

R-1

Deployment Platform

Failover Guide

Software Version 5.0

For Windows/UNIX operating systems

March 28, 2011

Copyright © 2011 RepliWeb® Inc., All Rights Reserved

The information in this manual has been compiled with care, but RepliWeb, Inc. makes no warranties as to its accuracy or completeness. The software described herein may be changed or enhanced from time to time. This information does not constitute a commitment or representation by RepliWeb and is subject to change without notice. The software described in this document is furnished under license and may be used and/or copied only in accordance with the terms of this license and the End User License Agreement.

No part of this manual may be reproduced or transmitted, in any form, by any means (electronic, photocopying, recording or otherwise) without the express written consent of RepliWeb, Inc.

Windows, Windows XP and Windows Vista are trademarks of Microsoft Corporation in the US and/or other countries. UNIX is a registered trademark of Bell Laboratories licensed to X/OPEN.

Any other product or company names referred to in this document may be the trademarks of their respective owners.

Please direct correspondence or inquiries to:

RepliWeb, Inc.
6441 Lyons Road
Coconut Creek, Florida 33073
USA

Telephone: (954) 946-2274
Fax: (954) 337-6424

Sales & General Information: info@repliweb.com
Documentation: docs@repliweb.com
Technical Support: <http://support.repliweb.com>
Website: <http://www.repliweb.com>

Table of Contents

1. OVERVIEW.....	1
2. PREREQUISITES	2
3. CONSOLE GUI.....	3

1. Overview

The Failover process enables R-1 administrators to maintain a Backup (standby) server that will be used as the R-1 Center in case the Primary Server is temporarily shut down.

Failover is an important fault tolerance function of mission-critical systems that rely on constant accessibility. The backup server is automatically and transparently to the user being updated with all R-1 Center files, so when the need arises, it can replace the primary server and act as an R-1 center, running all replication and distribution jobs as before.

The Failover process updates the Failover server with the following R-1 information:

- **Scheduled Jobs** - This includes all Continuous Update, Scheduled and On-Demand jobs
- **Templates** – Common and User-specific
- **Containers** – Common and User-specific
- **Configuration Files**
- **Preference Files**

NOTE: The Failover process does NOT copy any content from the primary server to the failover server. This is the responsibility of the user/administrator.

2. Prerequisites

- The Failover and Primary server should be of the same platform (Windows ↔ Windows, Linux ↔ Linux, etc.).

NOTE: The Failover server must be a different machine than the Primary server.

- The Failover Server must contain a complete R-1 installation, using the same R-1 version, installed on **exactly** the same path as the R-1 Center initiating the Failover process. Install at least the Center and Edge components; the Console is optional.

The Failover process will fail if these requirements are not met.

- The Failover process does NOT copy any content from the Primary server to the Failover server. This is the responsibility of the user/administrator.

NOTE: Content in the replication Source Directories is not replicated to the Failover server by the Failover process.

3. Console GUI

Using the R-1, REM or RCM Console, define and activate the Failover process.

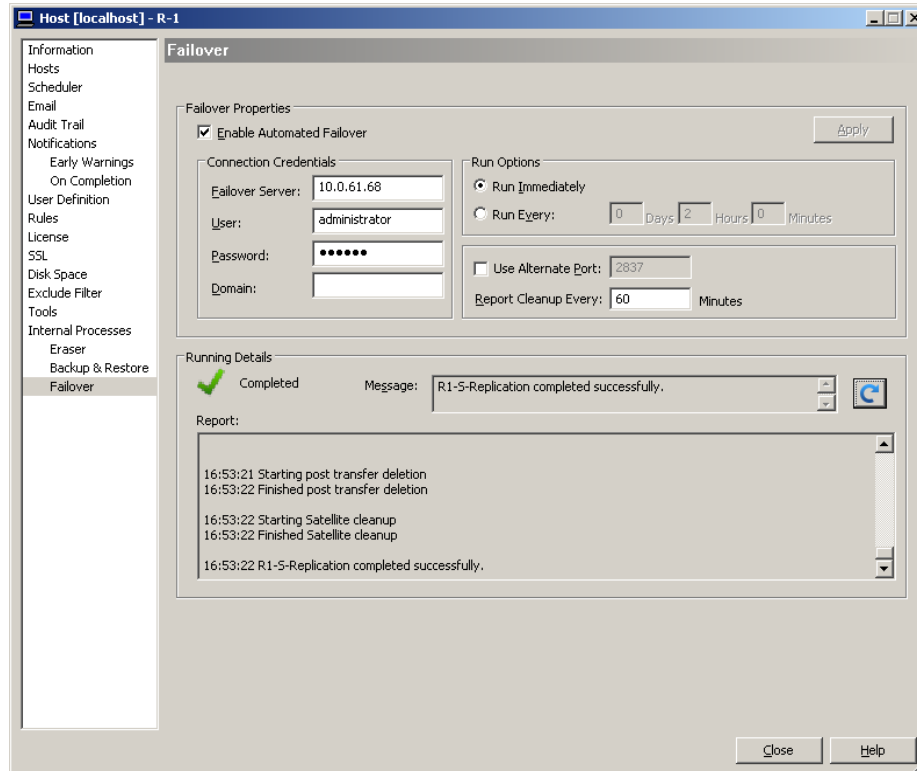


Figure: Internal Process – Failover

To configure the Failover process:

1. Install R-1 on both the primary and failover servers.

Make sure at least the Center and Edge components of the same R-1 version are installed on **exactly** the same path on both servers and that the failover and primary servers are of the same platform (Windows ↔ Windows, Linux ↔ Linux, etc.). The failover process will fail if these requirements are not met.

NOTE: The failover server must be a different machine from the primary server.

2. On the failover server, set the **RDS Scheduler** service to start manually, so it won't start automatically in case of a reboot.
3. On the primary server, open the **Failover** tab, select the **Enable Automated Failover** checkbox to run the Failover process manually and specify the connection credentials and run options. To disable the Failover process, unselect this checkbox.


The option selected will start when you click the **Apply** button.

- a. In the **Connection Credentials** area, specify the credentials that the Failover process will use to connect to the failover server, as follows:
 - i. In the **Failover Server** field, enter the host name or IP address of the R-1 Center being used as the failover server.
 - ii. In the **User** field, enter the user name that the Failover Process will use.
 - iii. In the **Password** field, enter the password for the account specified in the **User** field.
 - iv. In the **Domain** field, enter the domain if the user is part of a domain. If it is not, leave this field blank.

NOTE: Jobs submitted to the failover server should be configured to use the same machine name as specified in the Failover tab.

- b. In the **Run Options** area, specify how to run the Failover process:
 - i. Schedule the Failover process to either run immediately or at fixed intervals, as follow:
 - **Run Immediately** – Run the Failover process immediately. This option only runs once meaning that all the jobs/templates that currently exist on the Center will be transferred to the Failover Center just once.
 - **Run Every** – Run the Failover process at a specific time interval (days/hours/minutes).
 - ii. To change the port being used for the failover job, select the **Use Alternate Port** checkbox and specify the new port number to use.
 - iii. To specify the intervals at which the Failover report will be cleaned, select the **Report Cleanup** checkbox and specify the interval value, in minutes.
4. Click **Apply** to save changed settings. If the Failover is enabled, saved settings take effect immediately.

When the failover process is enabled on the primary server, the **RDS Scheduler** service is automatically stopped on the failover server. The failover process copies scheduled job definitions, Templates, Containers and all configuration files required for the failover server to act as the primary server when needed. The Failover process does NOT copy content from the primary server to the failover server.

5. In the **Running Details** pane, you can view the Failover job's state, completion message and report. To refresh the displayed information, click **Refresh** .

Using the Failover Server

The failover server should get the IP/DNS and NAME of the primary server. This should be done according to corporate policy to allow previously scheduled jobs to continue uninterrupted without resubmitting.

To use the failover server:

1. When the primary server goes down, start the R-1 Scheduler service on the failover server.

Using Windows:

- a. From the Windows **Start** menu select **Control Panel** → **Administrative Tools** → **Services**.
- b. Start the service named **RDS Scheduler**.
- c. Set the RDS Scheduler service to start automatically.

Using UNIX:

Start the repliweb_scheduler: `/usr/bin/repliweb_scheduler`

2. Open R-1 Console and connect to the failover server.
3. In the failover Center's Console GUI, go to **Manage/Center/Scheduler** and select the **Allow local admin to view all jobs** checkbox.

This is necessary because the primary server jobs are running under a different user context.

4. From the Console GUI, re-connect to the failover Center. The failover server now acts as the primary server.

For any additional information, contact us at support.repliweb.com