



# Assure iTERA HA 6.2 Failover Checklist

This document contains instructions to perform a live failover.

---

**IMPORTANT**

*If you have experienced an actual failure on your primary node, call Precisely [Support](#) **immediately** or contact your local provider of product support. As part of your support agreement, Precisely will help you work through a failover situation. By providing this document, Precisely does not imply, recommend, or suggest that you perform a failover on your own. As part of your support agreement, Precisely will help you work through a failover situation. By providing this document, Precisely does not imply, recommend, or suggest that you perform a failover on your own.*

---

## Before executing the failover

- Make sure you are signed on with the iTERA admin profile. (Identified in 30.21, F7.)
- Make sure you know the user IP address before you continue. The user IP addresses is located on the 30.21 (*User (Takeover) IP* field toward the top) or 30.22.
- The failover will not process correctly if communications are detected between machines. *Do not attempt the failover if communications are active.*

---

**IMPORTANT**

After you start the failover process, do not access menu options 30.21 or 30.22 on either system unless instructed to do so.

---

## Definitions

The following terms are used throughout this procedure.

Node	Definition
Primary	The node that is <i>currently</i> defined to iTERA as the primary.
Backup	The node that is <i>currently</i> defined to iTERA as the backup.
Old Primary	The node that <i>was</i> acting as primary <i>prior to</i> starting the failover.

Node	Definition
Old Backup	The node that <i>was</i> acting as backup <i>prior to</i> starting the failover.
New Primary	Upon failover completion, the node <i>now</i> acting as primary.
New Backup	Upon failover completion, the node <i>now</i> acting as backup.

---

**IMPORTANT**

Never manually end the subsystems, except where specifically indicated in this document. The failover process will automatically end the subsystems when appropriate.

Once the issue with the failed system has been resolved, when ready to start replication from the new primary **DO NOT start the iTERA subsystem or user jobs on the new backup node.** These processes will occur automatically as part of the steps in the “Resume Replication” section. On the new backup, check QSTRUP while in restricted mode to ensure E2SBS is not started automatically.

---



---

**IMPORTANT**

This checklist should only be used in cases where the primary is down. However, it may also be used for simulating a failover but Precisely does not encourage this practice. Instead, we recommend that the role swap be tested at least quarterly, as it uses the same processes as the failover and includes additional steps that validate role swap readiness. If you must test the failover process, first perform the steps in the “Simulate the Failover” section.

---



---

**IMPORTANT**

Never override the sequence number of any apply jobs if the primary is down. Doing so will delete the journal receivers on the backup, which will result in the loss of data that cannot be recovered.

---



---

**IMPORTANT**

If you attempt a failover then need to abort it for any reason, contact Support immediately for assistance.

---

# Failover Procedure

## Section One: Failover Prechecks

Perform the following pre-check steps before executing the actual failover procedure.

Section One: Failover Pre-Checks	Assigned	Done
1. Hold all scheduled jobs on the backup node that may submit during the time the environment is failed over.		
2. Create a document (e.g., in MS Word, Notepad, etc.) for recording selected information about your failover process. Copy and paste information from your IBM i screens into the document as directed.		
3. Check the Objects Requesting Sync screen (1.1, F6). Copy the names of the files into the document.		
4. Check for problems in the Mirrored Object Maintenance screen (1.21, backup). Record any objects that display when the daily filtering checks are performed (this is a step in the “Monitoring Checklist” in the Assure iTERA HA User Guide.)		
5. Check the Work with Audits (WRKAUD) screen on the target node. Select option 7=History on any unsuccessful audit (such as those with a status of *FAILED or *NOTRCVD), then press Enter. Select option 8=View difference details on the most recent audit to review the list of objects not recovered.		
6. Check the Heal Monitor screen (3.7, backup) for pending entries. Record in the document any objects that have pending entries.		
7. Check the <i>System Monitor</i> for latency (1.1, backup). The values are current as of the last update. Record the number of <i>Journal Entries Not Applied</i> . If the number of entries not applied concerns you, check F16=Process Monitor to see how many are exposed entries (created on the primary but not found on the backup). Record the information in the document.		
<p><b>IMPORTANT</b></p> <p>If there are exposed entries, then a decision must be made on whether the transactions should be allowed to be lost. <u>The entries will be lost if the failover is executed.</u>  <u>If it is too costly to lose them then do not execute the failover. Repair the primary system and restart it.</u></p>		

Section One: Failover Pre-Checks	Assigned	Done
8. Check the subsystem (E2SBS, backup). The iTERA and MIMIX jobs must be running.		
9. Review the list of out-of-sync objects (from steps 3, 4, and 5, above) with the application personnel to check whether you need the data contained in these files.  <div style="border: 1px solid black; border-radius: 10px; padding: 2px; display: inline-block; background-color: black; color: white; font-weight: bold;">IMPORTANT</div> The loss of these objects may result in data loss or data integrity. These objects may not have all of the information or data.		
If you have not already done so, contact Precisely <a href="#">Support</a> for assistance in finishing the failover.		

## Section Two: Execute the Failover

Section Two: Execute the Failover	Assigned	Done
<p>10. Execute the failover (40.30, F16, F16; backup).</p> <p>The Role Swap Readiness tests are automatically executed. If all tests return an OK status, no results are displayed. However, if a test returns a status of <i>ERR</i>, <i>WRN</i>, <i>SEV</i>, or <i>FTL</i>, the test results are displayed on the screen. Investigate and resolve issues prior to continuing. While some issues cannot be resolved, any questionable results should be discussed with your Support Representative. Resolution for any test in <i>SEV</i> or <i>FTL</i> status requires assistance from Support.</p> <ul style="list-style-type: none"> <li>• If the test in error is a target system test, use the F7 key to display the Role Swap Readiness Monitor and resolve the issue(s).</li> <li>• Since the primary is down, you will not be able to resolve issues.</li> </ul> <p>After all issues have been investigated then on the Role Swap Readiness Test Errors screen on the target system press F22=Continue to continue failover processing.</p> <hr/> <p><b>IMPORTANT</b></p> <p>Review the confirmation screen. Verify that the screen indicates “Failover”. If it indicates “Role Swap”, <u>do not continue with the Failover</u>. The most likely cause is that communications are still active between systems. Contact Support immediately in order to troubleshoot and resolve the issue.</p> <hr/> <p>If “Failover” is indicated on the screen, review the parameters and set as needed, then press F10 to initiate the failover. The current backup will be converted to the new primary.</p>		
<p>11. If there are unapplied entries, a window is displayed with the heading “WARNING - UNAPPLIED ENTRIES FOUND” displayed at the top. Press F5=Refresh periodically and note the change in the number of unapplied entries. When Unapplied Entries is zero, role swap processing is automatically resumed. (You do not have to wait on this screen for the entries finish processing. You may press F22=Continue with Role Swap to continue. The role swap process will automatically monitor for the completion of the apply jobs before continuing.) If desired, you may press F7 to review the Journal Apply Statistics screen and then use F12 to return to this screen.</p>		

Section Two: Execute the Failover	Assi-g ned	Done
<p>12. The Open Commit Cycles screen is displayed. Processing for these cycles can be rolled back to the beginning of the commit cycle (the data will be lost) (option 1=Rollback), or the rollback can be blocked (the data in the open commit cycle is retained and a partial commit cycle is applied to the data; a manual check of the commit cycle will be required) (option 4=Cancel Rollback).</p> <ul style="list-style-type: none"> <li>• For open commit cycles that should be rolled back, enter option 1. (The remote receivers containing the entries for the open commit cycles must be on the target in order to roll back.)</li> <li>• For open commit cycles that should not be rolled back enter option 4. This will block rollback processing for these cycles.</li> <li>• If an option is not entered for a commit cycle, it will fail validation and none of the commit cycles can be processed.</li> <li>• When all lines have been validated, press F10=Continue with Failover.</li> </ul> <p><b>NOTE</b></p> <p>If F3 or F12 is pressed, the failover will be aborted and the display returned to the Role Swap Monitor. Commit cycles will not be rolled back.</p>		
<p>13. When control of the system has returned, do the following:</p> <ul style="list-style-type: none"> <li>– Refresh the main iTERA menu (F5). Verify that the role indicates it is the primary machine.</li> <li>– Enter E2SBS to check the iTERA subsystem. It should <i>not</i> be active. However, if there are additional target nodes defined, it will be active.</li> <li>– Check 40.30 F8. Step 10 should indicate <i>Complete</i>.</li> </ul>		

Section Two: Execute the Failover	Assigned	Done
<p>14. <b>Optional:</b> If using ODBC, JDBC, or any other communication protocol, you may need to change the new primary *LOCAL RDB entry to be the same as it was on the old primary.</p> <hr/> <p><b>IMPORTANT</b></p> <p>Changing the *LOCAL RDB entry will cause security changes in CHGDDMTCPA. Verify the settings in CHGDDMTCPA prior to removing the entry so that these settings can be restored after the RDB change is made. Failure to verify these security settings will cause communication failure between systems in the replication environment.</p> <hr/> <p>To change the *LOCAL RDB entry on the new primary, do the following:</p> <ol style="list-style-type: none"> <li>a. Select 30.3.</li> <li>b. Delete *LOCAL (opt 4).</li> <li>c. Respond to the message with G, then press Enter.</li> <li>d. Create *LOCAL on the new primary using the value that was in *LOCAL on the old primary.</li> <li>e. If you are not using passwords on DDM files then in the <i>Check DDM Attributes</i> screen (30.4), verify that the <i>Lowest Authentication Method</i> (OS 6.1 and later) is set to *USRID. (A service authorization entry on all nodes is required.) (In V5R4 verify that the <i>Password Required</i> field is set to *NO.</li> </ol>		

Section Two: Execute the Failover	Assigned	Done
<p>15. <b>Optional:</b> If necessary, change the system name.</p> <p><b>IMPORTANT</b></p> <p>Precisely does <u>not</u> recommend you change the system name; most applications will work without making the change. JDBC and ODBC drivers do not use the system name, they use the *LOCAL entry in the relational database. If your software vendor has told you to change the system name, use the instructions below. If you do change the system name, follow the instructions below in the order listed. Failure to do so will require an entire system resync!</p> <hr/> <p>If using APPC connections you must change LCLCPNAME and LCLLOCNAME.</p> <p>To change the system name on the new primary to the system name of the old primary, do the following:</p> <ol style="list-style-type: none"> <li>Check to see if you are using passwords on DDM files (30.4).</li> <li>Use CHGNETA to change the system name on the new primary to the old primary's name.</li> <li>IPL the new primary using your company's established protocol.</li> <li>If you are using passwords on DDM files you will need to update the Server Authority Entries (DSPSVRAUTE to display, CHGSVRAUTE to change).</li> <li>Check TCP/IP Domain Information to see if you need to change the host name (CFGTCP, opt 12).</li> </ol>		
<p>16. If you are replicating WebSphere MQ, do the following on the new primary:</p> <ol style="list-style-type: none"> <li>Start the QMQM subsystem and all MQ manager(s) using your company's established protocol.</li> <li>If MQ fails to start, remove WebSphere MQ markers using the command E2DLTMQCPT.</li> </ol>		
<p>17. Move any necessary devices from the old primary to the new primary.</p>		
<p>18. If the primary and backup are on different subnets, redirect users' IP address to the new primary.</p>		
<p>19. Start any user processes and applications, and allow users to access the new primary.</p>		
<p>20. <b>Optional:</b> If running on the new primary for an extended period of time, release the scheduled jobs.</p>		

## Section Three: Resume Replication

**NOTE**

You must wait to start replicating to the new backup until it is available.

**NOTE**

The iTERA subsystem on the new primary will not start until replication to the new backup has been initiated. However, if there are multiple target nodes defined, the subsystems will be active.

**IMPORTANT**

If the old primary system has been lost or needs to be reloaded, you will need to reinstall and reconfigure iTERA! Do not attempt this on your own. You must contact [Support](#) for assistance.

**IMPORTANT**

When the new backup is available it **MUST** be brought back up in a restricted state so that the user IP address does not become active and the iTERA subsystem is not automatically started. **DO NOT** start user jobs or the iTERA subsystem on the new backup node.

**IMPORTANT**

Menu options 30.21 and 30.22 should not be accessed on any node until the failover has completed and replication to the new backup node has been started and only when instructed to do so in this section.

Section Three: Resume Replication	Assi-g ned	Done
1. On the new backup, activate the replication IP (iTERA) (CFGTCP option 1).		
2. On the new backup, hold all scheduled jobs.		

Section Three: Resume Replication	Assi-g ned	Done
<p>3. Optional: If you need to change the *LOCAL RDB entry, do the following:</p> <ol style="list-style-type: none"> <li>Select 30.3</li> <li>Delete *LOCAL (opt 4).</li> <li>Respond to the message with G, then press Enter.</li> <li>Create *LOCAL on the new backup using the value that was in *LOCAL on the old backup.</li> <li>If you are not using passwords on DDM files then in the <i>Check DDM Attributes</i> screen (30.4), verify that the <i>Lowest Authentication Method</i> (OS 6.1 and later) is set to *USRID. (A service authorization entry on all nodes is required.) (In V5R4 verify that the <i>Password Required</i> field is set to *NO.)</li> </ol>		
<p>4. If you changed the system name on the new primary, you must change the system name on the new backup to an unused name (preferably, to the name of the old backup).</p> <p><b>IMPORTANT</b></p> <p>After you have performed the name change on the new backup, contact Support for assistance with completing the failover and restarting replication from the new primary to the new backup. <u>DO NOT CONTINUE WITH THIS CHECKLIST WITHOUT ASSISTANCE FROM SUPPORT!</u> If you did NOT change the system name on the new primary, continue working through this checklist.</p> <hr/> <p>If using APPC connections you must change LCLCPNAME and LCLLOCNAME.</p>		
<p>5. Execute the communication tests (30.7) on the node indicated to verify they are all working in both directions.</p> <ol style="list-style-type: none"> <li>On all target nodes, execute 30.7 <u>to</u> the new primary and all other nodes. For the Scope parameter, specify *FULL.</li> <li>On the new primary, execute 30.7 <u>to</u> all target nodes. For the Scope parameter, specify *FULL.</li> <li>On all nodes, check E2MSGLOG for test results.</li> </ol>		
<p>6. On the new primary, select menu 30.21, then F3=Exit.</p>		
<p>7. On the new primary, select menu 30.22, F14=iTERA IP Interfaces, then F3=Exit.</p>		
<p>8. On the new backup, execute the command EDTRBDAP. Rebuild the access paths (refer to IBM instructions, if needed). Monitor the rebuild of access paths until completed. Ensure production access paths have been rebuilt prior to starting replication.</p>		

Section Three: Resume Replication	Assi-g ned	Done
<p>9. Start the new backup:</p> <ul style="list-style-type: none"> <li>a. <u>On the new primary</u>, select menu option 40.30, F16. Select option 16 on the subfile record for the new backup node, then press Enter. If the IP address of the new backup is active and visible to the new primary, the <i>Replication Startup</i> screen is displayed and the backup system communications will indicate “Ready”.</li> <li>b. Press F16 to start replication. Be patient—this step may take several minutes. If control has not returned after five minutes, investigate the issue with Support.</li> <li>c. When control is returned to your screen, press F5 to refresh. The system roles (i.e., primary and backup) should be correct.</li> </ul>		
<p>10. On all nodes, select menu 3.32. Review the values in the <i>JrnRcv RtnHrs</i> fields. Due to limited disk space on some target systems, the retention hours may have been set lower than on the primary. Since the backup node is now the primary, the retention hours should be reviewed and may need to be adjusted. To change the default retention hours for a journal receiver, select option 2=Change on the desired journal, specify the <i>Journal Receiver Retention</i> value as desired, then press Enter to accept the change.</p>		

Section Three: Resume Replication	Assigned	Done
<p>11. Verify that the process completed correctly.</p> <ul style="list-style-type: none"> <li>a. On the new primary, verify the following: <ul style="list-style-type: none"> <li>– The iTERA subsystem is active. If not, activate it: E2STRSBS SYS (*LOCAL)</li> <li>– The menu indicates it is the primary.</li> </ul> </li> <li>b. On all target nodes, verify the following: <ul style="list-style-type: none"> <li>– The iTERA subsystem is active. If not, activate it: E2STRSBS SYS (*LOCAL)</li> <li>– The main menu indicates the correct system role (e.g., BACKUP1, REPLICATE, etc.).</li> <li>– Select option 3.4. All apply jobs (except for the comm journal) will be in *SEQMAP status. This is normal.</li> <li>– For EACH apply job (except the comm journal apply job), do the following: <ul style="list-style-type: none"> <li>a. Select option 9=Display Jrn Seq. For the first entry listed on the Display Journal Entries screen, verify that the <i>Code</i> is U and the <i>Type</i> is 12.</li> </ul> </li> </ul> </li> </ul> <hr/> <p style="text-align: center;"><b>IMPORTANT</b></p> <p style="text-align: center;"><b>IF THE ENTRY IS NOT U-12, DO NOT CONTINUE.</b> Contact Support immediately to troubleshoot the issue. (If you continue, data integrity issues may occur.)</p> <hr/> <ul style="list-style-type: none"> <li>b. Select option 13=Activate.</li> <li>c. Select option 14=Restart.</li> </ul> <p>12. On the new primary, update the System Monitor twice (1.1, F10). Verify that everything is OK.</p> <hr/> <p style="text-align: center;"><b>NOTE</b></p> <p style="text-align: center;">For three-node environments, execute WRKAUD on the replicate node, then press F5 to refresh the audit compliance status.</p> <hr/>		
<p>13. When replication to the new backup has caught up (verify in the System Monitor screen on the primary) then you may perform a role swap (using the Assure iTERA HA v6.2 Role Swap Checklist) to return to the original primary.</p>		

## Simulate a Failover on the Primary (Failover Test)

### NOTE

Precisely does not encourage simulating the failover process. Instead, we recommend that the role swap be tested at least quarterly, as it uses the same processes as the failover and includes additional steps that validate role swap readiness.

The following items should be executed when performing a test of the failover process.

Simulate a Failover on the Primary (Failover Test)	Assi-g ned	Done
1. Perform the steps on the Assure iTERA HA v6.2 Role Swap Checklist up to—but not including—the step to do the actual role swap.		
2. Perform the steps in the first section of the Failover Checklist. See “Section One: Failover Pre-Checks”.		
3. Inactivate the replication (iTERA) IP addresses on the current primary node. iTERA will not allow a failover if the primary replication IP address is active.		
4. Verify that the replication (iTERA) IP address is not active (check in 30.21). a. From a command line, enter NETSTAT. b. Select option 3, Work with TCP/IP connection status. c. Use option 4 to end any process that is using the replication (iTERA) IP address (primary).		
5. Optional: Inactivate the user IP Address. Verify the user IP address to ensure that it is not active (check in 30.22). a. From a command line, enter NETSTAT. b. Select option 3, Work with TCP/IP connection status. c. Use option 4 to end any process that is using the primary user or iTERA IP (primary).		
6. End the iTERA subsystem on the primary node using the IBM ENDSBS command, as follows:  ENDSBS SBS (SBSNAME) OPTION (*IMMED)		
7. Verify the iTERA subsystem is active on the backup node (E2SBS).		
8. Hold all jobs in the Job Scheduler that may submit during the time you are rolled over (primary).		

<b>Simulate a Failover on the Primary (Failover Test)</b>	<b>Assi-g ned</b>	<b>Done</b>
9. Ensure that you have quiesced the primary system so that your users cannot access it. Verify there are no new journal entries being processed in 40.22.		
10. Complete the following sections: <ul style="list-style-type: none"> <li>– “Section Two: Execute the Failover”</li> <li>– “Section Three: Resume Replication”</li> </ul>		