

You can optimize resource usage by the way you designate port types and manage port scheduling. The following sections describe managing ports.

Contents

- [1 Designating Port Types](#)
 - ◆ [1.1 Figure: Distribution of Port Types](#)
- [2 Designating Overbook Ports](#)
- [3 Configuring Ports and Port Groups to Improve Performance](#)
- [4 Configuring Individual Ports](#)
- [5 Configuring Port Groups](#)
 - ◆ [5.1 Table: Configuring Port Groups](#)
- [6 Managing Port Scheduling with Guard Times](#)
 - ◆ [6.1 Table: Resources Managed with Guard Times](#)
 - ◆ [6.2 Figure: How Guard Times Affect Users](#)
- [7 Setting Guard Time Parameters](#)
 - ◆ [7.1 To Specify Parameters for Consumable Resources](#)
 - ◆ [7.2 Table: Guard Time Parameters](#)
- [8 Managing Port Scheduling with Run Times](#)
 - ◆ [8.1 Table: Managing Port Scheduling with Run Times](#)
 - ◆ [8.2 Figure: Interaction of Release Time Parameters](#)
- [9 Setting Immediate Meeting Parameters](#)
 - ◆ [9.1 To Set Parameters for Immediate Meetings](#)
- [10 Collecting Data About Port Utilization](#)
 - ◆ [10.1 Configuration](#)

Designating Port Types

The connection between Cisco Unified MeetingPlace and the phone network is called a port. The ports that connect Cisco Unified MeetingPlace to the phone network are classified as access ports. Access ports are the total number of possible simultaneous connections to Cisco Unified MeetingPlace. All ports are automatically designated as access ports.

A subset of your access ports are pre-configured as conference ports. There is no difference between an access port and a conference port. Access ports that are not configured as conference ports are used for scheduling meetings and listening to recorded meetings.

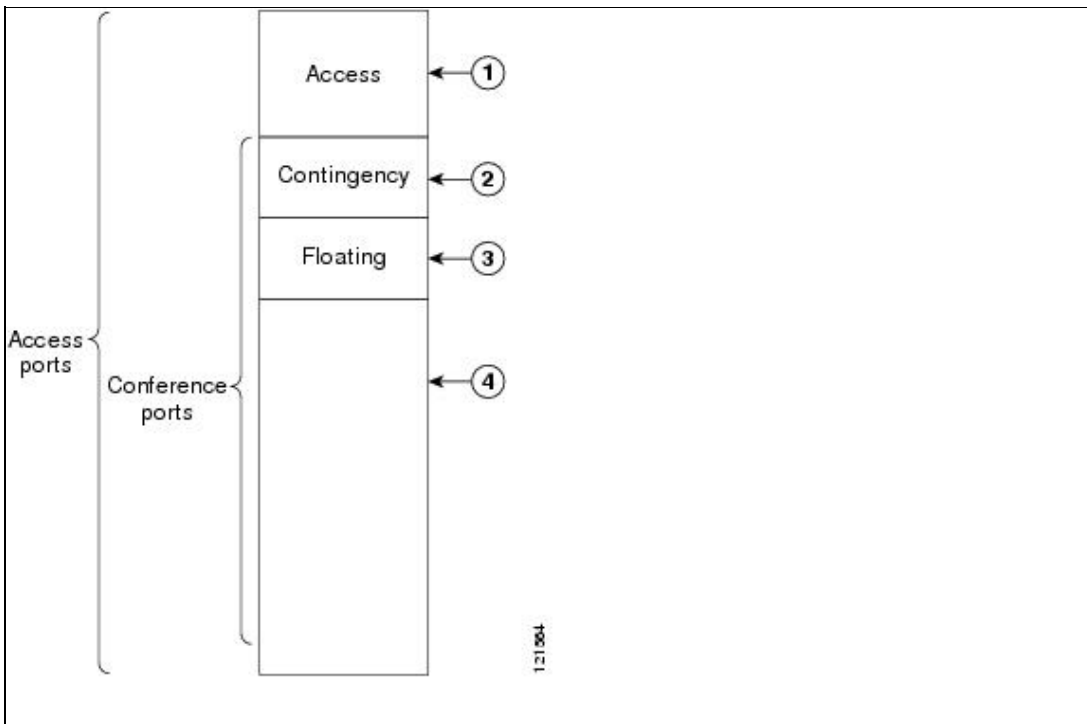
Video ports reside physically on the Cisco Unified Videoconferencing MCU, but they are a resource that can be scheduled in advanced within Cisco Unified MeetingPlace. Unlike access and conference ports, video ports are not licensed individually. Instead, they are all enabled when the Cisco Unified MeetingPlace Video license is enabled.

Figure: Distribution of Port Types shows the following:

- All access ports are not necessarily pre-configured as conference ports. The number of conference port licenses you acquire (as listed on your sales order) is the total number of conference ports in your system.
- Distribution of port types illustrates the distribution of ports in a Cisco Unified MeetingPlace system.
- You can reserve some conference ports to handle call transfers. These are contingency ports. Contingency ports are ports that the system keeps in reserve, making it possible for meeting participants to reach a contact or attendant for assistance during a meeting and for the system administrator to dial in to meetings. The values entered in the Number of Contingency Ports field (Configure tab, Server Configuration topic) determine the number of contingency ports in your system.
- The remaining conference ports, called floater ports, are for unexpected port needs. Floater ports can float between meetings, taking up the slack when an extra person attends a meeting that is already full and when ports that can be scheduled in advance are not available. For example, if someone joins a meeting that is already full, the server tries to reschedule the meeting for an additional port. If all ports are reserved for other meetings, the server tries to find a floater port for the person joining the meeting. The values entered in the Number of Floating Ports field (in the Configure tab, Server Configuration topic) determine how many ports are floater ports.

Note Floater port parameters apply only to audio ports; they do not apply to video ports. There is no concept of floater ports with respect to video.

Figure: Distribution of Port Types



1	Access ports reserved to schedule and listen to recorded meetings	3	Floater ports, reserved to handle unexpected meeting attendance
2	Contingency ports reserved to handle call transfers to contacts for attendant	4	Ports in use or reserved for meetings

We recommend that you dedicate 1-2 conference ports as contingency ports and 15 percent of the total number of conference ports as floater ports.

For reservationless meetings floater point recommendations, see the [About Reservationless Meetings](#).

Designating Overbook Ports

You designate overbook ports for audio conference needs. The values entered in the Overbook Ports field (which is in the Configure tab, Server Configuration topic) allow the system to schedule more ports than are actually available. Suppose you have 96 conference ports and all 96 are booked for a meeting. If you set Overbook to 2 , a user can still schedule another two-person meeting at the same time.

Overbooking assumes that all users who are scheduled to attend often do not attend, which usually leaves unused ports available. In the rare case that all ports are scheduled and all people attend the meeting-including those who are overbooked-the last people to call into any meeting would not be able to get through. If such a meeting were critical, the system administrator or attendant could ask the overbooked participants to reschedule for later.

For reservationless meetings overbook ports recommendations, see the [About Reservationless Meetings](#).

Overbook port parameters apply only to audio ports; they do not apply to video ports. There is no concept of overbook ports with respect to video.

Configuring Ports and Port Groups to Improve Performance

Cisco Unified MeetingPlace provides several parameters with which you control the port configuration to maximize users' access to the system. In most cases the default parameters (or the parameters set by a Cisco Network Consulting Engineer (NCE) during installation) will suffice. However, as users increase their interest in the system, you may want to alter the port configuration to improve performance.

The Cisco Unified MeetingPlace database includes individual ports and port groups. By organizing ports into groups, you can configure multiple ports at one time.

Configuring Individual Ports

Port configuration information includes a default port access type. When the system receives dialed number information (DID/DNIS or DDI), the DID access plan determines which port access type is assigned to a call.

With DID access, the number that the user dials causes a particular set of digits to pass to the Cisco Unified MeetingPlace server. Each digit set is assigned to one of the service types (scheduling, profile access, access to recorded meetings, participation in a meeting). The call can come in on any port, so that the ports are dynamically allocated. If, for example, all users call the number that corresponds to recorded meetings access, then all unscheduled ports are allocated to recorded meetings, and no new meetings can be scheduled.

With direct access, ports are assigned to a particular service type. For example, ports 1-3 to scheduling, 4-5 to recorded meetings, and 6-12 to meetings. The number that users dial determines which ports they access, and thus which service they receive.

When the system does not receive dialed number information, the default port access type assigned to a call is the type assigned in the Ports topic (Configure tab) to the port that is handling the call. If your system does not receive dialed number information, you may want to map out the connections between Cisco Unified MeetingPlace and the phone network. Then, either decide which access type to assign to each line or group of lines that connects to Cisco Unified MeetingPlace, or choose to use all ports universally with combined access.

If users cannot access a particular service, you can reconfigure the number of ports assigned to each Cisco Unified MeetingPlace service using the Ports topic.

Remember the following information:

- If you assign an individual port to a port group, you do not have to reenter data that is already provided for the port group.
- You can configure ports to allow or disallow outdialing. Verify that the trunk is configured in Cisco Unified MeetingPlace the same way your telephony provider has configured it (it can either make outgoing calls or cannot make outgoing calls). The system can block certain numbers and dialing patterns system-wide, by group, or by individual user. Call blocking is controlled by an internal Cisco Unified MeetingPlace table.
- You can limit the number of ports for scheduled meetings; however, you cannot limit the number of ports for reservationless meetings.

Configuring Port Groups

Because multiple ports may use identical settings, the Cisco Unified MeetingPlace database uses port groups to define the settings for multiple ports simultaneously. You define a port group by selecting the Port Groups topic in the Configure tab.

Some information for individual ports is included in the port group definition. You do not have to provide

this information twice when a port is a member of a port group. (Group information has precedence over individual port information.)

The configuration information you enter to define a port group, which includes a port access type, applies by default to all ports that are assigned to the group in the absence of individual port configuration information or a DID access plan. [Table: Configuring Port Groups](#) describes the port groups.

Table: Configuring Port Groups

Attribute	Description	Choices and Recommendations
ID number	A number from 0 to 31 that identifies this port group. This number is predefined.	Assign every line coming into Cisco Unified MeetingPlace to a port group. The port group ID number is your choice.
Active	Whether to use this port group definition. You may want to define port groups now for later use.	Yes or No Most likely, port groups are always active. You may wish to make port groups inactive during servicing.
Provider	An abbreviation that identifies the service provider.	Alpha, Numeric, or Both
Circuit ID	A number that identifies the circuit.	Assigned by the circuit provider
Card Type	Whether this trunk is a digital T1, IP, or unassigned.	T1, IP, or None (unassigned)
Signaling Protocol	The signaling type used by this trunk.	Loop start, ground start, E&M wink start, immediate start, DID/DDI, clear channel, IP, protocol table, any E1
Protocol Table	Contains the configuration information for the type of signaling used.	All T1 PRI systems are shipped from the factory with protocol table 2 set to use the default setting of ATT PRI protocol; protocol 3 to use Nortel PRI; and table 4 to use Bell PRI.
# of DID Digits Expected	Indicates the number of digits sent by the PBX or network.	Your choice (0 or a number from 2 to 6; 1 is not valid)
Mtg. ID for Direct Mtg. Access or Default Digits for System	Indicates which meeting callers to this port will attend when DID/DDI digits = 0 and the default access type = "meeting."	Meeting ID or DID/DDI number of the meeting in which you want callers to be placed.
Default Access Type	Which access type is applicable to all ports in the group in the absence of dialed number information.	DID Meeting, Profile, or Combined Your choice depends on the level of access you want users to have on these ports.
Language	The language in which prompts will play.	English (USA), English (UK), Japanese, French-Canadian, French (France), German, Portuguese (Brazil), Spanish

		(Americas), or No Language When No Language, callers hear a language prompt when they dial in to the Cisco Unified MeetingPlace Audio Server system.
Human Assistance	When someone in a meeting needs help, Cisco Unified MeetingPlace uses another port to connect that person (via a call transfer) to a contact or attendant. The value in this field determines whether the ports in this group can be used for call transfer during a meeting.	Yes or No Choose Yes, assuming that an individual who can provide assistance to users can be reached on this port group.
Flash Transfer	Whether these ports initiate a call transfer by using a hook flash.	Yes or No Your telephony network provider can provide this answer.
Outdial	Whether these ports can be used for outgoing calls.	Yes or No For most applications, you should set ports to handle outdial.

Managing Port Scheduling with Guard Times

Guard times ensure that meetings do not overlap each other or that two meetings are not scheduled back-to-back with the same meeting ID.

Guard times become part of the scheduled meeting record. For example, if you schedule a meeting with the system start and end guard time parameters set at 0 minutes, then change the system start and end guard times to 60 minutes, your meeting will not take on the new guard time settings. To ensure the meeting in the example takes on the current system start and end guard time minutes, you must reschedule your meeting. Back-to-back reservationless meetings are permitted regardless of guard times. In effect, the guard time for reservationless meetings is equal to zero.

Cisco Unified MeetingPlace uses guard times to manage three resources, as shown in [Table: Resources Managed with Guard Times](#).

Table: Resources Managed with Guard Times

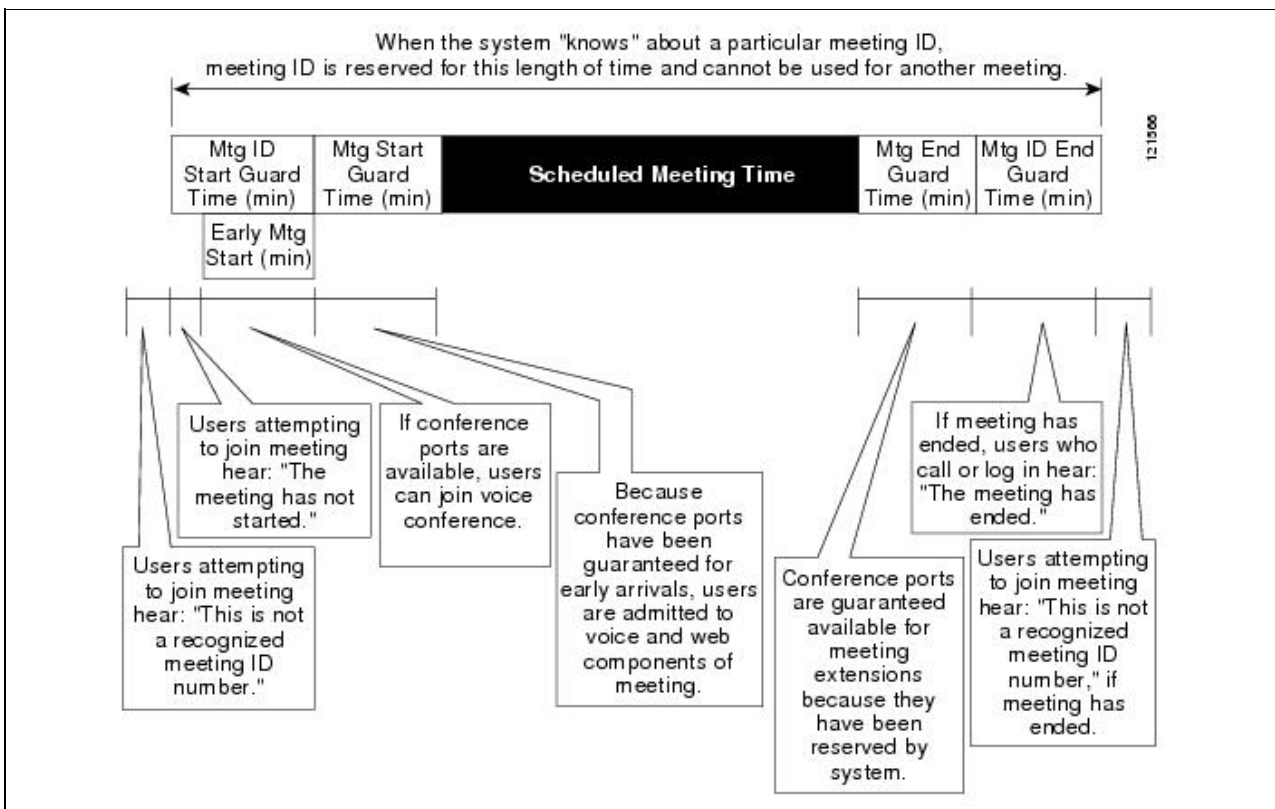
Resource	Description
Conference ports	Every Cisco Unified MeetingPlace conference server has a finite number of ports available for users, as described earlier. The parameters described in this section determine how long before, during, and after a scheduled meeting the ports are reserved (and therefore when the ports become available to other users).
Meeting IDs	Every meeting has an ID number, which is how users identify the meeting they want to attend. During any given time, a meeting ID must be unique, so that the system knows which meeting

	to put a caller into. Meeting guard time fields determine how long before and after a meeting the meeting ID is protected. This protection prevents another person from scheduling another meeting at the same time with that particular meeting ID number. The meeting ID guard time also defines the period before and after a meeting that the caller hears that the meeting has not yet started or has ended, rather than hearing that Cisco Unified MeetingPlace does not recognize the meeting.
Voice storage	Guard times reserve space for recording a meeting (if the organizers have elected to record the meetings at the time they scheduled the meetings).

Caution! When using the Reservationless Meetings feature, profile numbers cannot match existing meeting IDs, because reservationless meetings use profile numbers as reservationless meeting IDs.

Figure: How Guard Times Affect Users shows how guard time settings affect users.

Figure: How Guard Times Affect Users



Setting Guard Time Parameters

To Specify Parameters for Consumable Resources

1. In the MeetingTime Configure tab, select the **Scheduling Parameters** view.
2. Set values for the parameters, as shown in [Table: Guard Time Parameters](#).

Table: Guard Time Parameters

Parameter	Description
Mtg. Start Guard Time (min)	<p>Number of minutes in advance of the requested meeting start time that ports are reserved for the meeting. Cisco Unified MeetingPlace automatically adds this advance time to the requested start time. This time guarantees ports for early arrivals to the voice component of a meeting.</p> <p>Set this parameter to 0 minutes and use Mtg. ID Start Guard Time to ensure that Cisco Unified MeetingPlace acknowledges the meeting and users hear "The meeting has not started" when they try to enter the meeting.</p>
Mtg. End Guard Time (min)	<p>Number of minutes after the scheduled meeting end time that ports, or a "zero-port" meeting, remains reserved. The system adds this time to the reserved end time of the meeting. This time guarantees ports when a voice meeting overruns its scheduled end time.</p> <p>Set this parameter to 0 minutes to be able to schedule back-to-back meetings. Set this parameter for a longer period if your users tend to have longer meetings and/or your system is less busy. When your guard time expires and no ports are available, you cannot extend the voice component of your meeting.</p> <p>Although non-zero guard times guarantee ports before and after a meeting, most users simply schedule some "padding" into their meetings. For example, if a user wants an hour-long meeting and wants to pad it to allow callers time to enter the meeting, listen to the roll call and agenda, and greet one another, the user would simply schedule the meeting for 70 minutes. In this case, managing guard times can be an unnecessary complication.</p> <p>Note: In Cisco Unified MeetingPlace Audio Server, future meetings scheduled with a non-zero value for audio or video ports will not extend if there are less than two audio or video participants in the conference, regardless of the number of active Data Conference participants at extension time. To enable a future meeting of only Data Conference participants to extend, schedule the meeting with zero audio and video ports. Reservationless meeting extension behavior has not changed.</p>
Mtg. ID Start Guard Time (min)	<p>Number of minutes before a meeting that the associated meeting ID or DID/DDI number is reserved. Both this parameter and Mtg. ID end guard time control when meeting IDs become available to be re-used. This is to ensure that two meetings do not have the same ID at the same time, and callers do not inadvertently enter the wrong meeting.</p> <p>If users try to enter the meeting before the meeting ID start guard time begins, they are told, "This is not a recognized meeting ID number." Callers in the voice conference hear it as a voice prompt. Users logging into a web conference see it on their screen.</p> <p>After the Mtg. start guard time begins, callers who try to enter the voice conference hear the message "The meeting has not started." The recommended value for this parameter is 15 minutes.</p>
Mtg. ID End Guard Time (min)	<p>Number of minutes after a meeting that the associated Meeting ID or DID/DDI number is reserved. If callers call into a voice conference during this period, they hear "The meeting has ended." Users logging into a web conference see it on their screen. After this period,</p>

users are informed, "This is not a recognized meeting ID number." The recommended value for this parameter is 15 minutes.

Remember the following information:

- If two people are on a zero-port meeting and forget to log out, the meeting continues extending itself for up to 24 hours. The meeting expires if it encounters a meeting ID conflict.
- Meeting guard time minutes and meeting ID guard times are added to the meeting start and end times. For example, if the start and end guard times are both 10 minutes, the meeting itself is scheduled for an hour, and the Start and End Meeting ID guard times are set to 10 minutes, the total time reserved for the meeting ID is one hour and forty minutes.

Managing Port Scheduling with Run Times

Run times determine how early someone can call into a meeting before its scheduled start time, how long a meeting can be extended, and when ports should be released.

Run-time parameters take effect at the start of each meeting. As each new meeting begins, the system looks to the current run-time settings to determine how to manage the meeting.

Table: Managing Port Scheduling with Run Times describes the parameters used to manage port scheduling with run times.

Table: Managing Port Scheduling with Run Times


Parameter	Description
End of Mtg. Warning (min)	Amount of time before meeting will end that users are first warned. (Recommended time is 10 minutes.) Users receive another warning two minutes before the meeting ends.
Extend Meeting (min)	Amount of time a meeting is extended if the meeting continues past scheduled end time and ports are available (0 to 60 minutes). Callers receive warnings informing participants when the meeting will end. The system can extend the meeting time if ports are not reserved for other meetings. The recommended value for this parameter is 15 minutes. End of meeting announcements and meeting extension announcements can be configured for individual meetings. However, the number of minutes set for each parameter is system wide. Neither warning is supported during a zero-port meeting.
Early Mtg. Start (min)	Maximum amount of time before a meeting starts that early arrivals can enter the meeting. This time is added to the meeting start guard time, but must be less than the meeting ID start

	<p>guard time. The recommended value for this parameter is 15 minutes.</p> <p>For example, if the conference start time is 9:00, the Mtg. start guard time is 15 minutes, and the early meeting start minutes is 15 minutes (and ports are available), callers can enter the meeting at 8:30.</p> <p>For video conferences, this parameter must match the value on the Video Administration for Cisco Unified MeetingPlace server.</p>
Disconnect Empty Port (min)	<p>Amount of time that ports are held after everyone hangs up or one person remains on a web conference. This parameter accommodates longer meetings where people might want to take breaks and re-enter the meeting.</p> <p>This parameter also determines how long to hold ports after the scheduled start time, if no one shows up for the meeting.</p>
Early Meeting Release (min)	Amount of time before the scheduled end time of a meeting when ports are released if none are in use for that meeting.

The early meeting release feature takes effect in a zero-port meeting when only one person remains on the web conference.

Figure: Interaction of Release Time Parameters illustrates the interaction of the release time parameters.

Figure: Interaction of Release Time Parameters



Setting Immediate Meeting Parameters

Immediate meetings (meetings that users start right away) also use ports. Immediate meetings do not use pre-meeting guard times; however, they do use resources while they are in progress and use Mtg End Guard Time, Mtg End ID guard time, Disconnect Empty Port, Early Mtg Release, Extend Meeting parameters, and the parameters for reserving space for recording meetings.

The system administrator sets the default port and meeting length, and the meeting scheduler can change the

defaults.

To Set Parameters for Immediate Meetings

1. In the MeetingTime Configure tab, select the **Scheduling Parameters** view (under Company Specific Information).
2. Scroll to the Immediate Meetings attributes, and set the following attributes:
 - ◆ **# of Ports to Schedule** -The default number of conference ports offered by the system when callers ask for immediate meetings. These ports are not reserved for general use by immediate meetings. A caller can only schedule an immediate meeting if ports are available.
 - ◆ **Length of Meeting (Min)** -The default meeting length offered by the system when callers ask for immediate meetings.
3. Click **Save Changes** .

Remember the following information:

- If you cannot schedule back-to-back meetings, make sure that the meeting start and end guard times are set 0 (zero). These values must be 0 to allow scheduling back-to-back meetings.
- To ensure that users hear or see "The meeting has not yet started" rather than "That is not a valid meeting ID" if they call or log in too early, set the meeting ID start and end guard times to be large values, such as 60 minutes.
- Reservationless meetings use immediate meeting parameters to determine the length and size of the meeting. When the Reservationless Meetings feature is enabled and you try to schedule an immediate meeting (or a scheduled meeting with the start time of now or earlier) and do not specify the meeting ID, the meeting becomes a reservationless meeting. (For more information about reservationless meetings, see the [About Reservationless Meetings](#).)

Collecting Data About Port Utilization

You can collect port information using the port data collection utility. The utility collects port data and transfers it to a gateway unit configured for file transfers via the GWSIM, where the data can be collected and manipulated as desired. Collection and transfer occur at a specified time interval based on configured settings. Transfer is accomplished using the "gwputfile" utility.

The utility consists of a script "/lat/bin/getportdat.sh" which drives the data collection and transfer operations. The script is invoked as the last step of the Audio Server startup.

NOTE: The script is not meant to be run interactively, and will exit with an error if invoked from the command line.

Logging is done to the file `?/tmp/cpserver/getportdat.log?` on the Audio Server. During startup, if a previous log exists, it is moved to a `?old?` version.

Configuration

To enable the utility to collect the port data, you must configure various parameters in the `/lat/etc/getportdat.cfg` file. The startup script initially generates this file and for the initial configuration, it uses the default values. These default values render the script inactive; that is, it exits without performing any actions.

The default values are as follows (this data is taken directly from the configuration file):

UNIT=xx

REMOTEFILE=""

LOCALFILE=/tmp/getportdat.out

INTERVAL=60

Parameter	Description
UNIT	Must contain a unit number between 16 and 31. If you enter an invalid unit, the script will log an error and exit.
REMOTEFILE	A valid path on the gateway unit; must be configured in the transfer destination field of the GWSIM on that unit. If you enter an invalid entry, gwputfile will return an error. The return code of the first call to gwputfile (but only the first, to avoid flooding the log) is recorded in the log file and should be 0; non-zero error codes need to be investigated. The assumption here is that a particular configuration will be verified and correct initially, and that error conditions will be presumed to be transient.
LOCALFILE	A valid path on the Audio Server; the default shown here is acceptable. If you enter an invalid entry, the script will log an error and exit.
INTERVAL	Must contain a valid interval between 60 and 300. If you enter an invalid interval, the script will use the minimum interval (60s).