The Cisco Unified MeetingPlace server is a carrier-class conferencing platform with several built-in redundancy and reliability features. This section describes such features and provides recommendations for disaster recovery planning when the following conditions occur:

- Data network failure
- Cisco Unified MeetingPlace 8106 or Cisco Unified MeetingPlace 8112 failure
- Site disaster

Contents

- 1 Data Network Failure
- 2 Cisco Unified MeetingPlace Server Failure
  - 2.1 Table: Server Redundancy Features
- 3 Using the Shadow Server to Recover Site Disaster
- 4 Importing Meetings to Recover Site Disaster
  - 4.1 To Generate a Report
  - 4.2 To Import Meeting Details into a Target Server
- 5 Importing and Deleting Meetings in Batch Mode
  - 5.1 To Import Meetings in Batch Mode
  - 5.2 Table: Batch Mode Import Options
  - 5.3 Table: Batch Mode Import Parameters
- 6 Deleting Meetings in Batch Mode

Data Network Failure

When a data network failure occurs, assuming that the voice network is still in operation, the Cisco Unified MeetingPlace server continues to operate. This means that previously scheduled meetings can start and meetings that are taking place can continue.

Scheduling future or immediate meetings, however, is limited to the voice user interface only, and no notifications are available until the data network is backed up. Web conferences are affected because the data network is critical to web conferencing.

Cisco Unified MeetingPlace Server Failure

The Cisco Unified MeetingPlace server can handle several critical component failures without affecting service. Table: Room Conditions describes the types of redundancy features.

Table: Server Redundancy Features

Features	Description
N+1 hot swappable redundant power supply/fans	If a power supply fails, the system continues to operate with no experienced downtime. In addition, a new replacement power supply can be hot swapped (inserted while the system is up and running).
Duplicate databases	The server database is stored on two separate disk drives (Disk 1 and Disk 2). One is established as the primary drive, the other as a backup. Assuming that the database on Disk 1 is the active one, the system copies all database contents to the database on Disk 2. This process occurs in the background, while the system is operating. It does not affect system performance. If Disk 1 fails, the system can be restarted using Disk 2 with the latest copy of the database. If Disk 2 fails, the system continues to operate. However, voice files on Disk 2 can be lost in this case. **Note:** Make sure that you remove the failing drive and replace it as soon as possible. For troubleshooting, see the *MeetingPlace Server Disk Failure Recovery* tech note at http://www.cisco.com/en/US/products/sw/ps5664/ps5669/prod_tech_notes_list.html. For more information about replacing a drive, see the *Hard Disk Replacement Procedure for Cisco Unified MeetingPlace 8106* or the *Hard Disk Replacement Procedure for Cisco Unified MeetingPlace 8112* , at http://www.cisco.com/en/US/products/sw/ps5664/ps5669/prod_maintenance_guides_list.html.
Network backup	We recommend that network backups be performed once each night, in addition to available disaster recovery mechanisms. **Caution!** Before performing a network backup, Cisco Unified MeetingPlace Network Backup must be installed. (For more information about backing up the database, see Cisco Unified MeetingPlace Network Backup Gateway, Release 5.3.)

Using the Shadow Server to Recover Site Disaster

One site disaster recovery method is the Cisco Unified MeetingPlace Shadow Server. The Shadow Server is a backup system that can replace the conference server in case of system or site failure. Switching over to the shadow server ensures that complete functionality can be returned with minimal loss of time and disruption in service.

For more information about the shadow server, see Configuring a Shadow Server.

Importing Meetings to Recover Site Disaster

The Import Meeting feature allows a Raw Meeting Details report containing all meeting information-including the scheduler, date, time, the Meeting ID and number of reserved ports-to be imported into a spare conference server in a separate location. The import process is handled through MeetingTime.

After the information is imported, the system creates a meeting with the same meeting ID and sends notifications to the meeting scheduler with the correct phone number.

To prepare for disaster recovery, generate the Raw Meeting Details report at least once a day and store it on the network in a location different from the physical server location.

This feature supports importing meetings from a Cisco MeetingPlace 2001 (4.3) system to a Cisco Unified MeetingPlace Audio Server 5.0 system and later, but the reverse is not true.

User profiles and meeting categories must be synchronized on the active and target servers before attempting to import meetings. If a meeting is imported onto the target server without a corresponding scheduler profile, the meeting schedules using the user ID in the Scheduler User ID field, which defaults to the user ID of the system administrator running the import.

For information about best results when importing, see the Importing User Profile and Group Information.

Do the following procedures, in the order presented, to import meetings:

- To Generate a Report
- To Import Meeting Details into a Target Server

To Generate a Report

We recommend that you establish a folder into which to import these reports.

1. In the MeetingTime Report tab, for **Report Type** , select **Raw Meeting Details Info** .
2. Enter values for the **Start Date** and **End Date** attributes.
   Details for all meetings scheduled to start between these dates are exported.
3. For the **Destination** attribute, choose **File** and click **OK** .
4. For the **Output File** attribute, choose **Browse** . In the Save As window, choose the folder to hold the output file, and click **Save** . Then click **OK** to accept your output file destination.
5. Click **Generate Report** .

The Import Meetings feature only imports meetings scheduled to start within 30 minutes of import meeting completion and after. To import details of all future meetings scheduled beyond the current date, you must set the **End Date** with the following parameter in mind: **Max adv days to schedule** determines how many days before a proposed meeting you can schedule the meeting (in the Configure tab, Scheduling Parameters topic). To include all future scheduled meetings, make the end date equal to the current date plus the **Max adv days to schedule** parameter.

For example, if today is June 1, 2005 and your Max Adv Days to Schedule parameter is set to 20, your End Date will be June 21, 2005 to ensure that all scheduled meetings are captured.

To Import Meeting Details into a Target Server

1. In the MeetingTime **System** tab of the target server, for **Actions** , select **Import Meetings** .
   If the target server is a network server rather than a standalone server, the system administrator has access to four additional fields: Preferred server, Force on preferred server, Preferred site, and Force on preferred site.
   Recurring meetings are imported as separate, one-time meetings rather than chains of meetings. Continuous meetings are imported as continuous meetings with the import meeting date as the start date.
2. For the **Data File to Use** attribute, choose **Browse** .
3. In the Open window, choose the generated report file, then click **OK** .
4. For the **Override w/Sched. ID** attribute, choose a value:
   - When Yes, the user identified in the Scheduler User ID field schedules all the meetings in the import file.
   - When No, the system schedules the meetings using the scheduler that originally scheduled the meeting. If the scheduler does not exist for any meeting, the user identified in the Scheduler User ID field schedules the meeting.
5. For the **Send Log Info To** attribute, choose an option.
6. For the **Log File Name** attribute, choose **Browse** . Enter the log file name, then click **OK** .
7. For the **Error Threshold** attribute, enter the number of errors you expect in the field.
   If the number of errors encountered exceeds the number indicated in the Error Threshold field, the system stops importing meetings.
8. Click **Execute** to complete importing meetings.
   The target server generates meetings and sends e-mail notifications to meeting organizers affected by the import. These notifications are standard meeting notifications containing the name of the meeting organizer, the meeting ID, and the phone number of the target server so that participants can access their meeting.
9. To view a record of your meeting import, open your directory to locate the appropriate file.

Remember the following information:

- If you selected **Screen** as your destination for log information, your Meeting Import Information Record displays after you click **Execute** .
- Import Meetings generates new meetings for all meetings in progress and those started within 30 minutes of import meeting completion. If a meeting started more than 30 minutes before meeting import completion, it appears in the Past Meetings field of the Meeting Import Information Record.
- The notification feature is limited by the space available in your notification queue. You can check your notification queue status in the System tab in MeetingTime. When 90 percent of the notification

queue is full, the system sends a major alarm. When the notification queue reaches its limit, the system sends another major alarm (see sample below) and drops incoming notifications:

```
REFNO SEV CODE COUNT FIRST LAST UNIT
---------------------------------------------------------------------
481684) MAJ 0f0088 1 Aug 3 00:28 Aug 3 00:28 0 SW MODULE=18 PO
mailbox full!
Mbox 2, NumMsgs 10000, MaxMsgs 10000 /tt>
<tt> 477777) MAJ 0f0089 1 Aug 2 23:53 Aug 2 23:53 0 SW MODULE=18 PO mailbox
90 percent full! Mbox 2, NumMsgs 9000, MaxMsgs 10000
```

- To ensure that you receive notifications with your imported meetings, we recommend that you limit the number of meetings imported to 10,000 per import. (If you do not care whether you receive notifications, import everything as suggested in the previous steps.) The target server does not disperse e-mail notifications to participants. We recommend that meeting organizers forward their e-mail notification to all meeting participants.

Importing and Deleting Meetings in Batch Mode

To avoid tying up your machine, import and delete meetings overnight in batch mode. You can perform all import operations in batch mode using command line options.

To Import Meetings in Batch Mode

1. Create the import file.
   The import file is the output from the Raw Meeting Details Report. For information about generating the Raw Meeting Details Report in batch mode, see the To Generate Batch Reports in Windows Batch Mode. The only required field for importing is StartDateTimeOfConfGMT, the date and time of the meeting:
   Any fields not included in the import file derive their value from the default meeting preferences of the scheduler's profile.
2. Run MeetingTime using a command line option shown in Table: Server Redundancy Features.
   Begin the command line string with your login information, unless you have stored this information in the MeetingTime Settings. Use the following syntax:
   ```
   C:\Program Files\Latitude\Mtgtime\mtgtime.exe [user ID][user
   password][hostname]
   For example: C:\Program Files\Latitude\Mtgtime\ mtgtime.exe
   tsmith
   MyPassword Server-CA
   ```
   Or, enter a command line using Windows 95, 98 or Windows 2000 (choose Windows **Start** menu > **Run**).

Table: Server Redundancy Features shows the command line options. Parameter information is explained after the table.

Remember the following information:

- ImportMeetings is a keyword, which must be followed by an import file name in your command line string.

- All other parameters are optional, but positional, which means that parameters occurring later in the command line (for example, Error threshold) must be preceded by all of the preceding optional parameters.

Table: Batch Mode Import Options

Options	Sample Import Strings
Importing Meetings	importMeetings [Data file to use] [Add/Delete flag] For example: importMeetings rawmtdet.csv yes
Importing Meetings with all optional fields defined	importMeetings [Data file to use] [Add/Delete flag] [Scheduler user ID] [Override?] [Log file name] [Error threshold] For example: importMeetings rawmtdet.csv yes tsmith no mylogfile.txt 50
Deleting Meetings	ImportMeetings [Data file to use] [Add/Delete flag] For example: ImportMeetings rawmtdet.csv no
Deleting Meetings with all optional fields defined (When deleting meetings, the Scheduler User ID field is not used but still needs to be specified)	importMeetings [Data file to use] [Add/Delete flag] [Scheduler user ID] [Override?] [Log file name] [Error threshold] For example: importMeetings rawmtdet.csv yes tsmith no mylogfile.txt 50

Caution! Enter all the parameters shown, including import file name, as one line (do not press Enter until after you enter the import file name).

Use Table: Batch Mode Import Parameters when replacing parameters with values.

Table: Batch Mode Import Parameters

Parameters	Sample Import Strings
Add/Delete flag	Used to specify whether the meetings in the import file should be added (scheduled) or deleted (canceled). Default for this field is Yes. • If this flag is Yes, the meetings are scheduled. • If this flag is No, the meetings are deleted.

Scheduler User ID	The scheduler user ID command line field is used for overriding or filling in the scheduler ID for some or all records in the import file.
Override w/Sched. ID flag	Used when the specified scheduler ID should schedule all the meetings in an import file. If the override flag is not specified, it defaults to no. • When the override flag is yes, the user identified in the scheduler ID command line field schedules all the meetings in the import file. • When the override flag is no, the scheduler ID command line field is used to fill in the scheduler ID for any records in the import file that have a missing or invalid scheduler ID. All other meetings in the import file are scheduled by the original scheduler as found in the import file. • If a scheduler ID is not specified, it defaults to the ID of the person performing the batch operation. • If the scheduler ID specified on the command line is invalid or does not have sufficient rights to schedule a meeting, the ID of the person performing the batch operation is used.
Log file name	The log file name allows users to specify a file into which error and status information is logged. If the log file name is not specified, it defaults to loginfo.txt. The log file name and path cannot contain spaces.
Error threshold	The error threshold specifies the number of errors that MeetingTime allows before aborting the import process. If not specified, the error threshold defaults to *50* .

Deleting Meetings in Batch Mode

When deleting meetings, the import file must have been generated from the same system. The following fields are required when performing a delete:

- ConfNum
- DialableConfID
- StartDateTimeOfConfGMT