RELMOD on f N User define Y *CFG sync'c Y *CFG sync'c Y Product lit Y Product lit Y Product lit Y *USRPRF syr Y *AUTL sync' Y RSFUSER syr Y QSTRUPPGM's	Ption PROD and *BACKUP target and source are equal d before audit exit point d to BACKUP o sync'd to BACKUP o sync'd from BACKUP oc'd to BACKUP d to BACKUP oc'd to BACKUP a are different	Step Value		Recover N N N N Y Y Y Y N More
F3=Exit F4=Pro	ompt F12=Cancel F21=System Com	mand		

Columns

Run

This specifies if the step will run. Where allowed, enter Y or N to change the behavior for this audit.

Step Description

The text description of the step.

Step Value

This specifies the step value. This value must be selected from a predefined list by using F4 prompt. Currently, SWAP-READY is the only value.

Failure - Stop

This specifies if the step should stop audit processing upon failure.

Failure - Recover

This specifies if the step should recover upon failure.

Function Keys

- F3 (Exit): Exits the panel without processing any pending changes.
- F4 (Prompt): Displays a list of possible values from which you may select one.

F12 (Cancel): Returns to the previous screen.

F21 (System Command): Displays a command line.

Add Group Sync Attributes

The Add Group Sync Attributes panel allows you to add synchronization attributes for a replication group.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, press F11 to toggle until you

get to the Group Sync Attributes.

3. Press **F6**, Create. The Add Group Sync Attributes screen appears.

RHA1032 Add Group Sync Attril	butes 14:19:01
(1 of 4) Group	Name Name, F4=Prompt
Default synchronization: Interval <u>10</u> Units <u>*MINUTES</u>	1-32767, *NONE *HOURS, *MINUTES, *SECONDS
Journaling Information	
Robot HA manages the journal <u>*YES</u> Library management options <u>*ENDJOB</u>	*YES, *NO *ENDJOB, *ENDRMTJRN, *DLTJRN
Local journal <u>*GROUP</u> Library <u>JRNLIB</u> Remote journal <u>*JRN</u> Library <u>RMTJRNLIB</u>	Name, *GROUP, *NONE, F4=Prompt _ Name Name, *NONE, *JRN, F4=Prompt _ Name
F3=Exit F4=Prompt F5=Refresh F12=Cancel F2	More 1=System Command

Options

Group

Enter the name of the replication group.

The possible values are:

Group-name: Enter a valid group name.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your machine.

Press **F4= Prompt** to select from a list of possible values.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize Groups (SYNCGRPRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the group is currently running. When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

None: The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

***HOURS:** The repeat interval is specified in hours.

***MINUTES:** The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the library being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Library Management Options

Indicates how to manage items associated with library sync attribute entries that are added to the group.

The possible values are:

***ENDJOB:** Any replication jobs started for library entries before they were added to the group are ended.

ENDRMTJRN:** In addition to the actions associated with the **ENDJOB option, remote journaling defined in library sync entries is stopped if it is not also being used by other sync attribute entries.

DLTJRN:** In addition to the actions associated with the **ENDJOB and *****ENDRMTJRN options, local, remote, and standby journals associated with the library sync entries are deleted if they are not also being used by other sync attribute entries.

***NONE:** No action is taken to manage items associated with new library sync entries added to the group.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level and data areas are synchronized at the byte level. Only the records or bytes that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file or data area results in the entire object being sent to the target.

The possible values are:

***GROUP:** The journal has the same name as the group. This is the recommended value.

***NONE:** Detail changes are not tracked. Changes to files and data areas result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per library and name the journal after the library to which it corresponds. Specifying the Robot HA product library is not recommended as journal objects stored in this library will be lost the next time Robot HA is upgraded to a new release.

Press **F4= Prompt** to select from a list of possible values.

The possible values for the library are:

JRNLIB: The journal is stored in a separate library called JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target machine.

If a value other than *NONE is specified for this parameter and the Journal parameter, Robot HA creates the remote journal library and all other needed objects automatically.

This parameter is ignored if *NONE is specified for the Journal parameter.

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- 2. The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

Disadvantages of remote journaling support:

- 1. All journal entries are sent, even those that are not needed for replication.
- 2. Filtering is not supported. *NO is assumed for Filter Journal Entries.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each syncinterval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the library the journal corresponds to and store it in a common library such as RMTJRNLIB.

Press **F4= Prompt** to select from a list of possible values.

The possible library values are:

RMTJRNLIB:This is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.

NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library for remote journal name and library. Any other value for Remote Journal will result in an error.

Library Information

Specify the library sync attributes entries to include in the group. A library entry is uniquely identified by from-lib, to-lib, server-ID and set-name.

From Library

Specify the library sync attributes entries to include in the group. A library entry is uniquely identified by from-lib, to-lib, server-ID and set-name.

The possible values are:

Library-name: Enter a valid library name.

To Library

Enter the name of the target library.

The possible values are:

Library-name: Enter a valid library name on the target system.

Library Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter a valid server ID. While a different server ID may be needed to identify the desired library sync entry, the group server ID overrides the library server ID at run time.

Set Name

Specify a set name to help identify the entry. A combination of from-library, to-library, server-ID, and set-name uniquely identifies an entry.

The possible values are:

***DFT:** The default set name is used.

Set-name: Enter a valid set name to help identify the entry.

Exists

Indicates whether this library sync attribute entry already exists.

The possible values are:

***YES:** This library sync attribute entry already exists.

***NO:** This library sync attribute entry does not exist. It will be created.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this group must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this group must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

End Journaling for Excluded Objects

Indicate whether Robot HA should end journaling for objects that are journaled to the synchronization journal but are listed as objects to omit from synchronization.

A value of *NO is assumed for this parameter if the journal is used by more than one sync attributes entry. In that case, you must end journaling for excluded objects manually after confirming that the objects are excluded from all sync attributes entries using this journal.

The possible values are:

***YES:** Robot HA will end journaling for excluded objects if they are journaled to the synchronization journal.

***NO:** Robot HA does not end journaling for excluded objects. This could result in some of these objects being mirrored, despite their being listed as excluded.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Apply Journal Entries By

Specify whether Robot HA should apply journal entries to target objects by unique key or relative record number.

The possible values are:

***KEY:** Robot HA looks for a suitable unique key access path over each target file. If one exists, record updates and deletes are done by key. Otherwise, updates are done by relative record number. Updating by key requires a small amount of extra overhead, but may be more reliable overall as relative record numbers could get jumbled between the source and target files. What's more, updating by key is a necessity if two-way or N-way synchronization is to be used. Robot HA's keyed update option allows you to use N-way synchronization if you wish. If you do not like the access path that Robot HA chooses automatically for updating a particular file, you can use the Change Unique Key Association (**CHGRSFUKEY**) command on the target machine to tell Robot HA which access path to use.

NOTE: Selecting this option requires that before and after images be journaled. If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA handles this automatically.

***RRN:** All record changes are done by relative record number.

Journal Images

Indicate the type of journal images to capture.

NOTE: If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA sets the journaling for each object automatically.

The possible values are:

***DFT:** Robot HA captures only those journal entries needed for replication. If ***KEY** is specified the **Apply Journal Entries By** parameter, both before and after images are needed. Otherwise, only after images are needed.

***BOTH:** Both before and after images are captured, regardless of the value specified for the APPLY parameter.

Disable Triggers on Target

Specify whether to disable PF trigger programs on the target machine when entire *FILE objects are sent from the source to the target machine.

When replicating systems, it is often advisable to disable PF triggers on the target machine because the changes caused by triggers on the source machine are already being replicated to the target.

Any triggers that are disabled are automatically re-enabled in the event of a role swap.

The possible values are:

***YES:** Robot HA will disable PF triggers on the target whenever an entire file is sent to the target.

***NO:** PF triggers are not disabled on the target.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

• The local journal lives on the production machine where it constantly captures the changes to be replicated.

- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES:** Robot HA creates a standby journal on the target machine.

***NO:** Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this group. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected and for the library containing the object. For physical files, member and record locks are also shown.

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Apply Retry Type

Indicate the types of operations to retry on the target machine. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

***DBOPEN:** Retry database file opens. Robot HA must open files on the target in order to replicate record-level changes.

ALLDB:** Retry the operations covered by **DBOPEN, plus all other database operations, including record-level, member-level and file-level changes.

***ALL:** Retry the operations covered by *ALLDB, plus all data area and data queue operations.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this group on the Synchronize Directories (SYNCGRPRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the grou name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this group in batch.

Remote journal validity check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate the remote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Remote journal image filter

Indicates whether the system should refrain from sending certain entries from the local to the remote journal. If you specify a value other than *NONE, IBM licensed program "HA Journal Performance" (5770SS1 feature 42) is required. This value is ignored unless journaling and remote journaling are used. *NONE must be specified for this parameter if *KEY is specified for the Apply Journal Entries By parameter.

The possible values are:

***NONE:** No filtering is done. All entries are sent to the remote journal.

*BEFORE: Record change before images are not sent to the remote journal.

NOTE: If you change this value, Robot HA must temporarily delete the remote journal from the target machine. Robot HA will recreate the remote journal automatically at the next sync interval.

SAVACT record change wait

Indicates whether to wait for a commit boundary when saving objects within this library for replication. For more information, see the SAVACTWAIT parameter on the various system SAVxxx commands.

The possible values are:

***DFT:** The default value specified for SAVACTWAIT on system save commands is used.

***LOCKWAIT:** SAVACT operations for this library wait for a commit boundary before continuing with the save. The maximum wait time is the same as the value specified for the object lock wait time on the save command, typically 120 seconds.

***NOCMTBDY:** SAVACT operations for this library do not wait for a commit boundary. This option can be safely used for replication because all entries in the journal are replicated to the target, including any rollback operation that might follow a commit. Specifying ***NOCMTBDY** may improve performance of both replication jobs and user jobs that run concurrently with replication.

***NOMAX:** SAVACT operations for this library wait forever for a commit boundary before continuing with the save.

Ignore remote error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

The possible values are:

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

YES:** Only specify **YES for this parameter if advised to do so by Help Systems Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Add IFS Sync Attributes

The Add IFS Sync Attributes panel allows you to add synchronization attributes for a source directory, target directory, and server combination.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, press **F11** to toggle until you get to the IFS Sync Attributes.
- 3. Press **F6**, Create. The Add IFS Sync Attributes screen appears.

RHA1033 (1 of 4)	Add IFS Sync Attribu	tes 14:19:3
Server ID	*FROMDIR	Name, F4=Prompt
Default synchronization: Interval Units	<u>10</u> <u>*MINUTES</u>	1-32767, *NONE *HOURS, *MINUTES, *SECONDS
Journaling Information		
Robot HA manages the journa	al <u>*YES</u>	*YES, *NO
Local journal Library Remote journal Library	· · · <u>*DIR</u> · · · <u>JRNLIB</u> · · · <u>*JRN</u> · · · <u>*JRN</u> RMTJRNLIB	Name, *DIR, *NONE, F4=Prompt Name Name, *NONE, *JRN, F4=Prompt Name
F3=Exit F4=Prompt F5=Ref	resh F12=Cancel F21	=System Command F22=Full path

Options

Directory

Enter the name of the directory on the local machine that is to be synchronized. The attributes set with this command apply only to a synchronization operation for the given Directory, To Directory, Server ID combination.

The possible values are:

Directory-name: Enter a valid directory name of up to 256 characters.

To type a path longer than 50 characters, press **F22** to display a full path and type the directory into the window provided.

Press **F4= Prompt** to select from a list of possible directories.

NOTE:

- If your 'Directory' is /QDLS, your 'To Directory' must also be /QDLS. Redirection cannot be used.
- /QDLS is considered a folder. When it is saved or restored you need to have directory entries for the user profiles ROBOTHA and RSFSRV on the production and backup systems. Use WRKDIRE to create these entries. If these entries do not exist, you will see errors in the Swap Audit.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your machine.

Press **F4= Prompt** to select from a list of possible values.

To Directory

Enter the name of the target directory for synchronization.

The possible values are:

Directory-name: Enter a valid directory path on the target system. The name may be up to 256 characters.

To type a path longer than 50 characters, press F22=Full Path and type the directory into the window provided.

***FROMDIR:** The path of the target directory is the same as the source directory.

NOTE:

- If the **To Directory** path doesn't match the **From Directory**, you are using redirection. If you are using redirection, we will not restore authorities for the parent directories. You will have to apply authorities manually on your target system.
- If the parent directories of the attribute you are defining do not exist on the target side when we start syncing, we will create the parent directories. If you are using redirection, we will not restore authorities for the parent directories. If you are not using redirection we will restore the authorities.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize

Directories (SYNCIFSRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the directory is currently running. When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

***HOURS:** The repeat interval is specified in hours.

*MINUTES: The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the directory being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level. Only the records that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file results in the entire object being sent to the target.

***NONE:** Detail changes are not tracked. Changes to files result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per directory and name the journal after the directory to which it corresponds.

***DIR:** The name of the local journal is the upper case interpreted value of the directory at top of page. The name is "I" plus the upper case alphabetic characters up to a maximum of 10 characters.

The possible library values are:

JRNLIB: The journal is stored in library JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target system

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

The disadvantage of remote journaling support is that all journal entries are sent, even those that are not needed for replication.

If a value other than *NONE is specified for this value and Local Journal, Robot HA creates the remote journal library and all other needed objects automatically. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each sync interval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the directory the journal corresponds to and store it in a common library such as RMTJRNLIB.

NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library. Any other value for Remote Journal will result in an error.

The possible library values are:

RMTJRNLIB: RMTJRNLIB is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this directory must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this directoy must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Refresh on Journal Apply Error

Specify whether to refresh objects when an error is detected applying journal entries for the object.

Though journal apply errors should not occur often, it is possible for them to occur, particularly with multi-way synchronization where a record is updated simultaneously on more than one machine.

When refreshed, the entire object is resent to the target machine, replacing the existing object. Refreshing is the easiest way to automatically re-synchronize objects that get out of sync due to failed journal entries.

The possible values are:

***DATA**: Robot HA will refresh objects only when an error is encountered applying a journal entry for the object that would change the contents of the object. Examples of entries that change the object contents are record add, change and delete; member clear, etc. Examples of entries that do not change the object contents are authority changes and changes to the object attributes.

***ALL**: Robot HA will refresh an object when any error is encountered applying changes to the target object.

***NONE**: Robot HA does not refresh objects with journal entry apply errors. A message is sent to the system operator's message queue on the target machine with information about the entry that could not be applied. It is then the operator's responsibility to re-synchronize the object manually.

Ignore move/rename errors

Indicate how to handle errors that occur when renaming or moving objects on the target machine.

For renames: If a file is renamed on the source machine immediately after it's created, the original file name may be unavailable to the journal apply job on the target. In that case, the file is created with the new name on the target and a subsequent attempt to rename the file from the original to the new name will fail. As the target file already has the new name, this error can be safely ignored in most cases. A value of *YES for this parameter will ignore this error.

For moves: Since the target directory for the move may be outside the directory being replicated, an error moving an object may or may not be important. The classic case is when an object on the source machine is moved to some directory like '/trash' which doesn't exist on the target machine. If the target directory should exist on the target machine, then this should be treated as a hard error. But if the target directory is purposefully omitted from the target machine, the failed move should be handled differently. A value of *YES for this parameter will ignore this error.

The possible values are:

*NO: All move and rename errors are treated as hard errors.

***YES:** Most move and rename errors are ignored.

If an object cannot be moved on the target machine:

- Messages are placed in the job log on the target machine noting the condition.
- The object that could not be moved is deleted from the target machine.
- No error is reported to the source machine.

If an object cannot be renamed on the target machine, and if the to-file exists and the fromfile does not:

- Messages are placed in the job log on the target machine noting the condition.
- No error is reported to the source machine.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

- The local journal lives on the production machine where it constantly captures the changes to be replicated.
- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES**: Robot HA creates a standby journal on the target machine.

***NO**: Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this directory. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected .

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this directory on the Synchronize Directories (SYNCIFSRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the directory name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this directory in batch.

Remote Journal Validity Check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the

remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate the remote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Ignore Remote Error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

YES:** Only specify **YES for this parameter if advised to do so by Fortra Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Objects to Omit

Specify objects to omit. These objects will not be synchronized. Up to 20 omit specifications may be entered.

The possible object values are:

***NONE:** No objects are omitted.

Path-name: Enter a path name for an object or objects to omit.

You can include the special characters ? and * in the last element of the path name. The ? will match any single character at that position in the name. The * will match any string of zero or more characters. Multiple instances of ? and * are allowed.

To eliminate all files in a subdirectory, end the specification with the subdirectory name as in this example:

/my_directory/notes

which will omit all files in subdirectory

'/my_directory/notes'.

NOTE: If you are syncing /QDLS with no subfolder defined, you can specify folders within QDLS to omit. If your 'From Directory' includes a subfolder in QDLS ('/QDLS/XXXX), you cannot specify an omit value.

Server ID for src from target

Enter the server ID to use on the target machine to connect back to the source machine. This value is ignored if *NONE is specified for Remote Journal.

The possible values are:

***CURRENT:** This is the recommended value. If the target machine needs to send a request back to the source machine, it uses the same connection the source machine used to initially connect to the target machine. *CURRENT may not work as expected with certain TCP/IP proxy configurations. For these cases, you can designate a server ID defined on the target machine to be used to contact this machine back when *CURRENT is specified. See the Associated Server ID (SERVER) parameter on the Add Requester Directory Entry (ADDRSFRDE) command for more information.

Server-ID: Enter a server ID that exists on the target machine and can be used to contact the source machine.

Max items to refresh

Enter the maximum number of objects Robot HA should refresh individually before refreshing the whole directory. When not using remote journaling, the value specified here is ignored and a value of 8 is used.

The possible values are:

***NOMAX** All objects with errors are refreshed individually without refreshing the whole directory.

Number: Specify the maximum number of objects to refresh individually.

QDLS ASP number

When replicating /qdls, you can specify an auxiliary storage pool number to limit the folders that are included. This value is ignored for all directories except /Qdls.

The possible values are:

***ANY:** All folders are included.

1-32: Enter an ASP number. The ASP must already be defined on your system. Only folders in the specified ASP are included.

Function Keys

- F3 (Exit): Exit the current panel without processing any pending changes.
- F4 (Prompt): Displays a list of possible values from which you may select one.
- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F12 (Cancel): Exit the current panel without processing any pending changes.
- F21 (System Command): Displays a command line.

F22 (Full path): Displays the Full Path screen, which allows you to enter an IFS path up to 256 characters.

Page Up/Page Down: Displays the previous or next parameters.

Add Library Sync Attributes

The Add Library Sync Attributes panel allows you to add synchronization attributes for a source library, target library, server ID, and set name combination.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, if you are not already on Library Sync Attributes, press **F11** to toggle until you get to the Library Sync Attributes.
- 3. Press **F6**, Create. The Add Library Sync Attributes screen appears.
| (HA1031 Add Librar | y Sync Attributes | 14:17: | 36 |
|----------------------------------|-------------------|------------------------|----|
| (1 01 4)
Library | Name, | F4=Prompt | |
| Server ID | Name, | F4=Prompt | |
| To library *FROMLI B | Name | | |
| Set name <u>*DFT</u> | Name, | *DFT | |
| Default synchronization: | | | |
| Interval | 10 1-327 | 67, *NONE | |
| Units | *MINUTES *HOUR | S, *MINUTES, *SECONDS | |
| Journaling Information | | | |
| Robot HA manages the journal | <u>*YES</u> *YES, | *N0 | |
| Local journal | *LIB Name, | *LIB, *NONE, F4=Prompt | |
| Library | JRNLIB Name | | |
| Remote journal | <u>*JRN</u> Name, | *NONE, *JRN, F4=Prompt | |
| Library | RHIJKNLIB Name | | |
| | | | |
| | | More | • |
| F3=Exit F4=Prompt F5=Refresh F12 | =Cancel F21=Syste | m Command | |

NOTE: Robot HA will not replicate any journals or journal receivers that reside within the library being synchronized by Robot HA.

Options

Library

Enter the name of the library on the local machine to be synchronized. The attributes set apply only to a synchronization operation for the given library, to-library, server-ID, and setname combination.

The possible values are:

Library-name: Enter a valid library name.

Press **F4= Prompt** to select from a list of possible values.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your machine.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your system

Press **F4=Prompt** to select from a list of available values.

To Library

Enter the name of the target library for synchronization. The attributes set apply only to a synchronization operation for the given library, to-library, server-ID, and set-name combination.

The possible values are:

Library-name: Enter a valid library name on the target system.

Set Name

Specify a set name to help identify the entry. A combination of from-library, to-library, server-ID, and set-name uniquely identifies an entry.

***DFT** : The default set name is used.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize Libraries (SYNCLIBRSF) command.

Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the library is currently running.

When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

***HOURS:** The repeat interval is specified in hours.

*MINUTES: The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the library being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level and data areas are synchronized at the byte level. Only the records or bytes that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file or data area results in the entire object being sent to the target.

The possible values are:

***LIB:** The name of the local journal will have the same name as the library in this sync attribute.

***NONE:** Detail changes are not tracked. Changes to files and data areas result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per library and name the journal after the library to which it corresponds. Specifying the Robot HA product library is not recommended as journal objects stored in this library will be lost the next time Robot HA is upgraded to a new release.

The possible library values are:

JRNLIB: The journal is stored in library JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target machine.

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- 2. The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

Disadvantages of remote journaling support:

- 1. All journal entries are sent, even those that are not needed for replication.
- 2. Filtering is not supported. *NO is assumed for Filter Journal Entries.

If a value other than *NONE is specified for this value and Local Journal, Robot HA creates the remote journal library and all other needed objects automatically. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each syncinterval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the library the journal corresponds to and store it in a common library such as RMTJRNLIB.

NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library. Any other value for Remote Journal will result in an error.

The possible library values are:

RMTJRNLIB: This is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this library must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this library must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

End Journaling for Excluded Objects

Indicate whether Robot HA should end journaling for objects that are journaled to the synchronization journal but are listed as objects to omit from synchronization.

A value of *NO is assumed for this parameter if the journal is used by more than one sync attributes entry. In that case, you must end journaling for excluded objects manually after confirming that the objects are excluded from all sync attributes entries using this journal.

The possible values are:

***YES:** Robot HA will end journaling for excluded objects if they are journaled to the synchronization journal.

***NO:** Robot HA does not end journaling for excluded objects. This could result in some of these objects being mirrored, despite their being listed as excluded.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Filter Journal Entries

Indicate whether Robot HA should examine journal entries to ensure that only entries for included objects are sent to the target system

In some cases, objects excluded from replication will still be journaled to the synchronization journal for the library. This can occur when a value other than *ALL is specified for the **Objects to Include** parameter, or a value other than *NONE is specified for the **Objects to Omit** parameter and any of the following apply:

- The journal used for synchronization is managed outside of Robot HA.
- The journal is used by another application besides Robot HA.
- *NO was specified for the **Robot HA Manages the Journal** parameter.
- *NO was specified for the End Jrn for Excluded Objects parameter.
- *NO was specified for the **Fix Objects Using Wrong Journal** parameter.

When *YES is specified for this parameter, Robot HA filters the journal entries before they are sent to the target, sending only the appropriate entries.

NOTE: Filtering journal entries may slightly effect performance.

The possible values are:

***YES:** Robot HA filters the journal entries.

***NO:** Robot HA does not filter the journal entries.

Apply Journal Entries By

Specify whether Robot HA should apply journal entries to target objects by unique key or relative record number.

The possible values are:

***KEY:** Robot HA looks for a suitable unique key access path over each target file. If one exists, record updates and deletes are done by key. Otherwise, updates are done by relative record number. Updating by key requires a small amount of extra overhead, but may be more reliable overall as relative record numbers could get jumbled between the source and target files. What's more, updating by key is a necessity if two-way or N-way synchronization is to be used. Robot HA's keyed update option allows you to use N-way synchronization if you wish. If you do not like the access path that Robot HA chooses automatically for updating a particular file, you can use the Change Unique Key Association (**CHGRSFUKEY**) command on the target machine to tell Robot HA which access path to use.

NOTE: Selecting this option requires that before and after images be journaled. If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA handles this automatically.

***RRN:** All record changes are done by relative record number.

Journal Images

Indicate the type of journal images to capture.

NOTE: If *YESIS specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA sets the journaling for each object automatically.

The possible values are:

***DFT:** Robot HA captures only those journal entries needed for replication. If ***KEY** is specified the **Apply Journal Entries By** parameter, both before and after images are needed. Otherwise, only after images are needed.

***BOTH:** Both before and after images are captured, regardless of the value specified for the APPLY parameter.

Refresh on Journal Apply Error

Specify whether to refresh objects when an error is detected applying journal entries for the object.

Though journal apply errors should not occur often, it is possible for them to occur, particularly with multi-way synchronization where a record is updated simultaneously on more than one machine.

When refreshed, the entire object is resent to the target machine, replacing the existing object. Refreshing is the easiest way to automatically re-synchronize objects that get out of sync due to failed journal entries.

The possible values are:

***DATA:** Robot HA will refresh objects only when an error is encountered applying a journal entry for the object that would change the contents of the object. Examples of entries that change the object contents are record add, change and delete; member clear, etc. Examples of entries that do not change the object contents are authority changes and changes to the object attributes.

***ALL:** Robot HA will refresh an object when any error is encountered applying changes to the target object.

***NONE:** Robot HA does not refresh objects with journal entry apply errors. A message is sent to the system operator's message queue on the target machine with information about the entry that could not be applied. It is then the operator's responsibility to re-synchronize the object manually.

Disable Triggers on Target

Specify whether to disable PF trigger programs on the target machine when entire *FILE objects are sent from the source to the target machine.

When replicating systems, it is often advisable to disable PF triggers on the target machine because the changes caused by triggers on the source machine are already being replicated to the target.

Any triggers that are disabled are automatically re-enabled in the event of a role swap.

The possible values are:

***YES:** Robot HA will disable PF triggers on the target whenever an entire file is sent to the target.

***NO:** PF triggers are not disabled on the target.

Sync Object Authorities

Indicate whether authority changes to objects should be synchronized to the target machine.

Authority changes synchronized include changes to object ownership (CHGOBJOWN) as well as changes to *PUBLIC authority, private authorities and association with an authorization list (GRTOBJAUT and RVKOBJAUT commands.)

This is not a valid option for RBTHALIB.

NOTE:

- Synchronizing object authorities will be done during the normal synchronization cycle. When the value is set to *YES, any save/restore needed for this library or objects within this library will use PVTAUT(*YES) during the synchronization cycle. When you change the value from *NO to *YES, we will process authorities for every object in the library to ensure they are correct. This will cause the first synchronization cycle after changing this value to run longer. When changing this field from *YES to *NO, any authority changes not yet synchronized will be lost. To ensure all private authority changes are applied, sync the system *PVTAUT job, then sync the library and ensure there are no authority errors for this library.
- When synchronizing object authorities, system auditing must be active and *SECRUN or *SECURITY must be specified for the QAUDLVL system value. If *YES is specified for Robot HA Manages the Journal, Robot HA will ensure system auditing is properly set.

The possible values are:

*NO: Changes to object authorities are not synchronized.

***YES:** Changes to object authorities are synchronized for this library. This is the default value.

Sync Spooled Files

Indicate whether spooled files on output queues in this library should be synchronized to the target machine.

When spooled files are synchronized:

- New spooled files are sent to the target machine.
- Deleting a spooled file from the source machine causes the matching spooled file to be deleted from the target machine.
- Moving a spooled file to or from an output queue being synchronized causes the spooled file to be moved on the target machine as well.

The target output queue name for a new spooled file is the same as the name of the queue on which the spooled file was created. The target output queue library is as specified in the "To Library" (TOLIB) parameter.

NOTE:

1. Synchronizing spooled files may make each synchronization cycle for the library run longer, especially the first time spooled files are synchronized for a library. The exact impact will depend on your configuration.

2. When synchronizing spooled files, system auditing must beactive and *SPLFDTA must be specified for the QAUDLVL system value. If *YES is specified for the "HA manages the journal" (MNGRSF) parameter, Robot HA will ensure system auditing is properly set.

The possible values are:

***NO:** Spooled files are not synchronized.

***YES:** Spooled files on output queues in this library are synchronized.

Spooled File Error Threshold

Indicate the types of spooled file errors that can occur in a given replication interval before a replication error is reported.

The possible values are:

***LOW:** Any spooled file replication error is flagged as a replication error for the library.

***MEDIUM:** Errors sending or deleting spooled files are ignored if the spooled file no longer exists on the source machine. In addition, the first occurrence of an error sending an individual spooled file is ignored. Any spooled file that could not be sent is flagged to be processed again at the next interval.

User jobs can create and delete spooled files quickly. It is possible that some spooled files could be created and deleted before Robot HA has had a chance to send them to the target machine. To avoid having a replication error reported in such a case, specify *MEDIUM or higher for this parameter.

A spooled file that is open and still being processed by the job that created it cannot be sent. To ensure that Robot HA tries to send these spooled files again later, specify *MEDIUM or higher for this parameter.

HIGH:** In addition to **MEDIUM errors, any errors moving a spooled file to a new queue are ignored. Errors sending individual spooled files are also ignored, even if the error occurs on the second or subsequent attempts to send the spooled file.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

- The local journal lives on the production machine where it constantly captures the changes to be replicated.
- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES:** Robot HA creates a standby journal on the target machine.

***NO:** Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this library. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected and for the library containing the object. For physical files, member and record locks are also shown.

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Apply Retry Type

Indicate the types of operations to retry on the target machine. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

***DBOPEN:** Retry database file opens. Robot HA must open files on the target in order to replicate record-level changes.

ALLDB:** Retry the operations covered by **DBOPEN, plus all other database operations, including record-level, member-level and file-level changes.

***ALL:** Retry the operations covered by *ALLDB, plus all data area and data queue operations.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this library on the Synchronize Libraries (SYNCLIBRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the library name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this library in batch.

Objects To Include

Specify objects to include. These objects will be the only ones synchronized. Up to 100 include specifications may be entered.

NOTE:

1. Robot HA uses the values you specify for objects to include and omit when deciding whether to start journaling for an object. However, if journaling for excluded objects is started outside of Robot HA, journal entries for those objects will be sent to the target. For performance reasons, Robot HA does not check include/omit specifications when sending journal entries to the target.

2. Omitted objects are sent to the target machine when the whole library is sent to establish an initial synchronization boundary unless a) CLEAR(*NO) is specified on the Synchronize Libraries (SYNCLIBRSF) command, or b) the source and target library names are not the same.

The possible single values are:

***JRN:** All objects in the library being replicated that are journaled to the journal specified on the Journal parameter (JRN) are included. The set of included objects updates dynamically as objects are added to or removed from the specified journal.

NOTE: Note: If this library entry is included in a group, then the group journal overrides the library journal. In that case, the group journal is used to determine what objects are included. You cannot specify *JRN for a library included in a group.

When *JRN is specified for this parameter, the following are required:

- *NONE cannot be specified for the Journal (JRN) parameter.
- *NONE cannot be specified for the Remote journal (RMTJRN) parameter.
- *NONE must be specified for the "Objects to omit" (OMITOBJ) parameter.

***OTHER:** Objects in the library being replicated are included in this set if they are not included in any other specification for the set collection. A set collection is all specifications with the same from-lib, to-lib and server-ID values.

When *OTHER is specified for this parameter, the following is required:

 If a value other than *NONE is specifiec for the Journal (JRN) parameter, then a value other than *NONE must be specified for the Remote journal (RMTJRN) parameter.

The possible object values are:

***ALL:** All objects are included. All objects in the library that are not omitted by the "Objects to Omit" (OMITOBJ) parameter will be synchronized.

Generic-name: Enter a generic name, ending with an asterisk (*). Objects whose names begin as specified are included.

Name: Enter the name of a specific object to include.

The possible object type values are:

***ALL:** All object types are included if their name matches

the specified value.

Object-type: Enter a valid object type. Objects are included if their name and type match the values specified.

Objects to Omit

Specify objects to omit. These objects will not be synchronized. Up to 100 omit specifications may be entered.

NOTE: 1. *DTAQ objects are not sent unless the entire library is sent and *YES is specified for the Clear Target Library First (CLEAR) parameter on the Synchronize Libraries (SYNCLIBRSF) command.

2. Robot HA uses the values you specify for objects to include and omit when deciding whether to start journaling for an object. However, if journaling for excluded objects is started outside of Robot HA, journal entries for those objects will be sent to the target. For performance reasons, Robot HA does not check include/omit specifications when sending journal entries to the target.

3. Omitted objects are sent to the target machine when the whole library is sent to establish an initial synchronization boundary unless a) CLEAR(*NO) is specified on the Synchronize Libraries (SYNCLIBRSF) command, or b) the source and target library names are not the same.

The possible object values are:

***NONE:** No objects are omitted.

***ALL:** All objects of the specified type are omitted.

Generic-name: Enter a generic name, ending with an asterisk (*). Objects whose names begin as specified are omitted.

Name: Enter the name of a specific object to omit.

The possible object type values are:

***ALL:** All object types are omitted if their name matches the specified value.

Object-type: Enter a valid object type. Objects are omitted if their name and type match the values specified.

Max items to refresh

Enter the maximum number of objects Robot HA should refresh individually before refreshing the whole library. When not using remote journaling, the value specified here is ignored and a value of 12 is used.

The possible values are:

***NOMAX** All objects with errors are refreshed individually without refreshing the whole library.

Number: Specify the maximum number of objects to refresh individually.

Sync Non-Journaled Objects

Indicates whether Robot HA should replicate non-journaled objects in this library. Only database files, data areas and data queues can be replicated with journaling. All other types of objects must be replicated using save/restore operations. In general, you should specify *YES for this parameter to ensure that all objects in the library are replicated. However, for large libraries with many objects and few non-journaled objects that change, specifying *NO can improve performance.

The possible values are:

***YES:** Non-journaled objects are replicated.

***NO:** Only journaled objects are replicated for this library.

Remote Journal Validity Check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate theremote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Remote Journal Image Filter

Indicates whether the system should refrain from sending certain entries from the local to the remote journal. If you specify a value other than *NONE, IBM licensed program "HA Journal Performance" (5770SS1 feature 42) is required. This value is ignored unless journaling and remote journaling are used. *NONE must be specified for this parameter if *KEY is specified for the Apply Journal Entries By parameter.

The possible values are:

***NONE:** No filtering is done. All entries are sent to the remote journal.

*BEFORE: Record change before images are not sent to the remote journal.

NOTE: If you change this value, Robot HA must temporarily delete the remote journal from the target machine. Robot HA will recreate the remote journal automatically at the next sync interval.

Object Count Optimization

Specifies the optimization to use when processing objects in the library. Some steps in the replication cycle require internally listing the objects in the library. It is most efficient to list objects to a user space, but a user space can accommodate a maximum of about 74k objects.

The possible values are:

***NORMAL:** When objects in the library are listed, they are listed first to a user space. If there are too many objects to fit in the user space, a second attempt is automatically made to list the objects to a database file. This will be the most efficient option for most libraries.

***LARGE:** When objects in the library are listed, they are listed directly to a database file. No attempt is made to list them to a user space. Use this option if you know the library contains a large number of (> 74k) objects.

Object Option

Indicates the class of objects to be sent from this machine to a target machine for replication. This value should be set to *ALL except when performing two-way replication for a library. When a library is being replicated in both directions, it may be desirable to limit the objects that can be sent to the other machine to those originally created on this machine. Various conditions can trigger the sending of whole objects, including adding new objects to the library, refreshing objects that are out of sync, etc.

The possible values are:

***ALL:** All objects due to be sent are sent, regardless of the system on which they were originally created.

***LCL:** Only objects originally created on the local machine might be sent. Objects originally created on another machine are never sent.

SAVACT Record Change Wait

Indicates whether to wait for a commit boundary when saving objects within this library for replication. For more information, see the SAVACTWAIT parameter on the various system SAVxxx commands.

The possible values are:

***DFT:** The default value specified for SAVACTWAIT on system save commands is used.

***LOCKWAIT:** SAVACT operations for this library wait for a commit boundary before continuing with the save. The maximum wait time is the same as the value specified for the object lock wait time on the save command, typically 120 seconds.

***NOCMTBDY:** SAVACT operations for this library do not wait for a commit boundary. This option can be safely used for replication because all entries in the journal are replicated to the target, including any rollback operation that might follow a commit. Specifying ***NOCMTBDY** may improve performance of both replication jobs and user jobs that run concurrently with replication.

***NOMAX:** SAVACT operations for this library wait forever for a commit boundary before continuing with the save.

To ASP Device

Use this value when the target library is in a different iASP than the source library.

The possible values are:

***SAVASPDEV:** The target library is not in an iASP, or the name of the iASP containing the target library is the same as the iASP of the source library.

Name: Enter the name of the iASP that contains the target library.

To ASP

Use this value when the target library is in a different user ASP than the source library.

The possible values are:

***SAVASP:** The target library is in the same ASP as the source library.

1-16: Enter 1 for the system ASP, or 2-16 for a user ASP.

NOTE: If a name is specified for To ASP Device, a value of *SAVASP is assumed.

Replication Group

Allows you to associate this entry with a replication group. When associated with a group, sync attributes at the group level override similar attributes specified at the library level. You can also use the Change Group Sync Attributes to associate library sync entries with groups.

***NONE:** The entry is not associated with a group.

Name: Enter an existing replication group name.

Press **F4= Prompt** to select from a list of possible values.

Ignore Remote Error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

The possible values are:

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

YES:** Only specify **YES for this parameter if advised to do so by Help Systems Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Add System Sync Attributes

The Add System Sync Attributes panel allows you to add synchronization attributes for system information such as user profiles and system values.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization attributes.
- 2. From the Work with Synchronization Attributes panel, press **F11** to toggle until you get to the System Sync Attributes.
- 3. Press **F6**, Create. The Add System Sync Attributes screen appears.

RHA1034 Add System Sync Attributes 1 (1 = 5.2) 1		
(1 of 3) Server ID	Name, F4=Prompt *AUTL, *CFG, *NETA, *SYSVAL	
Default synchronization: Interval <u>10</u> Units <u>*MINUTES</u>	1-32767, *NONE *HOURS, *MINUTES, *SECONDS	
Journaling Information		
Robot HA manages the journal <u>*YES</u>	*YES, *NO	
F3=Exit F4=Prompt F5=Refresh F12=Cancel F21	=System Command	
	-	

Options

Available options may differ depending on the type of system synchronization attribute being added.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your machine.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your system.

Press **F4= Prompt** to select from a list of possible values.

Type of Information

Indicate the type of information to be synchronized.

The attributes set with this command apply only to a synchronization operation for the given type and server ID combination. The possible values are:

***USRPRF:** User profile synchronization attributes are set.

NOTE: User profiles with TOTP keys for multifactor authentication, introduced in IBM i 7.6, are synchronized with Robot HA if both systems are at a configuration that supports TOTP keys.

***SYSVAL:** System value synchronization attributes are set.

***NETA:** Network attribute synchronization attributes are set.

***PVTAUT:** Private authority synchronization attributes are set.

***AUTL:** Authorization list synchronization attributes are set.

***CFG:** Configuration synchronization attributes are set.

NOTE:

- Configuration information is handled differently than other types of system synchronization. Configuration information is extracted to the library specified by the **Config Extract Library** parameter and replicated to a library by the same name on the target machine. However, while it is stored on the target machine, the configuration is not applied to the target machine until a role swap is performed or the Apply Configuration Info is run.
- The *PVTAUT sync job works with the library and group sync jobs to ensure object authority is synchronized. It gathers object authority changes from QAUDJRN and then the library and group sync process will apply the changes to the backup.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize System Info (SYNCSYSRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the attribute is currently running. When a non-zero repeat interval is used, the synchronization repeats at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

*HOURS: The repeat interval is specified in hours.

***MINUTES:** The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the system information being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *****NONE is specified for Local Journal.

***YES:** Robot HA manages the journal and journal receivers.

*NO: The journal must be managed outside Robot HA.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this entry must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this entry must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this library on the Synchronize System Information (SYNCSYSRSF) command. Defaults to S_nnnnnn where nnnnnn is the type of information.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this entry in batch.

Include or Omit Items

Indicate whether items listed on the **Items to Include/Omit** parameter should be included in or omitted from synchronization.

NOTE: Include and Omit parameters are not valid with *PVTAUT, *CFG, and *NETA.

The possible values are:

***OMIT:** The items listed are omitted. All other items are included.

***INCLUDE:** The items listed are included. All other items are omitted.

Items to Include/Omit

Specify items to include or omit from synchronization. Whether the specified items are included or omitted is determined by the **Include or Omit Items** parameter.

The possible object values are:

Generic-name: Enter a generic name, ending with an asterisk (*). Items whose names begin as specified are included or omitted.

Name: Enter the name of a specific item to omit.

IMPORTANT: The following is a list of system values that are always omitted from syncing along with the earliest version of Robot HA where the omit was applicable.

System Value Name	Earliest Robot HA Version
QABNORMSW	12.01
QACGLVL	13.09
QAUDCTL	13.09
QCONSOLE	12.01
QDAYOFWEEK	12.01
QIGC	12.01
QIPLSTS	12.01
QLOCALE	13.09
QMODEL	12.01
QPRCFEAT	12.01
QPWDLVL	13.09
QPWRRSTIPL	13.09
QRMTIPL	13.09
QSRLNBR	12.01
QSTRPRTWTR	12.01
QUTCOFFSET	12.01

IMPORTANT: Some system values have a special handling starting in Robot HA version 13.09.

QPWDLMTAJC, QPWDLMTCHR, QPWDLMTREP, QPWDMAXLEN, QPWDMINLEN, QPWDPOSDIF, QPWDRQDDGT

The syncing of these is determined based on the contents of QPWDRULES. If QPWDRULES is set to *PWDSYSVAL, we will sync these, otherwise they will be ignored.

QSSLCSL

The syncing of this is determined based on the contents of QSSLCSLCTL. If QSSLCSLCTL is set to *OPSYS, QSSLCSL is treated as though it were read-only, and is therefore not synced.

Max Journal Entries to Process

Enter the maximum number of outstanding QAUDJRN entries to process for this type of information before completely refreshing all items. If the processing of journal entries gets too far behind, it may be more efficient to use current values to completely refresh the information for all items of the specified type. The processing of journal entries may fall behind for various reasons, including:

- Replication for the system item has not run for a while.
- Replication for the item is repeatedly ending in error.
- Your settings for system values QAUDCTL, QAUDLVL and QAUDLVL2 are causing the system audit journal, QAUDJRN, to grow very rapidly.

The possible values are:

*NOMAX: All journal entries are processed, regardless of how many remain outstanding.

Number: Enter the maximum number of unprocessed journal entries allowed before Robot HA discards the entries and uses current values to completely refresh the items.

NOTE: Normal processing of journal entries will resume with the cycle following the one in which all items are refreshed.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Change Library Sync Attributes

The Change Library Sync Attributes panel is used to change the synchronization attributes for a given source library, target library, server-ID, and set-name combination.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, if you are not already on Library Sync Attributes, press **F11** to toggle until you get to the Library Sync

Attributes.

3. Enter option **2** next to a library sync attribute and press Enter.

RHA1031 Change	Library Sync Att	ributes	10:12:31
(1 of 4)			
Library <u>ARC</u> H	IIVE	Name, F4=Prompt	
Server ID MYBA	ACKUP	Name, F4=Prompt	
To library ARCH		Name	
Set name <u>SETR</u>	REST	Name, *DFT	
Default sunchronization:			
Interval	. 10	1-32767, *NONE	
Units	. * MINUTES	*HOURS, *MINUTES, >	*SECONDS
Journaling Information			
	**		
Robot HH Manages the journal .	. <u>*1E5</u>	*YES, *NU	
Local journal	. ARCHREST	Name, *LIB, *NONE,	F4=Prompt
Library	. JRNLIB	Name	
Remote journal	. ARCHREST	Name, *NONE, *JRN,	F4=Prompt
Library	. RMTJRNLIB	Name	
Last receiver processed	. <u>ARCHRE0006</u>	Name, *NONE	
Last sequence number processed	. 97	Number	
			More
F3=Exit F4=Prompt F5=Refresh	F12=Cancel F21	=System Command	

NOTE: Robot HA will not replicate any journals or journal receivers that reside within the library being synchronized by Robot HA.

Options

Library

Enter the name of the library on the local machine to be synchronized. The attributes set with this panel apply only to a synchronization operation for the given library, to-library, server-ID, and set-name combination.

The possible values are:

Library-name: Enter a valid library name.

Press **F4= Prompt** to select from a list of possible values.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your system

Press **F4=Prompt** to select from a list of available values.

To Library

Enter the name of the target library for synchronization. The attributes set with this panel apply only to a synchronization operation for the given library, to-library, server-ID, and setname combination.

The possible values are:

Library-name: Enter a valid library name on the target system.

Set Name

Specify a set name to help identify the entry. A combination of from-library, to-library, server-ID, and set-name uniquely identifies an entry.

The possible values are:

***DFT:** The default set name is used.

Set-name: Enter a valid set name to help identify the entry.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize Libraries (SYNCLIBRSF) command.

Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the library is currently running.

When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

***HOURS:** The repeat interval is specified in hours.

*MINUTES: The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the library being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level and data areas are synchronized at the byte level. Only the records or bytes that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file or data area results in the entire object being sent to the target.

The possible values are:

***LIB:** The name of the local journal will have the same name as the library in this sync attribute.

***NONE:** Detail changes are not tracked. Changes to files and data areas result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per library and name the journal after the library to which it corresponds. Specifying the Robot HA product library is not recommended as journal objects stored in this library will be lost the next time Robot HA is upgraded to a new release.

The possible library values are:

JRNLIB: The journal is stored in library JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target system

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- 2. The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

Disadvantages of remote journaling support:

- 1. All journal entries are sent, even those that are not needed for replication.
- 2. Filtering is not supported. *NO is assumed for Filter Journal Entries.

If a value other than *NONE is specified for this value and Local Journal, Robot HA creates the remote journal library and all other needed objects automatically. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each syncinterval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the library the journal corresponds to and store it in a common library such as RMTJRNLIB.

NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library for remote journal name and library. Any other value for Remote Journal will result in an error.

The possible library values are:

RMTJRNLIB: RMTJRNLIB is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.

Last receiver processed.

The last journal receiver processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

***NONE:** No journal entries have been processed. The next synchronization point begins with the first available journal entry.

Name: Enter the name of the last journal receiver processed. This, along with the "Last Sequence Number Processed" determines where the next synchronization point begins.

Last Sequence Number Processed

The last sequence number processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

Number: Enter the number of the last journal sequence number processed. This, along with the "Last Receiver Processed" determines where the next synchronization point begins.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this library must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this library must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

End Journaling for Excluded Objects

Indicate whether Robot HA should end journaling for objects that are journaled to the synchronization journal but are listed as objects to omit from synchronization.

A value of *NO is assumed for this parameter if the journal is used by more than one sync attributes entry. In that case, you must end journaling for excluded objects manually after confirming that the objects are excluded from all sync attributes entries using this journal.

The possible values are:

***YES:** Robot HA will end journaling for excluded objects if they are journaled to the synchronization journal.

***NO:** Robot HA does not end journaling for excluded objects. This could result in some of these objects being mirrored, despite their being listed as excluded.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Filter Journal Entries

Indicate whether Robot HA should examine journal entries to ensure that only entries for included objects are sent to the target system
In some cases, objects excluded from replication will still be journaled to the synchronization journal for the library. This can occur when a value other than *ALL is specified for the **Objects to Include** parameter, or a value other than *NONE is specified for the **Objects to Omit** parameter and any of the following apply:

- The journal used for synchronization is managed outside of Robot HA.
- The journal is used by another application besides Robot HA.
- *NO was specified for the **Robot HA Manages the Journal** parameter.
- *NO was specified for the End Jrn for Excluded Objects parameter.
- *NO was specified for the **Fix Objects Using Wrong Journal** parameter.

When *YES is specified for this parameter, Robot HA filters the journal entries before they are sent to the target, sending only the appropriate entries.

NOTE: Filtering journal entries may slightly effect performance.

The possible values are:

***YES:** Robot HA filters the journal entries.

***NO:** Robot HA does not filter the journal entries.

Apply Journal Entries By

Specify whether Robot HA should apply journal entries to target objects by unique key or relative record number.

The possible values are:

***KEY:** Robot HA looks for a suitable unique key access path over each target file. If one exists, record updates and deletes are done by key. Otherwise, updates are done by relative record number. Updating by key requires a small amount of extra overhead, but may be more reliable overall as relative record numbers could get jumbled between the source and target files. What's more, updating by key is a necessity if two-way or N-way synchronization is to be used. Robot HA's keyed update option allows you to use N-way synchronization if you wish. If you do not like the access path that Robot HA chooses automatically for updating a particular file, you can use the Change Unique Key Association (**CHGRSFUKEY**) command on the target machine to tell Robot HA which access path to use.

NOTE: Selecting this option requires that before and after images be journaled. If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA handles this automatically.

***RRN:** All record changes are done by relative record number.

Journal Images

Indicate the type of journal images to capture.

NOTE: If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA sets the journaling for each object automatically.

The possible values are:

DFT:** Robot HA captures only those journal entries needed for replication. If **KEY is specified the **Apply Journal Entries By** parameter, both before and after images are needed. Otherwise, only after images are needed.

***BOTH:** Both before and after images are captured, regardless of the value specified for the APPLY parameter.

Refresh on Journal Apply Error

Specify whether to refresh objects when an error is detected applying journal entries for the object.

Though journal apply errors should not occur often, it is possible for them to occur, particularly with multi-way synchronization where a record is updated simultaneously on more than one machine.

When refreshed, the entire object is resent to the target machine, replacing the existing object. Refreshing is the easiest way to automatically re-synchronize objects that get out of sync due to failed journal entries.

The possible values are:

***DATA:** Robot HA will refresh objects only when an error is encountered applying a journal entry for the object that would change the contents of the object. Examples of entries that change the object contents are record add, change and delete; member clear, etc. Examples of entries that do not change the object contents are authority changes and changes to the object attributes.

***ALL:** Robot HA will refresh an object when any error is encountered applying changes to the target object.

***NONE:** Robot HA does not refresh objects with journal entry apply errors. A message is sent to the system operator's message queue on the target machine with information about

the entry that could not be applied. It is then the operator's responsibility to re-synchronize the object manually.

Disable Triggers on Target

Specify whether to disable PF trigger programs on the target machine when entire *FILE objects are sent from the source to the target machine.

When replicating systems, it is often advisable to disable PF triggers on the target machine because the changes caused by triggers on the source machine are already being replicated to the target.

Any triggers that are disabled are automatically re-enabled in the event of a role swap.

The possible values are:

***YES:** Robot HA will disable PF triggers on the target whenever an entire file is sent to the target.

***NO:** PF triggers are not disabled on the target.

Sync Object Authorities

Indicates whether authority changes to objects should be synchronized to the target machine. Authority changes synchronized include changes to object ownership (CHGOBJOWN) as well as changes to *PUBLIC authority, private authorities and association with an authorization list (GRTOBJAUT and RVKOBJAUT commands).

NOTE:

- Synchronizing object authorities will be done during the normal synchronization cycle. When the value is set to *YES, any save/restore needed for this library or objects within this library will use PVTAUT(*YES) during the synchronization cycle. When you change the value from *NO to *YES, we will process authorities for every object in the library to ensure they are correct. This will cause the first synchronization cycle after changing this value to run longer. When changing this field from *YES to *NO, any authority changes not yet synchronized will be lost. To ensure all private authority changes are applied, sync the system *PVTAUT job, then sync the library and ensure there are no authority errors for this library.
- When synchronizing object authorities, system auditing must be active and *SECRUN or *SECURITY must be specified for the QAUDLVL system value. If *YES is specified for Robot HA Manages the Journal, Robot HA will ensure system auditing is properly set.

The possible values are:

***NO:** Changes to object authorities are not synchronized.

***YES:** Changes to object authorities are synchronized for this library.

Sync Spooled Files

Indicate whether spooled files on output queues in this library should be synchronized to the target machine.

When spooled files are synchronized:

- New spooled files are sent to the target machine.
- Deleting a spooled file from the source machine causes the matching spooled file to be deleted from the target machine.
- Moving a spooled file to or from an output queue being synchronized causes the spooled file to be moved on the target machine as well.

The target output queue name for a new spooled file is the same as the name of the queue on which the spooled file was created. The target output queue library is as specified in the "To Library" (TOLIB) parameter.

NOTE:

1. Synchronizing spooled files may make each synchronization cycle for the library run longer, especially the first time spooled files are synchronized for a library. The exact impact will depend on your configuration.

2. When synchronizing spooled files, system auditing must beactive and *SPLFDTA must be specified for the QAUDLVL system value. If *YES is specified for the "HA manages the journal" (MNGRSF) parameter, Robot HA will ensure system auditing is properly set.

The possible values are:

***NO:** Spooled files are not synchronized.

***YES:** Spooled files on output queues in this library are synchronized.

Spooled File Error Threshold

Indicate the types of spooled file errors that can occur in a given replication interval before a replication error is reported.

The possible values are:

*LOW: Any spooled file replication error is flagged as a replication error for the library.

***MEDIUM:** Errors sending or deleting spooled files are ignored if the spooled file no longer exists on the source machine. In addition, the first occurrence of an error sending an individual spooled file is ignored. Any spooled file that could not be sent is flagged to be processed again at the next interval.

User jobs can create and delete spooled files quickly. It is possible that some spooled files could be created and deleted before Robot HA has had a chance to send them to the target machine. To avoid having a replication error reported in such a case, specify *MEDIUM or higher for this parameter.

A spooled file that is open and still being processed by the job that created it cannot be sent. To ensure that Robot HA tries to send these spooled files again later, specify *MEDIUM or higher for this parameter.

HIGH:** In addition to **MEDIUM errors, any errors moving a spooled file to a new queue are ignored. Errors sending individual spooled files are also ignored, even if the error occurs on the second or subsequent attempts to send the spooled file.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

- The local journal lives on the production machine where it constantly captures the changes to be replicated.
- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES:** Robot HA creates a standby journal on the target machine.

***NO:** Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this library. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected and for the library containing the object. For physical files, member and record locks are also shown.

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Apply Retry Type

Indicate the types of operations to retry on the target machine. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

***DBOPEN:** Retry database file opens. Robot HA must open files on the target in order to replicate record-level changes.

ALLDB:** Retry the operations covered by **DBOPEN, plus all other database operations, including record-level, member-level and file-level changes.

***ALL:** Retry the operations covered by *ALLDB, plus all data area and data queue operations.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the

end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this library on the Synchronize Libraries (SYNCLIBRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the library name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this library in batch.

Objects To Include

Specify objects to include. These objects will be the only ones synchronized. Up to 100 include specifications may be entered.

NOTE:

1. Robot HA uses the values you specify for objects to include and omit when deciding whether to start journaling for an object. However, if journaling for excluded objects is started outside of Robot HA, journal entries for those objects will be sent to the target. For performance reasons, Robot HA does not check include/omit specifications when sending journal entries to the target.

2. Omitted objects are sent to the target machine when the whole library is sent to establish an initial synchronization boundary unless a) CLEAR(*NO) is specified on the Synchronize Libraries (SYNCLIBRSF) command, or b) the source and target library names are not the same.

The possible single values are:

***JRN:** All objects in the library being replicated that are journaled to the journal specified on the Journal parameter (JRN) are included. The set of included objects updates dynamically as objects are added to or removed from the specified journal.

NOTE: Note: If this library entry is included in a group, then the group journal overrides the library journal. In that case, the group journal is used to determine what objects are included. You cannot specify *JRN for a library included in a group.

When *JRN is specified for this parameter, the following are required:

- *NONE cannot be specified for the Journal (JRN) parameter.
- *NONE cannot be specified for the Remote journal (RMTJRN) parameter.
- *NONE must be specified for the "Objects to omit" (OMITOBJ) parameter.

***OTHER:** Objects in the library being replicated are included in this set if they are not included in any other specification for the set collection. A set collection is all specifications with the same from-lib, to-lib and server-ID values.

When *OTHER is specified for this parameter, the following is required:

 If a value other than *NONE is specifiec for the Journal (JRN) parameter, then a value other than *NONE must be specified for the Remote journal (RMTJRN) parameter.

The possible object values are:

***ALL:** All objects are included. All objects in the library that are not omitted by the "Objects to Omit" (OMITOBJ) parameter will be synchronized.

Generic-name: Enter a generic name, ending with an asterisk (*). Objects whose names begin as specified are included.

Name: Enter the name of a specific object to include.

The possible object type values are:

***ALL:** All object types are included if their name matches

the specified value.

Object-type: Enter a valid object type. Objects are included if their name and type match the values specified.

Objects to Omit

Specify objects to omit. These objects will not be synchronized. Up to 100 omit specifications may be entered.

NOTE: 1. *DTAQ objects are not sent unless the entire library is sent and *YES is specified for the Clear Target Library First (CLEAR) parameter on the Synchronize Libraries (SYNCLIBRSF) command.

2. Robot HA uses the values you specify for objects to include and omit when deciding whether to start journaling for an object. However, if journaling for excluded objects is started outside of Robot HA, journal entries for those objects will be sent to the target. For performance reasons, Robot HA does not check include/omit specifications when sending journal entries to the target.

3. Omitted objects are sent to the target machine when the whole library is sent to establish an initial synchronization boundary unless a) CLEAR(*NO) is specified on the Synchronize Libraries (SYNCLIBRSF) command, or b) the source and target library names are not the same.

The possible object values are:

***NONE:** No objects are omitted.

***ALL:** All objects of the specified type are omitted.

Generic-name: Enter a generic name, ending with an asterisk (*). Objects whose names begin as specified are omitted.

Name: Enter the name of a specific object to omit.

The possible object type values are:

*ALL: All object types are omitted if their name matches the specified value.

Object-type: Enter a valid object type. Objects are omitted if their name and type match the values specified.

Max items to refresh

Enter the maximum number of objects Robot HA should refresh individually before refreshing the whole library. When not using remote journaling, the value specified here is ignored and a value of 12 is used.

The possible values are:

***NOMAX:** All objects with errors are refreshed individually without refreshing the whole library.

Number: Specify the maximum number of objects to refresh individually.

Sync Non-Journaled Objects

Indicates whether Robot HA should replicate non-journaled objects in this library. Only database files, data areas and data queues can be replicated with journaling. All other types of objects must be replicated using save/restore operations. In general, you should specify *YES for this parameter to ensure that all objects in the library are replicated. However, for large libraries with many objects and few non-journaled objects that change, specifying *NO can improve performance.

The possible values are:

***YES:** Non-journaled objects are replicated.

***NO:** Only journaled objects are replicated for this library.

Remote Journal Validity Check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate theremote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Remote Journal Image Filter

Indicates whether the system should refrain from sending certain entries from the local to the remote journal. If you specify a value other than *NONE, IBM licensed program "HA Journal Performance" (5770SS1 feature 42) is required. This value is ignored unless journaling and remote journaling are used. *NONE must be specified for this parameter if *KEY is specified for the Apply Journal Entries By parameter.

The possible values are:

***NONE:** No filtering is done. All entries are sent to the remote journal.

*BEFORE: Record change before images are not sent to the remote journal.

NOTE: If you change this value, Robot HA must temporarily delete the remote journal from the target machine. Robot HA will recreate the remote journal automatically at the next sync interval.

Object Count Optimization

Specifies the optimization to use when processing objects in the library. Some steps in the replication cycle require internally listing the objects in the library. It is most efficient to list objects to a user space, but a user space can accommodate a maximum of about 74k objects.

The possible values are:

***NORMAL:** When objects in the library are listed, they are listed first to a user space. If there are too many objects to fit in the user space, a second attempt is automatically made to list the objects to a database file. This will be the most efficient option for most libraries.

***LARGE:** When objects in the library are listed, they are listed directly to a database file. No attempt is made to list them to a user space. Use this option if you know the library contains a large number of (> 74k) objects.

Object Option

Indicates the class of objects to be sent from this machine to a target machine for replication. This value should be set to *ALL except when performing two-way replication for a library. When a library is being replicated in both directions, it may be desirable to limit the objects that can be sent to the other machine to those originally created on this machine. Various conditions can trigger the sending of whole objects, including adding new objects to the library, refreshing objects that are out of sync, etc.

The possible values are:

***ALL:** All objects due to be sent are sent, regardless of the system on which they were originally created.

***LCL:** Only objects originally created on the local machine might be sent. Objects originally created on another machine are never sent.

SAVACT Record Change Wait

Indicates whether to wait for a commit boundary when saving objects within this library for replication. For more information, see the SAVACTWAIT parameter on the various system SAVxxx commands.

The possible values are:

***DFT:** The default value specified for SAVACTWAIT on system save commands is used.

***LOCKWAIT:** SAVACT operations for this library wait for a commit boundary before continuing with the save. The maximum wait time is the same as the value specified for the object lock wait time on the save command, typically 120 seconds.

***NOCMTBDY:** SAVACT operations for this library do not wait for a commit boundary. This option can be safely used for replication because all entries in the journal are replicated to the target, including any rollback operation that might follow a commit. Specifying ***NOCMTBDY** may improve performance of both replication jobs and user jobs that run concurrently with replication.

***NOMAX:** SAVACT operations for this library wait forever for a commit boundary before continuing with the save.

To ASP Device

Use this value when the target library is in a different iASP than the source library.

The possible values are:

***SAVASPDEV:** The target library is not in an iASP, or the name of the iASP containing the target library is the same as the iASP of the source library.

Name: Enter the name of the iASP that contains the target library.

To ASP

Use this value when the target library is in a different user ASP than the source library.

The possible values are:

***SAVASP:** The target library is in the same ASP as the source library.

1-16: Enter 1 for the system ASP, or 2-16 for a user ASP.

NOTE: If a name is specified for To ASP Device, a value of *SAVASP is assumed.

Replication Group

Allows you to associate this entry with a replication group. When associated with a group, sync attributes at the group level override similar attributes specified at the library level. You can also use the Change Group Sync Attributes to associate library sync entries with groups.

***NONE:** The entry is not associated with a group.

Name: Enter an existing replication group name.

Press **F4= Prompt** to select from a list of possible values.

Ignore Remote Error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

The possible values are:

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

YES:** Only specify **YES for this parameter if advised to do so by Help Systems Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Last SPLF receiver processed.

The last SPLF journal receiver processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

NONE:** No journal entries have been processed. If **YES is specified for "Sync Spooled Files" parameter, all spooled files on all output queues in this library will be refreshed at the next sync point.

Name: Enter the name of the last journal receiver processed. This, along with the "Last SPLF Sequence Number Processed" parameter, determines where the next synchronization point begins.

Last SPLF sequence number processed.

The last spooled file sequence number processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

Number: Enter the number of the last journal sequence number processed. This, along with the "Last SPLF Receiver Processed" parameter, determines where the next synchronization point begins.

Synchronization date

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

***NONE:** The library synchronization date is reset. The entire library will be refreshed at the next synchronization point.

Date: Enter a date in your job date format. The library is considered synchronized before the date specified. Only objects changed on or after the date specified will be sent to the target machine.

Synchronization time

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

Time: Enter a time in HHMMSS format. The library is considered synchronized before the time specified. Only objects changed at or after the time specified will be sent to the target machine.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Change IFS Sync Attributes

The Change IFS Sync Attributes panel is used to change the synchronization attributes for a given IFS directory.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, toggle (**F11**) until you get to the Work with Synchronization Attributes-IFS panel.
- 3. Enter option **2** next to an IFS sync attribute and press Enter.

RHA1033	Change I	FS Sync Attri	butes	10:57:09
(1 of 4) Directory	/help s	ystems/hsload	mgr/easy view/output	
Server ID To directory	<u>MYBACKU</u> /help s	ystems/hsload	Name, F4=Prompt mgr/easy view/output	
Default synchronization: Interval Units	· · · ·	10 *MINUTES	1-32767, *NONE *HOURS, *MINUTES, *SEC	CONDS
Journaling Information				
Robot HA manages the journ	al	*YES	*YES, *NO	
Local journal Library		IHELPSYSTE JRNLIB	Name, *DIR, *NONE, F4= Name	=Prompt
Remote journal	· · · · · ·	IHELPSYSTE RMTJRNLIB	Name, *NONE, *JRN, F4= Name	-Prompt
Last receiver processed . Last sequence number proce	ssed .	IHELPS0009 253	Name, *NONE Number	
F3=Exit F4=Prompt F5=Ref	resh F1	2=Cancel F21	=System Command F22=Fu	More ull path

Options

Directory

Enter the name of the directory on the local machine that is to be synchronized. The attributes set with this command apply only to a synchronization operation for the given Directory, To Directory, Server ID combination.

The possible values are:

Directory-name: Enter a valid directory name of up to 256 characters.

To type a path longer than 50 characters, press **F22** to display a full path and type the directory into the window provided.

Press **F4= Prompt** to select from a list of possible directories.

NOTE: If your 'Directory' is /QDLS, your 'To Directory' must also be /QDLS. Redirection cannot be used.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your machine.

Press **F4= Prompt** to select from a list of possible values.

To Directory

Enter the name of the target directory for synchronization.

The attributes set with this panel apply only to a synchronization operation for the given Directory, To Directory, Server ID combination.

The possible values are:

Directory-name: Enter a valid directory path on the target system. The name may be up to 256 characters.

To type a path longer than 50 characters, press F22=Full Path and type the directory into the window provided.

***FROMDIR:** The path of the target directory is the same as the source directory.

NOTE:

- If the **To Directory** path doesn't match the **From Directory**, you are using redirection. If you are using redirection, we will not restore authorities for the parent directories. You will have to apply authorities manually on your target system.
- If the parent directories of the attribute you are defining do not exist on the target side when we start syncing, we will create the parent directories. If you are using redirection, we will not restore authorities for the parent directories. If you are not using redirection we will restore the authorities.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize

Directories (SYNCIFSRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the directory is currently running. When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

***HOURS:** The repeat interval is specified in hours.

*MINUTES: The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the directory being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level. Only the records that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file results in the entire object being sent to the target.

The possible values are:

***NONE:** Detail changes are not tracked. Changes to files result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per IFS directory and name the journal after the directory to which it corresponds. Specifying the Robot HA product library is not recommended as journal objects stored in this library will be lost the next time Robot HA is upgraded to a new release.

***DIR:** The name of the local journal is the upper case interpreted value of the directory at top of page. The name is "I" plus the upper case alphabetic characters up to a maximum of 10 characters.

The possible library values are:

JRNLIB: The journal is stored in library JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target system

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

The disadvantage of remote journaling support is that all journal entries are sent, even those that are not needed for replication.

If a value other than *NONE is specified for this value and Local Journal, Robot HA creates the remote journal library and all other needed objects automatically. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each sync interval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the directory the journal corresponds to and store it in a common library such as RMTJRNLIB.

NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library. Any other value for Remote Journal will result in an error.

The possible library values are:

RMTJRNLIB: RMTJRNLIB is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.

Last receiver processed.

The last journal receiver processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

***NONE:** No journal entries have been processed. The next synchronization point begins with the first available journal entry.

Name: Enter the name of the last journal receiver processed. This, along with the "Last Sequence Number Processed" determines where the next synchronization point begins.

Last Sequence Number Processed

The last sequence number processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

Number: Enter the number of the last journal sequence number processed. This, along with the "Last Receiver Processed" determines where the next synchronization point begins.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *****NONE is specified for the Journal parameter or *****NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this directory must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

*INTERVAL: Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *****NONE is specified for the Journal parameter or *****NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this directoy must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Refresh on Journal Apply Error

Specify whether to refresh objects when an error is detected applying journal entries for the object.

Though journal apply errors should not occur often, it is possible for them to occur, particularly with multi-way synchronization where a record is updated simultaneously on more than one machine.

When refreshed, the entire object is resent to the target machine, replacing the existing object. Refreshing is the easiest way to automatically re-synchronize objects that get out of sync due to failed journal entries.

The possible values are:

***DATA:** Robot HA will refresh objects only when an error is encountered applying a journal entry for the object that would change the contents of the object. Examples of entries that change the object contents are record add, change and delete; member clear, etc. Examples of entries that do not change the object contents are authority changes and changes to the object attributes.

***ALL:** Robot HA will refresh an object when any error is encountered applying changes to the target object.

***NONE:** Robot HA does not refresh objects with journal entry apply errors. A message is sent to the system operator's message queue on the target machine with information about the entry that could not be applied. It is then the operator's responsibility to re-synchronize the object manually.

Ignore move/rename errors

Indicate how to handle errors that occur when renaming or moving objects on the target machine.

For renames: If a file is renamed on the source machine immediately after it's created, the original file name may be unavailable to the journal apply job on the target. In that case, the file is created with the new name on the target and a subsequent attempt to rename the file from the original to the new name will fail. As the target file already has the new name, this error can be safely ignored in most cases. A value of *YES for this parameter will ignore this error.

For moves: Since the target directory for the move may be outside the directory being replicated, an error moving an object may or may not be important. The classic case is when an object on the source machine is moved to some directory like '/trash' which doesn't exist on the target machine. If the target directory should exist on the target machine, then this should be treated as a hard error. But if the target directory is purposefully omitted from the target machine, the failed move should be handled differently. A value of *YES for this parameter will ignore this error.

The possible values are:

***NO:** All move and rename errors are treated as hard errors.

*YES: Most move and rename errors are ignored.

If an object cannot be moved on the target machine:

- Messages are placed in the job log on the target machine noting the condition.
- The object that could not be moved is deleted from the target machine.
- No error is reported to the source machine.

If an object cannot be renamed on the target machine, and if the to-file exists and the fromfile does not:

- Messages are placed in the job log on the target machine noting the condition.
- No error is reported to the source machine.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

• The local journal lives on the production machine where it constantly captures the changes to be replicated.

- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES:** Robot HA creates a standby journal on the target machine.

***NO:** Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this directory. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected .

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this directory on the Synchronize Directories (SYNCIFSRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the directory name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this directory in batch.

Remote Journal Validity Check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate the remote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Ignore Remote Error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

***YES:** Only specify *YES for this parameter if advised to do so by Fortra Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Objects to Omit

Specify objects to omit. These objects will not be synchronized. Up to 20 omit specifications may be entered.

The possible object values are:

***NONE:** No objects are omitted.

Path-name: Enter a path name for an object or objects to omit.

You can include the special characters ? and * in the last element of the path name. The ? will match any single character at that position in the name. The * will match any string of zero or more characters. Multiple instances of ? and * are allowed.

To eliminate all files in a subdirectory, end the specification with the subdirectory name as in this example:

/my_directory/notes

which will omit all files in subdirectory

'/my_directory/notes'.

NOTE: If you are syncing /QDLS with no subfolder defined, you can specify folders within QDLS to omit. If your 'From Directory' includes a subfolder in QDLS ('/QDLS/XXXX), you cannot specify an omit value.

Server ID for src from target

Enter the server ID to use on the target machine to connect back to the source machine. This value is ignored if *NONE is specified for Remote Journal.

The possible values are:

***CURRENT:** This is the recommended value. If the target machine needs to send a request back to the source machine, it uses the same connection the source machine used to initially connect to the target machine. *CURRENT may not work as expected with certain TCP/IP proxy configurations. For these cases, you can designate a server ID defined on the target machine to be used to contact this machine back when *CURRENT is specified. See the Associated Server ID (SERVER) parameter on the Add Requester Directory Entry (ADDRSFRDE) command for more information.

Server-ID: Enter a server ID that exists on the target machine and can be used to contact the source machine.

Max items to refresh

Enter the maximum number of objects Robot HA should refresh individually before refreshing the whole directory. When not using remote journaling, the value specified here is ignored and a value of 8 is used.

The possible values are:

***NOMAX** All objects with errors are refreshed individually without refreshing the whole directory.

Number: Specify the maximum number of objects to refresh individually.

QDLS ASP number

When replicating /qdls, you can specify an auxiliary storage pool number to limit the folders that are included. This value is ignored for all directories except /Qdls.

The possible values are:

***ANY:** All folders are included.

1-32: Enter an ASP number. The ASP must already be defined on your system. Only folders in the specified ASP are included.

Synchronization date

The last date and time the directory was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

***NONE:** The directory synchronization date is reset. The entire directory will be refreshed at the next synchronization point.

Date: Enter a date in your job date format. The directory is considered synchronized before the date specified. Only objects changed on or after the date specified will be sent to the target machine.

Synchronization time

The last date and time the directory was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

Time: Enter a time in HHMMSS format. The directory is considered synchronized before the time specified. Only objects changed at or after the time specified will be sent to the target machine.

Function Keys

- F3 (Exit): Exit the current panel without processing any pending changes.
- F4 (Prompt): Displays a list of possible values from which you may select one.
- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F12 (Cancel): Exit the current panel without processing any pending changes.
- F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Change Group Sync Attributes

The Change Group Sync Attributes panel is used to change the synchronization attributes for a given group.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, toggle (**F11**) until you get to the Work with Synchronization Attributes-Group panel.
- 3. Enter option **2** next to a group sync attribute and press Enter

RHA1032 Change Gr	oup Sync Attr	ibutes 13:16:26
Group ORDERGR Server ID <u>MYBACKU</u>	P	Name Name, F4=Prompt
Default synchronization: Interval	10 *MINUTES	1-32767, *NONE *HOURS, *MINUTES, *SECONDS
Robot HA manages the journal Library management options	<u>*YES</u> <u>*ENDJOB</u>	*YES, *NO *ENDJOB, *ENDRMTJRN, *DLTJRN
Local journal	ORDERGRP JRNLIB ORDERGRP RMTJRNLIB ORDERG0010 279	Name, *GROUP, *NONE, F4=Prompt Name Name, *NONE, *JRN, F4=Prompt Name Name, *NONE Number
F3=Exit F4=Prompt F5=Refresh F1	2=Cancel F21	=System Command

Options

Group

Enter the name of the replication group.

The possible values are:

Group-name: Enter a valid group name.

Group Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your machine.

Press **F4= Prompt** to select from a list of possible values.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize Groups (SYNCGRPRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the group is currently running. When a non-zero repeat interval is used, the synchronization will repeat at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

None: The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

*HOURS: The repeat interval is specified in hours.

***MINUTES:** The repeat interval is specified in minutes.

*SECONDS: The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with the group being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA changes journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA manages the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA

Library Management Options

Indicates how to manage items associated with library sync attribute entries that are added to the group.

The possible values are:

***ENDJOB:** Any replication jobs started for library entries before they were added to the group are ended.

ENDRMTJRN:** In addition to the actions associated with the **ENDJOB option, remote journaling defined in library sync entries is stopped if it is not also being used by other sync attribute entries.

DLTJRN:** In addition to the actions associated with the **ENDJOB and *****ENDRMTJRN options, local, remote, and standby journals associated with the library sync entries are deleted if they are not also being used by other sync attribute entries.

***NONE:** No action is taken to manage items associated with new library sync entries added to the group.

Local Journal

Enter the qualified name of a journal to use on the source machine to track detailed changes. If the journal does not exist, and if *YES is specified for **Robot HA Manages the Journal**, the journal and journal receiver are created for you. When a value other than *NONE is specified, files are synchronized at the record level and data areas are synchronized at the byte level. Only the records or bytes that have changed are sent to the target system. On the other hand, when *NONE is specified, any change to a file or data area results in the entire object being sent to the target.

The possible values are:

***GROUP**: The journal has the same name as the group. This is the recommended value.

***NONE:** Detail changes are not tracked. Changes to files and data areas result in whole objects or members being sent.

Name: Enter the name of the journal to use to track detailed changes. The recommended standard is to use one journal per library and name the journal after the library to which it corresponds. Specifying the Robot HA product library is not recommended as journal objects stored in this library will be lost the next time Robot HA is upgraded to a new release.

Press **F4= Prompt** to select from a list of possible values.

The possible values for the library are:

JRNLIB: The journal is stored in a separate library called JRNLIB. This is the recommended value.

Library-name: Enter the name of the library containing the journal.

NOTE: Specifying library RBTHALIB is not recommended. Journal objects stored in library RBTHALIB will be lost the next time Robot HA is upgraded to a new release.

Remote Journal

Indicates whether remote journaling support should be used to transmit journal changes from the source to the target machine.

If a value other than *NONE is specified for this parameter and the Journal parameter, Robot HA creates the remote journal library and all other needed objects automatically.

This parameter is ignored if *NONE is specified for the Journal parameter.

Advantages of remote journaling support:

- 1. The journal changes are sent more quickly and efficiently.
- 2. The journal changes are sent continuously and applied continuously. At synchronization intervals, Robot HA checks the progress of the journal apply process but does not need to send the journal entries.

Disadvantages of remote journaling support:

- 1. All journal entries are sent, even those that are not needed for replication.
- 2. Filtering is not supported. *NO is assumed for Filter Journal Entries.

The possible values are:

***NONE:** Remote journaling is not used. The journal changes are sent by Robot HA at each syncinterval.

***JRN:** Remote journaling is used. The journal changes are sent continuously. The remote journal has the same name as the local journal specified on the JRN parameter.

Name: Enter the name of the remote journal to use. If the journal does not exist on the remote machine, it will be created. A good practice is to name the remote journal after the library the journal corresponds to and store it in a common library such as RMTJRNLIB.

Press **F4= Prompt** to select from a list of possible values.

The possible library values are:

RMTJRNLIB: This is the recommended name to use for the remote journal library (except when the source and target systems are the same.)

Name: Enter the name of the remote journal library to use. If the library does not exist on the remote machine, it will be created.
NOTE: When using remote journaling, the library for the remote journal should be different than the library for the local journal to facilitate role swaps. There is one exception to this rule: when the source and target systems/partitions are the same. In that case, the remote journal name and library must match the local journal name and library for remote journaling replication logic to be used. When the source and target libraries are on the same machine (as implied by the server ID specified), you must specify either *NONE for Remote Journal or the remote journal name and library must match the local journal name and library. Any other value for Remote Journal will result in an error.

Last receiver processed.

The last journal receiver processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

***NONE:** No journal entries have been processed. The next synchronization point begins with the first available journal entry.

Name: Enter the name of the last journal receiver processed. This, along with the "Last Sequence Number Processed" determines where the next synchronization point begins.

Last Sequence Number Processed

The last sequence number processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

Number: Enter the number of the last journal sequence number processed. This, along with the "Last Receiver Processed" determines where the next synchronization point begins.

Library Information

Specify the library sync attributes entries to include in the group. A library entry is uniquely identified by from-lib, to-lib, server-ID and set-name.

From Library

Specify the library sync attributes entries to include in the group. A library entry is uniquely identified by from-lib, to-lib, server-ID and set-name.

The possible values are:

Library-name: Enter a valid library name.

To Library

Enter the name of the target library.

The possible values are:

Library-name: Enter a valid library name on the target system.

Library Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter a valid server ID. While a different server ID may be needed to identify the desired library sync entry, the group server ID overrides the library server ID at run time.

Set Name

Specify a set name to help identify the entry. A combination of from-library, to-library, server-ID, and set-name uniquely identifies an entry.

The possible values are:

***DFT:** The default set name is used.

Set-name: Enter a valid set name to help identify the entry.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this group must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

*INTERVAL: Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this group must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

End Journaling for Excluded Objects

Indicate whether Robot HA should end journaling for objects that are journaled to the synchronization journal but are listed as objects to omit from synchronization.

A value of *NO is assumed for this parameter if the journal is used by more than one sync attributes entry. In that case, you must end journaling for excluded objects manually after confirming that the objects are excluded from all sync attributes entries using this journal.

The possible values are:

***YES:** Robot HA will end journaling for excluded objects if they are journaled to the synchronization journal.

***NO:** Robot HA does not end journaling for excluded objects. This could result in some of these objects being mirrored, despite their being listed as excluded.

Fix Objects Using Wrong Journal

Indicate whether objects journaled to a journal other than the specified synchronization journal should be changed to the synchronization journal.

NOTE: Journaling for objects cannot be changed if they are in use.

The possible values are:

***YES:** Robot HA will attempt to change journaling for objects that are not journaled to the correct journal. If an object is in use, journaling cannot be changed and Robot HA will try again at the next synchronization interval.

***NO:** Robot HA will not change objects that are journaled to the wrong journal. These objects will not be synchronized.

Apply Journal Entries By

Specify whether Robot HA should apply journal entries to target objects by unique key or relative record number.

The possible values are:

***KEY:** Robot HA looks for a suitable unique key access path over each target file. If one exists, record updates and deletes are done by key. Otherwise, updates are done by relative record number. Updating by key requires a small amount of extra overhead, but may be more reliable overall as relative record numbers could get jumbled between the source and target files. What's more, updating by key is a necessity if two-way or N-way synchronization is to be used. Robot HA's keyed update option allows you to use N-way synchronization if you wish. If you do not like the access path that Robot HA chooses automatically for updating a particular file, you can use the Change Unique Key Association (**CHGRSFUKEY**) command on the target machine to tell Robot HA which access path to use.

NOTE: Selecting this option requires that before and after images be journaled. If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA handles this automatically.

***RRN:** All record changes are done by relative record number.

Journal Images

Indicate the type of journal images to capture.

NOTE: If *YES is specified for the **Fix Objects Using Wrong Journal** parameter, Robot HA sets the journaling for each object automatically.

The possible values are:

***DFT:** Robot HA captures only those journal entries needed for replication. If ***KEY** is specified the **Apply Journal Entries By** parameter, both before and after images are needed. Otherwise, only after images are needed.

***BOTH:** Both before and after images are captured, regardless of the value specified for the APPLY parameter.

Disable Triggers on Target

Specify whether to disable PF trigger programs on the target machine when entire *FILE objects are sent from the source to the target machine.

When replicating systems, it is often advisable to disable PF triggers on the target machine because the changes caused by triggers on the source machine are already being replicated to the target.

Any triggers that are disabled are automatically re-enabled in the event of a role swap.

The possible values are:

***YES:** Robot HA will disable PF triggers on the target whenever an entire file is sent to the target.

***NO:** PF triggers are not disabled on the target.

Create Standby Journal

Indicate whether Robot HA should create a standby journal on the target machine. The standby journal is not used during normal replication, but having a standby journal in place can improve the efficiency of a role swap.

- The local journal lives on the production machine where it constantly captures the changes to be replicated.
- The remote journal lives on the backup (target) machine. Any changes that show up in the local journal are automatically mirrored to the remote journal. Robot HA uses the information in the remote journal to keep the target machine objects in sync.
- The standby journal also lives on the target machine. Most of the time, the standby journal is not used and the journal changes it collects are thrown away. However, if a role swap results in the target machine moving into production, the standby journal is ready to assume the role of the local journal.
- Because the standby journal's main job is to take over as the local journal in the event of a role swap, the name and library of the standby journal on the target machine will match the name and library of the local journal on the production machine.
- What saves time during a role swap is not just the existence of the standby journal, but the fact that journaling has already been started for all of the objects that will use the journal for replication

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter.

The possible values are:

***YES:** Robot HA creates a standby journal on the target machine.

***NO:** Robot HA does not create a target standby journal. In the event of a role swap, the needed journal will be created at that time.

Clear Standby Journal Every

Indicate how often Robot HA should clear the standby journal. The entries collected in the standby journal are not needed until a role swap. Clearing the journal saves space on the target machine.

This parameter is ignored if *NONE is specified for the Journal (JRN) parameter or *NO is specified for the Create Standby Journal (CRTSTBY) parameter.

IMPORTANT: If you do not have Robot HA clear the standby journal, you are responsible for managing it. If it is not managed, it will continue to grow and could cause disk space issues.

The possible values are:

***DAY:** Robot HA clears the standby journal once a day.

NOTE: The daily clear is initiated by the replication job for this group. If the replication job does not run in a given day, the standby journal is not cleared that day.

***INTERVAL:** Robot HA clears the standby journal at each sync interval.

***NEVER:** Robot HA does not clear the standby journal. The journal will continue to grow until it is cleared manually.

Apply Error Log Level

Specify the type of information to log on the target machine in the event of a replication error.

In a perfect world, there would be no replication errors. When they do occur, it is often due to an unexpected job that is running on the target machine and interfering with replication. Capturing log information at the time of the error can help diagnose and correct any problems.

The possible values are:

0: No log information is captured.

1: Lock information is captured. Locks are shown for the object for which the replication error was detected and for the library containing the object. For physical files, member and record locks are also shown.

2: In addition to all level 1 information, a list of non-replication changes is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were not part of the replication process are listed.

3: In addition to all level 2 information, a job log is printed. If the target apply job has been running for a long time, the job log could be large.

4: In addition to all level 3 information, a list of changes from triggers and constraints is generated. If the object for which the replication error was detected is journaled on the target machine (ex. if you specified *YES for "Create Standby Journal" (SRTSTBY)), any changes to the object which were generated by triggers or constraints are listed.

Apply Error Retry

Specify the number of additional times to try an operation on the target machine when the first attempt fails.

Locks and other activity by non-replication jobs on the target machine can cause a replication step to fail. In such a case, retrying failed operations may reduce the need for objects to be refreshed.

The possible values are:

***NONE:** No additional attempt is made to retry a failed operation.

Number: Enter the number of retry attempts.

Apply Retry Delay

Indicate the number of seconds to wait between successive retry attempts. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

Seconds: Enter the number of seconds to wait between retry attempts.

Apply Retry Type

Indicate the types of operations to retry on the target machine. This parameter is ignored if *NONE is specified for "Error Retry" (RETRY).

The possible values are:

***DBOPEN:** Retry database file opens. Robot HA must open files on the target in order to replicate record-level changes.

ALLDB:** Retry the operations covered by **DBOPEN, plus all other database operations, including record-level, member-level and file-level changes.

***ALL:** Retry the operations covered by *ALLDB, plus all data area and data queue operations.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this group on the Synchronize Directories (SYNCGRPRSF) command. Defaults to S_nnnnnnn where nnnnnnn is the first eight characters of the grou name.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this group in batch.

Remote journal validity check

Indicates whether the system should use validity checking to verify that remote journal information received by the target is identical to what was sent by the source. For performance reasons, remote journaling does not do integrity checking by default. For a typically reliable network connection, remote journaling will faithfully transfer data from the source to the target. When using a noisy or unreliable connection, however, the data could be corrupted in transmission. This value is ignored unless journaling and remote journaling are used.

The possible values are:

***NO:** Validity checking is not enabled. This is the recommended value.

***YES:** Validity checking is enabled for the remote journal. The system ensures that what was received matches what was sent. If an error is detected, the data is not written to the remote journal and the remote journal connection is ended. Robot HA will attempt to restart the remote journal automatically at the next sync interval.

NOTE: Changing this value will temporarily deactivate the remote journal. Robot HA will reactivate the remote journal automatically at the next sync interval.

Remote journal image filter

Indicates whether the system should refrain from sending certain entries from the local to the remote journal. If you specify a value other than *NONE, IBM licensed program "HA Journal Performance" (5770SS1 feature 42) is required. This value is ignored unless journaling and remote journaling are used. *NONE must be specified for this parameter if *KEY is specified for the Apply Journal Entries By parameter.

The possible values are:

***NONE:** No filtering is done. All entries are sent to the remote journal.

*BEFORE: Record change before images are not sent to the remote journal.

NOTE: If you change this value, Robot HA must temporarily delete the remote journal from the target machine. Robot HA will recreate the remote journal automatically at the next sync interval.

SAVACT record change wait

Indicates whether to wait for a commit boundary when saving objects within this library for replication. For more information, see the SAVACTWAIT parameter on the various system SAVxxx commands.

The possible values are:

***DFT:** The default value specified for SAVACTWAIT on system save commands is used.

***LOCKWAIT:** SAVACT operations for this library wait for a commit boundary before continuing with the save. The maximum wait time is the same as the value specified for the object lock wait time on the save command, typically 120 seconds.

***NOCMTBDY:** SAVACT operations for this library do not wait for a commit boundary. This option can be safely used for replication because all entries in the journal are replicated to the target, including any rollback operation that might follow a commit. Specifying ***NOCMTBDY** may improve performance of both replication jobs and user jobs that run concurrently with replication.

***NOMAX:** SAVACT operations for this library wait forever for a commit boundary before continuing with the save.

Ignore remote error

Indicates whether to change the sync attributes on the local system, even if the remote system is unavailable and corresponding attributes on the remote system could not be changed.

The possible values are:

***NO:** If the remote system cannot be contacted and any of the attributes to be changed affect remote journaling, the request ends in error. No sync attributes are changed.

YES:** Only specify **YES for this parameter if advised to do so by Help Systems Technical Support. The sync attributes on the local system are changed, even if some of the changes affect remote journaling and the remote system is unavailable.

Synchronization date

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

***NONE:** The library synchronization date is reset. The entire library will be refreshed at the next synchronization point.

Date: Enter a date in your job date format. The library is considered synchronized before the date specified. Only objects changed on or after the date specified will be sent to the target machine.

Synchronization time

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

Time: Enter a time in HHMMSS format. The library is considered synchronized before the time specified. Only objects changed at or after the time specified will be sent to the target machine.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Change System Sync Attributes

The Change System Sync Attributes panel is used to change the synchronization attributes for system information such as user profiles, system values, and so on.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. From the Work with Synchronization Attributes panel, press **F11** to toggle until you get to Work with Sync Attributes-System.
- 3. Enter **2** next to a System Sync Attribute and press Enter.

ibutes 13:17:10
Name, F4=Prompt
*HUIL, *CFG, *NETH, *SYSVHL
1-22767 *NONE
*HOURS, *MINUTES, *SECONDS
*YES, *NO
Date, *NONE
Name, *NONE
Number
System Command

Options

Available options may differ depending on the type of system synchronization attribute being added.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current system are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your system.

Press **F4= Prompt** to select from a list of possible values.

Type of Information

Indicate the type of information to be synchronized.

The attributes set with this command apply only to a synchronization operation for the given type and server ID combination.

The possible values are:

*USRPRF: User profile synchronization attributes are set.

NOTE: User profiles with TOTP keys for multifactor authentication, introduced in IBM i 7.6, are synchronized with Robot HA if both systems are at a configuration that supports TOTP keys.

***SYSVAL:** System value synchronization attributes are set.

***NETA:** Network attribute synchronization attributes are set.

***PVTAUT:** Private authority synchronization attributes are set.

***AUTL:** Authorization list synchronization attributes are set.

***CFG:** Configuration synchronization attributes are set.

NOTE:

- Configuration information is handled differently than other types of system synchronization. Configuration information is extracted to the library specified by the **Config Extract Library** parameter and replicated to a library by the same name on the target machine. However, while it is stored on the target machine, the configuration is not applied to the target machine until a role swap is performed or the Apply Configuration Info is run.
- The *PVTAUT sync job works with the library and group sync jobs to ensure object authority is synchronized. It gathers object authority changes from QAUDJRN and then the library and group sync process will apply the changes to the backup.

Default Synchronization

Specify the default synchronization interval. This will be used as the synchronization interval if *ATTR is specified for the "Repeat Every" (REPEAT) parameter on the Synchronize System Info (SYNCSYSRSF) command. Changes to this value take effect at the next sync interval. Changes to this value are ignored if synchronization for the attribute is currently running. When a non-zero repeat interval is used, the synchronization repeats at regular intervals until the job is ended. CPU cycles are not used while waiting for the next interval.

The possible Interval values are:

***NONE:** The synchronization is done once and not repeated.

Number: Enter a number for the repeat interval.

The possible Units values are:

*HOURS: The repeat interval is specified in hours.

*MINUTES: The repeat interval is specified in minutes.

***SECONDS:** The repeat interval is specified in seconds.

Robot HA Manages the Journal

Indicate whether Robot HA should manage the journal and journal receivers associated with this system information being synchronized. When Robot HA manages the journal, journals and journal receivers that do not exist are created. Also, Robot HA will change journal receivers and delete receivers when they are no longer needed for synchronization. This value is ignored if *NONE is specified for Local Journal.

The possible values are:

***YES:** Robot HA will manage the journal and journal receivers.

***NO:** The journal must be managed outside Robot HA.

Synchronization Date.

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

***NONE:** The library synchronization date is reset. The entire library will be refreshed at the next synchronization point.

Date: Enter a date in your job date format. The library is considered synchronized before the date specified. Only objects changed on or after the date specified will be sent to the target machine.

Synchronization Time.

The last date and time the library was synchronized are shown. You can change these values to control which changed objects are sent to the target machine.

The possible values are:

Time: Enter a time in HHMMSS format. The library is considered synchronized before the time specified. Only objects changed at or after the time specified will be sent to the target machine.

Last receiver processed.

The last journal receiver processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

***NONE:** No journal entries have been processed. The next synchronization point begins with the first available journal entry.

Name: Enter the name of the last journal receiver processed. This, along with the "Last Sequence Number Processed" determines where the next synchronization point begins.

Last Sequence Number Processed

The last sequence number processed is shown. You can change this value if needed to reposition the start of the next synchronization point.

The possible values are:

Number: Enter the number of the last journal sequence number processed. This, along with the "Last Receiver Processed" determines where the next synchronization point begins.

Change Receiver Every

Indicate how often Robot HA should change the journal receiver for the journal associated with this entry.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

Replication for this entry must be running for the receiver to be changed by Robot HA. The journal is checked at the end of each sync interval.

The possible values are:

***DAY:** Robot HA changes the journal receiver once a day.

***INTERVAL:** Robot HA changes the journal receiver at each sync interval.

***SYS:** Robot HA does not change the journal receiver. The system will change the journal receiver when it reaches its threshold size. You must ensure that the **Manage receivers** parameter for the journal is set to ***SYSTEM**.

Days to Keep Receivers

Indicate how many days to keep old journal receivers on the system before they are deleted by Robot HA.

This parameter is ignored if *NONE is specified for the Journal parameter or *NO is specified for **Robot HA manages the journal**.

NOTE:

- 1. Replication for this entry must be running for any receivers to be deleted. Receivers are checked at the end of each sync interval.
- 2. When remote journaling is used, remote journal receivers are deleted automatically by Robot HA, shortly after the corresponding receiver on the source machine is deleted. The request to delete the remote receiver is logged as a new journal entry which must be processed in order, after any other entries that precede it are processed on the target system

The possible values are:

***SYNC:** Journal receivers are deleted when they are no longer needed for replication.

Number: Enter a number greater than zero. Journal receivers are deleted if they are no longer needed for replication and they only contain entries that are older than the number of days specified.

Restrict to Window

Limits when the sync job associated with this entry runs. If you specify start and end times such that start <= end, the job can start any time between the start and end times each day. If you specify values such that start > end, the job can start any time except between the end and start times specified. The values specified for Default Synchronization determine how frequently the job runs within the specified window.

To run a job once a day at a specific time, specify a start and end time for the window and a value of 24 *HOURS for Default Synchronization. A longer value will have the same effect.

The possible Start values are:

***NONE:** The job can start at any time without restrictions.

Time: Enter the start time of the window.

The possible End values are:

Time: Enter the end time of the window.

Synchronization Job Name

The name specified here is used when JOB(*DFT) is specified for this library on the Synchronize System Information (SYNCSYSRSF) command. Defaults to S_nnnnnn where nnnnnn is the type of information.

NOTE: If you have a non-IBM tape management system installed, the default job name must begin with the two characters "S_" so that save operations initiated by Robot HA can be distinguished. You can use the command WRKREGINF EXITPNT(QIBM_QTA_TAPE_TMS) to determine if a tape management system is installed.

The possible values are:

Name: Specify the default job name to use when synchronizing this entry in batch.

Include or Omit Items

Indicate whether items listed on the **Items to Include/Omit** parameter should be included in or omitted from synchronization.

NOTE: Include and Omit parameters are not valid with *PVTAUT, *CFG, and *NETA.

The possible values are:

***OMIT:** The items listed are omitted. All other items are included.

*INCLUDE: The items listed are included. All other items are omitted.

Items to Include/Omit

Specify items to include or omit from synchronization. Whether the specified items are included or omitted is determined by the **Include or Omit Items** parameter.

The possible object values are:

Generic-name: Enter a generic name, ending with an asterisk (*). Items whose names begin as specified are included or omitted.

Name: Enter the name of a specific item to omit.

IMPORTANT: The following is a list of system values that are always omitted from syncing along with the earliest version of Robot HA where the omit was applicable.

System Value Name	Earliest Robot HA Version
QABNORMSW	12.01
QACGLVL	13.09
QAUDCTL	13.09
QCONSOLE	12.01
QDAYOFWEEK	12.01
QIGC	12.01
QIPLSTS	12.01
QLOCALE	13.09
QMODEL	12.01
QPRCFEAT	12.01
QPWDLVL	13.09
QPWRRSTIPL	13.09
QRMTIPL	13.09
QSRLNBR	12.01
QSTRPRTWTR	12.01
QUTCOFFSET	12.01

IMPORTANT: Some system values have a special handling starting in Robot HA version 13.09.

QPWDLMTAJC, QPWDLMTCHR, QPWDLMTREP, QPWDMAXLEN, QPWDMINLEN, QPWDPOSDIF, QPWDRQDDGT

The syncing of these is determined based on the contents of QPWDRULES. If QPWDRULES is set to *PWDSYSVAL, we will sync these, otherwise they will be ignored.

QSSLCSL

The syncing of this is determined based on the contents of QSSLCSLCTL. If QSSLCSLCTL is set to *OPSYS, QSSLCSL is treated as though it were read-only, and is therefore not synced.

Max Journal Entries to Process

Enter the maximum number of outstanding QAUDJRN entries to process for this type of information before completely refreshing all items. If the processing of journal entries gets too far behind, it may be more efficient to use current values to completely refresh the information for all items of the specified type. The processing of journal entries may fall behind for various reasons, including:

- Replication for the system item has not run for a while.
- Replication for the item is repeatedly ending in error.
- Your settings for system values QAUDCTL, QAUDLVL and QAUDLVL2 are causing the system audit journal, QAUDJRN, to grow very rapidly.

The possible values are:

*NOMAX: All journal entries are processed, regardless of how many remain outstanding.

Number: Enter the maximum number of unprocessed journal entries allowed before Robot HA discards the entries and uses current values to completely refresh the items.

NOTE: Normal processing of journal entries will resume with the cycle following the one in which all items are refreshed.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F21 (System Command): Displays a command line.

Page Up/Page Down: Displays the previous or next parameters.

Change RSF Defaults (CHGRSFDFT)

The Change RSF Defaults panel is used to tailor Robot HA default values to reflect the preferences of your installation.

How to Get There

From the Robot HA Main Menu, select option **3**. The Setup Menu displays. Then, select option **3**, System defaults.

Change	RSF	Defaults	(CHGRSFDFT)
Type choices, press Enter.			
Max times access allowed Pre-processing program		<u>*Nomax</u> *None	_ 1-9999997, *NOMAX, *SAME Name, *SAME, *NONE
Library	•	<u>*NONE</u>	Name, *LIBL, *CURLIB Name, *SAME, *NONE Name, *LIBL, *CURLIB
Post-processing program Library			Name, *SAME, *NONE Name, *LIBL, *CURLIB
Normal message queue Library			Name, *SAME, *NUNE Name, *LIBL, *CURLIB Name. *SAME. *NONE
Library		*LIBL *NONE	Name, *LIBL, *CURLIB Name, *SAME, *NONE
Library		*Nomax *Nomax	Name, *LIBL, *CURLIB 1-9999997, *NOMAX, *SAME 1-9999997, *NOMAX, *SAME
F3=Exit F4=Prompt F5-Pofros	h	E10=0ddit	More
F13=How to use this display	,,,	F24=More	keys

Options

Max times access allowed

The maximum number of times each remote IBM i system can access a Robot HA package on your system. This value may be overridden for individual packages. If a requester attempts to access a package more than the number of times allowed, an error message is sent and the request ends abnormally.

The possible values are:

***NOMAX:** There is no system-wide limit to the number of times Robot HA packages can be accessed by each requester.

1-9999997: The system default for the number of times Robot HA packages can be accessed by each requester.

Pre-processing program

The qualified name of a user program to call before granting access to an Robot HA package. This value can be overridden for individual packages. If specified, this program is

called upon receiving a request to download data for an Robot HA package, before an acknowledgment is sent to the requester. Depending upon how the program returns, access to the package is either granted or denied.

The possible values are:

***NONE:** No pre-processing program is called.

program name: The name of a system default pre-processing program to call

The possible library values are:

*LIBL: The job library list is used to locate the program.

***CURLIB:** The current library is used to locate the program.

library name: The name of an existing library which contains the specified program.

Catalog pre-processing program

The qualified name of a user program to call to determine whether information about Robot HApackages should be sent to the requester for a *CATALOG request. This value can be overridden for individual packages. If specified, this program is called for each Robot HApackage on your system upon receiving a request for a catalog. Depending upon how the program returns, catalog information about the package is either sent to the requester or not.

The possible values are:

***NONE:** No catalog pre-processing program is called.

program name: The name of a system default catalog pre-processing program to call.

The possible library values are:

*LIBL: The job library list is used to locate the program.

***CURLIB:** The current library is used to locate the program.

library name: The name of an existing library which contains the specified program.

Post-processing program

The qualified name of a user program to call to process Robot HA packages sent to your system by remote requesters with the **SNDRSFPKG** command. This value can be overridden

for individual requesters. If specified, this program is called after all of the data associated with the Robot HA package has been received and only if the data is received successfully. This program can be used to restore objects to your system from the save file sent for the Robot HA package.

The possible values are:

***NONE:** No program is called when Robot HA data is received.

program name: The name of a system default program to call when an Robot HA package arrives on your system, sent by remote requesters with the **SNDRSFPKG** command.

The possible library values are:

*LIBL: The job library list is used to locate the program.

***CURLIB:** The current library is used to locate the program.

library name: The name of an existing library which contains the specified program.

Normal message queue

The qualified name of a message queue to receive normal transaction messages. This value can be overridden for individual Robot HA packages. If specified, information and completion messages generated during a successful Robot HA request are sent to the message queue.

The possible values are:

***NONE:** Information and completion messages generated from successful transactions are not sent.

Name: The name of a system default message queue to which information and completion messages generated from successful transactions are sent.

The possible library values are:

*LIBL: The job library list is used to locate the message queue.

***CURLIB:** The current library is used to locate the message queue.

library name: The name of an existing library that contains the specified message queue.

Error message queue

The qualified name of a message queue to receive abnormal transaction messages. This value can be overridden for individual Robot HA packages. If specified, error messages generated during an unsuccessful Robot HA request are sent to the message queue.

The possible values are:

***NONE:** Error messages generated from unsuccessful transactions are not sent. Error message ids will still appear in the Robot HA log.

Name: The name of a system default message queue to which error messages generated from unsuccessful transactions are sent.

The possible library values are:

*LIBL: The job library list is used to locate the message queue.

***CURLIB:** The current library is used to locate the message queue.

library name: The name of an existing library that contains the specified message queue.

Requester message queue

The qualified name of a message queue to receive ad hoc requester messages. This value can be overridden for individual Robot HA packages. If specified, ad hoc messages sent by requesters using the **MSG** parameter of one of the Robot HA commands that initiate a transaction are sent to the message queue.

The possible values are:

***NONE:** Ad hoc requester messages are not sent. The messages are still logged in the Robot HA log and can be viewed with the **WRKRSFLOG** command.

Name: The name of a system default message queue to which ad hoc requester messages are sent.

The possible library values are:

*LIBL: The job library list is used to locate the message queue.

***CURLIB:** The current library is used to locate the message queue.

library name: The name of an existing library that contains the specified message queue.

Max spooled file pages to receive

The maximum number of spooled file pages or cover letter records that requesters can send to your system. This value may be overridden for individual requesters. If a requester attempts to send more spooled file pages than allowed, an error message is sent and the request ends abnormally.

The possible values are:

***NOMAX:** There is no system-wide limit to the number of spooled file pages or cover letter records that requesters can send to your system.

1-9999997: The system default for the maximum number of spooled file pages and cover letter records a requester can send to your system, per request.

Max save file records to receive

The maximum number of save file records that requesters can send to your system. This value may be overridden for individual requesters. If a requester attempts to send more save file records than allowed, an error message is sent and the request ends abnormally.

The possible values are:

***NOMAX:** There is no system-wide limit to the number of save file records requesters can send to your system.

1-9999997: The system default for the maximum number of save file records a requester can send to your system, per request.

Encryption Key Data Area

Enter the qualified name of a data area containing the basic encryption key to use. The data area must be at least 128 bytes long.

This parameter is only used for TCP/IP connections where the client requests *BASIC encryption. The client specifies TCP/IP encryption options with the Add/Change Server directory Entry commands (ADDRSFSDE/CHGRSFSDE).

The first 128 bytes of the data area are used as the encryption key. Trailing blanks in the key are ignored. The key can be from 1 to 128 bytes long and may contain any valid hex or character data. Longer keys provided greater security.

The data area does not need to exist at the time this command is run but it must exist at the time of transmission.

The possible values are:

***NONE:** No data area is used. A default key is supplied by Robot HA.

data-area-name: The name of a data area to use.

The possible library values are:

*LIBL: The job library list is used to locate the data area.

library-name: The name of the library containing the data area.

SSL Status

Specify whether to enable or disable Robot HA's Secure Sockets Layer (SSL) support. SSL is used to encrypt TCP/IP connections.

When SSL support is enabled, Robot HA modules that handle TCP/IP connections are linked to SSL run-time routines. This linkage will only be successful if the SSL run-time routines exist on your system. The SSL run-time routines are included with the Internet Connection Secure Server licensed program from IBM (US version: 5769NC1; International version: 5769NCE.) Therefore, Robot HA's SSL support can only be enabled if the ICSS licensed program is installed.

NOTE: If you enable Robot HA's SSL support, and later delete the ICSS licensed program, be sure to then disable Robot HA's SSL support. Otherwise you will not be able to connect via TCP/IP.

The possible values are:

***DISABLED:** SSL encryption is not available.

***ENABLED:** SSL encryption can be used over TCP/IP connections.

Client Authentication

Specify whether to perform client as well as server authentication for TCP/IP connections encrypted using the Secure Sockets Layer (SSL) protocol.

This parameter is only used for TCP/IP connections where the client requests *SSL encryption. The client specifies TCP/IP encryption options with the Add/Change Server directory Entry commands (ADDRSFSDE/CHGRSFSDE).

NOTE: In order to use SSL encryption, The Internet Connection Secure Server licensed program must be installed on both the client and server machines. (US version: 5769NC1. International version: 5769NCE.) Contact your IBM sales representative for more information.

The possible values are:

***NO:** Only server authentication is performed during the SSL handshake.

***YES:** Client and server authentication are performed.

TCP/IP Send Buffer Size

Specify the send buffer size to use when sending or retrieving objects over TCP/IP connections. Increasing the buffer size requires more memory during a transmission, but it can improve performance. This value affects the requester job when sending data, and the server job when retrieving data.

The possible values are:

65536: The default value is used.

number: Enter a valid send buffer size, from 16384 to 8388608.

TCP/IP Receive Buffer Size

Specify the receive buffer size to use when sending or retrieving objects over TCP/IP connections. Increasing the buffer size requires more memory during a transmission, but it can improve performance. This value affects the server job when sending data, and the requester job when retrieving data.

The possible values are:

65536: The default value is used.

number: Enter a valid receive buffer size, from 16384 to 8388608.

TCP/IP Record Block Size

Specify the number of records to transmit at once when sending or retrieving objects over TCP/IP connections. Increasing the block size can improve performance. However, specifying too large a block size could degrade performance if it becomes difficult for TCP/IP to keep up. You must experiment to determine the best setting for your environment.

Setting this parameter on a given machine controls the block size to use when sending data from or retrieving data from that machine.

The possible values are:

1000: This is a reasonable starting point for block size. Experiment to determine the optimal value for your environment.

number: Enter a valid block size, from 1 to 5000.

TCP/IP Keep-Alive Interval

Specify the number of minutes between keep-alive messages. Keep-alive messages help ensure that a TCP/IP connection is not dropped by a network router before the transmission is complete, even though the connection may appear to be idle for a time. Keep-alive messages also help Robot HA detect broken or dropped connections.

It is normal for there to be idle periods during a transmission.

For example:

- Waiting for a save file to be processed after it has been sent.
- Waiting for a long-running request on the target machine to complete.
- Waiting for the next in a series of RSFLINK requests to be received.

This value only affects Robot HA jobs and does not affect other TCP/IP applications on your system.

NOTE: This parameter is ignored for machines running an earlier version of the operating system than V6R1. In such a case, use the Change TCP/IP Attributes (CHGTCPA) command on the source and target machines to set the keep-alive interval globally.

The possible values are:

10: A keep-alive signal is sent every ten minutes when the Robot HA transmission is otherwise idle for that long.

***DFT:** The default keep-alive interval set by the Change TCP/IP Attributes (CHGTCPA) command is used.

number: Enter the number of minutes between keep-alive messages.

Log Local Commands

Specify whether an entry should be placed in the Robot HA log each time RTVRSFPKG, SNDRSFPKG, SNDSPLFRSF or STRPASRSF commands are executed on your machine.

The possible values are:

***SAME:** The current value is not changed.

***NO:** Local commands are not logged. Only requests received from remote machines are logged.

***YES:** Local commands and requests received from remote machines are both logged.

Log Switched Connections

Specify whether an entry should be placed in the Robot HA log each time Robot HA dials a remote server machine for one of the Robot HA requester

commands.

The possible values are:

***SAME:** The current value is not changed.

***NO:** Switched connections are not logged.

***YES:** An entry is placed in the Robot HA log for each switched Robot HA connection initiated by your machine. The log entry shows the number called, the start date, start time and end time.

Server ID For This Machine

Enter the server ID that other machines would typically use to contact this machine. The server ID must still be defined in the server directories of other machines. Certain Robot HA functions will pass this ID to remote machines so that those machines can contact this machine back at a later time.

At install time, Robot HA initially sets this value to the system name.

The possible values are:

***SAME:** The current value is not changed.

name: Enter the ID that other systems should use to contact this one. This should match the value specified for server ID in the server directory of other machines, for the entry that points to this machine.

Outgoing Mail Server

Specify the address of your outgoing mail server.

The possible values are:

address: Enter a valid mail server address.

For example: smtpauth.earthlink.net.

Outgoing Mail Port

Specify the port to use to connect to the outgoing mail server.

The possible values are:

***SMTP:** The default SMTP port of 587 is used.

***SMTPSSL:** Port 465 is used and an encrypted connection is required.

NOTE: Most SMTP servers prefer to handle both encrypted and non-encrypted connections using port 587 (*SMTP).

***SMTPOLD:** The older SMTP port of 25 is used.

port: Enter a valid port number.

Outgoing Mail User ID

Specify the user ID to use when verification is required by the mail server.

The possible values are:

***NONE:** No user ID or password are sent. Depending on the mail server, a user ID and password may be required.

user-ID: Enter a character string of up to 128 characters. Case is significant.

Outgoing Mail Password

Specify the password to use when verification is required by the mail server.

The possible values are:

password: Enter a character string of up to 128 characters. Case is significant.

Outgoing Mail From Address

Specify the default email address from which mail from this system is sent. This is the email address to which replies to your outgoing email will be directed.

NOTE: You can use the Change RSF User Options (CHGRSFUO) command to override this value for individual users.

The possible values are:

email-address: Enter the email address to which replies should be directed.

For example: fred.jones@your_company.com

Outgoing Mail Encryption

Indicate whether to use SSL/TLS to encrypt the connection to your mail server.

NOTE:

- 1. Some mail servers, like smtp.gmail.com, require encryption; most mail servers do not.
- Because internet mail is relayed from one server to another as it makes its way from the source to the target, an encrypted connection to the initial mail server does not guarantee that the message will remain encrypted as it moves through the internet.
- 3. To use an encrypted connection, you must configure SSL on your IBM i for use with Robot HA mail.

The possible values are:

***NO:** Outgoing mail is not encrypted.

***YES:** Outgoing mail is encrypted.

Alert Email Address

Specify up to eight email addresses to which to send alerts and status updates. The information entered here may be used by Robot HA system monitors and other Robot HA functions.

You can enter more than eight email addresses by directly editing file RSPF056 in library RBTHALIB.

The possible values are:

***NONE:** No email addresses are specified.

email-address: Enter one or more valid email addresses.

Alert Text Message

Specify up to eight phone numbers to which to send text messages for alerts and status updates. The information entered here may be used by Robot HA system monitors and other Robot HA functions.

You can enter more than eight phone numbers by directly editing file RSPF057 in library RBTHALIB.

The possible numbervalues are:

*NONE: No text message recipients are specified.

phone-number: Enter a valid cell phone number, as you would dial it from another cell phone. Spaces, dashes and other non-numeric characters may be inserted wherever you like for clarity.

email-to-phone-number: Enter the the complete email address to use to send a text message to a specific phone via email.

NOTE: Note: This format should only be used if *OTHER is specified for the "Carrier" element.

The possible carrier values are:

***SPRINT:** The message is formatted to send to the cell phone of a Sprint user.

***VERIZON:** The message is formatted to send to the cell phone of a Verizon user.

***TMOBILE:** The message is formatted to send to the cell phone of a T-Mobile user.

***ATT:** The message is formatted to send to the cell phone of an AT&T user.

***BELLSOUTH:** The message is formatted to send to the cell phone of a Bell South user.

***OTHER:** All the formatting necessary to email a text message to the desired phone must be specified in the "Phone number" element.

See http://en.wikipedia.org/wiki/List_of_SMS_gateways forinformation about emailing a text message (SMS) to various carriers.

Sync Start Program

Enter the qualified name of a user-written program to call to start High Availability synchronization.

The possible program values are:

***NONE:** No default synchronization start program is specified.

name: Enter a valid program name.

The possible library values are:

*LIBL: The library list is used to find the program.

***CURLIB:** The program is found in the current library.

name: Enter the name of the library containing the program.

Sync Job Display Program

Enter the qualified name of a user-written program to call to work with active and submitted synchronization jobs. The specified program is called when F9 is pressed from the Work With Sync Attributes display. The program should accept no parameters.

A model program, WRKSYNCJOB, is provided in library RBTHALIB. The source can be found in RBTHALIB/RHASAMPLE. To modify the program, copy the source to another library (such as RSFUSER) and make your changes to the copy.

The possible program values are:

WRKSYNCJOB: This is the default program provided in RBTHALIB.

name: Enter a valid program name.

The possible library values are:

RBTHALIB: The default program is located in library RBTHALIB.

***LIBL:** The library list is used to find the program.

***CURLIB:** The program is found in the current library.

name: Enter the name of the library containing the program.

Role Swap to Backup

Enter the qualified name of a user-written program to call to perform a role swap from production to backup.

The possible program values are:

NONE: No default production-to-backup role swap program is specified.

name: Enter a valid program name.

The possible library values are:

*LIBL: The library list is used to find the program.

***CURLIB:** The program is found in the current library.

name: Enter the name of the library containing the program.

Role Swap to Production

Enter the qualified name of a user-written program to call to perform a role swap from backup to production.

The possible program values are:

***NONE:** No default backup-to-production role swap program is specified.

name: Enter a valid program name.

The possible library values are:

*LIBL: The library list is used to find the program.

***CURLIB:** The program is found in the current library.

name: Enter the name of the library containing the program.

Current Replication Role

Specify the replication role of this machine. This value is changed automatically during a role swap.

The possible values are:

***PROD:** This machine is currently acting in the production (source) role.

***BACKUP:** This machine is currently acting in the backup (target) role.

Software Release/Mod Level

The Robot HA software release level is shown. This value cannot be changed.

The possible values are:

***SAME:** The current value is not changed.

Authority Error Threshold

Enter the maximum number of times an authority error can happen on an object before we error your synchronization job. Authority errors may occur when the authority for an object is unable to be retrieved or granted due to locks on the object. Each synchronization cycle will attempt to set any new authority changes and reprocess any previous errors, so if locks are removed, the authorities will be applied at that time.

The possible values are:

1-9999997:Enter the number of times an authority error can happen.

***SAME:** The current value is not changed.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F10 (Additional parameters): Displays additional parameters for the command.

F12 (Cancel): Exit the current panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F24 (More keys): Shows additional function keys that can be used for this panel.

Page Up/Page Down: Displays the previous or next parameters.

Change RSF User Options (CHGRSFUO)

The Change RSF User Options panel is used to change the defaults that control the operation of the Work with RSF Catalog and other Robot HA displays. Different default values can be set for each user.

Change RSF L	lser Options	(CHGRSFUO)
Type choices, press Enter.		
Target save file	<u>*PKG</u> QGPL *NO *YES *NO *YES RSFHA RSFUSER	Name, *PKG Name, *CURLIB *YES, *NO *YES, *NO *YES, *NO Name, *USRPRF Name, *LIBL, *CURLIB
Earliest date to include	<u>*FIRST</u> *LAST	Date, *FIRST, *NEW Date, *LAST
Entry status to include t for more values	*ALL	*ALL, *NOTBLANK, *NOTINL
Text compare string	*ALL	
Option file	<u>RSFOPT</u> <u>*LIBL</u>	Name Name, *LIBL, *CURLIB More
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	F13=How to use this display

Options

User

Specify the user whose default values you want to change.

The possible values are:

***CURRENT:** Defaults for the user running the command are changed.

user-ID: Enter a valid user ID.

Target Save File

Enter the qualified name of a save file on your system into which data for a retrieved package should be placed. If the save file does not exist, it will be created.

The possible values are:
***PKG:** The name of the package being retrieved is used for the save file name.

save-file-name: Enter the name of a save file to contain retrieved package data. If the file does not exist, it will be created.

The possible library values are:

***CURLIB:** Your job's current library is used to locate the save file

library-name: Enter the name of the library which contains or will contain the save file.

Replace Existing Save File

Specify whether Robot HA should replace the contents of the save file if it already exists.

The possible values are:

***NO:** If the save file already exists, the request ends in error.

***YES:** The contents of the save file are replaced with the data being downloaded.

Automatic Installation

Specify whether Robot HA should attempt to install a package once it has been retrieved.

For automatic installation, Robot HA attempts to restore program RSFINST to library QTEMP from the save file that was sent with the package.

If program RSFINST can be restored, it is called with the following parameter list:

- A data structure describing the request. CHAR(*)
- Return Message data CHAR(128)
- Return Message type CHAR(7)
- Save lib CHAR(10)
- Save command CHAR(10)

It is the server's responsibility to write installation program RSFINST and include it in the package save file if automatic installation for this package is to be supported.

The possible values are:

***YES:** Automatic installation is attempted. If no automatic installation program is found in the save file, the request still ends normally.

***NO:** No attempt is made to automatically install the package.

Retrieve Package If Installed

Specify whether to retrieve packages that are already flagged as installed on your system.

The possible values are:

***NO:** A request to retrieve an installed package ends in error.

***YES:** Installed packages may be retrieved again.

Run In Batch

Specify whether requests to retrieve and install packages should run interactively or in batch.

The possible values are:

***NO:** Requests run interactively.

***YES:** Requests are submitted to batch.

Job Description

Enter the qualified name of a job description to use when submitting requests to batch.

The possible values are:

***USRPRF:** Use the job description from the user profile of the users whose defaults are being changed.

job-description-name: Enter the name of an existing job description to use.

The possible library values are:

*LIBL: Your job's library list is used to locate the job description.

***CURLIB:** Your job's current library is used to locate the job description.

library-name: Enter the name of the library containing the job description.

Earliest Date To Include

Enter the earliest date to include in the catalog list. Only catalog entries that were created or changed on or after the date specified are included.

The possible values are:

*FIRST: No minimum date is used to filter catalog entries.

***NEW:** Include only entries that were added to the catalog since the last time the catalog was retrieved.

date: Enter an early cutoff date to use.

Last Date To Include

Enter the latest date to include in the catalog list. Only catalog entries that were created or changed on or before the date specified are included.

The possible values are:

*LAST: No maximum date is used to filter catalog entries.

date: Enter a minimum cutoff date to use.

Entry Status To Include

Use this parameter to select which entries to include by status.

The possible single values are:

***ALL:** All entries are included regardless of status.

***NOTBLANK:** Entries with any status other than blank are included.

***NOTINL:** Entries for packages that have not been installed are included.

The other possible values are:

***BLANK:** Entries for packages that have not been retrieved or installed are included.

***RTV:** Entries for packages that have been retrieved but not installed are included.

***INL:** Entries for packages that have been installed are included.

***OPN:** Entries for packages that have been retrieved again after they were last installed are included.

Text Compare String

Enter characters to compare to catalog entry text to determine which entries should be included in the list. An entry is included in the list if the entry text contains the string specified.

The possible library values are:

***ALL:** All entries are included, regardless of their text.

characters: Enter a character string to use to filter catalog entries by text. The search is not case sensitive.

Option File

Enter the qualified name of a user option file to use to process user-defined options.

The possible values are:

file-name: Enter the name of an existing option file to use.

The possible library values are:

*LIBL: Your job's library list is used to locate the file.

***CURLIB:** Your job's current library is used to locate the file.

library-name: Enter the name of the library containing the file.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

- F4 (Prompt): Displays a list of possible values from which you may select one.
- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F12 (Cancel): Exit the current panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F24 (More keys): Shows additional function keys that can be used for this panel.

Page Up/Page Down: Displays the previous or next parameters.

Check Defined Items (CHKATTRSF)

The Check Defined Items panel is used to check the objects in some or all of the libraries and IFS directories defined on your <u>Work with Synchronization Attributes panel</u>. The libraries and directories are compared to their counterparts on the backup system to determine if they match.

How to Get There

Run the command RBTHALIB/CHKATTRSF.

Check Def:	ined Items (CHKATTRSF)
Type choices, press Enter.	
Server ID	Name
Libries to compare	<pre>*ALL Name, *All, *NONE, *generic*</pre>
Omit libraries	<u>*NONE</u> Name, *NONE, *generic*
+ for more values	
Compare IFS directories	<u>*NO</u> *YES, *NO
Compare system objects	<u>*NO</u> *YES, *NO
Object name	<u>*ALL</u> Name, generic*, *ALL
Object type	<u>*ALL </u>
+ for more values	
Compare file data	<u>*NO</u> *YES, *NO
Compare while active	<u>*AUTO</u> *YES, *NO, *AUTO
Compare active wait time	<u>120 Number</u>
Save active wait time	<u>120 </u>
Allow differences	<u>*NONE</u> *ALL, *NONE, *FID
+ for more values	
Output	<pre>* *, *PRINT, *OUTFILE, *NONE</pre>
	More
F3=Exit F4=Prompt F5=Refresh	F10=Additional parameters F12=Cancel
F13=How to use this display	F24=More keys
l	

Options

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current system are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source system. For an Robot HA pre-processing program running on a server system, *CURRENT refers to the requester system. *CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server system, or a batch job submitted from one of the above two job types.

NOTE: *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. It is allowed within other Robot HA server jobs that connect via TCP/IP.

server-ID: Enter the name of an entry in the server directory on your system.

Libraries to Compare

Enter the name or generic name for the libraries to be included.

The possible values are:

***ALL:** All libraries for the given server that are defined on the Work With Synchronization Attributes panel are included.

***NONE:** No libraries are included.

generic-name: Enter a generic name for the libraries to include. An asterisk (*) in the generic specification will match any string of zero or more characters in the name. An underscore (_) in the generic specification will match any single character in the name.

library-name: Enter a library name.

Omit Libraries

Enter the name or generic name for libraries to omit. Up to 20 specifications can be entered.

The possible values are:

***NONE:** No libraries are omitted.

generic-name: Enter a generic name for libraries to omit. An asterisk (*) in the generic specification will match any string of zero or more characters in the name. An underscore (_) in the generic specification will match any single character in the name.

library-name: Enter a library name.

Compare IFS Directories

Indicate whether to submit compare jobs for IFS entries in the list.

The possible values are:

***YES:** All IFS entries on the Work With Synchronization Attributes panel that are defined for the given server are included.

***NO:** No IFS compare jobs are submitted.

Compare System Objects

Indicate whether to submit compare jobs for entries in the list such as *USRPRF and *AUTL.

The possible values are:

***YES:** All system-type entries on the Work With Synchronization Attributes panel that are defined for the given server are included.

***NO:** No system-type compare jobs are submitted.

Object Name

Indicate which objects to check.

The possible values are:

***ALL:** All objects in the selected libraries are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object Type

Indicate the type of objects to check. The possible values are:

*ALL: All types of objects are checked.

object-type: List the types of objects you want to check. Compare Data (CMPDTA) Indicate whether Robot HA should compare the data (contents) of the files or only the file attributes. The possible values are:

*NO: Only file attributes such as these are compared :

- Current number of members
- Maximum number of members
- Record length
- Number of fields
- Total number of records
- Number of deleted records
- Reuse deleted records attribute
- Create date

This is faster than comparing file data.

***YES:** In addition to the file attributes, the file data is compared. A hashing algorithm is used to compare the file data. This option takes longer to run but it will detect small differences with a high degree of accuracy.

Compare File Data

Indicate whether Robt HA should compare the data (contents) of the files or only the file attributes.

The possible values are:

***NO:** Only the following file attributes are compared:

- Current number of members
- Maximum number of members
- Record length
- Number of fields
- Total number of records
- Number of deleted records

- Reuse deleted records attribute
- Create date

***YES:** In addition to the file attributes, the file data is compared. A hashing algorithm is used to compare the file data. This option takes longer to run but it will detect small differences with a high degree of accuracy. This option requires operating system level V5R3 or greater.

Compare While Active

Indicate whether steps should be taken to ensure accuracy when comparing objects that are being updated.

Restrictions:

- This parameter is only valid when comparing database files, data areas, data queues and IFS stream files. A value of *NO is assumed for all other object types.
- The libraries and directories being compared must be defined by an existing sync attributes entry.
- Remote journaling must be specified in the related sync attributes entry.
- The remote journal apply job for the sync attributes entry must be active on the target system.
- The remote journal apply lag for the library should be small. You can key option **5** beside an entry on the Work With Synchronization Attributes panel to see the current journal lag.

Considerations:

- Objects that are compared while active use space in temporary storage. The temporary storage required may be up to twice the size of the source object being compared. Temporary work space is allocated for one object at a time, not for all objects in the library or directory simultaneously. So, the temporary space required to compare a whole library or directory is the same as would be required to compare the largest object in the library or directory.
- Due to additional processing, a comparison with CMPACT (*YES) may run longer.

The possible values are:

***NO:** Comparison for all libraries and directories is done without compare-while-active processing. Objects that are changing while being compared may report false mismatches due to the small delay inherent in replicating changes to the target object.

***YES:** Compare-while-active processing is used for all libraries and directories that meet the CMPACT requirements listed above. For those libraries and directories, objects are

compared reliably, even if the source object is being changed concurrently. Libraries and directories that do not meet the CMPACT requirements are processed as if CMPACT(*NO) were specified.

*AUTO: For all libraries and directories, each object is compared using the logic of CMPACT (*NO). If the source and target objects appear to differ, and if the object's type and journaling status make it eligible for comparing while active, a second comparison is done for the object using CMPACT(*YES). In this way, CMPACT(*YES) is attempted only for objects reporting a mismatch from an initial CMPACT(*NO) check. If most objects in the source and target libraries or directories match, and if most of the source objects are not being changed during the compare, CMPACT(*AUTO) will run significantly faster than CMPACT(*YES), with no loss in accuracy.

Compare Active Wait Time

Specify the number of seconds to wait for information about a target object to be returned from the target system. If the information is not returned in the allotted time, the compare for that object fails and processing continues with the next object to be compared. This parameter is ignored unless *YES is specified for the **Compare While Active** parameter.

Several factors can affect the time it takes to receive the required information from target system including:

- Journal apply lag on the target.
- The time required to extract the needed information from the target object. Extracting information takes longer for multi-member files and when CMPDTA(*YES) is specified to perform a byte-by-byte comparison of file data.

The possible values are:

120: Robot HA waits up to 2 minutes for the target system to return the required information for each object being compared.

seconds: Enter the number of seconds to wait for a reply from the target.

Save Active Wait Time

Specify the number of seconds to wait for a source object to reach a comparison checkpoint. If a checkpoint is not reached in the allotted time, the compare for that object fails and processing continues with the next object to be compared. This parameter is ignored unless *YES is specified for Compare While Active (CMPACT).

The possible values are:

120: Robot HA waits up to 2 minutes for the source object to reach a compare checkpoint.

seconds: Enter the number of seconds to wait for the source object to reach a checkpoint.

Allow Differences

Indicate whether certain differences between the objects should be ignored. If ignored, a difference in these attributes will not signal an error. You can specify up to four values for this parameter.

The possible single values are:

*NONE: No differences are allowed (ignored.)

***ALL:** Differences in File ID, REUSEDLT, Journal and Journal receivers, Change Date, and Number of Parameters are ignored.

The other possible values are:

***FID:** Differences in file ID (create date/time) are ignored.

***REUSEDLT:** Differences in the Reuse Deleted Records attribute are ignored.

***CHGDATE:** Differences in the last changed date/time for objects are ignored.

RLS:** Differences in **PGM, *****MODULE, *****SRVPGM and *****USRPRF objects due to the release of the operating system to which they were restored are ignored.

*JRN: Differences in journals and journal receivers are ignored.

***PWD:** Differences in user profile passwords are ignored.

***STATUS:** Differences in user profile status are ignored.

***TEXT:** Differences in the text description for user profiles and authorization lists are ignored.

Output

Indicate the type of output to generate.

The possible values are:

***NONE:** No output is generated but completion and error messages are placed in the job log.

*: A list of unmatched objects is displayed on the screen.

***PRINT:** A list of unmatched objects is printed.

***OUTFILE:** A list of unmatched objects is placed in an output file. The model outfile for a library comparision is RSPF046. The model outfile for an IFS directory comparison is RSPF047.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check IFS Directory (CHKDIRRSF)

The Check IFS Directory panel is used to determine if two IFS directories, their contents and directories match. A report or an output file is generated, listing and unmatched objects.

In a replication environment, you can use the **Refresh unmatched objects** parameter to cause any unmatched objects to be automatically refreshed.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **2** to display the Integrity Reports Menu.
- 2. Select option 1, Check directories for equality, to view the Check IFS Directory panel.

Check IFS	Directory (C	HKDIRRSF)			
Type choices, press Enter.					
Directory					
To directory	*FROMDIR				
Server ID	<u>*LOCAL</u> <u>*NO</u>	Name, *LOCAL *YES, *NO *YES, *NO, *AUTO Number *, *PRINT, *OUTFILE, *NONE			
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additio F24=More ke	Bottom nal parameters F12=Cancel ys			

Options

Directory

Enter the path for the local directory to check.

The possible values are:

path: Enter the full path name of an existing directory.

To Directory

Enter the path for the remote directory to check.

***FROMDIR:** The remote directory has the same name as the local directory.

path: Enter the full path name of the remote directory.

Server ID

Indicate the system containing the remote directory.

The possible values are:

***LOCAL:** The remote directory is on the same system as the local directory.

server-ID: Enter the Robot HA server ID of the system containing the remote directory.

Compare Data

Indicate whether Robot HA should compare the data (contents) of files in the directories or only the file attributes.

The possible values are:

***NO:** Only file attributes such as type, length and last change date are compared. This is faster than comparing file data.

***YES:** In addition to file attributes, file data is compared. A hashing algorithm is used to compare file data. This option takes longer to run but it will detect small differences with a high degree of accuracy.

Compare While Active

Indicate whether steps should be taken to ensure accuracy when comparing objects that are being updated.

Restrictions:

- This parameter is only valid when comparing stream files. A value of *NO is assumed for all other object types.
- The directories being compared must be defined by an existing sync attributes entry. See the Change IFS Sync Attributes (CHGRSFSA) command for more info.
- Remote journaling must be specified in the related sync attributes entry.
- The remote journal apply job for the sync attributes entry must be active on the target system.

• The remote journal apply lag for the library should be small. You can key option 5 beside an entry on the Work With Sync Attributes (WRKRSFSA) display to see the current journal lag.

Considerations:

- Objects that are compared while active use space in temporary storage. The temporary storage required may be up to twice the size of the source object being compared. Temporary work space is allocated for one object at a time, not for all objects in the directory simultaneously. So, the temporary space required to compare the whole directory is the same as would be required to compare the largest object in the directory.
- Due to additional processing, a comparison with CMPACT(*YES) may run longer.

The possible values are:

***NO:** If an object is changing while it is being compared, a false mismatch may be reported. This is due to the small delay inherent in replicating changes to the target object.

***YES:** The objects are compared reliably, even if the source object is being changed during the compare.

***AUTO:** Each object is compared using the logic of CMPACT(*NO). If the source and target objects appear to differ, and if the object's type and journaling status make it eligible for comparing while active, a second comparison is done for the object using CMPACT(*YES).

In this way, CMPACT(*YES) is attempted only for objects reporting a mismatch from an initial CMPACT(*NO) check. If most objects in the source and target directories match, and if most objects in the source directory are not being changed during the compare, CMPACT (*AUTO) will run significantly faster than CMPACT(*YES), with no loss in accuracy.

Compare Active Wait Time

Specify the number of seconds to wait for information about a target object to be returned from the target system. If the information is not returned in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless *YES is specified for **Compare While Active** (CMPACT).

Several factors can affect the time it takes to receive the required information from target system including:

- Journal apply lag on the target.
- The time required to extract the needed information from the target object. Extracting information takes longer when CMPDTA(*YES) is specified to perform a byte-by-byte comparison of file data.

120: Robot HA waits up to 2 minutes for the target system to return the required information for each object being compared.

seconds: Enter the number of seconds to wait for a reply from the target.

Output

Indicate the type of output to generate.

The possible values are:

***NONE:** No output is generated but completion and error messages are placed in the job log.

*: A list of unmatched objects is displayed on the screen.

***PRINT:** A list of unmatched objects is printed.

***OUTFILE:** A list of unmatched objects is placed into an output file.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check IFS Object (CHKIFSRSF)

The Check IFS Object panel is used to determine if two IFS objects are identical.

How to Get There

This panel can be viewed by running the command RBTHALIB/CHKIFSRSF.

Check IFS	Object (CHKIFSRSF)
Type choices, press Enter.	
From path	
To path	*FROMPATH
Server ID	<u>*LOCAL</u> Name, *LOCAL <u>*NO</u> *YES, *NO
	Bottom
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additional parameters F12=Cancel F24=More keys

Options

From Path

Enter the path for the local object to check.

The possible values are:

path: Enter the full path name of an existing object.

To Path

Enter the path for the remote object to check.

The possible values are:

***FROMPATH:** The remote object has the same path name as the local object.

path: Enter the full path name of the remote object.

Server ID

Indicate the machine containing the remote object.

The possible values are:

***LOCAL:** The remote object is on the same machine as the local object.

server-ID: Enter the Robot HA server ID of the machine containing the remote object.

Compare Data

Indicate whether Robot HA should compare the data (contents) of the objects or only the object attributes.

The possible values are:

***NO:** Only file attributes such as type and length are compared. This is faster than comparing file data.:

***YES:** In addition to the file attributes, the object data is compared. A hashing algorithm is used to compare the object data. This option takes longer to run but it will detect small differences with a high degree of accuracy.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check Libraries (CHKLIBRSF)

The Check Libraries is used to determine if the objects in two libraries match. A report or an output file is generated, listing any unmatched objects.

In a replication environment, you can use the **Refresh unmatched objects** parameter to cause any unmatched objects to be automatically refreshed.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **2** to display the Integrity Reports Menu.
- 2. Select option **3**, Check libraries for equality, to view the Check Libraries panel.

Check Li	braries (CHKLI	BRSF)
Type choices, press Enter.		
From library	*FROMLIB *LOCAL *DFT *ALL *ALL *NO 120 120 *NONE *NONE	Name Name, *FROMLIB Name, *LOCAL Name, *DFT Name, generic*, *ALL *ALL, *ALRTBL, *BNDDIR, *YES, *NO *YES, *NO *YES, *NO, *AUTO Number Number *ALL, *NONE, *FID *, *PRINT, *OUTFILE, *NONE
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additiona F24=More keys	Bottom al parameters F12=Cancel a

Options

From Library

Enter the name of the local library to check.

The possible values are:

name: Enter the name of an existing library.

To Library

Enter the name of the remote library to check.

The possible values are:

***FROMLIB:** The remote library has the same name as the local library.

name: Enter the name of the remote library.

Server ID

Indicate the system containing the remote library.

***LOCAL:** The remote library is on the same system as the local library.

server-ID: Enter the Robot HA server ID of the system containing the remote library.

Set Name

Specify the library set name to check. The set name is used together with the from lib, to lib and server ID to determine if the library is defined for replication.

In a non-replication environment, always specify ***DFT** for this value.

The possible values are:

***DFT:** The default set name is used.

set-name: Enter a valid set name to check.

Object Name

Indicate which objects to check.

The possible values are:

***ALL:** All objects in the library are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object Type

Indicate the type of objects to check.

The possible values are:

***ALL:** All types of objects are checked.

object-type: List the types of objects you want to check.

Compare Data

Indicate whether Robot HA should compare the data (contents) of the files or only the file attributes.

***NO:** Only the following file attributes are compared. This is faster than comparing file data.

- Current number of members
- Maximum number of members
- Record length
- Number of fields
- Total number of records
- Number of deleted records
- Reuse deleted records attribute
- Create date

***YES:** In addition to the file attributes, the file data is compared. A hashing algorithm is used to compare the file data. This option takes longer to run but it will detect small differences with a high degree of accuracy. This option requires operating system level V5R3 or greater.

Compare While Active

Indicate whether steps should be taken to ensure accuracy when comparing objects that are being updated.

Restrictions:

- The remote journal apply job for the sync attributes entry must be active on the target system to work properly. If the remote journal apply job is not active this can incorrectly report errors stating the object is unavailable on the target.
- This parameter is only valid when comparing database files, data areas and data queues. A value of *NO is assumed for all other object types.
- The libraries being compared must be defined by an existing sync attributes entry. See the Change Library Sync Attributes panel (Option **2** from the Work with Sync Attributes - Libraries panel), for more info.
- Remote journaling must be specified in the related sync attributes entry.
- The remote journal apply lag for the library should be small. You can key option 5 beside an entry on the Work With Sync Attributes (WRKRSFSA) display, or use the Display Library Sync Status (DSPRSFSS) command to see the current journal lag.

Considerations:

- Objects that are compared while active use space in temporary storage. The temporary storage required may be up to twice the size of the source object being compared. Temporary work space is allocated for one object at a time, not for all objects in the library simultaneously. So, the temporary space required to compare the whole library is the same as would be required to compare the largest object in the library.
- Due to additional processing, a comparison with CMPACT(*YES) may run longer.
- If an object is locked and we are unable to get the information to compare, we will indicate on the report that the object was unavailable to compare.

***NO:** If an object is changing while it is being compared, a false mismatch may be reported. This is due to the small delay inherent in replicating changes to the target object.

***YES:** The objects are compared reliably, even if the source object is being changed during the compare.

***AUTO:** Each object is compared using the logic of CMPACT(*NO). If the source and target objects appear to differ, and if the object's type and journaling status make it eligible for comparing while active, a second comparison is done for the object using CMPACT(*YES).

In this way, CMPACT(*YES) is attempted only for objects reporting a mismatch from an initial CMPACT(*NO) check. If most objects in the source and target libraries match, and if most objects in the source library are not being changed during the compare, CMPACT (*AUTO) will run significantly faster than CMPACT(*YES), with no loss in accuracy.

Compare Active Wait Time

Specify the number of seconds to wait for information about a target object to be returned from the target system. If the information is not returned in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless ***YES** is specified for **Compare While Active** (CMPACT).

Several factors can affect the time it takes to receive the required information from target system including:

- Journal apply lag on the target.
- The time required to extract the needed information from the target object. Extracting information takes longer for multi-member files and when CMPDTA(*YES) is specified to perform a byte-by-byte comparison of file data.

The possible values are:

120: Robot HA waits up to 2 minutes for the target system to return the required information for each object being compared.

seconds: Enter the number of seconds to wait for a reply from the target.

Save Active Wait Time

Specify the number of seconds to wait for a source object to reach a comparison checkpoint. If a checkpoint is not reached in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless ***YES** is specified for **Compare While Active** (CMPACT).

The possible values are:

120: Robot HA waits up to 2 minutes for the source object to reach a compare checkpoint.

seconds: Enter the number of seconds to wait for the source object to reach a checkpoint.

Allow Differences

Indicate whether certain differences between the objects should be ignored. If ignored, a difference in these attributes will not signal an error.

You can specify up to five values for this parameter.

The possible single values are:

***NONE:** No differences are allowed (ignored.)

***ALL:** Differences in File ID, REUSEDLT, Journal and Journal receivers, Change Date, and Number of Parameters are ignored.

The other possible values are:

***FID:** Differences in file ID (create date/time) are ignored.

***REUSEDLT:** Differences in the Reuse Deleted Records attribute are ignored.

***CHGDATE:** Differences in the last changed date/time for objects is ignored.

RLS:** Differences in **PGM, *****MODULE, and *****SRVPGM objects due to the release of the operating system to which they were restored are ignored.

*JRN: Differences in journals and journal receivers are ignored.

***NBRPARMS:** Differences in the max and min number of parameters for program objects is ignored.

NOTE: This option is retained for backwards compatibility only. Use the *RLS option instead. Also, when comparing programs on different systems running different versions of the operating system, the number of parameters may be reported differently by the operating system, even if the programs are identical.

Output

Indicate the type of output to generate.

The possible values are:

***NONE:** No output is generated but completion and error messages are placed in the job log.

*: A list of unmatched objects is displayed on the screen.

***PRINT:** A list of unmatched objects is printed.

***OUTFILE:** A list of unmatched objects is placed into an output file.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check Objects (CHKOBJRSF)

The Check Objects panel is used to determine if objects match.

How to Get There

This panel can be viewed by running the command **RBTHALIB/CHKOBJRSF**. Library objects can also be checked by entering a specific name for the Object Name parameter on the Check Libraries panel.

Check	Objects (CHKOB	BJRSF)
Type choices, press Enter.		
From object	*LIBL *FROMOBJ *FROMLIB *ALL *LOCAL *DFT *NO *NO 120 120 *NONE	Name Name, *LIBL, *CURLIB *ALRTBL, *BNDDIR, *CHTFMT Name, *FROMOBJ Name, *FROMLIB, *LIBL Name, *ALL, *FIRST Name, *LOCAL Name, *DFT *YES, *NO *YES, *NO Number Number *ALL, *NONE, *FID
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	More F13=How to use this display

Options

From Object

Enter the qualified name of the local object to check.

The possible object values are:

name: Enter the name of an existing object.

The possible library values are:

*LIBL: The library list is used to find the object.

***CURLIB:** The object is found in the job's current library.

name: Enter the name of the library containing the object.

Object Type

Specify the type of object to check.

The possible values are:

type: Enter a valid object type.

To Object

Enter the qualified name of the remote object to check.

The possible object values are:

***FROMOBJ:** The remote object has the same name as the local object.

name: Enter the name of the remote object.

The possible library values are:

***FROMLIB:** The remote library is the same as the local library.

*LIBL: The library list is used to find the remote object.

***CURLIB:** The remote object is found in the remote job's current library.

name: Enter the name of the library containing the remote object.

Member Name

Enter the name of the member to compare for database files. If a value other than *ALL is specified for Member Name and *YES is specified for Compare Data, then any file-level differences are ignored. Only differences in the specified file members are noted.

The possible values are:

***ALL:** All database file members are checked.

***FIRST:** Only the first member in database files is checked.

name: Enter the name of the database file member to check.

Server ID

Indicate the machine containing the remote object.

The possible values are:

***LOCAL:** The remote object is on the same machine as the local object.

server-ID: Enter the Robot HA server ID of the machine containing the remote object.

Set Name

Specify the library set name to check. The set name is used together with the from lib, to lib and server ID to determine if the library is defined for replication.

In a non-replication environment, always specify ***DFT** for this value.

The possible values are:

***DFT:** The default set name is used.

set-name: Enter a valid set name to check.

Compare Data

Indicate whether Robot HA should compare the data (contents) of the files or only the file attributes.

The possible values are:

***NO:** Only the following file attributes are compared. This is faster than comparing file data.

- Current number of members
- Maximum number of members
- Record length
- Number of fields
- Total number of records
- Number of deleted records
- Reuse deleted records attribute
- Create date

***YES:** In addition to the file attributes, the file data is compared. A hashing algorithm is used to compare the file data. This option takes longer to run but it will detect small differences with a high degree of accuracy. This option requires operating system level V5R3 or greater.

Compare While Active

Indicate whether steps should be taken to ensure accuracy when comparing objects that are being updated.

Restrictions:

- The remote journal apply job for the sync attributes entry must be active on the target system to work properly. If the remote journal apply job is not active this can incorrectly report errors stating the object is unavailable on the target.
- This parameter is only valid when comparing database files, data areas and data queues. A value of *NO is assumed for all other object types.
- The libraries being compared must be defined by an existing sync attributes entry. See the Change Library Sync Attributes panel (Option **2** from the Work with Sync Attributes - Libraries panel) for more info.
- Remote journaling must be specified in the related sync attributes entry.
- The remote journal apply lag for the library should be small. You can key option **5** beside an entry on the Work With Sync Attributes display, or use the Display Library Sync Status (DSPRSFSS) command to see the current journal lag.

Considerations:

 Objects that are compared while active use space in temporary storage. The temporary storage required may be up to twice the size of the source object being compared. Temporary work space is allocated for one object at a time, not for all objects in the library simultaneously. So, the temporary space required to compare the whole library is the same as would be required to compare the largest object in the library.

- Due to additional processing, a comparison with CMPACT(*YES) may run longer.
- If an object is locked and we are unable to get the information to compare, we will indicate on the report that the object was unavailable to compare.

***NO:** If an object is changing while it is being compared, a false mismatch may be reported. This is due to the small delay inherent in replicating changes to the target object.

***YES:** The objects are compared reliably, even if the source object is being changed during the compare.

***AUTO:** The objects are compared using the logic of CMPACT(*NO). If the source and target objects appear to differ, and if the object's type and journaling status make it eligible for comparing while active, a second comparison is done for the object using CMPACT(*YES).

In this way, CMPACT(*YES) is attempted only for objects reporting a mismatch from an initial CMPACT(*NO) check. If most objects in the source and target libraries match, and if most objects in the source library are not being changed during the compare, CMPACT (*AUTO) will run significantly faster than CMPACT(*YES), with no loss in accuracy.

Compare Active Wait Time

Specify the number of seconds to wait for information about a target object to be returned from the target system. If the information is not returned in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless ***YES** is specified for **Compare While Active** (CMPACT).

Several factors can affect the time it takes to receive the required information from target system including:

- Journal apply lag on the target.
- The time required to extract the needed information from the target object. Extracting information takes longer for multi-member files and when CMPDTA(*YES) is specified to perform a byte-by-byte comparison of file data.

The possible values are:

120: Robot HA waits up to 2 minutes for the target system to return the required information for each object being compared.

seconds: Enter the number of seconds to wait for a reply from the target.

Save Active Wait Time

Specify the number of seconds to wait for a source object to reach a comparison checkpoint. If a checkpoint is not reached in the allotted time, the compare aborts.

This parameter is ignored unless ***YES** is specified for **Compare While Active** (CMPACT).

The possible values are:

120:Robot HA waits up to 2 minutes for the source object to reach a compare checkpoint.

seconds: Enter the number of seconds to wait for the source object to reach a checkpoint.

Allow Differences

Indicate whether certain differences between the objects should be ignored. If ignored, a difference in these attributes will not signal an error.

You can specify up to five values for this parameter.

The possible single values are:

***NONE:** No differences are allowed (ignored.)

***ALL:** Differences in File ID, REUSEDLT, Journal and Journal receivers, Change Date, and Number of Parameters are ignored.

The other possible values are:

***FID:** Differences in file ID (create date/time) are ignored.

***REUSEDLT:** Differences in the Reuse Deleted Records attribute are ignored.

***CHGDATE:** Differences in the last changed date/time for objects is ignored.

RLS:** Differences in **PGM, *****MODULE, and *****SRVPGM objects due to the release of the operating system to which they were restored are ignored.

*JRN: Differences in journals and journal receivers are ignored.

***NBRPARMS:** Differences in the max and min number of parameters for program objects is ignored.

NOTE: This option is retained for backwards compatibility only. Use the *RLS option instead. Also, when comparing programs on different systems running different versions of the operating system, the number of parameters may be reported differently by the operating system, even if the programs are identical.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check Replication Group (CHKGRPRSF)

The Check Replication Group panel is used to determine if the objects in a replication group match. A report or an output file is generated, listing any unmatched objects.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **2** to display the Integrity Reports Menu.
- 2. Select option **2**, Compare items in replication group, to view the Check Replication Group panel.

Check Poplication Group (CHKGPPPSE)			
Type choices, press Enter.			
Group	*ALL *ALL *NO *AUTO 120 120 *NONE *	Name Name, generic*, *ALL *ALL, *ALRTBL, *BNDDIR, *YES, *NO *YES, *NO, *AUTO Number Number *ALL, *NONE, *FID *, *PRINT, *OUTFILE, *NONE	
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additiona F24=More keys	Bottom al parameters F12=Cancel s	

Options

Group

Enter the name of the group to check.

The possible values are:

name: Enter the name of an existing replication group.

Object Name

Indicate which objects to check.

The possible values are:

***ALL:** All objects in the group are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object Type

Indicate the type of objects to check.

***ALL:** All types of objects are checked.

object-type: List the types of objects you want to check.

Compare Data

Indicate whether Robot HA should compare the data (contents) of the files or only the file attributes.

The possible values are:

***NO:** Only the following file attributes are compared. This is faster than comparing file data.

- Current number of members
- Maximum number of members
- Record length
- Number of fields
- Total number of records
- Number of deleted records
- Reuse deleted records attribute
- Create date

***YES:** In addition to the file attributes, the file data is compared. A hashing algorithm is used to compare the file data. This option takes longer to run but it will detect small differences with a high degree of accuracy. This option requires operating system level V5R3 or greater.

Compare While Active

Indicate whether steps should be taken to ensure accuracy when comparing objects that are being updated.

Restrictions:

- The remote journal apply job for the sync attributes entry must be active on the target system to work properly. If the remote journal apply job is not active this can incorrectly report errors stating the object is unavailable on the target.
- This parameter is only valid when comparing database files, data areas and data queues. A value of *NO is assumed for all other object types.
- Remote journaling must be specified in the related sync attributes entry.

• The remote journal apply lag for the group should be small. You can enter option **5** beside an entry on the Work With Sync Attributes (WRKRSFSA) display, or use the Display Group Sync Status (DSPRSFSS) command to see the current journal lag.

Considerations:

- Objects that are compared while active use space in temporary storage. The temporary storage required may be up to twice the size of the source object being compared. Temporary work space is allocated for one object at a time, not for all objects in the group simultaneously. So, the temporary space required to compare the whole group is the same as would be required to compare the largest object in the group.
- Due to additional processing, a comparison with CMPACT(*YES) may run longer.
- If an object is locked and we are unable to get the information to compare, we will indicate on the report that the object was unavailable to compare.

The possible values are:

***NO:** If an object is changing while it is being compared, a false mismatch may be reported. This is due to the small delay inherent in replicating changes to the target object.

***YES:** The objects are compared reliably, even if the source object is being changed during the compare.

***AUTO:** Each object is compared using the logic of CMPACT(*NO). If the source and target objects appear to differ, and if the object's type and journaling status make it eligible for comparing while active, a second comparison is done for the object using CMPACT(*YES).

In this way, CMPACT(*YES) is attempted only for objects reporting a mismatch from an initial CMPACT(*NO) check. If most objects in the group match, and if most objects in the source libraries are not being changed during the compare, CMPACT(*AUTO) will run significantly faster than CMPACT(*YES), with no loss in accuracy.

Compare Active Wait Time

Specify the number of seconds to wait for information about a target object to be returned from the target system. If the information is not returned in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless ***YES** is specified for **Compare While Active** (CMPACT).

Several factors can affect the time it takes to receive the required information from target system including:

• Journal apply lag on the target.

 The time required to extract the needed information from the target object. Extracting information takes longer for multi-member files and when CMPDTA(*YES) is specified to perform a byte-by-byte comparison of file data.

The possible values are:

120: Robot HA waits up to 2 minutes for the target system to return the required information for each object being compared.

seconds: Enter the number of seconds to wait for a reply from the target.

Save Active Wait Time

Specify the number of seconds to wait for a source object to reach a comparison checkpoint. If a checkpoint is not reached in the allotted time, the compare for that object fails and processing continues with the next object to be compared.

This parameter is ignored unless ***YES** is specified for Compare While Active (CMPACT).

The possible values are:

120: Robot HA waits up to 2 minutes for the source object to reach a compare checkpoint.

seconds: Enter the number of seconds to wait for the source object to reach a checkpoint.

Allow Differences

Indicate whether certain differences between the objects should be ignored. If ignored, a difference in these attributes will not signal an error.

You can specify up to five values for this parameter.

The possible single values are:

***NONE:** No differences are allowed (ignored.)

***ALL:** Differences in File ID, REUSEDLT, Journal and Journal receivers, Change Date, and Number of Parameters are ignored.

The other possible values are:

***FID:** Differences in file ID (create date/time) are ignored.

***REUSEDLT:** Differences in the Reuse Deleted Records attribute are ignored.

***CHGDATE:** Differences in the last changed date/time for objects is ignored.
RLS:** Differences in **PGM, *****MODULE, and *****SRVPGM objects due to the release of the operating system to which they were restored are ignored.

*JRN: Differences in journals and journal receivers are ignored.

***NBRPARMS:** Differences in the max and min number of parameters for program objects is ignored.

NOTE: This option is retained for backwards compatibility only. Use the *RLS option instead. Also, when comparing programs on different systems running different versions of the operating system, the number of parameters may be reported differently by the operating system, even if the programs are identical.

Output

Indicate the type of output to generate.

The possible values are:

***NONE:** No output is generated but completion and error messages are placed in the job log.

*: A list of unmatched objects is displayed on the screen.

***PRINT:** A list of unmatched objects is printed.

***OUTFILE:** A list of unmatched objects is placed into an output file.

File to Receive Output

Enter the qualified name of the file to receive the unmatched object information.

This parameter is required if *OUTFILE is specified for the Output parameter. The model outfile record format is based on RSPF046.

The possible values are:

name: Enter the name of the output file. If the file does not exist it will be created.

The possible library values are:

*LIBL: The file is found using the library list. If the file does not exist, *CURLIB is assumed.

***CURLIB:** The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Check System Items (CHKSYSRSF)

The Check System Items panel is used to determine if system objects, such as user profiles or authorization lists, match. A report or an output file is generated, listing any unmatched objects.

In a replication environment, you can use the **Refresh unmatched objects** parameter to cause any unmatched objects to be automatically refreshed.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **2** to display the Integrity Reports Menu.
- 2. Select option **4**, Check system items for equality, to view the Check System Items panel.

Check Syst	tem Items (CHKSYSRSF)
Type choices, press Enter.	
Server ID	Name <u>*ALL</u> Name, generic*, *ALL <u>*ALL</u> *ALL, *USRPRF, *AUTL
Allow differences	*NONE *ALL, *NONE, *PWD, *STATUS
Output	* *, *PRINT, *OUTFILE, *NONE
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	Botto F10=Additional parameters F12=Cancel F24=More keys

Options

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source system. For an Robot HA pre-processing program running on a server system, *CURRENT refers to the requester system. *CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server system, or a batch job submitted from one of the above two job types.

- 1. *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. *CURRENT is allowed within other Robot HA server jobs that connect via TCP/IP.
- 2. For SNA/APPC connections, the user profile and password used to connect back to the *CURRENT system are taken from the server entry that is associated with the job in which this command is run. If no server entry is associated with the job, the default Robot HA signon is used. Use the WRKRSFRQS command to view or change the server IDs that are associated with requests sent to this system.

server-ID: Enter the name of an entry in the server directory on your system.

Object Name

Indicate which objects to check.

The possible values are:

***ALL:** All system objects are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object Type

Indicate the type of objects to check.

The possible values are:

*ALL: User profiles and authorization lists are checked.

*USRPRF: User profiles are checked.

***AUTL:** Authorization lists are checked.

Allow Differences

Indicate whether certain differences between the objects should be ignored. If ignored, a difference in these attributes will not signal an error.

You can specify up to four values for this parameter

The possible single values are:

***NONE:** No differences are allowed (ignored.)

***ALL:** Differences in password, status, text, and release-dependent information are ignored. Not all differences apply to all object types.

The other possible values are:

***PWD:** Differences in user profile passwords are ignored.

***STATUS:** Differences in user profile status are ignored.

***TEXT:** Differences in the text description for user profiles and authorization lists are ignored.

***RLS:** Differences in release-dependent information for user profiles are ignored. This may be needed when comparing objects on systems running different OS release levels.

Output

Indicate the type of output to generate.

The possible values are:

***NONE:** No output is generated but completion and error messages are placed in the job log.

*: A list of unmatched objects is displayed on the screen.

***PRINT:** A list of unmatched objects is printed.

***OUTFILE:** A list of unmatched objects is placed into an output file.

File to Receive Output

Enter the qualified name of the file to receive the unmatched object information.

This parameter is required if ***OUTFILE** is specified for the **Output** parameter. The model outfile record format is based on RSPF061.

The possible values are:

name: Enter the name of the output file. If the file does not exist it will be created.

The possible library values are:

LIBL:** The file is found using the library list. If the file does not exist, **CURLIB is assumed.

***CURLIB:** The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Clean Defined Items (CLNATTRSF)

The Clean Defined Items panel is used to clean some or all of the libraries, IFS directories and system items defined on your Work With Sync Attributes (WRKRSFSA) list. The indicated libraries, directories and system items on the server machine are cleaned, using libraries and directories on this machine for reference.

The Clean Library (CLNLIBRSF) command is run for each selected library. The Clean Directory command (CLNDIRRSF) is run for each selected IFS directory. The Clean System Items command (CLNSYSRSF) is run for system items.

How to Get There

Run the command **RBTHALIB/CLNATTRSF.**

Clean Defi	ned Items (CLN	ATTRSF)	
Type choices, press Enter.			
Server ID	*ALL *NONE	Name Name, *All, *NONE, *gen Name, *NONE, *generic*	eric*
Object name	<u>*ALL</u>	Name, generic*, *ALL	
<pre>UDject type</pre>	<u>*HLL</u>	*HLL, *HLRIBL, *BNDDIR,	
Remove orphan members	*N0 *YES *YES *N0 *USRPRF	*YES, *NO *YES, *NO *YES, *NO *YES, *NO *YES, *NO *YES, *NO Name, *NONE, *USRPRF Name, *LIBL, *CURLIB	
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additiona F24=More keys	l parameters F12=Cance	Bottom l

Options

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current machine are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source machine. For an Robot HA pre-processing program running on a server machine, *CURRENT refers to the requester machine.

*CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server machine, or a batch job submitted from one of the above two job types.

NOTE: *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. It is allowed within other Robot HA server jobs that connect via TCP/IP.

server-ID: Enter the name of an entry in the server directory on your machine.

Libraries to Include

Enter the name or generic name for the libries to be included.

The possible values are:

***ALL:** All libraries for the given server that are defined on the Work With Sync Attributes panel are included.

***NONE:** No libraries are included.

generic-name: Enter a generic name for the libraries to include. An asterisk (*) in the generic specification will match any string of zero or more characters in the name. An underscore (_) in the generic specification will match any single character in the name.

Example 1. This specification will include all names that begin with "A":

'A*'

Example 2. This specification will include all names that contain "ABC":

'*ABC*'

Example 3. This specification will include all names that end with "A", followed by any single character, followed by "C":

'*A_C'

library-name: Enter a library name.

Omit Libraries

Enter the name or generic name for libraries to omit. Up to 20 specifications can be entered.

The possible values are:

***NONE:** No libraries are omitted.

generic-name: Enter a generic name for libraries to omit. An asterisk (*) in the generic specification will match any string of zero or more characters in the name. An underscore (_) in the generic

specification will match any single character in the name.

Example 1. This specification will include all names that begin with "A":

'A*'

Example 2. This specification will include all names that contain "ABC":

'*ABC*'

Example 3. This specification will include all names that end with "A", followed by any single character, followed by "C":

'*A_C'

library-name: Enter the name of a library to omit.

Object Name

Indicate which objects to check.

The possible values are:

***ALL:** All objects in the selected libraries are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object Type

Indicate the type of objects to check.

The possible values are:

***ALL:** All types of objects are checked.

object-type: List the types of objects you want to check.

Remove Orphan Members

For physical and logical files, indicate whether Robot HA should remove orphan members.

The possible values are:

***NO:** Orphan members are not removed.

***YES:** Members in files on the target machine are removed if there is no matching member in the corresponding file on this machine.

Delete Orphan Libraries

Indicate whether orphan libraries should be deleted.

The possible values are:

***NO:** Orphan libraries are not deleted.

***YES:** If the library being compared to does not exist on this machine, it is deleted from the server machine.

Clean IFS Directories

Indicate whether to submit cleanup jobs for IFS entries in the list.

The possible values are:

***YES:** All IFS entries on the Work With Sync Attributes panel that are defined for the given server are included.

*NO: No IFS cleanup jobs are submitted.

Clean System Items

Indicate whether to submit cleanup jobs for system items in the list.

The possible values are:

***YES:** System items defined on the Work With Sync Attributes panel that meet the criteria specified in the Server (SERVER), Object Name (OBJ) and Object Type (OBJTYPE) parameters are included.

***NO:** No system cleanup jobs are submitted.

Delete Orphan Sync Attributes

Indicate whether to delete sync attributes entries that define replication for a library or IFS directory that no longer exists on this machine.

The possible values are:

***NO:** Orphan sync attribute entries are not deleted.

***YES:** Sync attribute entries that define replication for a library or IFS directory which no longer exists on this machine are deleted.

Allow on Production

Indicate whether objects can be deleted from a machine that is not set to the *BACKUP replication role.

The replication role is set automatically during a role swap and can be changed manually with the Change RSF Defaults command (CHGRSFDFT).

NOTE: This does not affect the deletion of orphan sync attribute entries. See the **Delete Orphan Sync Attributes** parameter.

The possible values are:

***NO:** This command will end in error if attempting to delete objects from the production machine.

***YES:** The command can be used to delete objects from the production machine. Use caution when selecting this option. One reason to allow the deleting of objects from the production machine would be when the "from" and "to" libraries are both on the same machine.

Job Description

Specify an optional job description to use. When a value other than ***NONE** is specified, a separate compare job is submitted for each selected library and directory.

The possible single values are:

***NONE:** No job description is used. The cleanup for each selected item runs in the current job.

***USRPRF:** The job description from the user profile of the user running this command is used.

The possible Job Description values are:

name: Enter the name of an existing job description.

The possible library values are:

***LIBL:** The library list is used to find the job description.

***CURLIB:** The job description in the current library is used.

name: Enter the name of the library containing the job description.

Command Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Clean Directory (CLNDIRRSF)

The Clean Directory panel is used to delete objects from the FROM system that do not exist in the corresponding directory on the REF system.

Objects that are replicated without journaling are not always deleted from the backup system when the associated object is deleted from the production system. This panel provides a way to periodically clean up those orphan objects.

NOTE: Using remote journals is always recommended for directories and is used by default. This command does not need to be run for directories unless you want your omitted directories cleaned.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** to display the <u>Work with Sync</u> <u>Attributes</u> panel.
- 2. Enter option 16 and prompt (F4) in front of a directory.
- 3. Click the command entry section of the sbmjob and prompt (F4).

	Clean Directory	(CLNDIRRSF)	
Type choices, press Enter			
Delete orphan objects from Directory Server ID	m: 	Name, *LOCAL	
To directory Server ID Allow delete on production	 n <u>*NO_</u>	Name, *LOCAL, > *YES, *NO	CUR RENT
			Bottom
F3=Exit F4=Prompt F5= F24=More keys	Refresh F12=Can	cel F13=How to use	this display

Options

Delete orphan objects from

Enter the name of the directory and the system from which to delete the orphan objects.

The possible directory values are:

path: Enter the path for an existing directory.

The possible Server ID values are:

***LOCAL:** The objects are deleted from the specified directory on this system (the system on which this command is run).

server-ID: Enter the Robot HA server ID of the system from which to delete the objects.

Compare to

Enter the names of the directory and system to check. Objects in the directory and system specified by the FROM parameter are deleted if an object with the same name and type does not exist in the directory and system specified by the REF parameter.

The possible directory values are:

***FROMDIR:** The name of the reference directory is the same as the name of the directory in which the objects will be deleted.

name: Enter the name of the directory to check.

The possible Server ID values are:

***LOCAL:** The reference directory is on the same system as the directory from which objects will be deleted.

***CURRENT:** The reference directory is on this system (the system on which this command is run) but the directory from which objects will be deleted in on another system.

server-ID: Enter the Robot HA server ID of the system which contains the objects you want to keep. Objects will be deleted from the other system if they do not exist in the indicated directory on the system referred to here.

Allow delete on production

Indicate whether the directory from which objects will be deleted can be on a system that is not set to the *BACKUP replication role.

The replication role is set automatically during a role swap and can be changed manually with the Change RSF Defaults command (CHGRSFDFT).

The possible values are:

***NO:** This command will end in error if attempting to delete objects from the production system.

***YES:** The command can be used to delete objects from the production system. Use caution when selecting this option. One reason to allow the deleting of objects from the production system would be when the "from" and "to" directories are both on the same system.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Clean Library (CLNLIBRSF)

The Clean Library panel is used to delete objects on the FROM system that do not exist in the associated library on the REF system.

Objects that are replicated without journaling are not always deleted from the backup system when the associated object is deleted from the production system. This panel provides a way to periodically clean up those orphan objects.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** to display the <u>Work with Sync</u> <u>Attributes</u> panel.
- 2. Enter option **16** and prompt (**F4**) in front of a library.
- 3. Click the command entry section of the sbmjob and prompt (F4).

Clean L	_ibrary (CLNL	IBRSF)
Type choices, press Enter.		
Delete orphan objects from: Library	*LOCAL *ALL *ALL *NO *NO *NO	Name Name, *LOCAL Name, *FROMLIB Name, *LOCAL, *CURRENT Name, generic*, *ALL *ALL, *ALRTBL, *BNDDIR, Character value, *YES, *NO Character value, *YES, *NO *YES, *NO
F3=Exit F4= <u>P</u> rompt F5=Refresh F24=More keys	F12=Cancel	Bottom F13=How to use this display

Options

Delete orphan objects from

Enter the name of the library and the system from which to delete the orphan objects.

The possible library values are:

path: Enter the path for an existing library.

The possible Server ID values are:

***LOCAL:** The objects are deleted from the specified library on this system (the system on which this command is run).

server-ID: Enter the Robot HA server ID of the system from which to delete the objects.

Compare to

Enter the names of the library and system to check. Objects in the library and system specified by the FROM parameter are deleted if an object with the same name and type does not exist in the library and system specified by the REF parameter.

The possible library values are:

***FROMLIB:** The name of the reference library is the same as the name of the library in which the objects will be deleted.

name: Enter the name of the library to check.

The possible Server ID values are:

***LOCAL:** The reference library is on the same system as the library from which objects will be deleted.

***CURRENT:** The reference library is on this system (the system on which this command is run) but the directory from which objects will be deleted in on another system.

server-ID: Enter the Robot HA server ID of the system which contains the objects you want to keep. Objects will be deleted from the other system if they do not exist in the indicated library on the system referred to here.

Object name

Indicate which objects to check.

The possible values are:

*ALL: All objects in the library are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object type

Indicate the type of objects to check.

The possible values are:

***ALL:** All types of objects are checked.

object-type: List the types of objects you want to check.

Remove orphan members

For physical and logical files that exist in both the "from" and "to" libraries, indicate whether Robot HA should compare and remove members in the "from" library that do not exist in the "to" library.

The possible values are:

***NO:** Orphan members are not removed.

***YES:** Members in files in the "from" library are removed in there is no matching member in the corresponding file in the "to" library.

Delete orphan library

Indicate whether the "from" library should be deleted when the "to" library does not exist.

The possible values are:

***NO:** An orphan "from" library is not deleted.

***YES:** If the "to" library does not exist, the "from" library is deleted.

Allow delete on production

Indicate whether the library from which objects will be deleted can be on a system that is not set to the *BACKUP replication role.

The replication role is set automatically during a role swap and can be changed manually with the Change RSF Defaults command (CHGRSFDFT).

The possible values are:

***NO:** This command will end in error if attempting to delete objects from the production system.

***YES:** The command can be used to delete objects from the production system. Use caution when selecting this option. One reason to allow the deleting of objects from the production system would be when the "from" and "to" libraries are both on the same system.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Clean Replication Group (CLNGRPRSF)

The Clean Replication Group panel is used to clean the libraries in a replication group. Libraries on the server are cleaned, using libraries and directories on this system for reference.

The Clean Library command (CLNLIBRSF) is run for each selected library in the group.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** to display the <u>Work with Sync</u> <u>Attributes</u> panel.
- 2. Enter option **16** and prompt (**F4**) in front of a group.
- 3. Click the command entry section of the sbmjob and prompt (F4).

Clean Replic	ation Group (CLNGRPRSE)
Type choices, press Enter.	
Group	Name <u>*ALL</u> Name, generic*, *ALL *ALL *ALRTBL, *BNDDIR,
+ for more values	<u></u> ,,,
Remove orphan members Delete orphan library Delete orphan sync attributes . Allow delete on production Job description Library	*N0 Character value, *YES, *N0 *N0 Character value, *YES, *N0 *N0 *YES, *N0 *USRPRF Name, *NONE, *USRPRF Name, *LIBL, *CURLIB
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	Bottom F10=Additional parameters F12=Cancel F24=More keys

Options

Group

Enter the name of the replication group.

The possible values are:

group-name: Enter a replication group name.

Object name

Indicate which objects to check.

The possible values are:

***ALL:** All objects in the group libraries are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object type

Indicate the type of objects to check.

The possible values are:

***ALL:** All types of objects are checked.

object-type: List the types of objects you want to check.

Remove orphan members

For physical and logical files, indicate whether Robot HA should remove orphan members.

The possible values are:

***NO:** Orphan members are not removed.

***YES:** Members in files on the target system are removed if there is no matching member in the corresponding file on this system.

Delete orphan library

Indicate whether orphan libraries should be deleted.

The possible values are:

***NO:** Orphan libraries are not deleted.

***YES:** If the library being compared to does not exist on this system, it is deleted from the server.

Delete orphan sync attributes

Indicate whether to delete sync attribute entries that define replication for a library that no longer exists on this system.

The possible values are:

***NO:** Orphan sync attribute entries are not deleted.

***YES:** Sync attribute entries that define replication for a library which no longer exists on this system are deleted.

Allow delete on production

Indicate whether objects can be deleted from a system that is not set to the *BACKUP replication role.

The replication role is set automatically during a role swap and can be changed manually with the Change RSF Defaults command (CHGRSFDFT).

NOTE: This does not affect the deletion of orphan sync attribute entries. See the **Delete orphan sync attributes** parameter.

The possible values are:

***NO:** This command will end in error if attempting to delete objects from the production system.

***YES:** The command can be used to delete objects from the production system. Use caution when selecting this option. One reason to allow the deleting of objects from the production system would be when the server ID associated with the group points to the same system.

Job description

Specify an optional job description to use. When a value other than *NONE is specified, a separate clean job is submitted for library in the group.

The possible values are:

***NONE:** No job description is used. The cleanup for each selected item runs in the current job.

***USRPRF:** The job description from the user profile of the user running this command is used.

name: Enter the name of an existing job description.

The possible library values are:

***LIBL:** The library list is used to find the job description.

***CURLIB:** The job description in the current library is used.

name: Enter the name of the library containing the job description.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Clean System Items (CLNSYSRSF)

The Clean System Items panel is used to delete authorization lists or user profiles on the FROM system that do not exist on the REF system.

Objects that are replicated without journaling are not always deleted from the backup system when the associated object is deleted from the production system. This command provides a way to periodically clean up those orphan objects.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** to display the <u>Work with Sync</u> <u>Attributes</u> panel.
- 2. Enter option **16** and prompt (**F4**) in front of *AUTL or *USRPRF.
- 3. Click the command entry section of the sbmjob and prompt (F4).

Clean Syst	em Items (CL	LNSYSRSF)
Type choices, press Enter.		
Delete orphan objects from: Server ID Compare to: Server ID	*ALL	Name, *LOCAL Name, *LOCAL, *CURRENT Name, generic*, *ALL
Object type	*ALL	*ALL, *AUTL, *USRPRF
Allow delete on production	*N0	*YES, *NO
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	Bottom F13=How to use this display

Options

Delete orphan objects from

Enter the system from which to delete the orphan objects.

The possible values are:

***LOCAL:** The objects are deleted from this system.

server-ID: Enter the Robot HA server ID of the system from which to delete the objects.

Compare to

Enter the name of the system to check. Objects in the system specified by the FROM parameter are deleted if an object with the same name and type does not exist on the system specified by the REF parameter.

The possible values are:

***LOCAL:** The reference system is the same system as the one from which objects will be deleted.

***CURRENT:** The reference system is the system on which the command is run.

server-ID: Enter the Robot HA server ID of the system which contains the objects you want to keep. Objects will be deleted from the other system if they do not exist on the system referred to here.

Object name

Indicate which objects to check.

The possible values are:

***ALL:** All objects are checked.

generic-name: Enter a name or a generic name (ending in an asterisk (*)) for the objects to check.

Object type

Indicate the type of objects to check.

The possible values are:

*ALL: User profiles and authorization lists are checked.

***AUTL:** Authorization lists are checked.

***USRPRF:** User profiles are checked.

Allow delete on production

Indicate whether objects can be deleted from a system that is not set to the *BACKUP replication role.

The replication role is set automatically during a role swap and can be changed manually with the Change RSF Defaults command (CHGRSFDFT).

The possible values are:

***NO:** This command will end in error if attempting to delete objects from the production system.

***YES:** The command can be used to delete objects from the production system. Use caution when selecting this option. One reason to allow the deleting of objects from the production system would be when the server ID associated with the group points to the same system.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F10 (Additional parameters): Shows entry fields for the parameters that are not commonly used.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Display Defined Item Status (DSPRSFDSS)

The Display Defined Item Status panel is used to display the status for some or all of the libraries, replication groups, IFS directories, and system information defined on your Work with Synchronization Attributes panel.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option **1**, Display status for defined items, to view the Display Defined Item Status panel.

Display Define	d Item Status	(DSPRSFDSS)
Type choices, press Enter.		
Server ID	*ALL *NONE *NONE *NONE *NO *NO *	Name Name, *All, *NONE, *generic* Name, *NONE, *generic* Name, *All, *NONE, *generic* Name, *NONE, *generic* *YES, *NO *YES, *NO *, *PRINT, *OUTFILE
F9=All parameters _ F11=Keywords	F14=Command s	Bottom string F24=More keys

Options

For option descriptions, press **F1** to display help text for the Display Defined Item Status panel.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F15 (Error messages): Shows all error messages that pertain to the command being entered.

F16 (Command complete): Indicates that all values needed have been entered, and requests the system to run the command without showing additional displays.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Display Group Sync Status (DSPRSFGSS)

The Display Group Sync Status panel is used to display the status of a replication group.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option **2**, Display group status, to view the Display Group Sync Status panel.

	Display Group	Sync Status	(DSPRSFGSS)	
Type choices, press	Enter.			
Group Output		*	Name *, *PRINT, *OUTFILE	
F9=All parameters	F11=Ke <u>y</u> words	F14=Command	string F24=More keys	Bottom

Options

Library

Enter the name of the group for which to display the status.

The possible values are:

group-name: Enter a valid group name.

Output

Indicate the type of output to generate.

The possible values are:

***:** The information is displayed on the screen.

***PRINT:** The information is printed.

***OUTFILE:** The information is sent to an output file.

File to receive output

If you specify ***OUTFILE** for the **Output** parameter, you must enter a qualified name of the file to receive generated information.

The possible values are:

name: Enter the name of the output file. If the file does not exist, it'll be created.

The possible library values are:

*LIBL: The file is found using the library list. If the file does not exist, *CURLIB is assumed.

***CURLIB:** The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Output member options

Enter the name of the output member and indicate whether to replace or add to existing data. If the member does not exist it will be added.

The possible values are:

***FIRST :** The first member in the file is used.

name: Enter the name of the member to use.

The possible replace values are:

***REPLACE:** Existing data in the member is replaced by the new data.

***ADD:** New data is added to existing data in the member.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F15 (Error messages): Shows all error messages that pertain to the command being entered.

F16 (Command complete): Indicates that all values needed have been entered, and requests the system to run the command without showing additional displays.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Private Authority Status

The Private Authority Status panel is used to display information about private authority status

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option 6, Private Authority Status.



Options

Target Server

Enter the name of a server directory entry. The entry must have been previously added to the server directory on your machine.

The possible values are:

target-server: Enter a valid target server name.

From Library

The name of an existing library synchronization attribute for which to display private authority status.

The possible values are:

*ALL: Display private authority status for all the library synchronization attributes.

name: The name of an existing library synchronization attribute.

***NONE:** No private authority status will be displayed for the library synchronization attributes.

Group

The name of the group synchronization attribute for which to display private authority status.

The possible values are:

*ALL: Display private authority status for all the group synchronization attributes.

Name: The name of a group synchronization attribute.

***NONE:** No private authority status will be displayed for the group synchronization attributes.

Output Queue and Library

The output queue and library to be used for the report. Use F4=Prompt to select from a list. If you enter a library name beforepressing F4=Prompt, only output queues from that library will be listed.

The possible values are:

*: The output queue will be the default queue for the job running this request.

Name: The name of the queue and library where the output will be sent.

Function Keys

- F3 (Exit): Returns to the previous screen.
- F4 (Prompt): Displays a list of possible values from which you may select one.
- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F12 (Cancel): Returns to the previous screen.
- F21 (System Command): Displays a system command line window.

Display IFS Sync Status (DSPRSFISS)

The Display IFS Sync Status panel is used to display information about an IFS synchronization task.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option **3**, Display IFS status, to view the Display IFS Sync Status panel.

Display IFS	Sync Status (DSPRSFISS)
Type choices, press Enter.	
Directory	
Server ID	Name *FROMDIR
Output	<u>*</u> *, *PRINT, *OUTFILE
F3=Exit F4=Prompt F5=Refresh F24=More keys	Bottom F12=Cancel F13=How to use this display

Options

Directory

Enter the name of the directory on the local machine for which to display the status.

The possible values are:

directory-name: Enter a valid directory name of up to 256 characters.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current system are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source system. For an Robot HA pre-processing program running on a server system, *CURRENT refers to the requester system. *CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server system, or a batch job submitted from one of the above two job types.

NOTE: *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. It is allowed within other Robot HA server jobs that connect via TCP/IP.

server-ID: Enter the name of an entry in the server directory on your system.

To directory

Enter the name of the target directory.

The possible values are:

***FROMDIR:** The name of the target directory is the same as the source directory.

directory-name: Enter a valid directory name on the target system. The name may be up to 256 characters.

Output

Indicate the type of output to generate.

The possible values are:

***:** The information is displayed on the screen.

***PRINT:** The information is printed.

***OUTFILE:** The information is sent to an output file.

File to receive output

If you specify ***OUTFILE** for the **Output** parameter, you must enter a qualified name of the file to receive generated information.

The possible values are:

name: Enter the name of the output file. If the file does not exist, it'll be created.

The possible library values are:

*LIBL: The file is found using the library list. If the file does not exist, *CURLIB is assumed.

***CURLIB:** The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Output member options

Enter the name of the output member and indicate whether to replace or add to existing data. If the member does not exist it will be added.

The possible values are:

***FIRST :** The first member in the file is used.

name: Enter the name of the member to use.

The possible replace values are:

***REPLACE:** Existing data in the member is replaced by the new data.

***ADD:** New data is added to existing data in the member.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.
F15 (Error messages): Shows all error messages that pertain to the command being entered.

F16 (Command complete): Indicates that all values needed have been entered, and requests the system to run the command without showing additional displays.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Display Library Sync Status (DSPRSFSS)

The Display Library Sync Status panel is used to display information about a library synchronization task.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option 4, Display library status, to view the Display Library Sync Status panel.

Display Librar	y Sync Status	(DSPRSFSS)
Type choices, press Enter.		
Library	*FROMLIB *DFT *GROUP *	Name Name, *FROMLIB Name Name, *DFT *GROUP, *LIB *, *PRINT, *OUTFILE
F3=Exit F4=Prompt F5= <u>R</u> efresh F24=More keys	F12=Cancel	Bottom F13=How to use this display

Library

Enter the name of the library on the local system for which to display the status.

The possible values are:

library-name: Enter a valid library name.

To library

Enter the name of the target library.

The possible values are:

***FROMLIB:** The name of the target library is the same as the source library.

library-name: Enter a valid library name on the target system.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current system are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source system. For an Robot HA pre-processing program running on a server system, *CURRENT refers to the requester system. *CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server system, or a batch job submitted from one of the above two job types.

NOTE: *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. It is allowed within other Robot HA server jobs that connect via TCP/IP.

server-ID: Enter the name of an entry in the server directory on your system.

Set name

Specify the library set name. The set name is used together with the from library, to library, and server ID to uniquely define a replication entry.

The possible values are:

***DFT:** The default set name is used.

set-name: Enter a valid set name.

Group option

Indicate what to display when the library sync entry is a member of a group.

The possible values are:

***GROUP:** If the entry is a member of a group, display the group status. Otherwise, a value of *LIB is assumed for this parameter.

*LIB: Display status information for the library sync entry.

Output

Indicate the type of output to generate.

The possible values are:

*: The information is displayed on the screen.

***PRINT:** The information is printed.

***OUTFILE:** The information is sent to an output file.

File to receive output

If you specify ***OUTFILE** for the **Output** parameter, you must enter a qualified name of the file to receive generated information.

The possible values are:

name: Enter the name of the output file. If the file does not exist, it'll be created.

The possible library values are:

*LIBL: The file is found using the library list. If the file does not exist, *CURLIB is assumed.

***CURLIB:** The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Output member options

Enter the name of the output member and indicate whether to replace or add to existing data. If the member does not exist it will be added.

The possible values are:

***FIRST :** The first member in the file is used.

name: Enter the name of the member to use.

The possible replace values are:

***REPLACE:** Existing data in the member is replaced by the new data.

*ADD: New data is added to existing data in the member.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F15 (Error messages): Shows all error messages that pertain to the command being entered.

F16 (Command complete): Indicates that all values needed have been entered, and requests the system to run the command without showing additional displays.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Display Synchronization Log

The Display Synchronization Log panel is used to display or print selected synchronization errors. These errors are logged on the source system whenever a problem occurs synchronizing to the target.

How to Get There

From the Robot HA Main Menu, select option **2**, then select option **13**, Synchronization error log.

RHA1012 Dis	play Synchronization Log	14:30:31
Starting date <u>*PRV</u> Ending date <u>*CURREN</u> Starting time <u>*FIRST</u> Ending time <u>*CURREN</u>	_ Date, *FIRST, *CURRENT, *PRV T Date, *CURRENT _ Time, *FIRST T Time, *CURRENT	
Library or Type <u>*ALL</u> To library <u>*ALL</u> To server ID <u>*ALL</u>	<pre>*ALL, name, *generic*, F4=Prompt *ALL, name, *generic*, F4=Prompt *ALL, name, *generic*, F4=Prompt</pre>	
Library set <u>*ALL</u>	<pre>*ALL, name, *generic*</pre>	
Replication group . <u>*ALL</u>	<pre>*ALL, name, *generic*, F4=Prompt</pre>	
Object <mark>*ALL</mark> Object type <u>*ALL</u> Output <u>*</u>	<pre>*ALL, name, *generic*, F4=Prompt Type, *ALL, F4=Prompt *, *PRINT</pre>	
F3=Exit F4=Prompt F5=Refree	sh F10=Advanced options F12=Cancel	

Options

For option descriptions, press **F1** to display help text for the Display Synchronization Log panel.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): If the cursor is in a field which supports **F4 prompt**, displays a list with valid entries for that field.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F10 (Advanced Options): Switches to advanced mode and displays the first page of advanced options. Use the page keys to switch between the various pages.

F12 (Cancel): Exits the panel without processing any pending changes.

Display System Sync Status (DSPRSFSSS)

The Display System Sync Status panel is used to display information about a system synchronization task.

How to Get There

- 1. From the Robot HA Main Menu, select option **4**, then select option **1** to display the Status Reports Menu.
- 2. Select option **5**, Display system status, to view the Display System Sync Status panel.

	Display System	ı Sync Status	(DSPRSFSSS)
Type choices, press	Enter.		
Type of information Server ID Output	· · · · · · · ·	*	*AUTL, *CFG, *NETA Name *, *PRINT, *OUTFILE
			Bottom
F3=Exit F4=Prompt F24=More keys	F5=Refresh	F12=Cancel	F13=How to use this display

Type of information

Specify the type of information for which to display the status.

The possible values are:

***USRPRF:** User profile synchronization information is displayed.

***SYSVAL:** System value synchronization information is displayed.

***NETA:** Network attribute synchronization information is displayed.

***PVTAUT:** Private authority synchronization attributes are set.

***AUTL:** Authorization list synchronization information is displayed.

***CFG:** Configuration synchronization information is displayed.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

NOTE: For SNA/APPC connections, the user profile and password used to connect back to the current system are taken from the server entry that is associated with the current job. If no server entry is associated with the job, the default Robot HA sign on is used. Use the **WRKRSFRQS** command to view or change the server IDs that are associated with requests sent to this system.

The possible values are:

*CURRENT: Contact the server that is currently associated with the job. For pass-through target jobs, *CURRENT refers to the source system. For an Robot HA pre-processing program running on a server system, *CURRENT refers to the requester system. *CURRENT is only valid for this parameter if the job executing the command is an Robot HA target pass-through job, an Robot HA pre-processing program executing on the server system, or a batch job submitted from one of the above two job types.

NOTE: *CURRENT is not valid for Server ID for Robot HA commands run in a target Telnet session. It is allowed within other Robot HA server jobs that connect via TCP/IP.

server-ID: Enter the name of an entry in the server directory on your system.

Output

Indicate the type of output to generate.

The possible values are:

*: The information is displayed on the screen.

***PRINT:** The information is printed.

***OUTFILE:** The information is sent to an output file.

File to receive output

If you specify ***OUTFILE** for the **Output** parameter, you must enter a qualified name of the file to receive generated information.

The possible values are:

name: Enter the name of the output file. If the file does not exist, it'll be created.

The possible library values are:

*LIBL: The file is found using the library list. If the file does not exist, *CURLIB is assumed.

*CURLIB: The file is found in the job's current library.

library-name: Enter the name of the library that contains or will contain the output file.

Output member options

Enter the name of the output member and indicate whether to replace or add to existing data. If the member does not exist it will be added.

The possible values are:

***FIRST :** The first member in the file is used.

name: Enter the name of the member to use.

The possible replace values are:

***REPLACE:** Existing data in the member is replaced by the new data.

***ADD:** New data is added to existing data in the member.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (All parameters): Shows entry fields for all parameters, including those not selected by entries on previous parameters and those not commonly used. It does not show parameters which have been defined with the selective prompt character ?-.

F11 (Keywords): Toggles between the version of the prompt display that shows possible choices and the version that shows parameter keywords.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F14 (Command string): Shows the resulting command as it would run with the parameter values currently entered.

F15 (Error messages): Shows all error messages that pertain to the command being entered.

F16 (Command complete): Indicates that all values needed have been entered, and requests the system to run the command without showing additional displays.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Email Address List

The Email Address List panel displays the email addresses that will get status notifications for the current audit.

How to Get There

1. From the Add or Change Audit Attributes panel, tab down to the **Robot HA email address** field.

2. Press **F4** Prompt. The Email Address List displays. Any email addresses that you have previously added to this audit will display.

HA8450		E	mail Address	List		11:04:47
Enter v	alid Email	Addresses				
sample1	@sample.com					
sample2	@sample.com	1				
						More
3=E×it	F4=Prompt	F5=Refresh	F17=Clear A	ill F21=System	m Command	F22=Expand

Function Keys

- F3 (Exit): Exits the panel without processing any pending changes.
- F4 (Prompt): Displays a list of possible values from which you may select one.

- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F17 (Clear All): Deletes all email addresses.
- F21 (System Command): Displays a command line.
- F22 (Expand): Displays the entire selected email address.

End Synchronization Job (ENDSYNCRSF)

Use this command to end one or more Robot HA synchronization jobs or Remote Apply jobs in a controlled fashion. The job must have previously been started with the Synchronize Libraries (SYNCLIBRSF), Synchronize System Info(SYNCSYSRSF), or Synchronize IFS Directories (SYNCIFSRSF) command.

If the job is currently active, the job will end when processing completes for the current synchronization cycle.

If the job has a state of "SCD" on the job queue, the job is set to run immediately.

A job waiting on a job queue may perform one final synchronization cycle when it becomes active, or simply end normally without a final synchronization cycle depending on the value specified for the OPTION parameter.

If the job specified for this command is not a Robot HA synchronization job, the job is not ended. However, the SCDDATE and SCDTIME for the job may be changed.

How to Get There

Enter the command **ENDSYNCRSF** and press F4 to prompt.

End Synchroni	zat <u>i</u> on Job (E	ENDSYNCRSF)
Type choices, press Enter.		
Job	*IMMED *NONE *LOCAL 0 10	Name, generic*, *ALL, *NONE _ Name 000000-999999 *SYNC, *IMMED Name, generic*, *ALL, *NONE Name, *LOCAL Minutes, *NOMAX Seconds
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	Bottom F13=How to use this display

Job

Enter the qualified name of the job to end. You must specify a job name but the user name and job number are optional.

If you do not specify a user name and job number and only one job exists on the system with the given job name, that job is ended.

NOTE: Note: The machine on which the jobs are ended is indicated by the "Server ID" (SERVER) parameter.

The possible job values are:

***ALL:** All Robot HA synchronization jobs are ended. This includes active jobs and those waiting on the job queue. Active jobs complete the synchronization cycle before ending.

generic: Enter a generic name, ending in an asterisk (*).

Name: Enter a valid job name.

The possible user values are:

Name: Enter a valid user name or leave this field blank to select all jobs that match the specified job name.

The possible number values are:

Name: Enter a valid six-digit job number or leave this field blank to select all jobs that match the job name specified.

End Option

Indicate how the job should end.

NOTE: This parameter effects synchronization jobs and remote journal apply jobs differently.

The possible values are:

***IMMED:** Sync jobs on the job queue are ended without performing a final synchronization. Active sync jobs end when they complete their current synchronization cycle. Remote journal apply jobs end immediately.

***SYNC:** Sync jobs on the job queue are released to perform a final synchronization. Active sync jobs end when they complete their current synchronization cycle. Remote journal apply jobs process all remaining entries in the remote journal, and then end.

NOTE: Note: If changes to production files continue to be made and remote journaling from the source to the target machine remains active, a remote journal apply job ended with the *SYNC option will never end. To end the remote journal apply job in such a case, you must use the *IMMED option.

Remote Journal Apply Job

Enter the name or generic name of the remote journal apply job to end. If the job is active, it will be ended in a controlled fashion.

NOTE: Note: The machine on which the jobs are ended is indicated by the "Server ID" (SERVER) parameter.

The possible values are:

***NONE:** No Robot HA remote journal apply jobs are ended.

***ALL:** All active Robot HA remote journal apply jobs are ended.

Generic-name: Enter a generic name, ending in an asterisk (*). All remote journal apply jobs with matching names are ended.

Name: Enter the name of a specific job to end.

Server ID

Indicate the machine on which to end the jobs.

*LOCAL: Replication and remote journal apply jobs are ended on the local machine.

Server-ID: Enter the server ID of the machine on which to end replication and remote journal apply jobs

Wait

Enter the number of minutes to wait for the specified synchronization jobs to end.

The possible values are:

0: Do not wait. The jobs end asynchronously.

***NOMAX:** Wait as long as it takes for all specified synchronization jobs to end. Control does not return to the program running this command until all of the jobs have ended.

Minutes: Enter the number of minutes to wait for all specified synchronization jobs to end. Control does not return to the program running this command until all of the jobs have ended or the time limit expires.

Delay Between Checks

Enter the number of seconds to wait before checking again to see if all of the specified synchronization jobs have ended.

This parameter is ignored if 0 is specified for the WAIT parameter

The possible values are:

Seconds: Enter a number of seconds, greater than zero.

The WAIT and DELAY parameters work together.

EXAMPLE: If the command is set to WAIT(10) DELAY(5), the command will attempt to end the synchronization job or remote apply jobs. After the first attempt, if the maximum wait of 10 minutes has not been reached, then it will delay 5 seconds and check again. After this next attempt, if the maximum wait of 10 minutes has not been reached, then it will delay 5 seconds and check again. This process will repeat until the maximum wait time has been reached or all the jobs have ended. A message indicating the results will be returned.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F24 (More keys): Shows additional function keys that can be used for this panel.

Initialize Robot HA Setup (INZRSFHA)

The Initialize Robot HA Setup panel is used to perform initial product setup.

For information on how to run this command, see "Initializing Robot HA" in the Robot HA Implementation Guide.

How To Get There

From the Robot HA Main Menu, press **F21** to display a command line and enter the following command:

INZRFHA

NOTE: This command can only be executed on the production system.

Initialize Ro Type choices, press Enter. Check disk space Replace sync start program Replace sync job program Replace monitor program Replace role swap programs Replace job schedule entries Role swap type Journal library	*NO_ *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO *AUTO	(INZRSFHA) *YES, *NO *AUTO, *YES, *NO *AUTO, *YES, *NO *AUTO, *YES, *NO *AUTO, *YES, *NO *AUTO, *YES, *NO *NOSWAPIP, *SWAPIP Name, JRNLIB, *SAME Name, RMTJRNLIB, *SAME
Sync new directories	<u>*NO</u> 'Robot HA'	*YES, *NO, *SAME
		More
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	F13=How to use this display

For option descriptions, press **F1** to display help text for the Initialize Robot HA Setup panel.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (How to use this display): Shows help for the prompt display or associated display you are currently using.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Purge History

The Purge History panel is used to purge synchronization logs. You can choose to run the purge report immediately or schedule it to run using Robot Schedule.

How to Get There

From the Robot HA Main Menu, select option **5**, then select option **3**, Purge synchronization log.

RHA1354	Purge History	14:36:28
Options 1=Submit Delete 2=Schedule	on Robot SCHEDULE	
Opt Report to Run		
_ Purge synchronization log	Days of data to retain: Purge entries for: Set name <u>*ALL</u> Replication group <u>*ALL</u>	31
F3=Exit		

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

Sync Authorities

The Sync Authorities (RHASYNCAUT) panel is used to run the private authority update when the sync job has ended. The sync job can be restarted.

How to Get There

From a command line, enter the command RHASYNCAUT and press F4 to prompt.

Sync Auth	norities (RHAS	SYNCAUT)
Type choices, press Enter.		
Server ID	*LIBRARY *FROMLIB *DFT *YES	Name, F4=Prompt *ALL, *LIBRARY, *GROUP *ALL, Name, F4=Prompt Name, *FROMLIB Name, *DFT *ALL, Name, F4=Prompt *YES, *NO
F3=Exit F4=Prompt F5=Refresh F24=More keys	F12=Cancel	Bottom F13=How to use this display

The options vary slightly depending on whether *LIBRARY or *GROUP is selected for **Library** or **Group** parameter.

Server ID

Enter the name of a server directory entry to be referenced. The entry must have been previously added to the server directory on your system.

The possible values are:

Server-ID: Enter the name of an entry in the server directory on your system

Press **F4=Prompt** to select from a list of available values.

Library or Group

Enter the attribute type to process

The possible values are:

*ALL: Process all library and group synchronization attributes.

*LIBRARY: Process library synchronization attribute(s).

*GROUP: Process group synchronization attribute(s).

From Library

Enter the name of an existing library synchronization attribute.

The possible values are:

name: Enter the name of an existing library attribute.

***ALL:** Process all library sync attributes.

NOTE: The **Sync object authorities** parameter on the library attribute must be set to *YES for authorities to be processed.

To Library

Enter the name of the target library for synchronization. The attributes for a library synchronization definition are for a given from-library, to-library, server-ID, and set-name combination.

The possible values are:

***FROMLIB:** The name of the target library is the same as the source library.

Library-name: Enter a valid library name on the target system.

Set Name

Specify the library set name for synchronization. The set name is used together with the from-library, to-library, and server-ID to determine if the library is defined.

The possible values are:

***DFT:** The default set name is used.

Set-name: Enter a valid set name for synchronization.

Group

Enter the name of an existing group synchronization attribute.

The possible values are:

name: Enter the name of an existing group synchronization attribute.

NOTE: The **Sync object authorities** parameter on the library attribute within the group must be set to *YES for authorities to be processed.

Restart Sync Job

Indicate whether to restart synchronization after processing object authorities.

The possible values are:

***NO:** The synchronization job(s) will not be restarted.

***YES:** The synchronization job(s) will be restarted.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F12 (Cancel): Exit the current panel without processing any pending changes.

Page Up/Page Down: Displays the previous or next parameters.

Select Email Address

The Select Email Address panel displays the email addresses that have been added to previously defined audits and are available for selection in the current audit.

How to Get There

1. From the Add or Change Audit Attributes panel, tab down to the **Robot HA email address** field.

2. Press **F4** Prompt. The Email Address List displays. Any email addresses that you have previously added to this audit will display.

3. Press **F4** Prompt again. The Select Email Addresses panel displays. Any email addresses that you have added to previously defined audits will display

RHA8000	Select Email Addresses	11:23:59
1=Multi-Select Select Email Addresses <u>bob@bob.com</u>	Start:	
F3=Exit F15=Select All	F17=Clear All F22=Expand	Bottom

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F15 (Select All): Inserts a '1' in front of all listed email addresses.

F17 (Clear All): Deletes all email addresses.

F22 (Expand): Displays the entire selected email address.

Synchronization Attributes Report Panel

The Synchronization Attributes Report panel is used to set up a report that will display the synchronization attributes for some or all of the libraries, groups, IFS directories and system information defined. A report or an output file is generated.

How to Get There

- 1. From the Robot HA Main Menu, select option **4** to display the Reports Menu.
- 2. Select option **4** Synchronization Attributes Report

RHA3404 Synchron	ization Attributes Report	12:30:22 MOHA01
Target server From library *NONE Group *NONE System type Directory *NONE	Name, F4=Prompt *ALL, Name, *NONE, F4=Prompt *ALL, Name, *NONE, F4=Prompt *ALL, Name, *NONE, F4=Prompt *ALL, Name, *NONE, F4=Prompt	
Output queue <u>*</u> Library ASP Group *SYSBA	*, Name, F4=Prompt Name, F4=Prompt S	
F3=Exit F4=Prompt F5=Refresh	F12=Cancel F21=System Command	

Target Server

Enter the name of a server directory to be referenced. The entry must have been previously added to the server directory on your machine.

NOTE: This is a required parameter.

The possible values are:

Name

Enter the name of the target server

From Library

Specify the name of the from library to be included in the synchronization attributes report.

NOTE: Libraries that are part of a group are not included with this parameter. See the Group parameter.

The possible values are:

*ALL

All Libraries for the given server that are defined in Robot HA will be included in the synchronization attributes report.

Name

Specify the name of an existing library to be included in the synchronization attributes report.

*NONE

No libraries will be selected for the synchronization attributes report.

Group

Specify the name of the group to be included in the synchronization attributes report.

NOTE: Libraries that are part of a group will be printed with the specified group.

The possible values are:

*ALL

All Groups for the given server that are defined in Robot HA will be included in the synchronization attributes report.

Name

Specify the name of an existing group to be included in the synchronization attributes report

*NONE

No groups will be selected for the synchronization attributes report.

System Type

Specify the name of the system type to be included in the synchronization attributes report.

The possible values are:

*ALL

All system types for the given server that are defined in Robot HA will be included in the synchronization attributes report.

Name

Specify the name of an existing system type to be included in the synchronization attributes report.

*NONE

No system types will be selected for the synchronization.

Directory

Specify the name of an IFS directory to be included in the synchronization attributes report.

The possible values are:

*ALL

All IFS directories for the given server that are defined in Robot HA will be included in the synchronization attributes report.

Name

Specify the name of an existing IFS directory to be included in the synchronization attributes report.

*NONE

No IFS directories will be selected for the synchronization attributes report.

Output Queue and Library

This is the output queue and library to be used for the report. Use **F4** prompt to select from a list. If you enter a library name before pressing **F4** prompt, only output queues from that library will be listed.

The possible values are:

*

The output queue will be the default queue for the job running this request.

Name

The output will be sent to this named queue and library.

Function Keys

- F3 (Exit): Exits the panel without processing any pending changes.
- F4 (Prompt): Displays a list of possible values from which you may select one.
- F5 (Refresh): Refreshes the panel and resets all available text fields.
- F12 (Cancel): Exits the panel without processing any pending changes.
- F21 (System Command): Displays a system command line window.

Start Synchronizing

The Start Synchronizing panel is used to synchronize your Robot HA servers.

How to Get There

From the Robot HA Main Menu, select option **2**, then select option **3**, Start synchronizing.

RHA1023 Start synchronizing	15:12:31
Enter the RSF Server to start synchronizing, then press enter.	
RSF Server Name, F4=Prompt	
F3=Exit F4=Prompt	

Options

RSF Server

Enter the required server ID and press Enter. If you do not know the server ID, press **F4** to prompt and select one from a list of possible values.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

Work with Active Jobs

The Work with Active Jobs panel is used to show the performance and status information for jobs that are currently active on the system. All information is gathered on a job basis.

How to Get There

From the Robot HA Main Menu, select option **2**, then select option **11**, Synchronization jobs.

```
Work with Active Jobs
                                                            09/13/16 14:38:12
CPU %:
           1.2
                   Elapsed time:
                                   00:57:43
                                                Active jobs:
                                                               182
Type options, press Enter.
  2=Change
           3=Hold 4=End
                              5=Work with
                                            6=Release
                                                        7=Display message
  8=Work with spooled files
                             13=Disconnect ...
                   Current
                                Type CPU % Function
Opt
    Subsystem/Job User
                                                             Status
                                                              DEOW
    RSFHA
                   OSYS
                                SBS
                                         . 0
                                                                        Bottom
Parameters or command
===>
         F5=Refresh
                           F7=Find
                                        F10=Restart statistics
F3=Exit
F11=Display elapsed data
                           F12=Cancel
                                        F23=More options
                                                           F24=More keys
```

Options

For option descriptions, press **F1** to display help text for the Work with Active Jobs panel.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F7 (Find): Allows the user to search for information on the display.

F8 (Repeat find): Repeats the previous find operation. If no previous find was performed, you are prompted for the find information.

F9 (Retrieve): Shows the last command you typed on the command line, along with any parameters you included. You can press F9 multiple times to cycle through your previously executed/prompted commands.

F10: Restart statistics: Shows the list again with the start time set to the previous display time. The measurement time interval is the time between when the previous display was presented and the time F10 was pressed.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (Reset statistics): Shows the list again, but with the start time set to the current time. The measurement time interval is zero.

F23 (More options): Shows additional list options on the panel. The options can be used even if they are not currently shown.

F24 (More Keys): Shows additional function keys that can be used for this panel.

Work with All Spooled Files

The Work with All Spooled Files panel is used to show all (or a specified portion) of the spooled files that are currently on the system.

How to Get There

From the Robot HA Main Menu, select option **2**, then select option **14**, Synchronization job logs.

	Work with All Spooled Files								
Type 1= 8=	Type options, press Enter. 1=Send 2=Change 3=Hold 4=Delete 5=Display 6=Release 7=Messages 8=Attributes 9=Work with printing status								
Opt _ _ _ _	File QPJOBLOG QPJOBLOG QPJOBLOG QPJOBLOG	User RBTUSER RBTUSER RBTUSER RBTUSER	Device or Queue QEZJOBLOG QEZJOBLOG QEZJOBLOG QEZJOBLOG	User Data S_RSFUSER S_CFG S_RBTHALIB S_CFG	Sts RDY RDY RDY RDY	Total Pages 3 1 4 9	Cur Page	Copy 1 1 1	
Bottom Parameters for options 1, 2, 3 or command ===>									
F3=E	xit F10=V	iew 4 F11=\	/iew 2 F12=	Cancel F22	=Print	ters F	24=More	keys	

You can type an option number next to one or more files. When you then press the Enter key, the function associated with the selected option is performed for each of the selected files.

If you press the Enter key from this display without typing any other data, you are returned to the previous menu or display.

You can type parameters in the command line if you want to override the defaults for the options you typed. You can type parameters only if you used options 1 (Send), 2 (Change), or 3 (Hold).

1=Send: Sends the spooled file. When you choose this option, you cause the Send Network Spooled File command to process.

2=Change: Runs the Change Spooled File Attributes command to change the attributes of the spooled file.

3=Hold: Holds the spooled file. When you choose this option, you cause the Hold Spooled File command to process.

4=Delete: Deletes the specified spool file. When you choose this option, the Confirm Delete of Spooled Files panel displays.

5=Display: Displays the data in the spooled file. When you choose this option, you cause the Display Spooled File command to process.

6=Release: Releases the specified spool file from a hold status. When you choose this option, you cause the Release Spooled File command to process.

7=Messages: Displays messages related to the printing status of a spooled file.

8=Attributes: Displays the attributes of the specified spooled file. When you choose this option, you cause the Work with Spooled File Attributes command to process.

9=Work with printing status: Displays the printing status of the specified spool file. When you choose this option, you cause the Work with Printing Status command to process.

Function Keys

F3 (Exit): Exit the current panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F10 (View 4): Switches to the fourth view of the display.

F11 (View 2): Switches to the second view of the display.

F12 (Cancel): Exit the current panel without processing any pending changes.

F22 (Printers): Displays the Work with All Printers panel.

F24 (More keys): Shows additional function keys that can be used for this panel.

Work with User Jobs

The Work with User Jobs panel is used to show the performance and status information for RSFSRV jobs. All information is gathered on a job basis.

How to Get There

From the Robot HA Main Menu, select option **2**, then select option **12**, Synchronization jobs.

ſ		Work wit	th User Jobs	
Type options 2=Change 8=Work wit	3, press Enter. 3=Hold 4=Er th spooled fild	nd 5=Work es 13=Disc	with 6=Release connect	08/16/17 14:22:50 CUT 7=Display message
0pt Job RHAAUDH RSFSRV1	User 10N RSFSRV ICP_ RSFSRV	Туре ВАТСН ВАТСН	Status ACTIVE ACTIVE	Function PGM-RHA2133 PGM-RSCL115
Parameters (===>	or command			Bottom
F3=Exit F12=Cancel	F4=Prompt F F17=Top F	⁻ 5=Refresh F18=Bottom	F9=Retrieve F11 F21=Select assist	EDisplay schedule data ance level

For option descriptions, press **F1** to display help text for the Work with User Jobs panel.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F9 (Retrieve): Allows the user to search for information on the display.

F8 (Repeat find): Repeats the previous find operation. If no previous find was performed, you are prompted for the find information.

F9 (Retrieve): Shows the last command you typed on the command line, along with any parameters you included. You can press F9 multiple times to cycle through your previously executed/prompted commands.

F11: Display schedule data: Shows the list again with the dates and times for scheduled jobs.

F12 (Cancel): Exits the panel without processing any pending changes.

F17 (Top): Positions the display to the top of the list.

F18 (Bottom): Positions the display to the bottom of the list.

F21 (Select assistance level): Allows you to select the level of assistance you want when interacting with the system.

Work with RSF Servers

The Work with RSF Servers panel is used to view, add, change, copy, rename, and remove Robot HA server directory entries on your system.

How to Get There

From the Robot HA Main Menu, select option **3**, then select option **1**, Server definitions.

Work With RSF Servers						
		Position to				
Type o 2=Ch 7=Re	options, pres nange (name (ss Enter. 3=Copy 4=Delete 5=Display 12=Work with RSF catalog 15=Start pass-through _				
Opt 	Opt Server Text BACKUP Backup machine LOCAL For loop back SDLC testing LOOPBACK For loop back TCP/IP testing PIGGYBACK Piggyback on incoming call on line RSFSRV PROD Production machine TEST My SDLC connection to TEST					
F3=Exi	t F4=Prompt	t F5=Refresh F6=Create F12=Cancel F13=Defaults F HelpSystems (C) Cop	Bottom 21=Cmd yright			

Options

Server

Indicate which entries to include in the list.

The possible values are:

- *ALL: No entries are excluded based on name.
- **generic-name:** Enter a generic name for the servers to be included in the list. Case is significant. An asterisk (*) in the generic specification will match any string of zero or more characters in the name. An underscore (_) in the generic specification will match any single character in the name.

Text Compare String

Enter characters to compare to entry text to determine which entries should be included in the list. An entry is included in the list if the entry text contains the string specified. Case is not significant.

The possible values are:

- ***ALL:** No entries are excluded based on text.
- String: Enter any text string of up to 50 characters.

Output

Indicates whether the list should be displayed or printed.

The possible values are:

- *: Display the list.
- ***PRINT:** Print the list.

Detail

Controls the amount of detail shown when printing the list.

The possible values are:

- ***BASIC:** A subset of the total information is shown for each directory entry. Multiple entries are printed per page.
- ***FULL:** All information is shown for each directory entry. One entry is printed per page.

Position to

Enter a value in the **Position to** field and press Enter to position the list to a specific entry. The cursor is positioned to the first entry in the list that is greater than or equal to the value you specify.

Opt

Enter an option number in the "Opt" column beside a list entry, and press **Enter** to perform a function on the list entry. You may enter options beside several list entries before pressing **Enter**. The options for the list entries are processed in turn when you press **Enter**. The following is a list of options and their functions.

- **2=Change:** The CHGRSFSDE command prompt is displayed with the current values for the entry filled in.
- **3=Copy:** A display is presented with which you specify the new names for entries to be copied.
- **4=Delete:** A display is presented allowing you to confirm your choices for delete. When you press Enter a second time, the entries are deleted.
- **5=Display:** Detailed information about the entry is displayed.
- **7=Rename:** A display is presented with which you specify new names for the entries to be renamed.
- **12=Work with RSF catalog:** The WRKRSFCAT display is presented for the selected server. The WRKRSFCAT display allows you to view a list of Robot HA packages available from the server and easily retrieve any package.
- **15=Start pass-through:** A pass-through session is started with the server machine. The server must grant your machine permission to pass-through, by adding an entry to the requester directory on their machine.
- **user-defined-option:** You may key a user-defined option. To define new options, select option **9** from the Start PDM (STRPDM) menu.

See the online help text for this display for more information about user-defined options.

Server

The unique server ID is shown in this column.

Text

Text describing the server is shown. Type over the text and press Enter to change the text.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F6 (Create): Adds a new server directory entry. The Add Server Directory Entry panel (**ADDRSFSDE**) is displayed.

F12 (Cancel): Exits the panel without processing any pending changes.

F13 (Defaults): Changes the user defaults.

F21 (System Command): Presents a system command line window.

Work with Synchronization Attributes

The Work with Synchronization Attributes panel is used to view, add, change, and delete synchronization attribute definitions on your system. You can also display error logs and launch synchronization jobs.

How to Get There

- 1. From the Robot HA Main Menu, select option **1** Synchronization Attributes.
- 2. Work with Synchronization Attributes panel displays.

RHA300501			Work With	Synchronizati	on Attributes - Librari	es	15:49:10
Position to L	ibrary	:			Sort by:	FR (F4) From Library	*PROD
Type options,	press Ent	er. (F4)				Conditions exist	
2=Chg Attr 4	=Delete 5	=Display 6=Sta	tus 7=List	Obj 8=Sync	9=Sync Auth 10=DSPJRN	12=WRKJRNH 14=Work with	15=Dsp Err Log
Upt From	Server	10	Set Name	10 Sync/Hpply	Local Journal	Last SyncLast	Error/Type CDN
H	BHCKUP	H		*_H	JRNLIB/H		
HBCNU	BHCKUP	HBCNU		*_HBCNU	JRNLIB/HBCNU		
HBCYES	BHCKUP	HBCYES		*_HBCYES	JRNLIB/HBCYES		
ABC1	BACKUP	ABC1		*_ABC1	JRNLIB/ABC1 _		
HBC2	BACKUP	ABC2		I *_ABC2	JRNLIB/ABC2		
HBC2	BHCKUP	HBC2	OTHER	I *_HBC2H	JRNLIB/ABC20THER		
B1	BHCKUP	B1		*_B1	JRNLIB/B1	03/15/21 13:46	L
B1	BHCKUP2	B1		*_B1H	JRNLIB/B1	03/15/21 13:55	L
^{B2}	BACKUP	B2		*_B2	JRNLIB/DLTGRP		
_ ^{CCC}	BACKUP	CCC		*_CCC	JRNLIB/CCC		
DEMONOAUT	BACKUP	DEMONOAUT		*_DEMONOAU	JRNLIB/DEMONOAUT		
DEMOYESAUT	BACKUP	DEMOYESAUT		0 *_DEMOYESA	JRNLIB/DEMOYESAUT	03/15/21 13:54	L
DIFF	BACKUP	DIFFLIB		*_DIFF	JRNLIB/DIFF		
DTAARA	BACKUP	DTAARA		*_DTAARA	JRNLIB/DTAARA		
FIXNOAFT	BACKUP	FIXNOAFT		*_FIXNOAFT	JRNLIB/FIXNOAFT		
FIXNOBOTH	BACKUP	FIXNOBOTH		*_FIXNOBOT	JRNLIB/FIXNOBOTH		
							More
F3=Exit F4=P	rompt F5=	Refresh F6=Crea	ate F9=Job	s F11=Toggle	F13=Defaults F17=Subs	et F21=System Command F	23=More Options

Position to Library

Enter the library name you want to position to in the list and press enter. The panel will redisplay with the matching library at the top of the page. If no matching entry exists the top line will be the nearest match possible.

The possible values are:

Name: Enter the library name, or beginning character(s) of the name.

Sort by

Select the sort sequence you would like the list displayed in.

F4 prompt for a list of sort options.

The possible values are:

FR: From Library

TO: To Library

SV: Server

LS: Last Sync Date

LE: Last Error Date

Columns

From

This is the library, IFS directory, group, or system attribute that is being synchronized.

Server

This is the server that the attribute is being replicated to.

То

This is the library, IFS directory, group, or system attribute that is being synchronized to.

Set Name

If a library is separated into sets, the set name will be listed here.

10

If a sync attribute contains includes or omits, an I or O will be listed here.

Sync/Apply

This value is used to associate this sync attribute with jobs that are produced for it. For instance, $*_$ will be replaced with S_ for the sync job on the local system. $*_$ will be replaced with J_ for the remote apply job on the remote system. C_ is used for the Compare Job, and D_ is used for Clean.

Local Journal

This is the name of the local journal being used for the attribute.

Last Sync

The last date/time the item was synchronized is shown. If this is highlighted, the last sync point was later than expected.

Last Error

The last date/time an error was logged for the item is shown. If this is highlighted, the last error occurred after the last sync point.

Last Error Type

This value indicates the type of error that exists.

The possible values are:

S: Sync error exists.

A: Authority error exists.

CDN

This gives an indication of which condition(s) are responsible for the sync date displayed.

The possible values are:

L: Late - If any items are late reaching a sync check point. The time since the last sync point is compared to the default sync interval defined in the entry. If the time since the last sync
point is greater than it should be, the condition is flagged. The value specified for threshold is interpreted as a percent. If a threshold of 20 is specified, an item with a sync check point that is late by 20% or more causes the condition to be flagged. For example, when (*LATE 20) is specified, an item with a default sync interval of 10 minutes is considered late if it has not had a sync check point within the last 12 minutes.

A: Apply Lag - If, for any item, the count of journal changes generated on the source system and not yet applied on the target is greater than the threshold specified, the condition is flagged.

P: *PVTAUT Lag - If the number of authority entries not yet processed for the server is greater than the threshold specified, the condition is flagged

T: Transfer Lag - If, for any item, the backlog of journal changes to be sent from the source to the target is greater than the threshold specified, the condition is flagged.

S: Spool Lag - If, for any item, the number of QAUDJRN entries not yet processed for spooled file changes is greater than the threshold specified, the condition is flagged.

Q: QAUDJRN Lag - If, for any item, the number of QAUDJRN entries not yet processed is greater than the threshold specified, the condition is flagged.

O: Object Authority Lag - If, for any item, the number of object authority entries not yet processed is greater than the threshold specified, the condition is flagged.

Options

Key an option number in the "Opt" column beside a list entry and press Enter to perform a function for the entry. You can key options beside several list entries before pressing Enter.

Press F4 instead of Enter to prompt for additional parameters.

The possible values are:

2 - Chg Attr

The Change Synchronization Attributes panel for the current type of Sync Attribute displays with the current values filled in.

Change Library Sync Attribute

Change IFS Sync Attribute

Change Group Sync Attribute

Change System Sync Attribute

4 - Delete

The <u>Delete Synchronization Attributes</u> panel displays with the current values filled in.

5 - Display

The <u>Display Synchronization Attribute</u> displays for the selected attribute.

6 - Status

Displays detailed status for the selected attribute.

7 - List Objects

Lists objects specified by this attribute's synchronization specifications. A physical file or a spool file can be created that shows the objects which will be synchronized when Option 8 is run. This can be useful if special values *JRN or *OTHER are used in the attribute specifications.

NOTE: This option will not list IFS objects.

8 - Synchronize

Begins synchronization for the specified item. If a job for the entry is already active or on the job queue, option 8 is treated as if user-defined option **RN** were entered. Otherwise, a new job for the item is submitted.

9 - Sync Auth

Brings up the Confirm Sync Authority of Synchronization Attributes window. Press **F8** to synchronize the authorities for the selected attributes, or press **F12** to cancel.

NOTE: This will call the <u>RHASYNCAUT</u> command. It will end any active sync jobs and remote journal apply jobs for the attribute before running the authority synchronize job. In the confirmation window you have the option to restart the sync jobs after.

10 - DSPJRN

If a journal is associated with an entry, the contents of the journal display. You can press **F4** to access additional parameters that allow you to filter the journal entries. Any filtering specified remains in effect for this sync entry for the current job. Press **F12** while prompting to remove all filtering and revert to the initial DSPJRN display.

12 - WRKJRNA

If a journal is associated with an entry, the journal attributes display.

14 - Work With

Works with the objects in the library or directory associated with the entry.

15 - Display Log

Error log information for the entry displays.

16 - Compare

Check the source against the target library or directory to see if they match.

18 - Check Journal

Checks the journal associated with the item. The oldest journal receiver still in use by Robot HA is determined and the entry using the receiver is shown.

20 - Clean

Deletes orphan objects from the target library or directory.

HA - Hold Attribute

This option places the attribute on hold and should be used with caution. When an attribute is held:

- The attribute will not process
- A swap will not run
- We will not flag that attribute as not being processed on the conditions report
- The swap audit will not start syncing

RA - Release Attribute

Releases the hold on an attribute. RA should not be used to release a hold that was placed using the BT option as it can interfere with the process. ET will release the hold

BT - Begin Tape

Begins the tape process for manually replicating a library or IFS directory. See <u>Using Tape to</u> <u>Manually Replicate a Library or IFS directory in Robot HA</u> for more information on BT.

ET - End Tape

Ends the tape process for manually replicating a library or IFS directory. See <u>Using Tape to</u> <u>Manually Replicate a Library or IFS directory in Robot HA</u> for more information on ET.

For more information on the Work with Synchronization Attributes panel, press **F1** to view the help text.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Refreshes the panel and resets all available text fields.

F6 (Create): Displays the Add Attribute screen for the attribute type currently being displayed.

F9 (Jobs): Displays the current synchronization jobs using the program specified in the user defaults.

F11 (Toggle): Cycles the display through each attribute type. Attribute types *IFS, *GROUP and *SYSTEM may be omitted from the cycle by setting the appropriate user default or subset value.

F13 (Defaults): Change user defaults. Some default values for other options are taken from the values you specify for your user defaults.

F17 (Subset): Presents a window which allows certain values to be altered to change the content of the display. The values remain in effect until changed by another F17=Subset, or by leaving the screen entirely. See <u>Subset Window</u> for more information.

F21 (Cmd entry): Displays a command line.

F23 (More options): Shows additional options that can be used for this panel.

Display Synchronization Attributes

This screen shows the values you have setup for the sync attribute in display only mode. Journal and additional setup data are separated for ease in viewing.

How to Get There

From the <u>Work with Sync Attributes</u> panel, select option **5** in front of one or more sync attributes you would like to display.

1A3009		Display	Synchronization Attribute - Library	16:04:
°om: A	To: A	Server: BACKUP	Set Name: *DFT	*PROD
efault Job Nar epeat Every.	me:S_A :10 *MINUTES		Synchronization Date/Time : 03/10/21 08:51:58 Restrict to Window :	
Local jour Last ree Last see Remote jou HA manages Change ree Days to ke End jrn fo Fix obj us Filter jou Apply jour Journal ir Create sta Clear star	Journal So rnal	etup : JRNLIB/A : AA0106 ed: 30 : RMTJRNLIB/A : *YES : *DAY : *SYNC : *YES : *YES : *NO : *RRN : *BOTH : *YES : *YES : *YES		
}=Exit F7=Ind	cludes F8=Omits F12: ot allowed.	=Cancel		More.

Function Keys

- F3 (Exit): Returns to the Work With Sync Attributes screen.
- F7 (Includes): Displays the Includes for the attribute.
- F8 (Omits): Displays the Omits for the attribute.

F12 (Cancel): Either continues on to display the next Option 5=Display selected, or return to the Work With Sync Attributes screen if no more have been selected to display.

Confirm Delete of Synchronization Attributes

The confirm Delete of Synchronization Attributes window displays the attributes you have selected to delete using option 4=Delete. You can choose here whether to delete related journals and whether to ignore individual object errors that may occur during the deletion of each attribute.

How to Get There

From the <u>Work with Sync Attributes</u> panel, select option **4** in front of one or more sync attributes.

RHA300505	Confirm	Delete of Syn	chronization	Attributes	11:29:01
Press Ente	r to confirm y	your choices f	or 4=Delete.		*PROD
Opt From 4 ARCHIVE 4 ARCHIVE 4 ARCHIVE 4 ARCHIVE	to return to o Server MYBACKUP MYBACKUP MYBACKUP	change your ch To ARCHIVE ARCHIVE ARCHIVE	Set SET2017 SET2018 SET2019	Dlt Rlt Jrn (F4) <u>*ALL</u> *ALL *ALL	Ignore Rlt Obj Err (F4) <u>*NO</u> <u>*NO</u> <u>*NO</u>
F4=Prompt	F12=Cancel				Bottom

Options

Delete Related Journals

Indicate whether related replication journals should be deleted before the synchronization entry is deleted.

The possible values are:

*ALL: All related local, remote and standby journals are deleted.

***REMOTE:** Only related remote and standby journals are deleted. The local journal is not deleted.

***NONE:** No journals are deleted.

Ignore Related Object Errors

Specify whether to delete the sync entry even if errors occur ending related sync jobs or deleting related journals.

The possible values are:

***NO:** The sync entry is not deleted if errors occur deleting related objects.

***YES:** The sync entry is deleted, even if errors occur deleting related objects.

Function Keys

F4 (Prompt): Displays a list of possible values to select.

F12 (Cancel): Cancels the delete.

Subset Window

Use this window to select a subset of the sync attributes you would like to display. Each time this window is displayed, you will see the selections that were used to display the current list. The very first time this screen is displayed, the default selections will be shown.

How to Get There

From the Robot HA Main Menu, select option **1** Synchronization Attributes, then press function **F17** Subset

RHA3006	Subset and Addit	ional Display Options
Include servers	*ALL	Name, *ALL (F4)
Libraries to include	*DEFINED	*DEFINED, *UNDEFINED
Include group attributes .	*YES	*YES, *NO
Include IFS attributes	¥YES	*YES, *NO
Include system attributes.	*YES	*YES, *NO
Include errors	<u>*All</u>	*ALL, *ERROR, *NOERROR
Limit to	<u>*Sync</u>	Date, *ALL, *DAY, *WEEK, *MONTH, *SYNC
Display sync dates Threshold Conditions:	<u>*Day</u>	Date, *ALL, *DAY, *WEEK, *MONTH
Late	*CDN	Number, *CDN, *NONE
Apply Lag	*NONE	Number, *CDN, *NONE
Transfer Lag	*NONE	Number, *CDN, *NONE
QAUDJRN Lag	*NONE	Number, *CDN, *NONE
*PVTAUT Lag	*NONE	Number, *CDN, *NONE
Spool Lag	<u>*NONE</u>	Number, *CDN, *NONE
Object Authority Lag	*NONE	Number, *CDN, *NONE

Options

Include servers

Select the server(s) you would like included on the sync attributes display.

The possible values are:

Name: Enter the name of a specific server you want to see the sync attributes for. F4 is available on this selection.

***ALL:** Entries for all servers will be included on the sync attributes display. This is the default selection.

Libraries to include

Select the libraries you would like included on the sync attributes display.

The possible values are:

***DEFINED:** Only libraries with synchronization attributes defined are included. This is the default selection.

***UNDEFINED:** Only libraries that do not have synchronization attributes defined are included. When this option is selected, all other parameters related to libraries are moot. For example, server subset, conditions, and others related to libraries would not apply.

Include group attributes

Specify whether or not you would like to see group sync attributes included on the sync attributes display.

The possible values are:

***YES:** Include group sync attributes on the sync attributes display. This is the default selection.

***NO:** Do not include group sync attributes on the sync attributes display.

Include IFS attributes

Specify whether or not you would like to see IFS sync attributes included on the sync attributes display.

The possible values are:

***YES:** Include IFS sync attributes on the sync attributes display. This is the default selection.

***NO:** Do not include IFS sync attributes on the sync attributes display.

Include system attributes

Specify whether or not you would like to see system sync attributes included on the sync attributes display.

The possible values are:

***YES:** Include system sync attributes on the sync attributes display. This is the default selection.

***NO:** Do not include system sync attributes on the sync attributes display.

Include Errors

Specify the errors you would like included on the sync attributes display.

The possible values are:

***ALL:** All sync attributes, whether or not they are currently in error. This is the default selection.

***ERROR:** Only sync attributes that are currently in error and meet the subset criteria from "Limit to" are included on the display

***NOERROR:** Only sync attributes that are not currently in error and meet the subset criteria from "Limit to" are included on the display.

NOTE: A sync attribute is considered currently in error if the last error date and time are more current than the last sync date and time.

Limit to

Specify the date for the earliest errors to show on the sync attributes display. Items with no errors before this date will show blanks for their error date if *ALL or *NOERROR was selected for "Include Errors".

The possible values are:

Date: Enter a date in your system date format (WRKSYSVAL QDATE). This is the earliest date errors will show on the sync attributes display.

***ALL:** All sync attributes that have any errors will show their error date. Items with no errors will show blanks for their error date if *ALL or *NOERROR was selected for "Include Errors".

***DAY:** The earliest errors shown are one day old.

***WEEK:** The earliest errors shown are one week old.

***MONTH:** The earliest errors shown are one month old.

***SYNC:** Only errors that occurred after the last synchronization point are shown. This is the default selection.

Display sync dates

Specify the earliest sync date to show in the list. Items last synchronized before this date will show blanks for their sync date.

The possible values are:

Date: Enter a date in your system date format (WRKSYSVAL QDATE). This is the earliest sync dates will show on the sync attributes display.

***ALL:** All sync attributes that have any sync dates will show their sync date. Items with no sync dates will show blanks for their sync date.

***DAY:** The earliest sync dates shown are one day old. This is the default selection.

***WEEK:** The earliest sync dates shown are one week old.

***MONTH:** The earliest sync dates shown are one month old.

Threshold Conditions

Late

Items that are late reaching a sync check point are flagged by an (L) on the Work With Synchronization Attributes panel under the CDN column. The time since the last sync point is compared to the default sync interval defined in the entry. If the time since the last sync point is greater than it should be, the entry is flagged.

For *LATE, the value specified for Threshold is interpreted as a percent. If a threshold of 10 is specified, items with sync check points that are late by 10% or more are flagged.

For example, when (*LATE 10) is specified, an item with a default sync interval of 10 minutes would be flagged if it had not had a sync check point within the last 11 minutes.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the condition specified is used. The default value for threshold late is: 10

***NONE:** Do not flag sync attributes that are late.

Apply Lag

An item is flagged with (A) on the Work With Synchronization Attributes panel under the CDN column if the count of journal changes generated on the source machine and not yet applied on the target is greater than the threshold specified.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the condition specified is used. The default value for threshold apply lag is 1,000.

***NONE:** Do not flag sync attributes with apply lag.

Transfer Lag

An item is flagged with (T) on the Work With Synchronization Attributes panel under the CDN column if the backlog of journal changes to be sent from the source machine to the target machine is greater than the threshold specified.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the condition specified is used. The default value for threshold transfer lag is 1,000

***NONE:** Do not flag sync attributes with transfer lag.

QAUDJRN Lag

An item is flagged with (Q) on the Work With Synchronization Attributes panel under the CDN column if the number of QAUDJRN entries not yet processed for *USRPRF, *AUTL, (system values other than *PVTAUT, spooled files) is greater than the threshold specified.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the Condition specified is used. The default value for QAUDJRN lag is 10,000.

*NONE: Do not flag sync attributes with QAUDJRN lag.

*PVTAUT Lag

An item is flagged with (P) on the Work With Synchronization Attributes panel under the CDN column if the number of authority entries not yet processed for the server is greater than the threshold specified.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the Condition specified is used. The default value for *PVTAUT lag is 50.

NONE:** Do not flag sync attributes with **PVTAUT lag.

Spool Lag

An item is flagged with (S) on the Work With Synchronization Attributes panel under the CDN column if the number of QAUDJRN entries not yet processed for spooled file changes is greater than the threshold specified.

The possible values are:

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the Condition specified is used. The default value for threshold spool lag is: 2,000

***NONE:** Do not flag sync attributes with spool lag.

Object Authority Lag

An item is flagged with (O) on the Work With Synchronization Attributes panel under the CDN column if the number of object authority entries not yet processed for the library is greater than the threshold specified.

Number: Enter a threshold value to use.

***CDN:** A default value appropriate for the Condition specified is used. The default value for object authority lag is: 50

***NONE:** Do not flag sync attributes with object authority lag.

Function Keys

F3 (Exit): Exits the panel without processing any pending changes.

F4 (Prompt): Displays a list of possible values from which you may select one.

F5 (Refresh): Resets all fields to the last subset in effect.

Appendixes

The topics in this section include reference information on various Robot HA subject matter.

Appendix A: Terminology

This user guide uses a lot of different terms to describe certain functions and features in Robot HA. Refer to these definitions if you have any questions.

Role

Each system set up with Robot HA is assigned to a role, like production or backup, and can only have one role at a time. The role is defined during the initial product configuration.

• Production system

The system assigned to the *PROD role. Your original data is located here.

• Backup system

The system assigned to the *BACKUP role. It contains a copy of all replicated items it receives. In the event the production system becomes unavailable, the backup will act as the production system and have a full copy of data to work with.

° Server ID

The name you wish to give your system. You should avoid using the actual system name or IP address for the server ID, as a role swap could make things confusing. Instead, we recommend PROD and BACKUP as your system names, which reflects their original roles.

Replication

When using a remote journal, replication is the process that sends objects between systems in real time.

Synchronization

The process that sends objects between systems periodically, as defined by the sync interval set for each synchronization attribute.

System Name

The actual system name where Robot HA is installed (for example, if you installed Robot HA on ARIES, the system name would be ARIES).

IP Address

The IP address you wish to use for product installation. If you use DNS lookup, it could be the same as the system name.

Appendix B: Defining System Connections

When you set up Robot HA using the Initialize Robot HA command (INZRSFHA), a TCP/IP connection between your production and backup systems is created automatically.

However, if you would like to create connections between your backup system and a third, fourth, or even fifth system, use this section to learn how.

Overview

Before you can use any functions between systems, you must define a connection. Two concepts are important for connections: the server directory and the requester directory. Each system has both directories.

- **The server directory** is used for outbound connections. This is where you store information and options that tell Robot HA how you want to connect to other systems or partitions. You give each entry in the server directory a ten-character name or server ID, which can then be used with various product functions to indicate the system you want to connect to.
- **The requester directory** is used for inbound connections. It lets you control the rights and options that apply to other systems when they contact you.

The term "requester system" is used to refer to the system that initiates a transaction. The server system is the one that responds to the request. A given physical system can act as both a requester and a server.

Referring to systems as source and target can be confusing in some contexts, particularly when a requester system is retrieving objects from a server system. In this case, the server system is the target of the request but the requester system is the target for the objects.

Defining TCP/IP for LAN/WAN

Use the following steps to connect to another IBM i system with TCP/IP over a LAN or WAN.*

To add server directory entries:

- 1. From the Robot HA Main Menu, select option **3**, System Setup.
- 2. Select option **1**, Server definitions, to display the Work With RSF Servers panel.
- 3. Press F6 to define a new connection.

Add Server Dir	rectory Entry	(ADDRSFSDE)				
Type choices, press Enter.						
Server ID	→ PROD *TCPIP '10.3.4.5'	Name *SDLCDIAL, *SDLC, *TCPIP				
RSF data port	602 2 *NONE *BASIC *NONE *LIBL *NONE SDE for proc	Number Seconds, *NONE 1-65534, *JOB, *NONE *NONE, *BASIC, *SSL Name, *NONE Name, *LIBL *NONE, *BASIC, *MAX duction'				
F3=Exit F4=Prompt F5=Refresh F13=How to use this display	F10=Additiona F24=More keys	Bottom al parameters F12=Cancel 5				

- 4. On the Add Server Directory Entry panel, specify a name for the Server ID.
- 5. Select ***TCPIP** for the connection method, then press Enter to display more prompts.
- 6. Enter an IP address or TCP/IP network name of the other system for **Remote System**, then press Enter to display more prompts.
- 7. Accept all the defaults, then press Enter.
- 8. Enter a value for **Text 'description'**, accept the other defaults, and press Enter one last time to complete the entry.

Repeat steps 1-3 above on the other system if you want to be able to initiate the connection from either system.

Once you are done setting up TCP/IP over a LAN, you can start the server function and test your connection.

Defining TCP/IP Over the Internet

NOTE: Robot HA transmissions are encrypted by default.

Use the following steps to connect to another IBM i system with TCP/IP over the internet.

On the target (server) system:

- 1. Determine the public IP address of your network or firewall. You will need to give this IP address to anyone who will be contacting your IBM i over the internet.
- 2. Open port 602 in your firewall and forward incoming traffic to your IBM i. Or, you can ask for the public IP address of those allowed to contact your IBM i system and open port 602 to traffic from those specific IP addresses.

NOTE: All Robot HA traffic passes through port 602, so there is no need to open FTP or Telnet ports.

3. Start the server function. No one can use Robot HA to connect to your system unless the server function is started. Ending the server function is an effective way to block incoming connections.

On the source (requester) system:

- 1. Obtain the public IP address for the system you wish to contact.
- 2. Specify an IP address for "Remote System."
- 3. Create an RSF Server Directory Entry using the following steps:
 - a. Access the Work With RSF Servers panel by selecting option **1** from the System Setup menu.
 - b. Press F6 to define a new connection.
 - c. On the Add Server Directory Entry panel, specify a name for the Server ID.
 - d. Select ***TCPIP** for Connection Method, then press Enter to display more prompts.
 - e. Enter an IP address or TCP/IP network name of the other system for **Remote System**, then press Enter to display more prompts.
 - f. Accept all the defaults, then press Enter.
 - g. Enter a value for **Text 'description'**, accept the other defaults, and press Enter one last time to complete the entry.

TIP: For internet connections, be sure to press **F10** when adding the server entry and specify ***JOB** (recommended) or some value other than ***NONE** for Local Tunnel Access Port. This ensures that Robot HA Telnet and other traffic is routed to port 602. 4. Provide your public IP address to the target location so that it can use this to restrict access to its firewall.

When you sre finished, check to make sure that the TCP/IP server function has started at the target location. Then test your connection.

Appendix C: Running a System 36 Environment

If you are running a System 36 environment on your systems, keep the following in mind:

- The job used to submit replication tasks should not t run in the S/36 environment. To make sure it is not, specify *NONE for the Special Environment parameter on the profile that is used to submit replication jobs.
- The RSFSRV user profile should have *NONE specified for the Special Environment parameter on the production *and* backup system.
- Make sure to omit objects of type *S36 when replicating #LIBRARY. When it is replicated for the first time, the library on the backup system may not be able to be cleared completely if it contains an *S36 object.

NOTE: If inquiry message RSF3136 is sent and an*S36 object is the only one that could not be cleared, select option **I** to ignore the message and continue replication.