



About Cisco Unity Failover

This chapter contains the following sections:

- [How Failover Works in Cisco Unity, page 4-1](#)
- [Requirements for Cisco Unity Failover, page 4-4](#)
- [Effects of Using or Not Using the Force Failover Setting, page 4-5](#)
- [Effects of Failover and Failback on Calls in Progress, page 4-6](#)
- [Status Monitoring and File Replication, page 4-6](#)
- [Events When Failover Occurs, page 4-9](#)
- [Events When Failback Occurs, page 4-9](#)
- [Automatic and Manual Failback, page 4-10](#)
- [Causes of Failover and Failback, page 4-10](#)
- [Intervals for Failover and Failback, page 4-12](#)
- [Causes of Both Servers Becoming Active at the Same Time, page 4-14](#)
- [Effects of Shutting Down and Restarting the Primary and Secondary Servers, page 4-15](#)
- [Licensing Restrictions on Using a Secondary Server Without a Primary Server, page 4-15](#)

How Failover Works in Cisco Unity

Failover is a feature that provides a simple redundancy, allowing voice messaging functions to continue if the Cisco Unity server fails or when you need to perform maintenance. To set up failover, you install and configure Cisco Unity on two servers, a primary server and a secondary server.



Caution

The failover feature cannot be used for continuing Cisco Unity service (voice messaging) on one server while upgrading the Cisco Unity software on the other server. Both primary and secondary servers must be out of service while the Cisco Unity software is upgraded. The secondary server cannot handle voice messaging while the primary server is being upgraded.

[Figure 4-1](#) shows a failover configuration for an IP (Internet Protocol) phone system. [Figure 4-2](#) shows a failover configuration for a circuit-switched phone system in a DTMF integration with voice cards. [Figure 4-3](#) shows one possible failover configuration for a circuit-switched phone system in a serial

integration (with voice cards) that uses a data splitter on the RS-232 serial cable to provide a data link to both Cisco Unity servers. For information on making the line connections, see [Appendix B, “Line Connections Between the Phone System and the Cisco Unity Servers.”](#)

Figure 4-1 Failover Configuration for an IP Phone System




Figure 4-2 Failover Configuration for a Circuit-Switched Phone System in a DTMF Integration with Voice Cards


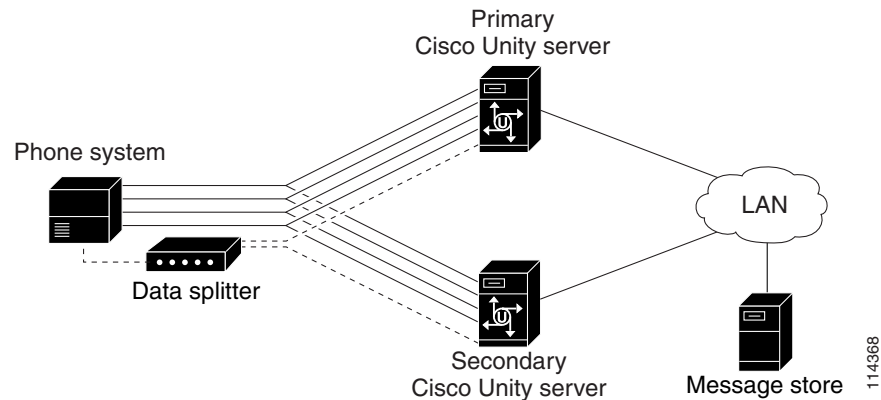


Figure 4-3 One Possible Failover Configuration for a Circuit-Switched Phone System in a Serial Integration with Voice Cards



Under normal circumstances, the primary server is active—Cisco Unity answers phone calls and takes messages, sends message notifications, and turns message waiting indicators (MWIs) on and off. The secondary server is inactive—Cisco Unity is running, but it does not perform any voice messaging functions.

Table 4-1 shows the behavior of the two servers when the primary server is active.

Table 4-1 Cisco Unity Behavior When the Primary Server Is Active

Cisco Unity Behavior	Primary Server	Secondary Server
Running	Yes	Yes ¹
Active (handles voice messaging processes)	Yes	No
Answers phone calls	Yes	No
Sends message notifications	Yes	No
Turns on and off MWIs	Yes	No

- Typically, Cisco Unity on the secondary server is running unless the server is shut down (for example, when the server is undergoing maintenance).

If the primary server fails or if the Cisco Unity service on the primary server stops, the secondary Cisco Unity server automatically becomes active and starts performing standard Cisco Unity operations. This shift from primary to secondary Cisco Unity servers is called failover. If you want to stop the primary Cisco Unity server for maintenance, you can also initiate failover manually.

Table 4-2 shows the behavior of the two servers when the secondary server is active.

Table 4-2 Cisco Unity Behavior When the Secondary Server Is Active

Cisco Unity Behavior	Primary Server	Secondary Server
Running	Depends ¹	Yes
Active (handles voice messaging processes)	No	Yes
Answers phone calls	No	Yes
Sends message notifications	No	Yes
Turns on and off MWIs	No	Yes

1. Under certain circumstances, Cisco Unity on the primary server continues to run after failover occurs. Under other circumstances, Cisco Unity is not running. For details, see the “[Intervals for Failover and Failback](#)” section on page 4-12.

When the secondary server is active, Cisco Unity functions normally, with the following exceptions:

- For a few seconds after a failure occurs and before the secondary server becomes active, subscribers hear a fast-busy tone when they dial the internal phone number to log on to Cisco Unity by phone.
- Subscribers cannot use the phone as a recording or playback device with the Media Master unless the Cisco Unity server name specified in the Media Master is updated accordingly. Depending on the application and the version of Cisco Unity, the server name may update automatically or subscribers may need to manually update the Media Master with the correct server name. (For more information, see the “[Notifying Subscribers to Update the Server Name in the Media Master](#)” section on page 2-1.)
- Reports can be generated only while the primary server is active.

You can configure the secondary server to initiate failback daily. When failback succeeds, the primary server becomes the active server again. Alternatively, you can configure failover so that the secondary server fails back only when you manually initiate failback by using the Failover Monitor.

For the behavior of failover during outages of network components, see [Appendix C, “Behavior of Cisco Unity Failover During Outages of Network Components.”](#)

Requirements for Cisco Unity Failover

- The primary and secondary servers must both be qualified for Cisco Unity. Refer to Part 1 of *Cisco Unity 4.0 System Requirements, and Supported Hardware and Software* at http://www.cisco.com/univercd/cc/td/doc/product/voice/c_unity/sysreq/40_sysrq.htm.
- Both Cisco Unity servers must have the same platform overlay.
- Both Cisco Unity servers must be member servers of the same domain (they cannot be domain controllers). Do not install Active Directory on either Cisco Unity server.
- The Cisco Unity server names must be 15 characters or fewer, and the server names must not be identical for the primary and secondary servers.
- Both Cisco Unity servers must be connected to the same message store.
- The only Exchange software installed on the Cisco Unity servers is the administration software. All other Exchange services (including subscriber mailboxes) are on a computer other than the Cisco Unity servers. Messages must not be stored on the Cisco Unity servers. (This separation allows the availability of messages when either the primary or secondary server is not functioning.)
- SQL Server 2000 Standard Edition must be installed on both Cisco Unity servers. (MSDE 2000 is not supported on either server with Cisco Unity failover.)
- One Cisco Unity server is designated the primary server, and the other Cisco Unity server is designated the secondary server.
- Both Cisco Unity servers must have the same enabled features and configurations.
- Both Cisco Unity servers must be connected to the network and have a reliable connection of 100 Mbps minimum. There is no option for installing failover without a network connection.
- Failover can be used with any supported Cisco Unity configuration except with one that has no network connection.

- Cisco Unity and SQL Server 2000 must be installed on both the primary and secondary servers with the same domain account.
- MSSQLSERVER and SQLSERVERAGENT services on both Cisco Unity servers must be configured to use the same domain account that is a member of the Local Administrators group on both servers. These services cannot be configured to run as Local System. SQLSERVERAGENT on the primary server must be able to log on to SQL Server on the secondary server by using Windows NT authentication.

Effects of Using or Not Using the Force Failover Setting

How Cisco Unity failover responds to calls that are not answered on the primary server depends on the Force Failover If Call Arrives on Inactive Secondary setting in the Failover Monitor. Causes for the primary server not answering calls are:

- One or more voice messaging ports on the primary server lock up and fail to respond to calls.
- An insufficient number of voice messaging ports are available to handle phone traffic.
- The primary server experiences high CPU utilization.
- *Cisco CallManager integration only:* The voice messaging ports on the primary server unregister with the Cisco CallManager server.

By checking or not checking the Force Failover check box, you can control whether an unanswered call on the primary server causes failover to occur.

Events When the Force Failover If Call Arrives on Inactive Secondary Check Box Is Checked

When a call is not answered on the primary server and the check box is checked in the Failover Monitor, the following events occur:

1. The call is presented to the secondary server.
2. The call causes an Event log entry on the secondary server, which can alert the system administrator to the potential problem on the primary server. (For details on setting up event notification, see the [“Setting Up Notification of When Failover Occurs”](#) section on page 1-5.)
3. The secondary server answers the call, causing failover to occur.

Customers may want the secondary server to become active when a voice messaging port on the primary server is locked up. However, other customers who install Cisco Unity servers with a large number of ports may be annoyed that failover occurs when just one port is locked.

Events When the Force Failover If Call Arrives on Inactive Secondary Check Box Is Not Checked

When a call is not answered on the primary server and the check box is not checked in the Failover Monitor, the following events occur:

1. The call is forwarded to the next available voice messaging port on the primary server. (Failover does not occur, and the primary server remains active.)
2. The call causes an Event log entry on the secondary server, which can alert the system administrator to the potential problem on the primary server. (For details on setting up event notification, see the [“Setting Up Notification of When Failover Occurs”](#) section on page 1-5.)

Failover occurs only when the primary server fails or when the system administrator manually initiates failover.

Customers who install Cisco Unity servers with a large number of voice messaging ports do not experience failover when just one port is locked. However, if all ports are locked or the CPU utilization is high, Cisco Unity stops answering calls until the system administrator manually initiates failover.

Effects of Failover and Failback on Calls in Progress

When failover or failback is initiated, calls in progress are maintained or dropped, depending on the status of the Cisco Unity service (AvCsMgr.exe) and of the network.


Table 4-3 Failover and Failback Effects on Calls in Progress

Cisco Unity Service or Network Status	Action
The Cisco Unity service is stopped or crashes	Calls in progress are dropped. When the Cisco Unity service (AvCsMgr.exe) on the primary server becomes inactive, the secondary server becomes active, but the dropped calls cannot be recovered.
The Cisco Unity service is active when failover or failback occurs	Calls in progress are maintained until the callers hang up.
Network connections are lost	Calls in progress from circuit-switched phone systems are maintained until the callers hang up. Calls in progress from IP phone systems (for example, Cisco CallManager) may be dropped, depending on the nature of the network problem. (For the behavior of failover during outages of network components, see Appendix C, “Behavior of Cisco Unity Failover During Outages of Network Components.”)

Status Monitoring and File Replication

When failover is configured and Cisco Unity is functioning normally, monitoring and replication take place as shown in [Figure 4-4](#).

Figure 4-4 Monitoring and Replication Actions



69854

For the behavior of failover during outages of network components, see [Appendix C, “Behavior of Cisco Unity Failover During Outages of Network Components.”](#)

Details of the monitoring and replication actions are described below.

Node Manager Service

- The Node Manager service (AvCsNodeMgr in the Services window) on the active server sends its status to the inactive server (it sends keep-alive events to the inactive server) at a frequency the installer configures. The default is 1 second. The TCP packet is 785 bytes in size and uses port 3653, which the installer can change as a hidden registry setting.
- The Node Manager service in the inactive server sends its status to the active server (it sends keep-alive events to the active server) at a frequency the installer configures. The default is 1 second. The TCP packet is 785 bytes in size and uses port 3653, which the installer can change as a hidden registry setting.
- The Node Manager service monitors the Cisco Unity service (AvCsMgr) on the server.

Cisco Unity Files

- The frequency of replication between the active and inactive servers depends on the value of the File Replication Interval field in the Failover Monitor and how often the Cisco Unity files change.
- The Node Manager service monitors and replicates the files in the directories Localize\DefaultConfiguration, Localize\Prompts, Snapshot, Stream Files, Support, and UnityMTA (which contains the Unity Messaging Repository, or UMR) on the Cisco Unity server.

- The Node Manager service creates a snapshot in the Snapshot directory of the files in the replicated directories during each replication cycle. In subsequent replication cycles, the Node Manager service compares the current snapshot with the previous snapshot to determine the files that should be replicated.
- If changed files have not been replicated to the secondary server when the AvCsMgr service is stopped or the Cisco Unity server crashes, the changes may be lost. Under such circumstances, the new information is not on the secondary server and may be lost or corrupted on the primary server.
- Replication of changed files is disabled when the Node Manager service is set to manual mode.

SQL Server Database (UnityDb)

- Replication occurs every minute if there are SQL data changes.
- The database on the primary server is configured as the publisher and distributor.
- The database publication is UnityDbPublication.
- The database on the secondary server is configured as the subscriber.
- Two-way replication between the two servers is enabled to let the primary server receive all changes made during periods it is online and inactive.
- If changed data has not been replicated to the secondary server when the AvCsMgr service is stopped or the Cisco Unity server crashes, the new data may be lost. Under such circumstances, the new data is not on the secondary server and may be lost or corrupted on the primary server.

Voice Messages

- Voice messages are not replicated because failover requires that the message store (Exchange server) be on a separate computer from the Cisco Unity server. When voice messages are kept on a separate server, they are available to subscribers no matter which Cisco Unity server is active.
- Voice messages saved in the Unity Messaging Repository (UMR) when the message store (Exchange server) is off line are replicated to the inactive server with other Cisco Unity files. (The UMR is located in the UnityMTA directory.)

Data That Is Not Replicated

Changes to the following Cisco Unity settings are not replicated between the primary and secondary servers. You must manually change values on both servers.

- Registry settings
- Recording settings
- Phone language settings
- GUI language settings
- Port settings
- Integration settings
- Media Master server name settings
- *In Cisco Unity 4.0(1) through 4.0(3):* AMIS restriction table selection. (In Cisco Unity 4.0(4) and later, the AMIS restriction table selection is replicated.)

Events When Failover Occurs

1. Status monitoring and file replication occur when the primary server is active:
 - a. The Node Manager service monitors the Cisco Unity service and other services on the server to determine whether failover should be initiated.
 - b. The Node Manager services on both the primary and secondary servers send status to each other on a regular basis.
 - c. Certain Cisco Unity files when changed are replicated from the primary server to the secondary server.
 - d. Data from the SQL Database (UnityDb) when changed is replicated from the primary server to the secondary server.
 - e. The primary server answers calls, sends message notification, and turns MWIs on and off. The secondary server is inactive and performs none of these functions.
 - f. Voice messages are stored in the message store (Exchange server), which is on a separate computer from the primary and secondary servers.
2. Failover occurs when any of the following events take place:
 - The Node Manager service on the secondary server does not receive a status update from the primary server within the period specified (for example, when the primary server crashes, the Node Manager service is stopped, or there are network connectivity problems).
 - The Cisco Unity service on the primary server is stopped. The Node Manager service on the primary server reports the stoppage to the Node Manager service on the secondary server.
 - A port on the secondary server receives a call when the Force Failover If Call Arrives on Inactive Secondary check box is checked in the Failover Monitor.
3. The Node Manager service on the secondary server assumes that the Cisco Unity service is stopped on the primary server and activates the Cisco Unity service on its server. (In the case of the secondary server receiving a call, the Node Manager service on the secondary server instructs the Node Manager service on the primary server to initiate failover.) The secondary server starts answering calls, sending message notification, and turning MWIs on and off.
4. The Node Manager service on the secondary server writes a warning to the Event log.

For the behavior of failover during outages of network components, see [Appendix C, “Behavior of Cisco Unity Failover During Outages of Network Components.”](#)

Events When Failback Occurs

1. Status monitoring and file replication occur when the secondary server is active and the primary server is running (but inactive):
 - a. The Node Manager service monitors the Cisco Unity service and other services on the server to determine whether failback should be initiated.
 - b. The Node Manager services on both the secondary and primary servers send status to each other on a regular basis.
 - c. Certain Cisco Unity files when changed are replicated from the secondary server to the primary server.
 - d. Data from the SQL Database (UnityDb) when changed is replicated from the secondary server to the primary server.

- e. The secondary server answers calls, sends message notification, and turns MWIs on and off. The primary server is inactive and performs none of these functions.
 - f. Voice messages are stored in the message store (Exchange server), which is on a separate computer from the secondary and primary servers.
2. Failback occurs when either of the following events takes place:
 - When automatic failback is enabled, the time specified for failback arrives, and the Node Manager service on the secondary server notifies the Node Manger service on the primary server that the Cisco Unity service on the secondary server is inactive.
 - Manual failback is initiated by using the Failover Monitor.
 3. The Node Manager service on the primary server assumes that the Cisco Unity service is stopped on the secondary server and activates the Cisco Unity service on its server. (In the case of the primary server receiving a call, the Node Manager service on the primary server instructs the Node Manger service on the secondary server to initiate failback.) The primary server then starts answering calls, sending message notification, and turning MWIs on and off.
 4. The Node Manager service on the primary server writes a warning to the Event log.

For the behavior of failover during outages of network components, see [Appendix C, “Behavior of Cisco Unity Failover During Outages of Network Components.”](#)

Automatic and Manual Failback

Failback can occur automatically at a time you specify, or you can manually initiate failback. The Failback Type setting is in the Failover Monitor and has the following options:

Scheduled	Use this option if you want the secondary server to fail back automatically. In the Failover Monitor, enter values for the Scheduled Failback Start and Scheduled Failback End fields to set the time for automatic failback. This option is useful when you want failback to occur, for example, at night. You can initiate failback manually before the scheduled time by using the Failover Monitor.
Manual	Use this option if you do not want the secondary server to automatically fail back to the primary server when the primary server becomes available. The secondary server fails back to the primary server only when you manually fail back by using the Failover Monitor.

For details on customizing the failback type in the Failover Monitor, see the [“Customizing Failover and Failback Settings for the Cisco Unity System \(Optional\)”](#) section on page 1-12.

Causes of Failover and Failback

Failover Causes

Failover from the primary server to the secondary server can occur for the following reasons:

- Manual failover is initiated from the Failover Monitor.

- The network connection to the primary server is lost. However, both primary and secondary servers remain active.
- The network connection between the primary and secondary servers experiences latency, which causes the Node Manager service on the secondary to miss 30 keep-alive events from the primary server (the default setting in the Failover Monitor).
- If the Force Failover If Call Arrives on Inactive Secondary check box is checked in the Failover Monitor and a call is answered by the secondary server.
- The primary server loses power while the secondary server is still powered.
- The primary server is turned off while the secondary server remains on.
- On the primary server, the Cisco Unity service is stopped or crashes.
- On the primary server, the Node Manager service stops.
- On the primary server, the Node Manager service fails to send keep-alive events to the Node Manager service on the secondary server as configured in the Failover Configuration dialog box of the Failover Monitor.
- There is excessive CPU usage on either the primary or secondary server. When this problem occurs on the primary server before failover and the secondary server after failover, the primary and secondary servers will failover and failback repeatedly until the CPU usage problem is resolved.
- The Event Monitoring Service (EMS) initiates manual failover when a user-specified event occurs.

Failback Causes

Failback from the secondary server to the primary server can occur for the following reasons:

- The problem on the primary server is resolved, and manual failback is initiated from the Failover Monitor.
- The problem on the primary server is resolved, and the scheduled time for failback has arrived as configured in the Failover Configuration dialog box of the Failover Monitor.
- The secondary server loses power while the primary server is still powered but inactive.
- The secondary server is turned off while the primary server remains on but is inactive.
- On the secondary server, the Cisco Unity service is stopped or crashes, while the primary server is on but inactive.
- On the secondary server, the Node Manager service stops, while the primary server is on but inactive.
- On the secondary server, the Node Manager service fails to send keep-alive events to the Node Manager service on the primary server as configured in the Failover Configuration dialog box of the Failover Monitor, while the primary server is on but inactive.

Intervals for Failover and Failback

Failover Interval

When failover occurs, the time it takes for the secondary server to begin answering calls typically depends on what causes failover as described in [Table 4-4](#).

Table 4-4 *Interval for Failover*

Failover Cause	Time
Manual failover	<ul style="list-style-type: none"> The secondary server begins answering calls immediately. Calls in progress on the primary server are not dropped but are maintained until the callers hang up.
Network outage	<ul style="list-style-type: none"> The secondary server begins answering calls after waiting for the number of keep-alive events set in the Missed Events Before Failover field in the Failover Monitor. The time required depends on the value of the field and on the value of the Interval (ms) field. The default settings for the fields result in a 30-second delay for the secondary server to become active. A network outage may result in both primary and secondary servers being active at the same time. When the network outage is resolved, the primary server remains active and the secondary server becomes inactive. Calls in progress from a circuit-switched phone system are not dropped but are maintained until the callers hang up. Calls in progress from an IP phone system (for example, Cisco CallManager) may be dropped, depending on the nature of the network outage.
Failure of the AvCsMgr service on the primary server	<ul style="list-style-type: none"> The secondary server begins answering calls immediately. Calls in progress on the primary server are dropped.
Failure of the operating system on the primary server	<ul style="list-style-type: none"> The secondary server begins answering calls after waiting for the number of keep-alive events set in the Missed Events Before Failover field in the Failover Monitor. The time required depends on the value of the field and on the value of the Interval (ms) field. The default settings for the fields result in a 30-second delay for the secondary server to become active. Calls in progress on the primary server are dropped.
Inactive secondary server answers a call	<ul style="list-style-type: none"> The secondary server begins answering calls immediately. Calls in progress on the primary server are not dropped but are maintained until the callers hang up.

Failback Interval

When failback occurs, the time it takes for the primary server to begin answering calls typically depends on what causes failback as described in [Table 4-5](#). File replication occurs in the background and does not affect the failback interval.

Table 4-5 *Interval for Failback*

Failback Cause	Time
Automatic failback, and the primary server is on but inactive	<ul style="list-style-type: none"> • Failback is initiated only when the time scheduled in the Failover Monitor arrives. • When Cisco Unity is integrated with a circuit-switched phone system, the voice messaging ports on the primary server that are connected to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with a SIP phone system, the voice messaging ports on the primary server that are registered to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with Cisco CallManager, the primary server begins answering calls as soon as the first voice messaging port on the primary server connected to Cisco CallManager is registered after failback is initiated. All voice messaging ports connected to Cisco CallManager are registered in about 15 seconds. • Calls in progress on the secondary server are not dropped but are maintained until the callers hang up.
Manual failback, and the primary server is on but inactive	<ul style="list-style-type: none"> • Failback is initiated immediately. • When Cisco Unity is integrated with a circuit-switched phone system, the voice messaging ports on the primary server that are connected to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with a SIP phone system, the voice messaging ports on the primary server that are registered to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with Cisco CallManager, the primary server begins answering calls as soon as the first voice messaging port on the primary server connected to Cisco CallManager is registered after failback is initiated. All voice messaging ports connected to Cisco CallManager are registered in about 15 seconds. • Calls in progress on the secondary server are not dropped but are maintained until the callers hang up.

Table 4-5 Interval for Failback (continued)

Failback Cause	Time
Failure of the AvCsMgr service on the secondary server, and the primary server is on but inactive	<ul style="list-style-type: none"> • Failback is initiated immediately. • When Cisco Unity is integrated with a circuit-switched phone system, the voice messaging ports on the primary server that are connected to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with a SIP phone system, the voice messaging ports on the primary server that are registered to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with Cisco CallManager, the primary server begins answering calls as soon as the first voice messaging port on the primary server connected to Cisco CallManager is registered after failback is initiated. All voice messaging ports connected to Cisco CallManager are registered in about 15 seconds. • Calls in progress on the secondary server are dropped.
Failure of the operating system on the secondary server, and the primary server is on but inactive	<ul style="list-style-type: none"> • The primary server initiates failback after waiting for the number of keep-alive events set in the Missed Events Before Failover field in the Failover Monitor. The time required depends on the value of the field and on the value of the Interval (ms) field. The default settings for the fields result in a 30-second delay for the primary server to become active. • When Cisco Unity is integrated with a circuit-switched phone system, the voice messaging ports on the primary server that are connected to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with a SIP phone system, the voice messaging ports on the primary server that are registered to that phone system begin answering calls immediately after failback is initiated. • When Cisco Unity is integrated with Cisco CallManager, the primary server begins answering calls as soon as the first voice messaging port on the primary server connected to Cisco CallManager is registered after failback is initiated. All voice messaging ports connected to Cisco CallManager are registered in about 15 seconds. • Calls in progress on the secondary server are dropped.

Causes of Both Servers Becoming Active at the Same Time

Under some circumstances, both the primary server and secondary server can become active at the same time, as indicated in the Failover Monitor. The condition can occur for the following reasons:

- The network connection to the primary server is lost, so both servers become active.

With Cisco CallManager and SIP phone systems, Cisco Unity will receive no calls. However, if the primary server loses its connection to the Cisco CallManager server or the SIP proxy server while the secondary server retains its connection, calls can reach the secondary server and initiate failover.

With circuit-switched phone systems, calls will be answered by both the primary server and the secondary server simultaneously.

- The Node Manager service on the primary server is stopped. The secondary server becomes active, but the components on the primary server remain active and continue to answer calls.

Effects of Shutting Down and Restarting the Primary and Secondary Servers

When you shut down or restart either the primary or secondary server, the servers behave in the following ways:

- When both servers are running and the active server is shut down, the inactive server becomes active.
- When neither server is running, the first server started becomes the active server.
- When the secondary server is active and configured for automatic failback, and the primary server is also running, the secondary server attempts failback on the failback schedule.

Licensing Restrictions on Using a Secondary Server Without a Primary Server

To prevent the secondary server from being used as the primary server in another location, the secondary server stops answering calls:

- Four hours after you restart the secondary Cisco Unity server, if you have never run the Configure Cisco Unity Failover wizard.
- Four hours after you restart the secondary Cisco Unity server, if you have run the Configure Cisco Unity Failover wizard but the secondary server has never been able to contact the primary server.
- 60 days after the last time that the secondary server was able to contact the primary server, if the secondary has contacted the primary server at least once. In this case, the only way to make the secondary server start answering calls again is to reconnect the primary server to the network.

There are no licensing restrictions on a primary server, so a primary server will operate without a secondary server. However, without a secondary server, the failover feature cannot work, and an error appears in the Event log whenever you restart the Cisco Unity server. To stop the errors from appearing, disable the Node Manager service (AvCsNodeMgr) in the Services window.

