

opentext™ | Replication

Administrator's Guide

Notices

Notices

OpenText Replication Web Administrator's Guide, version 8.5.8, Thursday, May 1, 2025

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](#).

© 2025 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
- Managing jobs 17**
 - Viewing job details 19
 - Start Mirroring Process 21

OpenText Replication Web Overview

OpenText Replication Web is a web-based user interface where you can view and monitor servers and jobs that are protected in the OpenText Replication Console. You can add servers from various environments to the Replication Web UI and view these servers and the associated jobs in the interface.

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



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

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**—Once 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.
 - **License Agreement** page—Review the license agreement. You must accept the license agreement in order to continue with the installation program. Click **Next**.
 - **Customer Information** page—Accept the default **User Name** or enter your own values. Click **Next**.
 - **Destination Folder** page—Accept the default folder or click **Change>** to select a different folder. Click **Next**.
 - **Network Settings** page—Accept the default ports for **HTTP** and **HTTPs** for an SSL connection, or enter your own port numbers. Click **Next**.
 - **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 has been 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 17.
 - **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 where 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 number of servers that you want to poll. 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**.

- **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**.
- **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 your servers that you have created in your Replication Console and 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.
 - **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.

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.
 - **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 19.

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.

Start Mirroring Process

On the **Jobs** page, you can start a mirroring process. In the table overflow menu on the right of a table row, select the **Mirror** and then **Start**.

Mirror Options—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file. Intellimirror changes are to compare attributes to determine which files to compare instead of unconditionally comparing every file. It monitors file activity/open files and falls back to checksum for those files only.

1. **Mirror** page—Right click the More options button beside the job that you want to run a mirroring process for.
2. **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.
3. 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.

Notices

OpenText Replication Web Administrator's Guide, version 8.5.8, Thursday, May 1, 2025

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](#).

© 2025 Open Text.