



CHAPTER 8

Configuring the Cisco Emergency Responder 2.0 Disaster Recovery System

These topics describe how to configure the Cisco Emergency Responder (Cisco ER) 2.0 Disaster Recovery System:

- [What is the Disaster Recovery System?, page 8-1](#)
- [Quick-Reference Tables for Backup and Restore Procedures, page 8-2](#)
- [Supported Features and Components, page 8-4](#)
- [System Requirements, page 8-4](#)
- [How to Access the Disaster Recovery System, page 8-4](#)
- [Master Agent Duties and Activation, page 8-4](#)
- [Local Agents, page 8-4](#)
- [Adding Backup Devices, page 8-5](#)
- [Creating and Editing Backup Schedules, page 8-6](#)
- [Enabling, Disabling, and Deleting Schedules, page 8-7](#)
- [Starting a Manual Backup, page 8-7](#)
- [Checking Backup Status, page 8-7](#)
- [Restoring a Backup, page 8-8](#)
- [Restoring a Cluster, page 8-9](#)
- [Viewing the Backup and Restore History, page 8-11](#)
- [Trace Files, page 8-12](#)
- [Command Line Interface, page 8-13](#)

What is the Disaster Recovery System?

The Disaster Recovery System (DRS), which can be invoked from the main Cisco Emergency Responder (Cisco ER) 2.0 web interface, provides full data backup and restore capabilities for all servers in a Cisco ER cluster. The Disaster Recovery System allows you to perform a regularly scheduled automatic or user-invoked data backup. DRS supports multiple backup schedules.

The Cisco Disaster Recovery System performs a cluster-level backup, which means that it collects backups for all servers in a Cisco ER cluster to a central location and archives the backup data to physical storage device.

When performing a system data restoration, you can choose which servers in the cluster you want to restore.

The Disaster Recovery System includes the following capabilities:

- A user interface for performing backup and restore tasks
- A distributed system architecture for performing backup and restore functions
- Scheduled backups
- Archive backups to a physical tape drive or remote sftp server



Note The tape device must be attached to the Publisher.

The Disaster Recovery System contains two key functions, *Master Agent* (MA) and *Local Agent* (LA). The Master Agent coordinates backup and restore activity with all the Local Agents.

The system automatically activates both the Master Agent and the Local Agent on all nodes in the cluster.



Note

The Disaster Recovery System does not migrate data from Windows to Linux or from Linux to Linux. You must run a restore on the same product version as the backup. For information on data migration from a Windows-based platform to a Linux-based platform, refer to the *Data Migration Assistant User Guide*.

The following is a list of supported/recommended SFTP servers (freeware) that can be used as backup storage locations:



Note

This list is not comprehensive and is offered as a recommendation only. You may use any other SFTP server of your choosing.

- OpenSSH (for UNIX)
- Sygwin (<http://sshwndows.sourceforge.net/>)
- freeFTPD (<http://www.freeftpd.com/?ctt=download>)

Quick-Reference Tables for Backup and Restore Procedures

The following tables provide a quick reference for the backup and restore procedures.

Backup Quick Reference

[Table 8-1](#) provides a quick, high-level reference to the major steps, in chronological order, that you must perform to do a backup procedure using the Disaster Recovery System.

**Note**

The Disaster Recovery System does not migrate data from Windows to Linux or from Linux to Linux. A restore must run on the same product version as the backup. For information on data migration from a Windows-based platform to a Linux-based platform, refer to the *Data Migration Assistant User Guide* before following the steps in [Table 8-1](#).

Table 8-1 High-Level Tasks for Performing a Backup Procedure

Task	Reference
Create backup devices on which to back up data.	“Adding Backup Devices” section on page 8-5
Create and edit backup schedules to back up data on a schedule. Note Either a manual or a scheduled backup backs up the whole cluster.	“Creating and Editing Backup Schedules” section on page 8-6
Enable and disable backup schedules to back up data.	“Enabling, Disabling, and Deleting Schedules” section on page 8-7
Optionally, run a manual backup.	“Starting a Manual Backup” section on page 8-7
Check the Status of the Backup—While a backup is running, you can check the status of the current backup job.	“Checking Backup Status” section on page 8-7

Restore Quick Reference

[Table 8-2](#) provides a quick, high-level reference to the major steps, in chronological order, that you must follow to perform a restore procedure using the Disaster Recovery System.

Table 8-2 High-Level Tasks for Performing a Restore Procedure

Task	Reference
Choose Storage Location—You must first choose the storage location from which you want to restore a backup file.	“Restoring a Backup” section on page 8-8
Choose the Backup File—From a list of available files, choose the backup file that you want to restore.	“Restoring a Backup” section on page 8-8
Choose Features—From the list of available features, choose the features that you want to restore.	“Restoring a Backup” section on page 8-8
Choose Nodes—If the feature was backed up from multiple nodes, you must choose the nodes that you want to restore.	“Restoring a Backup” section on page 8-8
Check the Status of the Restore—While the restore process is running, you can check the status of the current restore job.	“Viewing the Restore Status” section on page 8-11

Supported Features and Components

For the Cisco ER 2.0 release, you can back up and restore the following feature:

- CER

When you choose a feature for backup, the system backs up all of its subcomponents automatically.

System Requirements

Make sure that Cisco ER 2.0 is running on all servers in the cluster.

How to Access the Disaster Recovery System

To access the Disaster Recovery System, select **Disaster Recover System** from the pulldown **Navigation** menu on the main Cisco ER 2.0 web interface. Log in to the Disaster Recovery System by using the same Administrator username and password that you use for the Cisco Unified OS Administration web interface.

Master Agent Duties and Activation

The system automatically activates the Master Agent on all nodes in the cluster, but only the Master Agent running on the publisher server is fully active.

The Master Agent (MA) performs the following duties:

- The MA stores systemwide component registration information.
- The MA maintains a complete set of scheduled tasks in the Cisco ER database. When it receives updates from the user interface, the MA sends executable tasks to the applicable Local Agents, as scheduled. (Local Agents execute immediate-backup tasks without delay.)
- You access the MA through the Disaster Recovery System user interface to perform activities such as scheduling backups, adding a new backup task for a specific server or a defined cluster, updating or reviewing an existing entry, displaying status of executed tasks, and performing system restoration.
- The MA stores backup sets on a locally attached tape drive or a remote network location.

Local Agents

Each server in a Cisco ER cluster, including the server that contains the Master Agent, must have its own Local Agent to perform backup and restore functions for its server.

**Note**

By default, a Local Agent automatically gets activated on each node of the cluster.

The Local Agent runs backup and restore scripts on each node in the cluster.

Adding Backup Devices

Before using the Disaster Recover System, you must configure the locations where you want the backup files to be stored. You can configure up to 10 backup devices.

To configure backup devices, follows these steps:

Procedure

Step 1 From the main Disaster Recovery System web page, select **Backup>Backup Device**.

The Backup Device List page appears.

Step 2 To configure a new backup device, click **Add New**.

The Backup Device window appears.

Step 3 Enter the backup device name in the **Backup device name** field.

Step 4 Choose one of the following backup devices and enter the appropriate field values in the Select Destination area:

- **Tape Device**—Stores the backup file on a locally attached tape drive. Choose the appropriate tape device from the list.



Note You cannot span tapes or store more than one backup per tape.

- **Network Directory**—Stores the backup file on a networked drive that is accessed through an SFTP connection. Enter the following required information:
 - **Server name:** Name or IP address of the network server
 - **Path name:** Path name for the directory where you want to store the backup file
 - **User name:** Valid username for an account on the remote system
 - **Password:** Valid password for the account on the remote system
 - **Number of backups to store on Network Directory:** The number of backups to store on this network directory.



Note You must have access to an SFTP server to configure a network storage location. The SFTP path must exist prior to the backup. The account that is used to access the SFTP server must have write permission for the selected path.

Step 5 To update these settings, click **Save**.



Note For network directory backups, after you click the Save button, the DRS Master Agent will validate the selected SFTP server. If the user name, password, server name, or directory path is invalid, the save will fail.

Creating and Editing Backup Schedules

You can create up to 10 backup schedules. Each backup schedule has its own set of properties, including a schedule for automatic backups, the set of features to back up, and a storage location.

To manage backup schedules, follow these steps:

Procedure

Step 1 From the main Disaster Recovery System web page, select **Backup>Scheduler**.

The Schedule List window appears.

Step 2 Do one of the following steps to add a new schedule or edit an existing schedule:

- a. To create a new schedule, click **Add New**.
- b. To configure an existing schedule, click its name in the **Schedule List** column.

The scheduler window appears.

Step 3 Enter a schedule name in the **Schedule Name** field.



Note You cannot change the name of the default schedule.

Step 4 Select the backup device in the **Select Backup Device** area.

Step 5 Select the features to back up in the **Select Features** area. You must choose at least one feature.

Step 6 Choose the date and time when you want the backup to begin in the **Start Backup at** area.

Step 7 Choose the frequency at which you want the backup to occur in the **Frequency** area: Once, Daily, Weekly, or Monthly. If you choose Weekly, you can also choose the days of the week when the backup will occur.



Tip To set the backup frequency to Weekly, occurring Tuesday through Saturday, click **Set Default**.

Step 8 To update these settings, click **Save**.

Step 9 To enable the schedule, click **Enable Schedule**.

The next backup occurs automatically at the time that you set.



Note Ensure that all servers in the cluster are running the same version of Cisco ER and are reachable through the network. Servers that are not running at the time of the scheduled backup will not be backed up.

Step 10 To disable the schedule, click **Disable Schedule**.

Enabling, Disabling, and Deleting Schedules

To enable, disable, or delete schedules, follow these steps:

Procedure

- Step 1** From the main Disaster Recovery System web page, select **Backup>Scheduler**.
The Schedule List window appears.
- Step 2** Select the check boxes next to the schedules that you want to modify:
- To select all schedules, click **Select All**.
 - To clear all check boxes, click **Clear All**.
- Step 3** To enable the selected schedules, click **Enable Selected Schedules**.
- Step 4** To disable the selected schedules, click **Disable Selected Schedules**.
- Step 5** To delete the selected schedules, click **Delete Selected**.
-

Starting a Manual Backup

To start a manual backup, follow these steps:

Procedure

- Step 1** From the main Disaster Recovery System web page, select **Backup>Manual Backup**.
The Manual Backup page appears.
- Step 2** Select a backup device in the **Select Backup Device** area.
- Step 3** Select the features to back up in the **Select Features** area.
- Step 4** To start the manual backup, click **Start Backup**.
-

Checking Backup Status

You can check the status of the current backup job and cancel the current backup job. To view the backup history, see the [“Viewing the Backup and Restore History”](#) section on page 8-11.

To check the status of the current backup job, follow these steps:

Procedure

- Step 1** From the main Disaster Recovery System web page, select **Backup>Current Status**.
The Backup Status page appears.
- Step 2** To view the backup log file, click the log filename link.

Step 3 To cancel the current backup, click **Cancel Backup**.



Note The backup is canceled after the current component has completed its backup operation.

Restoring a Backup

The Restore Wizard leads you through the steps that are required to restore a backup.



Tip

To restore all servers in a cluster, see the [“Restoring a Cluster” section on page 8-9](#).



Caution

Before you restore Cisco ER, ensure that the Cisco ER version that is installed on the server matches the version of the backup file that you want to restore.

To perform a restore, follow these steps:

Procedure

Step 1 From the main Disaster Recovery System web page, select **Restore > Restore Wizard**.

The first page of the Restore Wizard (Step1 Restore—Choose Backup Device) appears.

Step 2 Choose the backup device from which to restore in the **Select Backup Device** area.

Step 3 Click **Next**.

The Step2 Restore—Choose the Backup Tar File page appears.

Step 4 Choose the backup file that you want to restore.



Note The backup filename indicates the date and time that the system created the backup file.

Step 5 Click **Next**.

The Step3 Restore—Select the Type of Restore page appears.

Step 6 Choose the features that you want to restore.



Note Only the features that were backed up to the chosen file display.

Step 7 Click **Next**. The Step4 Restore—Final Warning for Restore page appears.

Step 8 To start restoring the data, click **Restore**.

You are prompted to choose the node to restore.

Step 9 Choose the appropriate node.

**Caution**

After you choose the node to which you want the data restored, any existing data on that server gets overwritten.

Step 10 Your data is restored on the nodes that you chose. To view the status of the restore, see the “[Viewing the Restore Status](#)” section on page 8-11.

Step 11 Restart the server.

**Note**

Depending on the size of your database and the components that you choose to restore, the system can require one hour or more to restore.

Restoring a Cluster

If a major failure or a hardware upgrade occurs, you may need to restore all nodes in the cluster. To restore a whole cluster, you must first restore the Publisher server and then restore the Subscriber server.

The following procedures describe how to perform full-cluster restore process.

Restoring the Publisher

To restore the Publisher server, follow these steps:

Procedure

Step 1 Perform a fresh installation of Cisco ER 2.0 on the Publisher server. See the “[Installing the Cisco Emergency Responder Publisher](#)” section on page 2-4 for more information.

**Caution**

Before you restore Cisco ER 2.0, ensure that the Cisco ER version that is installed on the Publisher server matches the version of the backup file to be restored.

Step 2 From the main Disaster Recovery System web page, select **Restore > Restore Wizard**.
The first page of the Restore Wizard (Step1 Restore—Choose Backup Device) appears.

Step 3 Choose the backup device from which to restore in the **Select Backup Device** area.

Step 4 Click **Next**.

The Step2 Restore—Choose the Backup Tar File page appears.

Step 5 Choose the backup file that you want to restore.

**Note**

The backup filename indicates the date and time that the system created the backup file.

Step 6 Click **Next**.

The Step3 Restore—Select the Type of Restore page appears.

Step 7 Choose the features that you want to restore.



Note Only the features that were backed up to the chosen file display.

Step 8 Click **Next**.

The Step4 Restore—Final Warning for Restore page appears.

Step 9 To start restoring the data, click **Restore**.

Step 10 When you are prompted to choose the nodes to restore, choose only the first node (the Publisher).

Step 11 Your data is restored on the Publisher server. To view the status of the restore, see the [“Viewing the Restore Status”](#) section on page 8-11.



Note During the restore process, do not perform any tasks with Cisco ER Administration or User Pages.

Step 12 Restart the server.



Note Depending on the size of your database and the components that you choose to restore, the system can require one hour or more to restore.

Step 13 After the first node restarts, continue with the [“Restoring Subsequent Cluster Nodes”](#) section on page 8-10.

Restoring Subsequent Cluster Nodes

To restore subsequent nodes in the cluster, follow these steps:

Procedure

Step 1 Perform a fresh installation of Cisco ER 2.0 on the Subscriber server. See the [“Installing the Cisco Emergency Responder Subscriber”](#) section on page 2-8 for more information.



Caution Before you restore Cisco ER, ensure that the Cisco ER version that is installed on the server matches the version of the backup file to restore.

Step 2 From the main Disaster Recovery System web page, select **Restore > Restore Wizard**.

The first page of the Restore Wizard (Step 1 Restore—Choose Backup Device) appears.

Step 3 Choose the backup device from which to restore in the **Select Backup Device** area:

Step 4 Click **Next**.

The Step2 Restore—Choose the Backup Tar File page appears.

Step 5 Choose the backup file that you want to restore.

**Caution**

To restore Subscriber nodes in the cluster, you must choose the same backup file that you used to restore the Publisher.

Step 6

Click **Next**.

The Step3 Restore—Select the Type of Restore page appears.

Step 7

Choose the features that you want to restore.

**Note**

Only the features that were backed up to the chosen file display.

Step 8

Click **Next**.

The Step4 Restore—Final Warning for Restore page appears.

Step 9

To start restoring the data, click **Restore**.

Step 10

When you are prompted to choose the nodes to restore, choose only the subsequent (Subscriber) nodes.

Step 11

Your data is restored on the subsequent nodes. To view the status of the restore, see the [“Viewing the Restore Status” section on page 8-11](#).

Step 12

Restart the server.

**Note**

Depending on the size of your database and the components that you choose to restore, the system can require one hour or more to restore.

Viewing the Restore Status

To check the status of the current restore job, follow these steps:

Procedure**Step 1**

From the main Disaster Recovery System web page, select **Restore>Status**.

The Restore Status page appears.

Step 2

To view the restore log file, click the log filename link.

Viewing the Backup and Restore History

These topics describe how you can see the last 20 backup and restore jobs:

- [Backup History](#)
- [Restore History](#)

Backup History

To view the backup history, follow these steps:

Procedure

-
- Step 1** From the main Disaster Recovery System web page, select **Backup > History**.
The Backup History page appears.
- Step 2** From the Backup History page, you can view the backups that you have performed, including filename, storage location, completion date, result, and features that are backed up.



Note The Backup History page displays only the last 20 backup jobs.

Restore History

To view the restore history, follow these steps:

Procedure

-
- Step 1** From the main Disaster Recovery System web page, select **Restore > History**.
The Restore History page appears.
- Step 2** From the Restore History page, you can view the restores that you have performed, including filename, storage location, completion date, result, and the features that were restored.



Note The Restore History page displays only the last 20 restore jobs.

Trace Files

Trace files for the Master Agent, the GUI, and each Local Agent are written to the following locations:

- For the Master Agent, the trace file is *platform/drf/trace/drfMA0**
- For each Local Agent, the trace file is *platform/drf/trace/drfLA0**
- For the GUI, the trace file is *platform/drf/trace/drfConfLib0**
- For completed backups, the trace files are *platform/drf/log/<timestamp>_b.log*
- For completed restores, the trace files are *platform/drf/log/<timestamp>_r.log*

You can view trace files by using the command line interface. For more information, see [Appendix F, “Command Line Interface.”](#)

Command Line Interface

The Disaster Recovery System also provides command-line access to a subset of backup and restore functions, as shown in [Table 8-3](#). For detailed information on these commands and for information on using the command line interface, see the [Appendix F, “Command Line Interface.”](#)

Table 8-3 *Disaster Recovery System Command Line Interface*

Command	Description
utils disaster_recovery backup	Starts a manual backup by using the features that are configured in the Disaster Recovery System interface
utils disaster_recovery restore	Starts a restore and requires parameters for backup location, filename, features, and nodes to restore
utils disaster_recovery status	Displays the status of ongoing backup or restore job
utils disaster_recovery show_backupfiles	Displays existing backup files
utils disaster_recovery cancel_backup	Cancels an ongoing backup job
utils disaster_recovery show_registration	Displays the currently configured registration
utils disaster_recovery show_tapeid	Displays the tape identification information

