

FORTRA

Robot HA
13.11

Role Swap User Guide

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.

202601281245

Welcome to Robot HA	4
Role Swap Overview	5
Preparing for a Role Swap	6
Step 1: Create or modify your role swap test plan	6
Step 2: Check and compare your systems	7
Step 3: Ensure your libraries and IFS directories are being replicated	8
Step 4: Synchronize your configuration information, user profile, and authority list	8
Step 5: Ensure the RSFUSER library is synchronized to the backup system	10
Step 6: Ensure the RBTHALIB library is synchronized on both systems	10
Step 7: Ensure system start programs exist on both systems	11
Step 8: Run the Swap Audit Job	11
The Role Swaps	12
Swap-While-Active Role Swap	13
Before You Begin	13
The Swap-While-Active Role Swap	14
Full Role Swap	17
Before You Begin	17
The Full Role Swap	19
Unplanned Role Swaps	24
Role Swap Steps	27
Appendix A: Terminology	29

Welcome to Robot HA

Preparing for and Executing Successful Role Swaps

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 two 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.

NOTE: These instructions are intended as a guide for preparing and executing role swaps on your systems. You can supplement this information, where needed, with the *Robot HA Implementation Guide* and the *Robot HA User Guide*.

Role Swap Overview

In a two-system HA environment, one system is designated as the production system and the other 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 ARIES is the production system and TAURUS is the backup system. If a role swap is triggered, TAURUS then acts as the production system and ARIES (when available) acts as the backup system.

Role swaps fall into two categories: planned and unplanned. Under normal circumstances, a role swap will be initiated to a) test the backup system and the integrity of the replication environment or b) allow software and hardware maintenance to occur on the production system. However, if a hardware error (or other errors, such as natural disasters, data center outages, and overheating) forces the production system out of service, an unplanned role swap will be necessary.

The following sections and steps prepare you for either scenario.

Preparing for a Role Swap

Before you perform a role swap, there are several things you should do to prepare your systems. The more thorough your planning is, the smoother the role swap will be.

NOTE: If you're planning to perform a role swap, contact Technical Support before you begin. They'll guide you through the process and make sure everything is ready on both systems.

TIP: Determine which third-party products you have running on your system that should be ended before attempting to run a role swap. Generally, any product that would need to be ended before going into restricted state should be ended before performing a role swap.

NOTE: As of Robot HA version 12.06 and above, if you are an existing PowerHA customer and wish to use Robot HA and PowerHA together in your role swap plan, see the [Using PowerHA with Robot HA - *NOSWAPIP Role Swap Checklist](#) and the [Using PowerHA with Robot HA – One to Many Role Swap Checklist](#) for more information. As part of this integration, be aware that Robot HA does not end all subsystems.

Step 1: Create or modify your role swap test plan

Before you attempt a role swap, create your role swap test plan. This plan should be read and maintained on a regular basis. If you already have one, review it carefully and make any necessary changes.

Be sure to check that the following are documented and tested:

- Your system and application startup scripts
- Your replication integrity
- Ensure IBM installed products are the same on both systems.
- Your applications

NOTE: Documenting and testing your applications is important to satisfy business requirements and allow users back on the system to resume normal business processes. Ensure license codes are applied for all applications on your backup system. This includes third party applications and IBM LPPs.

Everyone's test plan and setup is different. It is up to you to determine if the role swap you attempt completes successfully.

Step 2: Check and compare your systems

Follow these steps to check and compare your systems:

- Ensure that RBTHALIB is at the same relmod on both systems.
- To ensure that triggers are enabled correctly during a role swap, the following PTFs must be installed:
 - 7.4 - SI85372
 - 7.5 - SI85373
- If the OS versions are different (production is at 7.3 and backup is at 7.2, for example), call Technical Support before proceeding. They'll discuss what steps you should take to ensure the role swap completes successfully.
- If you are doing a one to many role swap, the primary and backup systems cannot have the same name.
- 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.
- Ensure the server security for DDM TCP/IP network entries are made. To do this prompt the CHGDDMTCPA command from a command line. The value for the "lowest authentication method" should be the same on both systems. If you are using a value for "lowest authentication method" that requires a password, verify that it is set up the same for both systems. Then when swapped, Robot HA will continue to sync properly and other applications will continue to work correctly.
 - To see what authentication entries exist, run the following SQL statement as supplied by IBM. This will list the entries defined on your system. Note the values from the existing source and target systems.
SELECT * from QSYS2.DRDA_AUTHENTICATION_ENTRY_INFO ORDER BY AUTHORIZATION_NAME, SERVER_NAME
 - Ensure all entries exist on both systems.
 - For Robot HA entries, refer to [Using a Secure DDM TCP/IP Connection with Robot HA](#).
- To ensure optimum synchronization, ensure that the port used by PING requests is open on both systems.
- 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 HA will not maintain the changes to an SQL trigger on the backup systems after the initial sync. Any changes made on the production system must also be done on the backup systems.

Step 3: Ensure your libraries and IFS directories are being replicated

This is the most important step to take when preparing for a role swap. Do the following to check the status of your synchronization attributes:

1. From the Robot HA Main Menu, select option **1** Synchronization attributes. The Work With Sync Attributes panel displays.
2. Review the parameters and make any changes, then press Enter.
3. If any sync attributes have Sync Object Authorities set to *YES:
 - a. Ensure that the *PVTAUT job is running and up to date.
 - b. If the sync attributes have not synced recently, enter option **8** next to the ones you want to sync.
 - c. Run the **Private Authority Status** report to review any errors.
 - d. Resolve 'A' type errors by running option **9** to sync authorities.
4. Review the **Last Sync** dates for each attribute listed. If they are not being replicated or have not been synced recently, enter option **8** next to the ones you want to sync.
5. Press Enter.

NOTE: After pressing Enter, if there is no active or scheduled sync job, the job will run and, depending on your setup, it will be rescheduled automatically. If the job has previously been submitted and option 8 is taken, you will see a list of job(s) displayed. If this happens, you may have an active job running or the job may be in a SCD (scheduled) status. If it is in a SCD status, you need to select this with a 1 and press Enter to get the job to run immediately.

Step 4: Synchronize your configuration information, user profile, and authority list

TIP: This *CFG information needs to be synchronized from the production system to the backup system and from the backup system to the production system. The *USRPRF and *AUTL information needs to be synchronized only from the production system to the backup system.

During initial product setup, the Initialize Robot HA command (INZRSFHA) sets *AUTL, *CFG, *USRPRF, RSFUSER and RBTHALIB on PROD and *CFG and RBTHALIB on BACKUP with specific attributes that make role swapping possible.

The following default definitions are set for your configuration information, user profile, and authority list. You can check these and change their parameters by viewing the Change System Sync Attributes panel for *CFG, *USRPRF, *AUTL, and *PVTAUT:

*CFG

- *CFG is synchronized to a different library on each system. On the production system, the default to-library is RSFCFGPRD. On the backup system, the default to-library is RSFCFGBKP.
- No objects are omitted (**Include or omit items:** *OMIT; **Items to include/omit:** *NONE) from synchronization. All of *CFG is replicated.
- Synchronization is set to repeat every four hours by default. This ensures that configuration information synchronizes from production to backup *and* from backup to production consistently. The interval value can be changed after INZRSFHA sets it.

*USRPRF

- User profiles starting with Q* are omitted. This is an IBM recommendation.
- Synchronization is set to repeat every two hours. This ensures that the user profiles synchronize from production to backup consistently.

*AUTL

- No objects are omitted (**Include or omit items:** *OMIT; **Items to include/omit:** *NONE) from synchronization. All of *AUTL is replicated.
- Synchronization is set to repeat every two hours. This ensures that the authority list synchronizes from production to backup consistently.

*PVTAUT

- No objects are omitted (**Include or omit items:** *OMIT; **Items to include/omit:** *NONE) from synchronization. All of *PVTAUT is replicated.
- Synchronization is set to repeat every 10 minutes. This ensures that the private authority list synchronizes from production to backup consistently.

Synchronization should have started on the production system and backup system as part of the initial product setup. Ensure that *CFG, *USRPRF, *AUTL, and *PVTAUT have been synchronized on a regular (and recent) basis. For more information on initial synchronization, see "Setting Up Synchronization" in the *Robot HA Implementation User Guide*.

Step 5: Ensure the RSFUSER library is synchronized to the backup system

The RSFUSER library needs to be synchronized from the production system to the backup system.

During the initial product setup, the Initialize Robot HA command (INZRSFHA) set the following default definitions for RSFUSER (you can check these by viewing the Change Library Sync Attributes panel for RSFUSER):

- Journaling and remote journaling are set to *NONE. This is done because after initial setup, the data changes less frequently, making it a good candidate for sync only.
- Synchronization is set to repeat every four hours. This ensures that RSFUSER synchronizes from production to backup consistently.

Synchronization should have started on the production system and backup system as part of the initial product setup. Ensure that RSFUSER has been synchronized on a regular (and recent) basis. For more information on initial synchronization, see "Setting Up Synchronization" in the *Robot HA Implementation User Guide*.

Step 6: Ensure the RBTHALIB library is synchronized on both systems

TIP: The RBTHALIB library needs to be synchronized from the production system to the backup system **and** from the backup system to the production system. This ensures that the rules for replication on the production system are also on the backup system in the event of a role swap.

During initial product setup, the Initialize Robot HA command (INZRSFHA) sets definitions on both systems with specific attributes that make role swapping possible. The following definitions are set for RBTHALIB (you can check these by viewing the Change Library Sync Attributes panel for RBTHALIB):

- Journaling and remote journaling are set to *NONE. This is done because after initial setup, the data changes less frequently, making it a good candidate for sync only.
- RBTHALIB is synchronized to a different library on each system. On the production system, the default to-library is RSFPRD. On the backup system, the default to-library is RSFBKP.
- All programs (*PGM) are omitted from synchronization. Only data is needed for a role swap; RBTHALIB programs are already on both systems.

- Synchronization is set to repeat every hour. This ensures that RBTHALIB synchronizes from production to backup *and* from backup to production consistently.

Synchronization should have started on the production system and backup system as part of the initial product setup. Ensure that RBTHALIB has been synchronized on a regular (and recent) basis. For more information on initial synchronization, see "Setting Up Synchronization" in the *Robot HA Implementation Guide*.

Step 7: Ensure system start programs exist on both systems

In order to perform a role swap, you need system start programs on the production and backup systems. These programs ensure that your replication environment will start consistently. These programs are specified when running the Initialize Robot HA command (INZRSFHA). The system start program for both Prod and Backup should have been set at this time.

The system start program on the backup system only needs to start synchronization for configuration information and the RBTHALIB library.

Step 8: Run the Swap Audit Job

After performing steps 1-7, you can run a Swap Audit job to provide a more in-depth layer of assurance that you are ready to successfully perform a role swap. See the Swap Audit sections in the Robot HA User Guide for details on how to create and run this audit.

The Role Swaps

A role swap can take two different forms, each with its own benefits.

The following tables summarize these options and benefits. View them carefully to decide which role swap is right for your environment.

Table 1: Role swap types and descriptions

Type	Description	Benefits
Swap-While-Active	The swap-while-active role swap allows work to continue on the production system during the test.	<p>A swap-while-active role swap allows you to perform a test swap during working hours. During this time, it will continue to update files on the production system.</p> <p>TIP: We recommend you use the Swap-While-Active method first when testing a role swap. It is the least invasive way to perform a role swap. Your production system can still run during the swap, and there is no disruption to any users on the system.</p>
Full	Performs a complete swap of both systems.	<p>A full planned swap may take a while to complete, but it exercises role swaps in both directions on both systems. Any user data changes made during the role swap are preserved.</p> <p>NOTE: The production system is not available during the swap.</p>

Table 2: Role swap comparisons

Type	Swaps PROD	Swaps BACKUP	Work Continues on PROD	Allows Data Change on BACKUP	Preserves Data Changes from Swap	Restores Data on BACKUP After Swap
Swap-While-Active		✓	✓	✓		✓ - Via PROD
Full	✓	✓		✓	✓	

Swap-While-Active Role Swap

A swap-while-active role swap allows you to perform a test swap during working hours and continue to update files on the production system.

TIP: We recommend you use the swap-while-active method before attempting a full role swap, as it is the least invasive role swap of the two. Your production system can still run during the swap, and there is no disruption to any users on the system.

Before You Begin

Read the following considerations before you perform a swap-while-active role swap.

- Read [Preparing for a Role Swap](#). There are several steps you must take to ensure your data is replicated and up-to-date.
- Determine which third-party products you have running on your system that should be ended before attempting to run a role swap. Generally, any product that would need to be ended before going into restricted state should be ended before performing a role swap.
- If you are running a role swap with systems at different OS level, contact Technical Support to discuss your role swap plan. There are additional steps to take to ensure a successful swap.
- Run a Swap Audit to ensure that you are prepared for a successful role swap.
- Check the Initialize Robot HA (INZRSFHA) command before you attempt this role swap. If *SWAPIP is listed as the role swap type, ensure that the 'Set test flag' field is set to *YES when you execute the role swap on the Role Swap panel (option 7, Role Swap Menu from the Robot HA Main Menu. Then, select 10, Role Swap). In this scenario, setting the test flag to *YES will not swap IP and you will be able to run the swap test without the IP address changing.

NOTE: When running a swap-while-active role swap and specifying *SWAPIP you want to ensure you specify *YES for the 'Set test flag' field. During a swap-while-active role swap, the production system is not "down," only invisible. Because of this, *SWAPIP with 'Set test flag' set to *YES will not change the IP address, system name and will not allow TCP/IP to start on the backup system.

- The role swap must be run with QSECOFR or user profile with *SECOFR user class.
- Decide whether you want to run the role swap interactively or in batch:

- If running interactively, you must be signed on to the system console.

NOTE: Display size must be 27x132 to display Work With Sync Attributes Panel. HMC 5250 Console does not support 132 columns.

- If running in batch, the role swap should be running in a job queue in the controlling subsystem (QBASE or QCTL, for example).
- We recommend you sync library RBTHALIB before attempting a swap-while-active role swap.
- Contact Technical Support to ensure everything is ready on both systems. They'll guide you through the process and make sure the role swap completes successfully.

The Swap-While-Active Role Swap

NOTE: Users can continue working and making changes to the production system throughout all of the following steps.

To perform the role swap:

Make sure all sync jobs have ended, then continue with the following steps.

On the production system:

1. Sign on as QSECOFR or user profile with *SECOFR user class.
2. From the Robot HA Main Menu, select option **2**, Control Menu.
3. Select option 6 to end audit monitor.
4. Select option 4 to end all synchronization jobs.
5. Select option 2 to end the TCP/IP servers.

The production system should now be invisible to the backup system because the communication link has been ended (Step 5 above).

On the backup system:

1. Sign off the system prior to running the role swap. This releases any locks you might have from working in the product.
2. Select option 7, Role Swap Menu from the Robot HA Main Menu. Then, select 10, Role Swap.
3. Do the following on the Role Swap to Production panel:

- a. Choose whether you want to run the role swap from the system console, or specify *YES for **Run in batch** to submit the job to a job queue.

NOTE: Either choice will bring your system down into restricted state.

- b. Change **Set test flag** to *YES if you are using *SWAPIP so Robot HA does not change the IP address and system name. If you are not running with *SWAPIP, this flag can be set to *YES or *NO.

NOTE: When using *NOSWAPIP, *YES will run the startup program you specified on the INZRSFHA command - inzrsfha parm 'System test start program', otherwise *NO will run the program specified for 'System start program'.

- c. Specify *YES for **Are you sure?**, otherwise the role swap will not proceed.
 - d. Press Enter. This swaps the backup system to the production role.
4. After submitting the role swap, watch for messages in the following places:
 - If running **interactively**, the messages will show up on the screen or in your interactive joblog.
 - If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

When the role swap completes, users can sign on. The production system should be hidden (the communication link was ended), and the backup system is in the production role.

Verify your data is correct. It is recommended that you do not allow changes to the backup system (currently in the production role) during the test.

TIP: We recommend you do not attempt to run any synchronization or audit jobs during the test.

To end the role swap:

NOTE: Make sure all users are signed off the backup system (currently in the production role).

On the system in the production role or the backup system in the production role:

1. Sign off the system prior to running the role swap. This releases any locks you might have from working in the product.
2. On the system in the production role or the backup system in the production role, execute the following command to clear the RSFHA job queue and remove any synchronization jobs that might not have been submitted in a previous step:

CLRJOBQ JOBQ(RSFUSER/RSFHA) LOG(*NONE)

This ensures that any journal changes made on the backup system during the swap-while-active role swap are not applied to the production system.

3. Select option 7, Role Swap Menu from the Robot HA Main Menu. Then, select 10, Role Swap.
4. Do the following on the Role Swap to Backup panel:
 - a. Choose whether you want to run the role swap from the system console, or specify *YES for **Run in batch** to submit the job to a job queue.

NOTE: If you choose to run in batch, the system will be brought down into restricted state.

- b. Specify *YES for **Are you sure?**, otherwise the role swap will not proceed.
 - c. Press Enter. This swaps the system back to the backup role.
5. After submitting the role swap, watch for messages in the following places:
 - If running **interactively**, the messages will show up on the screen or in your interactive joblog.
 - If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

6. Then, ensure the TCP/IP server has started and subsystem RSFHA is active. You may need to start them manually.
 - To check the server, navigate to the Control Menu and view the "TCP/IP Server is" status.
 - To check the subsystem, execute the **WRKSBS** command and look for RSFHA in the list of subsystems.
7. Job queue RSFUSER/RSFHA is automatically placed on hold after a role swap. To resume synchronization, release the job queue by executing the **RLSJOBQ RSFUSER/RSFHA** command.
8. After you release the job queue, sign off the system.

On the production system:

Return to the hidden production system and restart it using the following steps:

1. Sign on to the system.
2. From the Robot HA Main Menu, select option **2**, Control Menu.
3. Select option **1** to start the TCP/IP server.
4. Select option **3** to start all synchronization jobs.
5. Select option **5** to start the audit monitor.

On the backup system:

1. Sign on and open Robot HA.
2. From the Robot HA Main Menu, select option **2**, Control Menu.
3. Select option **1** to start the server, option **3** to restart all synchronization jobs, and option **5** to restart the audit monitor.

When you are finished, the systems should be in their normal roles again.

Full Role Swap

A full role swap is used in the event of a planned or [unplanned role swap](#). It may take a while to complete, but it exercises role swaps in both directions on both systems. Any user data changes made during the role swap are preserved.

Before You Begin

Read the following considerations before you perform a full role swap.

- Read [Preparing for a Role Swap](#). There are several steps you must take to ensure your data is replicated and up-to-date.
- Determine which third-party products you have running on your system that should be ended before attempting to run a role swap. Generally, any product that would need to be ended before going into restricted state should be ended before performing a role swap.
- If you are running a role swap with systems at different OS level, contact Technical Support to discuss your role swap plan. There are additional steps to take to ensure a successful swap.
- The role swap must be run with QSECOFR or user profile with *SECOFR user class.
- Decide whether you want to run the role swap interactively or in batch:
 - If running interactively, you must be signed on to the system console.

NOTE: Display size must be 27x132 to display Work With Sync Attributes Panel. HMC 5250 Console does not support 132 columns.

- If running in batch, the role swap should be running in a job queue in the controlling subsystem (QBASE or QCTL, for example).
- If *SWAPIP is listed as the role swap type in the Initialize Robot HA (INZRSFHA) command, the system will IPL to allow the system name to change. Ensure that the 'Set test flag' field is set to *NO when you execute the role swap on the Role Swap panel (option 7, Role Swap Menu from the Robot HA Main Menu. Then, select 10, Role Swap).

NOTE: *SWAPIP with 'Set test flag' field set to *YES will NOT change the backup system's IP address and system name to match the production system's IP address and system name during Role Swap execution. *YES would be specified for a swap-while-active role swap test.

- Before running a *SWAPIP role swap, make note of the "lowest authentication method" set on the CHGDDMTCPA command. This value may change during the role swap. After the role swap, check this value again to ensure it is set correctly. In addition, if you are using the value *USRDPWD for "lowest authentication method" on the CHGDDMTCPA command, refer to [Using a Secure DDM TCP/IP Connection with Robot HA](#)
- If using Insite, bring down the Robot HA HTTP server.
- Contact Technical Support to ensure everything is ready on both systems. They will guide you through the process and make sure the role swap completes successfully.

The Full Role Swap

To perform the role swap:

NOTE: Make sure all users are signed off the production system. It is also important to verify that all synchronization jobs are complete and any subsystems or jobs associated with the production system are ended.

On the production system: Role Swap To Backup

1. Sign off the system and sign back on prior to running the role swap. This releases any locks you might have from working in the product.
2. Select option **7**, Role Swap Menu from the Robot HA Main Menu. Then, select **10**, Role Swap.
3. Do the following on the Role Swap to Backup panel:
 - a. Choose whether you want to run the role swap from the system console, or specify `*YES` for **Run in batch** to submit the job to a job queue.

NOTE: Either choice will bring your system down into restricted state.

- b. Leave Set test flag as the default `*NO`. This flag must be set to `*NO` in order for us to change the IP address and system name if the role swap is set to `*SWAPIP`.
 - c. Specify `*YES` for **Are you sure?**, otherwise the role swap will not proceed.
 - d. Press Enter. This swaps the production system to the backup role.
4. After starting the role swap, watch for messages in the following places:
 - If running **interactively**, the messages will show up on the screen or in your interactive joblog.
 - If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

5. Then, ensure the TCP/IP server has started and subsystem RSFHA is active. You may need to start them manually.

- To check the server, navigate to the Control Menu and view the "TCP/IP Server is" status.
 - To check the subsystem, execute the **WRKSBS** command and look for RSFHA in the list of subsystems.
6. Job queue RSFUSER/RSFHA is automatically placed on hold after a role swap. To resume synchronization, release the job queue by executing the **RLSJOBQ RSFUSER/RSFHA** command.

On the backup system: Role Swap To Production

Complete the following steps on the backup system to swap it to the production role.

1. Sign off the system and sign back on prior to running the role swap. This releases any locks you might have from working in the product.
2. Select option **7**, Role Swap Menu from the Robot HA Main Menu. Then, select **10**, Role Swap.
3. Do the following on the Role Swap to Production panel:
 - a. Choose whether you want to run the role swap from the system console, or specify ***YES** for **Run in batch** to submit the job to a job queue.

NOTE: Either choice will bring your system down into restricted state.

- b. If you want to flag the role swap as a test run, change **Set test flag** to ***YES**. Otherwise, leave the default as ***NO**. The option ***NOCHG** will use the value specified during the last role swap (this is helpful if your role swap needs to be restarted).
- EXAMPLE:** Specifying ***YES** allows you to run different system startup programs, since it is flagged as a test role swap instead of a real one.
- c. Specify ***YES** for **Are you sure?**, otherwise the role swap will not proceed.
 - d. Press Enter. This swaps the backup system to the production role.
4. After starting the role swap, watch for messages in the following places:

- If running **interactively**, the messages will show up on the screen or in your interactive joblog.
- If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

5. Then, ensure the TCP/IP server has started and subsystem RSFHA is active. You may need to start them manually.
 - To check the server, navigate to the Control Menu and view the "TCP/IP Server is" status.
 - To check the subsystem, execute the **WRKSBS** command and look for RSFHA in the list of subsystems.
6. Job queue RSFUSER/RSFHA is automatically placed on hold after a role swap. To resume synchronization, release the job queue by executing the **RLSJOBQ RSFUSER/RSFHA** command.

When all steps in the role swap have been followed, users can sign back on and resume working. The systems should now be swapped, with the production system in the backup role and the backup system in the production role.

NOTE: If you want to synchronize your changes to the production system during your swap, navigate to the Control Menu and take option **11** to see if the sync jobs have been submitted. If they have not, navigate to the Control Menu and take option **3**.

To end the role swap:

NOTE: Make sure all users are signed off the system in the production role.

On the system in the production role:

1. Sign off the system and sign back on prior to running the role swap. This releases any locks you might have from working in the product.
2. Select option **7**, Role Swap Menu from the Robot HA Main Menu. Then, select **10**, Role Swap.
3. Do the following on the Role Swap to Backup panel:

- a. Choose whether you want to run the role swap from the system console, or specify *YES for **Run in batch** to submit the job to a job queue.

NOTE: If you choose to run in batch, the system will be brought down into restricted state.

- b. Specify *YES for **Are you sure?**, otherwise the role swap will not proceed.
 - c. Press Enter. This swaps the system back to the backup role.
4. After starting the role swap, watch for messages in the following places:
 - If running **interactively**, the messages will show up on the screen or in your interactive joblog.
 - If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

5. Then, clear job queue RSFUSER/RSFHA.
6. Ensure the TCP/IP server has started and subsystem RSFHA is active. You may need to start them manually.
 - To check the server, navigate to the Control Menu and view the "TCP/IP Server is" status.
 - To check the subsystem, execute the **WRKSBS** command and look for RSFHA in the list of subsystems.
7. Job queue RSFUSER/RSFHA is automatically placed on hold after a role swap. To resume synchronization, release the job queue by executing the **RLSJOBQ RSFUSER/RSFHA** command.

On the system in the backup role:

Complete the following steps on the system in the backup role to swap it to production.

1. Sign off the system and sign back on prior to running the role swap. This releases any locks you might have from working in the product.

2. Select option **7**, Role Swap Menu from the Robot HA Main Menu. Then, select **10**, Role Swap.
3. Do the following on the Role Swap to Production panel:
 - a. Choose whether you want to run the role swap from the system console, or specify *YES for **Run in batch** to submit the job to a job queue.
4. After starting the role swap, watch for messages in the following places:
 - If running **interactively**, the messages will show up on the screen or in your interactive joblog.
 - If running **in batch**, watch for messages in QSYSOPR until the swap completes successfully.

NOTE: If you choose to run in batch, the system will be brought down into restricted state.

NOTE: Because running the role swap in batch sends you into restricted state, you will lose your interactive session and need to wait until the system comes out of restricted state to see the current progress of the role swap.

5. Then, clear job queue RSFUSER/RSFHA.
6. Ensure the TCP/IP server has started and subsystem RSFHA is active. You may need to start them manually.
 - To check the server, navigate to the Control Menu and view the "TCP/IP Server is" status.
 - To check the subsystem, execute the **WRKSBS** command and look for RSFHA in the list of subsystems.
7. Job queue RSFUSER/RSFHA is automatically placed on hold after a role swap. To resume synchronization, release the job queue by executing the **RLSJOBQ RSFUSER/RSFHA** command.

On both systems:

Take option **1** on the Control Menu to start the servers, then take option **3** to resume synchronization.

When all steps in the role swap have been followed, users can sign back on and resume working. The systems should be in their normal roles again.

If users cannot telnet following the role swap, follow these steps:

1. Log in to the console for the system having the telnet issue.
2. Enter **CFGTCP** and take menu option **1**.
3. Find the TCP/IP address for the system you are currently on and use option **10** to end the TCP/IP interface.
4. You should receive a message that TCP/IP interface was successfully ended.
5. Use option 9 with the same system to start the TCP/IP interface.
6. You should now be able to successfully telnet to this system.

These steps are necessary because when doing a *SWAPIP, the TCP/IP interface needs to be restarted to use the updated address.

Unplanned Role Swaps

Even in pristine HA environments, the production system can go down unexpectedly. In these cases, an unplanned role swap may be necessary to continue work until the issue is resolved.

NOTE: You must be signed onto the system console as QSECOFR or equivalent to perform an unplanned role swap. If you want to run it in batch, it must go onto a job queue in the controlling subsystem (QBASE or QCTL, for example).

Performing an unplanned role swap:

Follow these steps to swap the backup system into the Production role.

1. Sign off the system prior to running the role swap. This releases any locks you might have from working in the product.

- Execute the following commands to check RBTHALIB and RSFUSER for object locks:

WRKOBJLCK OBJ(RBTHALIB) OBJTYPE(*LIB)

WRKOBJLCK OBJ(RSFUSER) OBJTYPE(*LIB)

If there are locks, ensure they are removed.

- Select option **7**, Role Swap Menu from the Robot HA Main Menu. Then, select **10**, Role Swap. The Role Swap to Production panel displays.

```

Role Swap to Production (SWAPPRDRSF)

Type choices, press Enter.

Run in batch . . . . . *NO          *YES, *NO
Set test flag . . . . . *NO          *YES, *NO, *NOCHG
Are you sure? . . . . . *YES        *YES, *NO
Current role . . . . . *BACKUP
    
```

- Choose whether you want to run the role swap from the system console, or specify ***YES** for 'Run in batch' to submit the job to a job queue.

NOTE: Regardless of if you choose to run interactively or in batch, the system will be brought down into restricted state.

- If you are doing a ***SWAPIP** role swap (***SWAPIP** is listed as the role swap type in the Initialize Robot HA (INZRSFHA), ensure that the 'Set test flag' field is set to ***NO**.

When 'Set test flag' is set to ***NO**, we will:

- End all subsystems
- Change the system name
- Change the IP address
- IPL the system

NOTE: In order to allow for the IPL, ***SWAPIP** must be listed as the role swap type in the Initialize Robot HA (INZRSFHA) command, AND the 'Set test flag' field must be set to ***NO**.

If the 'Set test flag' is set to ***YES**, the system will not IPL and the system name will not change.

- Specify ***YES** for 'Are you sure?', otherwise the role swap will not proceed.
- Press **Enter**. This swaps the backup system to the Production role.

If running interactively, status messages at the bottom of the screen will keep you informed of the role swap progress

After the Swap to Production

1. Ensure your Robot HA server is not active on the swapped backup system (now in the Production role). This needs to remain inactive until after you have your original production system active and you have swapped it to a Backup role. If you have changed your start up program to include Robot HA commands on your original production it will attempt to make Robot HA active. After an unplanned outage, this needs to be controlled.
2. Your sync jobs do not need to be started because there is nowhere to send things. You will need to wait for your unexpected outage on your original production system to be resolved.
3. Once it is resolved, you must first swap your original production system to a Backup role and sync the changes made during the unplanned downtime to your original production system.

Ending an unplanned role swap:

Once the problem with the production system is resolved and becomes available for use again, follow the instructions below to end the role swap.

1. Swap the original production system to a temporary Backup role so that you can sync the changes that have been made since your unplanned down time started.
 - a. Sign off the system prior to running the role swap. This releases any locks you might have from working in the product.
 - b. Execute the following commands to check RBTHALIB and RSFUSER for object locks:
WRKOBJLCK OBJ(RBTHALIB) OBJTYPE(*LIB)
WRKOBJLCK OBJ(RSFUSER) OBJTYPE(*LIB)
 If there are locks, ensure they are removed.
 - c. Select option **7**, Role Swap Menu, from the Robot HA Main Menu. The Role Swap to Backup panel displays.

```

Role Swap to Backup (SWAPBKPRSF)

Type choices, press Enter.

Run in batch . . . . . *NO          *YES, *NO
Set test flag . . . . . *NO          *YES, *NO, *NOCHG
Are you sure? . . . . . *YES        *YES, *NO
Current role . . . . . *PROD
  
```

- d. Choose whether you want to run the role swap from the system console, or specify ***YES** for 'Run in batch' to submit the job to a job queue.

NOTE: Regardless of if you choose to run interactively or in batch, the system will be brought down into restricted state.
 - e. Specify ***YES** for 'Are you sure?', otherwise the role swap will not proceed.
 - f. Press **Enter**. This swaps the production system to the Backup role.
2. Once you have your original production system in a Backup role, you need to start syncing from the system currently in the Production role to the system currently in the Backup role.
 - a. Ensure your servers are active on both your production and backup
 - b. Ensure the RSFUSER/RSFHA subsystem is started
 - c. Ensure the RSFUSER/RSFHA jobq is released
 - d. Start syncing. You need to ensure all the changes made during the unplanned downtime have been synced. This time frame can vary depending on how long your downtime was.
 3. You are now replicating from the backup system to the production system, with the backup system in the Production role.
 4. Plan a time to swap your systems back to their normal roles because both systems will be affected. To swap your systems back to their normal roles, wait until all replication jobs are complete. Then, follow the steps listed for [ending a full planned role swap](#).

Role Swap Steps

The Role Swap runs a number of distinct steps. The steps defined below are the actions that Robot HA performs while running a role swap.

NOTE: Should an error occur at any point, the swap program, when re-run, will automatically reposition itself at the step in which the error occurred. When rerunning the role swap, a message is displayed to the user (RSF7203) to confirm or change the restart step.

Role Swap Step Definitions:

Step Name	Description
ENDSYNC	Ends all running swap-ready and sync jobs.
RMVRMTJRN	Removes remote journal linkage.
DLTRMTJRN	Deletes remote journals.
SAVRSF	Saves current configuration and reset Robot HA settings.
SAVCFG	Saves the Robot HA configuration.
ENDRSFSRV	Ends Robot HA server jobs.
IMPORTRSF	Imports Robot HA configuration.
SETRSFSDE	Sets server IDs with proper IP addresses.
MANUALSYNC	Sets Sync point.
ENDSBS	Ends all subsystems in preparation for resetting configuration.
SETUSRPRF	Sets user profiles.
SETTRG	Sets triggers.
SETCST	Sets constraints.
APYTCP	Currently NA.
APYSCD	Applies job schedule entries.(SWAPTOPRD only)
RMVSCD - Remove schedule entries. (SWAPTOBKP only)	Applies job schedule entries. Then, removes RSF scheduled jobs.
APYRDB	Applies remote directory entries.
APYSNADS	Currently NA.
APYSYS	Applies SYSVAL, NETA, etc.
COMPLETION	Clears data areas, check point. Checks test mode, then either RE-IPL or restarts controlling subsystem.

NOTE: All of the swap programs run the same commands with either backup or production values interchanged. Not all steps necessarily run in all programs.

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.