



To prepare a replication job

1. Click REPLICATE on the toolbar.
2. In the Network list on the Source tab, choose the server or server object to replicate.
3. Uncheck **Include deleted files** if you *don't* want to replicate previously deleted files.
4. On the Schedule tab, specify whether to start the job now or later.

Note:

If you select Scheduled, specify on which day and at what time to replicate data. You can also specify when to reuse weekly tapes.

5. In the **Target drives** list on the Destination tab, choose the device on which to replicate data.
6. Click Start.
7. Give the replication job a name and description, then click OK.



To replicate data now

1. Click REPLICATE on the toolbar.
2. On the Source tab, click the server (or server object) you want to replicate.
3. Uncheck **Include deleted files** if you *don't* want to replicate previously deleted files.
4. Select the Destination tab.
5. In the **Target drives** list, choose the device on which to replicate data.
6. Click Start.
7. Give the replication job a name and description, then click OK.

Note:

The Status window opens to show the job's progress.



To schedule a job to run later

1. Click REPLICATE on the toolbar.
2. On the Source tab, click the server (or server object) you want to replicate.
3. Uncheck **Include deleted files** if you *don't* want to replicate previously deleted files.
4. On the Schedule tab, click **Scheduled**.
5. Specify on which day(s) to replicate this job, and at what time.
6. In the **Target drives** list on the Destination tab, choose the device on which to replicate data.
7. Click Submit.
8. Give the replication job a name and description, then click OK.

Note:

The Status window opens to show the job's progress.



To choose a device on which to replicate data

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. In the **Target drives** list on the Destination tab, choose the device on which to replicate data.
4. Start or Submit the job.



To eject a tape after replication

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. On the Options tab, check **Eject tape after replication.**
4. Start or Submit the job.



To name a replication job

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. Start or Submit the job.
4. In the Start/Submit Job dialog box, give the replication job a name and description, then click OK.



To overwrite data on a tape

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. On the Options tab, select **Overwrite**.
4. Start or Submit the job.

Note:

The Overwrite field is selected by default.



To replicate data while a volume is in use

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. On the Options tab, select **7/24 Live Server Replication**.

Note:

This field is checked by default.

4. Start or Submit the job.



To specify when to reuse weekly tapes

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. In **Reuse weekly tapes after** on the Schedule tab, specify after how many weeks to reuse the tape.
4. Click Submit.



To verify data can be read from tape

1. Click REPLICATE on the toolbar.
2. [Prepare a replication job.](#)
3. On the Options tab, check **Verify replication.**
4. Start or Submit the job.



To bring up a server from a disaster

Remote server

1. Make sure the server you want to bring up is on the network.
2. Start the server with its Disaster Recovery disks.
3. [Recover the server's data](#) from tape.

Host Replica Server

1. Insert your latest full server replication tape in the device.
2. Start the server with its Disaster Recovery disks.
3. Follow the online instructions.



To import a tape

1. Insert the tape into the Host Replica Server's tape device.
2. Click UTILITIES on the toolbar.
3. On the Device tab, click Import Tape.

Note:

Replica reads the latest database information on the tape and imports it into the current Replica database.



To recover files or directories

1. Click RECOVER on the toolbar.
2. On the Source tab, check **Selective file recovery (Mount tape)**.
3. Select the server (or server object) to recover.
4. Choose the tape from the **Tape library**.
5. In the **Mount tape** field on the Destination tab, type the volume name you want to mount this tape as (or accept the default volume name Replica assigns).
6. Click Mount Tape.
7. Switch to File Manager (or other utility) and [map a drive letter](#) to the mounted tape volume and access data.



To recover data from tape

1. Click RECOVER on the toolbar.
2. Insert the tape in the Host Replica Server's tape device.
3. On the Source tab, select the server (or server object) to recover.
4. Choose the tape from the **Tape library**.
5. On the Destination tab, select the destination location.
6. Click Start.



To recover remote Replica servers

1. Start the downed server with its Disaster Recovery disks.
2. On the Administrative PC, run Replica.
3. Click RECOVER on the toolbar.
4. On the Source tab, select the server you want to recover.
5. In the **Host Replica Server** field, make sure the Host Replica Server is selected.
6. Choose the tape from the **Tape library**.
7. On the Destination tab select the server to recover. You may need to log into it.
8. Click Start.



To recover Replica databases

1. See *Console Commands* in the Replica User's Guide.



To recover the Host Replica Server

1. Insert the latest full server replication tape in the device.
2. Start the Host Replica Server with its Disaster Recovery disks.
3. Follow the online instructions.



To specify what to recover

1. Click RECOVER on the toolbar.
2. On the Source tab, click a server (or server object) to select it.

Note:

With Replica, you can recover one server (or server object) at a time.



To make Disaster Recovery disks

Note:

You will need at least two formatted (or unformatted) floppy disks.

1. Insert the first disk in the Administrative PC's drive.
2. Click UTILITIES on the toolbar.
3. On the Disaster Recovery tab, select the server.
4. In **Make disks on Drive**, select the appropriate drive letter and click Create Disks.
5. When done, label the disks and store them someplace safe.



To recover remote Replica servers

1. Start the downed server with its Disaster Recovery disks.
2. On the Administrative PC, run Replica.
3. Click RECOVER on the toolbar.
4. On the Source tab, select the server you want to recover.
5. In the **Host Replica Server** field, make sure the Host Replica Server is selected.
6. Choose the tape from the **Tape library**.
7. On the Destination tab select the server to recover. You may need to log into it.
8. Click Start.



To recover the Host Replica Server

1. Insert the latest full server replication tape in the device.
2. Start the Host Replica Server with its Disaster Recovery disks
3. Follow the online instructions.



To delete a tape from the database

1. Click UTILITIES on the toolbar.
2. On the Tape Database tab, choose the tape and click Delete.



To eject a tape

1. Click UTILITIES on the toolbar.
2. On the Device tab, make sure the correct tape is selected.

Note:

The current tape's name is shown in **Current tape**. If the device is a tape changer, choose the correct tape.

3. Click Eject.

Note:

Some devices require you to manually eject the tape.



To erase data on a tape

1. Click UTILITIES on the toolbar.
2. On the Device tab, make sure the correct tape is selected.

Note:

The current tape's name is shown in **Current tape**. If the device is a tape changer, choose the correct tape.

3. Click Erase.



To import a tape

1. Insert the tape into the Host Replica Server's tape device.
2. Click UTILITIES on the toolbar.
3. On the Device tab, click Import Tape.

Note:

Replica reads the database information on the tape and imports it into the current Replica database.



To label a tape

Note:

When you relabel a tape, you will not be able to access any existing data on the tape.

1. Click UTILITIES on the toolbar.
2. On the Device tab, make sure the correct tape is selected.

Note:

The current tape's name is shown in **Current tape**. If the device is a tape changer, choose the correct tape.

3. Click Label to open the Assign New Tape Label dialog box.
4. In **New label**, type the new name and click OK.



To make Disaster Recovery disks

Note:

You will need at least two 1.44 megabytes formatted (or unformatted) floppy disks.

1. Insert the first disk in the Administrative PC's drive.
2. Click UTILITIES on the toolbar.
3. On the Disaster Recovery tab, select the server.
4. In **Make disk on**, select the appropriate drive letter and click Create Disks.
5. When done, label the disks and store them someplace safe.



To prepare for a disaster

1. [Create Disaster Recovery disks](#) for all servers.
2. Replicate all servers regularly.
3. Be sure to keep additional Disaster Recovery disks and replication copies off-site in case of irreparable property damages.



To set password security

1. On the File menu, choose Preferences.
2. On the Security tab, check (or uncheck) **Remember passwords for current Replica server**, and click OK.

Note:

When you check this field, Replica will not request that you enter your password each time you log into the server.



To retension a QIC tape

1. Click UTILITIES on the toolbar.
2. On the Device tab, make sure the correct tape is loaded.

Note:

The current tape's name is shown in **Current tape**. If the device is a tape changer, choose the correct tape.

3. Click Retension.



To view tape details

1. Click UTILITIES on the toolbar.
2. On the Tape Database tab, select the tape for which you want to see database details.
3. Click Tape Details.



To abort a job

1. Click STATUS on the toolbar.
2. On the Device tab, select the active job.
3. Click Abort.



To delete a job from the queue

1. Click STATUS on the toolbar.
2. On the Device tab, select the waiting (or holding) job you want to delete.
3. Click Delete.

Note:

You can't get back deleted jobs.



To delete the history of all jobs

1. Click STATUS on the toolbar.
2. On the History tab, click Delete History.

Note:

You can't get back deleted job histories. You may want to print a copy of the job history before you delete it.



To look at a job's progress

1. Click STATUS on the toolbar.

Note:

The Status window opens automatically when you run a replicate or recover job.

2. Click Progress.



To look at job status

1. Click STATUS on the toolbar.

Note:

The Status window opens automatically when you run a replicate or recover job.



To look at the history of all jobs

1. Click STATUS on the toolbar.
2. Select the History tab.



To place a job on hold

1. Click STATUS on the toolbar.
2. On the Device Status tab, select the waiting job and click Hold.



To show or hide the toolbar

1. On the File menu, choose Preferences.
2. Select **Show Toolbar** on the Window tab.

To show or hide the status bar

1. On the File menu, choose Preferences.
2. Select **Show status bar** on the Window tab.

To show or hide Tooltips

1. On the File menu, choose Preferences.
2. Select **Show Toolbar tips** on the Window tab.



To change the default startup window

1. On the File menu, choose Preferences.
2. In **Startup window** on the Window tab select the window you want to start automatically when you start Replica.



To log into the Host Replica Server

1. Start Replica.
2. In the Server Login dialog box, select the Host Replica Server.
3. Type the name and password you use to log into this server, and click OK.



To change the current Host Replica Server

1. In the **Host Replica Server** field of any Replica window, select the new Host Replica Server.
2. In the Server Login dialog box, type the name and password you use to log into this server.



To mount a tape

1. Click RECOVER on the toolbar.
2. On the Source tab, check **Selective file recovery (Mount tape)**.
3. Select the volume from which data was replicated.
4. Choose the tape from the **Tape library**.
5. In the **Mount tape** field on the Destination tab, type the volume name you want to mount this tape as (or accept the default name Replica assigns).
6. Click Mount Tape.



To map a drive letter to recover data on a mounted tape

- ▶ [Windows 3.1](#)
- ▶ [Windows 95](#)
- ▶ [Windows for Workgroups](#)



To map a drive letter in Windows for Workgroups

1. Switch to Windows' File Manager.
2. Choose Connect Network Drive on the Disk menu.
3. Click NetWare.
4. In the Network - Drive Connections dialog, select the drive letter to which you want to map the volume.
5. Click Browse to open the Browse Connections dialog box.
6. Select the Replica mounted tape volume (and any directory) you want mapped to the drive letter, and click OK.

Note:

The Replica mounted tape volume will look something like this: FS1/^VOL1. Notice the ^ symbol as the first character of the volume's name. The tape is mounted as a NetWare Read-Only volume.

7. Click Map in the Network - Drive Connections dialog box.
8. Close all open dialogs.
9. Select the drive you recently mapped to the tape, and copy (or use) files as usual.



To map a drive letter in Windows 3.1

1. Switch to Windows' File Manager.
2. Choose Network Connections on the Disk menu.
3. Click Browse in the Network - Drive Connections dialog box.
4. Select the mounted tape volume (for example, ^VOL1) and click OK.
5. Select the desired drive letter and click Map.
6. Close all open dialogs.
7. In File Manager, select the drive you recently mapped to the tape, and copy (or use) files as usual.



To map a drive letter in Windows 95

1. Right-mouse click Network Neighborhood.
2. Choose Map Network Drive from the context-menu.
3. Select the drive letter and path to the mounted tape volume and click OK.

Note:

The mounted tape volume is listed with a preceding Caret symbol (for example, ^VOL1).

4. In Explorer, select the drive you recently mapped to the tape, and copy (or use) files as usual.



To sort the tape database list

1. Click UTILITIES on the toolbar.
2. On the Tape Database tab, click on a column title to sort the list by that selection.



To resume a job on hold

1. Click STATUS on the toolbar.
2. On the Device Status tab, select a holding job.
3. Click Activate.



To edit an existing replication job

1. Click STATUS on the toolbar.
2. On the Device Status tab, select the replicate job you want to change and click Edit.
3. Make the appropriate changes in the Replicate window that appears.
4. Click Resubmit.



To sort the jobs in the queue

1. Click STATUS on the toolbar.
2. On the Device Status tab, click on a column title to sort the list by that selection.



To print job information

1. Make sure the Host Replica Server is connected to a printer.
2. Click STATUS on the toolbar.
3. Click the History tab.
4. On the File menu, choose Print.

Click **OK** to accepts all entries.

Click **Apply** to accept the entries you've made in the window or dialog box. The window or dialog box remains open.

Click **Help** to show general help information for the active dialog box.

Click **Help** to show general help information for the active tab or dialog box.

Click **Help** to show general help information for the Volume Name dialog box.

Click **Help** to show general help information for the Start/Submit dialog box.

Click **Start** to open the Start/Submit dialog box and give this job a name. Once you click OK, Replica sends the job to the queue, and opens the Status window to show you the job's progress.

Click **Submit** to open the Start/Submit dialog box and give this job a name. Once you click OK, Replica sends the job to the queue to await its scheduled start time and opens the Status window to show you the job's progress.

Click **Resubmit** to resubmit the job you've just edited. Replica opens the Start/Submit dialog box so you can give the edited job a new name (you can use the previous name, if you prefer). Once you click OK, Replica sends the job to the queue to await its scheduled start time, and opens the Status window to show you the job's progress.

Click **Cancel** or **Close** to close the window or dialog box.

Host Replica Server shows the name of the current Host Replica Server. Click the arrow if you want to change it.

In **Host Replica Server**, select the server that has the storage device on which you will replicate data.

In **User name**, type the name you use to attach to this server.

In **Password**, type the password you use to attach to this server.

Click **Show Toolbar** to show or hide the Replica toolbar.

Click **Show Tooltips** to turn Tooltips Help on or off. Tooltips Help appears when you pause the mouse pointer on a button.

Click **Show status bar** to show or hide the status bar.

Click **Date** to show or hide the date in the status bar.

Click **Time** to show or hide the time in the status bar.

In **Startup window**, select the window you want automatically loaded when you start Replica.

Check **Remember passwords for current Replica session** to make Replica remember the password you entered the last time you accessed a server. This is also an effective way of protecting your servers so that only people with the appropriate security rights can access the server.

In the **network list**, select the server (or server object) to replicate. Click or

to collapse or expand the list. You can replicate a server, or any (or all) of its objects. Click a server or object to select or deselect it.

The **Total** field shows the size of the Host Replica Server's hard disk.

Used by files shows the amount of disk space used by all files on the current server.

Deleted files shows the amount of disk space occupied by deleted files.

The **To replicate** field shows the size of object(s) currently selected for replication.

Estimated time shows approximately how long it will take to replicate the selected object(s).

Check **Include deleted files** to replicate all files or directories that have been deleted, but not purged, from this server. This field is checked by default.

Host Replica Server shows the name of the current Host Replica Server.

In **User name**, type the name you use to attach to the current server.

In **Password**, type the password you use to attach to the current server.

Select **Immediate** to replicate data now, then click the Destination tab to tell Replica where to replicate the data. When you click Start, the Job Name dialog box opens so that you can give this job a name and description.

Select **Scheduled** to schedule a job to run later. Fields appear so you can select at what time and on which day(s) to run the job. When you click Submit, the Job Name dialog box opens where you can give this job a name and description.

In the **Sun.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Mon.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Tue.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Wed.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Thurs.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Fri.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In the **Sat.** field, specify whether or not to replicate data on this day.

Choose **None** if you don't want to replicate data on this day. Choose **Daily** to replicate data on this day at the scheduled time. Choose **Weekly** if you want to keep your replication tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse your weekly tapes. You can set only one day to Weekly.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape. If the wrong tape is in the device, Replica will not run the job.

In **Time to replicate**, specify at which *time* to start replicating data for the day(s) you selected.

In **Time to replicate**, specify at which *time of day* to start replicating data for the day(s) you selected.

In the **Reuse weekly tapes after** field, specify when to reuse weekly tapes. If, for example, you schedule a Weekly job on Friday and select 4 in this field, every Friday you'll use a new tape, but on the fifth Friday you'll start over with the first Friday tape.

By default, Replica overwrites data on the tape. To append data on the tape, choose **Append** from the Options tab. You can also have Replica verify the correct tape is in the device before it writes to the tape.

In **Target drives**, select the device (tape or disk) on which to replicate data. The list shows all available devices for the current Host Replica Server (including free space). If you selected a single volume as the source, the list shows the available unmounted disk volumes on which you can replicate data. If a device is dimmed in the list, it is not available for the current server.

Tape name shows the tape's label (name). If you don't label a tape, Replica labels it `TAPE_time` for you automatically. You can always change a tape's label by clicking the New Tape button. Be careful about changing a tape's label. Replica will not associate existing data on the tape with the new label.

The **Seq.** field shows the tape's sequence number. Replica starts the sequence number at 1. If the data spans across tapes, Replica automatically increments the tape sequence number. You can change the number by clicking the New Tape button.

Description shows a description for the replication tape. Replica automatically gives the tape a description when you run an Immediate or Scheduled backup. You can change this by clicking the New Tape button.

Click **New Tape** to open the New Tape dialog box where you can change (or specify) the tape's name, description, and sequence number.

The **Tape library** shows a list of all tapes created (or imported) on the selected Host Replica Server's device.

The tape list shows the name, sequence number, type, and description of all tapes created (or imported) on the selected tape device.

In the **Name** field, type the tape's name. Use names that will help you identify the tape's contents (for instance, SERVERNAME_DATE, SERVERNAME_DAYOFWEEK, SERVERNAME_REPLICATIONTYPE).

In the **Description** field, type a description for this tape.

In **Sequence number**, enter a sequence number for this tape (or use the default number Replica assigns). If data spans across multiple tapes, Replica increments the sequence number automatically.

Click **Help** to open Help for the New Tape dialog box.

Check **Verify replication** if you want Replica to verify data can be read from tape once it is replicated.

Check **Eject tape after replication** to eject the tape from the drive once the job is complete. You might want to use this feature if you replicate data at the end of the day, or if you are replicating data on tapes in a tape changer.

Select **7/24 Live Server Replication** to replicate data while the volumes are in use. Replica protects open files and adds a second control block to the end of the tape. Using this method allows unlimited use of the server during replication. It does not allow tape spanning. This field is selected by default.

Select **Limited live server access with spanning** to allow some server access during replication but let the data span tapes if necessary. If user activity causes the cache buffers to fill, Replica suspends user activity on volume until its data is replicated.

Select **Dismount volume before replicating** to take the volume off-line while data is being replicated.

Select **Overwrite** to start replicating data at the beginning of the tape and write over any data that currently exists.

Select **Append** if you don't want Replica to overwrite existing data on the tape during a replication job.

Check **Verify tape labels** if you want Replica to check the tape to make sure it is the correct tape before writing to it.

In **Job name**, type a name for this replication job. Use names that will help identify the job (for instance, SERVERNAME_VOLUMENAME, SERVERNAME_REPLICATIONTYPE).

In **Description**, type a brief description to identify this job.

In **Device**, select the tape or disk device whose job status you want to monitor.

The job list shows whether a job is in progress, was submitted and is waiting to start, or has been placed on hold. This also shows the server being replicated, the job's name, whether the job is scheduled to run daily or weekly, when it is scheduled to start, and its description.

The job list shows whether a job is in progress, was scheduled and is waiting to start, or has been placed on hold. This also shows the server being replicated, the job's name, whether the job is scheduled to run daily or weekly, when it is scheduled to start, and its description.

Click **Hold** to put a waiting job on hold. Once on hold, you can activate, edit, or delete the job.

Click **Activate** to place a job on hold in a Waiting mode so it will run when the time comes. You can edit, delete, or place on hold a job that is waiting.

Click **Delete** to remove the selected job from the queue. You can delete a job that is on hold or is waiting.

Click **Abort** to stop a running job. The job will be placed on hold. To restart it, click Activate.

Click **Edit** to open the Replicate window and view (or change) what you've defined for the selected job.

[Click this to view details for the running job.](#)

Start time shows when the job started.

Server shows which server is being replicated.

Volume shows which server volume is being replicated.

Tape name shows the name of the tape on which data is being replicated.

Data rate shows how fast data is being written to tape.

Bytes written shows how much data has been written to tape.

Total vol. size shows the actual size of the recovered data, not the total size of the media. During replication, this field shows the **estimated size** of the data being replicated.

Percent complete shows the percentage of data that has been written to tape.

Click **Delete History** to delete the history of all jobs from the database. You cannot recover deleted job histories.

The history list shows a report of all replicate and recover jobs you've submitted or started.

Click **Selective file recovery (Mount tape)** to recover individual files from a tape. You'll need to select the volume and the tape the volume is on, then open the Destination tab and enter a volume name (or use the default name *^volume name*) so NetWare can mount the tape. Click Mount Tape, and NetWare will present the tape as a Read-Only volume. Map a drive letter to the volume so you can access its files using File Manager or other utility.

In **Tape device**, select the tape device from which to recover data. The Tape library field shows all tapes and jobs created for the selected object. If a tape is dimmed in the list, it is not compatible with the selected device.

In the network list, select the object to recover (server, volume, and so on). You can recover data from one server at a time.

The **Tape library** lists all tapes created for the selected object. If a tape is dimmed in the list, this means it is not compatible with the selected device.

The tape list shows the available tapes for the selected device. For each tape, it shows the tape's name, sequence number, when the job was created, the type of tape, and the job's name and description.

Click **Start** to begin recovering data to the location you specified for the tape you selected on the Destination tab.

Click **Mount Tape** to mount the selected tape. NetWare presents the tape as a Read-Only volume. Map a drive letter to the volume so that you can use files from the volume using File Manager or other utility.

When you use Selective file recovery (Mount tape), you can enter a non-default volume name in the **Mount tape** field so NetWare can mount the tape (the default name is *^volume name*), then click Mount Tape. NetWare presents the tape as a Read-Only volume. Map a drive letter to the volume so that you can use files from it using File Manager or other utility.

If you are recovering to free space, in **Free Space name**, type the name you want to assign to free space. Replica converts free space into usable server space.

In the network list, select the server where you want to recover data.

Destination shows the path to the new location where you will recover data.

Destination size shows how much space is available at the destination location.

Source size shows how much data is being recovered.

SCSI ID shows the SCSI ID of the tape device connected to the Host Replica Server and the name of the current device type.

The **Type** field shows the type of tape for the current device (for instance, 4 millimeter, DAT, and so on).

Current tape shows the name of the current tape.

Click **Label...** to open the Assign New Tape Label dialog box and assign a new label (name) to a tape; or change a tape's label, description, and sequence number.

Click **Erase** to erase everything on the tape. The tape is removed from the database.

Click **Import Tape** to recover, into the current database, information about a tape that was created on a remote Replica server. Replica adds the tape information to the current database.

Click **Eject** to dismount and eject the tape in the current device.

Click **Retension** to fast-forward the quarter-inch cartridge (QIC) tape in the current device to the end of the tape and then back to the beginning of the tape. This retensions the tape to ensure it is correctly aligned on the spool.

Old label shows this tape's current label (name).

Old sequence shows this tape's current sequence number.

In **New label**, type the new label you want to give this tape.

In **Description**, type the new description for this tape.

Sequence shows the tape's sequence number. You can change this number if you prefer. If the data you are replicating requires multiple tapes, Replica automatically increments the tape sequence number for each tape. Each tape in the set has the same label and a different sequence number.

Click the scrolling arrows to change the sequence number for this tape. If the information on the tape requires multiple tapes, Replica automatically increments the tape sequence number for each tape. Each tape in the set the same label and a different sequence number.

The tape database list shows all the tapes in the Replica database, their sequence number when they were last accessed, and the last time their names were changed.

The tape database list shows all the tapes in the Replica database, their sequence number when they were last accessed, and the last time their names were changed.

Click **Delete** to remove the selected tape from the Replica database.

Click **Tape Details** to open the Tape Details dialog box for the selected tape and see how many times the tape was written to and read from.

Tape name shows the current tape's name.

Sequence number shows the current tape's sequence number.

Description shows the current tape's description.

Last access shows the last time the current tape was accessed.

Last replication shows when the current tape was last written to.

Last device shows the name of the device last used.

Last device type show the name of the last device used.

Number of labels shows the number of times the current tape's label was changed.

Minutes written to shows the number of times the current tape was written to.

Minutes read from shows the number of times the current tape was read from.

Number of times used shows the number of times the current tape has been used.

In **Create disaster disks for server**, select the Replica server for which you want to create Disaster Recovery disks.

In the **Make disks on Drive** field, select the drive letter where you inserted the Disaster Recovery floppy disk.

The **Created on** field shows when you last created disaster recovery disks for the selected server.

Click **Create Disks** to start creating Disaster Recovery disks for this server.

Shows the progress of the Disaster Recovery disk creation process.


In **Free space**, type the name you want to assign to free space. Replica converts free space into usable server space.



Logging into the Host Replica Server

To use Replica, you must first attach to the *Host Replica Server*. This server contains the [Replica engine](#), storage device (tape or disk), and [Replica Server Agent](#). In the Server Login dialog box, select and log into the Host Replica Server.

The **Host Replica Server** field lists all Replica servers (host or [remote](#) server) on the network. Once you log into the Host Replica Server, you can replicate or recover data, manage replication [jobs](#) and tapes, or view the status and progress of your jobs.




What to Replicate

Use the Replicate Source tab to specify what to replicate. The **Host Replica Server** field shows the name of the current Host Replica Server—the server where the tape device is located, and where you will replicate data.

The *network list* presents all servers and server objects on the network available for replication. Click



or



to expand or collapse objects in the list. To select a server or server object in the list, click it to check it. To uncheck it, click it again. When you select a server or object, details about it appear to the right in **System information**.

Note:

If a server is dimmed in the list, it is not a Replica server (it doesn't have the [Replica Server Agent](#)) and cannot be selected for replication.


When you replicate a server, Replica processes all volumes (for instance, VOL1, VOL2, and so on), the Bindery (or NDS if you are running Novell NetWare 4.x), NetWare partition, non-NetWare partition (DOS), and volume SYS.

Selecting objects


You can replicate one server on the network at a time. Select:



A server, and Replica automatically selects partition data, the Bindery (or NDS), non-NetWare partition, volume SYS, and all other volumes



Partition data or non-NetWare partition, and Replica automatically selects the Bindery (or NDS) and volume SYS.




Bindery, SYS, or a volume. Replica lets you replicate these objects individually or in any combination.

Replica also gives you the option of replicating files that have been deleted but not yet purged. In **Include deleted files**, tell Replica whether or not to replicate these files.

Once you've specified what to replicate, you can [start the job now](#) or [schedule it](#) to run later.

See also:



[Where to Replicate](#)



Starting a Job Now

Once you've selected what to replicate, click Start on the Replicate window to begin the job immediately. Replica selects **Immediately** on the Schedule tab by default, and it processes data onto the tape in the current Host Replica Server's device. If your device is a tape changer, and you want to replicate data on a different tape, open the Destination tab and choose the tape.


When you click Start, the Start/Submit Job dialog box opens from which you can give the job a name and description. Use names and descriptions that will help you (or someone else) identify the data on the tape. Here are some typical job naming formats:

SERVERNAME_VOLUMENAME

SERVERNAME_REPLICATIONTYPE

When you start a job, Replica opens the Status window to show you the job's progress.

See also:




[Labeling a New Tape](#)



[Preparing for a disaster](#)



[Viewing the status of jobs in the queue](#)



Scheduling a Job

If you don't want a job to run immediately, use **Scheduled** on the Schedule tab to specify when to run the job.

Choose **Daily** to replicate a job on the day selected. Choose **Weekly** if you want to keep your tapes longer than one week. In the **Reuse weekly tapes after** field, specify after how many weeks to reuse the tapes. You can set only one day to Weekly. Any other day set to Weekly automatically changes to Daily. Choose **None** if you don't want to replicate data on a specific day.

Replica automatically replicates data on the tape it finds in the current device. By default, it overwrites data on the tape. To append data on the tape, choose **Append** on the Options tab.

Here are some examples of what happens when you use Daily and Weekly:




If you select Daily for all days and **did not** check the Overwrite tape option on the Options tab, data will be appended on the tape each day.



If you select Daily for all days and checked the Overwrite tape option, data on the tape will be overwritten. Each day, the tape contains that day's replication.



If you select Weekly for a day and did not check the Overwrite tape option, data will be appended on the tape when you reuse the tape after the number of weeks specified in the **Reuse weekly tapes after** field have gone by.



If you select Weekly for a day and checked the Overwrite tape option, data will be overwritten on the tape when you reuse the tape after the number of weeks specified in the **Reuse weekly tapes after** field have gone by.

After you define scheduling options, click Submit. Replica opens the Job Name dialog box so you can give the job a name and description. Replica appends the day or week to the job's name to make it easy for you to identify the job on the tape (for example, FS5_VOL1_SUN or FS5_VOL1_WEEKLY1). Replica also labels the tape it finds in the current device TAPE_ *time*. *Time* is the actual *hour:minutes:seconds* when the job started.

Once you click OK, Replica submits the job to the queue and will run it on the specified day and at the specified time.


See also:




[Labeling a New Tape](#)



[Preparing for a Disaster](#)



[Viewing the Status of Jobs in the Queue](#)



[Overwriting Data on the Tape](#)



Where to Replicate Data

On the Replicate Destination tab, choose the device where you want to replicate data, then choose the tape from the Tape library. By default, data is replicated to the tape in the current device. If the device is a tape changer, you can easily select a different tape.


Note:

You can tell Replica on which tape to replicate data when you run immediate jobs only.


The **Target drives** field lists all Host Replica Server devices, as well as the names of all unmounted volumes, on which you can replicate data. When you select a device, the **Tape Library** field lists the tapes that were created for the selected device. If you have an unlabeled tape in the device, the tape is labeled <<default>>.

When you are ready to submit a job, click Submit. Replica opens the Job Name dialog box so you can give the job a name and description. Once you click OK, Replica submits the job to the queue and will run it on the specified day and at the specified time.


See also:



[Labeling a New Tape](#)



[Preparing for a Disaster](#)



Managing Data on the Tape

Use the Options tab on the Replicate window to manage replication jobs on the tape. Replica can verify that data can be read from tape after replication. After the tape is written to, Replica reads it to verify data integrity.

You can have Replica eject the tape after each job, append data on the tape or overwrite existing data. Replica can even verify the correct tape is in the device before it writes to it.

You can choose how much access users have to the server volumes during replication.

7/24 Live Server Replication gives users full access to the server during the process. Replica keeps track of changes to open files and writes a second control block at the end of the tape. Jobs that use this method cannot span tapes.

Limited live server access with spanning allows tape spanning because it uses a single control block, at the beginning of the tape. This requires a different method of tracking changes to files on the disk, which causes Replica's cache buffers to fill up if there is heavy use of file copies and writes. If the buffers fill up, Replica suspends all access until the job is complete. Replica can span tapes resulting from jobs run using this method.

Dismount volume before replicating prohibits all access during the job. Tape spanning is allowed. Replica mounts the volume again when the data is replicated.

See also:



[Changing a Tape's Label](#)



[Managing Replication Tapes](#)



Logging into a Different Server

If you have several Host Replica Servers on the network, you can switch servers at any time from any Replica window.

In the **Host Replica Server** field from the Server Login dialog box, select the server and log into it as usual.




Labeling a New Tape

Use the New Tape dialog box to give a tape a name and description. By default, Replica labels the tape `TAPE:time`. The *time* is the actual *hour:minutes:seconds* when you started the job.

Use names that will help you (or someone else) easily identify data on the tape. Here are some typical formats:


```
SERVERNAME_DATE  
SERVERNAME_DAYOFWEEK  
SERVERNAME_REPLICATIONTYPE
```

Replica automatically gives a tape a sequence number  starting with 1. If the data being replicated spans across multiple tapes, Replica automatically increments the sequence number.

See also:



[Managing Replication Tapes](#)



What to Recover

On the Recover Source tab, select the server or object(s) on the tape to recover. If replicated, you can recover:


[A server](#)

[A specific server object](#)


[Enough information to bring a server up quickly from a disaster](#)


[One or more files or directories](#)

When you select a server or object to recover, the **Tape library** lists all available tapes. If a tape is

dimmed in the list, it is available for the object but is not compatible with the current device.  the device shown in the **Tape device** field.

Recovering a server or object

To recover an entire server or a specific server object, select the server's (or object's) name from the network list, then select the appropriate tape from the Tape library. Click  or



to expand or collapse the network list.

Note:

You can recover one server (or object) at a time when you perform a normal recovery.

Recovering a server from a disaster

If you need to recover a server from a disaster, Replica allows you to perform a full server or a minimal disaster recovery. You'll need to start the server with its Disaster Recovery disk, then recover data from tape.


For a full server replication, Replica recovers all volumes, the Bindery (or NDS), partition data, non-NetWare partition, and SYS volume.

If you select partition data or non-NetWare partition, Replica automatically assumes you want to do a minimal disaster recovery and will select the appropriate objects for you.


Recovering files or directories

To recover individual files from tape, you'll need to mount the tape. Check **Selective file recovery (Mount tape)**, select the server or object to recover, then select the tape from the Tape library. On the Destination tab, you'll need to give the mounted tape a name so it can be mounted as a Read-Only NetWare volume. Once mounted, you can map a drive letter to the mounted tape volume using File Manager or other utility, and then copy or use files as usual.

See also:



[Mounting a Tape](#)



[Where to Recover Data](#)



Mapping a Drive Letter to a Mounted Tape



Where to Recover Data

On the Recover Destination tab, specify where to recover data. You can recover data to the same location from where it was replicated, or you can recover data someplace else. Select the destination location from the *network list*. If you select a different location, it must have at least as much disk space as the original location.

Note:

What appears in the network list depends on the selections you've made on the Source tab. If you are performing a disaster recovery, the list shows only the servers that have been started with their Disaster Recovery disks.

If you are recovering individual files, you must check the Selected file recovery (Mount tape) field on the Recover Source tab, then enter a volume name in the Recover Destination tab to [mount](#) the tape.

The **Destination** field shows the path to the new location where you will recover data, the **Destination size** field shows how much room is available at the destination location to recover data, and the **Source size** field shows the original size of the data.

See also:



[Mounting a Tape](#)



[Mapping a Drive Letter to a Mounted Tape](#)



Managing Replication Tapes

Use the Utilities Device tab to manage the current tape. The current device's name and [SCSI ID](#) are shown in the SCSI ID field. If the device is a [tape changer](#), you can switch tapes using **Current tape**.

You can change a tape's [label](#), [sequence number](#), or description; erase everything on the tape; eject the current tape from the device; and [retension](#) QIC tapes. You can also import tape database information created on remote Replica servers into the current [database](#).

See also:



[Importing a Tape](#)



Viewing and Removing a Tape's Database Information

The Tape Database tab shows all the tapes in the current server's [database](#). You can remove a tape from the database by selecting it and clicking Delete.

When you delete a tape from the database, the tape becomes inaccessible in Replica. Use **Import Tape** on the Utilities Device tab to make it accessible again. Replica reads the tape's label and replication information and adds it to the database.

To view details about a tape, select the tape and click Details. You can also sort the list by tape name, sequence number, last access date, or by last time the tape was written to or read from. Click the list's header to sort it by that format.


See also:



[Importing a Tape](#)



[Getting More Information About a Tape](#)



[Printing the Tape Database](#)



Getting More Information About a Tape

When you select a tape from the Tape Database tab in the Utilities window and click Details, the Tape Details dialog box gives you even more information about the tape.

The tab shows the tape's name, description, sequence number, the type of device used, the number of times the tape was written to or read from, and the number of times the tape's label was changed.



Preparing for a Disaster

Always be prepared for a disaster. With a set of Disaster Recovery disks and your latest replication tapes, you can bring your downed server back up in minutes.

The Disaster Recovery tab on the Utilities window helps you create the disks you'll need to recover the server from a disaster. Select a server from the network list and create Disaster Recovery disks for it. You need at least two floppy disk (3.5-inch). Choose the disk drive from the **Make disks on Drive** field and click Create Disks to start the process. If the tapes are unformatted, Replica will format them for you.

If you've already created Disaster Recovery disks for a server, the **Created on** field shows when they were last created. You should create Disaster Recovery disks for a server if its configuration has changed (NetWare, DOS, Replica, hardware, and so on).

When the time comes to recover from a disaster, simply restart the server with its Disaster Recovery disks, and recover data from your most recent tapes.

See also:



[Recovering from a Disaster](#)



[Making Disaster Recovery Disks](#)



Changing a Tape's Label

Use the Assign New Tape Label dialog box to change the tape's current label (name).

IMPORTANT:

If you already have data on a tape and you relabel the tape, the new name will not be associated with any existing data.

Use names that will help you (or someone else) easily identify data on the tape. Here are some typical formats:

SERVERNAME_DATE

SERVERNAME_DAYOFWEEK

SERVERNAME_REPLICATIONTYPE

See also:



[Managing Replication Tapes](#)



Making Disaster Recovery Disks

To make Disaster Recovery disks for a server, you'll need to select the server from the server list on the Utilities window's Disaster Recover tab, and create the disks from the Administrative PC. You'll also need two floppy disks (formatted or unformatted).

When you click Create Disks on the Disaster Recovery tab, Replica will start creating the disks and will prompt you to insert the second disk when appropriate. If the disks are not formatted, Replica will format them for you.

As a general rule, create Disaster Recovery disks when a server's configuration has changed (NetWare, DOS, Replica, or hardware).

See also:



[Preparing for a Disaster](#)



Viewing the Status of Jobs in the Queue

The Status window opens automatically when you start or submit a replication or recover job. The Device Status tab shows each job's progress. The **Device** field shows the name of the current device, and the job list shows all jobs processed for that device.

If a job is running, you can abort it. If it is waiting to start, you can place it on hold, delete or edit it. You can also activate, delete, or edit a job that is on hold.

For the running job, click Progress to find out how fast data is being written to tape, how much data has been written to tape, and how much space on the tape the data is expected to use.

In addition, you can sort the list by status, server, scheduling type, start time, or description. Click on a header category to sort the list by that category.

See also:



[Viewing Job Progress](#)



Viewing Job Progress


In addition to showing details about a selected job, the Progress view of the Device Status tab shows how fast data is being written to tape, how much data has been written to tape, how much space on the tape the data is expected to use.

Note:

If you are viewing the progress of a recovered job, the **Total vol. size** field shows the actual size of the data being recovered.

If a job is running, you can find out how much data has been written to tape. The progress bar shows the progression of data as it is being written to tape.

See also:



[Viewing the Status of Jobs in the Queue](#)



[Viewing the History of All Jobs](#)



Viewing the History of All Jobs

The History tab on the Status window shows the history and status of all jobs you've submitted. When you click Delete History, you delete the history of **all** jobs. Once you delete a job's history, you cannot recover it.

You can also print the history list using the Print command from the File menu.

See also:



[Printing Job Histories](#)



Printing Job Histories

The Print dialog box allows you to send the information on the History tab (from the Status window) to a printer. Choose Print from the File menu to print the contents of the History tab. Make sure the printer is ready for printing.



Printing the Tape Database

The Print dialog box allows you to send the information on the Tape Database tab (from the Utilities window) to a printer. Choose Print from the File menu to print the contents of the Tape Database tab. Make sure the printer is ready for printing.

Disaster Recovery Disk Creation

Once you start creating the Disaster Recovery disks, the Boot Disk Creation dialog box appears to show the percentage of completion. When prompted, insert the appropriate disk in the floppy disk drive.

After creating the disks, store them someplace safe. You'll need them to recover the server from a disaster.

Formatting Disaster Recovery Disks

You can insert a pre-formatted (or an unformatted) disk in the drive. If you insert an unformatted disk in the drive, Replica will format the disk for you when you click Create Disks on the Utilities window's Disaster Recovery tab.

Starting or Submitting a Job

When you start or submit a job, Replica opens the Start/Submit dialog box so you can give the job a name and description. Use names and descriptions that will help you easily identify the job, for instance:

SERVERNAME_VOLUME

SERVERNAME_REPLICATIONTYPE

Once you've named and described a job, Replica sends the job to the queue and presents the Status window to show you the job's progress.

Setting Window Preferences

The Preferences Window tab allows you to:

[Show or hide the toolbar](#)

[Show or hide the status bar](#)

[Turn on or off Tooltips](#)

[Select a startup window](#)



Password Security

On the Preferences Security tab, check **Remember passwords for current session** if you want Replica to remember a server's password so you don't have to keep entering it each time you access the server. This field is checked by default.

This is useful if you switch servers frequently. Replica remembers the server's password and doesn't request it the next time you log into the server.




Overwriting Data on the Tape

By default, Replica writes over existing data on the tape during replication. You can have Replica append data on the tape (not overwrite data) by selecting Append on the Replicate Options tab.

Whether you overwrite or append data on the tape, you can have Replica verify the correct tape is in the drive before it writes to tape by checking **Verify tape labels**.


Note:

When data is overwritten on the tape, it cannot be recovered.




Bringing up a Server Quickly from a Disaster

If a server is down, you can bring it back up quickly using disaster recovery disks you previously created for the server. Replica gives you the option of performing a full server (or a minimal) disaster recovery. You'll need to:

 Start the downed server with its disaster recovery disks.



Load the latest replication tape you created for the server into the Host Replica Server's device.



Use the Administrative PC to recover the server's data from tape.

Note:

If you didn't create disaster recovery disks for the server and it has the same configuration as another server (hardware, software, and so on), you may be able to recover NetWare using the other server's disaster recovery disks. Otherwise, you'll have to reinstall NetWare.

Full server disaster recovery

When you perform a full server disaster recovery, Replica automatically recovers all volumes, the NetWare partition, the Bindery (or NDS), non-NetWare partition, and SYS volume. Use the [Source tab](#) on the Administrative PC's Recover window to select the tape from the Tape library, and then select the server on the [Destination tab](#).

Minimal disaster recovery

If you only want to recover the minimal files necessary to bring the server back up, you can perform a minimal disaster recovery. Select either NetWare partition or non-NetWare partition, and Replica will automatically select the appropriate files to get this server up and running quickly.

Use the [Source tab](#) on the Administrative PC's Recover window to select the tape from the Tape library, and then select the server on the [Destination tab](#).



Specifying which Window will Start Automatically

When you start Replica, it opens the Replicate window by default. If you want to open a different window (or if you don't want any window to open) during startup, make your selection on the Preferences Window tab.

You can also have Replica open the window that was open before you exited Replica.



Mounting a Tape

When you need to recover a specific file or directory from tape, Replica makes it easy for you do so you. Use Replica Recover to mount the tape. (Replica mounts the tape as a Read-Only NetWare volume). Then, give the tape volume a name or use the default name Replica assigns. Replica normally use the original volume name and precedes the name with the ^ symbol (for example ^VOL1).

Once mounted, switch to Windows File Manager (or other utility) to map a drive letter to the volume. From there, copy or use files as usual.

See also:



[Mapping a Drive Letter to a Mounted Tape](#)



[Unmounting a Tape](#)



Unmounting a Tape

When a tape is mounted, you can unmount it at any time. Use the Status Device tab on the Status window to unmount the tape. Simply select the tape from the list, and click Unmount.

See also:



[Mounting a Tape](#)



Mapping a Drive Letter to a Mounted Tape

Map a drive letter to a mounted tape volume to recover files or directories on the tape. Click the How To button to display the procedures list, and choose **Map a drive letter to a mounted tape** for instructions on mounting the tape from Windows 3.1, Windows 95, or Windows for Workgroups.

Once mapped, use DOS, File Manager, or your application's File Open command to recover data on the tape. Replica mounts the tape as a Read-Only NetWare volume, and adds the ^ symbol to the volume's name (for example, ^VOL1) so you can easily identify it.

Once you map a drive letter to a mounted tape volume, copy or use files on the tape as usual.

See also:



[Mounting a Tape](#)



Importing a Tape

You can import the tape database created on another Replica server into the current Host Replica Server using Replica's Import Tape on the Utilities window.

Replica adds the new information to the current database.

Naming Free Space Volumes

When you replicate or recover data to free space, you must give the intended free space volume a name. The free space volume you are replicating or recovering data to must be at least the same size as the original replicated volume.

In the Volume Name Required dialog box, enter the name you want to assign to the intended free space volume.

Using Free Space to Duplicate Server Objects

Replica makes it easy for you to duplicate a server object. If you have unused (free) space on the server, you can convert it into a volume and replicate the the selected object into it. The new free space volume will have all the characteristics of the

Note:

Free space must be at least the same size as the original object being replicated.

To replicate data onto free space

1. On the Source tab in the Replicate window, choose the volume to replicate.
2. Select Free Space on the Destination tab and click Start.
3. In the Volume Name Required dialog box, type the new volume name (up to 15 characters) for free space. The name must be a valid NetWare name.
4. Click OK then give the job a name and description.

Once data has been replicated, the new free space volume will appear in the network list.

High-Performance Option

When you have the Replica High-Performance Option, you can use as many supported devices as are attached to a server. The devices can be of the same type (for instance, DAT drives), or a combination of different device types (for instances, DAT, DLT, and so on). With the High-Performance Option, Replica recognizes them all.

For more information, see the High-Performance Option documentation.

Multi-Server Agent

When you have the Replica Multi-Server Agent option, you can protect additional servers without dedicated tape drives. It lets you replicate and recover objects from remote servers using any Host Replica Server.

For more information, see the Multi-Server Agent documentation.

How To Reach Us

If you have questions or problems, click the icon for details on how to get help. The SUPPORT.WRI file is stored in the Replica directory.



Product
Support

A *job* defines the replication task you want to accomplish. It includes the server (or server objects) to replicate, the day and time to start replication, details about devices and tape rotations, and special replication options you specify.

A job also has a name and description. You can use the default name Replica assigns, or select a unique one (for instance, SERVERNAME_VOLUMENAME, SERVERNAME_REPLICATIONTYPE).

A *Replica Server Agent* is required on all servers you want to replicate. It allows Replica to see the server on the network. The agent reads data from the Host Replica Server.

Replica lets you *mount* a tape as a Read-Only NetWare volume so that you (or your users) can easily recover files or directories from Windows File Manager, an application, or other utilities.

When you mount a tape as a volume to recover files or directories, you access data directly from the tape.

SCSI ID (Small Computer System Interchange Identification) tells NetWare where your tape device is located.

A *tape changer* is a device that holds multiple tapes. If your device is a tape changer, you can prepare jobs to replicate on any tape in the changer.

A tape's *label* is its name. If you don't assign a unique name to a tape, Replica automatically assigns a default name for you.

When you run an immediate replication job, Replica brings up a dialog box where you can assign a label to the current tape. When you schedule a job, you can assign a label from the Utilities window.

Here are some example tape labels: SERVERNAME_DATE, SERVERNAME_DAYOFWEEK, SERVERNAME_REPLICATIONTYPE.

Replica assigns a *sequence number* to a tape to help you identify tapes in series in the event data spans across multiple tapes.

If a job fits on one tape, the tape's sequence number is 1. If it fits on two tapes, the first tape's sequence number is 1 and the second tape is 2. And so on.

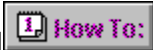
With Replica, you can easily describe data being replicated. Use *descriptions* that will make it easier for you (or someone else) to identify data on the tape when you need to recover data later.

Replica retensions QIC (Quarter-Inch Cartridge) tapes by fast-forwarding the tape to the end and then rewinding it. This ensures the tape is correctly aligned on the spool.

Replica stores information about a job in a *database*. It maintains two sets of databases: scheduling and tape. The scheduling database includes information about when the job started, and when to rotate tapes. The tape database includes details about the tape (name, sequence number, device name, last access and written to date, description of the job, and replication type).

The *Replica engine* is the component that makes Replica work. It is loaded on the server that contains the tape device (or disk drive, if you are replicating data to a server's disk).

A *remote server* is any server that is not the Host Replica Server. It has the Replica Server Agent so that Replica can recognize it on the network, but doesn't need/have a tape device or the [Replica engine](#). You replicate a remote server's data onto the Host Replica Server's device.

The *Host Replica Server* is where data is replicated.  think of it as a *tape server*. It has the [Replica Server Agent](#), [Replica engine](#), and storage device. You can replicate this server's (as well as [remote servers](#)) data onto its storage device.

If you have any *free space* on a server that has not been designated for use, you can recover data to it. Replica converts free space into a volume before recovering data to it. You choose the volume's name.

Welcome to Replica for NetWare Help

To use Help, place the mouse pointer on a topic and click. To learn how to use Help, press F1.



Introducing...



Replicating



Recovering




Disaster Recovery




Utilities



Status



Replica for NetWare Overview


Replica protects your data and makes it easy for you to recover it when you need it  whether your network is large, small, or somewhere in between.

With Replica, you decide which file server (or server object) to protect. You also decide when and where to replicate data. When it comes time to recover data, the only thing that can slow down the recovery process is the speed of your server and/or tape device. In fact, it sometimes takes only a few minutes to recover a NetWare server from a major disaster. You don't have to spend hours figuring out what happened or, even worse, reinstalling NetWare.

Whether you need to bring up a server quickly from a disaster, recover a file or directory that was accidentally deleted, or recover a server volume that became unstable, Replica simplifies the process.

Using Replica

Here's how you generally use Replica:




You use a workstation (Administrative PC) running Microsoft Windows and the Replica software to prepare replication jobs or to recover data. You can perform some tasks at the server console, but most of the time you'll use the Administrative PC.




You use a file server with a dedicated storage device (tape or disk), the [Replica engine](#), and [Replica Server Agent](#) to replicate and recover data. This server is called the *Host Replica Server*. Other Replica servers on the network that are not Host Replica Servers are called [remote servers](#).

Note:


All servers on the network whose data you want to protect must have the Replica Server Agent loaded.

Data is replicated to the Host Replica Server's storage device (tape or disk). For example, assuming FS1 (file server 1) is the Host Replica Server and you want to replicate FS2 (a remote server), data will be replicated onto the FS1 storage device  the current Host Replica Server.

See also:



[Replicating data](#)



[Recovering data](#)



Replicating

Overview

What would you like to do?

 How To:

[Choose a device on which to replicate data](#)

 How To:

[Eject a tape after replication](#)

 How To:

[Name a replication job](#)

 How To:

[Overwrite data on a tape](#)

 How To:

[Prepare a replication job](#)

 How To:


[Replicate data now](#)

 How To:

[Replicate data while a volume is in use](#)

 How To:

[Schedule jobs to run later](#)

 How To:

[Specify when to reuse weekly tapes](#)

 How To:

[Verify data can be read from tape](#)



Recovering



What would you like to do?



[Bring up a server from a disaster](#)



[Import a tape](#)



[Map a drive letter to a mounted tape](#)



[Mount a tape](#)



[Recover files or directories](#)



[Recover from a disaster](#)



[Recover remote Replica servers](#)



[Recover Replica databases](#)



[Recover the Host Replica Server](#)



[Specify what to recover](#)



Disaster Recovery



What would you like to do?



[Make Disaster Recovery disks](#)



[Recover remote Replica servers](#)



[Recover the Host Replica Server](#)



Utilities



What would you like to do?



[Delete a tape from the database](#)



[Eject a tape](#)



[Erase data on a tape](#)



[Import a tape](#)



[Label a tape](#)



[Make Disaster Recovery disks](#)




[Prepare for a disaster](#)



[Retention a QIC tape](#)



[Set password security](#)



[View tape details](#)



Status



What would you like to do?



[Abort a job](#)



[Delete a job from the queue](#)




[Delete the history of all jobs](#)



[Look at a job's progress](#)




[Look at job status](#)




[Look at the history of all jobs](#)



[Place a job on hold](#)




[Print job information](#)




Replicating Data

With Replica, you can protect:




An entire server. Replica automatically protects the NetWare partition, the Bindery (or NDS), non-NetWare partition, volume SYS, and all other volumes.




A specific server object (for instance, a volume, the Bindery, SYS volume, NetWare partition, and so on).

The replication process


Replica backs up data to the Host Replica Server's device. This device can be a tape or disk. If the device is a tape changer, you can specify on which tape to replicate data. Whether you are replicating a server or a specific server object, you must specify:




What to replicate. You can replicate data from any server that has a Replica Server Agent



whether the server is [remote](#) or is the [Host Replica Server](#). Assuming you have two servers on the network (FS1 and FS2). Let's say that FS1 is the Host Replica Server and FS2 is a remote server, you can replicate both servers' data but the data will be replicated onto the Host Replica Server's device (that is, FS1). Use the [Source tab](#) on the Replicate window to specify what to replicate.




When to replicate it (now or later). Once you've selected what to replicate, tell Replica when to start the replication job. You can start the job immediately or schedule it to run later. When you schedule a job, Replica lets you select the day(s) and time when the job will run. It also lets you specify when to reuse weekly tapes. Use the [Schedule tab](#) on the Replicate window to specify when to replicate data. To truly protect your data, replicate your servers regularly.




Where to replicate it. After you specify what to replicate and when to replicate it, specify where to replicate data. Replica backs up your data to the Host Replica Server's device. If the device is a tape changer, data will be replicated onto the tape you specify. If all the tapes in the changer are not labeled, data will be replicated onto the first tape Replica finds.

When you replicate data to the Host Replica Server's disk, the disk volume must be unmounted. You can have Replica dismount the volume for you, or you can do it at the server console. Use the [Destination tab](#) on the Replicate window to specify where to replicate data.

See also:



[Recovering Data](#)




[Recovering from a Disaster](#)




Recovering Data

Replica helps you recover your data quickly and easily. With Replica, you can recover:



A server. Replica recovers all volumes, the Bindery (or NDS), non-NetWare partition, the NetWare partition, and SYS volume.



A specific server object. For instance, you can recover a server volume, the Bindery, and so on.



One or more files or directories. You can mount the tape as a volume, and then map a drive letter to the mounted tape volume and use Windows File Manager (or other utility) to recover data. From there, copy or use files as usual.

The recovery process


Replica recovers data to the location you specify. You can recover data to its original location or someplace else. If you recover data to a different location, make sure there's enough storage space at the new location to store the data. (the total space must be at least the same size as the original space).

To recover data, you must specify:




What to recover. Tell Replica whether to recover a file, directory, server, or server object. Use the [Source tab](#) on the Recover window to specify what to recover. When you recover files or directories, you'll need to mount the appropriate replication tape. (NetWare mounts the tape as a Read-Only volume). You can then map a drive letter to the mounted volume using File Manager (or other utility), and copy or use files as usual.

If a server crashes and you can't access its hard disk, you can [recover it quickly](#) from a disaster. You'll need to start the server with its [Disaster Recovery disks](#) and then recover the data using the server's latest replication tape(s). Replica lets you specify whether to recover the full server or just the necessary files to bring the server back up in minutes.




Where to recover the data. Using the [Destination tab](#) on the Recover window, tell Replica whether to recover data to its original location or someplace different. You can also recover data to available [free space](#) on the server.


See also:



[Recovering from a Disaster](#)



[Mapping a Drive Letter](#)




Recovering from a Disaster

When you have a complete set of Replica Disaster Recovery disks and your latest replication tapes, you can recover any Replica server (host or remote) from a disaster.

Recovering the Host (or Remote) Replica Server

To recover a Host (or remote) Replica Server from a disaster, start the server with its Replica disaster recover disks, then recover the server's data from its latest replication tape.


For example, if you have a Host Replica Server on the network (FS1) and a remote server (FS2). Let's say you want to recover FS2. Here's what you'll need to do:



Start FS2 with its Disaster Recovery disks so that Replica can recognize it on the network.



From the Administrative PC, make sure FS1 is the Host Replica Server.



Select the FS2 object(s) you want to recover (for instance, partition data, SYS volume, VOL1, and so on).



Select the appropriate tape and click Start.

To recover the Host Replica Server, just load the tape in the device, and start the server with its Disaster Recovery disks. Replica will take care of the rest.

See also:



[Making Disaster Recovery Disks](#)



Managing Jobs Using Utilities

With Replica, you can manage any replication job with a simple point and click. If you are replicating data on QIC tapes, you can retension the tape. When you use tape changers, you can switch tapes easily, you can also change a tape's label.

But that's not all! You can add a tape to (or remove all tapes from) the current Replica database, or eject a tape from the current device. (Not all devices support this functionality). You can also bring database information for tapes created on other servers into the current database, or view tape database information (access, creation time and date, description, and last time it was updated).

At the bottom of the [Utilities window](#) are several buttons that make these tasks easier.



Looking at the Status of Jobs

When you submit or start a job, the Status window opens automatically to show how the job is progressing. You can change a job's status, or remove a job from the queue if you no longer want it to run.

You can look at the status of a job from two different views: Job List and Progress. The Job List view shows the status of all jobs, whereas the Progress view shows the status of the currently running job.

Replica also makes available the [history](#) of all jobs you've submitted. You can delete the history at any time. When you delete the job history, you can't recover it.

