

opentext™ | Replication

Administrator's Guide

Notices

OpenText Replication Web Administrator's Guide, version 8.5.11, Thursday, April 2, 2026

The only warranties for products and services of Open Text and its affiliates and licensors (“Open Text”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

One or more patents may cover this product. For more information, see <http://www.opentext.com/patents>.

For terms and conditions, see [Terms and Conditions](#).

© 2026 Open Text.

Contents

- OpenText Replication Web Overview 4**
- Installing Replication Web 5**
 - Requirements for Replication Web 6
 - Upgrading Replication Web 7
- Replication Web administration 8**
 - Replication Web interface 9
- Managing users 11**
 - Adding users 12
 - Editing and deleting users 13
- Managing servers 14**
 - Adding servers 16
 - Viewing server details 17
 - Server Properties 18
 - Server Jobs 19
 - Server Event Logs 20
- Managing jobs 21**
 - Creating a files and folders job 23
 - Editing a files and folders job 32
 - Viewing job details 33
 - Starting a Mirroring Process 35
 - Running a Failover job 36

OpenText Replication Web Overview

OpenText Replication Web is a web-based user interface that allows you to add, view, and monitor servers, and create **Files and Folders** protection jobs. You can add servers from different environments to the Replication Web UI, and view these servers and the associated jobs in the interface.

For more information on OpenText Core Replication, see [Replication Documentation Library](#).



OpenText Replication Web supports versions 8.3, 8.4, and 8.5 of OpenText Replication console.

If this is your first time using Replication Web, follow this procedure.

1. **Requirements**—Review the Replication Web requirements to make sure that your environment meets the minimums for deployment. See *Requirements for Replication Web* on page 6.
 2. **Installation**—Install Replication Web. See *Installing Replication Web* on page 5.
 3. **Creating an admin user**—After you have installed Replication Web, you will be ready to create an admin user. See *Replication Web administration* on page 8.
-



[Click here](#) to open a PDF version of this documentation.

Installing Replication Web

1. **Replication Web installation**—Download the latest installation file from the download site.
2. Run the Replication Web setup file.
 1. **License Agreement** page—Review the license agreement. You must accept the license agreement in order to continue with the installation program. Click **Next**.
 2. **Customer Information** page—Accept the default **User Name** or enter your own values. Click **Next**.
 3. **Destination Folder** page—Accept the default folder or click **Change...** to select a different folder. Click **Next**.
 4. **Network Settings** page—Accept the default ports for **HTTP** and **HTTPs** for an SSL connection, or enter your own port numbers. Click **Next**.
 5. **Ready to install the program** page—Click **Install**. When the installation is complete, an icon for OpenText Replication Web appears on your desktop.

If you have already installed Replication Web, you can simply upgrade to the latest version of the product. See *Upgrading Replication Web* on page 7.

Requirements for Replication Web

Your Replication Web environment must meet the following requirements:

- **Web browser**—You require a web browser to access the Replication Web. Preferred browsers are any recent versions of
 - **Google Chrome™**
 - **Mozilla Firefox**
 - **Microsoft® Edge**
- **Replication Web server**—You must have a server where you can install the Replication Web component. The server must meet the following requirements.
 - **Windows server**—Windows 2016 or newer
 - **Memory**—8GB of RAM
 - **Processors**—2 CPUs or more
 - **Disk space**—30 GB

Upgrading Replication Web

You can upgrade your version of Replication Web.

1. **Replication Web upgrade**—Download the latest installation file from the download site.
2. **Run the installation file**—Double click the `RepWeb.exe` file that you saved from the download site.
3. **Upgrade dialogue**—In upgrade confirmation dialogue, click **Yes**. Follow the instructions to upgrade your Replication Web product.

Replication Web administration

After Replication Web is installed, you will create an administrative user.

1. On the Create Admin User page

- **Open Replication Web**—Double click on the Replication Web icon on your desktop. Alternatively, you can open a browser window and type the URL for your server.

`https://IP address:port where`

- **IP address**—refers to the IP address or a server name that is accessible on the network,
- **Port**—refers to the port number specified during installation.
- **Email**—Enter the email address of the administrator.
- **Password**—Enter a password.
- **Confirm password**—Re-enter the password.
- **First name**—Enter the first name of the administrator.
- **Last name**—Enter the last name of the administrator and click **Save**.

After you have created the administrative user, a sign-in page appears, and you can sign in using the admin username and password that you just created.

Once you have signed into OpenText Replication Web, you can make changes to your administrative user.

1. Click the **Gear icon** and select **Edit user account**.
2. On the **Update Credentials** page
 - **Current password**—Enter the current password.
 - **New password**—Enter a new password.
 - **Confirm password**—Re-enter the password and click Save.

Replication Web interface

- **Dashboard**—The dashboard page displays initially each time you log in to Replication Web. The dashboard provides a high-level overview and summarizes the status of your users, servers, and jobs.
 - **Server Status**—This section updates dynamically as server status changes. A filter will automatically be applied to the **Jobs** page showing only the jobs that match the state tile you had selected. Click **Columns** to change the number and name of the columns that are displayed. See *Managing servers* on page 14.
 - **Health**—The health of the server.
 - **Server**—Name of the server. you can filter by server name.
 - **Last Response**—When Replication Web last attempted to contact the server.
 - **Last Successful Response**—When Replication Web last received a successful response from the server.
 - **Job Status**—This section updates dynamically as job status changes. You can click on the tiles in this section to view another table that separates the number of jobs in that state tile by the companies. A filter will automatically be applied to the **Jobs** page showing only the jobs that match the state tile you had selected. Click **Columns** to change the number and name of the columns that are displayed. See *Managing jobs* on page 21.
 - **Health**—Indicates the health of the job. You can filter by health status.
 - **Job**—Name(s) of the jobs in your queue. You can filter by job.
 - **Job Type**—Type(s) of the jobs in your queue. Filter by the type of job that you want to view.
 - **Job Status**—Status of the job(s). You can filter by job status.
- **Gear icon**—If you click the gear icon in the upper right corner, you will find another menu of options.
 - **System Settings**—Refers to the Replication Web service reaching out to the servers that were added to the console to obtain information about the servers and their associated jobs.
 - **Interval (in seconds)**—Enter the number of seconds between polling cycles. The default is 5 seconds.
 - **Maximum servers**—Refers to the number of servers that are polled simultaneously in the polling cycle. In environments in which there are a large number of servers and the Replication Web server has sufficient RAM and CPUs allotted, this number can be increased for better performance with a maximum of 100. The default is 10.
 - **Idle minutes**—If noon is interacting with a OpenText Replication Web browser session for the specified minutes, the OpenText Replication Web service will stop polling servers for updates to the their associated job.
 - **Edit User Account**—You can update your user credentials on this page. Enter your current password, and then a new password twice, and click **Save**.
 - **Sign out**—This option immediately logs you out of the Replication Web interface. If there is a period of inactivity, you will automatically be logged out. You will be given a warning

90 seconds before the system signs you out. Any jobs that you have started will continue to run, even when you are logged out.

- **Documentation**—This option will open a new browser window to the Replication Web User's Guide.
- **About**—This option provides information about the Replication Web product, including version number.

Managing users

The **Users** page provides high-level information and controls for your users.

If you were assigned the Administrator role for your user account, an additional tab in the user interface is available. The **Users** tab allows you to create users and assign a role to them. See *Adding users* on page 12.

- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Name**—Name of the user.
- **Role**—Role of the user, **Administrator** or **User**.
- **Email**—Email address of the user.
- **Actions**—In the table overflow menu on the right of a table row, you can select the following actions.
 - **Edit**—Select this option to update the user account. Only administrators can perform this action. See *Editing and deleting users* on page 13.
 - **Delete**—Select this option to delete the user. Only administrators can perform this action. See *Editing and deleting users* on page 13.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.

Adding users

As administrator, you can add users. Click the Users tab on the top navigation pane of the Users page.

- **Add user** page—Click **Add user**.
 1. **Email**—Enter the email address of the user.
 2. **Password**—Enter a password.
 3. **Confirm password**—Re-enter the password.
 4. **Role**—Select a role for the user, **Administrator** or **User**.
 - **Administrator**—The administrator can manage servers, jobs, and other users.
 - **User**—The user can manage servers, and jobs; users cannot manage other users.
 5. **First name**—Enter the first name of the user.
 6. **Last name**—Enter the last name of the user and click **Save**.

The new user appears in the list of servers table on the **Users** page.

Editing and deleting users

As administrator, you can edit and delete users. Click the **Users** tab on the top navigation pane. Alternatively, in the table overflow menu on the right of a table row of the user you want to change, select **Edit** or **Delete**.

1. **Edit user**—Right click the **More options** button beside the user that you want to edit.
 1. **Email**—Enter a new email address.
 2. **Role**—Change the role of the user.
 - **Administrator**—The administrator can manage servers, jobs, and users.
 - **User**—The user can manage servers, and jobs; users cannot manage other users.
 3. **First name**—Enter the first name of the user.
 4. **Last name**—Enter the last name of the user, and click **Save**.
2. **Delete user**—Right click the **More options** button beside the user that you want to delete.
 1. **Delete**—Select **Delete**.
 2. **Confirm delete** dialogue—Select **Delete** or **Cancel**.



You are not able to delete your own user.

As user, you can edit the **First name** and **Last name** of your user. In the table overflow menu on the right of a table row of the user you want to change, select **Edit**.

1. **Email**—Enter a new email address.
2. **First name**—Enter the first name of the user.
3. **Last name**—Enter the last name of the user, and click **Save**.

Managing servers

The **Servers** page provides high-level information and controls for the servers that you have added to your Replication Web **Servers** page.

The following controls and statuses are available on the **Servers** page.

- **Add server**—Click this button to add servers. For information about adding servers, see: *Adding servers* on page 16.
- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Expand and collapse**—Click the right arrow to expand a server to see high-level details. Click the down arrow to collapse the details.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.
- **Health**—Health indicates, by color and description, the health of the server
 - **Green**—A green circle indicates a healthy status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, OpenText Replication Web Service is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Gray**—A gray circle indicates an unprotected server.
 - **Black**—A black circle indicates that the status is unknown.
- **Server**—The **Server** field lists all the servers that you have added from your Replication Console(s). The server filter allows you to filter by name. The up and down arrow allows you change the order of the servers shown, alphabetically or reverse alphabetically. Click the link on the name of the server and the View Details page opens.
 - **Properties** page—Displays the server properties.
 - **Operating system**—Type of OS that is installed on the server.
 - **Status**—Health of the server.
 - **Activity**—Online or offline.
 - **Polling address**—IP address used for polling.
 - **IP address**—IP address of the server.
 - **Version**—Version of the Replication Console installed on the server.
 - **Licensing**—License information.
 - **Jobs** page—Displays the job details.

- **Server Status**—The **Server Status** field indicates the status of each of the servers in your list.
 - **Online**—The server is online and available.
 - **Cannot connect to service**—The server is not available.
- **Platform**—Type of OS that is installed on the server.
 - **Windows**—Servers on a Windows platform.
 - **Linux**—Servers on a Linux platform.
 - **Unknown**—Servers that are in an error state, where the platform is unknown.
- **Product**—Name of the Replication Console product installed on the server.
 - **Doubletake Availability**
 - **Doubletake Migrate**
 - **Doubletake target**
- **Product Version**—Version of the Replication Console product installed on the server.
- **Licensing Status**—Type of license and expiration date.
- **Actions**—In the table overflow menu on the right of a table row, you can select the following actions.
 - **Update credentials**>—Select this option to update the username or password for the server. On the Update credentials page:
 - **Username**—Enter a new username for this server.
 - **Password**—Enter a new password for this server, and click **Confirm** or **Cancel**.
 - **Remove server**—Click **Remove** to delete this server. On the **Remove Server** page, and click **Confirm** or **Cancel**.
- **Expand**—Click **Expand** on the right of a table row, and you can see the server details.
 - **Operating system version**—The operating system of the server. This field will not be displayed if the console cannot connect to OpenText Availability on the server.
 - **Polling address**—IP address used for polling.
 - **IP address**—IP address of the server.
 - **Serial number**—Serial number of the server.
 - **Additional information**—License information, for example.

Adding servers

You can add servers to Replication Web from any number of Replication Consoles.

1. On the **Servers** page, click **Add server**.
 - **Server**—Enter the IP address of the server in your Replication Console that you want to add to Replication Web.
 - **Username**—Enter the username of the administrator of that server.
 - **Password**—Enter the password.
2. Click **Save** to add one server. Click **Save+** to save this server and add another server.

The new server appears in the list of servers table on the **Servers** page.

Viewing server details

On the **Servers** page, in the **Server**, click the link for the server that you want to view the details for. The **View Details** page allows you to view details about that particular server. The server details vary depending on the type of server or appliance you are viewing.



View server details is not available for Linux servers.

Server Properties

- **Operating system version:**—The operating system of the server. This field will not be displayed if the console cannot connect to OpenText Availability on the server.
- **Health**—Indicates what state the server is in: **Information**, **Warning**, **Error** or **Unknown**.
- **Status**—Status of the licence for this server.
- **Polling address**—IP address used for polling.
- **IP address**—IP address of the server.
- **Version**—Product version information.
- **Licensing**—Type of licence licence for this server.
- **Additional information**—License information, for example.

Server Jobs

- **Job Name**—A list of any jobs from this server.
- **Source Server**—Name of the source server for this job.
- **Target Server**—Name of the target server for this job.
- **Job Type**—Type of job.

Server Event Logs

- **Type**—This column in the table indicates if the type of log is **Information**, **Warning**, **Error** or **Unknown**.
You can filter the column by **Type**.
- **Time**—This column indicates the date and time when the message was logged.
You can filter the column by **Time**.
- **Source**—This column if the message is from the Double-Take log or the Management Service log.
You can filter the column by **Source**, either **Double-Take** log or **Management Service** log.
- **Description**—This column displays the actual message that was logged.
You can filter the column by **Time**.

Managing jobs

The **Jobs** page provides high-level information and controls for the jobs that you have created in your Replication Console and have added to your Replication Web Jobs page.

The following controls are available on the Jobs page.

- **Sort**—Sort the table by clicking a column heading. When the arrow is pointing up, the table is sorted by that column in ascending order. When the arrow is pointing down, the table is sorted by that column in descending order.
- **Filter**—Text entered in a filter box and selected from a filter drop-down list will narrow the list displayed to only those rows that contain the search text and selected item.
- **Select All and Clear All**—Click the checkbox in the column heading to toggle between selecting all items on that page of the table or clearing all selections on that page of the table. This option will not apply to items on a page that are hidden by a search filter..
- **Expand and collapse**—Click the right arrow to expand a job to see high-level details. Click the down arrow to collapse the details.
- **Table paging**—At the bottom of the table you will see the row numbers you are currently viewing and the total number of table rows. Paging buttons allow you to move between pages of the table. The single arrow buttons move forward or backward one page. The double arrow buttons move to the first or last page.
- **Health**—Health indicates, by color and description, the health of the server
 - **Green**—A green circle indicates a healthy status.
 - **Yellow**—A yellow circle indicates a pending or warning status. Generally, OpenText Replication Web Service is working, waiting on a pending process, or attempting to resolve the warning state.
 - **Red**—A red circle indicates an error status. You will need to investigate and resolve the error.
 - **Gray**—A gray circle indicates an unprotected server.
 - **Black**—A black circle indicates the status is unknown.
- **Job**—The name of the job.
- **Source**—The name of the source. This could be a name or IP address of the server.
- **Target**—The name of the target. This could be a name or IP address of the server.
- **Job Type**
 - **Files and folders protection**
 - **Full server protection**
 - **Full server to ESX protection**
 - **SQL protection**
- **Job Status**
 - **Protection**—Data is being replicated to the target.
 - **Ready**—There is no data to replicate.
 - **Pending**—Replication is pending.
 - **Stopped**—Replication has been stopped.
 - **Out of Memory**—Replication memory has been exhausted.

- **Failed**—The Doubletake service is not receiving replication operations from the Replication driver.
- **Unknown**—The console cannot determine the status.
- **Mirror Status**
 - **Calculating**—The amount of data to be mirrored is being calculated.
 - **In Progress**—Data is currently being mirrored.
 - **Waiting**—Mirroring is complete, but data is still being written to the target.
 - **Idle**—Data is not being mirrored.
 - **Paused**—Mirroring has been paused.
 - **Stopped**—Mirroring has been stopped.
 - **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
 - **Verifying**—Data is being verified between the source and target.
 - **Restoring**—Data is being restored from the target to the source.
 - **Unknown**—The console cannot determine the status.
- **Operating System**—The job type operating system
- **Actions**—In the table overflow menu on the right of a table row, you can select the following actions.
 - **Failover**—Select this option to begin live or recovery point failover.
 - **Stop**—Select this option to stop the job. Data changes will not queue on the source (if you are protecting) or replica virtual machine (if you are restoring). Data synchronization will restart from the beginning when the job is restarted.
 - **Pause**—Select this option to pause the job. Data changes will queue on the source (if you are protecting) or replica virtual machine (if you are restoring). The changes will be transmitted once the job is resumed.
 - **Start**—Select this option to start a job or to resume a job that you have paused.
 - **Delete**—Select this option to delete the job. This option may not be available depending on the state of your job.
 - **Mirror**—Select this option to start a mirroring process.
- **Expand**—Click **Expand** on the right of a table row, and you can view the job details. See *Viewing job details* on page 33.

Creating a files and folders job

Use these instructions to create a files and folders job.

1. **Servers** page—On the **Servers** page, hover over the server that you want to protect, and in the table overflow menu on the right of the row for that server, select **Protect**.
2. **Job Type** page—In the **Select job type** box, select **Files and Folders**, and click **Next**.
3. **Target** page—On the **Target** page, select the target server. This is the server that will store the replica data from the source. Click **Next**.
4. **Workload** page—On the **Workload** page, select type of workload that you want to protect.
 - a. **Items**—Displays files and shares with data that you can protect. Select the files and folders that you want to protect.
 - b. **Advanced**—Provides granular control over what files and folders to include or exclude depending on the rules chosen. Expand **Advanced Rules** to add extra criteria for your protection job:
 - A. Expand the file to select individual file(s) that you want to protect.
 - B. **Path**—The sub-folder that you selected appears in the **Path** box.
 - C. **Rule**—Select the rule that you want to be applied to the files and folders that you selected:
 - Include > Recursive
 - Include > Not recursive
 - Exclude > Recursive
 - Exclude > Not recursive
 - D. Confirm your selection, or click **Delete**, and then click **Next**.
5. **Job Options** page—On the **Job Options** page, select the options for this job. Once you have selected all the job options, click **Validate**.
 - a. **Job name**—Specify a unique name for your job.
 - b. **Failover monitor**—Specify your failover monitor options:

Failover monitor details

- A. **Total time to failure**—Specify the number and the unit. for example, 5 Hours/Minutes/Seconds from the dropdown in units. This determines how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.
- B. **Monitor interval**—Specify the number and the unit. for example, 5 Hours/Minutes/Seconds from the dropdown in units. This determines how often

the longer the target will try to contact the source.

Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- C. **Network monitoring**—Select each **Source IP Address** that you want the target to monitor.
- i. **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - ii. **Monitoring Method**—This option determines the type of failover monitoring used. The **Network service** option tests source availability using an ICMP ping to confirm that the route is active. The **Management service** option opens a socket connection to confirm that the Double-Take service is active. If you are using a NAT environment, **Management service** is the only available option.
 - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
 - **Management service (Windows Only)**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
 - **Network and management services (Windows Only)**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. If either monitoring method fails, failover will be triggered.
 - iii. **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- c. **Failover options**—Specify your failover options:

Failover option details

- A. **Wait for user to initiate failover**—When this field is cleared, failover will occur automatically when the failover conditions is met. When checked, user will need to trigger failover. The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.
- B. **Failover shares**—Select this option to failover shares to the target. Only the shares that you selected on the **Workload** page will be protected and failed over.



Share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target.

If you are failing over Windows shares but your source and target do not have the same drive letters, you must use the **All to One** selection under **Path Mapping** when establishing your job. Otherwise, the shares will not be created on the target during failover.

Windows share information is automatically updated on the target every 30 seconds.

- d. **Failover identity**—Specify your failover identity options (Windows only):

Failover identity details



If you have selected to failover shares under the **Failover Options** sections, the source NetBIOS name will automatically be failed over so that the shares can be accessed after failover.

If you want to disable the ability to failover, you must select **Retain target network configuration**. Additionally, in the **Failover Options** section, you must disable **Failover shares**, **Failover host name**, and **Failback host name**, and do not specify anything for the failover and failback scripts. If you select any of these options or choose **Apply source network configuration to the target**, the ability to failover will not be disabled.

- A. **Apply source network configuration to the target**—If you select this option, you can configure the source IP addresses to failover to the target. If your target is on the same subnet as the source (typical of a LAN environment), you should select this option. Do not select this option if you are using a NAT environment that has a different subnet on the other side of the router.
 - i. **Failover server name**—Select this option to failover the server name to the target. Replication Web checks the hosts file and uses the first name there. If there is no hosts file, Replication Web will use the first name in DNS. (Keep in mind, the first name in DNS may not always be the same each time the

DNS server is rebooted.) Lastly, if there is no DNS server, Replication Web will use the failover monitor name created by the OpenText Replication Console. If you have selected to failover shares under the Failover Options sections, the server name will automatically be failed over so that the shares can be accessed after failover.

- ii. **Add these addresses to the selected target adapter after failover**— Select which IP addresses you want to failover and select the **Target Network Adapter** that will assume the IP address during failover.



If you have inserted your source server into the console using a reserved IP address, do not select the reserved IP address for failover.

If you configured failover to be triggered when all monitored IP addresses fail and are failing over more IP addresses than you are monitoring, you may have IP address conflicts after failover. For example, if you monitor two out of three addresses, and those two addresses fail but the third one does not, and you failover all three IP addresses, the third address that did not fail may exist on both the source and the target, depending on the cause of the failure. Therefore, when a source is failing over more IP addresses than are being monitored, there is a risk of an IP address conflict.

- B. **Retain target network configuration**— If you select this option, the target will retain all of its original IP addresses. If your target is on a different subnet (typical of a WAN or NAT environment), you should select this option.
 - **Failover server name**—Select this option if you want to failover the NetBIOS name.
 - **Update DNS server**—Specify if you want OpenText Availability to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the job requirements for updating DNS.



DNS updates are not available for Server Core servers or source servers that are in a workgroup.

Make sure that port 53 is open for DNS protocol from the target to the DNS servers so the target can discover the source DNS records.

Expand the DNS Options section to configure how the updates will be made. The DNS information will be discovered and displayed. If your servers are in a workgroup, you must provide the DNS credentials before the DNS information can be discovered and displayed.

- **Change**—If necessary, click this button and specify a user that has privileges to access and modify DNS records. The account must be a

member of the DnsAdmins group for the domain, and must have full control permissions on the source's A (host) and PTR (reverse lookup) records. These permissions are not included by default in the DnsAdmins group.

- **Remove**—If there are any DNS servers in the list that you do not want to update, highlight them and click **Remove**.
- **Update these source DNS entries with the corresponding target IP address**—For each IP address on the source, specify what address you want DNS to use after failover.
- **Update TTL**—Specify the length of time, in seconds, for the time to live value for all modified DNS A records. Ideally, you should specify 300 seconds (5 minutes) or less.



DNS updates will be disabled if the target server cannot communicate with both the source and target DNS servers.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non-Microsoft Active Directory integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

- e. **Mirror and Orphaned Files**—Specify your Mirror and Orphaned Files options:

Mirror and Orphaned Files

Mirror Options—Select a comparison method and whether to mirror the entire file or only the bytes that differ in each file.

- A. **Mirror Options**—Select a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
- **Do not compare files. Send the entire file.**—Replication Web will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This option requires no time for comparison, but the mirror time can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the entire file.**—Replication Web will compare file attributes and will mirror those files that have different attributes, sending the entire file. This option is the fastest comparison method, but the mirror time can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions

of small files that are mostly static and not changing. You may also use this option if configuration management policies require sending the entire file.

- **Compare file attributes. Send the attributes and bytes that differ.**—Replication Web will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror speed. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
- **Compare file attributes and data. Send the attributes and bytes that differ. (Windows Only)**—Replication Web will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror speed.



If a file is small enough that mirroring the entire file is faster than comparing it and then mirroring it, Replication Web will automatically mirror the entire file.

B. General Options—Select your general mirroring options.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target.



Orphaned file configuration is a per target configuration. All jobs to the same target will have the same orphaned file configuration.

If **Delete orphaned files** is enabled, carefully review any replication rules that use wildcard definitions. If you have specified wildcards to be excluded from protection, files matching those wildcards will also be excluded from orphaned file processing and will not be deleted from the target. However, if you have specified wildcards to be included in your protection, those files that fall outside the wildcard inclusion rule will be considered orphaned files and will be deleted from the target.

f. Network Route—Specify your network route options:

Network Route

- A. Send data to this target IP address**—By default, Replication Web will select an IP address on the target for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Replication Web traffic. For example, you can separate regular network traffic and Replication Web traffic on a machine with multiple IP addresses. You can also select or manually enter a public IP address (which is the

public IP address of the server's router) if you are using a NAT environment. If you enter a public IP addresses, you will see additional fields allowing you to disable the default communication ports and specify other port numbers to use, allowing the target to communicate through a router. The **Management Service port** is used to persist the source share configuration when shares are being protected. The **Replication Service port** is used for data transmission.



If you change the IP address on the target which is used for the target route, you will be unable to edit the job. If you must make any modifications to the job, it will have to be deleted and re-created.

- B. **Use default source route**—Select this option to disable the drop-down list that allows you to select the route from the target server. When this option is enabled, the default route will automatically be used.
- g. **Path mapping**—Specify your path mapping options:

Path Mapping

Path Mapping—Specify the location on the target where the replica of the source data will be stored. By default, the replica source data will be stored in the same directory structure on the target. Make sure you update this location if you are protecting multiple sources or jobs to the same target. You have two pre-defined locations as well as a custom option that allows you to set your path.

- **All To one**—Click this button to set the mapping so that the replica source data will be stored on a single volume on the target. The pre-defined path is:
 - **Windows**—\source_name\volume_name
 - **Linux**—\source_name\volume_name
 - **One to one**—Click this button to set the mapping so that the replica source data will be stored in the same directory structure on the target.
 - **Windows**—For example, C:\data and D:\files on the source will be stored in C:\dataC:\data and D:\files, respectively, on the target.
 - **Linux—Linux**—For example, /data and /home on the source will be stored in /data and /home, respectively, on the target.
- h. **Compression**—Specify your compression options:

Compression

Compression—Specify the level of compression that you want to apply. To help reduce the amount of bandwidth needed to transmit Replication data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. You can set the level to suit your needs:

- **Off**—No compression
- **Low**
- **Medium**
- **High**

Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether you should enable compression.

- If data is being queued on the source at any time, consider enabling compression.
- If the server CPU utilization is averaging over 85%, be cautious about enabling compression.
- The higher the level of compression, the higher the CPU utilization will be.
- Do not enable compression if most of the data is inherently compressed. Many image (.jpg, .gif) and media (.wmv, .mp3, .mpg) files, for example, are already compressed. Some images files, such as .bmp and .tif, are decompressed, so enabling compression would be beneficial for those types.
- Compression may improve performance even in high-bandwidth environments.
- Do not enable compression in conjunction with a WAN Accelerator. Use one or the other to compress Replication Web data.



All jobs from a single source connected to the same IP address on a target will share the same compression configuration.

6. **Summary page**—On the **Summary** page, the results of the validation step appear.
 - a. **Summary page**—Review the job options that you have selected.
 - **Errors**—designated by a white X inside a red circle.
 - **Warnings**—designated by a black exclamation point (!) inside a yellow triangle.
 - **Success**—designated by a white checkmark inside a green circle.
Errors are sorted at the top of the list, then warnings, then success messages.
 - b. Click any of the validation items to see details. You must correct any errors before you can continue.
 - If you receive a path transformation error during job validation indicating a volume does not exist on the target server, even though there is no corresponding data being protected on the source, you must manually modify your replication rules. Go back to the **Workload** page and under the **Replication Rules**, locate the volume from the error message. Remove any rules associated with that volume. Complete the rest of the workflow and the validation should pass.
 - Before a job is created, the results of the validation checks are logged in the Doubletake Management Service log on the target. After a job is created, the results of the validation checks are logged to the job log. See the Replication Web Reference Guide for details on the various Replication Web log files.

- Replication Web validates that your source and target are compatible. The Summary page displays your options and validation items.
- c. When you have confirmed the details on the **Summary** page, click **Finish**, and you will automatically be taken to the **Jobs** page.



For information about the actions that you can take on this job, like pause and resume, see [Job actions](#).

Editing a files and folders job

You can edit a job on the **Jobs** page in Replication Web.

Use these instructions to edit your job.

1. **Jobs** page—Open the Jobs page in Replication Web.
Select the job that you want to edit and, in the table overflow menu on the right of a table row of the job that you want to edit, select **Edit** in the toolbar. (You will not be able to edit a job if you have removed the source of that job from your Replication Web session.)
2. **Job options** page—You will see the same options available for your files and folders job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job.
 - **Job options**—*Creating a files and folders job* on page 23.
3. **Jobs** page—Once you have edited the job options, click **Validate**.
Summary page—The results of the validation step appear here.
 - **Errors**—designated by a white X inside a red circle.
 - **Warnings**—designated by a black exclamation point (!) inside a yellow triangle.
 - **Success**—designated by a white checkmark inside a green circle.Errors are sorted at the top of the list, then warnings, then success messages.
Click any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Replication Web correct the problem for you. For those errors that Replication Web cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors. Click the **Common Job Validation Warnings and Errors** link to open a web browser and view solutions to common validation warnings and errors.
4. **Finish job editing**—Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing job details

On the **Jobs** page, click the name of the job to see the job details and statistics.



You will see details while a job is in the process of being deleted, but once it is deleted, the details will no longer be available.

- **Name**—Name of the job.
- **Target data state**
 - **OK**—The data on the target is in a good state.
 - **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
 - **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
 - **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
 - **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
 - **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
 - **Not Loaded**—Replication Console target functionality is not loaded on the target server. This may be caused by a license key error.
 - **Not Ready**—The replication agent drivers have not yet completed loading on the target.
 - **Unknown**—The console cannot determine the status.
- **Mirror remaining**—This is the amount of data remaining to be sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
- **Mirror skipped**—This is the amount of data that has been skipped because the data is not different on the source servers and target appliances (when protecting) or on the replica virtual machines and failback targets (when restoring).
- **Replication queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica virtual machines (when restoring).
- **Replication queue**—This is the amount of disk space being used to queue replication data on the source servers (when protecting) or replica virtual machines (when restoring).
- **Disk queue**—This is the amount of disk space being used to queue data on the source servers (when protecting) or on the replica virtual machines (when restoring).
- **Recovery point latency**—This is the longest length of time replication is behind on any one target appliance compared to the source server they are protecting or on any one failback target compared to the replica virtual machine they are restoring from. This is the longest time

period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include synchronization data. If you are synchronizing and failover, the data on the target appliance will be at least as far behind as the replication point latency. It could potentially be further behind depending on the circumstances of the synchronization. If synchronization is idle and you failover, the data will only be as far behind as the replication point latency time.

- **Bytes sent**—This is the total amount of data (in bytes) sent from the source servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring).
- **Bytes sent (compressed)**—This is the total amount of compressed data sent from the sources servers to the target appliances (when protecting) or from the replica virtual machines to the failback targets (when restoring). If compression is disabled, this statistic will be the same as bytes sent.
- **Connected since**—The source server date and time indicating when the current job was started. When this field is blank, it indicates that a TCP/IP socket is not present, or that the job is waiting on transmit options, or if the transmission has been stopped.
- **Transmit mode**
 - **Active**—Data is being transmitted to the target.
 - **Paused**—Data transmission has been paused.
 - **Scheduled**—Data transmission is waiting on schedule criteria.
 - **Stopped**—Data is not being transmitted to the target.
 - **Error**—There is a transmission error.
 - **Unknown**—The console cannot determine the status.
- **Recent activity**—There are many different Activity messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction.
- **Additional information**—Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Starting a Mirroring Process

Use these instructions to start a mirroring process.

1. **Jobs** page—In the table overflow menu on the right of a table row, select the **Mirror** and then **Start**.
2. **Mirror** page—Right click the More options button beside the job that you want to run a mirroring process for.
3. **Mirror options** page—Select a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Replication Console will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This option requires no time for comparison, but the mirror time can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the entire file.**—Replication Console will compare file attributes and will mirror those files that have different attributes, sending the entire file. This option is the fastest comparison method, but the mirror time can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are mostly static and not changing. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Replication Console will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror speed. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Replication Console will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror speed.
4. Click **Start**.



If a file is small enough that mirroring the entire file is faster than comparing it and then mirroring it, Replication Console will automatically mirror the entire file.

For information about the actions that you can take on this job, like pause and resume, see [Job actions](#).

Running a Failover job

When a failover condition has been met, failover will be triggered automatically if you disabled the **Wait for user to initiate failover** option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The target will stand in for the source by assuming the network identity of the failed source. User and application requests destined for the source server or its IP addresses are routed to the target.
 - **Perform test failover**—This option is not available for files and folders jobs.
 - **Failover to a snapshot**—This option is not applicable to files and folders jobs.
2. Select how you want to handle the data in the target queue.
 - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
3. When you are ready to begin failover, click **Failover**.



For information about the actions that you can take on this job, like pause and resume, see [Job actions](#).

Notices

OpenText Replication Web Administrator's Guide, version 8.5.11, Thursday, April 2, 2026

The only warranties for products and services of Open Text and its affiliates and licensors (“Open Text”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall

not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

One or more patents may cover this product. For more information, see <http://www.opentext.com/patents>.

For terms and conditions, see [Terms and Conditions](#).

© 2026 Open Text.