Cisco TelePresence Conductor

Administrator Guide

XC3.0

January 2015

Contents

Introduction to the Cisco TelePresence Conductor

This section introduces the Cisco TelePresence Conductor and its features. It highlights the new features that were added in this software version.

About the Cisco TelePresence Conductor

The Cisco TelePresence Conductor simplifies multiparty video communications. It lies within a video communications network, working in conjunction with at least one conference bridge and at least one call control device. The call control devices can be Cisco TelePresence Video Communication Servers (Cisco VCSs) or Cisco Unified Communications Managers (Unified CMs) or both. TelePresence Conductor allows the video network to be configured such that spontaneous or rendezvous* conferences may be easily provisioned, initiated, accessed and managed.

* personal conferences with a unique conference ID that are often referred to as “MeetMe” conferences

Cisco TelePresence Conductor features

The Cisco TelePresence Conductor is the focal point in a multiparty network and acts in a similar manner to the conductor of an orchestra. It knows all of the individual conferencing components that have been configured to be used with it and their capabilities intimately. It will control all of these individual elements to achieve the best possible performance. This will ensure intelligent conference placement and improved resource utilization and provides powerful, comprehensive administrator control.

It is tightly integrated with the industry-leading Cisco TelePresence MCU, the Cisco TelePresence Servers, the Cisco TelePresence Video Communication Servers (Cisco VCSs), the Cisco Unified Communications Managers (Unified CMs), Cisco TelePresence Management Suite (Cisco TMS) and Cisco TelePresence Management Suite Provisioning Extension (Cisco TMSPE). It works with all standards-compliant endpoints.

The TelePresence Conductor can be deployed in a triple-redundant cluster (with nodes that may be geographically distributed), providing true reliability – conferencing is always available. The resilient architecture ensures service availability even if individual conference bridges or TelePresence Conductors are taken out of service.

It provides a single interface for service provisioning, no matter how many conference bridges there are. As scale increases, more conference bridges may simply be added without increasing provisioning overhead.

The TelePresence Conductor scales from IP phone through to immersive meeting room and from small businesses to the largest enterprises.

Conference personalization is supported, allowing a consistent user experience that satisfies the users' personal preferences (layouts, PINs, encryption etc.), irrespective of the conference bridge on which a conference is hosted.

TelePresence conferencing is elevated to a new level by ensuring a reliable and faultless conference experience with all of the conferencing components working together in harmony.

What's new in this version?

Authentication required when changing an administrator account password

When you add a new administrator account or change the password for an existing administrator account, you are now required to authenticate yourself by entering your current administrator password. There are now two pages for editing administrator account details: one for changing the password and one for editing the remaining account details.

Ability to mark pools within a Service Preference to be used for scheduling

On the **Service Preference** page of the TelePresence Conductor user interface you can mark pools to be used for scheduling. Only marked pools will be included in Capacity Management API requests that clients such as Cisco TMS make. If you configure the marked pool to contain only a single conference bridge and you do not include the same pool in more than one Service Preference, Cisco TMS can use the pool for dedicated-bridge scheduling.

Note: After an upgrade to XC3.0, all existing pools in all Service Preferences are marked to be used for scheduling.

Support for WebEx calls to be included in scheduled conferences

It is now possible for scheduled conferences on TelePresence Conductor to include WebEx calls as well as TelePresence calls. Cisco TMS scheduling with TelePresence Conductor now supports Cisco Collaboration Meeting Rooms Hybrid (formerly known as WebEx enabled TelePresence).

Support for participant role determined by PIN

Conferences provisioned using the TelePresence Conductor Provisioning API, for example via Cisco TMSPE, now allow participant role to be determined by PIN. Hosts and guests dial the same alias and then experience different privileges based on the PIN they have entered. Guest PINs are optional. This feature is supported on TelePresence MCUs and on TelePresence Servers version 4.1 or later.

Ability to specify whether guests must wait for a host to join a conference first

It is now possible to specify in the TelePresence Conductor Provisioning API whether guests must wait for a host to join the conference before they are able to join. This setting is only applicable to

- Conferences provisioned through the Provisioning API (for example by Cisco TMSPE)
- Conferences hosted on TelePresence Servers

Support for multistream calls

Multistream calls are now supported when you are using endpoints and TelePresence Servers that also support this feature. TelePresence Conductor forwards the relevant SDP (session description protocol) information from the endpoints to the TelePresence Servers.

Support for up to three SIP trunk destinations

You can specify up to three SIP trunk destinations, consisting of an IP address and a SIP port, for each rendezvous Location defined on the TelePresence Conductor user interface. The TelePresence Conductor considers all SIP trunk destinations for a Location as equivalent and may use any one of the destinations for out-dial calls, as long as the destination is reachable. The TelePresence Conductor maintains only one of the destinations, it does not load balance the dial-out calls across the configured destinations. If the current destination becomes unreachable, it automatically chooses a new SIP trunk destination.

SIP trunk destinations cannot be specified for ad hoc Locations, unless the **Conference type** of the Location is set to *Both*.

TelePresence Conductor regularly polls its SIP trunk destinations and reports reachability changes

TelePresence Conductor uses a SIP OPTIONS ping to regularly poll all SIP trunk destinations configured for the call control devices that are connected to it. This includes all Unified CMs and any Cisco VCSs connected using the back-to-back user agent (B2BUA).

If there is a response from the SIP trunk destination, it is considered to be reachable. If there is a change in the reachability, either from reachable to unreachable or vice versa, the state is reported in an event log

message. If any SIP trunk destinations are unreachable, an alarm is raised. The alarm is lowered if all SIP trunk destinations are reachable again.

Syslog publish filter

You can now filter the logs that TelePresence Conductor sends to each remote syslog host by severity level.

For example, your syslog host is typically receiving syslog messages from multiple systems, so you may want to limit TelePresence Conductor to sending only "Error" messages (and anything more severe) to this host. If you want to leave the host untouched while troubleshooting a TelePresence Conductor problem, you could configure a second, temporary, host to receive "Debug" level (most verbose = messages of all severities). Then you could safely remove the configuration after resolving the issue, without risking your primary syslog host.

User interface change

The menu path to the **IP** and **Ethernet** pages has changed to **System > Network interfaces > IP** and **System > Network interfaces > Ethernet** respectively.

API changes

- The XML RPC API call **`factory.conferencecreate`** has a new optional parameter called **`factoryLayout`**. It takes the following values:
  - equal
  - active
  - prominent
  - single

  It allows you to override the layout specified within the **`layout`** parameter of the **`ConfBundle`** object in the Provisioning API.
- The XML RPC API has a new call - **`factory.conferencemodify`**. It allows you to modify the values of the parameters that were set when the conference was created using the call **`factory.conferencecreate`**.
- The XML RPC API calls **`conference.create`** and **`conference.modify`** have been deprecated. Use **`factory.conferencecreate`** and **`factory.conferencemodify`** instead.
- The API call **`factory.conferencecreate`** now returns the **`factoryConferenceId`** and **`conferenceName`** values where possible. This includes successful calls as well as some failed calls.
- The XML RPC API call **`participant.message`** now allows the API client to specify the position and duration of the message displayed on the participant's screen.
- It is now possible for API clients to lock and unlock a conference on TelePresence Conductor. If a conference is locked, it keeps running with its existing participants. No new participants can dial into a locked conference, but API clients, such as Cisco TMSPE, can add more participants to a conference via the API call **`participant.add`**. The XML RPC API calls **`factory.conferencemodify`** and **`conference.modify`** have a new Boolean parameter called **`locked`**.
- The XML RPC API call **`participant.enumerate`** has a new return value - **`factoryCallState`**. It is returned as part of the participant struct and takes the values disconnected, ringing, connected, awaitingTrigger, callLegFailed and retrying. The new **`factoryCallState`** allows TelePresence Conductor to provide to its API clients more detailed state information about the participants in a conference.
- The Provisioning API object **`ConfBundle`** has a new Boolean attribute called **`guests_wait_for_host`**. It allows you to specify whether guests must wait for a host before they can join a conference. It is only applicable to TelePresence Server hosted conferences. The attribute is ignored for TelePresence MCU hosted conferences. The default value is *False*.

- The strings **`factoryConferenceId`** and **`factory_conference_id`** that are returned after issuing the API call **`factory.conferencecreate`** have been modified. They are now opaque strings that no longer resemble UUIDs.

Other changes

The certificate signing request storage location changed in XC3.x.

When you generate a CSR in XC2.x, the application puts **csr.pem** and **privkey_csr.pem** into **/tandberg/persistent/certs**.

When you generate a CSR in XC3.x, the application puts **csr.pem** and **privkey.pem** into **/tandberg/persistent/certs/generated_csr**.

If you want to upgrade from XC2.x and have an unsubmitted CSR, then we recommend discarding the CSR before upgrade, and then regenerating the CSR after upgrade.

Using the web interface

This section provides information on how to use the TelePresence Conductor web interface. It also describes what the web interface layout looks like and which browsers and characters are supported.

Logging in to the web interface

1. Open a browser window and in the address bar type either:
   - the IP address of the system (this is `192.168.0.100` by default and should be changed during the commissioning process)
   - the FQDN (fully-qualified domain name) of the system.

   The **Administrator login** page appears.
2. Enter a valid administrator **Username** and **Password** (see the Configuring administrator accounts [p.128] section for details on setting up administrator accounts) and click **Login**.

The **Overview** page appears.

Note: The default username for the administrator user is **admin** and the default password is **TANDBERG**. The TelePresence Conductor's conference functionality is disabled until this password has been changed. It is important to select a secure password.

When logging in to the TelePresence Conductor web interface, you may receive a warning message regarding the TelePresence Conductor's security certificate. To avoid this, ensure that you replace the factory default certificate with your own valid certificate. See Managing the TelePresence Conductor's server certificate [p.182] for more information.

Web page features and layout

This section describes the features that can be found on some or all of the web interface pages.

Figure 1: Example web page

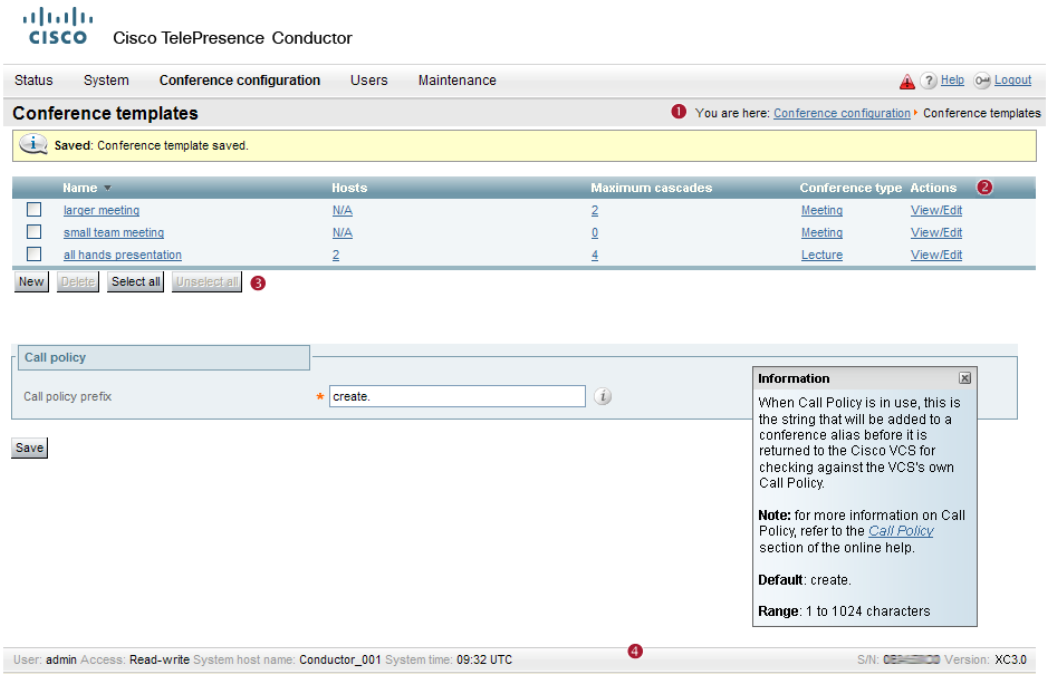

The elements included in the example web pages shown above are described in the table below.

Page element		Description
Page name and location	1	Every page shows the page name and the menu path to that page. Each part of the menu path is a link; clicking on any of the higher level menu items takes you to that page.
System alarm		This icon appears on the top right corner of every page when there is a system alarm in place. Click on this icon to go to the Alarms page which gives information about the issue and its suggested resolution. There should never be any active alarms on a fully functional, correctly configured system.
Help		This icon appears on the top right corner of every page. Clicking on this icon opens a new browser window with help specific to the page you are viewing. It gives an overview of the purpose of the page, and introduces any concepts related to the page.
Log out		This icon appears on the top right corner of every page. Clicking on this icon ends your administrator session.
Field level information		An information box appears on the configuration pages whenever you either click on the information icon or click inside a field. This box gives you information about the particular field, including where applicable the valid ranges and default value. To close the information box, click on the X at its top right corner.
Information bar		The TelePresence Conductor provides you with feedback in certain situations, for example when settings have been saved or when you need to take further action. This feedback is given in a yellow or red information bar at the top of the page.
Sorting columns	2	Click on column headings to sort the information in ascending or descending order.

Page element		Description
Select All and Unselect All	3	Use these buttons to select and unselect all items in the list.
Mandatory field	*	Indicates an input field that must be completed.
System Information	4	The name of the user currently logged in and their access privileges, the system name (or LAN 1 IPv4 address if no system name is configured), local system time, hardware serial number and TelePresence Conductor software version are shown at the bottom of the page.

Supported browsers and characters

Supported browsers

- Internet Explorer 8 or 9
- Firefox 3 or later
- Chrome

Later versions of these browsers may also work, but are not officially supported. It may work with Opera and Safari, but you could encounter unexpected behavior.

JavaScript and cookies must be enabled to use the TelePresence Conductor web interface.

Supported characters

The TelePresence Conductor supports the following characters when entering text in the web interface:

- the letters A-Z and a-z
- decimal digits (0-9)
- underscore (_)
- minus sign / hyphen (-)
- equals sign (=)
- plus sign (+)
- at sign (@)
- comma (,)
- period/full stop (.)
- exclamation mark (!)
- spaces

The following characters are allowed but we recommend that you do not use them:

- tabs
- angle brackets (< and >)

- ampersand (&)
- pipe symbol (|)

Case sensitivity

Most text items entered through the web interface are case-insensitive. The exceptions are passwords.

TelePresence Conductor capacity versions

From version XC2.2.1 onward there are three capacity versions available for the TelePresence Conductor:

- Full capacity - available as an appliance or a virtual machine
- TelePresence Conductor Select - only available as a virtual machine with an option key installed that supports up to 50 concurrent call sessions
- TelePresence Conductor Essentials - only available as a virtual machine running without a release key

The following limitations apply to the three different capacity versions of the TelePresence Conductor:

	TelePresence Conductor Essentials (free)	**TelePresence Conductor Select**	**Full capacity TelePresence Conductor**
Suitable deployment	Small Recommended for: ■ testing and reviewing new releases ■ proof of concept demonstrations	Small to medium-sized	Medium-sized to large
Total number of conference bridges supported	1 (standalone)	30	30
Maximum number of concurrent call sessions supported	The number of calls supported by the conference bridge	50	2400
Clustering of TelePresence Conductors supported for resilience	No	Yes (limited to 2 TelePresence Conductor Select)	Yes (up to 3 full capacity TelePresence Conductors)
Access to TAC support	No (for deployment in production environments we recommend upgrading to one of the other two capacity versions)	Yes	Yes
Available as virtual machine or appliance	Virtual machine only	Virtual machine only	Virtual machine and appliance
Release and option keys required to install	No release or option key required	Option key to support 50 concurrent call sessions required	Full capacity TelePresence Conductor release key required

Using the TelePresence Conductor without a release key

It is possible to use the TelePresence Conductor without a release key, as TelePresence Conductor Essentials. This results in limited system capacity. The limitations are:

- The TelePresence Conductor cannot be part of a cluster.
- Only one standalone conference bridge across all conference bridge pools can be enabled at a time. **Note:** conference bridges that are clustered are not supported. If a clustered conference bridge is enabled on the TelePresence Conductor it will be displayed as *Unusable* on the **Conference bridge status** page and the TelePresence Conductor will not use it.

When the TelePresence Conductor is running as TelePresence Conductor Essentials a banner is displayed along the top of the web interface, stating **Invalid or no release key installed:** The TelePresence Conductor is running with an invalid release key or without a release key; system capacity is limited to one standalone conference bridge with no clustering. Contact your Cisco account representative to obtain release and option keys to upgrade the supported capacity."

To buy a release key and/or option keys and take advantage of a broader feature set, contact your Cisco support representative with the serial number of the TelePresence Conductor (displayed under **Maintenance > Option keys**).

Before you start

This section provides information on the other devices that need to be configured to work with the TelePresence Conductor, how to design a dial plan and a configuration overview.

Configuring conference bridges for use with the TelePresence Conductor

A conference bridge is used for hosting the participants of a multipoint conference.

This version of the TelePresence Conductor supports the conference bridge types Cisco TelePresence MCU (version 4.2 or later) and Cisco TelePresence Server (version 3.0 or later). In order to support all features we strongly recommend that TelePresence MCUs and TelePresence Servers are running the latest software versions.

All conference bridges managed by a TelePresence Conductor must be added to a conference bridge pool that contains conference bridges of one type and with the same configuration. The TelePresence Conductor treats a pool of conference bridges as a single conference bridge resource, increasing the available capacity and providing redundancy. When the TelePresence Conductor receives a request for conference resources, it determines the conference bridge(s) that the conference should be hosted on.

All TelePresence Servers managed by the TelePresence Conductor must be configured to run in 'Remotely managed' mode.

All conference bridges in a single conference bridge pool must be configured identically before they are added to the TelePresence Conductor. The required configuration depends on whether the TelePresence Conductor is deployed in a setup using a Cisco VCS or a Unified CM.

For information on configuring a conference bridge in a Cisco VCS deployment, see *Cisco TelePresence Conductor with Cisco VCS (Policy Service) Deployment Guide* or *Cisco TelePresence Conductor with Cisco VCS (B2BUA) Deployment Guide*.

For information on configuring a conference bridge in a Unified CM deployment see *Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide*.

Configuring call control device(s) for use with the TelePresence Conductor

A call control device is used for managing the communication between SIP and/or H.323 devices.

This version of the TelePresence Conductor supports the call control devices Cisco Unified Communications Manager (Unified CM) version 8.6.2 or later, and Cisco TelePresence Video Communication Server (VCS) version X7.0 or later. In order to support all features we strongly recommend that Unified CMs and Cisco VCSs are running the latest software versions.

The Cisco VCS is supported in the following two types of deployments:

- Using the Cisco VCS's external policy service interface
  This method may be discontinued in future versions of the TelePresence Conductor software.
  For more information see *Cisco TelePresence Conductor with Cisco VCS (Policy Service) Deployment Guide*.
- Using the TelePresence Conductor's back-to-back user agent (B2BUA)
  This method requires a SIP trunk between the Cisco VCS and the TelePresence Conductor. It is the preferred method to use.
  For more information see *Cisco TelePresence Conductor with Cisco VCS (B2BUA) Deployment Guide*.

For information on configuring a Unified CM for use with the TelePresence Conductor see *Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide*.

Configuring endpoints for use with the TelePresence Conductor

No special endpoint configuration is required to enable endpoint users to dial conference aliases. As long as you have an appropriate dial plan in place and the endpoint can successfully register with Cisco TelePresence Video Communication Server (Cisco VCS) or Cisco Unified Communications Manager (Unified CM), it can make use of the TelePresence Conductor.

However, you should bear in mind that dialing behavior differs between SIP and H.323 endpoints. If you have a deployment that includes both types, ensure your conference aliases and dial plan are set up to support this.

Designing a dial plan

About dial plans

A dial plan defines dialed aliases and call routes within your video network. A well-designed dial plan is a key component of a successful audio/video network and should allow users to place calls simply and intuitively while retaining the ability to scale the network as more users and services are added.

General considerations

Before you add the Cisco TelePresence Conductor to your network, review your dial plan on your Cisco VCS or Unified CM to ensure it meets these requirements:

Area	Description
Conference aliases	These are the aliases that users will dial to create or join a conference. The conference aliases, which may be prefixes or exact matches, must be routed to the TelePresence Conductor. Ensure that your dial plan routes these prefixes appropriately. For more information, see Creating and editing conference aliases [p.104].
Conference names	These are the names that each conference will be known by on the host conference bridge. For more information, see Conference name [p.104].
Recording device	You can use the auto-dialed participants feature of the TelePresence Conductor to automatically add a recording device to a conference. To take advantage of this feature, your dial plan will need to forward particular aliases to your recording device. For more information, see Creating and editing auto-dialed participants [p.106].

Considerations in a Cisco VCS-only deployment

Area	Description
Call Policy prefix	(Only applicable if using the Cisco VCS's external policy server interface) This is used to allow or prevent specified users from creating conferences. The TelePresence Conductor adds this prefix to a conference alias and sends the request back to the Cisco VCS for checking against its own Call Policy. For more information, see Using Call Policy [p.123].
Conference bridge dial plan prefixes	(Only applicable if using the Cisco VCS's external policy server interface) These are used by the Cisco VCS to route calls to conference bridges in the TelePresence Conductor's pool. Each conference bridge must have a unique prefix. Each prefix must be configured on the Cisco VCS. For more information, see Dial plan prefix [p.65].

Area	Description
Deployments with both H.323 and SIP endpoints	SIP endpoints can only make calls in the form of URIs - for example **name@domain**. If the caller does not specify a domain when placing the call, the SIP endpoint automatically appends its own domain to the number that is dialed. So if you dial **meet.alice** from a SIP endpoint, the search will be placed for **meet.alice@domain**. H.323 endpoints do not append a domain, so if you dial **meet.alice** from an H.323 endpoint the call will be placed to **meet.alice**. If you have a deployment that includes both SIP and H.323 endpoints, you must ensure that your conference aliases and dial plan are set up so that users can dial the conference aliases from either type of endpoint. Some ways you can achieve this are: ■ Create all conference aliases in the form of a URI (for example **meet.alice@example.com**). All users will then have to dial the full URI to create or join a conference. ■ Using regular expressions, create conference aliases that will match an incoming alias with or without a domain appended. See Matching an alias with or without a domain appended [p.209] for an example. ■ Create all conference aliases in the form of a URI (e.g. **meet.alice@example.com**), and set up a transform on the Cisco VCS to append the domain to any alias that does not include one. All users will then have to dial just the name portion of the alias (e.g. **meet.alice**) to create or join a conference.
Avoiding a dial plan conflict	(Only applicable if using the Cisco VCS's external policy server interface) The Cisco VCS is responsible for routing calls to the appropriate destination (for example, TelePresence Conductor, TelePresence MCU, another Cisco VCS, or endpoint). It does this by employing search rules, which map search requests to zones and policy servers based on factors such as the alias being dialed and the source of the request. When creating your search rules on the Cisco VCS, you must ensure that they are specific enough so that: ■ all conference aliases will route to the TelePresence Conductor only ■ all calls beginning with any conference bridge dial plan prefix will route to the specified conference bridge only ■ all calls beginning with the TelePresence Conductor's call policy prefix route to the TelePresence Conductor only, and no conference aliases begin with the same prefix. Also make sure that no endpoints can register with a conference alias, conference bridge dial plan prefix or TelePresence Conductor call policy prefix; you can do this using registration allow or deny lists.

For full instructions on creating transforms and configuring search rules on the Cisco VCS, see *Cisco TelePresence Video Communication Server Administrator Guide*.

Overview of conference types

Conferences managed by TelePresence Conductor can have one of the following types:

- **Ad hoc conference**
  A non-scheduled, spontaneous meeting, where the user of an endpoint registered to Unified CM brings two or more endpoints together in a conference.
- **Multiway conference**
  A non-scheduled, spontaneous meeting, where the user of an endpoint registered to Cisco VCS brings two or more endpoints together in a conference.
- **Rendezvous conference / CMR**
  A non-scheduled conference for which the host knows the conference alias beforehand, and must share the alias with all participants. Participants can join at any time, and there is no scheduled end time. Rendezvous conferences can be configured in two different ways:
  - Via the TelePresence Conductor's web interface - by configuring a conference template and associated conference aliases that use regular expression matching
  - Via the TelePresence Conductor's Provisioning API - by using a management tool such as Cisco TMS to provision a CMR with associated direct-match aliases
- **Scheduled conference**
  A conference that has a fixed start and end time. Conferences cannot be scheduled directly through the TelePresence Conductor's web interface. Conferences can be configured and scheduled through a management tool such as Cisco TMS.

Provisioning conferences

A conference can be provisioned on TelePresence Conductor using a management tool such as Cisco TMS. The management tool must use the TelePresence Conductor's Provisioning API, defined in *Cisco TelePresence Conductor API Guide* for version XC2.3 or later.

The management tool provisioning a conference sets up a ConfBundle object, which on Cisco TMS is referred to as a Collaboration Meeting Room (CMR) with one or more associated aliases and optionally one or more associated auto-dialed participants.

Note: Conference aliases and auto-dialed participants configured via the web interface are separate from aliases and auto-dialed participants associated with CMRs, which are used to provision conferences via the TelePresence Conductor's Provisioning API.

On the **Collaboration meeting room** page (**Status > Provisioning > Collaboration meeting room**) you can search for one or more CMRs and view the details for specific CMRs. See Searching Collaboration Meeting Rooms [p.148]

The TelePresence Conductor's **Conferences status** page (**Status > Conferences**) displays all conferences currently managed by this TelePresence Conductor, whether they have been configured via the web interface or provisioned via the API.

The information for CMRs is separate from the information for conferences configured via the TelePresence Conductor's web interface. You cannot edit the information for CMRs via the TelePresence Conductor's web interface. You also cannot configure a conference on the TelePresence Conductor's web interface that uses objects created via the TelePresence Conductor's API. For example, a conference alias provisioned via the API cannot be used in a conference configured via the web interface.

An exception to this rule are Service Preferences, which are created via the TelePresence Conductor's web interface and can be used by either CMRs or conference templates.

All conference aliases provisioned via the Provisioning API apply an exact match to work out the conference name. All conference aliases configured via the web interface apply a regular expression (RegEx) to work out the conference name.

Exact match alias strings are case insensitive. They are stored by TelePresence Conductor as lower case strings and matched to dial strings entered into the endpoints using any case. Regex alias strings are case sensitive and match to dial strings entered into the endpoint using the same case as entered into the web interface.

Scheduling conferences

A scheduled conference can be created on TelePresence Conductor using a management tool such as Cisco TMS. The management tool must use the API for TelePresence Conductor, defined in *Cisco TelePresence Conductor API Guide*.

The TelePresence Conductor API allows API clients to:

- Request the conference bridge capacity that has been configured for a specific Service Preference (It includes conference bridges in pools that have been marked to be used for scheduling, without taking into account whether the bridges are currently used in a conference)
- Request the resources that would be required if a specific alias dialed into a conference using the specific Service Preference (It includes one-off per-conference costs and per-participant costs)
- Schedule the use of the resources returned in the Capacity API request (independently from the TelePresence Conductor)
- Create (and delete) conferences on the TelePresence Conductor at the scheduled time

It is not possible to schedule conferences directly on TelePresence Conductor.

TelePresence Conductor does not differentiate between conference bridges used for scheduled or non-scheduled conferences. It does, however, allow you to mark conference bridge pools to use for scheduling. This can be done on the **Service Preference** page on the TelePresence Conductor user interface. Only the pools marked for scheduling will be returned to the API client requesting capacity information. For more information see Marking pools to be used for scheduling [p.77].

Note: When configuring conference bridge pools dedicated for scheduling, we recommend the following:

- Give the conference bridge pool a name indicating that it should only be used for scheduled conferences.
- Check that the pool is only used in a single Service Preference.
- Check that the Service Preference is not used in a CMR or ad hoc conference.

Various scenarios can be configured to support scheduling on TelePresence Conductor. The required configuration, advantages and disadvantages for some example scenarios are detailed below.

Note: It is not possible to mark a conference bridge as a "backup" within the TelePresence Conductor. The TelePresence Conductor will automatically use the next available conference bridge within the Service Preference when a higher priority bridge fills up.

Table 1: Comparison of scheduling scenarios

	Service Preference contains ...	Configuration	Advantages	Disadvantages
Example 1	Dedicated bridge for scheduled conferences.	Single pool, with a single conference bridge. Pool marked to be used for scheduling in the TelePresence Conductor Service Preference. Pool is reported to Cisco TMS in capacity information requests.	Conference availability is guaranteed, subject to bridge failure (or full capacity). Maximizes use of resources, as Cisco TMS will book ports until the bridge is full.	Uses one conference bridge exclusively for scheduling. Cascaded conferencing does not occur: to avoid wasting resources, cascading should be disabled.
Example 2	■ Dedicated bridge for scheduled conferences ■ Dedicated backup bridge	Two pools. Both pools contain a single conference bridge. The second pool is used as a backup if the bridge in the highest priority pool fails. Only the first pool is marked for scheduling in the TelePresence Conductor Service Preference and reported to Cisco TMS.	As for Example 1, with added benefit of fallback in case of bridge failure.	Uses two conference bridges exclusively for scheduling. Consumes backup resources. To avoid wasting resources, cascading should be disabled.
Example 3	■ Dedicated bridge for scheduled conferences ■ Shared-use backup bridges for both scheduled and non-scheduled conferences	Two or more pools. Highest priority pool with one bridge only, used for scheduled conferences. Other pools contain bridges for both scheduled (as backup) and non-scheduled conferences. Only the first pool is marked for scheduling in the TelePresence Conductor Service Preference and reported to Cisco TMS.	As for Example 1, with possible benefit of fallback in case of bridge failure if the other pools have spare capacity.	Uses one conference bridge exclusively for scheduling. To avoid wasting resources on the dedicated bridge, cascading should be disabled.

Table 1: Comparison of scheduling scenarios (continued)

	Service Preference contains ...	Configuration	Advantages	Disadvantages
Example 4	■ Dedicated bridges for scheduled conferences ■ Shared-use backup bridges for both scheduled and non-scheduled conferences	Two or more pools. Highest priority pool with two or more bridges, used for scheduled conferences. Cascading enabled on the associated conference template. Other pools contain bridges for both scheduled (as backup and overflow) and non-scheduled conferences. For planned overflow you need to set Capacity Adjustment for the service preference to more than 100% in Cisco TMS. Only the first pool is marked for scheduling in the TelePresence Conductor Service Preference and reported to Cisco TMS.	As for Example 1, with possible benefit of fallback in case of bridge failure and overflow resource when cascading is used in a scheduled conference. Bridges in the backup pools are used for scheduling if: ■ A bridge in Pool 1 fails. ■ Cascading in Pool 1 uses up bridge resources that Cisco TMS expected to be available for scheduling.	Uses conference bridges exclusively for scheduling. If scheduled conferences are cascaded, they may need resources from a shared-use pool.
Example 5	Shared-use bridges for scheduled and non-scheduled conferences	One or more pools, shared for scheduled and non-scheduled conferences. All pools are marked for scheduling in the TelePresence Conductor Service Preference and reported to Cisco TMS.	Cascaded conferencing available (if enabled). Targeted management of bridge resources. Over time, monitoring of use patterns can identify the most appropriate pool configuration.	Resource availability for scheduled conferences not guaranteed (could be used up by non-scheduled conferences). This risk can be reduced by under-subscribing resources for the service preference in Cisco TMS using the Capacity Adjustment feature.

Figure 2: Illustration of scheduling scenarios

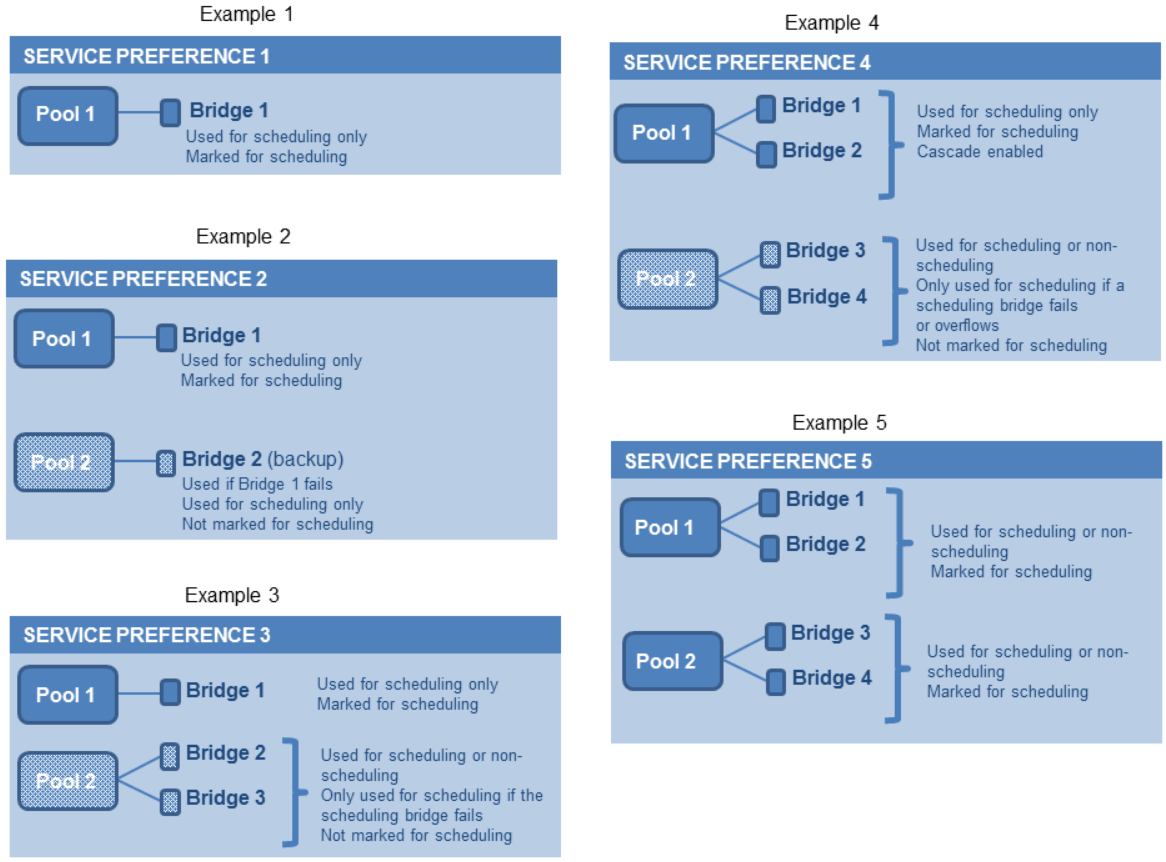

Configuration overview

TelePresence Conductor can be deployed with multiple Cisco VCSs or Unified CMs or a combination of the two. The configuration required on the TelePresence Conductor depends on the deployment.

Configuring TelePresence Conductor in a Cisco VCS deployment

In a Cisco VCS deployment a rendezvous conference is created when someone dials a pre-determined conference alias, e.g. **learn.math@example.com**. To enable this, you must first define the aliases that can be dialed, and the settings that will be applied to each conference when it is created.

Configuring rendezvous conferences that use the Cisco VCS's external policy service interface

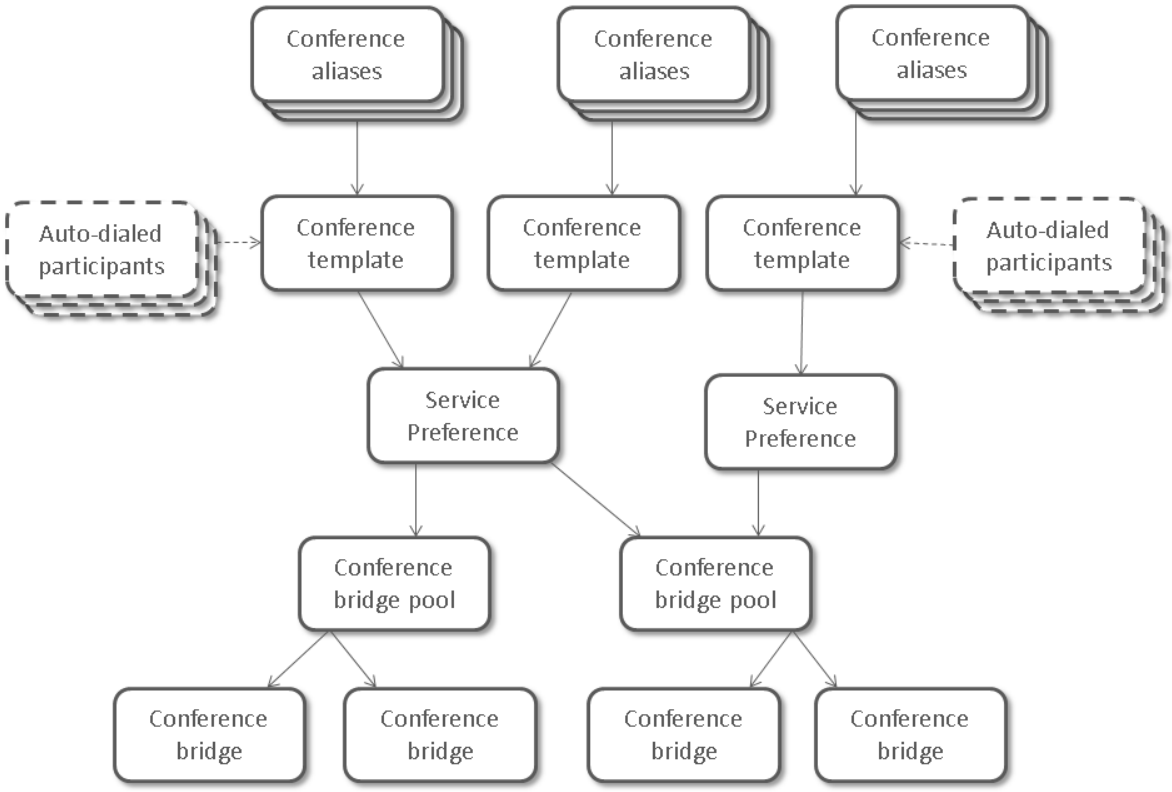

To configure a rendezvous conference in a Cisco VCS deployment using the Cisco VCS's external policy server interface:

1. Ensure that all the Cisco VCSs and conference bridges you will be using with this TelePresence Conductor are working properly together and have been configured in accordance with the information in the *Cisco TelePresence Conductor with Cisco TelePresence Video Communication Server Deployment Guide*.

2. Create one or more pools of conference bridges that the TelePresence Conductor will use for its conferences.
3. Add conference bridges to the pool. Each pool must contain at least one conference bridge.
4. Set up at least one Service Preference, defining the order in which conference bridge pools will be used when a large number of resources is needed.
5. Create a template for the conference. The template will determine whether the conference is a **Meeting** (where all participants dial in using the same conference alias and have the same privileges) or a **Lecture** (where the host(s) and the guests dial in using different aliases and are given different privileges).
6. Define a conference alias for the conference. A single conference can have more than one alias, and Lectures must have at least two aliases – one for the host(s) and one for the guests.
7. Optionally, define any auto-dialed participants whom you want to be called by the conference bridge when the conference starts. These participants can be individual endpoints, FindMe IDs, recording devices or even voice bridges.

Configuring rendezvous conferences that use the TelePresence Conductor's B2BUA

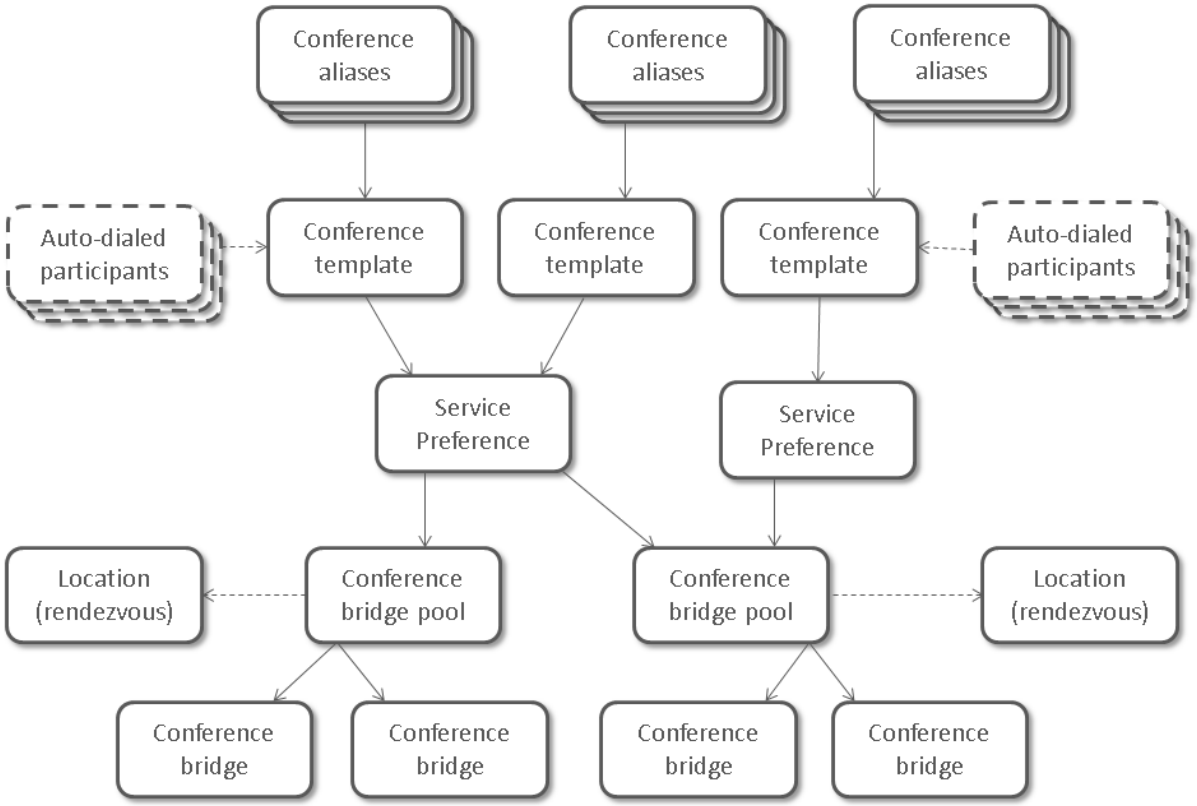

To configure a rendezvous conference in a Cisco VCS deployment using the TelePresence Conductor's back-to-back user agent (B2BUA):

1. Ensure that all the Cisco VCSs and conference bridges you will be using with this TelePresence Conductor are working properly together and have been configured in accordance with the information in the *Cisco TelePresence Conductor with Cisco TelePresence Video Communication Server (B2BUA) Deployment Guide*.
2. Add one local rendezvous IP address that the Cisco VCS connects to in order to create rendezvous

conferences on the TelePresence Conductor.

3. Add one Location for all Cisco VCSs that the TelePresence Conductor connects to via a SIP trunk. Each Location references a local rendezvous IP address and a trunk IP address, protocol and port on the Cisco VCS that is used for outbound rendezvous calls from the conference bridges.
4. Create one or more pools of conference bridges that the TelePresence Conductor will use for its conferences. Each pool, which contains conference bridges that make outbound calls to participants registered on a Cisco VCS, needs to reference a Location.
5. Add conference bridges to the pool. Each pool must contain at least one conference bridge. The conference bridge needs to have the correct SIP port for TLS defined. The default SIP port is 5061.
6. Set up at least one Service Preference, defining the order in which conference bridge pools will be used when a large number of resources is needed.
7. Create a template for the conference. The template will determine whether the conference is a **Meeting** (where all participants dial in using the same conference alias and have the same privileges) or a **Lecture** (where the host(s) and the guests dial in using different aliases and are given different privileges).
8. Define a conference alias for the conference. A single conference can have more than one alias, and Lectures must have at least two aliases – one for the host(s) and one for the guests.
9. Optionally, define any auto-dialed participants whom you want to be called by the conference bridge when the conference starts. These participants can be individual endpoints, FindMe IDs, recording devices or even voice bridges.

Configuring TelePresence Conductor in a Unified CM deployment

In a Unified CM deployment both rendezvous and ad hoc conferences can be configured. A rendezvous conference is initiated when someone dials a pre-determined conference alias. An ad hoc conference is initiated when two participants in a call add additional participants into a spontaneous conference.

Configuring rendezvous conferences

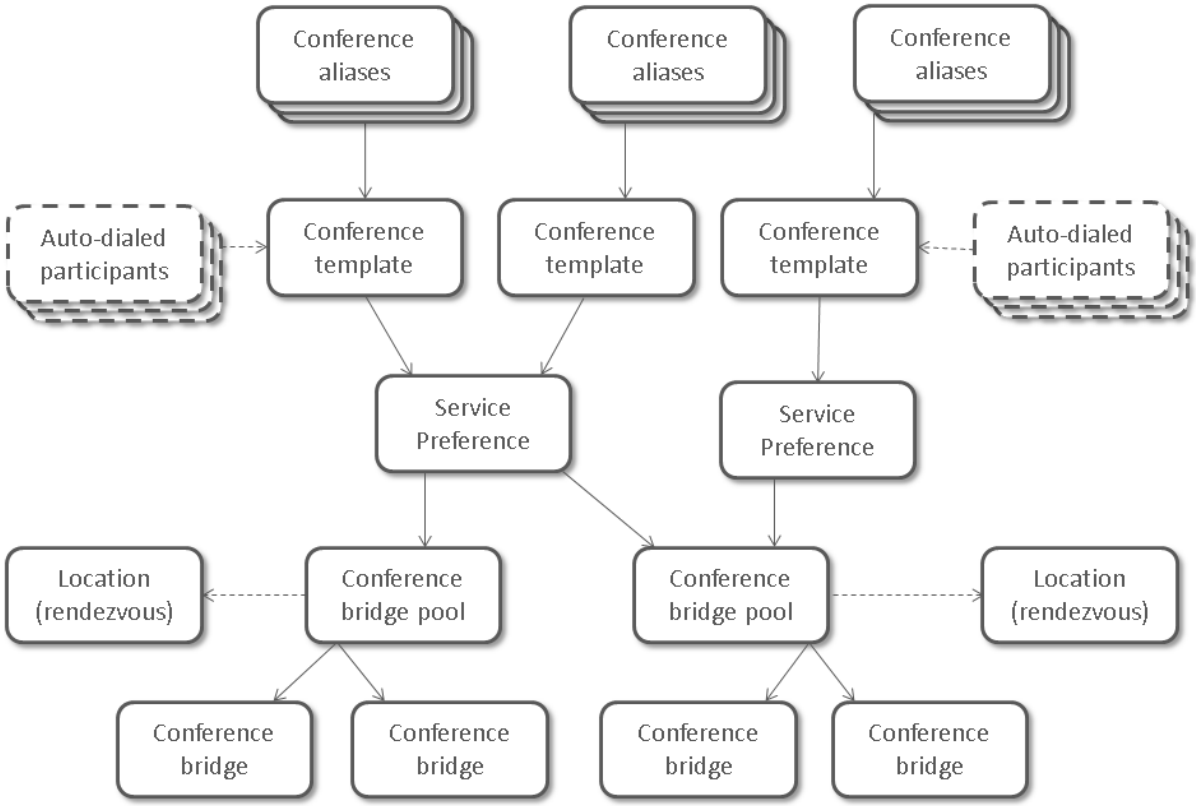

To configure a rendezvous conference in a Unified CM deployment:

1. Ensure that all the Unified CMs and conference bridges you will be using with this TelePresence Conductor are working properly together and have been configured in accordance with the information in the *Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide*.
2. Add all local rendezvous IP addresses that the Unified CM connects to in order to create rendezvous conferences on the TelePresence Conductor.
3. Add Locations for all locations in which rendezvous conferences will be created. Each Location references a local rendezvous IP address and a trunk IP address, protocol and port on the Unified CM that is used for outbound rendezvous calls from the conference bridges.
4. Create one or more pools of conference bridges that the TelePresence Conductor will use for its conferences. Each pool, which contains conference bridges that make outbound calls to participants registered on Unified CM, needs to reference a Location.
5. Add conference bridges to the pool. Each pool must contain at least one conference bridge. The conference bridge needs to have the correct SIP port for TLS defined. The default SIP port is 5061.
6. Set up at least one Service Preference, defining the order in which conference bridge pools will be used when a large number of resources is needed.
7. Create a template for the conference. The template will determine whether the conference is a **Meeting** (where all participants dial in using the same conference alias and have the same privileges) or a **Lecture** (where the host(s) and the guests dial in using different aliases and are given different privileges).
8. Define a conference alias for the conference. A single conference can have more than one alias, and Lectures must have at least two aliases – one for the host(s) and one for the guests.

9. Optionally, define any auto-dialed participants whom you want to be called by the conference bridge when the conference starts. These participants can be individual endpoints, FindMe IDs, recording devices or even voice bridges.

Configuring ad hoc conferences

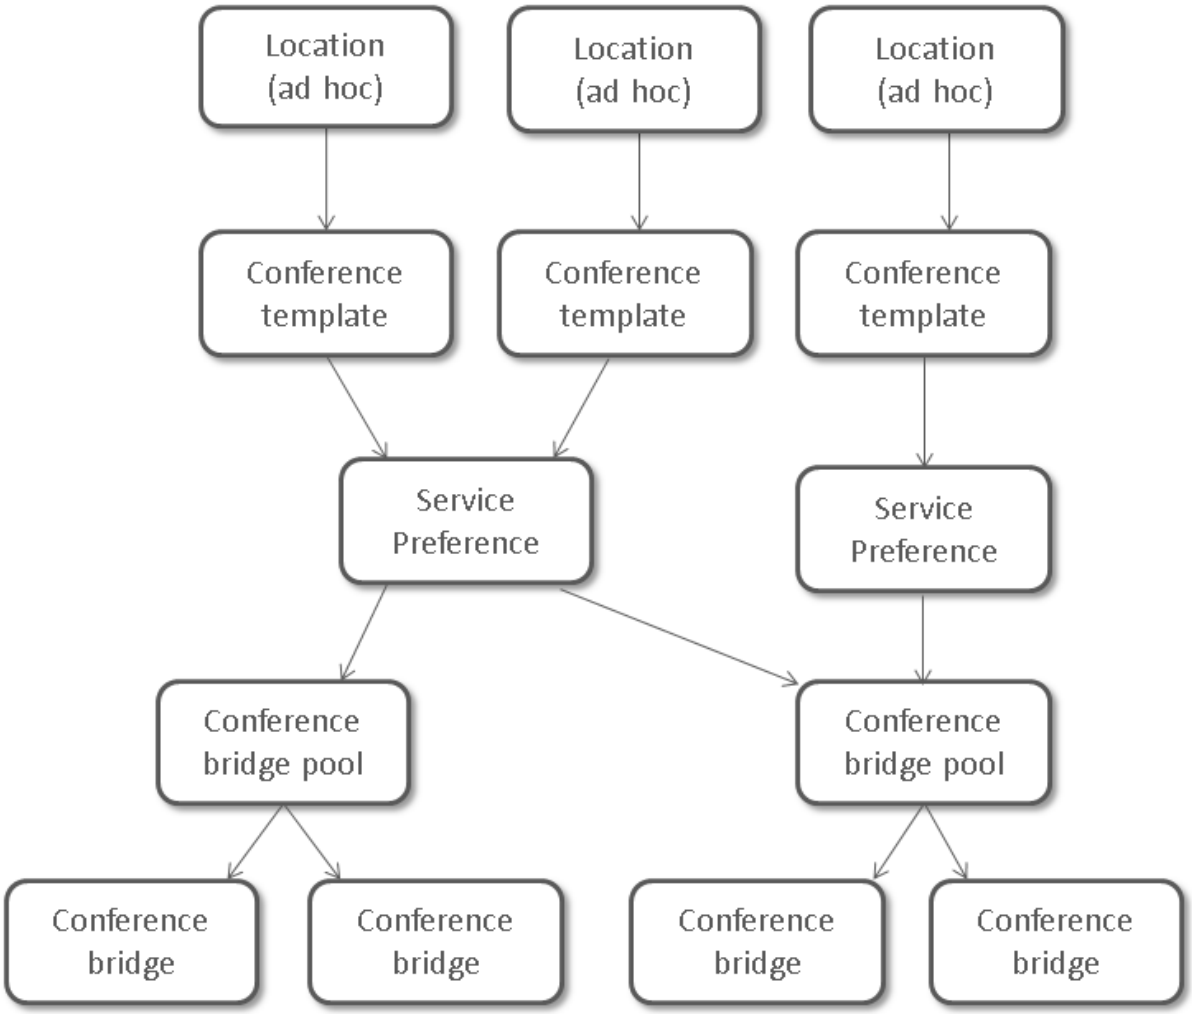

To configure an ad hoc conference in a Unified CM deployment:

1. Ensure that all the Unified CMs and conference bridges you will be using with this TelePresence Conductor are working properly together and have been configured in accordance with the information in the *Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide*.
2. Add all local ad hoc IP addresses that the Unified CM connects to in order to create ad hoc conferences on the TelePresence Conductor.
3. Create one or more pools of conference bridges that the TelePresence Conductor will use for its conferences.
4. Add conference bridges to the pool. Each pool must contain at least one conference bridge. The conference bridge needs to have the correct SIP port for TLS defined. The default SIP port is 5061.
5. Set up at least one Service Preference, defining the order in which conference bridge pools will be used when a large number of resources is needed.

6. Create a template for the conference.
7. Add or update Locations for all locations in which ad hoc conferences will be created. Each Location references a local ad hoc IP address and a conference template.

Configuration limits

General configuration limits

Configuration item	Limit
Conference bridges	Full capacity TelePresence Conductor: 30 TelePresence Conductor Select: 30 TelePresence Conductor Essentials: 1
Total number of calls	Full capacity TelePresence Conductor: 2,400 TelePresence Conductor Select: 50 TelePresence Conductor Essentials: the number of calls supported by the conference bridge
Maximum participants per conference	2,342

Configuration limits for conferences configured via the TelePresence Conductor web interface

Configuration item	Limit
Conference templates	1,000
Conference aliases (regex)	1,000
Auto-dialed participants	1,000
Locations	30

Configuration limits for conferences configured via the TelePresence Conductor's Provisioning API

Configuration item	Limit
Conference bundles (CMRs)	100,000
Direct match aliases	10 per ConfBundle/CMR 200,000 in total
Auto-dialed participants	10 per ConfBundle/CMR 100,000 in total

Configuring system settings

This section provides information on how to configure the system settings on the TelePresence Conductor, accessible via the **System** menu.

Configuring system administration settings

Most TelePresence Conductor administration can be performed using the web interface. The **System administration** page (**System > Administration**) is used to configure additional administration options available to administrators.

Note: tsh is not supported on the TelePresence Conductor and should not be used.

It is also possible to administer the TelePresence Conductor via a PC connected directly to the unit via a serial cable. Only root access is available via the serial cable.

Because access to the serial port allows the password to be reset, we recommend that you install the TelePresence Conductor in a physically secure environment.

The configurable options are:

Field	Description	Usage tips
Services		
Serial port / console	Whether the system can be accessed locally via either the serial port (for a physical system) or VMware console (for a virtual machine). Default is *On*.	The `pwrec` command to reset the root or administrator password is still available for one minute after a reboot, even if the serial port has been turned off.
SSH service	Whether the TelePresence Conductor can be accessed via SSH and SCP. Default is *On*.	
LCD panel	Whether any information will be displayed on the LCD panel on the front of the TelePresence Conductor unit (if you are using an appliance TelePresence Conductor). *On*: the LCD panel will display the product name (Cisco TelePresence Conductor), the LAN 1 IPv4 address, and any alarms that may apply to the unit. It is also possible to configure the IP settings and reboot the TelePresence Conductor unit via the LCD panel. *Off*: The LCD panel will display **Cisco**. Default is *On*.	See below for more information on configuring the TelePresence Conductor using the front panel.
Web interface (over HTTPS)	Whether the TelePresence Conductor can be accessed via the web interface. Default is *On*.	
Session limits		
Session time out	The number of minutes that an administration session (serial port, HTTPS or SSH) may be inactive before the session is timed out. Default is 30 minutes.	

Field	Description	Usage tips
Per-account session limit	The number of concurrent sessions that each individual administrator account is allowed on each TelePresence Conductor.	This includes web, SSH and serial sessions. A value of 0 turns session limits off.
System session limit	The maximum number of concurrent administrator sessions allowed on each TelePresence Conductor.	This includes web, SSH and serial sessions. A value of 0 turns session limits off.
System protection		
Automated protection service	Whether the automated protection service is active. Default is *Off*.	After enabling the service you must go and configure the specific protection categories.
Automatic discovery protection	Controls how management systems such as Cisco TMS can discover this TelePresence Conductor. *Off*: automatic discovery is allowed. *On*: Cisco TMS has to be manually configured to discover this TelePresence Conductor and must provide administrator account credentials. Default is *Off*.	You must restart the system for any changes to take effect.
Web server configuration		
Redirect HTTP requests to HTTPS	Determines whether HTTP requests are redirected to the HTTPS port. Default is *On*.	HTTPS must also be enabled for access via HTTP to function.
HTTP Strict Transport Security (HSTS)	Determines whether web browsers are instructed to only ever use a secure connection to access this server. Enabling this feature gives added protection against man-in-the-middle (MITM) attacks. *On*: the Strict-Transport-Security header is sent with all responses from the web server, with a 1 year expiry time. *Off*: the Strict-Transport-Security header is not sent, and browsers work as normal. Default is *On*.	See below for more information about HSTS.

HTTP Strict Transport Security (HSTS)

HTTP Strict Transport Security (HSTS) provides a mechanism where a web server forces a web browser to communicate with it using secure connections only.

As of August 2014, this mechanism is supported by the following browsers:

- Chrome, versions 4.0.211.0 and later
- Firefox, versions 4 and later

When HSTS is enabled, a browser that supports HSTS will:

- Automatically turn any insecure links to the website into secure links (for example, **`http://example.com/page/`** is modified to **`https://example.com/page/`** before accessing the server).
- Only allow access to the server if the connection is secure (for example, the server's TLS certificate is valid, trusted and not expired).

Browsers that do not support HSTS will ignore the Strict-Transport-Security header and work as before. They will still be able to access the server.

Compliant browsers only respect Strict-Transport-Security headers if they access the server through its fully qualified name (rather than its IP address).

Configuring the TelePresence Conductor using the front panel

(This is only applicable if you are using an appliance TelePresence Conductor.)

The LCD panel and buttons at the front of the TelePresence Conductor allow you to configure and check the IP settings as well as to reboot the system. We do not recommend that you perform the initial configuration using the front panel, but you may need to use this method if you do not have access to a PC and serial cable.

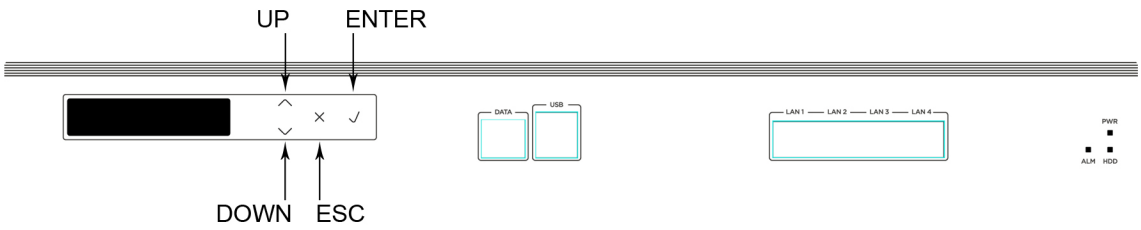

By default, during normal operation the front panel of the TelePresence Conductor unit shows a rotating display of the Cisco TelePresence Conductor's system name, the LAN 1 IPv4 address, and any alarms that may apply to the unit.

To control the display of status items:

- **`ENTER`** stops the display from automatically rotating through the status items. This is useful if you need to review all of the alarm messages. Press **`ENTER`** again to resume the rotating display.
- **`UP/DOWN`** displays the previous or next status item.

To access the front panel menu options, press **`ESC`**. The menu options are as follows:

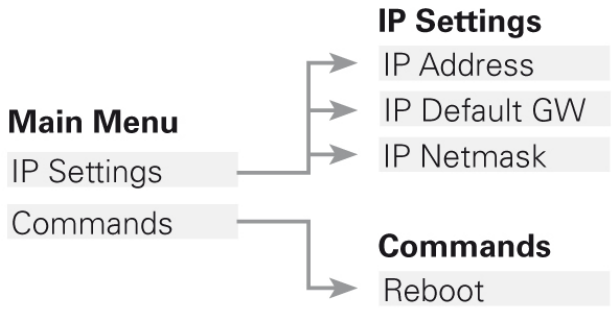

- **UP/DOWN** navigates to the next menu or submenu item.
- **ENTER** selects a menu or submenu item.
- When you have selected an IP setting from the menu:
  - **UP/DOWN** increases or decreases the value of the currently selected digit.
  - **ENTER** moves the cursor on to the next digit and saves your changes when you move off the final digit.
  - **ESC** cancels any changes; returns to the menu.
- When you are on the **Commands** submenu:
  - **ENTER** performs the command.
  - **ESC** returns to the main menu.
- To leave the menu options and return to the rotating display, press **ESC**.

Configuring Ethernet settings

Use the **Ethernet** page (**System > Network interfaces > Ethernet**) to configure the speed of the connection between the TelePresence Conductor and the Ethernet network to which it is connected. The speed and duplex mode must be the same at both ends of the connection.

The default **Speed** is *Auto*, which means that the TelePresence Conductor and the connected switch will automatically negotiate the speed and duplex mode.

Note: We recommend *Auto* unless the connected switch is unable to auto-negotiate. A mismatch in speed/duplex mode between the two ends of the connection will cause packet loss and could make the system inaccessible.

Configuring IP settings

The **IP** page (**System > Network interfaces > IP**) is used to configure the IP settings of the TelePresence Conductor.

A restart is required for any IP configuration changes to take effect.

Configuration

You can set the default **IPv4 gateway** used by the TelePresence Conductor. This is the gateway to which IP requests are sent for IP addresses that do not fall within the TelePresence Conductor's local subnet. The default **IPv4 gateway** is `127.0.0.1`, which should be changed during the commissioning process.

Primary LAN 1 IP address

LAN 1 is the primary network port on the TelePresence Conductor.

You can configure the primary **IPv4 address** and **subnet mask** for this port. Their default values are `192.168.0.100` and `255.255.255.0` respectively. The **IPv4 address** should be changed during the commissioning process. The **IPv4 subnet mask** should be changed if necessary.

The **Maximum transmission unit (MTU)** is the maximum Ethernet packet size that can be sent over the network interface. The default is 1500 bytes.

The primary LAN 1 IP settings cannot be changed, if this TelePresence Conductor is part of a cluster.

Additional addresses for LAN 1

When configuring the TelePresence Conductor to work with Unified CM or with the Cisco VCS in a deployment using the TelePresence Conductor's B2BUA, additional IP addresses must be configured on LAN 1 for ad hoc or rendezvous calls. Up to 64 IP addresses can be added on the **IP** page (**System > Network interfaces > IP**). The IP addresses must be in the same subnet as the primary IP address. Each cluster peer must have a unique set of IP addresses for each SIP trunk configuration.

A restart is required, after all additional IP addresses have been added, for the changes to take effect.

When the IP addresses have been added, they can be assigned to a Location.

Configuring DNS settings

The **DNS** page (**System > DNS**) is used to configure the TelePresence Conductor's DNS settings and DNS servers.

DNS settings

System host name

The **System host name** defines the DNS host name by which this TelePresence Conductor is known. This is not the fully-qualified domain name (FQDN), just the host label portion.

The name can only contain letters, digits, hyphens and underscores. The first character must be a letter and the last character must be a letter or a digit.

The table below shows where the **System host name** is used, and what will be shown instead if it is not configured.

Location	Notes
Web interface	If not configured, the system's IP address will be used instead.
Front panel of unit (so that you can identify it when it is in a rack with other systems)	If not configured, the system's IP address will be used instead. If the **System host name** is longer than 16 characters, only the last 16 characters are shown in the display on the front panel.
Remote log server	If not configured, the remote syslog server will show a default name of **TANDBERG**.

Note: The **System host name** must be unique for each peer in a cluster.

We recommend that you give the TelePresence Conductor a **System host name** that allows you to easily and uniquely identify it.

Domain name

The **Domain name** is used when attempting to resolve unqualified server addresses (for example `ldapserver`). It is appended to the unqualified server address before the query is sent to the DNS server. If the server address is fully qualified (for example `ldapserver.mydomain.com`) or is in the form of an IP address, the domain name is not appended to the server address before querying the DNS server.

It applies to the following configuration settings in the TelePresence Conductor:

- LDAP server
- NTP servers
- Remote logging server
- Conference bridges

We recommend that you use IP addresses for conference bridges, and an IP address or FQDN (Fully Qualified Domain Name) for all server addresses.

The **Domain name** may also be used along with the local **System host name** to identify references to this TelePresence Conductor in SIP messaging.

Note: The FQDN of the TelePresence Conductor is the **System host name** plus the **Domain name**.

DNS servers

You must specify at least one DNS server to be queried for address resolution if you want to use FQDNs (Fully Qualified Domain Names) instead of IP addresses when specifying external addresses (for example for LDAP and NTP servers, or conference bridges).

Note: If you do not configure any DNS servers, you must ensure that your NTP servers are configured using IP addresses so that NTP time can still be obtained. This is because NTP time is required for correct system operation.

Default DNS servers

Note: For reliability we recommend specifying at least two DNS servers, otherwise DNS could become a single point of failure for your deployment.

You can specify up to three default DNS servers. These default DNS servers are used if there is no **Per-domain DNS server** defined for the domain being looked up.

- The TelePresence Conductor only queries one server at a time; if that server is not available the TelePresence Conductor will try another server from the list.
- The order that the servers are specified is not significant; the TelePresence Conductor attempts to favor servers that were last known to be available.

Per-domain DNS servers

In addition to the three default DNS servers, you can specify three additional explicit DNS servers for specified domains. This can be useful in deployments where specific domain hierarchies need to be routed to their explicit authorities.

For each additional per-domain DNS server address you can specify up to two **Domain names**. Any DNS queries under those domains are forwarded to the specified DNS server instead of the default DNS servers.

You can specify redundant per-domain servers by adding an additional per-domain DNS server address and associating it with the same **Domain names**. In this scenario, DNS requests for those domains will be sent in parallel to both DNS servers.

Tip: You can also use the DNS lookup tool (**Maintenance > Tools > Network utilities > DNS lookup**) to check which domain name server (DNS server) is responding to a request for a particular hostname.

Configuring time settings

The **Time** page (**System > Time**) is used to configure the TelePresence Conductor's NTP servers and specify your local time zone.

Configuring the NTP servers

An NTP server is a remote server with which the TelePresence Conductor synchronizes its time in order to ensure that its time is accurate. The NTP server provides the TelePresence Conductor with UTC time.

Accurate time is necessary for correct system operation.

Note: It is essential for a TelePresence Conductor to have access to an NTP server if it is in a cluster of other TelePresence Conductors.

To configure the TelePresence Conductor with one or more NTP servers to be used when synchronizing system time, enter up to five **Address**es in one of the following formats, depending on the system's DNS settings. (You can check these settings on the **DNS** page, **System > DNS**):

- if there are no **DNS servers** configured, you must use an IP address for the NTP server
- if there are one or more **DNS servers** configured, you can use an FQDN or IP address for the NTP server
- if there is a DNS **Domain name** configured in addition to one or more **DNS servers**, you can use the server name, FQDN or IP address for the NTP server.

To configure the authentication method to use with the individual NTP servers use one of the following options for the **Authentication** field:

Authentication method	Description
Disabled	No authentication method
Private key	Private key authentication. Using this method automatically generates a private key in the background, with which messages sent to the NTP server are authenticated.
Symmetric key	Symmetric key authentication. When using this method the **Key ID**, **Hash** method and **Pass phrase** need to be specified. The values must match exactly the values on the NTP server. More than one NTP server can be configured to have the same combination of values. If a different **Pass phrase** is specified, the Key ID must also be unique and cannot be the same value as any Key ID already used on this device.

Displaying NTP status information

The synchronization status between the NTP server and the TelePresence Conductor is shown in the **Status** area as follows:

- *Starting*: the NTP service is starting.
- *Synchronized*: the TelePresence Conductor has successfully obtained accurate system time from an NTP server.
- *Unsynchronized*: the TelePresence Conductor is unable to obtain accurate system time from an NTP server.

- *Down*: the TelePresence Conductor's NTP client is not running.
- *Reject*: the NTP service is not accepting NTP responses.

Updates may take a few minutes to be displayed in the status table.

Other status information available includes:

Field	Description
NTP server	The actual NTP server that has responded to the request. This may be different to the NTP server in the NTP server address field.
Condition	Gives a relative ranking of each NTP server. All servers that are providing accurate time are given a status of *Candidate*; of those, the server that the TelePresence Conductor considers to be providing the most accurate time and is therefore using shows a status of *sys.peer*.
Flash	A code giving information about the server's status. `00 ok` means there are no issues. See the Flash status word reference table [p.224] for a complete list of codes.
Authentication	Indicates the status of the current authentication method. One of *ok*, *bad* or *none*. *none* is specified when the **Authentication** method is *Disabled*.
Event	Shows the last event as determined by NTP (for example *reachable* or *sys_peer*)
Reachability	Indicates the results of the 8 most recent contact attempts between the TelePresence Conductor and the NTP server, with a tick indicating success and a cross indicating failure. The result of the most recent attempt is shown on the far right. Each time the NTP configuration is changed, the NTP client is restarted and the **Reachability** field will revert to all crosses apart from the far right indicator which will show the result of the first connection attempt after the restart. However, the NTP server may have remained contactable during the restart process.
Offset	The difference between the NTP server's time and the TelePresence Conductor's time.
Delay	The network delay between the NTP server and the TelePresence Conductor.
Stratum	The degree of separation between the TelePresence Conductor and a reference clock. `1` indicates that the NTP server is a reference clock.
Ref ID	A code identifying the reference clock.
Ref time	The last time that the NTP server communicated with the reference clock.

For definitions of the remaining fields on this page, and for further information about NTP, see Network Time Protocol website.

TelePresence Conductor time display and time zone

Local time is used throughout the web interface. It is shown in the system information bar at the bottom of the screen and is used to set the timestamp that appears at the start of each line in the Event Log.

Note: A UTC timestamp is included at the end of each entry in the Event Log.

Internally, the TelePresence Conductor maintains its system time in UTC. It is based on the TelePresence Conductor's operating system time, which is synchronized using an NTP server if one is configured. If no NTP servers are configured, the TelePresence Conductor uses its own operating system time to determine the time and date.

Specifying your local **Time zone** lets the TelePresence Conductor determine the local time where the system is located. It does this by offsetting UTC time by the number of hours associated with the selected time zone. It also adjusts the local time to account for summer time (also known as daylight saving time) when appropriate.

Configuring SNMP settings

The **SNMP** page (**System > SNMP**) is used to configure the TelePresence Conductor's SNMP settings.

Tools such as Cisco TelePresence Management Suite (Cisco TMS) or HP OpenView may act as SNMP Network Management Systems (NMS). They allow you to monitor your network devices, including the TelePresence Conductor, for conditions that might require administrative attention.

The TelePresence Conductor supports the most basic MIB-II tree (.1.3.6.1.2.1) as defined in *RFC 1213*.

The information made available by the TelePresence Conductor includes the following:

- system uptime
- system name
- location
- contact
- interfaces
- disk space, memory, and other machine-specific statistics

By default, SNMP is *Disabled*, therefore to allow the TelePresence Conductor to be monitored by an SNMP NMS (including Cisco TMS), you must select an alternative **SNMP mode**. The configurable options are:

Field	Description	Usage tips
SNMP mode	Controls the level of SNMP support. *Disabled*: no SNMP support. *v3 secure SNMP*: supports authentication and encryption. *v3 plus TMS support*: secure SNMPv3 plus non-secure access to OID 1.3.6.1.2.1.1.2.0 only. *v2c*: non-secure community-based SNMP.	If you want to use secure SNMPv3 but you also use Cisco TMS as your external manager, you must select *v3 plus TMS support*.
Description	Custom description of the system as viewed by SNMP. The default is to have no custom description (empty field).	When you leave this field empty, the system uses its default SNMP description.
Community name	The TelePresence Conductor's SNMP community name. The default is *public*.	Only applies when using *v2c* or *v3 plus TMS support*.
System contact	The name of the person who can be contacted regarding issues with the TelePresence Conductor. The default is *Administrator*.	The **System contact** and **Location** are used for reference purposes by administrators when following up on queries.
Location	Specifies the physical location of the TelePresence Conductor.	
Username	The TelePresence Conductor's SNMP username, used to identify this SNMP agent to the SNMP manager.	Only applies when using *v3 secure SNMP* or *v3 plus TMS support*
v3 Authentication settings (only applicable to SNMPv3)		

Field	Description	Usage tips
Authentication mode	Enables or disables SNMPv3 authentication.	
Type	The algorithm used to hash authentication credentials. *SHA*: Secure Hash Algorithm. *MD5*: Message-Digest algorithm 5. The default is *SHA*.	
Password	The password used to encrypt authentication credentials.	Must be at least 8 characters.
v3 Privacy settings (only applicable to SNMPv3)		
Privacy mode	Enables or disables SNMPv3 encryption.	
Type	The security model used to encrypt messages. *DES*: Data Encryption Standard 56-bit encryption. *AES*: Advanced Encryption Standard 128-bit encryption. If available, the default and recommended setting is *AES*.	
Password	The password used to encrypt messages.	Must be at least 8 characters.

The TelePresence Conductor does not support SNMP traps or SNMP sets, therefore it cannot be managed via SNMP.

Note: SNMP is disabled by default, because of the potentially sensitive nature of the information involved. Do not enable SNMP on a TelePresence Conductor on the public internet or in any other environment where you do not want to expose internal system information.

Configuring firewall rules

Firewall rules provide the ability to configure IP table rules to control access to the TelePresence Conductor at the IP level. On the TelePresence Conductor, these rules have been classified into groups and are applied in the following order:

- Dynamic system rules: these rules ensure that all established connections/sessions are maintained. They also include any rules that have been inserted by the automated detection feature as it blocks specific addresses. Finally, it includes a rule to allow access from the loopback interface.
- Non-configurable application rules: this incorporates all necessary application-specific rules, for example to allow SNMP traffic.
- User-configurable rules: this incorporates all of the manually configured firewall rules (as described in this section) that refine — and typically restrict — what can access the TelePresence Conductor. There is a final rule in this group that allows all traffic destined for the TelePresence Conductor LAN 1 interface.

There is also a final, non-configurable rule that drops any broadcast or multicast traffic that has not already been specifically allowed or denied by the previous rules.

By default any traffic that is destined for the specific IP address of the TelePresence Conductor is allowed access, but that traffic will be dropped if the TelePresence Conductor is not explicitly listening for it. You have to actively configure extra rules to lock down the system to your specifications.

Note that return traffic from outbound connections is always accepted.

User-configured rules

The user-configured rules are typically used to restrict what can access the TelePresence Conductor. You can:

- Specify the source IP address subnet from which to allow or deny traffic.
- Choose whether to drop or reject denied traffic.
- Configure well known services such as SSH, HTTP/HTTPS or specify customized rules based on transport protocols and port ranges.
- Specify the priority order in which the rules are applied.

Setting up and activating firewall rules

The **Firewall rules configuration** page is used to set up and activate a new set of firewall rules.

The set of rules shown will initially be a copy of the current active rules. (On a system where no firewall rules have previously been defined, the list will be empty.) If you have a lot of rules you can use the **Filter** options to limit the set of rules displayed. Note that the built-in rules are not shown in this list.

You can then change the set of firewall rules by adding new rules, or by modifying or deleting any existing rules. Any changes made at this stage to the current active rules are held in a pending state. When you have completed making all the necessary changes you can activate the new rules, replacing the previous set.

To set up and activate new rules:

1. Go to **System > Protection > Firewall rules > Configuration**.
2. Make your changes by adding new rules, or by modifying or deleting any existing rules as required.

   You can change the order of the rules by using the up/down arrows and to swap the priorities of adjacent rules.

- New or modified rules are shown as **Pending** (in the **State** column).
- Deleted rules are shown as **Pending delete**.

3. When you have finished configuring the new set of firewall rules, click **Activate firewall rules**.
4. Confirm that you want to activate the new rules. This will replace the existing set of active rules with the set you have just configured.
   After confirming that you want to activate the new rules, they are validated and any errors reported.
5. If there are no errors, the new rules are temporarily activated and you are taken to the **Firewall rules confirmation** page.
   You now have 15 seconds to confirm that you want to keep the new rules:
   - Click **Accept changes** to permanently apply the rules.
   - If the 15 seconds time limit expires or you click **Rollback changes**, the previous rules are reinstated and you are taken back to the configuration page.

   The automatic rollback mechanism provided by the 15 seconds time limit ensures that the client system that activated the changes is still able to access the system after the new rules have been applied. If the client system is unable to confirm the changes (because it can no longer access the web interface) then the rollback will ensure that its ability to access the system is reinstated.

Rule settings

The configurable options for each rule are:

Field	Description	Usage tips
Priority	The order in which the firewall rules are applied.	The rules with the highest priority (1, then 2, then 3 and so on) are applied first. Firewall rules must have unique priorities. Rule activation will fail if there are multiple rules with the same priority.
IP address and **Prefix length**	These two fields together determine the range of IP addresses to which the rule applies.	The **Address range** field shows the range of IP addresses to which the rule applies, based on the combination of the **IP address** and **Prefix length**. The prefix length range is 0-32 for an IPv4 address.
Service	Choose the service to which the rule applies, or choose *Custom* to specify your own transport type and port ranges.	Note that if the destination port of a service is subsequently reconfigured on the TelePresence Conductor, for example from 80 to 8080, any firewall rules containing the old port number will not be automatically updated.
Transport	The transport protocol to which the rule applies.	Only applies if specifying a *Custom* service.
Start and **end port**	The port range to which the rule applies.	Only applies if specifying a UDP or TCP *Custom* service.

Field	Description	Usage tips
Action	The action to take against any IP traffic that matches the rule. *Allow*: Accept the traffic. *Drop*: Drop the traffic without any response to the sender. *Reject*: Reject the traffic with an 'unreachable' response.	Dropping the traffic means that potential attackers are not provided with information as to which device is filtering the packets or why. For deployments in a secure environment, you may want to configure a set of low priority rules (for example, priority 50000) that deny access to all services and then configure higher priority rules (for example, priority 20) that selectively allow access for specific IP addresses.
Description	An optional free-form description of the firewall rule.	If you have a lot of rules you can use the **Filter** by description options to find related sets of rules.

Current active firewall rules

The **Current active firewall rules** page (**System > Protection > Firewall rules > Current active rules**) shows the user-configured firewall rules that are currently in place on the system. There is also a set of built-in rules that are not shown in this list.

If you want to change the rules you must go to the **Firewall rules configuration** page from where you can set up and activate a new set of rules.

Configuring automated intrusion protection

The automated protection service can be used to detect and block malicious traffic and to help protect the TelePresence Conductor from dictionary-based attempts to breach login security.

It works by parsing the system log files to detect repeated failures to access specific service categories, such as SIP, SSH and web/HTTPS access. When the number of failures within a specified time window reaches the configured threshold, the source host address (the intruder) and destination port are blocked for a specified period of time. The host address is automatically unblocked after that time period so as not to lock out any genuine hosts that may have been temporarily misconfigured.

You can configure ranges of addresses that are exempted from one or more categories (see Configuring exemptions [p.56] below).

Automated protection should be used in combination with the firewall rules feature - use automated protection to dynamically detect and temporarily block specific threats, and use firewall rules to permanently block a range of known host addresses.

About protection categories

The set of available protection categories on your TelePresence Conductor are pre-configured according to the software version that is running. You can enable, disable or configure each category, but you cannot add additional categories.

The rules by which specific log file messages are associated with each category are also pre-configured and cannot be altered. You can view example log file entries that would be treated as an access failure/intrusion within a particular category by going to **System > Protection > Automated detection > Configuration** and clicking on the name of the category. The examples are displayed above the **Status** section at the bottom of the page.

Enabling automated protection

To enable intrusion protection on your TelePresence Conductor:

1. Go to **System > Administration**.
2. Set **Automated protection service** to *On*.
3. Click **Save**.
4. You must then ensure that the required protection categories are enabled and configured, and that any required exemptions are specified, as described below.
   All protection categories are disabled by default.

Configuring protection categories

The **Automated detection overview** page (**System > Protection > Automated detection > Configuration**) is used to enable and configure the TelePresence Conductor's protection categories, and to view current activity.

The page displays a summary of all available categories, showing:

- **Status**: this indicates if the category is configured to be *On* or *Off*. When *On*, it additionally indicates the state of the category: this is normally *Active*, but may temporarily display *Initializing* or *Shutting down* when a category has just been enabled or disabled. Check the alarms if it displays *Failed*.)

- **Currently blocked**: the number of addresses currently being blocked for this category.
- **Total failures**: the total number of failed attempts to access the services associated with this category.
- **Total blocks**: the total number of times that a block has been triggered. Note that:
  - The **Total blocks** will typically be less than the **Total failures** (unless the **Trigger level** is set to 1).
  - The same address can be blocked and released several times per category, with each occurrence counting as a separate block.
- **Exemptions**: the number of addresses that are configured as exempt from this category.

From this page, you can also view any currently blocked addresses or any exemptions that apply to a particular category.

Enabling and disabling categories

To enable or disable one or more protection categories:

1. Go to **System > Protection > Automated detection > Configuration**.
2. Select the check box alongside the categories you want to enable or disable.
3. Click **Enable** or **Disable** as appropriate.

Configuring a category's blocking rules

To configure a category's specific blocking rules:

1. Go to **System > Protection > Automated detection > Configuration**.
2. Click on the name of the category you want to configure.
   You are taken to the configuration page for that category.
3. Configure the category as required:
   - **State**: whether protection for that category is enabled or disabled.
   - **Description**: a free-form description of the category.
   - **Trigger level** and **Detection window**: these settings combine to define the blocking threshold for the category. They specify the number of failed access attempts that must occur before the block is triggered, and the time window in which those failures must occur.
   - **Block duration**: the period of time for which the block will remain in place.
4. Click **Save**.

Configuring exemptions

The **Automated detection exemptions** page (**System > Protection > Automated detection > Exemptions**) is used to configure any IP addresses that are to be exempted always from one or more protection categories.

To configure exempted addresses:

1. Go to **System > Protection > Automated detection > Exemptions**.
2. Click on the **Address** you want to configure, or click **New** to specify a new address.
3. Enter the **Address** and **Prefix length** to define the range of IPv4 addresses you want to exempt.
4. Select the categories from which the address is to be exempted.
5. Click **Add address**.

Note that if you exempt an address that is currently blocked, it will remain blocked until its block duration expires (unless you unblock it manually via the **Blocked addresses** page).

Managing blocked addresses

The **Blocked addresses** page (**System > Protection > Automated detection > Blocked addresses**) is used to manage the addresses that are currently blocked by the automated protection service:

- It shows all currently blocked addresses and from which categories those addresses have been blocked.
- You can unblock an address, or unblock an address and at the same time add it to the exemption list. Note that if you want to permanently block an address, you must add it to the set of configured firewall rules.

If you access this page via the links on the **Automated detection overview** page it is filtered according to your chosen category. It also shows the amount of time left before an address is unblocked from that category.

Investigating access failures and intrusions

If you need to investigate specific access failures or intrusion attempts, you can review all the relevant triggering log messages associated with each category. To do this:

1. Go to **System > Protection > Automated detection > Configuration**.
2. Click on the name of the category you want to investigate.
3. Click **View all matching intrusion protection triggers for this category**.
   The system will display all the relevant events for that category. You can then search through the list of triggering events for the relevant event details such as a user name, address or alias.

Automated protection service and clustered systems

When the automated protection service is enabled in a clustered system:

- Each peer maintains its own count of connection failures and the trigger threshold must be reached on each peer for the intruder's address to be blocked by that peer.
- Addresses are blocked against only the peer on which the access failures occurred. This means that if an address is blocked against one peer it may still be able to attempt to access another peer (from which it may too become blocked).
- A blocked address can only be unblocked for the current peer. If an address is blocked by another peer, you must log in to that peer and then unblock it.
- Category settings and the exemption list are applied across the cluster.
- The statistics displayed on the **Automated detection overview** page are for the current peer only.

Additional information

- When a host address is blocked and tries to access the system, the request is dropped (the host receives no response).
- A host address can be blocked simultaneously for multiple categories, but may not necessarily be blocked by all categories. Those blocks may also expire at different times.
- When an address is unblocked (either manually or after its block duration expires), it has to fail again for the full number of times as specified by the category's trigger level before it will be blocked for a second time by that category.

- A category is reset whenever it is enabled. All categories are reset if the system is restarted or if the automated protection service is enabled at the system level. When a category is reset:
  - Any currently blocked addresses are unblocked.
  - Its running totals of failures and blocks are reset to zero.
- You can view all Event Log entries associated with the automated protection service by clicking **View all intrusion protection events** on the **Automated detection overview** page.

Configuring the Login page

The **Login page configuration** page (**System > Login page**) is used to specify a message and image to appear on the login page.

The **Welcome message title** and **text** appears to administrators when attempting to log in using the CLI or the web interface.

You can upload an image that will appear above the welcome message on the login page when using the web interface.

- Supported image file formats are JPG, GIF and PNG.
- Images larger than 200x200 pixels will be scaled down.

Managing conference bridges

This section provides information on how to manage conference bridge pools and the conference bridges that are contained in them. It also explains how to change the settings for conference bridges.

About conference bridges

You must configure the TelePresence Conductor with one or more pools of conference bridges that it can use to host the conferences it creates. Conference bridge pools contain conference bridges that are of the same conference bridge type, have the same software versions installed and have identical configuration.

It is possible to combine a TelePresence Server and a Cisco TelePresence Server on Virtual Machine within the same pool. It is not possible to combine a TelePresence Server and a TelePresence MCU within the same pool. We recommend that you install the latest software versions; otherwise some features will not be supported.

The TelePresence Conductor periodically monitors all conference bridges in each of its pools for availability and resource usage. Upon receipt of a conference alias request from a Cisco VCS or Unified CM, the TelePresence Conductor checks the resource availability of the conference bridges in the preferred pool. It selects a suitable conference bridge and creates a conference on it. It may cascade the conference to one or more additional conference bridges if and when required. If the preferred pool cannot be used, the TelePresence Conductor will check the availability of the conference bridges in the pool that has the next highest preference.

Creating conference bridge pools

Each conference bridge must belong to a conference bridge pool. A single conference bridge can only belong to one pool. A single conference bridge pool can contain up to 30 conference bridges. Each TelePresence Conductor (or cluster of TelePresence Conductors) can use up to 30 conference bridges across all of its pools. For example, you could have a single pool with 30 conference bridges, or one pool with six conference bridges plus two pools with 12 conference bridges each.

The full capacity TelePresence Conductor supports up to 2400 concurrent calls, with each cluster of TelePresence Server blades supporting a maximum of 200 concurrent calls (104 for TelePresence Server version 3.x or earlier) and TelePresence MCUs supporting up to 80 concurrent calls. TelePresence Conductor Select supports up to 50 concurrent calls. See TelePresence Conductor capacity versions for more information.

To create a conference bridge pool:

1. Go to **Conference configuration > Conference bridge pools**.
   You will see a list of any existing conference bridge pools.
2. Click **New**.
   Enter the details of the new pool. The configurable options are:

Field	Description
Pool name	Descriptive name of the conference bridge pool.
Description	A free-form description of the conference bridge pool.
Conference bridge type	The type of conference bridges that can be assigned to this pool. All conference bridges within a pool must be of the same type and have the same configuration. This release of the TelePresence Conductor supports *Cisco TelePresence MCU* and *Cisco TelePresence Servers* only. Future versions may support other types of conference bridges.
Raise conference bridge resource alarm	Determines whether or not an alarm will be raised when the quantity of conference bridge resources being used and requested within this conference bridge pool exceeds a given percentage of the total quantity of resources available across all conference bridges in this pool. By default an alarm will be raised when 80% of resources are in use. For more information see Setting the threshold for raising conference bridge resource usage alarms [p.72]
Location	Contains a list of all Locations of types *Rendezvous* and *Both* that have been configured on the TelePresence Conductor. A conference bridge pool needs to reference a Location so that out-dial calls can be sent via the appropriate SIP trunk. Such calls could be initiated by • auto-dialed participants being configured on TelePresence Conductor. • Cisco TMS scheduling a conference with participants. • a user of conference control center in Cisco TMS adding a participant to an existing conference. Use *None* if no outbound calls are to be sent via the call control device. Change this field after you have created a Location, if there are no Locations in the list yet. The **default** is *None*.

3. Click **Create pool**.
   A new section **Conference bridges in this pool** will appear.

To add conference bridges to the pool, click **Create conference bridge**.

To save the new conference bridge pool, click **Save**.

Adding and editing conference bridges

A single conference bridge can only belong to one conference bridge pool.

Before adding a conference bridge, ensure that:

- you configure it in accordance with the relevant deployment guide:
  - *Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide*
  - *Cisco TelePresence Conductor with Cisco VCS (Policy Service) Deployment Guide*
  - *Cisco TelePresence Conductor with Cisco VCS (B2BUA) Deployment Guide*
- all conference bridges have the same software version installed (we recommend to install the latest version to ensure that all features are supported).
- all conference bridges in all pools are configured identically. Failure to do so will result in unpredictable behavior.
- all conference bridges in a pool are of the same conference bridge type (either TelePresence Server or TelePresence MCU).
- all conference bridges used by the TelePresence Conductor are reserved for its exclusive use and are not used by any other system, for example Cisco TelePresence MCU Conference Director.

Note: We strongly recommend that all conference bridges within a pool have the same capacity, so that conferences can be distributed efficiently across conference bridges. If there are conference bridges with different capacities in the same pool, this may lead to unbalanced conference placement in some scenarios.

When editing the configuration of conference bridges, be aware that new conferences may not be connected, as the conference bridge will temporarily be unreachable. We therefore recommend that you edit conference bridges at off-peak times.

If you make changes to any of the following options while conferences are in progress on the conference bridge, those conferences will be deleted:

- IP address or FQDN
- Protocol
- Port
- Conference bridge username
- Conference bridge password
- Dial plan prefix
- SIP port

If a conference bridge is not available, the TelePresence Conductor will wait for a set period of time before attempting to re-contact it. This period is configurable on the **Global conference bridge settings** page (**Conference configuration > Global conference bridge settings**).

To add a new conference bridge to a pool:

1. Go to **Conference configuration > Conference bridge pools**, then click on the name of the pool to which you wish to add a conference bridge.
   You will see a list of conference bridges (if any) currently belonging to the pool.
2. Click **Create conference bridge**.

When adding or editing a conference bridge, the configurable options are:

Field	Description
Name	Descriptive name of the conference bridge. For ease of reference, when using a Cisco VCS in your deployment, we recommend that you use the same name for the conference bridge both here and when adding each conference bridge as a neighbor zone.
Description	A free-form description of the conference bridge.
State	Determines whether the TelePresence Conductor will treat this conference bridge as available for use. When using the TelePresence Conductor without a valid release key it is only possible to enable one conference bridge across all conference bridge pools. This conference bridge cannot be clustered. If a TelePresence Conductor running without a release key finds more than one conference bridge enabled, for example when a release key existed, but is then deleted, the TelePresence Conductor will set all conference bridges to the *Busied out* state and only leave one conference bridge as *Enabled*. *Enabled*: the conference bridge will be used as and when required. *Busied out*: the conference bridge will not be used for any new conferences.
IP address or FQDN	The IP address or Fully Qualified Domain Name (FQDN) of the conference bridge. We strongly recommend that you use IP addresses in this field to ensure best performance.
Protocol	The protocol (either *HTTP* or *HTTPS*) that the TelePresence Conductor will use when communicating with the conference bridge. Because the conference bridge password is transmitted over the network, we recommend using HTTPS in deployments where interception of traffic between TelePresence Conductor and conference bridges could pose an unacceptable security risk. If you use HTTPS you must enable this feature on the conference bridge.
Port	The port on the conference bridge to which the TelePresence Conductor will connect. We recommend using *80* for HTTP and *443* for HTTPS.
Conference bridge username	The **User ID** of the account used by the TelePresence Conductor to log in to the conference bridge. This conference bridge account must have a **Privilege level** of *administrator*. We do not recommend that the TelePresence Conductor uses the conference bridge's default admin user account.
Conference bridge password	The password used to log in to the conference bridge.
Dial plan prefix	(Required only when TelePresence Conductor is deployed using the Cisco VCS's external policy server interface.) The prefix that has been configured as part of a Cisco VCS search rule to route calls to this conference bridge. The prefix must be unique for each conference bridge in the pool. It is alphanumeric. In a deployment using the Cisco VCS's external policy server interface, there must not be any conflict between any **Incoming alias** or **Conference name** (used when Creating and editing conference aliases [p.104]), the Call Policy prefix [p.22], and Conference bridge dial plan prefixes [p.22]. Otherwise you may experience unpredictable behavior. For more information, see Considerations in a Cisco VCS-only deployment [p.22]. For more information on dial plans and prefixes, see Designing a dial plan [p.22]. For more information on Cisco VCS search rule configuration, see *Cisco TelePresence Video Communication Server Administrator Guide*.

Field	Description
Conference bridge type	The type of this conference bridge. This field cannot be modified. This release of the TelePresence Conductor supports *Cisco TelePresence MCU* and *Cisco TelePresence Servers* only. Future versions may support other types of conference bridges.
Conference bridge pool	The pool to which this conference bridge belongs. This field cannot be modified.
Dedicated content ports	(Available when the **Conference bridge pool** has a conference bridge type of *TelePresence MCU*) Determines the number of dedicated content ports on the conference bridge. These content ports will be excluded from the calculation of how many ports to reserve. To see how many content ports your TelePresence MCU model has, see **MCU port matrix** within the Cisco TelePresence MCU online help. The **default** is 0.
SIP port	The port that the conference bridge is listening on for TLS SIP traffic. This is required when the TelePresence Conductor is connected to a Unified CM or Cisco VCS via a SIP trunk. The **default:** is 5061.
H.323 cascade call routing	(Available when the **Conference bridge pool** has a conference bridge type of *TelePresence MCU*) The method used for routing H.323 calls when a conference is cascaded from one conference bridge to another. *Gatekeeper routed*: Calls are routed via an H.323 gatekeeper, for example a Cisco VCS, which the conference bridge must be registered to. *Direct*: Calls are routed directly from one conference bridge to another without using an H.323 Gatekeeper. The conference bridges must be configured in such a way that IP traffic is routable between them. If you select this setting do not configure an H.323 gatekeeper on the conference bridge, because the setting on the conference bridge will take precedence over this setting and the call will therefore be routed via the conference bridge's H.323 gatekeeper. The **default** is *Gatekeeper routed*.

Busying out conference bridges

Busying out a conference bridge will make it temporarily unavailable, preventing the TelePresence Conductor from using it for new conferences. When a conference bridge that is currently in use is busied out, any existing conferences on that conference bridge will be unaffected and new callers will still be able to join the existing conference.

During the time that a conference bridge is in 'Busy out' state, the TelePresence Conductor will still continue to poll it for information about its resources. For this reason we recommend that a conference bridge that is either no longer required by a TelePresence Conductor, or required for use by another system, be completely removed from the pool when all existing conferences have been completed.

To busy out a conference bridge, temporarily preventing the TelePresence Conductor from using it:

1. Go to **Conference configuration > Conference bridges**.
2. Select the conference bridge you wish to busy out.
3. Click **Busy out**.

If a TelePresence Conductor is running without a release key and it finds more than one conference bridge enabled, for example when a release key existed, but is then deleted, the TelePresence Conductor will set all conference bridges to the *Busied out* state and only leave one conference bridge as *Enabled*.

Deleting conference bridges

Note: deleting a conference bridge from the conference bridge pool removes the conference bridge completely. Any conference running on a conference bridge that is deleted will be torn down.

To permanently delete a conference bridge from the pool:

1. Go to **Conference configuration > Conference bridges**.
2. Select the conference bridge you wish to delete.
3. Click **Busy out**.
4. Wait until all conferences on that conference bridge have finished (to check, go to **Status > Conference bridges**).
5. Go back to the **All conference bridges** page, select the conference bridge you wish to delete and click **Delete**.

Viewing all conference bridges across all pools

To view a list of all conference bridges currently being used by the TelePresence Conductor, go to **Conference configuration > Conference bridges**.

From this page you can click on any of the column headings to sort the list by, for example, **Conference bridge pool**, **Address**, or **Dial plan prefix**.

To edit details of any of the conference bridges click on the conference bridge name or on **View/Edit**.

Moving a conference bridge between pools

1. Go to **Conference configuration > Conference bridges** and click on the **Name** of the conference bridge you want to move.
2. From the **Conference bridge pool** drop-down list, select the pool to which the conference bridge is to be moved.
   The conference bridge types of the new pool and the conference bridge being moved have to match.
3. Select **Save**.

Adding or editing quality settings

(This feature is relevant only if using a conference bridge type of TelePresence Server.)

Quality settings consist of an audio quality level and a video quality level. You assign them to conference templates, auto-dialed participants and pre-configured endpoint codecs in order to allow for the required resources to be allocated on the corresponding TelePresence Server.

See Resource reservation and allocation on the TelePresence Server [p.98] for more information on how resources are allocated on a TelePresence Server.

The TelePresence Conductor has been pre-configured with the following quality settings, which can be deleted or edited:

- Full HD (1080p 30fps / 720p 60fps video, multi-channel audio)
- HD (720p 30fps video, stereo audio)
- SD (wide 448p / 480p 30fps video, mono audio)
- 360p (360p 30fps video, mono audio)
- Audio-only (no video, mono audio)

To add a new custom quality setting:

1. Go to **Conference configuration > Quality settings**.
2. Click **New**.

When adding or editing a quality setting, the configurable options are:

Field	Description
Description	A free-form description of this quality setting. The descriptions will appear as options in the following drop-down lists: ■ **Participant quality**, **Host quality** and **Guest quality** on the **Conference templates** page ■ **Maximum quality** on the **Auto-dialed participants** page ■ **Quality setting** on the **Codecs** page (accessible when adding codecs to pre-configured endpoints)
Audio quality level	The level of audio quality associated with this quality setting. The options are: ■ Multi-channel ■ Stereo ■ Mono The **default** is *Multi-channel*.

Field	Description
Video quality level	The level of video quality associated with this quality setting. The options are: ■ Full HD (1080p 30fps / 720p 60fps) ■ HD (720p 30fps) ■ SD (wide 448p / 480p 30fps) ■ 360p (360p 30fps) ■ None 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. The **default** is *Full HD (1080p 30fps / 720p 60fps)*.

Changing global conference bridge settings

Some settings on the TelePresence Conductor apply to all conference bridges in its pools. The **Global conference bridge settings** page (**Conference configuration > Global conference bridge settings**) allows you to perform the following tasks:

- Changing the conference bridge retry interval
- Setting the threshold for raising conference bridge resource usage alarms

Changing the conference bridge retry interval

The TelePresence Conductor constantly monitors its pool of conference bridges to check whether they are available. If a conference bridge is not contactable, the **Conference bridge retry interval** setting determines the number of seconds the TelePresence Conductor will wait before attempting to re-contact a conference bridge that was previously unavailable.

You can change this setting on the **Global conference bridge settings** page (**Conference configuration > Global conference bridge settings**).

Note: If you set this interval too high, it will take a long time for the TelePresence Conductor to start using a conference bridge it has previously experienced a problem with. If you set this interval too low, and there is a conference bridge with a persistent fault, the TelePresence Conductor will waste resource by trying to communicate with it.

The default is 300 seconds (5 minutes). You should not deviate from the default setting unless advised to do so by a Cisco Technical Support Representative. Setting the **Conference bridge retry interval** to *0* results in the TelePresence Conductor attempting to contact the conference bridge as often as possible, approximately once every second.

Setting the threshold for raising conference bridge resource usage alarms

The TelePresence Conductor constantly monitors all conference bridges in its pools to check the total number of resources that are available, and how many are currently in use.

By default, the TelePresence Conductor will raise an alarm when more than 80% of all the currently available conference bridge resources are in use and when 80% of all available resources within a conference bridge pool are in use. The alarm raised is one of the following:

- TelePresence MCU pool exceeds capacity: Unable to create conferences as insufficient TelePresence MCU resource is available in pool
- TelePresence MCUs exceed capacity: Unable to create conferences as insufficient TelePresence MCU resource is available
- TelePresence MCU pool resource warning: TelePresence MCU pool port usage is approaching or has reached full capacity
- TelePresence MCU resource warning: TelePresence MCU port usage is approaching or has reached full capacity
- TelePresence Server pool exceeds device capacity: Unable to create conferences as insufficient TelePresence Server device resource is available in pool
- TelePresence Server pool exceeds license capacity: Unable to create conferences as insufficient

TelePresence Server license resource is available

- TelePresence Servers license capacity exceeded: An individual TelePresence Server's license resource is insufficient to create a conference
- TelePresence Server license resource warning: An individual TelePresence Server's license usage is approaching or has reached full capacity
- TelePresence Server pool license resource warning: TelePresence Server pool license usage is approaching or has reached full capacity
- TelePresence Server device resource warning: An individual TelePresence Server's device usage is approaching or has reached full capacity

There are two situations when the alarm will be raised:

- When a new conference is created, or a new participant joins an existing conference, and this takes the resource usage above the configured threshold. In this situation one of the above alarms will be raised but participants can continue to create and join conferences until there are no more resources available.
- When a conference could not be created because the number of required resources exceeded the number of resources currently available. The number of required resources for a conference is the total of all auto-dialed participants, reserved hosts, cascade resources, and reserved content resources. In this situation, an event will appear in the event log in addition to the above alarm. The event is **Not enough conference bridge resource to handle request**.

To change whether and when this alarm is raised:

1. Go to the relevant page
   - For all conference bridges go to **Conference configuration > Global conference bridge settings**
   - For conference bridges within a pool go to **Conference configuration > Conference bridge pools**
2. Choose one of these options:
   - to change the threshold at which alarms are raised, select the **Raise conference bridge resource alarm** check box, and in the **Threshold (%)** field enter the desired value
   - to only raise alarms when a conference could not be created because the number of required resources exceeded the number of resources currently available, select the **Raise conference bridge resource alarm** check box and in the **Threshold (%)** field enter a value of *100*
   - to never raise alarms, clear the **Raise conference bridge resource alarm** check box

After the alarm has been raised, it can be acknowledged by the system administrator. If the alarm has been acknowledged, it will be raised again if the resource usage still exceeds the configured threshold after 24 hours.

The alarm will not be lowered until the system is restarted. It will not be lowered automatically by the system if the current resource usage drops back below the threshold.

Conference bridge response time

In networks with high latency or packet loss, a warning may be raised on the TelePresence Conductor indicating that the conference bridge has taken a long time to respond, seemingly indicating a fault on the conference bridge. However, there are several possible causes for this, including:

- packet loss
- high network latency
- conference bridge under high load
- slow-running conference bridge (although this is actually the least likely cause)

The warning will be seen in the event log and will have the following data:

Event="An error occurred while communicating externally."

Detail="mcu response took <x> seconds. It should take no longer than 1 second"

where 'x' is the number of seconds it took to respond.

Configuring conferences

This section describes the configuration steps required to configure ad-hoc and rendezvous conferences on the TelePresence Conductor.

Selecting the preferred conference bridges for a conference

For any particular conference, you can determine which conference bridge pools the TelePresence Conductor will attempt to use to host that conference, in order of preference. You do this by creating a Service Preference, and then assigning a Service Preference to a conference template.

A Service Preference is a prioritized list of conference bridge pools. If no conference bridges within the first pool can be used to host a conference (for example, if there are insufficient resources available for the requirements of the conference), the TelePresence Conductor will check whether the second pool in the list can be used, and so on.

The following limitations apply:

- A Service Preference can contain anywhere between 1 and 30 conference bridge pools.
- A single conference bridge pool can be used in any number of Service Preferences.
- All bridge pools within a Service Preference must be of the same conference bridge type.

Service Preferences can be used by

- conference templates that are configured via the TelePresence Conductor's web interface
- Collaboration Meeting Rooms set up via the TelePresence Conductor's Provisioning API, used to provision conferences

Note: Beware of deleting Service Preferences, because they may be in use by conference templates or by Collaboration Meeting Rooms configured via the TelePresence Conductor Provisioning API.

Creating a Service Preference

1. Ensure that you have created all the conference bridge pools that you want to include in the Service Preference.
2. Go to **Conference configuration > Service Preferences** and select **New**.
   You will be taken to the **Service Preference** page.
3. Enter details of the new Service Preference. The configurable options are:

Field	Description
Service Preference name	Descriptive name of the Service Preference. This name will appear in the **Service Preference** drop-down list when assigning the Service Preference to a template.
Description	A free-form description of the Service Preference.
Conference bridge type	The type of conference bridge that can be assigned to this Service Preference. All conference bridges within a pool, and all pools within a Service Preference, must be of the same type and have the same configuration. This release of the TelePresence Conductor supports *Cisco TelePresence MCU* and *Cisco TelePresence Servers* only. Future versions may support other types of conference bridges.

4. Click **Add Service Preference**.

Adding pools to the Service Preference

After you have created a Service Preference you can add pools to it.

1. In the **Pools** section, from the drop-down list select the conference bridge pool that you want to be used first for any conferences that use this Service Preference.
2. Click **Add selected pool**.
   The new Service Preference will be saved, with the selected conference bridge pool as the first pool to be used.
3. To assign additional conference bridge pools to this Service Preference, select another conference bridge pool from the drop-down list and click **Add selected pool**.
4. To change the order of priority of the conference bridge pools you have selected, use the arrows in the **Change order** column.
5. When all conference bridge pools have been added and are in the desired order, click **Save**.

You can add more conference bridge pools to a Service Preference later and update the preference order.

Marking pools to be used for scheduling

On the **Service Preference** page you can mark pools to be used for scheduling by API clients such as Cisco TMS. All pools that are marked to be used for scheduling will be reported to the API client via the TelePresence Conductor's Capacity Management API.

You can mark pools for scheduling within the **Pools** section of the **Service Preference** page. Click the marker under **Pools to use for scheduling** for each pool you want to mark for scheduling from the highest priority pool at the top of the list downwards. It is not possible to skip higher priority pools. When you change the priority order of a pool, the selected pools will adjust.

Note: After an upgrade to XC3.0, all existing pools in all Service Preferences are marked to be used for scheduling.

Beware of possible misconfigurations. To support dedicated-bridge scheduling:

- Include only a single conference bridge in a pool marked for scheduling
- Do not include pools marked for scheduling in other Service Preferences

See Scheduling conferences [p.26] for more information on how scheduled conferences are created on the TelePresence Conductor.

Cascading conferences across conference bridges and conference bridge pools

Cascading a conference results in resources being used on a secondary conference bridge when the primary conference bridge does not have enough resources available for all the participants.

When a conference is **created**, the TelePresence Conductor will check the resources of the preferred conference bridge pool (according to the Service Preference for the template being used) and where possible, use one of the conference bridges in that pool to host the primary conference. If there are insufficient resources available within that pool, the TelePresence Conductor will then check the availability of the conference bridges in each of the other pools within that Service Preference, in order of priority, until it finds a conference bridge on which it can host the conference.

When an existing conference is **cascaded**, regardless of the conference bridge pool being used to host the primary conference, the TelePresence Conductor will first check the resources of the conference bridges in the preferred conference bridge pool and where possible, use one of the conference bridges in that pool to host the cascade. This could mean that a single conference is cascaded between conference bridges in different conference bridge pools.

TelePresence Server cascading requires version 4.0(1.57) or later.

Cascading in an external policy deployment with TelePresence MCU

In an external policy deployment using TelePresence MCUs as conference bridges and Cisco VCS as the call control device, the Cisco VCS acts as an H.323 gatekeeper, via which all calls for conferences that are cascaded from one conference bridge to another are routed.

To configure this for TelePresence MCU:

1. Go to **Conference configuration > Conference bridges**
2. Select a conference bridge of type TelePresence MCU or click **New**
3. Set **H.323 cascade call routing** to *Gatekeeper routed*

Cascading in a B2BUA deployment or when using a TelePresence Server

In a Unified CM-based deployment, in a Cisco VCS-based deployment using the TelePresence Conductor's B2BUA, or in any deployment using TelePresence Servers as the conference bridges, calls for conferences that are cascaded from one conference bridge to another are routed directly.

No configuration is required for TelePresence Servers. To configure this for TelePresence MCUs:

1. Go to **Conference configuration > Conference bridges**
2. Select a conference bridge of type TelePresence MCU or click **New**
3. Set **H.323 cascade call routing** to *Direct*

Limitations of cascading

Cascading is not supported in ad hoc conferences, since the overhead would be too large for these typically small conferences.

Only single screen endpoints are supported on cascade links connecting TelePresence Servers. Therefore, if a multiscreen endpoint joins a conference on a cascade conference bridge, participants on the same cascade bridge will see all screens, whereas participants on the primary bridge and on other cascade bridges will only see one screen (the screen showing the loudest speaker).

Cascade links connecting TelePresence Servers support up to 720p/30fps video. Participants viewing video over a cascade link (that is, video from a participant hosted on a different conference bridge) will see a maximum video quality of 720p/30fps.

Participants on the same conference bridge will see full high quality video if all of the following apply:

- Higher quality video (1080p/30fps or 720p/60fps) has been configured on the TelePresence Conductor's conference template.
- The endpoint of the main displayed participant is providing that high quality video.
- The participants' own endpoint supports high quality video.

Creating and editing conference templates

Conference templates define the settings to be applied to different conferences when they are created. The same template can be used by more than one conference alias.

Note: Conference templates are configured via the web interface and are separate from Collaboration Meeting Rooms which are provisioned via the TelePresence Conductor's Provisioning API.

The **Conference templates** page (**Conference configuration > Conference templates**) lists all the existing conference templates and allows you to edit, delete and create new templates.

When creating or editing a conference template, the configurable options are:

Field	Description
Name	Descriptive name of the conference template.
Description	A free-form description of the conference template.
Conference type	Determines the nature of the conference that will be created when this template is used. *Meeting*: the conference will have one type of participant, and all participants will be given the same priority. *Lecture*: there will be two different types of participants with different levels of priority. Each participant type will use a different alias to dial in to the conference. The **default** is *Meeting*.
Number of hosts to reserve	(Available when **Conference type** is *Lecture*) The number of hosts to reserve resources for on the conference bridge. If using TelePresence MCUs, one port per host will be reserved on the primary TelePresence MCU. See About resource allocation [p.97] for more information.
Call Policy mode	(Applicable only to deployments using the Cisco VCS's external policy server interface) Determines whether you want to check whether users who have dialed a conference alias that uses this template have the right to create a conference. *Off*: no checks will be made. *On*: the TelePresence Conductor will check the Cisco VCS's Call Policy before allowing users to create a conference. If set to *On*, you must also configure the TelePresence Conductor's Call Policy prefix. See Using Call Policy [p.123] for more information. This should be set to *Off* in a Unified CM-based deployment or a Cisco VCS-based deployment using the TelePresence Conductor's B2BUA.
Service Preference	Determines the Service Preference that this template will use. A Service Preference is a prioritized list of conference bridge pools that the TelePresence Conductor will use for this conference. You must create your Service Preferences before you can create a template. For more information see Selecting the preferred conference bridges for a conference [p.76].

Field	Description
Maximum number of cascades	The maximum number of cascades that are allowed for this conference. This number affects the resources that are reserved on the primary conference bridge for connections to additional conference bridges. The resources will be used if the conference exceeds the capacity of the primary conference bridge and is cascaded to one or more conference bridges. On a TelePresence MCU, each connection between the primary conference bridge and an additional conference bridge requires one port. On a TelePresence Server, each connection between the primary conference bridge and an additional conference bridge requires the resources that would be used by a participant receiving 720p video, stereo audio and the **Content quality** that is selected on the conference template. If you want to prevent a conference from cascading across multiple conference bridges, you can set the **Maximum number of cascades** to 0. If you do so, be aware that this may prevent new participants from being able to join a conference. The **default** is '0'. For more information about cascading across conference bridges, see Cascading conferences across conference bridges and conference bridge pools [p.78] For more information about reserving cascade resources on a TelePresence MCU, see Reserving cascade resources [p.97]. For more information about reserving cascade resources on a TelePresence Server, see Reserving cascade resources [p.99].
Limit number of participants	Determines whether there will be a limit set on the total number of participants permitted in this conference. This number includes all auto-dialed participants (participants who are dialed in to the conference by the conference bridge) and all other participants (participants who dial in to the conference, including those who have had host resources reserved). The maximum number of participants must be more than the total number of auto-dialed participants plus the number of reserved hosts. **Note:** No preference is given to participants who have organized a conference. If the maximum number of participants is reached before the participant who organized the conference has dialed in, this participant is rejected. For more information see Limiting the number of participants in a conference [p.103].
Limit the conference duration (minutes)	Determines whether there will be a limit set on the maximum duration of conferences created using this template. When selected, specify the limit of the conference duration in minutes. On a TelePresence Server there will be a warning message displayed as overlaid text on the screen informing you that the conference is about to end. On a TelePresence MCU, depending on the configuration, there will be warnings issued - as an audio notification and/or as overlaid text - at varying intervals before the conference is about to end. See What warnings do I get on a Cisco TelePresence MCU that my conference is finishing? for information on how to turn the warnings off, and on the intervals at which the warnings will be displayed.

Field	Description
Participant quality	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server* and the **Conference type** is *Meeting*) The maximum quality setting to apply to participants using this conference template. Depending on the selected setting the required resources are allocated on the TelePresence Server that is associated with this conference template. The list of available quality settings can be changed on the **Quality settings** page. The pre-configured quality settings are: ■ Full HD (1080p 30fps / 720p 60fps video, multi-channel audio) ■ HD (720p 30fps video, stereo audio) ■ SD (wide 448p / 480p 30fps video, mono audio) ■ 360p (360p 30fps video, mono audio) ■ Audio-only (no video, mono audio) 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. When using a CTS3000 or TX9000 you must select *Full HD (1080p 30fps / 720p 60fps video, multi-channel audio)* or a custom quality setting that has an audio quality level of multi-channel, otherwise insufficient resources will be allocated to display multiple screens. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. The **default** is *HD (720p 30fps video, stereo audio)*.
Host quality	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server* and the **Conference type** is *Lecture*) The maximum quality setting to apply to hosts using this conference template. Depending on the selected setting the required resources are allocated on the TelePresence Server that is associated with this conference template. The list of available quality settings can be changed on the **Quality settings** page. The pre-configured quality settings are: ■ Full HD (1080p 30fps / 720p 60fps video, multi-channel audio) ■ HD (720p 30fps video, stereo audio) ■ SD (wide 448p / 480p 30fps video, mono audio) ■ 360p (360p 30fps video, mono audio) ■ Audio-only (no video, mono audio) 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. When using a CTS3000 or TX9000 you must select *Full HD (1080p 30fps / 720p 60fps video, multi-channel audio)* or a custom quality setting that has an audio quality level of multi-channel, otherwise insufficient resources will be allocated to display multiple screens. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. The **default** is *HD (720p 30fps video, stereo audio)*.

Field	Description
Guest quality	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server* and the **Conference type** is *Lecture*) The maximum quality setting to apply to guests using this conference template. Depending on the selected setting the required resources are allocated on the TelePresence Server that is associated with this conference template. The list of available quality settings can be changed on the **Quality settings** page. The pre-configured quality settings are: ■ Full HD (1080p 30fps / 720p 60fps video, multi-channel audio) ■ HD (720p 30fps video, stereo audio) ■ SD (wide 448p / 480p 30fps video, mono audio) ■ 360p (360p 30fps video, mono audio) ■ Audio-only (no video, mono audio) 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. When using a CTS3000 or TX9000 you must select *Full HD (1080p 30fps / 720p 60fps video, multi-channel audio)* or a custom quality setting that has an audio quality level of multi-channel, otherwise insufficient resources will be allocated to display multiple screens. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. The **default** is *HD (720p 30fps video, stereo audio)*.
Allow multiscreen	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server*) Whether or not the conference allows for multiscreen systems. *Yes*: The conference allows for resources to be made available for systems with more than one screen. This is also required if pre-configured endpoints should be checked. *No*: The conference only allows for single-screen systems or the primary screen of multiscreen systems. Pre-configured endpoints are not checked. **Note:** if a multiscreen endpoint joins a conference on a cascade conference bridge, participants on the same cascade bridge will see all screens, whereas participants on the primary bridge and on other cascade bridges will only see one screen (the screen showing the loudest speaker). The **default** is *No*.

Field	Description
Maximum screens	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server* and when **Allow multiscreen** is set to *Yes*) For TIP-compliant endpoints dialing into Rendezvous conferences using the TelePresence Conductor's B2BUA, this field specifies the maximum number of screens for which resources are allocated on the conference bridge. The TelePresence Conductor takes the lesser of the **Maximum screens** value and the number of screens specified by the TIP endpoint and allocates resources accordingly. For pre-configured endpoints this setting is ignored and the number of screens defined for the pre-configured endpoint are allocated. For endpoints that are neither TIP-compliant nor pre-configured, this setting is ignored and only a single screen is allocated, unless the endpoint is: ■ escalated into an ad hoc conference on the TelePresence Conductor ■ reserved as a host in a Lecture-type conference ■ using the Cisco VCS's external policy server interface to call into a rendezvous conference If the endpoint falls into one of the categories listed above, the **Maximum screens** defines the number of screens for which resources are initially allocated on the conference bridge. The **default** is *1*. For more information see Allocating resources for multiple screens [p.100].

Field	Description
Optimize resources	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server*) Whether resources are optimized for conferences that use this template. *Yes*: Resources that were initially allocated on the conference bridge and that the endpoints do not support, are freed up if five seconds pass and no new participant joins the conference. Resources are optimized and freed up on the TelePresence Server in the following situations: ■ If the maximum capability an endpoint advertises is a lower quality than defined in the conference template for one of the following settings: • **Participant quality** • **Guest quality** • **Host quality** ■ If an endpoint that uses the Cisco VCS's external policy server interface to dial into a rendezvous conference supports fewer screens than defined in the template under **Maximum screens**. (Endpoints in B2BUA deployments have resources for the correct number of screens allocated, because the TelePresence Conductor can detect the number of screens required from the SIP signaling.) ■ If an endpoint that is escalated into an ad hoc conference supports fewer screens than defined in the template under **Maximum screens**. Resources are optimized on the TelePresence Conductor, but not freed up on the TelePresence Server, if an endpoint that has been reserved as a host in a Lecture-type conference supports fewer screens than defined in the template under **Maximum screens**. The freed up resources can only be used for other hosts dialing into the same conference. Resources are not optimized for auto-dialed participants or for pre-configured endpoints, because when configuring these entities the desired quality is defined in the configuration, and overrides the capabilities defined in the conference template for incoming calls. *No*: The number of resources consumed by this conference is the same as the number of resources initially allocated on the conference bridge. Some of the resources may be wasted if there are endpoints that support fewer screens than defined or are supporting a lower quality level than defined. The **default** is *Yes*. For more information see Optimizing resources [p.101].
Allow content	(Available when the **Service Preference** has a conference bridge type of *TelePresence MCU)* Whether or not participants will be able to send content video, such as a presentation. *Yes*: a single port will be reserved on the primary TelePresence MCU and each cascade TelePresence MCU specifically for content. Use this setting for WebEx-enabled conferences. *No*: participants will not be able to send content, regardless of the number of ports available on the MCU. Content may still be displayed, since some endpoints provide content in their main video channel. The **default** is *Yes*. See Reserving a content port [p.98] for more information.

Field	Description
Content quality	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server*) The video quality level for content, such as presentations, associated with conferences based on this template. Depending on the quality level selected, the appropriate number of resources will be allocated on the conference bridge hosting the conference. The options are: ■ Full HD (1080p 30fps / 720p 60fps) - this setting supports wide UXGA 27fps ■ HD (720p 30fps) ■ 1280 x 720p 15fps ■ 1280 x 720p 5fps ■ Off Do not use *Off* for WebEx-enabled conferences. If **Maximum number of cascades** is greater than 0, resources for content are reserved on the primary conference bridge for each possible cascade link. Each cascade conference bridge uses content resources for the cascade link to the primary conference bridge as well as for all participants that dial in. For Cisco TelePresence System T3 (T3) and TelePresence Interoperability Protocol (TIP)-compliant endpoints to be treated as a multiscreen endpoint the associated conference template must have a **Content quality** of at least *1280 x 720p 5fps*. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. If *Off* has been selected, there will not be a separate content channel available for conferences based on this template. Content may still be displayed, since some endpoints provide content in their main video channel. The **default** is *Off*. For more information see Allocating resources for content [p.99].
Scheduled conference	Whether or not the conference may only be created by a scheduling application such as Cisco TMS using the API. *Yes*: the conference will be created at a pre-determined time by the scheduling application. Any participants dialing in to the conference before this time will be rejected and will have to call back after the conference has been created. Use this setting for WebEx-enabled conferences. *No*: the conference will be created as soon as the first participant dials in. The **default** is *No*.
Segment switching	(Available when the **Service Preference** has a conference bridge type of *TelePresence Server*) Whether or not the TelePresence Server associated with this conference template will have the 'Segment switching' feature enabled. Segment switching is supported on TelePresence Server version 4.0 or later. The setting is ignored when using an earlier software version. *Yes*: the default segment switching mode is enabled on the associated TelePresence Servers. Multiscreen endpoints can show just the screen containing the loudest speaker of another multiscreen system instead of all screens. This can lead to a mixture of single-screen endpoints and individual screens of a multiscreen system being displayed at the same time. Segment switching only works for multiscreen systems that provide loudest pane information. *No:* the room switching mode is enabled on the associated TelePresence Servers. If a multiscreen endpoint is the loudest speaker, all of its screens are displayed full-screen on other multiscreen endpoints (if they have enough screens). If the multiscreen endpoint is not the loudest speaker, none of its screens are displayed full-screen on the other multiscreen endpoints. The **default** is *Yes*.

Field	Description
Advanced parameters	**Advanced template parameters** are parameters that can be passed to the conference bridge via its API. The parameters can be edited after a conference template has been created. Click **Edit** to get to the **Advanced template parameters** page, where you can select the parameters you would like to send to the conference bridge. Advanced parameters that are not selected or specified will not be sent to the conference bridge and the conference bridge's default values will be used. See Adding and editing advanced template parameters [p.87] for more information. **CAUTION**: This feature is for advanced use only.

Note: If there is one or more auto-dialed participant defined for this conference template, the **Maximum quality** for the auto-dialed participant overrides the quality settings defined on the conference template (**Participant quality**, **Host quality** or **Guest quality**). This may result in the auto-dialed participant experiencing a different quality compared with the rest of the participants in the conference.

Adding and editing advanced template parameters

CAUTION: This feature is for advanced use only.

Cisco TelePresence MCUs support the *Cisco TelePresence MCU Remote Management API* and Cisco TelePresence Servers support the *Cisco TelePresence Server API*. These APIs enable third-party control of the relevant conference bridge via messages sent using the XML-RPC protocol. The TelePresence Conductor uses these APIs to manage conferences on the conference bridges in its pool. It supports the TelePresence MCU's API versions 2.8 or later and the TelePresence Server's API version 3.0 or later.

The TelePresence Conductor allows you to make use of the calls `conference.create` (for TelePresence MCU) and `flex.conference.create` (for TelePresence Server) of these APIs through the **Advanced template parameters** page. This page is accessible via the **Conference templates** page (**Conference configuration > Conference templates**) once a conference template has been created.

When a conference is created, a conference bridge will apply settings that have been configured on the conference bridge. However, these settings will be overridden by any values configured in the **Advanced template parameters** on the TelePresence Conductor.

To create or edit advanced template parameter settings:

1. Create a new conference template or select an existing conference template (**Conference configuration > Conference templates**)
2. In the **Advanced parameters** section click **Edit**.
   The **Advanced template parameters** page opens.
3. Select the check-boxes next to all the parameters that should be sent to the conference bridge.
   - For TelePresence Server there is only one check-box per setting. Selected settings are applied to the primary and to any cascade conference bridges.
   - For TelePresence MCU there are two check-boxes: the first one for the primary TelePresence MCU and the second one for any cascade TelePresence MCUs. Ensure that you supply the matching cascade value for all primary advanced parameters that you have selected. Also ensure that the values for primary and cascade advanced parameters are identical.
     Note: in a future release of the TelePresence Conductor the cascade parameters may be removed and the primary advanced parameters will be applied to any cascade conference bridges.
4. Enter or select the relevant parameter values.

Cisco TelePresence MCU parameters

Note: all TelePresence MCUs should be running the same release of software in order to guarantee consistent availability and behavior of the various parameters.

Field name	Parameter in API	Description
Automatic lecture mode	`automaticLectureMode`	Automatic lecture mode allows the lecturer (host) to be shown in full-screen view to the students. In this mode, the lecturer will continue to see their normal (continuous presence) view. That is, the lecturer will see the students (guests) and not himself. The TelePresence MCU identifies the lecturer and controls the layout seen by the other participant according to which mode is selected here. *Type 1*: The speaker sees continuous presence (or their custom layout) and all participants see the guest who is speaking (be they a host or a guest). *Type 2*: All guests, including the speaker, see the last host who spoke full screen. All hosts will see their custom layout. *Disabled*: Automatic lecture mode is disabled."
Timeout for automatic lecture mode type 1	`automaticLectureModeTimeout`	This parameter is applicable if an **Automatic lecture mode** of *Type 1* is selected. The value (in seconds) determines how quickly the loudest speaker will appear in full-screen view to the other participants.
Floor and chair control	`chairControl`	Controls floor and chair control settings for this conference. *None*: the use of floor and chair controls is not allowed in this conference. *Floor control only*: only floor control is allowed in this conference; chair control is not allowed. Any participant can 'take the floor' so long as no other participant has currently 'taken the floor'. *Chair and floor control*: both chair and floor control are allowed in this conference. Any participant can 'take the floor' and any host can 'take the chair' so long as no other participant has currently done so. *Default*: use the default setting on the TelePresence MCU. See the relevant TelePresence MCU's Online Help for more information on **Floor and chair control**.

Field name	Parameter in API	Description
Content mode	`contentMode`	Defines the content mode of the conference. Do not use when running TelePresence MCU version 4.2. *Transcoded*: content is always transcoded. The TelePresence MCU sends out a single, transcoded content stream. *Pass-through*: content is not decoded and is simply repackaged and sent out to each eligible endpoint in the conference. *Hybrid*: The TelePresence MCU sends out two content streams: a passed-through higher resolution stream, and a lower resolution stream transcoded and scaled down for any endpoints that are unable to support the higher stream.
Transmitted content resolutions	`contentTransmitResolutions`	The resolution for the content channel that will be transmitted to endpoints in conferences based on this template. *4-to-3 only*: the TelePresence MCU encodes the content and transmits it in a resolution of ratio 4:3. *16-to-9 only*: the TelePresence MCU encodes the content and transmits it in a resolution of ratio 16:9. *Allow all*: the TelePresence MCU decides on the most optimal resolution depending on information about capabilities sent by the endpoints in the conference.
Outgoing transcoded codec	`contentTxCodec`	The codec used to transmit content in conferences based on this template. If content is to be transcoded, this is the output format of the transcoder: *H.263+* or *H.264*. This setting does not apply in *Pass-through* mode.
Minimum bit rate to use for transmitted content	`contentTxMinimumBitRate`	Sets a lower limit on the bandwidth of the shared content video encoding sent to content receivers in a conference. Measured in bps.
Custom layout enabled	`customLayoutEnabled`	Whether a custom layout can be used for all participants in conferences based on this template. To use custom layouts, the field **New participants will see custom layout** must be set to *True* and **Custom layout** must be set to the appropriate value.
Custom layout	`customLayout`	The index number of the video layout seen by the conference participants. See Conference layouts [p.214] for a list of available layouts and corresponding index values. To use custom layouts, the fields **New participants will see custom layout** and **Custom layout enabled** must both be set to *True*. The **default** is *5*.
Description	`description`	Additional information about the conference.

Field name	Parameter in API	Description
Encryption required	`encryptionRequired`	The encryption setting for conferences based on this template. If *True*, encryption is required for these conferences; otherwise, encryption is optional. If encryption is required, the TelePresence MCU must have the encryption feature key enabled.
Guest PIN	`guestPin`	If a conference has a **Guest PIN** set, guest participants cannot join the conference or change its configuration without entering the correct PIN. The **Guest PIN** can only be set on the primary conference bridge. If set, the Guest PIN will be identical on any cascade conference bridges.
Audio muted initially	`joinAudioMuted`	Whether to initially mute audio from all participants when they join the conference.
Video muted initially	`joinVideoMuted`	Whether to initially turn video off from all participants when they join the conference.
Disconnect when last host leaves	`lastChairmanLeavesDisconnect`	Whether all other participants will be disconnected when the last participant with host status leaves the conference. **Note:** If **Disconnect when last host leaves** is set to *true* on your conference template and the conference bridge type is TelePresence MCU, beware of the following issue: If all host participants on the primary TelePresence MCU disconnect, any guest participants on the primary TelePresence MCU will be disconnected automatically. This occurs even if there are still hosts remaining on a cascade TelePresence MCU.
Layout control via FECC/DTMF	`layoutControlEx`	Defines how the view layout can be controlled. The setting can be overridden in an endpoint's individual configuration. *Disabled*: participants are not allowed to change their view layout using either far end camera control (FECC) or dual-tone multi frequency (DTMF). *FECC only*: participants can only change their view layout using FECC. *DTMF only*: participants can only change their view layout using DTMF. *FECC with DTMF fallback*: participants can change their view layout using FECC when it is available and via DTMF on endpoints, which do not have FECC. *Both FECC and DTMF*: participants can change their view layout using both FECC and DTMF.
New participants will see custom layout	`newParticipantsCustomLayout`	Whether new participants use the custom layout or not. To use custom layouts, the field **Custom layout enabled** must be set to *True* and **Custom layout** must be set to the appropriate value.

Field name	Parameter in API	Description
PIN	`pin`	If a conference has a **PIN** set, hosts cannot join the conference or change its configuration without entering the correct PIN. The **PIN** can only be set on the primary conference bridge. If set the PIN will be identical on any cascade conference bridges.
Mute in-band DTMF	`suppressDtmfEx`	Controls the muting of in-band dual-tone multi-frequency (DTMF) tones. *FECC*: in-band DTMF tones will be muted when DTMF is being used to control layout because far end camera control (FECC) is not available. *Always*: in-band DTMF tones will always be muted. *Never*: in-band DTMF tones will never be muted.
Template number	`templateNumber`	An index that uniquely identifies the TelePresence MCU template.
Template name	`templateName`	The name of the TelePresence MCU template.

Field name	Parameter in API	Description
Custom parameters		This field can be used to enter advanced parameters and their corresponding values in valid JSON. Do not use the following parameters. These are used by the TelePresence Conductor and changing them will result in a failure to create conferences: ■ **conferenceName** ■ **numericId** ■ **guestNumericId** ■ **startTime** ■ **maximumAudioPorts** ■ **reservedAudioPorts** ■ **maximumVideoPorts** ■ **reservedVideoPorts** ■ **enforceMaximumAudioPorts** ■ **enforceMaximumVideoPorts** ■ **repetition** ■ **weekday** ■ **whichWeek** ■ **weekDays** ■ **terminationType** ■ **terminationDate** ■ **numberOfRepeats** We also advise that you do not use the following parameters, which may also result in a failure to create conferences: ■ **cleanupTimeout** ■ **contentMode** (do not use when running TelePresence MCU version 4.2) ■ **contentContribution** ■ **h239Enabled** ■ **durationSeconds** ■ **private** The TelePresence Conductor does not perform in-depth checking of data in these fields. Any incorrect configuration of the settings may result in your conference failing to be created (or in a cascade failing to be created) and perhaps in a TelePresence MCU becoming temporarily unusable and excluded from the pool of available conference bridges.

For full information on using the MCU API, including the parameters that can be set using the **conference.create** call, see *Cisco TelePresence MCU Remote Management API*.

Cisco TelePresence Server parameters

Field name	Parameter in API	Description
Guest PIN	`guestPin`	If a conference has a Guest PIN set, guest participants cannot join the conference or change its configuration without entering the correct Guest PIN.
PIN	`pin`	If a conference has a PIN set, hosts cannot join the conference or change its configuration without entering the correct PIN.
Single-screen layout	`displayDefaultLayoutSingleScreen`	The default layout type on single-screen endpoints using this conference template. This setting can be overridden by a participant using far end camera control (FECC) or dual-tone multi-frequency (DTMF) keys when in the conference. *Single*: the active speaker is shown in one full-screen pane. *ActivePresence*: the active speaker is shown in a large pane with additional participants appearing in up to nine PIPs (picture-in-pictures) overlaid at the bottom of the screen. *Prominent*: the active speaker is shown in a large pane with additional participants appearing in up to four smaller panes at the bottom of the screen. *Equal*: conference participants are shown in a grid pattern of equal sized panes, up to 4x4.
Multiscreen layout	`displayDefaultLayoutMultiScreen`	The default layout type on multiscreen endpoints using this conference template. This setting can be overridden by a participant using far end camera control (FECC) or dual-tone multi-frequency (DTMF) keys when in the conference. *Single*: all screens of the endpoint with the active speaker are shown full-screen on the multiscreen endpoint. *ActivePresence*: all screens of the endpoint with the active speaker are shown full-screen on the multiscreen endpoint, additional participants appear in up to nine PIPs (picture-in-pictures) overlaid at the bottom of the screen.
Enable iX protocol	`iXEnabled`	If *True*, this field enables the iX protocol on the TelePresence Server associated with this conference. This allows the TelePresence Server to negotiate ActiveControl with endpoints that have this feature enabled.

Field name	Parameter in API	Description
Custom parameters		This field can be used to enter advanced parameters and their corresponding values in valid JSON. Do not use the following parameters. These are used by the TelePresence Conductor and changing them will result in a failure to create conferences: ■ `conferenceName` ■ `conferenceReference` ■ `startTime` ■ `metadata` We also advise that you do not use the following parameters, which may also result in a failure to create conferences: ■ `conferenceMediaTokens` ■ `conferenceMediaTokensUnlimited` ■ `conferenceMediaCredits` ■ `conferenceMediaCreditsUnlimited` ■ `waitForChair` ■ `duration` ■ `durationUnlimited` ■ `maxParticipants` ■ `maxParticipantsUnlimited` The TelePresence Conductor does not perform in-depth checking of data in these fields. Any incorrect configuration of the settings may result in your conference failing to be created and perhaps in a TelePresence Server becoming temporarily unusable and excluded from the pool of available conference bridges. For examples on how to configure common custom parameters see Example TelePresence Server custom parameters [p.94].

For more information on the parameters that can be configured on a TelePresence Server, see *Cisco TelePresence Server API Reference Guide*.

Example TelePresence Server custom parameters

Configuring the TelePresence Server's optimization profile

An optimization profile specifies how the TelePresence Server allocates media resources on endpoints in a particular conference. The options are listed in the table below.

Table 2: Optimization profiles enumerated type

optimizationProfile value	**Description**
maximizeEfficiency	Screen licenses are conserved aggressively. This value gives the most calls for the available resources.
favorEfficiency	This is a balance of efficiency and experience that favors conserving screen licenses over attempting to grant the requested resolution.
favorExperience	Default. This is a balance of efficiency and experience that favors granting the requested resolution over conserving screen licenses.
maximizeExperience	Screen licenses are more readily allocated. This value gives the best experience of the four profiles. If you disable the optimization by bandwidth (by setting **optimizationProfile** to **capabilitySetOnly**), calls will be capable of higher resolutions at lower bandwidths but the inefficiency in allocation could well outweigh the benefit.
capabilitySetOnly	This is the behavior of TelePresence Server 3.1. The TelePresence Server only considers the endpoint's maximum advertized resolution when reporting its screen license requirement to the managing system; it does not attempt to report resources based on the endpoint's advertized receive bandwidth.

To set the optimization profile you must modify the TelePresence Server's **optimizationProfile** API parameter. To do this enter the following JSON command into the **Custom parameters** field, specifying the appropriate option, for example:

```
{"optimizationProfile":"favorExperience"}
```

Configuring other common custom parameters

The following is an example of the JSON to enter into the **Custom parameters** field for other common TelePresence Server parameters:

```
{
        "disconnectOnChairExit":false,
        "welcomeScreen":true,
        "welcomeScreenMessage":"Welcome to your meeting. ミーティングへようこそ。",
        "useCustomOnlyVideoParticipantMessage":true,
        "customOnlyVideoParticipantMessage":"This screen will remain if you are the only
participant",
        "useCustomWaitingForChairMessage":true,
        "customWaitingForChairMessage":"Waiting for conference chairperson.",
        "useCustomPINEntryMessage":true,
        "customPINEntryMessage":"Please enter the security PIN followed by #.",
        "useCustomPINIncorrectMessage":true,
        "customPINIncorrectMessage":"PIN Incorrect - Please try again.",
        "useCustomConferenceEndingMessage":true,
```

```
        "customConferenceEndingMessage":"The conference is ending.",
        "callAttributes":
                {"displayShowEndpointNames":true}
}
```

The TelePresence Server parameters used in this example are:

- **disconnectOnChairExit** - boolean parameter that indicates whether callers are disconnected when the last host leaves
- **welcomeScreen** - boolean parameter that indicates whether to display a welcome screen for 5 seconds when a caller joins a conference
- **welcomeScreenMessage** - parameter that contains a string of up to 500 characters for the welcome message
- **useCustomOnlyVideoParticipantMessage** - boolean parameter that indicates whether to display a custom message when a participant is the only (active) video participant
- **customOnlyVideoParticipantMessage** - parameter that contains a string of up to 500 characters for the custom message displayed to the only video participant in a conference
- **useCustomWaitingForChairMessage** - boolean parameter that indicates whether to display a custom message when waiting for the host to join the conference
- **customWaitingForChairMessage** - parameter that contains a string of up to 500 characters for the custom message displayed to participants waiting for the host to join
- **useCustomPINEntryMessage** - boolean parameter that indicates whether to display a custom message in the PIN entry form
- **customPINEntryMessage** - parameter that contains a string of up to 200 characters for the custom message displayed in the PIN entry form
- **useCustomPINIncorrectMessage** - boolean parameter that indicates whether to display a custom message in the PIN entry form after an incorrect PIN has been entered
- **customPINIncorrectMessage** - parameter that contains a string of up to 100 characters for the custom message displayed after an incorrect PIN has been entered
- **useCustomConferenceEndingMessage** - boolean parameter that indicates whether to display a custom message when the conference is about to end
- **customConferenceEndingMessage** - parameter that contains a string of up to 100 characters for the custom message displayed when a conference is about to end
- **displayShowEndpointNames** - boolean parameter inside the **callAttributes** parameter that indicates whether endpoint names are displayed on the screen or not

For information on other parameters that can be configured see Cisco TelePresence Server API Reference Guide.

About resource allocation

Each conference is hosted on one or more conference bridges. When the TelePresence Conductor receives a request to create a new conference, it checks the resources available on all the conference bridges in the preferred pool to determine which conference bridge should host the conference. Before deciding which conference bridge to use, the TelePresence Conductor must know how many resources that conference will require, so it can assign the conference to a conference bridge that has sufficient resources.

Resource reservation and allocation on the TelePresence MCU

The TelePresence MCU uses the concept of ports to allocate resources. One TelePresence MCU port is allocated for each conference participant, not differentiating between quality levels.

Note: the reservation of different types of ports on the TelePresence Conductor, as described below, is independent of the **Media port reservation** setting on the TelePresence MCUs, which should be set to *Disabled*.

Reserving host resources

A Lecture-type conference can have more than one host. Each host requires conference bridge resources to send and receive audio and video.

Resources for hosts can be reserved on the primary conference bridge, by specifying the **Number of hosts to reserve** on the conference template. Additional hosts can dial into the conference if there are sufficient resources available on the primary or cascade conference bridge(s). If hosts are on a cascade conference bridge, they will not have all the capabilities that they would if they were on the primary conference bridge (for example, layout experience and control functionality may not behave as expected).

We therefore recommend that you enter a **Number of hosts to reserve** that is equal to the number of hosts. This will reserve resource on the primary conference bridge for all hosts, so that they all get the same experience.

Reserved host resources are reserved for the duration of the conference, and are reserved solely for the use of hosts of this conference.

By default guests dialing into a Lecture-type conference before the host will be kept waiting on an entry screen.

Reserving cascade resources

If a conference exceeds the capacity of the primary conference bridge, the TelePresence Conductor will bring in another conference bridge and use its resources as well - this is known as cascading. If the second conference bridge then runs out of resources, the conference can be cascaded from the primary conference bridge to a third conference bridge, and so on.

Each cascade (from the primary conference bridge to another conference bridge) will use one reserved cascade port on the primary conference bridge and one port on the cascade conference bridge. To reserve the required number of cascade ports on the primary conference bridge specify the **Maximum number of cascades** when creating a conference template on the TelePresence Conductor. For the duration of the conference, these ports will be reserved solely for the use of cascades.

It can be difficult to determine the correct number of cascade ports to reserve. If you reserve more cascade ports than are needed, you may unnecessarily reserve resources on the primary conference bridge that cannot be used for participants as a result. Conversely, if you reserve fewer ports than are needed, the conference may not be able to grow to the required size (especially when the network is heavily loaded).

If you want to prevent a conference from cascading across multiple conference bridges, you can set the **Maximum number of cascades** to 0. If you do so, be aware that this may prevent new participants from being able to join a conference.

Reserving a content port

Conference participants may want to send content video such as a presentation. Such content requires a separate port on the conference bridge.

To permit participants to send content, you must select the option to **Allow content**. This reserves a port on the primary conference bridge and, if the conference cascades, a port on each cascade conference bridge specifically for content. Any participant can receive content from these ports, but only one participant at a time can send content.

If you have not selected the option to **Allow content**, participants will **not** be able to send content, regardless of the number of ports available on the conference bridge. Content will only be displayed if the endpoint provides content in its main video channel.

The number of **Dedicated content ports** specified on the conference bridge is excluded from the calculation of how many ports to reserve for content.

Allocating ports

On the TelePresence MCUs one port is allocated for each participant joining a conference. Ports can be reserved for hosts and to allow cascading to additional TelePresence MCUs.

Some TelePresence MCUs have dedicated content and audio ports. When reserving or allocating resources for content the dedicated content ports are used first, and when all have been used, normal video ports are used for content. The TelePresence Conductor does not make use of the dedicated audio ports. It assumes that all participants require video at some point and allocates video ports for both video and audio-only calls.

TelePresence Conductor imposes the following limit on the number of conferences that can be created on a single TelePresence MCU:

```
Total number of conferences per MCU = Number of video ports on the MCU / 2
```

Resource reservation and allocation on the TelePresence Server

The TelePresence Server allocates resources based on the required audio, video and content quality level for a conference. The Quality settings page on TelePresence Conductor allows you to edit the pre-configured quality settings or define new quality settings.

Reserving host resources

A Lecture-type conference can have more than one host. There must be sufficient resources available on one single conference bridge for all hosts to join the conference. To ensure that there are enough resources

available for all hosts to join the conference, when creating a conference template you are asked how many hosts the TelePresence Conductor should reserve resources for on the conference bridge.

The **Number of hosts to reserve** is multiplied by the number of resources that are required for the **Host quality** defined on the same template:

```
Total amount of resources reserved = Number of hosts to reserve * ((Host
quality * number of screens) + Content quality)
```

The appropriate number of resources are reserved for the duration of the conference, solely for the use of hosts.

If the conference requires more host resources than have been reserved, a host may still be able to connect to the conference but only if the conference bridge hosting the conference has sufficient resources available.

Reserving cascade resources

If a conference exceeds the capacity of the primary conference bridge, the TelePresence Conductor will bring in another conference bridge and use its resources as well - this is known as cascading. If the second conference bridge then runs out of resources, the conference can be cascaded from the primary conference bridge to a third conference bridge, and so on.

Each cascade (from the primary conference bridge to another conference bridge) will use the following resources on the primary conference bridges and on the additional conference bridge:

```
[Resources that would be used by a participant receiving 720p video and
stereo audio] + [Resources that are allocated for content (selected under
Content quality on the conference template)]
```

To reserve the required number of cascade resources on the primary conference bridge specify the **Maximum number of cascades** when creating a conference template on the TelePresence Conductor. The total number of cascade resources are the **Maximum number of cascades** multiplied by the number of resources required for one cascade. For the duration of the conference, these resources will be reserved solely for the use of cascades.

It can be difficult to determine the correct maximum number of cascades. If you reserve more cascade resources than are needed, you may unnecessarily reserve resources on the primary conference bridge that cannot be used for participants as a result. Conversely, if you reserve fewer resources than are needed, the conference may not be able to grow to the required size (especially when the network is heavily loaded).

If you want to prevent a conference from cascading across multiple conference bridges, you can set the **Maximum number of cascades** to 0. If you do so, be aware that this may prevent new participants from being able to join a conference.

Allocating resources for participants and guests

Resources for participants in Meeting-type conferences and for guests in Lecture-type conferences are allocated as required when a participant joins a conference.

The number of resources that are allocated for each participant depends on the quality settings for audio and video. These are defined for the conference template using **Participant quality** and **Guest quality**.

Allocating resources for content

Conference participants may want to send content video such as a presentation. Such content requires separate resources to be allocated on the conference bridge.

The number of resources to allocate on the conference bridge is determined by the **Content quality** setting on the conference template.

The table below shows the proportions of resources allocated for each pre-configured content quality setting.

Content quality setting	Proportions of resources allocated
Full HD (1080p 30fps / 720p 60 fps)	1
HD (720p 30fps)	0.5 of the resources for Full HD (1080p 30fps / 720p 60fps)
1280 x 720p 15fps	0.5 of the resources for HD (720p 30fps)
1280 x 720p 5fps	0.33 of the resources for 1280 x 720p 15fps
Off	None

If you have selected *Off*, participants will **not** be able to send content, regardless of the number of resources available on the conference bridge. Content will only be displayed if the endpoint provides content in its main video channel.

Allocating resources for multiple screens

For TelePresence Server to support multiscreen endpoints you must:

- set **Allow for multiscreen** to *Yes* on the conference template, and
- either:
  - use a TelePresence Interoperability Protocol (TIP)-compliant endpoint in a TelePresence Conductor deployment that uses the TelePresence Conductor's B2BUA, or
  - pre-configure the endpoint on the Pre-configured endpoints page.

This release of the TelePresence Conductor does not support auto-dialed participants that are multiscreen endpoints.

Allocating resources for pre-configured multiscreen endpoints

For pre-configured multiscreen endpoints the number of resources that are allocated on the conference bridge is based on the quality settings for all codecs.

For example, if two endpoints, each with three screens/codecs, have been pre-configured, the resources required for the quality settings of each of the six codecs are added up.

The resources allocated for pre-configured endpoints take precedence over any quality settings on the associated conference template and the resources cannot be optimized.

Allocating resources for TIP-compliant multiscreen endpoints

TIP-compliant multiscreen endpoints that are using the Cisco VCS's external policy server interface must be pre-configured, otherwise they will be treated in the same way as endpoints that are neither TIP-compliant nor pre-configured.

For TIP-compliant endpoints using the TelePresence Conductor's B2BUA for rendezvous calls the number of resources that are allocated on the conference bridge is based on the lesser of:

- the number of screens advertised by TIP, and
- the **Maximum screens** configured on the conference template.

The conference bridge allocates the appropriate resources based on the quality setting for each host, participant or guest joining the conference, multiplied by the number of screens.

For example, if TIP advertises 3 screens, but **Maximum screens** is set to 1, the required resources for the endpoint, based on the quality settings, are multiplied by 1.

If the TIP-compliant endpoint has been pre-configured, the resources are allocated according to the quality settings for the codec(s) and the resources cannot be optimized.

Allocating resources for multiscreen endpoints that are neither TIP-compliant nor pre-configured

If a multiscreen endpoint is neither TIP-compliant nor pre-configured, it is assumed to have only a single screen, unless it is:

- escalated into ad hoc conferences on the TelePresence Conductor,
- reserved as a host in a Lecture-type conference, or
- using the Cisco VCS's external policy server interface to call into a rendezvous conference.

If the endpoint falls into one of the categories listed above, the number of resources that are initially allocated is based on the **Maximum screens** configured on the conference template. For example, if **Maximum screens** is set to 3, the required resources for the endpoint, based on the quality settings, are multiplied by 3.

If the endpoint does not fall into one of the categories listed above, the resources allocated on the conference bridge are purely based on the quality settings. There will not be any resources allocated for additional screens for this endpoint and only a single screen will be displayed in the conference.

Optimizing resources

If **Optimize resources** is set for the conference template, resources that were initially allocated for a particular conference, may be freed up.

Resources are optimized and freed up on the TelePresence Server in the following situations:

- If the maximum capability an endpoint advertises is a lower quality than defined in the conference template for one of the following settings:
  - **Participant quality**
  - **Guest quality**
  - **Host quality**
- If an endpoint that uses the Cisco VCS's external policy server interface to dial into a rendezvous conference supports fewer screens than defined in the template under **Maximum screens**. (Endpoints in B2BUA deployments have resources for the correct number of screens allocated, because the TelePresence Conductor can detect the number of screens required from the SIP signaling.)
- If an endpoint that is escalated into an ad hoc conference supports fewer screens than defined in the template under **Maximum screens**.

Resources are optimized on the TelePresence Conductor, but not freed up on the TelePresence Server, if an endpoint that has been reserved as a host in a Lecture-type conference supports fewer screens than defined in the template under **Maximum screens**. The freed up resources can only be used for other hosts dialing into the same conference.

Resources are not optimized for auto-dialed participants or for pre-configured endpoints, because when configuring these entities the desired quality is defined in the configuration, and overrides the capabilities defined in the conference template for incoming calls.

Resource optimization does not happen immediately when a call arrives in a conference - a short while is left for the resource negotiations to stabilize.

Resource optimization is performed in the following way:

1. Endpoint capability is negotiated between the endpoint and the TelePresence Server.
2. The required resources are reported to the TelePresence Conductor.
3. The TelePresence Conductor performs per-caller optimization 5 seconds after the latest attendee joined the conference.
4. Callers who leave the conference will have all their resources returned provided 30 seconds have passed since the latest attendee joined the conference.

Limiting the number of participants in a conference

You can limit the total number of participants in a conference. The total number of participants to which the limit applies includes all auto-dialed participants (i.e. participants who are dialed in to the conference by the conference bridge), participants for which resources have been reserved (i.e. hosts), plus all other participants (i.e. participants who dial in to the conference using one of the conference aliases). The number of participants does not include any resources reserved for content or cascades.

- To place a limit on the number of participants, select the **Limit number of participants** check box and in the **Maximum** field, enter the total number of participants for the conference.
- If you do not want to place a limit on the number of participants, clear the **Limit number of participants** check box.

The default is for no limit.

The maximum number of participants must be higher than:

- the total number of auto-dialed participants associated with the template, plus
- all the participants for whom resources have been reserved (i.e. the setting in the **Number of hosts to reserve** field for Lectures).

This is to ensure that all these participants will be able to access the conference. If the maximum number of participants is not higher than these two numbers added together, the conference will not be created.

The reason that the maximum number of participants must be more than (rather than equal to) the number of auto-dialed and reserved participants is because a conference is not created until the first participant dials in, so the conference already has one participant when it is created. You must therefore set the maximum number of participants to a value that will allow the first dial-in participant, plus all auto-dialed and reserved hosts, plus any additional participants, to dial in to the conference.

You will receive a warning in any of the following situations:

- an auto-dialed participant is assigned to an existing template, and doing so means that the number of auto-dialed participants plus reserved hosts is equal to or higher than the maximum
- the number of reserved hosts is changed, and doing so means that the number of auto-dialed participants plus reserved hosts is equal to or higher than the maximum
- the maximum number of participants is changed to a number that is equal to or lower than the current number of auto-dialed participants plus reserved hosts.

Note: No preference is given to participants who have organized a conference. If the maximum number of participants is reached before the participant who organized the conference has dialed in, this participant is rejected.

Creating and editing conference aliases

A conference alias maps dialed aliases to conferences using regular expressions and specifies the user's role in the conference (participant, host or guest).

Conference aliases are required for all rendezvous conferences in Cisco VCS and Unified CM deployments.

Note: Conference aliases configured via the web interface are separate from aliases associated with Collaboration Meeting Rooms, which are provisioned via the TelePresence Conductor's Provisioning API.

When configuring each conference alias, you must specify the template to use for the conference, the name of the conference, and whether that user will be admitted to the conference as a participant, host or guest. To create or join a conference, an endpoint user must dial a specified alias.

Conference aliases use regular expressions, allowing you to use pattern matching and wildcards to specify the alias that users dial to access the conference, and the name of the conference when it is created on the conference bridge.

- For more information about regular expressions, see About regular expressions [p.206].
- For specific examples of how regular expressions can be used to set up conference aliases see Regular expression examples - conference aliases [p.207] and Regular expression examples - Lectures [p.210].

The **Conference aliases** page (**Conference configuration > Conference aliases**) lists all the existing conference aliases and allows you to edit, delete and create new conference aliases.

For Meetings, you must configure at least one conference alias. However, you can set up two or more aliases for the same conference.

For Lectures, you must configure at least two conference aliases - one for the host and one for guests.

There must not be any conflict between any **Incoming alias** or **Conference name**.

In a deployment using the Cisco VCS's external policy server interface, there must not be any conflict between any **Incoming alias** or **Conference name**, the Call Policy prefix [p.22], and Conference bridge dial plan prefixes [p.22]. Otherwise you may experience unpredictable behavior. For more information, see Considerations in a Cisco VCS-only deployment [p.22].

When creating or editing a conference alias, the configurable options are:

Field	Description
Name	Descriptive name of the conference alias.
Description	A free-form description of the conference alias.
Incoming alias (must use regex)	A regular expression (regex) that matches one or more aliases that a user can dial to access a conference.
Conference name	A regular expression (regex) replace string that defines how the **Incoming alias** will be modified to result in the conference name. This will be the same conference name that is then used on the conference bridge. We recommend that you use conference names that are 31 characters or fewer. See Conference name length [p.105] for more information.

Field	Description
Priority	Assigns a priority to the conference alias. The priority must be unique for each conference alias. The priority is used if the alias that has been dialed matches the **Incoming alias** of more than one conference alias. In such cases, the conference alias with the highest priority (closest to 0) will be used. If the alias that was dialed matches only one conference alias, the **Priority** won't be used but is still a required field.
Conference template	The template that is used when the conference is created. This will determine whether the conference is a *Meeting* or a *Lecture*, and thus what **Role types** will be available in the following field.
Role type	Determines the privileges that will be assigned to a caller dialing in to the conference using this conference alias. The options that are available are determined by the settings of the **Conference template** that has been selected in the previous field. *Participant* (available when the template **Type** is *Meeting*): the caller will join the conference as a host. *Host* (available when the template **Type** is *Lecture*): the caller will join the conference as a host. *Guest* (available when the template **Type** is *Lecture*): the caller will join the conference as a guest. See About host and guest roles [p.114] for more information on the differences between the two roles.
Allow conference to be created	Whether participants dialing this conference alias can create the conference or not. *Yes*: the first conference participant who dials this conference alias will create the resulting conference. *No*: conference participants cannot create a conference by dialing this conference alias. They can only join the resulting conference if it exists already. If the conference does not exist the conference participants are rejected. You must ensure that there is a way in which the conference can be created, either via the API or via another conference alias that matches to the same conference. The **default** is *Yes*.

Conference name length

TelePresence MCUs support conference names of up to 31 characters and TelePresence Servers support conference names of up to 80 characters. If the TelePresence Conductor has a conference name that is longer than the maximum number of supported characters it will hash the name and pass the hash value to the conference bridge for it to use as the conference name. The TelePresence Conductor will continue to use the original name itself.

If a conference name is longer than 31 (for TelePresence MCU) or 80 (for TelePresence Server) characters, you can view the hashed value on the **Conferences status** page (**Status > Conferences**):

- **Name:** shows the conference name used by the TelePresence Conductor
- **Conference name:** shows the hashed value, i.e. the conference name used by the conference bridge.

To avoid hashing, we recommend that you use conference names that are 31 characters or fewer for TelePresence MCUs and 80 characters or fewer for TelePresence Servers. You will need to carefully consider any regular expressions that you use in the **Conference name** field to ensure that all resulting conference names do not exceed this length. For information on how to test your dial plan see Check dial plan [p.176].

Creating and editing auto-dialed participants

Auto-dialed participants are addresses that are automatically dialed by the conference bridge(s) when a conference starts. The address could relate to a device such as an endpoint or recording device, or could be a FindMe ID.

Note: Auto-dialed participants configured via the web interface are separate from auto-dialed participants associated with Collaboration Meeting Rooms, which are provisioned via the TelePresence Conductor's Provisioning API.

This release of the TelePresence Conductor does not support auto-dialed participants that are multiscreen endpoints.

Each auto-dialed participant is associated with a **Conference template** and has a **Conference name match**. Whenever a conference is created using the specified **Conference template**, the resulting conference name is compared with the **Conference name match**. If there is a match, the conference bridge will automatically dial the specified participant. In this sense the Conference name match acts as a filter, so that only conferences using a specific template and with a specific name will have the auto-dialed participant added.

The **Auto-dialed participants** page (**Conference configuration > Auto-dialed participants**) lists all the existing auto-dialed participants and allows you to edit, delete and create new participants.

When creating or editing an auto-dialed participant, the configurable options are:

Field	Description
Name	Descriptive name of the auto-dialed participant.
Description	A free-form description of the auto-dialed participant.
Conference template	The template that, when used for conference creation, will cause this participant to be dialed in to the conference (if there is a match with the conference name).
Conference name match (must use regex)	A filter defining which conferences this participant will be added to. You must use a regular expression (regex) that can match one or more conference names. Participants will be added to the conference only if the conference name matches the regular expression. The default for this field is `(.*)`, which is a regular expression that will match against all possible conference names. This will result in the **Address** being dialed for all conferences created using the specified **Conference template**. The default regular expression needs to be used for auto-dialed participants in Unified CM ad hoc conferences, because all conference names are unique and generated by the Unified CM.
Participant address	The address that is automatically dialed by the conference bridge when a matching conference starts. You can enter an explicit auto-dialed participant address, or a regular expression that is used to produce the auto-dialed participant address. If using a regular expression, the `\n` notations (`\1`, `\2`, and so on) are replaced by the bracketed portions of the **Conference name match** field. **Note:** SIP addresses must include the domain (in other words, be in the format `address@example.com`).

Field	Description
Protocol	Determines the protocol that the conference bridge will use to call this participant. The options are *H.323* or *SIP*. If this auto-dialed participant will be used with Unified CM or with Cisco VCS via the TelePresence Conductor's B2BUA, *SIP* must be selected.
Role type	Determines the privileges that will be assigned to the participant when it is dialed in to the conference by the conference bridge. The options that are available are determined by the settings of the **Conference template** that has been selected: *Participant* (available when the template **Type** is *Meeting*): the participant will join the conference as a host. *Host* (available when the template **Type** is *Lecture*): the participant will join the conference as a host. *Guest* (available when the template **Type** is *Lecture*): the participant will join the conference as a guest. For more information on the differences between the two roles for a Lecture, see About host and guest roles [p.114]
DTMF sequence	Specifies a series of DTMF tones that will be sent by the conference bridge to the auto-dialed participant after the call has been connected. This feature can be used where the auto-dialed participant is a device such as an audio bridge that has an audio menu navigated by DTMF. There is a two second pause after the call connects after which the conference bridge will send the DTMF tones, which are sent one every half second. The DTMF sequence can include the digits **0**-**9** and the characters ***** and **#**. It can also include a comma (**,**), which represents a two-second pause. You can insert as many additional two-second pauses as you want. For more information about using DTMF, see Sending DTMF tones to an auto-dialed participant [p.109].
Keep conference alive	Determines whether or not the conference will end when all other participants have left the conference. *Yes*: the conference will keep running when only this auto-dialed participant remains. *No*: the conference will automatically end when only this auto-dialed participant remains. Beware that: ■ if the auto-dialed participant is an endpoint that cannot terminate a call itself, such as a recording device (for example a TelePresence Content Server or ISDN), you must select *No*. Selecting *Yes* will result in the conference never being terminated. ■ if the auto-dialed participant is an ISDN endpoint that has been set to auto-answer, selecting *Yes* may result in an unexpectedly high ISDN bill. ■ this setting will be ignored if: • the TelePresence MCU's API parameter **lastChairmanLeavesDisconnect** is set to *true* (this setting can be configured via the **Disconnect when last host leaves** field under **Conference configuration > Conference templates > Advanced parameters**), and • this auto-dialed participant's **Role type** is set to *Guest*.

Field	Description
Maximum quality	(Available when the **Conference template** has a conference bridge type of *TelePresence Server*) The maximum quality setting to apply to this auto-dialed participant. Depending on the selected setting the required resources are allocated on the TelePresence Server that is associated with this auto-dialed participant. The quality settings can be changed on the **Quality settings** page. The initial list of quality settings includes: ■ Full HD (1080p 30fps / 720p 60fps video, multi-channel audio) ■ HD (720p 30fps video, stereo audio) ■ SD (wide 448p / 480p 30fps video, mono audio) ■ 360p (360p 30fps video, mono audio) ■ Audio-only (no video, mono audio) 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. This setting overrides the quality defined on the conference template, resulting in the auto-dialed participant potentially experiencing a different quality compared with other participants in the same conference. The **default** is *HD (720p 30fps video, stereo audio)*.
State	If *Enabled*, calls will be made to this auto-dialed participant when a conference is created using the selected template. If *Disabled*, the auto-dialed participant is ignored.
Advanced parameters	**Advanced parameters** are parameters that can be passed to the conference bridge hosting the conference this auto-dialed participant is configured for, via its API. The parameters can be edited after an auto-dialed participant has been created. Click **Edit** to get to the **Advanced auto-dialed participant parameters** page, where you can select the parameters you would like to send to the conference bridge. This release of TelePresence Conductor only supports advanced auto-dialed participant parameters for TelePresence MCUs, not for TelePresence Servers. See Adding and editing advanced auto-dialed participant parameters [p.109] for more information. **CAUTION**: This feature is for advanced use only.

Using auto-dialed participants and Multiway

When planning to use Multiway with TelePresence Conductor, refrain from adding auto-dialed participants that are or could be using Multiway. Add only devices that can be trusted not to use Multiway, such as for example recording servers or audio bridges.

Example

If you define the rendezvous conference **`meet.ben@domain.com`** with the auto-dialed participant **`ben@domain.com`**, you must ensure that

- **either** `ben@domain.com`'s endpoint has Multiway disabled,
- **or** Ben must promise not to use the Multiway call flows.

Sending DTMF tones to an auto-dialed participant

If the auto-dialed participant is a device such as an audio bridge that has an audio menu navigated by DTMF, you can use the **DTMF sequence** field to specify a series of DTMF tones to send to the device after the call has been connected.

Example

You want the conference bridge to dial out to a PIN-protected audio conference on an audio bridge. The conference ID is `555` and the PIN is `888`. The audio bridge requires that you press `#` after entering the ID and after entering the PIN.

In this example you would set the **DTMF sequence** to be `555#,,888#`. The two commas represent a four second pause which allows the audio bridge's automated menu system time to process the ID and request the PIN.

What if an auto-dialed participant cannot be reached?

Sometimes the call to the auto-dialed participant might not be successful (for example, if the participant is busy or does not answer). You can control what the TelePresence MCU does in such situations by following these steps:

1. Log into the TelePresence MCU as an administrator.
2. Go to **Settings > Conferences**).
3. In the **Advanced settings** section, select one of the following options from the **Failed preconfigured participants redial behavior** field:

Never redial	The TelePresence MCU never attempts to redial a failed connection to this participant.
Redial until connected	The TelePresence MCU redials this participant if it fails unexpectedly when first establishing a connection; the TelePresence MCU never retries the connection if it fails after being established.
Redial on unexpected disconnection	The TelePresence MCU redials this participant on any unexpected disconnection, whether it occurs while first being established or at any point thereafter. It does not attempt to redial if the participant deliberately ends the connection.
Redial on any disconnection	The TelePresence MCU redials this participant when the connection closes, irrespective of whether the call fails or is deliberately ended by the participant.

Note: Each conference bridge must be configured identically for this and all other settings. Failure to do so will result in unpredictable behavior.

Adding and editing advanced auto-dialed participant parameters

CAUTION: This feature is for advanced use only.

This release of TelePresence Conductor only supports advanced auto-dialed participant parameters for TelePresence MCUs. It does not support them for TelePresence Servers.

Advanced auto-dialed participant parameters are parameters that can be configured for an auto-dialed participant and passed via its API to the conference bridge that hosts the conference this auto-dialed participant is associated with.

The parameters are selected on the **Advanced auto-dialed participant parameters** page, which is accessible via the **Auto-dialed participants** page (**Conference configuration > Auto-dialed participants**, then click **Edit** in the **Advanced parameters** section), when an auto-dialed participant has been created.

Advanced parameters for auto-dialed participants are configured and passed on in a similar way to advanced template parameters.

To create or edit advanced auto-dialed participant parameter settings:

1. Create a new auto-dialed participant or select an existing auto-dialed participant (**Conference configuration > Auto-dialed participants**).
2. In the **Advanced parameters** section click **Edit**.
   The **Advanced auto-dialed participant parameters** page will open.
3. Select the check-boxes next to all the parameters that should be sent to the conference bridge.
4. Enter or select the relevant parameter values.

Cisco TelePresence MCU parameters

The configurable options for TelePresence MCUs are:

Field name	Parameter in API	Description
Appear as a recording device	`actAsRecorder`	Whether this participant appears as a recording device to other participants.
Apply fixed gain	`audioRxGainMillidB`	If **Adaptive gain control** is *Fixed*, this is the gain applied, in millidB. It can be a negative value.
Adaptive gain control	`audioRxGainMode`	Whether and how audio gain is applied. Choose from: *None*: no extra gain is applied. *Automatic*: automatic gain control is applied. *Fixed*: a fixed number of millidBs of gain is applied.
Audio muted initially	`audioRxMuted`	Whether audio from this participant will be muted, so that this participant cannot be heard by other conference participants.
View border size	`borderWidth`	Controls the width of the outer border of the auto-dialed participant's layout. *0* indicates that there is no outer border added. *1*, *2* and *3* indicate that borders are added matching the width defined for these values on the TelePresence MCU. *3* is the widest.
Custom layout	`cpLayout`	The type of video layout seen by the conference participants. See Conference layouts [p.214] for a list of available layouts.

Field name	Parameter in API	Description
Display name override status	`displayNameOverrideStatus`	Whether the participant uses the **Display name override value** to identify itself.
Display name override value	`displayNameOverrideValue`	If **Display name override mode** is *true*, this value overrides the participant's display name.
H.323 gateway address	`gatewayAddress`	The address of an H.323 gateway, if required. Applicable only if the auto-dialed participant's **Protocol** is *H.323*.
Content negotiation	`h239Negotiation`	Defines how the TelePresence MCU presents itself for H.239 token negotiation. *As master*: the TelePresence MCU acts as master in H.239 token negotiation. *As slave*: the TelePresence MCU acts as the slave in H.239 token negotiation and can send content to a master unit if it accepts the token request. *Mimic slave*: the TelePresence MCU acts as a mimic slave in H.239 token negotiation and will try to send content to all other endpoints/units even if this unit (i.e. the mimic slave) rejects the token request.
Layout control via FECC/DTMF	`layoutControlEx`	Defines how the video layout can be controlled for this auto-dialed participant. *Disabled*: layout control is disabled. *FECC only*: the auto-dialed participant can only change the view layout using far end camera control (FECC). *DTMF only*: the auto-dialed participant can only change the view layout using dual-tone multi-frequency (DTMF). *FECC with DTMF fallback*: the auto-dialed participant can change the view layout using FECC when it is available and via DTMF , when FECC is not available. *Both FECC and DTMF*: the auto-dialed participant can change the view layout using both FECC and DTMF.
Link type	`linkType`	Defines the type of auto-dialed participant. *Default*: the auto-dialed participant is an endpoint or recording device that is called into the conference. *Cascade slave to master*: the auto-dialed participant is a cascaded conference bridge that is called into a conference. This option allows for participants connected to a particular conference bridge to be joined into a conference that is hosted on another (master) conference bridge and for both to appear as one large conference bridge.
Preferred bandwidth from MCU	`maxBitRateFromMCU`	Maximum bandwidth from the TelePresence MCU in kbps.
Preferred bandwidth to MCU	`maxBitRateToMCU`	Maximum bandwidth to the TelePresence MCU in kbps.

Field name	Parameter in API	Description
Motion/sharpness tradeoff	`motionSharpnessTradeoff`	Defines the preference for motion versus sharpness. The options are: *Default*: use the global default setting. *Prefer motion*: prefer motion at the expense of sharpness. *Prefer sharpness*: prefer sharpness at the expense of motion. *Balanced*: balance the motion and sharpness trade-off.
Password	`password`	The password the TelePresence MCU uses to access VNC endpoints.
Mute in-band DTMF	`suppressDtmfEx`	Controls the muting of in-band dual-tone multi-frequency (DTMF) tones. *FECC*: in-band DTMF tones will be muted when DTMF is being used to control layout because far end camera control (FECC) is not available. *Always*: in-band DTMF tones will always be muted. *Never*: in-band DTMF tones will never be muted.
Transport protocol	`transportProtocol`	Defines the SIP transport protocol. Applicable only if the auto-dialed participant's **Protocol** is *SIP*. One of Default, TCP, UDP or TLS. Default is the default transport protocol configured on the TelePresence MCU.
Use SIP registrar	`useSIPRegistrar`	Whether the auto-dialed participant uses the SIP registrar. Applicable only if the auto-dialed participant's **Protocol** is *SIP*.
Received video resolutions	`videoRxMaxResolution`	The maximum resolution of the received video. The options are: *CIF*: this participant sends CIF or lower resolution to the TelePresence MCU. *4CIF*: this participant sends 4CIF or lower resolution to the TelePresence MCU. *Maximum*: this participant sends the maximum resolution that both sides can support.
Video muted initially	`videoRxMuted`	Whether video from this participant is muted and the participant cannot be seen by other conference participants.
Transmitted video resolutions	`videoTxMaxResolution`	The maximum resolution of the transmitted video. The options are: *CIF*: the TelePresence MCU sends CIF or lower resolution to this participant. *4CIF*: the TelePresence MCU sends 4CIF or lower resolution to this participant. *Maximum*: the TelePresence MCU sends the maximum resolution that both sides can support.
Transmit widescreen video	`videoTxWidescreen`	Whether the TelePresence MCU sends video in a form suitable for a widescreen 16:9 display to this participant.

Field name	Parameter in API	Description
Custom parameters		This field can be used to enter advanced parameters and their corresponding values in valid JSON. The TelePresence Conductor does not perform in-depth checking of data in these fields. Any incorrect configuration of the settings may result in your auto-dialed participant failing to be called into the conference.

Cisco TelePresence Server parameters

This release of TelePresence Conductor only supports advanced auto-dialed participant parameters for TelePresence MCUs. It does not support them for TelePresence Servers.

About host and guest roles

Assigning roles

In a TelePresence conference participants can have a role of either **Host** or **Guest**. Depending on the role, the participant has different capabilities and uses different amounts of resources.

TelePresence Conductor conferencing allows for these methods of determining the role of a participant:

- **By alias**: Host and guest participants dial separate aliases. The role is determined by the alias the participant dials. Each alias has a role of either host or guest associated with it. This method is supported for conferences configured via the TelePresence Conductor web interface and CMRs provisioned via the TelePresence Conductor Provisioning API.
- **By PIN**: Host and guest participants dial the same alias. There is a PIN defined for hosts and an optional PIN for guests. The role is determined by the PIN that the participant has entered. This method is supported only for CMRs provisioned via the TelePresence Conductor Provisioning API. It is not supported for conferences configured via the TelePresence Conductor web interface.

There are two ways in which a participant can join a conference:

- By dialing an alias.
- By being automatically dialed into an existing conference by the conference bridge. This type of participant is called an auto-dialed participant.

When a participant dials an alias, the role is either:

- Set to *Host* by TelePresence Conductor, because the alias was configured to have a role of host. These aliases have a role of host:
  - All aliases for Meeting-type conferences configured via the TelePresence Conductor web interface
  - Some aliases for Lecture-type conferences configured via the TelePresence Conductor web interface
  - Some aliases that are part of CMRs provisioned via the TelePresence Conductor Provisioning API
- Set to *Guest* by TelePresence Conductor, because the alias was configured to have a role of guest. These aliases have a role of guest:
  - Some aliases for Lecture-type conferences configured via the TelePresence Conductor web interface
  - Some aliases that are part of CMRs provisioned via the TelePresence Conductor Provisioning API
- Set to *Host* by the conference bridge hosting the conference, because the alias was configured to determine the role by PIN and the participant entered the valid host PIN.
- Set to *Guest* by the conference bridge hosting the conference, because the alias was configured to determine the role by PIN and the participant entered a valid guest PIN or no PIN (depending on the configuration).

When an auto-dialed participant is dialed into a conference by the conference bridge, the role is either:

- Determined by the role defined in the **Auto-dialed participant** associated with the relevant conference template or provisioned CMR
- Restricted to *Host*, if the provisioned CMR is configured to determine the role by PIN

Awareness of roles

The various devices involved in hosting and managing the TelePresence conference have a different level of awareness of a participant's role.

Conference bridges have full awareness of the participant's role. They can become aware of the role either when the role is determined by alias or by PIN.

The TelePresence Conductor has full awareness of the participant's role when the role was determined by alias. When the role is determined by PIN, the TelePresence Conductor does not see the PIN a participant enters and must treat all participants as having the same role, an undetermined role. When role is determined by PIN the following constraints apply:

- Host and guest quality must be specified to be the same.
- Host resources cannot be reserved.

Differences between host and guest roles

The detailed behavior and options available to **Host** and **Guest** participants are determined by the conference bridge that the conference is hosted on. The differences are described below.

Starting the conference

On a TelePresence MCU:

A conference will not begin until the first host joins. Guests who join a conference before the first host has joined will see a black screen with the on-screen text *Waiting for conference chairperson*. There will be no audio, apart from an audio prompt after five seconds and every minute thereafter. This behavior is not configurable.

On a TelePresence Server:

Conferences can be configured so that their guests must wait for the first host to join. This can be set using either of these methods:

- Via custom advanced template parameters on the TelePresence Conductor conference template
- Via a management tool such as Cisco TMSPE using the TelePresence Conductor's Provisioning API (attribute **`guests_wait_for_host`** on the **`ConfBundle`** object)

If only guests are present in a conference where guests must wait for a host, they remain in a lobby screen, and no guest can see or hear any other guest, until a host joins the conference. The lobby screen can have a customized message, which can be set via the custom advanced template parameters.

Ending the conference

On a TelePresence MCU:

You can control the behavior when the last host leaves the conference using the **When only guests remain** setting on the TelePresence MCU.

This setting can be modified via the advanced template parameter **Disconnect when last chairperson leaves** on the TelePresence Conductor under **Conference configuration > Conference template > Advanced parameters**.

The two options are:

- *true* (default) - all remaining guests will be disconnected
- *false* - all participants may continue the conference until the last one disconnects

The option *true* should only be set when the TelePresence Conductor's **Conference template** has a **Conference type** of *Lecture*.

Beware that if the TelePresence MCU parameter **Disconnect when last host leaves** is set to *true* then any auto-dialed participant with **Role type** of *Guest* will be disconnected when all other non-guest participants have left the conference. This applies even if **Keep conference alive** is set to *Yes*.

For more information see Adding and editing advanced template parameters [p.87].

On a TelePresence Server:

It is possible to set participants to automatically disconnect when only participants configured to automatically disconnect are left in a conference. This must be set explicitly via the TelePresence Server API parameter **autoDisconnect**. It can be applied to both host and guest auto-dial participants.

Taking the chair

(This section is only applicable to TelePresence MCUs.)

Only a host can "take the chair". On "taking the chair", a participant can:

- nominate a "broadcaster"; that is, they can choose which participant's video will be sent to all other participants in "1 x 1 view" (full-screen view)
- decide to disconnect any other participant(s)

This behavior is not configurable.

Note: "taking the chair" is only supported for H.323 calls and only if floor and chair control has been enabled on the TelePresence MCU. Not all endpoints support the H.243 floor and chair control functionality.

Conference layout in automatic lecture mode

(This section is only applicable to TelePresence MCUs.)

When automatic lecture mode is configured on the TelePresence MCUs, there is a difference in the conference layout for hosts and guests. The hosts see their custom layout, whereas guests see only the host who is currently speaking.

For more information on automatic lecture mode see *Understanding how participants display in layout views* in the *Cisco TelePresence MCU Online Help*.

Creating and editing Locations

TelePresence Conductor supports conferences between endpoints registered directly with Unified CM version 8.6.2 or later. A Location is needed to mimic the Unified CM's expectation that it is connecting to separate conference bridges in different locations. Both ad hoc conferences and rendezvous conferences are supported.

A Location is also required when the TelePresence Conductor is directly connected to one or more Cisco VCS(s) via the back-to-back user agent (B2BUA). One Location is used for all Cisco VCSs (or Cisco VCS clusters), which the TelePresence Conductor is connected to.

Before being able to create a new Location, you must:

- Configure sufficient additional IP addresses on the TelePresence Conductor
- Configure at least one conference template of type 'Meeting' (applicable to Locations where **Conference type** is *Ad hoc* or *Both*)

To create a new Location:

1. Go to **Conference configuration > Locations**.
2. Click **New**.

When creating or editing a Location, the configurable options are:

Field	Description
Location name	A descriptive name of the Location.
Description	A free-form description of the Location.
Conference type	Determines the type of conference supported by this Location. *Ad hoc*: a spontaneous meeting where a user on Unified CM brings two or more people together in a conference using the conference button on the phone. *Rendezvous*: a non-scheduled conference for which the host knows the conference dial-in beforehand and must share it with all participants. *Both*: ad hoc and rendezvous conference. **Note:** for conference types *Ad hoc* and *Both* at least one conference template of type 'Meeting' must have been configured prior to creating the Location.
Ad hoc IP address (local)	(Only applicable to ad hoc conferences) The IP address on TelePresence Conductor used for ad hoc conferences associated with this Location. The drop-down list contains IP addresses already configured for this TelePresence Conductor, but not yet assigned to a Location. If there are no IP addresses in the list, go to **System > Network interfaces > IP** to add more IP addresses to the TelePresence Conductor. Each cluster peer must be configured individually, as the IP address is unique per peer.
Template	(Only applicable to ad hoc conferences) The conference template to use for ad hoc conferences associated with this Location. Only 'Meeting'-type templates are displayed in the drop-down box. **Note:** at least one 'Meeting'-type conference template must have been configured to be able to create a Location of type 'Ad hoc' or 'Both'.

Field	Description
Rendezvous IP address (local)	(Only applicable to rendezvous conferences) The IP address on the TelePresence Conductor used for rendezvous conferences or for outbound calls, for example to auto-dialed participants. The drop-down list contains IP addresses already configured for this TelePresence Conductor, but not yet assigned to a Location. If there are no IP addresses in the list, go to **System > IP** to add more IP addresses to the TelePresence Conductor. Each cluster peer must be configured individually, as the IP address is unique per peer.
Trunk 1-3 IP address	(Only applicable to conferences that have out-dial participants) The far-end (call control device) IP address of the SIP trunk between TelePresence Conductor and the call control device. The trunk IP address is required if: ■ there are auto-dialed participants configured on the TelePresence Conductor. ■ Cisco TMS schedules a conference with participants. ■ a user of conference control center (CCC) in Cisco TMS adds a participant to an existing conference. If you specify more than one trunk IP address, the TelePresence Conductor considers all trunk IP addresses for a Location as equivalent. It may use any of the trunk IP addresses defined, as long as the destination is reachable. If the SIP trunk destination that the TelePresence Conductor currently uses becomes unreachable, it will automatically use another reachable destination. The TelePresence Conductor maintains only one of the destinations, it does not load balance the dial-out calls across the configured destinations. The TelePresence Conductor regularly polls all SIP trunk destinations for their reachability. It raises an alarm, if any destinations are unreachable. It also reports any reachability status changes in the event log.
Trunk 1-3 Port	(Only applicable to conferences that have out-dial participants) The far-end (call control device) port number of the SIP trunk between the TelePresence Conductor and the call control device. The **default** is *5061*.
Trunk transport protocol	(Only applicable to conferences that have out-dial participants) The transport protocol of the SIP trunk(s) between the TelePresence Conductor and the call control device. The options are either *TLS* or *TCP*. The **default** is *TLS*.

Creating and editing pre-configured endpoints

The TelePresence Conductor supports endpoints with more than one screen, provided the conferences they dial into are hosted on a TelePresence Server and the field **Allow for multiscreen** has been set to *Yes* on the **Conference templates** page. Supported endpoints are:

- Cisco TelePresence System T3 (T3) and multiscreen endpoints that are compliant with the TelePresence Interoperability Protocol (TIP), e.g. Cisco TelePresence System series (CTS)
- Custom endpoints, which are all other multiscreen endpoints that are compatible with TelePresence Server version 3.0 or later and have up to four codecs

For rendezvous calls using the TelePresence Conductor B2BUA, TelePresence Conductor can identify the number of screens supported by TIP-compliant endpoints through the signaling and so TIP-compliant endpoints in this deployment do not need to be added as pre-configured endpoints. T3s, custom endpoints and TIP-compliant endpoints using the Cisco Cisco VCS's external policy interface have to be pre-configured, otherwise only the codec that is dialing into the conference (one single screen) will be displayed.

It is optional to pre-configure single-screen endpoints. If they are pre-configured, the TelePresence Conductor ensures that sufficient resources are allocated on the conference bridge to guarantee the quality level.

TelePresence MCUs do not support multiscreen endpoints.

It is not possible to define the content quality for a pre-configured endpoint. The content quality for a conference that an endpoint dials into is defined in the conference template.

Pre-configured endpoints do not get their resources optimized after joining a conference.

Pre-configuring endpoints involves creating a pre-configured endpoint and then adding between one and four codecs to it.

To create a new pre-configured endpoint:

1. Go to **Conference configuration > Pre-configured endpoints**.
2. Click **New**.

When creating or editing a pre-configured endpoint, the configurable options are:

Field	Description
Name	A descriptive name of the endpoint.
Description	A free-form description of the endpoint.
Endpoint type	The type of endpoint to pre-configure. *Custom*: any endpoints that have more than one codec *Multiscreen TIP endpoint or T3*: multiscreen endpoints that are compliant with the Telepresence Interoperability Protocol, such as Cisco TelePresence System series, and Cisco TelePresence System T3 endpoints. For these endpoints to be treated as a multiscreen endpoint the associated conference template must have a **Content quality** of at least *1280 x 720p 5fps*. Only the details of the primary codec need to be pre-configured. **Note:** when configuring multiscreen endpoints you must also enable **Allow for multiscreen** on all conference templates that are to check for pre-configured endpoints.

Field	Description
State	Whether or not the pre-configured endpoint is enabled. When the pre-configured endpoint is *Disabled*, the endpoint can still dial into a conference, but resources are not allocated for it and it is seen as a single-screen endpoint. The **default** is *Enabled*.
Display name	Name to display in conference (in preference to endpoint URI) if endpoint name display is enabled on the conference bridge.
Receiver audio gain	Volume adjustment for incoming audio (used to adjust loudspeaker volume across codecs). *0* means that there is no change in volume of incoming audio. Negative values indicate a lower volume and positive values indicate a higher volume. Values are in dBm.
Transmitter audio gain	Volume adjustment for outgoing audio (used to adjust loudspeaker volume across codecs). *0* means that there is no change in volume of outgoing audio. Negative values indicate a lower volume and positive values indicate a higher volume. Values are in dBm.
Initial outgoing audio	When initially joining a conference the endpoint audio may be muted or active. The **default** is *Active*.
Initial incoming audio	When initially joining a conference the received audio may be muted or active. The **default** is *Active*.
Initial outgoing video	When initially joining a conference the endpoint video may be muted or active. The **default** is *Active*.
Initial incoming video	When initially joining a conference the received video may be muted or active. The **default** is *Active*.
Cameras crossed	Whether the cameras on the endpoint are crossed or not. The setting will be used to work out the order in which to display the screens of this endpoint on other endpoints in the conference. The **default** is *Not crossed*.
Bypass conference PIN entry	Whether to bypass the conference PIN entry screen. A number of older multiscreen endpoints do not allow the user to enter PINs and PIN entry has to be bypassed. We strongly recommend that you do not bypass PIN entry unless the endpoint does not support PIN entry. The **default** is *Do not bypass PIN entry*.

Field	Description
Legacy TIP endpoint	Whether TIP (Telepresence Interoperability Protocol) negotiation should be forced on this endpoint. We strongly recommend that you select *No*, unless the endpoint explicitly needs this configuration. If you select *Yes* for an endpoint that does not support TIP, the call will fail. If you select *Yes* for an endpoint that supports TIP, although the call will succeed, not all functionality may be available. The endpoints that require this setting to be *Yes* include: ■ CTS endpoints running versions 1.7.3 or earlier. ■ Polycom HDX/OTX systems. ■ CTS or TX endpoints that have their calls routed through a Unified CM running versions 8.0 or earlier. The **default** is *No*.
Codec layout order	(Only displayed when at least two codecs have been created for this endpoint) The order in which the codecs are arranged for this pre-configured endpoint. Drag the codecs into the correct order.

After creating a pre-configured endpoint, at least one codec has to be added to it.

If two or more codecs have been added you can configure the following settings:

- **Codec layout order**: the order in which the codecs are arranged for this pre-configured endpoint. In the **Configuration** section drag the codecs into the correct order.
- **Single screen audio**: the codec which receives audio from any single-screen endpoints in the conference. In the **Codecs** section select the appropriate codec.
- **Single screen content**: the codec which displays content from any single-screen endpoints in the conference. In the **Codecs** section select the appropriate codec.

Adding and editing pre-configured endpoint codecs

Each pre-configured endpoint must have at least one and at most four codecs configured.

Telepresence Interoperability Protocol (TIP) negotiated multiscreen endpoints, such as CTS and T3, only require the primary codec to be pre-configured.

To add a new codec to an existing pre-configured endpoint:

1. Go to **Conference configuration > Pre-configured endpoints**.
2. Select a pre-configured endpoint to add a codec to.
3. Click **Create codec** in the **Codec** section.

When adding or editing a pre-configured endpoint codec, the configurable options are:

Field	Description
Name	A descriptive name of the pre-configured endpoint codec.
Description	A free-form description of the codec.

Field	Description
Pre-configured endpoint	The pre-configured endpoint to which this codec belongs. This is not configurable.
Protocol	The protocol the codec uses. The options are *H.323* or *SIP*. The protocol is only required for outgoing codecs. The **default** is *SIP*.
Address	The SIP or H.323 address of the codec. The SIP address is the endpoint's SIP URI, and the H.323 address can be one of the endpoint's E164 number(s) or H.323 ID(s). All codec addresses must be unique.
Optional address 1-5	Additional SIP or H.323 addresses for the codec. The codec may appear to be at a different address, depending on the node in a cluster of call control devices to which the codec is registered. All possible addresses, from which calls for this codec could come in, must be added here. All codec addresses must be unique.
Direction	The direction of call signaling from the conference bridge's perspective. *Incoming*: The codec must call into the conference. *Outgoing*: The conference bridge will call the codec. At least one codec within a pre-configured endpoint should have a direction of *Incoming*, otherwise the endpoint cannot receive any calls.
Maximum quality	The video and audio quality setting that this codec will use. It will override the setting selected for the associated conference template. Depending on the selected setting the appropriate resources are allocated on the TelePresence Server. The quality settings can be changed on the **Quality settings** page. The initial list of quality settings includes: ■ Full HD (1080p 30fps / 720p 60fps video, multi-channel audio) ■ HD (720p 30fps video, stereo audio) ■ SD (wide 448p / 480p 30fps video, mono audio) ■ 360p (360p 30fps video, mono audio) ■ Audio-only (no video, mono audio) 360p video is only supported in TelePresence Server version 3.1 or later. If 360p is configured on a TelePresence Server that is running an earlier software version, SD video is used instead and the resource usage will be higher than expected. When using a CTS3000 or TX9000 you must select *Full HD (1080p 30fps / 720p 60fps video, multi-channel audio)* or a custom quality setting that has an audio quality level of multi-channel, otherwise insufficient resources will be allocated to display multiple screens. TelePresence Conductor allocates the same amount of resources on the TelePresence Server for both types of Full HD video quality settings (1080p 30fps and 720p 60fps). If 60fps is supported on the endpoint, the TelePresence Server will choose 720p 60fps over 1080p 30fps. The **default** is *HD (720p 30fps video, stereo audio)*.

Using Call Policy

About Call Policy

This feature is only applicable to deployments using the Cisco VCS's external policy server interface.

When Call Policy is in use, the TelePresence Conductor will check with the Cisco TelePresence Video Communication Server (Cisco VCS) to determine whether a user who is attempting to create a particular conference has the right to do so. Call Policy is enabled on a per-template basis, and requires a **Call Policy prefix** to be configured on the TelePresence Conductor. It also requires an appropriate Call Policy to be configured on the Cisco VCS.

In a deployment using the Cisco VCS's external policy server interface, there must not be any conflict between any **Incoming alias** or **Conference name** (used when Creating and editing conference aliases [p.104]), the Call Policy prefix [p.22], and Conference bridge dial plan prefixes [p.22]. Otherwise you may experience unpredictable behavior. For more information, see Considerations in a Cisco VCS-only deployment [p.22].

When to use Cisco VCS or TelePresence Conductor Call Policy

The TelePresence Conductor's Call Policy works in conjunction with the Cisco VCS's Call Policy and allows you to distinguish between those users you want to permit to **create** a conference based on a particular template, and those you want to permit to **join** the conference. Whether you use Call Policy on the Cisco VCS only, or on both the Cisco VCS and TelePresence Conductor depends on the desired result, as follows:

- **To prevent a user from creating or joining any conferences:** use Cisco VCS Call Policy only, so that the request never reaches the TelePresence Conductor.
- **To prevent a user from creating a particular conference, but allow them to join the existing conference:** enable Call Policy on the TelePresence Conductor in conjunction with an appropriate Call Policy on the Cisco VCS, as per the information in this section.
- **To allow a user to create and join a conference:** Call Policy is not required on the TelePresence Conductor or Cisco VCS.

Note: When Call Policy is enabled on the TelePresence Conductor, it applies only to users attempting to create a conference that does not already exist. After the conference has been created (by a user who is allowed to dial the prefix), a user who previously dialed the conference alias and had their call rejected because they did not have the right to create the conference will be able to dial the same conference alias and successfully join the conference.

Defaults

The default **Call Policy prefix** is **create.**. If no **Call Policy prefix** is configured, Call Policy on the TelePresence Conductor will not work, so this field cannot be left blank.

Note: In numeric dial plans, this field will need to be changed.

Configuring Call Policy

To use Call Policy on the TelePresence Conductor:

1. Enable or disable Call Policy on a per-template basis using the **Call Policy mode** setting on the **Conference templates** page (**Conference configuration > Conference templates**).
2. If any templates are using Call Policy, then you must configure the TelePresence Conductor with a prefix to use. This is done using the **Call Policy prefix** setting on the **Call Policy** page (**Conference configuration > Call Policy**). The same prefix is used for all templates that have Call Policy enabled.
3. Configure the Cisco VCS with an appropriate Call Policy that will allow only those users permitted to create conferences to place calls that start with the **Call Policy prefix**. See *Cisco TelePresence Video Communication Server Administrator Guide* and *Cisco TelePresence Conductor with Cisco TelePresence Video Communication Server Deployment Guide* for more information.

Example usage

In the following example:

- the TelePresence Conductor has been configured with a conference alias of **meet.alice** which creates a conference with the name **alice**
- the **meet.alice** alias uses a template that has Call Policy enabled
- the TelePresence Conductor's **Call Policy prefix** is **create.**
- the Cisco VCS is configured with a Call Policy that says Alice is the only person allowed to dial **create.meet.alice**

Call not allowed

Ben dials **meet.alice**.

The Cisco VCS forwards the request to the TelePresence Conductor.

The TelePresence Conductor looks up the **meet.alice** alias and sees that the conference name for that alias is **alice**. It then checks to see whether the **alice** conference already exists. It does not, so it adds the Call Policy prefix (**create.**) to the conference alias that it received and sends the resulting string (**create.meet.alice**) back to the Cisco VCS.

The Cisco VCS checks its Call Policy to see whether Ben is allowed to dial **create.meet.alice**. He is not, so the Cisco VCS rejects the call.

Call allowed

Alice dials **meet.alice**.

The Cisco VCS forwards the request to the TelePresence Conductor.

The TelePresence Conductor looks up the **meet.alice** alias and sees that the conference name for that alias is **alice**. It then checks to see whether the **alice** conference exists. It does not, so it adds **create.** to the conference alias and sends the resulting string (**create.meet.alice**) back to the Cisco VCS.

The Cisco VCS checks its Call Policy to see whether Alice is allowed to dial **create.meet.alice**. She is, so the Cisco VCS forwards the request for **create.meet.alice** to the TelePresence Conductor.

When the TelePresence Conductor receives a conference alias that begins with the same string as the **Call Policy prefix**, it interprets this as approval from the Cisco VCS to create the associated conference. So when it receives **create.meet.alice** it strips the **create.** prefix, looks up the resulting **meet.alice** conference alias and follows the settings for that alias to create the conference **alice**, with Alice as the first participant.

Ben then dials **meet.alice**.

The Cisco VCS forwards the request to the TelePresence Conductor.

The TelePresence Conductor looks up the **meet.alice** alias and sees that the conference name for that alias is **alice**. It then checks to see whether the **alice** conference already exists. It does, so Ben is allowed to join Alice in the **alice** conference.

Scheduling a WebEx conference on the TelePresence Conductor

A WebEx-enabled conference is created on the TelePresence Conductor with the **`factory.webex.add`** API call, which typically comes from a management tool such as Cisco TMS.

For a WebEx conference to be supported on the TelePresence Conductor:

- Signaling between the conference bridge and WebEx must be "early offer".
  - For a Cisco VCS deployment, "early offer" is supported by default.
  - For a Unified CM deployment, SIP trunks must be configured to support “early offer" and not to insert an MTP (for more information see the latest *Optimized Conferencing for Cisco Unified CM and Cisco VCS Solution Guide*)
- The conference bridges used to host the WebEx conference must be:
  - TelePresence MCUs running software version 4.4 or later, or
  - TelePresence Servers running software version 3.1 or later and configured in *Remotely managed* mode.
- The TelePresence Conductor must have a conference template configured with the **Scheduled conference** field set to *Yes*.
  This conference template must be used solely for scheduled conferences.
- The conference template must have either **Allow content** set to *Yes* (for TelePresence MCUs) or **Content quality** set to a quality setting that is not *Off* (for TelePresence Servers).

There are two ways in which a conference bridge may connect to a WebEx conference:

SIP video with TSP audio

In this method, video and content traffic use a SIP connection to WebEx, whereas audio traffic is routed via a telephony network (PSTN) to a telephony service provider (TSP).

On a TelePresence MCU, one port is used for audio traffic, one port is used for video traffic and an additional port is used for content.

On a TelePresence Server, two calls will be used, one for video only and one for audio only. The total number of resources allocated for these calls is defined in the **Host quality** (for Lecture-type conferences) or the **Participant quality** (for Meeting-type conferences) on the conference template that is used for the WebEx conference. Additional TelePresence Server resources are used for content.

SIP video and audio

All types of traffic, video, content and audio, use a SIP connection to WebEx. There is no telephony network involved in this method.

On a TelePresence MCU, a single port is used for video and audio traffic and an additional port is used for content.

On a TelePresence Server the number of resources allocated for video and audio traffic is defined in the **Host quality** (for Lecture-type conferences) or the **Participant quality** (for Meeting-type conferences) on the conference template that is used for the WebEx conference. Additional TelePresence Server resources are used for content.

Configuring users

This section provides information on how to configure the root account, administrator accounts, administrator groups and LDAP accounts.

Configuring administrator accounts

The **Administrator accounts** page (**Users > Administrator accounts**) lists all the local administrator accounts that have been configured on the TelePresence Conductor, and lets you add, edit and delete accounts. There is one pre-configured administrator account, which can have its username and password changed.

The account credentials for the administrator accounts can be used by an administrator to log in to the TelePresence Conductor web interface, by the Cisco TelePresence Video Communication Server when accessing the TelePresence Conductor policy service, or by third party applications such as Cisco TelePresence Management Suite (Cisco TMS) to access the TelePresence Conductor.

The administrator accounts can only be used when the **Administrator authentication source** on the **LDAP configuration** page (**Users > LDAP configuration**) has been set to *Local* or *Both*.

Note: The default username for the administrator user is **admin** and the default password is **TANDBERG**. The TelePresence Conductor's conference functionality is disabled until this password has been changed. It is important to select a secure password.

Adding local administrator accounts

You can add additional local administrator accounts. These accounts can be used to access the TelePresence Conductor over the web and API interfaces.

The configurable options are:

Field	Description	Usage tips
Name	The username for the administrator account.	Some names such as "root" are reserved. Local administrator account user names are case sensitive.
Access level	The access level of the administrator account: *Read-write*: allows all configuration information to be viewed and changed. This provides the same rights as the default **admin** account. *Read-only*: allows status and configuration information to be viewed only and not changed. Some pages, such as the **Upgrade** page, are blocked to read-only accounts. Default: *Read-write*	The access permissions of the currently logged in user are shown in the system information bar at the bottom of each web page.
Password	The password that this administrator will use to log in to the TelePresence Conductor.	All passwords on the TelePresence Conductor are encrypted, so you only see placeholder characters here. When entering passwords, the bar next to the **Password** field changes color to indicate the complexity of the password. You can configure the complexity requirements for local administrator passwords on the Password security page (**Users > Password security**). You cannot set blank passwords.
New password	Enter a new password for the account.	This field only appears when you are changing a password.

Field	Description	Usage tips
Confirm password	Re-enter the password for the account.	This field only appears when you create an account or when you change its password.
Web access	Select whether this account is allowed to log in to the system using the web interface. Default: *Yes*	
API access	Select whether this account is allowed to access the system's status and configuration using the Application Programming Interface (API). Default: *Yes*	This controls access to the XML and REST APIs by systems such as Cisco TMS.
State	Select whether the account is *Enabled* or *Disabled*. Disabled accounts are not allowed to access the system.	
Your current password	Enter your own, current password here if the system requires you to authorize a change.	To improve security, the system requires that administrators enter their own passwords when creating an account or changing a password.

Editing administrator account details

You can edit the details for the pre-configured administrator account and for additional local administrator accounts.

Go to **Users > Administrator accounts**. Under **Actions** for the relevant administrator account, click **Edit user**.

A new page is displayed, where you can edit all fields for the selected administrator account except for the password. To change the password, see below.

Changing the administrator password

You can change the password for the pre-configured administrator account and for additional local administrator accounts.

Go to **Users > Administrator accounts**. Under **Actions** for the relevant administrator account, click **Change password**.

A new page is displayed, where you can change the password for the selected administrator. Enter the new password and confirm it. You must also enter the password for the administrator account with which you are currently logged in to authorize the password change.

Configuring remote account authentication using LDAP

The **LDAP configuration** page (**Users > LDAP configuration**) is used to configure an LDAP connection to a remote directory service for administrator account authentication.

The configurable options are:

Field	Description	Usage tips
Remote account authentication: this section allows you to enable or disable the use of LDAP for remote account authentication.		
Administrator authentication source	Defines where administrator login credentials are authenticated. *Local only*: credentials are verified against a local database stored on the system. *Remote only*: credentials are verified against an external credentials directory. *Both*: credentials are verified first against a local database stored on the system, and then if no matching account is found the external credentials directory is used instead. The default is *Local only*.	*Both* allows you to continue to use locally-defined accounts. This is useful while troubleshooting any connection or authorization issues with the LDAP server. Note that you cannot log in using a locally-configured administrator account if *Remote only* authentication is in use.
LDAP server configuration: this section specifies the connection details to the LDAP server.		
FQDN address resolution	Defines how the LDAP server address is resolved. *SRV record*: DNS SRV record lookup. *Address record*: DNS A record lookup. *IP address*: entered directly as an IP address. The default is *Address record*.	
Host name and **Domain** or **Server address**	The way in which the server address is specified depends on the **FQDN address resolution** setting: *SRV record*: only the **Domain** portion of the server address is required. *Address record*: enter the **Host name** and **Domain**. These are then combined to provide the full server address for the DNS address record lookup. *IP address*: the **Server address** is entered directly as an IP address.	If using TLS, the address entered here must match the CN (common name) contained within the certificate presented by the LDAP server.
Port	The IP port to use on the LDAP server.	Typically, non-secure connections use *389* and secure connections use *636*.

Field	Description	Usage tips
Encryption	Determines whether the connection to the LDAP server is encrypted using Transport Layer Security (TLS). *TLS*: uses TLS encryption for the connection to the LDAP server. *Off*: no encryption is used. The default is *Off*.	When TLS is enabled, the LDAP server's certificate must be signed by an authority within the TelePresence Conductor's trusted CA certificates file.
Authentication configuration: this section specifies the TelePresence Conductor's authentication credentials to use when binding to the LDAP server.		
Bind DN	The distinguished name (case insensitive) used by the TelePresence Conductor when binding to the LDAP server. It is important to specify the DN in the order cn=, then ou=, then dc=	Any special characters within a name must be escaped with a backslash as per the LDAP standard (*RFC 4514*). Do not escape the separator character between names. The bind account is usually a read-only account with no special privileges.
Bind password	The password (case sensitive) used by the TelePresence Conductor when binding to the LDAP server.	The maximum plaintext length is 60 characters, which is then encrypted.
SASL	The SASL (Simple Authentication and Security Layer) mechanism to use when binding to the LDAP server. *None*: no mechanism is used. *DIGEST-MD5*: the DIGEST-MD5 mechanism is used. The default is *DIGEST-MD5*.	Enable Simple Authentication and Security Layer if it is company policy to do so.
Bind username	Username of the account that the TelePresence Conductor will use to log in to the LDAP server (case sensitive). Only required if SASL is enabled.	Configure this to be the sAMAccountName; Security Access Manager Account Name (in AD this is the account's user logon name).
Directory configuration: this section specifies the base distinguished names to use when searching for account and group names.		
Base DN for accounts	The ou= and dc= definition of the Distinguished Name where a search for user accounts should start in the database structure (case insensitive). It is important to specify the DN in the order ou=, then dc=	The Base DN for accounts and groups must be at or below the dc level (include all dc= values and ou= values if necessary). LDAP authentication does not look into sub dc accounts, only lower ou= and cn= levels.
Base DN for groups	The ou= and dc= definition of the Distinguished Name where a search for groups should start in the database structure (case insensitive). It is important to specify the DN in the order ou=, then dc=	If no **Base DN for groups** is specified, then the Base DN for accounts will be used for both groups and accounts.

Checking the LDAP server connection status

The status of the connection to LDAP server is displayed at the bottom of the page.

State = Active

No error messages are displayed.

State = Failed

The following error messages may be displayed:

Error message	Reason / resolution
DNS unable to do reverse lookup	Reverse DNS lookup is required for SASL authentication.
DNS unable to resolve LDAP server address	Check that a valid DNS server is configured, and check the spelling of the LDAP server address.
Failed to connect to LDAP server. Check server address and port	Check that the LDAP server details are correct.
Failed to setup TLS connection. Check your CA certificate	CA certificate, private key and server certificate are required for TLS.
Failure connecting to server. Returned code<return code>	Other non-specific problem.
Invalid Base DN for accounts	Check **Base DN for accounts**; the current value does not describe a valid part of the LDAP directory.
Invalid server name or DNS failure	DNS resolution of the LDAP server name is failing.
Invalid bind credentials	Check **Bind DN** and **Bind password**, this error can also be displayed if SASL is set to *DIGEST-MD5* when it should be set to *None*.
Invalid bind DN	Check **Bind DN**; the current value does not describe a valid account in the LDAP director. This failed state may be wrongly reported if the **Bind DN** is 74 or more characters in length. To check whether there is a real failure or not, set up an administrator group on the TelePresence Conductor using a valid group name. If TelePresence Conductor reports “saved” then there is not a problem (the TelePresence Conductor checks that it can find the group specified). If it reports that the group cannot be found then either the **Bind DN** is wrong, the group is wrong or one of the other configuration items may be wrong.
There is no CA certificate installed	CA certificate, private key and server certificate are required for TLS.
Unable to get configuration	LDAP server information may be missing or incorrect.

Configuring administrator groups

The **Administrator groups** page (**Users > Administrator groups**) lists all the administrator groups that have been configured on the TelePresence Conductor, and lets you add, edit and delete groups.

Administrator groups only apply if remote account authentication is enabled.

When you log in to the TelePresence Conductor web interface, your credentials are authenticated against the remote directory service and you are assigned the access rights associated with the group to which you belong. If the administrator account belongs to more than one group, the highest level permission is assigned.

The configurable options are:

Field	Description	Usage tips
Name	The name of the administrator group. It cannot contain any of the following characters: / \ [ ] : ; | = , + * ? > < @ "	The group names defined in the TelePresence Conductor must match the group names that have been set up in the remote directory service to manage administrator access to this TelePresence Conductor.
Access level	The access level given to members of the administrator group: *Read-write*: allows all configuration information to be viewed and changed. This provides the same rights as the default **admin** account. *Read-only*: allows status and configuration information to be viewed only and not changed. Some pages, such as the **Upgrade** page, are blocked to read-only accounts. *None*: no access is allowed. Default: *Read-write*	If an administrator belongs to more than one group, it is assigned the highest level permission for each of the access settings across all of the groups to which it belongs (any groups in a disabled state are ignored). See Determining the access level for accounts that belong in multiple groups [p.133] below for more information.
Web access	Determines whether members of this group are allowed to log in to the system using the web interface. Default: *Yes*	
API access	Determines whether members of this group are allowed to access the system's status and configuration using the Application Programming Interface (API). Default: *Yes*	This controls access to the XML and REST APIs by systems such as Cisco TMS.
State	Indicates if the group is enabled or disabled. Access will be denied to members of disabled groups.	If an administrator account belongs to more than one administrator group with a combination of both *Enabled* and *Disabled* states, their access will be *Enabled*.

Determining the access level for accounts that belong in multiple groups

If an administrator belongs to groups with different levels of access, the highest level of access is granted. Any groups in a disabled state are ignored.

For example, if the following groups were configured:

Group name	Access level	Web access	API access
Administrators	Read-write	-	-
Region A	Read-only	Yes	-
Region B	Read-only	-	Yes
Region C	Read-only	Yes	Yes

The following table shows examples of the access permissions that would be granted for accounts that belong in one or more of those groups:

Groups belonged to	Access permissions granted
Administrators and **Region A**	read-write access to the web interface but no API access
Administrators and **Region B**	read-write access to the API interface, but no web interface access
Administrators and **Region C**	read-write access to the web and API interfaces
Region A only	read-only access to the web interface and no API access

Viewing active administrator sessions

The **Active administrator sessions** page (**Users > Active administrator sessions**) lists all administrator accounts that are currently logged in to this TelePresence Conductor.

It displays details of their session including their login time, session type, IP address and port, and when they last accessed this TelePresence Conductor.

You can terminate active web sessions by selecting the required sessions and clicking **Terminate session**.

You may see many sessions listed on this page if a large **Session time out** value is configured. This typically occurs if an administrator ends their session by closing down their browser without first logging out of the TelePresence Conductor.

Configuring password security

The **Password security** page (**Users > Password security**) controls whether or not local administrator account passwords must meet a minimum level of complexity before they are accepted.

If **Enforce strict passwords** is set to *On*, all subsequently configured local administrator account passwords must conform to the following rules for what constitutes a strict password.

Configurable rules

The following rules apply by default but can be customized.

The password must contain at least 15 ASCII characters made up of at least:

- 2 numeric values ['0'..'9']
- 2 uppercase letters ['A'..'Z']
- 2 lowercase letters ['a'..'z']
- 2 special characters [such as '@' or '$']

You can also specify:

- the minimum number of the 4 character classes (numeric , lower case, upper case, and special characters) that must be present; use this setting if you want to mandate the use of 2-3 different character classes without requiring all of them to be present
- the maximum number of times the same character can be repeated consecutively; by default there is no restriction

Additional non-configurable rules

The following strict password rules always apply:

- Do not include dictionary words
- Do not include too many consecutive characters such as "abc" or "123"
- Avoid multiple instances of the same characters
- Do not use palindromes

If **Enforce strict passwords** is set to *Off*, no checks are made on administrator passwords.

Note:

- Regardless of this setting, it is not possible to set a blank password for any administrator account.
- This setting affects local administrator account passwords only. It does not affect any other passwords used on the TelePresence Conductor such as in the local authentication database, LDAP server, external registration credentials, user account passwords, or administrator account passwords stored on remote credential directories.
- All passwords and usernames are case sensitive.

Configuring the root account

The TelePresence Conductor provides a root account which can be used to log in to its operating system. This account has a username of **root** (all lower case) and a default password of **TANDBERG** (all upper case).

For security reasons you must change the password as soon as possible. Conference functionality is disabled and an alarm is displayed on the web interface until the **root** password is changed from the default.

Note: Do not use the **root** account in normal operation. Never perform system configuration using this account. Use the Configuring administrator accounts [p.128] instead.

Changing the root account password

To change the password for the **root** account:

1. Log in to the TelePresence Conductor as root. By default you can only do this using a serial connection or SSH.
   Note: If you have forgotten the **root** account password, see Resetting forgotten passwords [p.138] for instructions on how to reset it.
2. Type **`passwd`**.
   You will be asked for the new password.
3. Enter the new password and when prompted, retype the new password.
   You will receive the message:
   `passwd: password updated successfully`
4. Type **`exit`** to log out of the root account.

Enabling and disabling access over SSH

By default, the root account can be accessed over either a serial connection or SSH.

To enable and disable access to the root account using SSH:

1. Log in to the TelePresence Conductor as **`root`**.
2. Type one of the following commands:
   - **`rootaccess -s on`** to enable access using SSH
   - **`rootaccess -s off`** to disable access using SSH
3. Type **`exit`** to log out of the root account.

If you have disabled SSH access while logged in using SSH, your current session will remain active until you log out, but all future SSH access will be denied. The only way you can then re-enable SSH access is to log in using a serial connection and run the **`rootaccess -s on`** command.

Resetting forgotten passwords

Note: The username and password for the administrator account is replicated across peers in a cluster. Therefore if you change the username or password on one peer, it will be changed on all other peers.

The root account password is not replicated across peers.

Resetting your administrator password if you still have access to the root account

If you still have access to the **root** account, but you have forgotten your password for an administrator account, you can reset the administrator password using the following procedure:

1. Log in to the **root** account via a serial connection or SSH.
2. Enter the command **`passwd`** followed by the username of the administrator account.
3. When prompted enter the new password twice.

Resetting your administrator password or root password on a VM

If you have forgotten the password for either an administrator account or the **root** account and you are using a VM (Virtual Machine) TelePresence Conductor, you can reset it using the following procedure:

1. Open the vSphere client.
2. Click on the link **Launch Console**.
2. Reboot the TelePresence Conductor.
3. In the vSphere console log in with the username **`pwrec`**. No password is required.
4. When prompted, select the account (*root* or the username of the administrator account) whose password you want to change.
5. You will be prompted for a new password.

The **pwrec** account is only active for one minute following a reboot. After that time you will have to reboot the system again to reset the password.

Resetting your administrator password or root password on an appliance

If you have forgotten the password for either an administrator account or the **root** account and you are using an appliance TelePresence Conductor, you can reset it using the following procedure:

1. Connect a PC to the TelePresence Conductor using the serial cable as per the instructions in the *Cisco TelePresence Conductor Getting Started Guide*.
2. Reboot the TelePresence Conductor.
3. Log in from the PC with the username **`pwrec`**. No password is required.
4. When prompted, select the account (*root* or the username of the administrator account) whose password

you want to change.

5. You will be prompted for a new password.

The **pwrec** account is only active for one minute following a reboot. After that time you will have to reboot the system again to reset the password.

Viewing status

This section describes how to view status information on the TelePresence Conductor, accessible via the **Status** menu.

Getting a status overview

The **Overview** page (**Status > Overview**) gives an overview of the current status of the TelePresence Conductor.

The following information is displayed:

Field	Description	Notes
System host name	The name that has been assigned to the TelePresence Conductor by the system administrator.	This setting is configured on the **DNS** page (**System > DNS**).
IPv4 address	The TelePresence Conductor's IPv4 address.	This setting is configured on the **IP** page (**System > Network interfaces > IP**).
Hardware up time	The amount of time that has elapsed since the system last rebooted.	
Product	This will be Cisco TelePresence Conductor.	
Serial number	The serial number of the hardware or virtual machine on which the TelePresence Conductor software is installed.	
Software version	The version of software that is currently installed on the TelePresence Conductor.	To upgrade to a new version of software, see Upgrading software components [p.168].
Software build	The build number of this software version.	
Software release date	The date on which this version of the software was released.	
Number of conference bridges	The number of conference bridges that have been configured on the TelePresence Conductor. The list of conference bridges can be viewed and edited on the **All conference bridges** page (**Conference configuration > Conference bridges**).	For further information about each conference bridge, go to **Status > Conference bridges**. See About conference bridges [p.61] for more information regarding conference bridges.
Number of active conferences	The number of conferences currently taking place.	For further information about each conference, go to **Status > Conferences**.
Number of calls	The number of calls that are currently passing through the TelePresence Conductor.	For further information about the active calls, go to **Status > Calls > Calls**.

Alarms

Alarms occur when an event or configuration change has taken place on the TelePresence Conductor that requires some manual administrator intervention, such as a restart. Alarms may also be raised for hardware and environmental issues such as faulty disks and fans or high temperatures.

For a list of the alarm categories that can appear on the TelePresence Conductor, see Alarm categories [p.225].

For a list of the alarms that can be raised on the TelePresence Conductor see Alarms list [p.226].

Viewing alarms

The **Alarms** page (**Status > Alarms**, or by clicking on the red Alarm icon which appears at the top right of any page when an alarm is in place) provides a list of all the active alarms on your system (and, in the **Action** column where applicable, their proposed resolution). If your system is part of a cluster, this page will display all alarms across all peers in the cluster.

Actioning alarms

CAUTION: Do not run a system with unresolved alarms because functionality and performance may be affected.

Deal with each alarm immediately by clicking each **Action** and making the necessary configuration changes to resolve the problem.

If you are experiencing any problems with the TelePresence Conductor, immediately investigate and fix any active alarms.

If your system is part of a cluster, action each alarm on the peer to which it relates. The **Action** hyperlink will redirect you to the relevant peer. You may see multiple copies of the same alarm disappearing at the same time if they can be fixed on just one peer.

Acknowledging alarms

Acknowledging all alarms (by selecting the alarms and clicking on the **Acknowledge** button) removes the Alarm icon from the web interface, but the alarms will still be listed on the **Alarms** page with a status of *Acknowledged*. If a new alarm occurs, the Alarm icon will reappear.

After any configuration changes to the TelePresence Conductor, or following a restart of the system, any *Acknowledged* alarms that are still unresolved will reappear with a status of *Raised*, and must be re-acknowledged. Acknowledged alarms need to be re-acknowledged every 24 hours.

Deleting alarms

You cannot delete alarms from the **Alarms** page. Alarms are removed by the TelePresence Conductor only after the required action or configuration change has been made.

Alarm information

The table below details some of the fields that are included on the **Alarms** page:

Field	Description
Alarm	A description of the alarm.
State	**Raised**: a new alarm **Acknowledged**: the alarm is still in place but has been acknowledged by an administrator.
Severity	The severity of the condition that caused the alarm to be raised. See Alarm severity [p.143] for definitions.
Peer	If the TelePresence Conductor is part of a cluster, this indicates which peer the alarm relates to.
Action	How to resolve the situation that led to the alarm being raised. Where possible this will include a link to the page where any required configuration changes can be made.
ID	An identifier for the alarm. This can be provided to Cisco TAC engineers if required.

Alarm severity

The table below lists, in order of priority, each of the levels of severity that can be assigned to an alarm, and the definition of each.

Severity	Description
Emergency	A condition has occurred with the TelePresence Conductor hardware. Immediate action is required.
Alert	A condition has occurred with the TelePresence Conductor software. Immediate action is required.
Critical	The TelePresence Conductor has been configured in a way that will completely prevent it from working. Immediate action is required.
Error	A condition has occurred that will affect the performance of the TelePresence Conductor but it will continue to function to some extent.
Warning	A condition has occurred that may affect the performance of the TelePresence Conductor but it will continue to function to some extent.
Notice	A normal but significant condition has occurred. The TelePresence Conductor will continue to function normally.
Info	Information messages.
Debug	Information that Cisco TAC engineers may use for debugging.

Conference bridge status

The **Conference bridge status** page (**Status > Conference bridges**) shows all the conference bridges in your system's conference bridge pools, and their current status.

For each conference bridge the configuration status, the operational status, the conferences it is hosting and the resource usage is displayed.

When using a TelePresence Conductor without a release key, you can enable only one conference bridge across all conference bridge pools. All other conference bridges have a **State** of *Busied out*. The conference bridge that is enabled cannot be clustered. If the only enabled conference bridge is clustered, its **Status** is *Unusable* and the TelePresence Conductor is unable to communicate with any conference bridges.

The resource usage on a TelePresence MCU is measured in number of ports. The resource usage on a TelePresence Server is measured in number of screen licenses and calls.

The utilization is determined by the largest value out of the allocated resources and the used resources for a particular conference bridge. The utilization bar is shown in green when the conference bridge utilization is below the threshold at which a conference bridge resource alarm should be raised; which can be configured on the **Global conference bridge settings** page (**Conference configuration > Global conference bridge settings**). The utilization bar is shown in yellow when the conference bridge is approaching the threshold and in red when it has exceeded the threshold.

Allocated resources are those that have been requested on the TelePresence Conductor when the conference is first initiated. Used resources are those that the conference bridge actually uses when the conference is hosted on the conference bridge. Resource optimization can reduce the number of used resources on a TelePresence Server.

Conferences status

The **Conferences status** page (**Status > Conferences**) shows the number of conferences currently being managed by the TelePresence Conductor, and provides detailed information on each conference.

The list of active conferences includes

- conferences based on Collaboration Meeting Rooms (CMRs) provisioned via the TelePresence Conductor's API
- conferences based on conference templates that have been configured via the TelePresence Conductor's web interface

Depending on the type of conference, there is a link to the associated **Conference template** or **Collaboration meeting room**.

A conference can have one of the following states:

- Starting
- Running
- Stopping

If a conference has the word 'LOCKED' displayed next to it, it means that the conference is locked. Conferences can be locked and unlocked via the TelePresence Conductor's API. When a conference is locked:

- The conference continues to run with its existing participants
- No new participants can dial in
- More participants can be added to the conference via the TelePresence Conductor's API (for example using Cisco TMSPE)

For conferences provisioned via the Provisioning API the **Conference display name** and all **Host aliases** and **Guest aliases** are displayed. The aliases are displayed in lower case, although they are case insensitive when dialed by users.

For each conference the number of **Participant** resources that are reserved, requested or used are displayed. The information is displayed in summary, as well as for each conference bridge hosting the conference.

Host resources can be reserved on conference templates for Lecture-type conferences or via the Provisioning API. **Participant** resources, which apply to Meeting-type conferences and **Guest** resources, which apply to Lecture-type conferences cannot be reserved.

Auto-dialed participant resources can be requested by creating an auto-dialed participant with a particular role type and associating it with a conference template. Auto-dialed participants provisioned via the Provisioning API are not displayed.

The **Webex** counter **Used** shows resources that are either in use or waiting for a participant to be added.

The maximum number of **Webex** resources that are reserved and used per conference are:

- On a TelePresence MCU:
  - 1 port, if using SIP
  - 2 ports, if using SIP-TSP

- On a TelePresence Server:
  - 1 participant, if using SIP
  - 1 participant, if using SIP-TSP

  Where 1 participant uses the number of resources that have been defined for the participant quality (in a Meeting-type conference template) or for the host quality (in a Lecture-type conference)

Conference participants

The **Conference participants** page (**Status > Conferences**, then for a particular conference click **View the participants in this conference**) provides information about all the participants in a particular conference.

The information available includes:

Participant	The registered alias (H.323 ID, E.164, SIP AOR) of the endpoint being used by the conference participant.
State	*Connected*: this participant is currently connected to the conference. *Disconnected*: this participant is no longer connected to the conference. *Dormant*: this indicates an auto-dialed participant that was unable to be dialed in to the conference.
Conference name	The name of the conference as it appears on the conference bridge. If the name is longer than 31 characters (for TelePresence MCUs) or 80 characters (for TelePresence Servers), a hash value is displayed. If the conference name is an IP address, the IP address is displayed as a hash value that has been turned into a hexadecimal value.
Host	(Available when **Conference type** is *Lecture*) *Yes*: this participant has a role of Host. *No*: this participant has a role of Participant or Guest.
Conference bridge address	The address of the conference bridge (as it appears on the conference bridge pool page) to which this participant is connected. Click on the address to go to the web interface for this conference bridge.
Call direction	*Incoming*: the participant joined the conference by dialing a conference alias. *Outgoing*: the participant was auto-dialed in to the conference by the conference bridge.

Collaboration meeting rooms

Searching Collaboration Meeting Rooms

On the **Collaboration meeting room** page (**Status > Provisioning > Collaboration meeting room**) you can search for one or more Collaboration Meeting Rooms (CMR).

A CMR provides the details for a conference provisioned via the TelePresence Conductor's API. Apart from Service Preferences it does not contain any data that is configured via the TelePresence Conductor's web interface.

For more information on provisioning conferences see Provisioning conferences [p.25].

To search for one or more CMR, enter one of the following:

- **Conference display name** - the conference name passed to the Provisioning API that does not need to be unique and can therefore return results for more than one CMR. This string is case sensitive.
- **Conference name** - the unique conference name that the TelePresence Conductor generates from the conference display name when the CMR is created or updated. This string is case sensitive.
- **Alias** - the direct match alias string that the participant dials to get into the conference. This string is case insensitive.
  Note: even though this search string can be case insensitive, the direct match alias is stored by TelePresence Conductor in lower case.
- **Auto-dialed participant** - the address for an auto-dialed participant. This string is case sensitive.

Note: All details entered into the search must have been configured via the TelePresence Conductor's API, not via the web interface.

A search returns a list of CMRs that match the search term.

- Searching for the conference name will return at most one CMR.
- Searching for conference display name, alias or auto-dialed participant can return more than one CMR. If there are more than 200 CMRs matching the search term, only 200 CMRs are displayed in the list.

For each item in the list, you can

- click on the conference name or **View** link to view details about the CMR (see Viewing Collaboration Meeting Rooms [p.148])
- click on a specific alias link to view the details about the alias in a pop-up window
- click on a specific auto-dialed participant link to view the details about the alias in a pop-up window

Viewing Collaboration Meeting Rooms

The details page for a Collaboration Meeting Room (CMR) is a read-only display of the data related to a specific CMR.

A CMR is added as a ConfBundle object via the TelePresence Conductor's Provisioning API, using a management tool such as Cisco TMS. The data cannot be changed via the TelePresence Conductor's web interface.

The CMR fields displayed are:

Field	Description
Conference name	The unique conference name that the TelePresence Conductor generates from the conference display name when the CMR is created or updated. This consists of the initial characters from the **Conference display name** as received by the Provisioning API, followed by some pseudo-random characters. This string is case sensitive and must be entered in the **Collaboration meeting room** search page using the same case as stored by TelePresence Conductor.
Conference display name	The conference name passed to the Provisioning API that does not need to be unique. This string is case sensitive and must be entered in the **Collaboration meeting room** search page using the same case as provided to the API.
Aliases	The list of exact match aliases that the participant can dial to get into the conference. Each alias is displayed with corresponding details. The **Exact match alias** string is case insensitive. It is stored in lower case by the TelePresence Conductor Provisioning API. Participants can enter the alias using any case and the alias will still be matched. The string can be entered using any case on the **Collaboration meeting room** search page.
Service Preference	The Service Preference used. It defines the order in which conference bridge pools should be selected to host a conference. Service Preferences are configured via the TelePresence Conductor's web interface and accessed by the Provisioning API.
Maximum number of participants	The maximum number of participants allowed to join a conference based on this CMR. The number includes auto-dialed participants and reserved hosts.
Maximum conference duration (minutes)	The maximum limit on the duration of the conference based on this CMR. The conference bridges issue warnings that are displayed on the endpoints when the conference is about to end.
Allow content	Whether or not participants will be able to send content video, such as a presentation.
Layout	The video layout scheme to be seen by participants joining conferences based on this CMR. The layout will be one of the following types, which have been defined in the Provisioning API: *equal*: conference participants are shown in a grid pattern of equal sized panes, up to 4x4. (Not applicable to multiscreen endpoints) *active*: the active speaker is shown in a large pane with additional participants appearing in up to nine PIPs (picture-in-pictures) overlaid at the bottom of the screen. *prominent*: the active speaker is shown in a large pane with additional participants appearing in up to four smaller panes at the bottom of the screen. (Not applicable to multiscreen endpoints) *single*: the active speaker is shown in one full-screen pane. *Not set*: no layout has been set through the Provisioning API and no layout preference will be sent to the conference bridge. This will result in the conference bridge displaying its default layout. Depending on the conference bridge capabilities, the closest approximation to the specified layout will be used. Where applicable, multiscreen systems will be mapped to the closest approximation to the specified layout. See Conference layouts [p.214] for more information on layout options available on the conference bridge types.
Host quality	(Only applicable to TelePresence Server-hosted conferences) The video and audio quality level that participants with host privileges will experience.

Field	Description
Guest quality	(Only applicable to TelePresence Server-hosted conferences) The video and audio quality level that participants with guest privileges will experience.
Content quality	(Only applicable to TelePresence Server-hosted conferences) The content quality that participants viewing content video, such as a presentation, will experience.
Allow multiscreen	(Only applicable to TelePresence Server-hosted conferences) Whether or not the conference allows for multiscreen systems. If *No* was selected the conference only allows for single-screen systems or the primary screen of multiscreen systems.
Maximum screens	(Only applicable to TelePresence Server-hosted conferences) The maximum number of screens an endpoint in this conference is allowed to have; in the range of 1 to 4. This field is only applicable if **Allow multiscreen** is set to *Yes*.
Optimize resources	(Only applicable to TelePresence Server-hosted conferences) Whether or not to allow TelePresence Conductor to optimize the resources used by participants in the conference. For more information see Optimizing resources [p.101].
Number of reserved hosts	The number of hosts for whom resources should be reserved.
Maximum number of cascades	The maximum number of cascades allowed for this conference. For each cascade the appropriate number of resources are reserved on the primary conference bridge.
Scheduled conference	Whether or not conferences generated from this CMR are scheduled. If the value is *Yes* the conference can only be created via the API call `factory.conferencecreate` and not by participants dialing the conference alias.
Guests wait for host	Whether or not the guests must wait for a host to join the conference before they are able to join. This setting is only applicable to conferences hosted on TelePresence Servers. It is ignored for conferences hosted on TelePresence MCUs.
PIN	Separate PINs can be set for participants with host and with guest privileges. **Note:** you cannot set a separate PIN for the cascade portion of a conference.
Auto-dialed participants	The list of auto-dialed participants that the conference bridge dials into the conference when it starts. Each auto-dialed participant is displayed with corresponding details. The **Address** string is case sensitive and must be entered in the **Collaboration meeting room** search page using the same case as provided to the API.
Advanced parameters	JSON parameters that are sent to the conference bridge to change advanced configuration options.
Cascade advanced parameters	JSON parameters that are sent to the cascade conference bridges to change advanced configuration options. The parameters must be the same as for **Advanced parameters**.

Call status information

The **Status > Calls** pages provide information about the current and historic calls passing through the TelePresence Conductor. These pages list all calls from Unified CM for which TelePresence Conductor receives signaling, i.e. calls that Unified CM sends directly to TelePresence Conductor, without going through a SIP trunk between Unified CM and Cisco VCS, and calls from a Cisco VCS that do not use the Cisco VCS's external policy server interface.

Call status

Call status information can be displayed for both current and completed calls:

- **Current calls**: the **Call status** page (**Status > Calls > Calls**) lists all the calls currently passing through the TelePresence Conductor.
- **Completed calls**: the **Call history** page (**Status > Calls > Call history**) lists all the calls that are no longer active. The list is limited to the most recent 500 calls, and includes only calls that have taken place since the TelePresence Conductor was last restarted.

If the TelePresence Conductor is part of a cluster, all calls that apply to any peer in the cluster are shown, although the list is limited to the most recent 500 calls per peer.

Call summary information

The following summary information is displayed initially:

Field	Description
Start time	The date and time when TelePresence Conductor took the call.
End time	The date and time when the call ended on TelePresence Conductor (completed calls only).
Duration	The length of time of the call.
Source	The alias of the endpoint that placed the call.
Destination	The alias dialed from the endpoint. This will be different from the alias to which the call was placed, which will have been transformed.
Status	The reason the call ended (completed calls only). A call has a status of *Disconnected* when an administrator has terminated the call using the **Disconnect** button. A call has a status of *BYE* when either the caller or callee has terminated the call normally.
Peer	Identifies the cluster peer through which the call is being made.
Actions	Click **View** to see further information about the call.

Call detail information

After selecting a call from the primary list (as described above) you are shown further details of that call.

Disconnecting calls

Click **Disconnect** to disconnect the selected calls. If your TelePresence Conductor is part of a cluster you have to be logged into the peer through which the call is passing to be able to disconnect the call.

Event Log

About the Event Log

The TelePresence Conductor provides an event logging facility for troubleshooting and auditing purposes. This Event Log is a list of all the events that have occurred on your system since the last upgrade and records information about such things as conference creation and deletion, requests to join a conference, alarms raised, and conference bridge status changes. It may also contain system-level information.

The Event Log holds 22GB of data; when this size is reached, the oldest entries are overwritten. However, only the first 50MB of event log data can be displayed through the web interface. The entire event log is included in a system snapshot.

The **Event Log** page (**Status > Logs > Event Log > All events**) lets you view and filter the Event Log.

The other sub-menus under the **Status > Log > Event Log** menu provide you with a filtered view of the Event Log as follows:

- **Conference creation events** shows only those events relating to the creation of new conferences
- **Conference join events** shows only those events relating to users joining a conference
- **Conference destruction events** shows only those events relating to a conference being destroyed

Filtering the Event Log

The **Filter** area lets you view a subset of events based on words contained in the events.

By default, you can use the **Contains all of the words** field. Enter the words you want to search for and click **Filter**. Only those events that contain all the words you entered are shown.

To do more advanced filtering, click **more options**. This gives you additional filtering methods:

- **Contains the string**: only includes events containing the exact phrase entered here.
- **Contains any of the words**: includes any events that contain at least one of the words entered here.
- **Not containing any of the words**: filters out any events containing any of the words entered here.

Note:

- Use spaces to separate each word you want to filter by.
- You can use any combination of the above fields.

To reapply any modified filter conditions, click **Filter**.

To return to the complete Event Log listing, click **Reset**.

Reconfiguring the log settings

Clicking **Configure log settings** takes you to the Logging configuration page. From this page, you can set up one or more remote servers to which the event log can be copied.

Saving the results to a local disk

Click **Download this page** if you want to download the contents of the results section to a text file on your local PC or server.

Viewing events

The **Results** area shows all the events matching all the current filter conditions, with the most recent being shown first.

Many events contain hyperlinks in one or more of the fields (such fields change color when you hover over them). You can click on the hyperlink to show only those events that contain the same text string. For example, clicking on the text that appears after *Level*= filters the list to show only the events at that particular level.

Event Log color coding

Certain events in the Event Log are color-coded so that you can identify them more easily.

- **Green** indicates a successful event
- **Orange** acts as a warning, indicates an event about which you should be aware
- **Red** indicates a failure of some kind

Configuration Log

The **Configuration Log** page (**Status > Logs > Configuration Log**) provides a list of all changes to the TelePresence Conductor configuration.

The Configuration Log holds a maximum of 30MB of data; when this size is reached, the oldest entries are overwritten. The entire Configuration Log can be displayed through the web interface.

Filtering the Configuration Log

The **Filter** section lets you filter the Configuration Log. Enter the words you want to search for and click **Filter**. Only those events that contain all the words you entered are shown.

To do more advanced filtering, click **more options**. This gives you additional filtering methods:

- **Contains the string**: only includes events containing the exact phrase entered here.
- **Contains any of the words**: includes any events that contain at least one of the words entered here.
- **Not containing any of the words**: filters out any events containing any of the words entered here.

Use spaces to separate each word you want to filter by.

Click **Filter** to reapply any modified filter conditions. To return to the complete Configuration Log listing, click **Reset**.

Results section

The **Results** section shows all the web-based events, with the most recent being shown first.

Most events contain hyperlinks in one or more of the fields (such fields change color when you hover over them). You can click on the hyperlink to show only those events that contain the same text string. For example, clicking on the text that appears after **Event=** filters the list to show all the events of that particular type. Likewise, clicking on a particular **user** shows just those events relating to that particular administrator account.

Configuration Log events

Changes to the TelePresence Conductor configuration made by administrators using the web interface have an **Event** field of *System Configuration Changed*.

The **Detail** field of each of these events shows:

- the configuration item that was affected
- what it was changed from and to
- the name of the administrator user who made the change, and their IP address
- the date and time that the change was made

Clustering

This section provides information on how to configure a TelePresence Conductor to be part of a cluster of up to three TelePresence Conductor systems.

About clusters

A TelePresence Conductor can be part of a cluster of up to three full capacity TelePresence Conductor systems or up to two TelePresence Conductor Select systems. Each TelePresence Conductor in the cluster is a peer of every other TelePresence Conductor in the cluster. When a cluster has been created, any configuration changes made to one peer are shared immediately among all other peers.

Clusters are used to provide redundancy in the rare case that a TelePresence Conductor becomes unavailable (for example, due to a network or power outage).

Note: Clustering of the TelePresence Conductor is not supported when the TelePresence Conductor is running without a release key, as TelePresence Conductor Essentials.

To avoid confusion, we recommend that you make all configuration changes on one peer. The only exception to this is any Peer-specific configuration [p.157].

For more information, see the relevant deployment guide:

- *Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide*

Peer IP addresses

Note: Never change the IP address of a TelePresence Conductor that is part of a cluster. see Changing a peer's IP address [p.161] for more information.

Peers in a cluster are identified by IP address.

Cluster pre-shared key

The TelePresence Conductor uses IPSec (Internet Protocol Security) to enable secure communication between each cluster peer.

The **Cluster pre-shared key** is the common IPSec access key used by each peer to access every other peer in the cluster. This field is alphanumeric.

Each peer in the cluster must be configured with the same **Cluster pre-shared key**. This is a required field for peers in a cluster.

Note: A strong pre-shared key is important for system security and for the security of your video network.

Changing the pre-shared key

If you change the pre-shared key of one peer in the cluster, that peer will not be able to communicate with any other peers in the cluster that have a different pre-shared key. If you must change the cluster's pre-shared key, change it on all peers simultaneously.

Peer-specific configuration

Most items of configuration are applied to all peers in a cluster. However, the following items must be specified separately on each cluster peer.

Cluster configuration

The list of Peer IP addresses (including the peer's own IP address) that make up the cluster has to be specified on each peer and they **must** be identical on each peer (the order in which they appear is not important).

The cluster pre-shared key has to be specified on each peer and **must** be identical for all peers.

Ethernet

The Ethernet speed is specific to each peer. Each peer may have slightly different requirements for the connection to their Ethernet switch.

IP

Note: Never change the Primary LAN 1 IP address of a TelePresence Conductor that is part of a cluster. The only IP settings that can be changed when the system is part of a cluster are the additional IPv4 addresses.

The IPv4 address is specific to each peer. It **must** be different for each peer in the cluster.

The IPv4 subnet mask is specific to each peer. It can be different for each peer in the cluster.

The IPv4 gateway is specific to each peer. Each peer can use a different gateway.

Any additional IPv4 addresses added for use with Unified CM must be different for each peer in the cluster.

System host name and domain

The system host name is specific to each peer. We recommend that it is different for each peer in the cluster so that you can easily identify each system.

The DNS domain name is specific to each peer.

DNS servers

DNS servers are specific to each peer. Each peer can use a different set of DNS servers.

Time

The NTP servers are specific to each peer. Each peer may use one or more different NTP servers.

The time zone is specific to each peer. Each peer may have a different local time.

SNMP

SNMP settings are specific to each peer. They can be different for each peer.

Logging

The **Event Log** and **Configuration Log** on each peer will only report activity for the local TelePresence Conductor.

The list of remote syslog servers is specific to each peer. We recommend that you set up a remote syslog server to which the logs of all peers can be sent. This will allow you to have a global view of activity across all peers in the cluster. See Logging configuration [p.172] for further details.

Security certificates

The Trusted CA Certificate and Server Certificate used by the TelePresence Conductor are specific to each peer. They must be uploaded individually on each peer.

Administration access

The SSH service and LCD panel settings are specific to each peer. They can be different for each peer.

Root account password

The password for the root account is specific to each peer. Each peer may have a different password, and for security reasons we recommend that they do.

Note: The username and password for the administrator account is shared across peers.

Locations

All ad hoc or rendezvous IP addresses assigned to Locations must be different for each peer in the cluster.

Creating a new cluster

To create a cluster go to **System > Clustering**.

This page lists the **IP address**es of all the peers in the cluster to which this TelePresence Conductor belongs. It also allows you to set the common **Cluster pre-shared key** used by each peer in the cluster to access all other peers. The inline **Status** shows the current status of each peer.

Prerequisites

Before you create the cluster:

- Ensure that the TelePresence Conductor is running with a valid release key.
  Clustering is not supported on TelePresence Conductors that are running without a release key, as TelePresence Conductor Essentials.
- Note that only two TelePresence Conductor Select systems can be clustered.
  Up to three full capacity TelePresence Conductor peers can be clustered.
- Ensure that you can log in to the web interface of each TelePresence Conductor that is to be added to the cluster, and ensure that they each have the following settings configured as a minimum:
  - IPv4 address
  - IPv4 gateway
  - System host name (recommended so that you can easily differentiate between each peer in the cluster)
  - All systems to be clustered have their time synchronized using an NTP server
- Ensure that the initial TelePresence Conductor cluster peer does not have any unresolved alarms.
- If peers are deployed on different LANs, there must be sufficient connectivity between the networks to ensure a low degree of latency between the peers - a maximum delay of 15ms one way, 30ms round-trip.
- Note that deploying all peers in a cluster on the same LAN means they can be configured with the same routing information such as local domain names and local domain subnet masks.
- We recommend that you create a backup of each system.

Creating a cluster

1. Create a cluster of one peer.
   To do this you must decide which peer is to be the initial peer. The configuration of this peer will be shared with all other peers as they are added to the cluster.
   a. On the initial peer go to **System > Clustering**.
   b. Enter a password that is shared between the cluster peers under **Cluster pre-shared key**.
   c. Enter the initial peer's IP address under **Peer 1 IP address**.
   d. Click **Save**.
   e. Go to **Maintenance > Restart options**.
   f. Click **Restart**.
2. Add the remaining peer(s) to the cluster.
   To do this:
   a. On the initial peer go to **System > Clustering**.
   b. Enter the shared password into the **Cluster pre-shared key** field.
   c. Enter the second peer's IP address into the **Peer 2 IP address** field.
   d. Click **Save**.
   e. Go to **Maintenance > Restart options**.

f. Click **Restart**.
g. On the second peer go to **System Clustering**.
h. Enter the shared password into the **Cluster pre-shared key** field.
i. Enter the initial peer's IP address into the **Peer 1 IP address** field.
j. Enter the second peer's IP address into the **Peer 2 IP address** field.
k. Click **Save**.
l. Go to **Maintenance > Restart options**.
m. Click **Restart**.
n. Repeat the steps if there is a third peer, adding its IP address into the **Peer 3 IP address** field.

3. Update the configuration on the call control devices to accept the new peers.

For detailed information on the required configuration steps see the relevant deployment guide:

- *Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide*

Monitoring the status of the cluster

The inline status areas show you the current status of each peer in the cluster. To check that the cluster is not partitioned, make sure all peers have a status of *Active* on every peer.

Changing a peer's IP address

Note: Do not change the IP address of a peer while it is part of a cluster.

If you want to change the IP address of a peer that is part of an existing cluster, you must perform the following steps, in order:

1. Remove the peer from the cluster. See Removing a peer from an existing cluster [p.162] for instructions.
2. Change the IP address of the peer (go to **System > Network interfaces > IP** and change the entry in the **IPv4 address** field).
3. Re-add the peer to the cluster. See Creating a new cluster [p.159] for instructions.

Removing a peer from an existing cluster

After a cluster has been set up you can remove individual peers from it. When a peer has been removed from the cluster, it will retain the configuration it had at the moment it was removed.

Note: If you want to remove **all** peers from a cluster, see Disbanding a cluster [p.164].

The instructions for removing a peer from a cluster differ depending on the current status of the peer - that is, whether it is live or out-of-service.

Removing a live peer from a cluster

Removing a live peer from a cluster is a two-step process:

1. Placing the peer in standalone mode
2. Removing the peer from the cluster

Each of these steps is described below.

Removing an out-of-service peer from a cluster

If one of the peers in a cluster has become out of service and can no longer be accessed, you do not need to place it in standalone mode. However, you must still follow the instructions in Removing the peer from the cluster [p.163].

Note: If you want to place the out-of-service peer back into the cluster after successfully removing it, you must follow the instructions in Placing the peer in standalone mode [p.162] and Creating a new cluster [p.159].

Placing the peer in standalone mode

Before removing a live peer from a cluster, you must place the peer in standalone mode so that it no longer communicates with other peers in the cluster.

To do this:

1. On the peer to be removed, go to **System > Clustering**.
2. Delete the **Cluster pre-shared key**.
3. Delete all entries from the **Peer IP address** fields.
4. **Save** this configuration.
5. Restart the peer (**Maintenance > Restart**, then click **Restart system**).
6. Delete all entries from the conference bridge pool.
7. If using a Cisco VCS, update the policy service on the Cisco VCS so that it does not include the removed peer.

The peer will no longer consider itself part of the cluster. The peer will become unusable and no calls will go through it. You must now follow the instructions in Removing the peer from the cluster [p.163].

Removing the peer from the cluster

After the peer to be removed has been placed in standalone mode (or if the peer is out of service and cannot be contacted), you must update all other peers in the cluster so they no longer consider the removed peer to be part of their cluster.

To do this, on each remaining peer in the cluster:

1. Go to **System > Clustering**.
2. Delete the **Peer IP address** of the peer that has been removed from the cluster.
3. **Save** this configuration.
4. Repeat these steps for each remaining peer.

The removed peer will no longer be considered part of the cluster.

Disbanding a cluster

When a cluster is disbanded, all peers become standalone systems. They will retain the configuration they had at the moment the cluster was deleted.

Note: if you want to remove a single peer from a cluster without deleting the cluster, see Removing a peer from an existing cluster [p.162].

To delete a cluster, on each peer in the cluster:

1. Go to **System > Clustering**.
2. Delete the **Cluster pre-shared key**.
3. Delete all entries from the **Peer IP address** fields.
4. **Save** this configuration.
5. Restart the peer (**Maintenance > Restart**, then click **Restart system**).
6. Repeat the above steps for every peer in the cluster.

Upgrading a cluster

When the software of one peer in a cluster is upgraded, that peer is unable to communicate with any other peers in the cluster that are not running the same software version. This means that any configuration changes made on one peer in the cluster will not be replicated to any other peers in the cluster that are running a different version of software.

In order to maintain the stability of the cluster, we recommend that you:

- disband the cluster
- upgrade each peer in the cluster one by one, waiting until the upgraded peer is back in service before upgrading the next peer
- do not change any configuration on any peers in the cluster until all peers have been upgraded
- re-create the cluster with the upgraded peers

For detailed instructions on upgrading a cluster, see the relevant deployment guide:

- *Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide*

Note: You should backup the system configuration of each peer before upgrading. For more information, see Cluster backup and restore [p.166].

Cluster backup and restore

The backup and restore process saves all configuration information for a particular TelePresence Conductor.

We recommend that you regularly backup all peers in the cluster. This ensures that peer-specific configuration information (see Peer-specific configuration [p.157]) is saved and can be restored individually for each peer.

Note:

- Do not restore a backup made on one peer to another peer.
- Do not restore a backup made when running one version of software to the same peer running another version of software.

In all other aspects, the process for backing up and restoring peers in a cluster is the same as for standalone systems. For full instructions, see Backing up and restoring data [p.184].

Maintenance

This section provides information on how to perform the TelePresence Conductor maintenance tasks, accessible via the **Maintenance** menu.

Upgrading software components

You can install new releases of the TelePresence Conductor software on your existing hardware. Software upgrades can be performed in one of two ways:

- Using the web interface - this is the recommended process.
- Using secure copy.

This guide describes how both of these methods are used to perform upgrades.

Note: You should read the section Before you upgrade [p.168] prior to upgrading your software.

For information about upgrading peers in a cluster of TelePresence Conductors, see Upgrading a cluster [p.165].

Before you upgrade

To avoid any performance degradation we recommend that you upgrade the TelePresence Conductor while the system is inactive.

For specific information about upgrading peers in a cluster, see the relevant deployment guide:

- *Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide*

Prerequisites

The upgrade requires you to have:

- a valid **Release key**, if you are upgrading the major release of the TelePresence Conductor (for example from XC2.4 to XC3.0.
  A release key is not required for:
  - dot releases (for example XC2.3 to XC2.4)
  - systems that are running without a release key and with limited capacity (as TelePresence Conductor Essentials)

  Note: If you do not supply a valid release key when upgrading the major release, your system will run as TelePresence Conductor Essentials with limited capacity.
- a software image file for the component you want to upgrade, stored in a location that is locally accessible from your client computer.
- release notes for the software version you are upgrading to — additional manual steps may be required.

Backing up before upgrading

Before an upgrade we recommend that you back up your system configuration and/or create a system snapshot (for TelePresence Conductor VMs).

A backup can be created on the **Backup and restore** page (**Maintenance > Backup and restore**).

A snapshot can be created on the **System snapshot** page (**Maintenance > Diagnostics > System snapshot**). It is useful if you run into problems during the upgrade and you need to roll back the software version.

Installing and rebooting

Upgrading the TelePresence Conductor software is a two-stage process.

First, the new software image is uploaded onto the TelePresence Conductor. At the same time, the current configuration of the system is preserved, so that this can be restored after the upgrade. During this initial stage the system will continue running on its existing software version, and all normal system processes will continue.

The second part of the upgrade involves rebooting the system. It is only during the reboot that the TelePresence Conductor installs the new software version and restores the previous configuration.

Note: If a call goes through the TelePresence Conductor's back-to-back user agent, a reboot will cause all active calls to be terminated. If the TelePresence Conductor is using the Cisco VCS's external policy server interface, a reboot will not affect existing conferences; these will be left running. However, if the TelePresence Conductor is not a part of a cluster, while the system is rebooting users will not be able to create new conferences, or join or re-join existing conferences.

This two-stage process means that you can upload the new software to your system at any time, and then wait until a convenient moment to install the new version by rebooting the system.

Any configuration changes made between the start of the upgrade process and the reboot will be lost when the system restarts using the new software version.

Upgrading using the web interface

The **Upgrade** page (**Maintenance > Upgrade**) is used to install newer versions of the TelePresence Conductor software.

Note:

- You should backup your system configuration before upgrading. Click **System backup** to go to the **Backup and restore** page.
- See Before you upgrade [p.168] for full information about the upgrade process, prerequisites and how to backup your system.
- A system upgrade requires a system reboot to complete the process.
  Note: If a call goes through the TelePresence Conductor's back-to-back user agent, a reboot will cause all active calls to be terminated. If the TelePresence Conductor is using the Cisco VCS's external policy server interface, a reboot will not affect existing conferences; these will be left running. However, if the TelePresence Conductor is not a part of a cluster, while the system is rebooting users will not be able to create new conferences, or join or re-join existing conferences.
- For additional information about upgrading peers in a cluster, see Upgrading a cluster [p.165].

To upgrade using the web interface:

1. Review the relevant release notes to see if any special steps are required either before or after installing the software image file.
2. Go to **Maintenance > Upgrade**.
3. Click **Browse** and select the software image file for the software version to which you want to upgrade.
4. Enter the **Release key** if required.
5. Click **Upgrade**.
   The TelePresence Conductor will start loading the file. This may take a few minutes.

6. When the **Upgrade confirmation** page is displayed, check that the expected **New software version** and **Release key** are displayed.
7. Check that the **MD5 hash** and **SHA1 hash** (if available) values displayed on the **Upgrade confirmation** page match the values displayed on the cisco.com page from where you have downloaded the software image file.
8. Click **Continue with upgrade**.
   The **System upgrade** page opens and displays a progress bar while the software installs.
9. When the software has uploaded, the page will display:
   ```
   Software successfully upgraded
   The system needs to be rebooted for the new software to take effect.
   ```
10. Click **Reboot system**.
    If you made any configuration changes between starting the upgrade process and rebooting, those changes will be lost when the system restarts.
    After the reboot is complete you are taken to the **Login** page.
11. If you are using the TelePresence Conductor Provisioning API to provision conferences, e.g. via Cisco TMSPE, further configuration steps may be required. See After you upgrade for details.

The upgrade is now complete. The **Overview** and **Upgrade** pages now show the upgraded software component version numbers.

Upgrading using secure copy (SCP/PSCP)

To upgrade using a secure copy program such as SCP or PSCP (part of the PuTTY free Telnet/SSH package) you need to transfer two files to the TelePresence Conductor:

- A text file containing just the 16-character Release Key. Ensure there is no extraneous white space in this file.
- The file containing the software image.

To transfer these files:

1. Upload the Release Key file using SCP/PSCP to the **/tmp/** folder on the system. The target name must be **release-key**, for example:
   scp release-key root@10.0.0.1:/tmp/release-key.
2. Enter the root password when prompted.

   Note: the Release Key file must be uploaded before the image file.

3. Upload the software image using SCP/PSCP to the **/tmp** folder on the system. The target name must be **/tmp/tandberg-image.tar.gz**, for example:
   scp s42800xc2_3_0.tar.gz root@10.0.0.1:/tmp/tandberg-image.tar.gz
4. Enter the root password when prompted.
   The software installation begins automatically. Wait until the software has installed completely. This should not take more than five minutes.
5. Log in to the TelePresence Conductor and reboot the system. After about five minutes the TelePresence Conductor will be ready to use
6. If you are using the TelePresence Conductor Provisioning API to provision conferences, e.g. via Cisco TMSPE, further configuration steps may be required. See After you upgrade for details.

Note: if you make any further configuration changes before rebooting, those changes will be lost when the system restarts, so you are recommended to reboot your system immediately.

After you upgrade

Optional configuration step if TelePresence Conductor Provisioning API is used

We recommend that any existing provisioned conferences on TelePresence Conductor are re-provisioned after an upgrade to XC3.0.

If the TelePresence Conductor's Provisioning API has been used to provision CMRs (Collaboration Meeting Rooms) using Cisco TMSPE version 1.2 or later, we recommend that you follow these steps:

1. In Cisco TMS, go to **Systems > Provisioning > Users**
2. Click **TelePresence Conductor Settings**
3. Click the icon to *Purge CMRs on TelePresence Conductor* (hover over the icons for the tool tip description)
4. Click **Purge CMRs**
5. Close the **TelePresence Conductor Settings** window
6. Click **Regenerate CMRs** (if the option is grayed out, refresh the page)

Logging configuration

The **Logging configuration** page (**Maintenance > Logging**) lets you enable remote logging by configuring up to four **Remote syslog server**s to which copies of the TelePresence Conductor's Event Log are sent.

About the Event Log

The TelePresence Conductor provides an event logging facility for troubleshooting and auditing purposes. This Event Log is a list of all the events that have occurred on your system since the last upgrade and records information about such things as conference creation and deletion, requests to join a conference, alarms raised, and conference bridge status changes. It may also contain system-level information.

For more information see Event Log [p.152].

Remote logging of events

The Event Log is always stored locally on the TelePresence Conductor. However, it is often convenient to collect copies of all event logs from various systems in a single location. This is referred to as remote logging. This is particularly recommended for peers in a cluster.

- You can configure the TelePresence Conductor to copy event log messages to up to 4 remote syslog servers.
- The syslog server must support the BSD (as defined in *RFC 3164*) or IETF (as defined in *RFC 5424*) syslog protocols.

Configuring a remote syslog server

To enable remote logging, configure the TelePresence Conductor with the IP addresses or Fully Qualified Domain Names (FQDNs) of the **Remote syslog servers** to which the Event Log will be written.

For each server you must also specify the syslog protocol **Mode** to use when sending messages to that server, either *Legacy BSD format*, *IETF syslog format* or *IETF using TLS connection*. Alternatively, choose *Custom* to configure individually the **Transport**, **Port** and **Format** to use.

You can filter the events sent to each remote syslog server by severity. To do this set the **Log Level**.

Note:

- The remote server cannot be another TelePresence Conductor.
- A TelePresence Conductor cannot act as a remote log server for other systems.
- Events are always logged locally (to the Event Log) regardless of whether or not remote logging is enabled.
- If more than one remote syslog server is configured, the same information is sent to each server.
- The TelePresence Conductor uses the following facilities for remote logging. The software components / logs that map to the (local) facilities are emphasized:
  - 0 (kern)
  - 3 (daemon)
  - 16 (local0) Administrator
  - 17 (local1) Config
  - 19 (local3) Apache error
  - 20 (local4) etc/opt/apache2

- 21 (local5) Developer
- 22 (local6) Network

Adding option and release keys

The **Option keys** page (**Maintenance > Option keys**) lists all the existing options currently installed on the TelePresence Conductor, and allows you to add new options. It also displays the currently installed release key for this software version and allows you to enter a new release key.

Options are used to add additional features to the TelePresence Conductor. Option keys can either be valid for a fixed time period or have an unlimited duration.

The **System information** section displays the **Hardware serial number** and summarizes the existing features installed on the TelePresence Conductor. The options that you may see here include:

- Bridge enablement limit
- Call session limit
- Clustering

Adding option keys using the web interface

1. In the **Add option key** field, enter the key that has been provided to you for the option you want to add.
2. Click **Add option**.

Some option keys require that the TelePresence Conductor is restarted before the option key takes effect. In such cases an alarm is raised, which remains in place as a reminder until the system has been restarted.

About the Tools menu

The **Tools** menu contains a number of features that can assist with troubleshooting of your system.

- Check pattern allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.
- Check dial plan allows you to check what will happen when a particular alias is received by the TelePresence Conductor.
- Ping allows you to check that a particular host system is contactable from the TelePresence Conductor and that your network is correctly configured to reach it.
- Traceroute allows you to discover the details of the route taken by a network packet sent from the TelePresence Conductor to a particular destination host system.
- Tracepath allows you to discover the path taken by a network packet sent from the TelePresence Conductor to a particular destination host system.
- DNS lookup [p.178] allows you to check which domain name server (DNS server) is responding to a request for a particular hostname.

Check pattern

The **Check pattern** tool (**Maintenance > Tools > Check pattern**) allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.

Note: only incoming aliases and conference names configured via the web interface use regular expression matching. Incoming aliases and conference names configured via the Provisioning API are matched using exact matches and do not use regular expressions.

For more information about regular expressions, see Regular expression reference [p.206].

When using this tool, what you must enter in each of the fields will depend on what the regular expression you are checking is being used for, as follows:

Field	Input	to check a conference alias...	to check an auto-dialed participant...
Pattern	The string to be checked.	enter the alias that is received by the TelePresence Conductor.	enter the name of the conference to which this auto-dialed participant is to be added.
Regular expression	The regular expression against which the **Pattern** will be compared.	enter the string configured in the **Incoming alias** field.	enter the string configured in the **Conference name match** field.
Replacement	The regular expression replacement string that defines how the **Pattern** will be modified if there is a match.	enter the string configured in the **Conference name** field.	enter the string configured in the **Address** field.

When you have completed the fields, click **Check pattern**. The results of the regular expression will appear, as follows:

Match result	States whether or not there was a successful match.

Replacement result	■ When checking a conference alias, this will be the resulting conference name. ■ When checking an auto-dialed participant, this will be the address that will be dialed by the TelePresence Conductor.

Check dial plan

The **Check dial plan** tool (**Maintenance > Tools > Check dial plan**) allows you to check what will happen when a particular alias is received by the TelePresence Conductor. It checks whether the incoming alias matches any of the conference aliases configured via the web interface or any of the aliases associated with the Collaboration Meeting Rooms (CMRs) configured via the Provisioning API. If it finds a match, the results display information related to the conference.

To use this tool:

1. In the **Alias to check** field, enter the dial string that you want to check, exactly as it will be received by the TelePresence Conductor.
2. Click **Check dial plan**.

A new section will appear showing the results of the check. The results differ slightly depending on whether the dial string matches a regex conference alias or a provisioned CMR alias. The results state if there was no match found. If there was a successful match the results list the following fields:

Alias checked	The alias string as it was entered into the **Alias to check** field.
Successfully matched conference alias name	(Only applicable to regex conference aliases) The name of the regex conference alias that matched the alias being checked. The settings for this conference alias will be used by the TelePresence Conductor to determine how the call will be processed.
Successfully matched CMR alias	(Only applicable to provisioned CMR aliases) The alias string that the checked alias was matched to as it was defined via the Provisioning API. The settings for the corresponding CMR/ConfBundle will be used by the TelePresence Conductor to determine how the call will be processed. The CMR alias may differ from the alias that was entered into the **Alias to check** field. If this is the case, a transform of the SIP domain has taken place.
Alias type	Either *Regex Match Alias* or *Provisioned Direct Match Alias*.
Resulting conference name	The name of the conference that will be created on the conference bridge when this alias is dialed.
Role	The role that will be assigned to a caller dialing in to the conference using this alias. The role can be one of the following: ■ *Participant*: for a Meeting-type conference configured via the TelePresence Conductor web interface ■ *Host*: for a provisioned conference or for a Lecture-type conference configured via the TelePresence Conductor web interface ■ *Guest*: for a provisioned conference or for a Lecture-type conference configured via the TelePresence Conductor web interface

Incoming alias regular expression	(Only applicable to regex conference aliases) The regular expression that was configured in the **Incoming alias** field of the matching conference alias.
Conference name replacement string	(Only applicable to regex conference aliases) The replacement string that was configured in the **Conference name** field of the matching conference alias.
Alias can create conference	For regex conference aliases this is the value set for the **Allow conference to be created** field. For provisioned conferences this is the value set for the `allow_conference_creation` attribute of the `Alias` object. The value can be *True* or *False*. If the value is *True*, the first participant dialing the alias will create the conference. If the value is *False*, the conference must be created via the API call `factory.conferencecreate` or via a second alias that matches to the same conference and has the value set to *True*.
All aliases associated with the same CMR	(Only applicable to provisioned CMR aliases) The list of all aliases that are associated with the CMR that the returned alias is associated with.

Ping

The **Ping** tool (**Maintenance > Tools > Network utilities > Ping**) can be used to assist in troubleshooting system issues.

It allows you to check that a particular host system is contactable and that your network is correctly configured to reach it. It reports details of the time taken for a message to be sent from the TelePresence Conductor to the destination host system.

To use this tool:

1. In the **Host** field, enter the IP address or hostname of the host system you want to try to contact.
2. Click **Ping**.

A new section will appear showing the results of the contact attempt. If successful, it will display the following information:

Host	The hostname and IP address returned by the host system that was queried.
Response time (ms)	The time taken (in ms) for the request to be sent from the TelePresence Conductor to the host system and back again.

Traceroute

The **Traceroute** tool (**Maintenance > Tools > Network utilities > Traceroute**) can be used to assist in troubleshooting system issues.

It allows you to discover the route taken by a network packet sent from the TelePresence Conductor to a particular destination host system. It reports the details of each node along the path, and the time taken for each node to respond to the request.

To use this tool:

1. In the **Host** field, enter the IP address or hostname of the host system to which you want to trace the path.
2. Click **Traceroute**.

A new section will appear with a banner stating the results of the trace, and showing the following information for each node in the path:

TTL	(Time to Live). This is the hop count of the request, showing the sequential number of the node.
Response	This shows the IP address of the node, and the time taken (in ms) to respond to each packet received from the TelePresence Conductor. *** indicates that the node did not respond to the request.

The route taken between the TelePresence Conductor and a particular host may vary for each traceroute request.

Tracepath

The **Tracepath** tool (**Maintenance > Tools > Network utilities > Tracepath**) can be used to assist in troubleshooting system issues.

It allows you to discover the route taken by a network packet sent from the TelePresence Conductor to a particular destination host system.

To use this tool:

1. In the **Host** field, enter the IP address or hostname of the host system to which you want to trace the route.
2. Click **Tracepath**.

A new section will appear with a banner stating the results of the trace, and showing the details of each node along the path, the time taken for each node to respond to the request, and the maximum transmission units (MTU).

The route taken between the TelePresence Conductor and a particular host may vary for each tracepath request.

DNS lookup

The **DNS lookup** tool (**Maintenance > Tools > Network utilities > DNS lookup**) can be used to assist in troubleshooting system issues.

It allows you to query DNS for a supplied hostname and display the results of the query if the lookup was successful.

To use this tool:

1. In the **Host** field, enter either:
   - the name of the host you want to query, or
   - an IPv4 or IPv6 address if you want to perform a reverse DNS lookup
2. In the **Query type** field, select the type of record you want to search for:

(for reverse lookups the **Query type** is ignored - the search automatically looks for PTR records)

Option	Searches for...
All	any type of record
A (IPv4 address)	a record that maps the hostname to the host's IPv4 address
AAAA (IPv6 address)	a record that maps the hostname to the host's IPv6 address
SRV (SIP and H.323 servers)	SRV records (which includes those specific to either H.323 or SIP servers, see below)
NAPTR (Name authority pointer)	a record that rewrites a domain name (into a URI or other domain name for example)

3. Click **Lookup**.

A separate DNS query is performed for each selected **Query type**. The domain that is included within the query sent to DNS depends upon whether the supplied **Host** is fully qualified or not (a fully qualified host name contains at least one "dot"):

- If the supplied **Host** is fully qualified:
  - DNS is queried first for **Host**
  - If the lookup for **Host** fails, then an additional query for **Host.<system_domain>** is performed (where **<system_domain>** is the **Domain name** as configured on the **DNS** page)
- If the supplied **Host** is not fully qualified:
  - DNS is queried first for **Host.<system_domain>**
  - If the lookup for **Host.<system_domain>** fails, then an additional query for **Host** is performed

For SRV record type lookups, multiple DNS queries are performed. An SRV query is made for each of the following _service._protocol combinations:

- _h323ls._udp.<domain>
- _h323rs._udp.<domain>
- _h323cs._tcp.<domain>
- _sips._tcp.<domain>
- _sip._tcp.<domain>
- _sip._udp.<domain>

In each case, as for all other query types, either one or two queries may be performed for a <domain> of either **Host** and/or **Host.<system_domain>**.

Results

A new section will appear showing the results of all of the queries. If successful, it will display the following information:

Query type	The type of query that was sent by the TelePresence Conductor.
Name	The hostname contained in the response to the query.
TTL	The length of time (in seconds) that the results of this query will be cached by the TelePresence Conductor.

Class	**IN** (internet) indicates that the response was a DNS record involving an internet hostname, server or IP address.
Type	The record type contained in the response to the query.
Response	The content of the record received in response to the query for this **Name** and **Type**.

Managing trusted CA certificates

The **Trusted CA certificate** page (**Maintenance > Security certificates > Trusted CA certificate**) allows you to manage the list of certificates for the Certificate Authorities (CAs) trusted by this TelePresence Conductor. When a TLS connection to TelePresence Conductor mandates certificate verification, the certificate presented to the TelePresence Conductor must be signed by a trusted CA in this list and there must be a full chain of trust (intermediate CAs) to the root CA.

- To upload a new file containing one or more CA certificates, **Browse** to the required PEM file and click **Append CA certificate**. This will append any new certificates to the existing list of CA certificates. If you are replacing existing certificates for a particular issuer and subject, you have to manually delete the previous certificates.
- To replace all of the currently uploaded CA certificates with the system's original list of trusted CA certificates, click **Reset to default CA certificate**.
- To view the entire list of currently uploaded trusted CA certificates, click **Show all (decoded)** to view it in a human-readable form, or click **Show all (PEM file)** to view the file in its raw format.
- To view an individual trusted CA certificate, click on **View (decoded)** in the row for the specific CA certificate.
- To delete one or more CA certificates, tick the box(es) next to the relevant CA certificate(s) and click **Delete**.

Managing the TelePresence Conductor's server certificate

The **Server certificate** page (**Maintenance > Security certificates > Server certificate**) is used to manage the TelePresence Conductor's server certificate. This certificate is used to identify the TelePresence Conductor when it communicates with client systems using TLS encryption, and with web browsers over HTTPS. You can:

- view details about the currently loaded certificate
- generate a certificate signing request
- upload a new server certificate

Viewing the currently uploaded certificate

The **Server certificate data** section shows information about the server certificate currently loaded on the TelePresence Conductor.

- To view the currently uploaded server certificate file, click **Show (decoded)** to view it in a human-readable form, or click **Show (PEM file)** to view the file in its raw format.
- To replace the currently uploaded server certificate with the TelePresence Conductor's original certificate, click **Reset to default server certificate**.

Note: Do not allow your server certificate to expire as this may cause other external systems to reject your certificate and prevent the TelePresence Conductor from being able to connect to those systems.

Generating a certificate signing request (CSR)

The TelePresence Conductor can generate server certificate signing requests. This removes the need to use an external mechanism to generate and obtain certificate requests.

To generate a CSR:

1. Go to **Maintenance > Security certificates > Server certificate**.
2. Click **Generate CSR** to go to the **Generate CSR** page.
3. Enter the required properties for the certificate.
   - See Server certificates and clustered systems [p.183] if your TelePresence Conductor is part of a cluster.
   - The certificate request includes automatically the public key that will be used in the certificate, and the client and server authentication Enhanced Key Usage (EKU) extension.
4. Click **Generate CSR**. The system will produce a signing request and an associated private key.
   The private key is stored securely on the TelePresence Conductor and cannot be viewed or downloaded. You must never disclose your private key, not even to the certificate authority.
5. You are returned to the **Server certificate** page. From here you can:
   - **Download** the request to your local file system so that it can be sent to a certificate authority. You are prompted to save the file (the exact wording depends on your browser).
   - **View** the current request.

Note:

- Only one signing request can be in progress at any one time. This is because the TelePresence Conductor has to keep track of the private key file associated with the current request. To discard the current request and start a new request, click **Discard CSR**.
- The certificate signing request storage location changed in XC3.x.
  When you generate a CSR in XC2.x, the application puts **csr.pem** and **privkey_csr.pem** into **/tandberg/persistent/certs**.
  When you generate a CSR in XC3.x, the application puts **csr.pem** and **privkey.pem** into **/tandberg/persistent/certs/generated_csr**.
  If you want to upgrade from XC2.x and have an unsubmitted CSR, then we recommend discarding the CSR before upgrade, and then regenerating the CSR after upgrade.

Uploading a new server certificate

When the signed server certificate is received back from the certificate authority it must be uploaded to the TelePresence Conductor.

The **Upload new certificate** section is used to replace the TelePresence Conductor's current server certificate with a new certificate.

To upload a server certificate:

1. Go to **Maintenance > Security certificates > Server certificate**.
2. Use the **Browse** button in the **Upload new certificate** section to select and upload the **server certificate** PEM file.
3. If you used an external system to generate the Certificate Signing Request (CSR) you must also upload the **server private key** PEM file that was used to encrypt the server certificate. (The private key file will have been automatically generated and stored earlier if the TelePresence Conductor was used to produce the CSR for this server certificate.)
   - The **server private key** PEM file must not be password protected.
   - You cannot upload a server private key if a certificate signing request is in progress.
4. Click **Upload server certificate data**.

Server certificates and clustered systems

When a CSR is generated, a single request and private key combination is generated for that peer only.

If you have a cluster of TelePresence Conductors, you must generate a separate signing request on each peer. Those requests must then be sent to the certificate authority and the returned server certificates uploaded to each relevant peer.

You must ensure that the correct server certificate is uploaded to the appropriate peer, otherwise the stored private key on each peer will not correspond to the uploaded certificate.

Backing up and restoring data

This section provides information on backing up and restoring TelePresence Conductor data.

Backing up and restoring TelePresence Conductor data [p.184] gives information about when to create a backup, the contents of the backup file, and limitations you should be aware of.

Creating a system backup [p.185] describes how to backup TelePresence Conductor data.

Restoring a previous backup [p.185] describes how to restore TelePresence Conductor data.

For extra information about backing up and restoring peers in a cluster, see Cluster backup and restore [p.166].

Backing up and restoring TelePresence Conductor data

The **Backup and restore** page (**Maintenance > Backup and restore**) is used to create and restore backup files of your TelePresence Conductor data.

When to create a backup

You are recommended to create a backup in the following situations:

- before performing an upgrade
- before performing a system restore
- in demonstration and test environments if you want to be able to restore the TelePresence Conductor to a known configuration

Content of the backup file

The data in the backup includes:

- system configuration settings
- clustering configuration
- security certificates
- administrator account details

Log files are not included in the backup files.

Limitations

The following limitations apply:

- Backups can only be restored to a system running the same version of software from which the backup was made.
- You can create a backup on one TelePresence Conductor and restore it to a different TelePresence Conductor, for example if the original system has failed. However, before performing the restore you must install on the new system the same set of option keys that were installed on the old system. If you attempt to restore a backup made on a different TelePresence Conductor, you will receive a warning message, but you will be allowed to continue.

- Do not use backups to copy data between TelePresence Conductors, because system specific information, such as IP addresses, will be duplicated.

Note: We recommend that you take the TelePresence Conductor unit out of service before performing a restore.

For extra information about backing up and restoring peers in a cluster, see the Cluster backup and restore [p.166] section.

Creating a system backup

To create a backup of TelePresence Conductor system data:

1. Go to **Maintenance > Backup and restore**.
2. Optionally, enter an **Encryption password** with which to encrypt the backup file.
   If a password is specified, the same password will be required to restore the file.
3. Click **Create system backup file**.
4. After the backup file has been prepared, a pop-up window appears and prompts you to save the file (the exact wording depends on your browser). The default name is in the format:
   <software version>_<hardware serial number>_<date>_<time>_backup.tar.gz.
   (The file extension is normally **.tar.gz.enc** if an encryption password is specified. However, if you use Internet Explorer to create an encrypted backup file, the filename extension will be **.tar.gz.gz** by default. These different filename extensions have no operational impact; you can create and restore encrypted backup files using any supported browser.)
   The preparation of the system backup file may take several minutes. Do not navigate away from this page while the file is being prepared.
5. Save the file to a designated location.

Log files are not included in the system backup file.

Restoring a previous backup

To restore the TelePresence Conductor to a previous configuration of system data:

1. Go to **Maintenance > Backup and restore**.
2. In the **Restore** section, **Browse** to the backup file containing the configuration you want to restore.
3. In the **Decryption password** field, enter the password that was used to create the backup file, or leave it blank if the backup file was created without a password.
4. Click **Upload system backup file**.
5. The TelePresence Conductor checks the file and takes you to the **Restore confirmation** page.
   - If the backup file is not valid or an incorrect decryption password is entered, you will receive an error message at the top of the **Backup and restore** page.
   - You are shown the current software version and serial number.
6. Read all the warning messages that appear before proceeding with the restore.
7. Click **Continue with system restore** to continue with the restore process.
   This will restart your system, so ensure that there are no active calls.

After the system restarts, you are taken to the **Login** page.

Configuring diagnostic logging

The **Diagnostic logging** tool (**Maintenance > Diagnostics > Diagnostic logging**) can be used to assist in troubleshooting system issues.

It allows you to generate a diagnostic log of system activity over a period of time, and then to download the log so that it can be sent to your Cisco customer support representative. You can also take and subsequently download a tcpdump while logging is in progress.

To use this tool:

1. Go to **Maintenance > Diagnostics > Diagnostic logging**.
2. Optionally, select **Take tcpdump while logging**.
3. Click **Start new log**.
4. (Optional) Enter some **Marker** text and click **Add marker**.
   - The marker facility can be used to add comment text to the log file before certain activities are performed. This helps to subsequently identify the relevant sections in the downloaded diagnostic log file.
   - You can add as many markers as required, at any time while the diagnostic logging is in progress.
   - Marker text is added to the log with a "`DEBUG_MARKER`" tag.
5. Reproduce the system issue you want to trace in the diagnostic log.
6. Click **Stop logging**.
7. Click **Download log** to save the diagnostic log archive to your local file system. You are prompted to save the archive (the exact wording depends on your browser).

The downloaded diagnostic log archive contains the following files:

- loggingsnapshot.txt - containing log messages in response to the activities performed during the logging period
- xconf_dump.txt - containing information about the configuration of the system at the time the logging was started
- xstat_dump.txt - containing information about the status of the system at the time the logging was started
- (if relevant) diagnostic_logging_tcpdump.pcap - containing the packets captured during the logging period

These files can be sent to your Cisco support representative, if you have been requested to do so.

CAUTION: tcpdump files may contain sensitive information. Only send tcpdump files to trusted recipients. Consider encrypting the file before sending it, and also send the decrypt password out-of-band.

Note that:

- Only one diagnostic log can be produced at a time; creating a new diagnostic log will replace any previously produced log.
- The TelePresence Conductor continually logs all system activity to a unified log file. The diagnostic logging facility works by extracting a portion of this unified log. On busy systems the unified log file may become full over time and will discard historic log data so that it can continue logging current activity. This means that all or part of your diagnostic log could be overwritten. The system will warn you if you attempt to download a partial diagnostic log file.
- The diagnostic log will continue logging all system activity until it is stopped, including over multiple login

sessions and system restarts.
- The tcpdump has a maximum file size limit of 50 MB.

Clustered systems

Diagnostic logging can also be used if your TelePresence Conductor is a part of a cluster, however some activities only apply to the "current" peer (the peer to which you are currently logged in to as an administrator):

- The start and stop logging operations are applied to every peer in the cluster, regardless of the current peer.
- The taking a tcpdump operation is applied to every peer in the cluster, regardless of the current peer.
- Each cluster peer maintains its own unified log, and logs activity that occurs only on that peer.
- Marker text is only applied to log of the current peer.
- You can only download the diagnostic log from the current peer.
- To add markers to other peers' logs, or to download diagnostic logs from other peers, you must log in as an administrator to that other peer.

To collect comprehensive information for debugging purposes, we recommend that you extract the diagnostic log for each peer in a cluster.

Creating a system snapshot

The **System snapshot** page (**Maintenance > Diagnostics > System snapshot**) lets you create files that can be used for diagnostic purposes. The files should be sent to your support representative at their request to assist them in troubleshooting issues you may be experiencing.

You can create several types of snapshot file:

- **Status snapshot**: contains the system's current configuration and status settings.
- **Logs snapshot**: contains log file information (including the Event Log, and Configuration Log).
- **Full snapshot**: contains a complete download of all system information. The preparation of this snapshot file may take several minutes to complete and may lead to a drop in system performance while the snapshot is in progress.

To create a system snapshot file:

1. Click one of the snapshot buttons to start the download of the snapshot file. Typically your support representative will tell you which type of snapshot file is required.
   - The snapshot creation process will start. This process runs in the background. If required, you can navigate away from the snapshot page and return to it later to download the generated snapshot file.
   - When the snapshot file has been created, a **Download snapshot** button will appear.
2. Click **Download snapshot**. A pop-up window appears and prompts you to save the file (the exact wording depends on your browser). Select a location from where you can easily send the file to your support representative.

If your TelePresence Conductor is part of a cluster, it is recommended to create a system snapshot file on each individual cluster peer. Not all information is replicated across the cluster peers.

Note: System snapshots contain security-sensitive information. You should ensure that you handle and store them carefully in such a way as to prevent unauthorized access.

Incident reporting

The incident reporting feature of the TelePresence Conductor automatically saves information about critical system issues such as application failures. You can:

- Configure the TelePresence Conductor to send the reports automatically to Cisco customer support
- View the reports from the TelePresence Conductor web interface
- Download and send the reports manually to Cisco (usually at the request of Cisco customer support)

The information contained in these reports can then be used by Cisco customer support to diagnose the cause of the failures. All information gathered during this process will be held in confidence and used by Cisco personnel for the sole purpose of issue diagnosis and problem resolution.

Incident reporting caution: privacy-protected personal data

IN NO EVENT SHOULD PRIVACY-PROTECTED PERSONAL DATA BE INCLUDED IN ANY REPORTS TO CISCO.

Privacy-Protected Personal Data means any information about persons or entities that the Customer receives or derives in any manner from any source that contains any personal information about prospective, former, and existing customers, employees or any other person or entity. Privacy-Protected Personal Data includes, without limitation, names, addresses, telephone numbers, electronic addresses, social security numbers, credit card numbers, customer proprietary network information (as defined under 47 U.S.C. § 222 and its implementing regulations), IP addresses or other handset identifiers, account information, credit information, demographic information, and any other information that, either alone or in combination with other data, could provide information specific to a particular person.

PLEASE BE SURE THAT PRIVACY-PROTECTED PERSONAL DATA IS NOT SENT TO CISCO WHEN THE TELEPRESENCE CONDUCTOR IS CONFIGURED TO AUTOMATICALLY SEND REPORTS.

IF DISCLOSURE OF SUCH INFORMATION CANNOT BE PREVENTED, PLEASE DO NOT USE THE AUTOMATIC CONFIGURATION FEATURE. Instead, copy the data from the Incident detail page and paste it into a text file. You can then edit out any sensitive information before forwarding the file on to Cisco customer support.

Incident reports are always saved locally, and can be viewed via the Incident view page.

Enabling automatic incident reporting

Read the privacy-protected personal data caution before you decide whether to enable automatic incident reporting.

To configure the TelePresence Conductor to send incident reports automatically to Cisco customer support:

1. Go to **Maintenance > Diagnostics > Incident reporting > Configuration**.
2. Set the **Incident reports sending mode** to *On*.
3. Specify the **Incident reports URL** of the web service to which any error reports are to be sent. The default is `https://cc-reports.cisco.com/submitapplicationerror/`.
4. Optional. Specify a **Contact email address** that can be used by Cisco customer support to follow up any error reports.
5. Optional. Specify a **Proxy server** to use for the connection to the incident reporting server.

Use the format (http|https)://address:port/ such as **`http://www.example.com:3128/`**

6. Ensure that **Create core dumps** is *On*; this is the recommended setting as it provides useful diagnostic information.

Note: If the **Incident reports sending mode** is *Off*, incidents will not be sent to any URL but they will still be saved locally and can be viewed and downloaded from the **Incident detail** page.

Sending incident reports manually

Read the privacy-protected personal data caution before you decide whether to send an incident report manually to Cisco.

To send an incident report manually to Cisco customer support:

1. Go to **Maintenance > Diagnostics > Incident reporting > View**.
2. Click on the incident you want to send. You will be taken to the **Incident detail** page.
3. Scroll down to the bottom of the page and click **Download incident report**. You will be given the option to save the file.
4. Save the file in a location from where it can be forwarded to Cisco customer support.

Removing sensitive information from a report

The details in the downloaded incident report are Base64-encoded, so you will not be able to meaningfully view or edit the information within the file.

If you need to edit the report before sending it to Cisco (for example, if you need to remove any potentially sensitive information) you must copy and paste the information from the **Incident detail** page into a text file, and edit the information in that file before sending it to Cisco.

Viewing incident reports

The **Incident view** page (**Maintenance > Diagnostics > Incident reporting > View**) shows a list of all incident reports that have occurred since the TelePresence Conductor was last upgraded. A report is generated for each incident, and the information contained in these reports can then be used by Cisco customer support to diagnose the cause of the failures.

For each report the following information is shown:

Field	Description
Time	The date and time when the incident occurred.
Version	The TelePresence Conductor software version running when the incident occurred.
Build	The internal build number of the TelePresence Conductor software version running when the incident occurred.
State	The current state of the incident: *Pending*: indicates that the incident has been saved locally but not sent. *Sent*: indicates that details of the incident have been sent to the URL specified in the Incident reporting configuration page.

To view the information contained in a particular incident report, click on the report's **Time**. You will be taken to the Incident detail page, from where you can view the report on screen, or download it as an XML file for forwarding manually to Cisco customer support.

Incident report details

The **Incident detail** page (**Maintenance > Diagnostics > Incident reporting > View**, then click on a report's **Time**) shows the information contained in a particular incident report.

This is the information that is sent to the external web service if you have enabled **Incident reports sending mode** (via **Maintenance > Diagnostics > Incident reporting > Configuration**). It is also the same information that is downloaded as a Base64-encoded XML file if you click **Download incident report**.

The information contained in the report is:

Field	Description
Time	The date and time when the incident occurred.
Version	The TelePresence Conductor software version running when the incident occurred.
Build	The internal build number of the TelePresence Conductor software version running when the incident occurred.
Name	The name of the software.
System	The system name (if configured), otherwise the IP address.
Serial number	The hardware serial number.
Process ID	The process ID the TelePresence Conductor application had when the incident occurred.
Release	A true/false flag indicating if this is a release build (rather than a development build).
User name	The name of the person that built this software. This is blank for release builds.
Stack	The trace of the thread of execution that caused the incident.
Debug information	A full trace of the application call stack for all threads and the values of the registers.

CAUTION: for each call stack, the Debug information includes the contents of variables which may contain some sensitive information, for example alias values and IP addresses. If your deployment is such that this information could contain information specific to a particular person, read the caution regarding privacy-protected personal data before you decide whether to enable automatic incident reporting.

Viewing or deleting feedback receivers

TelePresence Conductor offers a feedback interface, which Cisco TMS and Prime Collaboration Manager (PCM) use to register for feedback about the current state of the TelePresence Conductor.

Whenever Cisco TMS or PCM register for feedback from TelePresence Conductor an entry for the Cisco TMS's or PCM's IP address is added to the TelePresence Conductor's database. The TelePresence Conductor will attempt to connect to the configured IP address(es) whenever there is an event, even if the IP address is not reachable anymore.

To view or delete defunct feedback receivers go to **Maintenance > Feedback receivers**. When deleting one or more feedback receiver, select the feedback receiver(s) and click **Delete**.

Restarting, rebooting and shutting down

The **Restart options** page (**Maintenance > Restart options**) allows you to restart, reboot or shut down the TelePresence Conductor without having physical access to the hardware.

CAUTION: do not restart, reboot or shut down the TelePresence Conductor while the red ALM LED on the front of the unit is on. This indicates a hardware fault. Contact your Cisco customer support representative.

Restarting

The restart function shuts down and restarts the TelePresence Conductor application software, but not the operating system or hardware. A restart takes approximately 2 minutes.

A restart is typically required in order for some configuration changes to take effect, or when the system is being added to, or removed from, a cluster. In these cases a system alarm is raised and will remain in place until the system is restarted.

If the TelePresence Conductor is part of a cluster and other peers in the cluster also require a restart, we recommend that you wait until each peer has restarted before restarting the next peer.

Rebooting

The reboot function shuts down and restarts the TelePresence Conductor application software, operating system and hardware. A reboot takes approximately 5 minutes.

Reboots are normally only required after software upgrades and are performed as part of the upgrade process. A reboot may also be required when you are trying to resolve unexpected system errors.

Shutting down

A shutdown is typically required if you want to unplug your unit, prior to maintenance or relocation for example. The system must be shut down before it is unplugged. Avoid uncontrolled shutdowns, in particular the removal of power to the system during normal operation.

After the system has been shut down, the only way it can be restarted (unless it is a virtual appliance) is by pressing the soft power button on the unit itself. You must therefore have physical access to the unit if you want to restart it after it has been shut down.

Effect on conferences

If a call goes through the TelePresence Conductor's back-to-back user agent, any of these restart options will cause all active calls to be terminated.

If the TelePresence Conductor is using the Cisco VCS's external policy server interface, none of the restart options will affect existing conferences; these will be left running. If the TelePresence Conductor is part of a cluster, users can still create, join or re-join conferences while an individual peer is restarting. If it is not part of a cluster, users will not be able to create new conferences, or join or re-join existing conferences.

Restarting, rebooting or shutting down using the web interface

To restart the TelePresence Conductor using the web interface:

1. Go to **Maintenance > Restart options**.
2. Click **Restart**, **Reboot** or **Shutdown** as appropriate and confirm the action.
   Sometimes only one of these options, such as **Restart** for example, may be available. This typically occurs when you access the **Restart options** page after following a link in an alarm or a banner message.

- Restart/reboot: the **Restarting**/**Rebooting** page appears, with an orange bar indicating progress. After the system has successfully restarted or rebooted, you are automatically taken to the **Login** page.
- Shutdown: the **Shutting down** page appears.
  This page remains in place after the system has successfully shut down but any attempts to refresh the page or access the TelePresence Conductor will be unsuccessful.

Developer resources

The TelePresence Conductor includes some features that are intended for the use of Cisco support and development teams only. Do not access these pages unless it is under the advice and supervision of your Cisco support representative.

CAUTION: Incorrect usage of the features on these pages could disrupt system operations, cause performance problems and corrupt the system configuration.

These features are:

- Debugging and system administration tools
- Experimental menu

Debugging and system administration tools

CAUTION: These features are not intended for customer use unless on the advice of a Cisco support representative. Incorrect usage of these features could cause the system operation to become unstable, cause performance problems and cause persistent corruption of system configuration.

You can use a number of debugging and system admin tools to inspect what is happening at a detailed level on a live system, including accessing and modifying configuration data and accessing network traffic.

To access these tools:

1. Open an SSH session.
2. Log in as root.
3. Follow the instructions provided by your Cisco support representative.

Experimental menu

The TelePresence Conductor web interface contains a number of pages that are not intended for use by customers. These pages exist for the use of Cisco support and development teams only. Do not access these pages unless it is under the advice and supervision of your Cisco support representative.

CAUTION: Incorrect usage of the features on these pages could disrupt system operations, cause performance problems and corrupt the system configuration.

To access these pages:

1. Go to `https://<TelePresence Conductor host name or IP address>/setaccess`
   The **Set access** page appears.
2. In the **Access password** field, enter `qwertsys`.
3. Click **Enable access**.

A new top-level **Experimental** menu will appear to the right of the existing menu items.

Reference

This section provides supplementary information regarding the administration of the TelePresence Conductor.

Software version history

XC2.4

Cascading of conferences hosted on TelePresence Servers

This version of the TelePresence Conductor supports cascading of conferences hosted on TelePresence Servers in a similar way to cascading of conferences hosted on TelePresence MCUs. TelePresence Server version 4.0(1.57) or later is required for this to work.

Cascading a conference results in resources being used on a secondary conference bridge when the primary conference bridge does not have enough resources available for all the participants.

The web interface allows you to specify the maximum number of cascades allowed for a conference. This number determines how many resources are reserved on the primary conference bridge purely for creating cascade links to other conference bridges.

If all available resources on the primary conference bridge are used up, a cascade link is created to one or more conference bridges to expand the size of the conference beyond the resource capabilities of the primary conference bridge. The resources used for each cascade link are equivalent to the resources that would be used by one participant receiving 720p/30fps video, stereo audio and the content quality selected on the conference template. These resources are allocated on the primary and on the cascade conference bridge.

Only single screen endpoints are supported on cascade links connecting TelePresence Servers. Therefore, if a multiscreen endpoint joins a conference on a cascade conference bridge, participants on the same cascade bridge will see all screens, whereas participants on the primary bridge and on other cascade bridges will only see one screen (the screen showing the loudest speaker).

Cascade links connecting TelePresence Servers support up to 720p/30fps video. Participants viewing video over a cascade link (that is, video from a participant hosted on a different conference bridge) will see a maximum video quality of 720p/30fps.

Participants on the same conference bridge will see full high quality video if all of the following apply:

- Higher quality video (1080p/30fps or 720p/60fps) has been configured on the TelePresence Conductor's conference template.
- The endpoint of the main displayed participant is providing that high quality video.
- The participants' own endpoint supports high quality video.

Improved conference placement algorithm

The algorithm that determines the conference bridge on which TelePresence Conductor places a new conference has been improved.

Note: We strongly recommend that all conference bridges within a pool have the same capacity, so that conferences can be distributed efficiently across conference bridges. If there are conference bridges with different capacities in the same pool, this may lead to unbalanced conference placement in some scenarios.

Minimum TelePresence Server version alarm has been updated

The minimum recommended version of TelePresence Server software is now 4.0. The alarm that is raised when an older version of TelePresence Server software is used has been updated accordingly.

Alarm raised when no Encryption feature key enabled

A new alarm has been added that is raised when one or more conference bridges used by the TelePresence Conductor do not have the Encryption feature key enabled. The Encryption feature key is required for back-to-back user agent (B2BUA) links and recommended for Policy Service links.

Alarm raised when the same conference bridge has been added to TelePresence Conductor more than once

A new alarm has been added that is raised when the same conference bridge has been added to the TelePresence Conductor more than once. The TelePresence Conductor check whether the conference bridges' serial numbers are identical.

Warning displayed when advanced template parameters on primary and cascade TelePresence MCUs are different

A new warning has been added that is displayed when the advanced parameters on the primary and cascade TelePresence MCUs are configured differently. In a future version of the TelePresence Conductor software, the cascade advanced parameters may be removed.

User interface changes

- The field **Number of cascade ports to reserve** on the **Conference templates** page has been renamed to **Maximum number of cascades**. It is now also applicable to conference templates based on TelePresence Server Service Preferences and the default value has been changed to '0'.
- It is no longer possible to turn off the administrator session timeout (for serial port, HTTPS or SSH) on the TelePresence Conductor. The session timeout must now be within the range of 1 to 65535 minutes.
- The **SNMP** page now displays a **Description** field, which allows you to define a description of the system as viewed by SNMP.

API changes

- A new optional parameter, **`bestEffort`**, has been added to the XML-RPC call **`factory.webex.add`**. This parameter determines whether the TelePresence Conductor should attempt to add a WebEx conference when insufficient resources were reserved at conference creation.
- The default value for the **`reserved_cascades`** attribute of the **`confBundle`** object has been changed to '0'. When the attribute is omitted, it is assumed to be '0', which means that cascading is disabled.

XC2.3

New Capacity Management API

A new API allows management applications (such as Cisco TMS) to obtain information about a conference and its associated resources. The API returns information about the capacity of a conference bridge that will be used for a conference with a given dialed alias.

New Provisioning API

A new API allows management applications (such as Cisco TMS) to provision conferences on TelePresence Conductor. The API allows the client to create a new ConfBundle on the TelePresence Conductor. A ConfBundle consists of information related to a conference and can have a number of aliases and auto-dialed participants associated with it. These aliases and auto-dialed participants are separate from conference aliases and auto-dialed participants configured via the TelePresence Conductor's web interface and can be edited via the Provisioning API only.

Direct match alias lookup

TelePresence Conductor version XC2.3 supports direct match alias lookup for conferences created via the Provisioning API. Allowing direct match alias configuration dramatically reduces the lookup time for tens of thousands of aliases.

In previous versions, conferences configured on TelePresence Conductor used regular expressions (RegEx) to match aliases. This allowed multiple aliases to create conferences with minimum configuration on TelePresence Conductor. However, when thousands of aliases were configured, conference lookup time and create time started to increase, and fine grain control of allowed conferences was difficult.

In version XC2.3 direct match alias lookup is supported only when using the new Provisioning API. All conferences configured directly via the TelePresence Conductor's web interface or XML RPC API continue to use RegEx lookup.

Numeric dialing with Unified CM

Unified CMs append the TelePresence Conductor's IP address (one of the additional IP addresses configured on TelePresence Conductor's user interface) or hostname instead of the domain to numeric dial strings. For example, when an endpoint dials the string **`1234`**, Unified CM will send the dial string **`1234@10.0.0.1`** to TelePresence Conductor. When TelePresence Conductor attempts to do an exact match of the dial string, it will not be able to match the dial string to an alias, because the user will have provisioned an alias that uses a domain, for example **`1234@domain.com`**.

A new API (the SIP Domain API) resolves this issue. You can set the SIP domain on TelePresence Conductor. TelePresence Conductor will transform incoming dial strings to include the SIP domain rather than an IP address or hostname, which will make it possible to compare the dial string with provisioned aliases.

Note that this modification to the URI is only internal to the TelePresence Conductor. The outgoing call URI does not change.

Collaboration meeting room information available on the web interface

You can use the **Collaboration meeting rooms** page to search for one or more Collaboration Meeting Rooms (CMRs) that have been configured via the TelePresence Conductor's Provisioning API using a management tool such as Cisco TMS. For each CMR, details on aliases, auto-dialed participant and other related data can be viewed. The data associated with a CMR is configured via the Provisioning API. It cannot be modified via the TelePresence Conductor's web interface and it cannot be used by conferences configured via the web interface.

Increased number of TelePresence Server calls supported

In older versions of TelePresence Server software it was possible to connect only up to 104 participants in total to either a standalone TelePresence Server or cluster of TelePresence Servers. In TelePresence Server version 4.0 it will be possible to connect up to 200 participants in total (up to 104 per conference). Changes to TelePresence Conductor allow support for these new limits.

Audio-only quality setting added

You can configure audio-only conferences on TelePresence Servers using a new predefined 'Audio-only (no video, mono audio)' conference quality setting.

Segment switching support

TelePresence Server version 4.0 supports segment switching, which allows multiscreen endpoints to switch just the screen of another multiscreen endpoint that contains the loudest speaker rather than all screens. This

feature works only for multiscreen endpoints that provide loudest pane information and for TelePresence Server version 4.0. It is ignored otherwise. The default is to have segment switching enabled.

The feature can be enabled or disabled via:

- TelePresence Conductor's web interface (on the **Conference template** page)
- Cisco TMS's web interface (via the **Custom Parameters** on the **Create new CMR Template** page)
- the **`advanced_parameters`** attribute of the **`ConfBundle`** object when using the TelePresence Conductor's Provisioning API directly

The TelePresence Server must be:

- connected to the TelePresence Conductor
- running version 4.0 or later
- configured in *Remotely managed* mode

H.264 - SVC signaling passthrough

The TelePresence Conductor back-to-back user agent (B2BUA) now passes through all H.264-SVC (scalable video codec) signaling. This will allow endpoints to use the hybrid conference and multistream endpoint support on the TelePresence Server when it is available. A hybrid conference includes some audio and video streams that are switched and some that are transcoded. See *TelePresence Server Release Notes* for more information.

H.265 passthrough

The TelePresence Conductor back-to-back user agent (B2BUA) now passes through all H.265 signaling. Like H.264 - SVC signaling this will allow endpoints to use the hybrid conference and multistream endpoint support on the TelePresence Server when it is available.

Encrypted iX passthrough

Previously the iX protocol, used for example to support the ActiveControl feature on the TelePresence Server, was passed through the TelePresence Conductor's B2BUA. Now the B2BUA also allows Encrypted iX to pass through.

SIP Remote Party ID (RPID) passthrough

The TelePresence Conductor's B2BUA now supports the passthrough of SIP RPID, which is a SIP header used by Unified CM to convey calling and connected line identity. SIP RPID is defined in the document draft-ietf-sip-privacy-04. Although RPID is non-standard, it is implemented by a large number of vendors and is included in most of Cisco's SIP products.

TelePresence Conductor forwards the SIP RPID header without checking the validity of the identity information contained in the header or the authority of the source. To indicate this, TelePresence Conductor sets the **`screen`** parameter to *no*.

TelePresence Conductor does not support the ability to set the display name field to 'anonymous' for endpoints requesting anonymity. All display names are forwarded on as they are.

Secure conference configuration passthrough

Previously when configuring a conference on Cisco TMS, the secure conference parameter was ignored. This has now changed and the configuration is passed on to the conference bridges.

Certificate management

- The management of CA certificates has been improved, allowing you to view, upload and delete individual CA certificates.
- New installations of TelePresence Conductor software now ship with a temporary trusted CA, and a server certificate issued by that temporary CA. We strongly recommend that you replace the server certificate with one generated by a trusted certificate authority, and that you install CA certificates for the authorities that you trust. When you upgrade to this release from an earlier installation of TelePresence Conductor software, your existing server and trusted CA certificates are retained, and will not be affected by this feature.

Other enhancements and usability improvements

- The online help has a new skin and an improved search capability.
- When configuring firewall rules:
  - You can choose whether to drop or reject denied traffic. On upgrade to XC2.3 or later, any existing "deny" rules will now drop the traffic; prior to XC2.3 the traffic would have been rejected.
  - If you have made several changes there is now an option to revert all changes. This discards all pending changes and resets the working copy of the rules to match the current active rules.
  - You can more easily change the order of the rules by using up/down arrow buttons to swap the priorities of adjacent rules.
- Improved web interface usability when switching between SRV and address record resolution modes when configuring the address of an LDAP server for remote user account authentication.
- You have the option to take a tcpdump while diagnostic logging is in progress.
- The diagnostic logging feature has been extended to include:
  - a tcpdump that can be enabled cluster-wide
  - an xconfig file
  - an xstatus file
  - an indication on the web administration page of which user / IP address initiated the logging

  The xconfig and xstatus files are taken at the start of the logging process.
- It is now possible to view all matching intrusion protection triggers for a particular category.

XC2.2.1

New option key supporting up to 50 concurrent call sessions

From Cisco TelePresence Conductor version XC2.2.1 a new option key is available for the Virtual Machine TelePresence Conductor which supports up to 50 concurrent call sessions. This option key can be obtained from your Cisco representative. It provides access to Cisco TAC (Technical Assistance Center) support and is suitable for small and medium-sized deployments.

XC2.2

Improved TIP-compliant endpoint support

Multiscreen endpoints that are compliant with the TelePresence Interoperability Protocol (TIP) do not need to be pre-configured any longer. The TelePresence Conductor is now able to retrieve the number of screens and the associated resources that are required on the conference bridge via TIP. The TelePresence Conductor no longer over-allocates resources on the conference bridge for multiscreen endpoints.

However, these improvements are only applicable to deployments where SIP signaling is routed via the TelePresence Conductor; Cisco VCS deployments using the external policy service continue to work as in previous releases. Additionally, the TelePresence Conductor still over-allocates resources initially for reserved chairperson participants and for escalated Unified CM ad hoc conferences.

360p video support

The TelePresence Conductor now supports 360p video for TelePresence Servers running software version 3.1 or later. There is a new pre-defined quality level with 360p video and mono audio defined that can be selected for conference templates and pre-configured endpoint codecs. New quality levels can be added that use 360p video.

Support for TelePresence Server software on new hardware

The TelePresence Conductor now supports new hardware platforms for the TelePresence Server software version 3.1. It supports the platforms Cisco Multiparty Media 310 and Cisco Multiparty Media 320, as well as the Cisco TelePresence Server on Virtual Machine.

Alarm for minimum conference bridge version

The TelePresence Conductor now raises a minimum version alarm when connected to a TelePresence MCU running version 4.3 or lower, and when connected to a TelePresence Server running version 3.0 in remotely managed mode. If a TelePresence Server is running version 2.x and/or is in locally managed mode, TelePresence Conductor will raise an alarm stating that the conference bridge is running in the wrong mode.

These older conference bridge versions do not support some of the new features for XC2.2.

Automated intrusion protection

An automated intrusion protection feature has been added. It can be used to detect and block malicious traffic and to help protect the TelePresence Conductor from dictionary-based attempts to breach login security.

Automated protection should be used in combination with the existing firewall rules feature - use automated protection to temporarily block specific threats and use firewall rules to block permanently a range of known host addresses.

Changes to B2BUA security status handling

The TelePresence Conductor back-to-back user agent (B2BUA) now modifies the conference security status to be unencrypted when the inbound SIP connection is over TCP and the outbound SIP connection is over TLS. For the conference security status to be encrypted, SIP signaling must be encrypted on all call legs.

ActiveControl support

The TelePresence Conductor now allows ActiveControl to be negotiated between an endpoint and TelePresence Servers that support this feature. To operate, ActiveControl must be enabled on a TelePresence Server version 3.1 or later by enabling the iX protocol on the TelePresence Conductor under **Conference templates > Advanced template parameters**. Information about capabilities and limitations of this feature is available in the *Cisco TelePresence Server Release Notes*.

Allow or disallow conference creation

A conference alias can now be configured to either allow or disallow conference creation. If the conference alias is configured to disallow conference creation, participants can only join the conference via that alias if the conference already exists. The conference can still be created via the API or by dialing a different conference alias defined for the same conference.

New XML-RPC API parameter added for maximum conference duration

A new XML-RPC parameter has been added to the TelePresence Conductor API that allows API users to override the maximum conference duration configured on the conference template with a lower value.

New XML-RPC API parameter added for number of endpoint screens

A new XML-RPC parameter has been added to the TelePresence Conductor API that allows API users to override the number of endpoint screens configured on the conference template with a lower value. This parameter is only applicable to conference templates that:

- point to a Service Preference containing TelePresence Server pools
- have **Allow multiscreen** set to *Yes*
- have a **Maximum screens** value that is greater than the value specified in the API parameter

Firewall rules configuration

When configuring the firewall rules priority, it is now easier to change the order of the rules by using up/down arrow buttons to swap the priorities of adjacent rules.

Managing trusted CA certificates

The TelePresence Conductor's server and trusted CA certificates can now be viewed in either a human-readable, decoded format, or in raw PEM format.

Administrator authentication source

When configuring the source for administrator account authentication, the *Remote* option is now labeled as *Remote only*. You can no longer access the TelePresence Conductor via a locally configured admin account if a *Remote only* authentication source is in use.

The *Local* option has also been renamed to *Local only*.

Improved user interface

Changes have been made to improve the TelePresence Conductor user interface.

XC2.1

Limited system capacity when running without a release key

The TelePresence Conductor can be run without a release key. In this mode the system capacity is limited; only a single un-clustered conference bridge can be enabled and the TelePresence Conductor cannot be clustered.

Where the TelePresence Conductor has no release key, only "community support" is available. This is a self / collaborative support effort, using technical forums like https://supportforums.cisco.com/community/netpro/collaboration-voice-video/telepresence. TAC support is only available for TelePresence Conductors that have a release key; for further details see https://www.cisco.com/web/services/portfolio/product-technical-support/index.html.

For deployments in production environments we recommend that customers upgrade to a fully licensed installation of TelePresence Conductor.

XC2.0

Cisco TelePresence Server support

A new conference bridge type of Cisco TelePresence Server is supported in this release of the TelePresence Conductor. Conference bridge pools can now be made up of either TelePresence Servers or TelePresence MCUs.

Cisco Unified Communications Manager support

The TelePresence Conductor now supports direct connection to Cisco Unified Communications Manager for ad hoc and rendezvous calls. Endpoints can be registered with either Unified CM or Cisco VCS and call into the same conference.

Addition of multiple IP addresses

Multiple IP addresses can be added on TelePresence Conductor. A different IP address is needed on the TelePresence Conductor for each ad hoc Unified CM location and each rendezvous Unified CM location. This allows the TelePresence Conductor to mimic Unified CM's expectation that it is connecting to separate conference bridges in each location, for ad hoc and rendezvous calls.

Known and unknown multiscreen endpoint support for TelePresence Server conferences

The TelePresence Conductor supports endpoints with more than one screen in conferences hosted on TelePresence Servers.

Cisco TelePresence System (CTS) series endpoints, including Cisco TelePresence System T3, can be pre-configured, in which case they will be allocated the resources defined for the endpoint, or supported without pre-configuration, in which case they will be allocated the resources defined for the conference template.

Other customized multiscreen endpoints have to be pre-configured if sufficient resources are to be allocated on the TelePresence Servers used in the relevant conferences.

Support for third-party and customized multiscreen TelePresence systems (i.e. those other than CTS3xxx, TX9000 or T3) require the optional third-party interop key on the TelePresence Server.

Resource optimization

Resource optimization allows resources that are initially over-allocated on a TelePresence Server to be recovered and re-allocated for other participants, allowing more participants to be handled by a single TelePresence Server.

XML-RPC API support to communicate between Cisco TMS and TelePresence Server 3.0

The TelePresence Conductor API now has support added to translate information being sent between Cisco TelePresence Management Suite and TelePresence Server version 3.0 running in 'Remotely managed' mode. See *Cisco TelePresence Management Suite Release Notes* for information on when this support has been added to the Cisco TMS.

Improvements to logging

Filtered event logs can now be downloaded from the UI.

It is possible to specify the remote syslog server mode as one of the following:

- Legacy BSD format
- IETF syslog format

- IETF syslog using TLS connection
- Custom

The Configuration Log page provides a list of all changes to the TelePresence Conductor configuration, providing users with an audit trail of the TelePresence Conductor configuration.

System Administration session timeout and limits

It is now possible to set a session time out, as well as limits for concurrent sessions and concurrent logins per administrator account for web, SSH and serial sessions.

Certificate signing request (CSR)

The TelePresence Conductor can now generate server certificate signing requests, which removes the need to use an external mechanism to generate and obtain certificate requests.

Firewall rules

Firewall rules can now be added to the TelePresence Conductor, which provide the ability to configure IP table rules to control access to the TelePresence Conductor at the IP level.

Addition of multiple administrator accounts

It is now possible to add multiple administrator accounts with pre-determined access level settings.

Other changes and improvements

Improvements have been made to the TelePresence Conductor web interface.

Regular expression reference

This section provides the following information about regular expressions:

- About regular expressions [p.206] provides a table of common regular expressions.
- Regular expression examples - conference aliases [p.207] shows how to use regular expressions to achieve some basic functions such as prefix and suffix matching, stripping, and replacing. These can be used when configuring conference aliases for either Meetings or Lectures.
- Regular expression examples - Lectures builds on the examples already given to show how to use regular expressions when configuring conference aliases for Lectures.
- Regular expression examples - auto-dialed participants shows how to use regular expressions to match a range of conference names and convert each into an address to be dialed.

Note: The **Check pattern** page (**Maintenance > Tools >Check pattern**) allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.

About regular expressions

Regular expressions can be used when Creating and editing conference aliases [p.104] and when Creating and editing auto-dialed participants [p.106].

With conference aliases, regular expressions can be used to specify a pattern for the **Incoming alias**, and a replace string can then be used to specify the way in which any alias that matches that pattern is transformed to create the **Conference name**.

With auto-dialed participants, regular expressions can be used to specify a pattern for the **Conference name match**, and a replace string can then be used to specify the way in which any conference name that matches that pattern is transformed to create the **Address** that is to be dialed.

The TelePresence Conductor uses Python format regular expression syntax. The table below provides a list of some commonly used special characters in regular expression syntax. This is only a subset of the full range of expressions available. For a detailed description of regular expression syntax see the publication *Regular Expression Pocket Reference*.

On the TelePresence Conductor, regular expressions are compared with the string being matched as a whole line. If the string includes other characters that follow after the matched characters, this will not be considered a match. For example, `meet.ben` will match `meet.ben` but not `meet.benjamin`.

For examples of regular expression usage with the TelePresence Conductor, see:

- Regular expression examples - conference aliases [p.207]
- Regular expression examples - Lectures
- Regular expression examples - auto-dialed participants

Common regular expressions

Character	Description	Example
.	Matches any single character.	

`*`	Matches 0 or more repetitions of the previous match.	`.*` will match against an empty string or any sequence of characters.
`+`	Matches 1 or more repetitions of the previous match.	`.+` will match against any sequence of characters.
`?`	Matches 0 or 1 repetition of the previous match.	`meet\.alice(@example\.com)?` will match either `meet.alice` or `meet.alice@example.com`
`\`	Escapes a regular expression special character.	`\.` will match against a full stop (i.e. `.`) only.
`\d`	Matches any decimal digit, i.e. 0-9.	`\d\d\d` will match any number that is 3 digits long.
`[...]`	Matches a set of characters. Each character in the set can be specified individually, or a range can be specified by giving the first character in the range followed by the `-` character and then the last character in the range. You cannot use special characters within the `[]` - they will be taken literally.	`[a-z]` will match against any lower case alphabetical character. `[a-zA-Z]` will match against any alphabetical character. `[0-9#*]` will match against any single E.164 character - the E.164 character set is made up of the digits 0-9 plus the hash key (#) and the asterisk key (*).
`(...)`	Groups a set of matching characters together. Groups can then be referenced in order using the characters `\1`, `\2`, etc. as part of a replace string.	A regular expression can be constructed to transform a URI containing a user's full name to a URI based on their initials. The regular expression `(.).*_(.).*(@example\.com)` would match against the user `john_smith@example.com` and with a replace string of `\1\2\3` would transform it to `js@example.com`.
`|`	Matches against one expression or an alternate expression.	`.*@example\.(net|com)` will match against any URI for the domain `example.com` or the domain `example.net`.
`^`	Signifies the start of a line. When used immediately after an opening brace, negates the character set inside the brace.	`^meet\.*` will match `meet.alice` but not `alice.meeting` `[^abc]` matches any single character that is NOT one of a, b or c.
`(?!...)`	Negative lookahead. Defines a subexpression that must not be present in order for there to be a match.	`(?!.*@example\.com$).*` will match any string that does not end with `@example.com`. `(?!alice).*` matches any string that does not start with `alice`.
`(?<!...)`	Negative lookbehind. Defines a subexpression that must not be present in order for there to be a match.	`.*(?<!net)` matches any string that does not end with `net`.

Regular expression examples - conference aliases

When configuring a conference alias, you can use a regular expression (regex) in the **Incoming alias** field in combination with a replace string in the **Conference name** fields. This allows you to use advanced pattern

matching and replacing functions of regular expressions.

The examples below show how to use regular expressions to achieve some basic functions such as prefix and suffix matching, stripping, and replacing. These can be used when configuring conference aliases for either Meetings or Lectures.

For specific examples of using regular expressions when configuring conference aliases for Lectures, see Regular expression examples - Lectures [p.210].

The **Check pattern** page (**Maintenance > Tools >Check pattern**) allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.

Note: When configuring conference aliases for Lectures, you **must** ensure that the Conference name for the host alias and the guest alias resolve to the same string. If you do not, they will end up in separate conferences.

Matching a prefix

To allow users to create a conference by dialing a given prefix followed by a string, and use exactly what they dialed as the conference name:

	Incoming alias	Conference name
regex	`(meet\..+)`	`\1`
example	meet.alice	meet.alice

Stripping a prefix

To allow users to create a conference by dialing a given prefix followed by a string, and use what they dialed minus the prefix as the conference name:

	Incoming alias	Conference name
regex	`meet\.(.+)`	`\1`
example	meet.alice	alice

Replacing a prefix

To allow users to create a conference by dialing a given prefix followed by a string, and replace the prefix they dialed with another string to create the conference name:

	Incoming alias	Conference name
regex	`meet\.(.*)`	`666\1`
example	meet.alice	666alice

Matching a suffix

To allow users to create a conference by dialing a string followed by a given suffix, and use exactly what they dialed as the conference name:

	Incoming alias	Conference name
regex	`(.*\.meet)`	`\1`
example	alice.meet	alice.meet

Stripping a suffix

To allow users to create a conference by dialing a string followed by a given suffix, and use what they dialed minus the suffix as the conference name:

	Incoming alias	Conference name
regex	`(.*)\.meet`	`\1`
example	alice.meet	alice

Replacing a suffix

To allow users to create a conference by dialing a string followed by a given suffix, and replace the suffix they dialed with another string to create the conference name:

	Incoming alias	Conference name
regex	`(.*)\.meet`	`\1.meeting`
example	alice.meet	alice.meeting

Adding a prefix or suffix

To allow users to create a conference by dialing a given string, and add another string before or after what they dialed to create the conference name:

	Incoming alias	Conference name
regex	`(meet)\.(.*)`	`conference.\1.\2`
example	meet.alice	conference.meet.alice
regex	`(meet)\.(.*)`	`\1.\2.conference`
example	meet.alice	meet.alice.conference

Matching an alias with or without a domain appended

To allow users to dial the same conference alias from either an H.323 endpoint (which will not append its domain) or a SIP endpoint (which will append its domain):

	Incoming alias	Conference name
regex	`meet\.([^@]*)(@example\.com)?`	`meet.\1`
example	meet.alice@example.com	meet.alice
example	meet.alice	meet.alice

Regular expression examples - Lectures

When configuring a conference alias for Lectures, you can use a regular expression (regex) in the **Incoming alias** field in combination with a replace string in the **Conference name** field for both the host and guest aliases. The conference name for both host and guest must resolve to the same string. Using the **Incoming alias** field allows you to apply advanced pattern matching and replacing functions of regular expressions.

The examples on this page build on the examples given in Regular expression examples - conference aliases [p.207], which show how to use regex to achieve some basic functions such as prefix and suffix matching, stripping, and replacing. These can be used when configuring conference aliases for either Meetings or Lectures.

The **Check pattern** page (**Maintenance > Tools >Check pattern**) allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.

Note: When configuring conference aliases for Lectures, you **must** ensure that the Conference name for the host alias and the guest alias resolve to the same string. If you do not, they will end up in separate conferences.

Stripping a prefix

One of the simplest ways to use regex when configuring conference aliases for Lectures is to allocate different prefixes to the host and guest conference aliases, and then strip the prefix to result in the conference name. Below is an example of how to do this using regular expressions:

	Incoming alias	Conference name
Host regex	`show\.(.*)`	`\1`
example	show.sales.team	sales.team
Guest regex	`watch\.(.*)`	`\1`
example	watch.sales.team	sales.team

Regular expression examples - auto-dialed participants

When configuring an auto-dialed participant, you can use a regular expression (regex) in the **Conference name match** in combination with a replace string in the **Address** field. This allows you to use wildcards and other advanced pattern matching and replacing functions of regular expressions to match a range of conference names and convert each into an address to be dialed.

The examples below show how to use regular expressions to achieve some basic functions such as matching a conference name, and adding and replacing prefixes.

The **Check pattern** page (**Maintenance > Tools >Check pattern**) allows you to check whether a regular expression you intend to use when configuring a conference alias or auto-dialed participant on the TelePresence Conductor will have the expected result.

Note: If you have used regular expressions when creating conference aliases, the conference names that are being generated will vary depending on the incoming alias. You must therefore ensure that any regex you use to match against potential conference names will cover all possible outcomes of the regex that was used to generate the conference name.

Adding a prefix to all conference names

In this example we want to record all conferences. Our dial plan is set up so that any calls to addresses that start with **`record.`** are routed to our recording device.

We create an auto-dialed participant that matches any conference name and adds the prefix **`record.`** as follows:

Field	Input	Explanation	Example
Conference name match	(.*)	This regex is the default for this field. It will match against all possible conference names. This will result in the **Address** always being dialed for any conferences created using the specified **Conference template**.	sales_ meeting
Address	record\.\1	This replace string will result in the conference name being prefixed with **`record.`** to create the address to be dialed.	record.sales_ meeting

Matching a prefix

In this example, we want to record all **all hands** conferences.

Our dial plan is set up so that:

- these conferences all begin with **allhands.**, for example **allhands.sales** and **allhands.operations**.
- any calls to addresses that start with **`record.`** are routed to our recording device.

We set up an auto-dialed participant that matches any conference name that starts with **allhands**, and replaces that prefix with **`record.`** as follows:

Field	Input	Explanation	Example
Conference name match	allhands\.(.*)	This regex will match against any conference names beginning with **`allhands.`**.	allhands.sales
Address	record\.\1	This replace string will result in **`allhands.`** being removed from the conference name and replaced with **`record.`** to create the address to be dialed.	record.sales

Combining the use of regular expressions in conference aliases and auto-dialed participants

The following example shows how you can combine the use of regular expressions when creating conference aliases and auto-dialed participants. In this example, our dial plan is configured so that:

- all conference aliases for Meetings start with **meet.**
- all users have a FindMe ID in the format **name.findme@domain.com**

We set up the TelePresence Conductor so that whenever anyone creates a conference based on a user's name (e.g. **meet.alice.findme@domain.com**), that user will automatically be dialed in to the conference via their FindMe ID (e.g. **alice.findme@domain.com**).

Step 1 - create a template

On the **Conference templates** page (**Conference configuration > Conference templates**, then click **New**):

Field	Input	Explanation
Name	Meeting with FindMe	This template will be used whenever we want to automatically dial in a user's FindMe ID.
Description	template to route Meetings to FindMe IDs	Descriptions are useful when you are managing a number of templates.
Conference type	*Meeting*	We want this template to apply to Meetings only.

The rest of the settings on this page will depend on your network configuration.

Step 2 - create a conference alias

On the **Conference aliases** page (**Conference configuration > Conference aliases**, then click **New**):

Field	Input	Explanation	Example
Incoming alias (can use regular expression)	meet\.(.*\.findme@.*)	This regex will match any conference alias that begins with **meet.** and ends in **.findme@** followed by any domain name.	meet.alice.findme@domain.com
Conference name (can use regular expression)	\1	This replace string will result in **meet.** being removed from the conference alias to create the conference name.	alice.findme@domain.com

The rest of the settings on this page will depend on your network configuration.

Step 3 - create an auto-dialed participant

On the **Auto-dialed participants** page (**Conference configuration > Auto-dialed participants**, then click **New**):

Field	Input	Explanation	Example
Name	FindMe user		
Description	dials user's FindMe ID in to Meeting	Descriptions are useful if you are managing a number of auto-dialed participants.	
Conference template	*meeting with FindMe*	This is the name of the template that we created in Step 1.	
Conference name match	(.*\.findme@.*)	This regex will match against all possible conference names that include **.findme@** followed by any domain name.	alice.findme@domain.com
Address	\1	This replace string takes the conference name to create the address to be dialed.	alice.findme@domain.com

The rest of the settings on this page will depend on your network configuration.

Other examples that apply with this configuration:

- if a user dials **meet.bob.findme@domain.com** then **bob.findme@domain.com** will be dialed in to the resulting conference
- if a user dials **meet.carol.jones.findme@domain.com** then **carol.jones.findme@domain.com** will be dialed in to the resulting conference

Conference layouts

When creating a conference template or an auto-dialed participant, one of the parameters that can be set and passed to the conference bridge is the conference layout. Below are three tables displaying the layouts that can be selected for TelePresence MCUs and TelePresence Servers.

TelePresence MCU layouts

Layout families

The <*index*> values for **`family<index>`** correspond to the following pane arrangements:

index	Example layouts
1	
2	
3	
4	
5	

Specific layouts

The <*index*> values for **`layout<index>`** correspond to the following pane arrangements:

index	Layout	index	Layout	index	Layout	index	Layout
1		16		31		46	
2		17		32		47	
3		18		33		48	
4		19		34		49	

index	Layout	index	Layout	index	Layout	index	Layout
5		20		35		50	
6		21		36		51	
7		22		37		52	
8		23		38		53	
9		24		39		54	
10		25		40		55	
11		26		41		56	
12		27		42		57	
13		28		43		58	
14		29		44		59	
15		30		45			

TelePresence Server layouts

There are four different layouts on the TelePresence Server, all of which are supported on single-screen endpoints and two of which are supported on multiscreen endpoints.

Layout	Name	Description
	Single	**On a single-screen endpoint**: the active speaker is shown in one full-screen pane. **On a multiscreen endpoint**: all screens of the endpoint with the active speaker are shown full-screen.

Layout	Name	Description
	ActivePresence	**On a single-screen endpoint**: the active speaker is shown in a large pane with additional participants appearing in up to nine PIPs (picture-in-pictures) overlaid at the bottom of the screen. **On a multiscreen endpoint**: all screens of the endpoint with the active speaker are shown full-screen, additional participants appear in up to nine PIPs (picture-in-pictures) overlaid at the bottom of the screen.
	Prominent	**On a single-screen endpoint**: the active speaker is shown in a large pane with additional participants appearing in up to four smaller panes at the bottom of the screen.
	Equal	**On a single-screen endpoint**: conference participants are shown in a grid pattern of equal sized panes, up to 4x4.

Port reference

The TelePresence Conductor uses different IP ports and protocols for different services and functions. The table below lists each of these services and functions. For each, it shows the default port(s) and protocol used and whether these ports are used for inbound or outbound communications.

Service/ function	Description	Local port	Remote port	Protocol	Direction
SSH	Used for encrypted command line administration.	22	TCP port from the ephemeral range	TCP	inbound
NTP	Used for updating the system time.	UDP port from the ephemeral range	123	UDP	outbound
SNMPv2	Used for network management.	161	UDP port from the ephemeral range	UDP	inbound
SNMPv3	Used for network management	162	TCP port from the ephemeral range	TCP	inbound
HTTP	Used for web administration (diverts to HTTPS at port 443)	80	TCP port from the ephemeral range	TCP	inbound
HTTP	Used for unencrypted communication with the conference bridge	TCP port from the ephemeral range	80	TCP	outbound
HTTPS	Used for: ■ encrypted web administration ■ CPL requests from the Cisco VCS ■ XML-RPC API requests from Unified CM for ad hoc calls ■ XML-RPC and REST API requests from external systems	443	TCP port from the ephemeral range	TCP	inbound
HTTPS	Used for: ■ encrypted communication with the conference bridge ■ communication with the conference bridge's XML-RPC API	TCP port from the ephemeral range	443	TCP	outbound
SIP	Used for making SIP calls to the conference bridges via the TelePresence Conductor's B2BUA.	5060 or 5061	5060 or 5061	TCP	inbound and outbound

Service/ function	Description	Local port	Remote port	Protocol	Direction
ISAKMP	Used for IPSec secure communication between cluster peers for PKI (Public Key Infrastructure) key exchange.	500	UDP port from the ephemeral range	UDP	inbound
DNS	Used for sending requests to DNS servers.	TCP port from the ephemeral range	53	TCP	outbound
Syslog	Used to send messages to the remote syslog server.	TCP port from the ephemeral range	514	TCP	outbound
Login authentication	Used to connect to an LDAP server for login account authentication.	TCP port from the ephemeral range	389	TCP	outbound

For information on ports used by TelePresence Conductor clusters see the relevant Clustering Deployment Guide:

- *Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide*
- *Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide*

Event Log reference

This section provides the following reference information about the Event Log:

- Event Log format describes the structure of the Event Log.
- Message details list all the possible elements within the **message_details** field of the Event Log, in the order that they would normally appear, along with a description of each.

Event Log format

The Event Log is displayed in an extension of the UNIX syslog format:

```
date time process_name: message_details
```

where:

Field	Description
date	The local date on which the message was logged.
time	The local time at which the message was logged.
process_name	The name of the program generating the log message. This could include: ■ web for all web login and configuration events ■ conferencefactory.controller ■ conferencefactory.switchboard
message_details	The body of the message (see Message details [p.219] for further information).

Message details

For most messages appearing in the event log, the **message_details** section, which contains the body of the message, consists of a number of human-readable **name=value** pairs, separated by a space.

The first elements within the **message_details** field are always **Level** (where applicable) and **Event**. Following the **Event** is a list of parameters with corresponding values. The last element is always **UTCTime**. The table below describes these elements.

Note: In addition to the events described below, a **syslog.info** event containing the string `MARK` is logged after each hour of inactivity to provide confirmation that logging is still active.

Name	Description
Level	The classification of the event. This could be one of: ■ **ERROR**: A condition has occurred that will affect the performance of the TelePresence Conductor but it will continue to function to some extent. ■ **WARNING**: A condition has occurred that may affect the performance of the TelePresence Conductor but it will continue to function to some extent. ■ **INFO**: Information messages. ■ **DEBUG**: Information that Cisco TAC engineers may use for debugging.
Event	The event which caused the log message to be generated.

Name	Description
Detail	Descriptive detail of the event.
UTCTime	Time the event occurred, using a full UTC timestamp in YYYY-MM-DD HH:MM:SS,SSS format. Using this format permits simple ASCII text sorting/ordering to naturally sort by time. This is included due to the limitations of standard syslog timestamps.

Restoring default configuration

You can restore the TelePresence Conductor to its default factory configuration, with the options of retaining the existing IP configuration and the root and administrator account passwords.

CAUTION: This procedure cannot be reversed and you will lose your current configuration. We recommend that you create a backup of the configuration before restoring the default configuration.

To restore the system to its default configuration:

1. Using SSH or a serial connection, log in to the TelePresence Conductor as **root**.
2. Type **factory-reset**
3. The following text appears:

```
**************************************************************************
Warning! This operation resets the unit to factory default settings!
**************************************************************************
To cancel operation before final confirmation press Ctrl+C
Keep option keys [YES/NO]?
```

4. Follow the prompts on the screen, typing **YES** or **NO** as appropriate to each option.

A description of each of the options is given in the table below.

Option	Description
Keep option keys	The TelePresence Conductor does not currently use option keys, so you can type either **YES** or **NO**.
Keep IP configuration	**YES** retains the system's IPv4 address, subnet mask and gateway.
Keep ssh keys	**YES** retains the system's SSH identity. Do this if you want to be sure that the system is identifying itself to other systems in the same way it did before the factory reset. **Note:** Changes to a system's SSH identity may lead to other systems thinking that its identity is being spoofed.
Keep root and admin passwords	**YES** retains the existing passwords for the root and administrator accounts.
Save log files	**YES** saves the system's log files, including the latest rotation of the Event Log [p.152] and Configuration Log [p.154], to the hard disk where they can be retrieved by Cisco customer support if required.
Replace hard disk	**YES** pauses the factory reset process so that you can replace the system's hard disk. This option should be selected only on advice from Cisco customer support.
Are you sure you want to continue	**YES** starts the factory reset process with the selected options. To abort the factory reset process, type **NO**.

Identifying calls across your network

Call Tags

Call Tags are UUIDs that are used to track calls passing through a network of Cisco TelePresence Conductors and Cisco TelePresence Video Communication Servers (Cisco VCSs). When a Cisco VCS receives a call, it checks to see if there is a Call Tag already assigned to it. If so, the Cisco VCS will use the existing Call Tag; if not, it will assign a new Call Tag to the call. This Call Tag is then included in the call's details when the call is forwarded on to another Cisco VCS or a TelePresence Conductor. A single call passing between two or more Cisco VCSs and TelePresence Conductors can be identified as the same call by use of the Call Tag.

Note: Call Tags are supported by Cisco TelePresence Video Communication Server version X3.0 or later and all versions of Cisco TelePresence Conductor. If a call passes through a system that is not a Cisco VCS or TelePresence Conductor, or a Cisco VCS that is running an earlier version of the software, the Call Tag information will be lost.

Password encryption

All passwords configured on the TelePresence Conductor are stored in encrypted form. This applies to the following, which all have usernames and passwords associated with them:

- the administrator accounts
- outbound connection credentials (used by the TelePresence Conductor when required to authenticate with another system)
- LDAP server (used by the TelePresence Conductor when binding to an LDAP server)

When entering or viewing passwords using the web interface, you will see placeholder characters (e.g. dots or stars, depending on your browser) instead of the characters you are typing.

Maximum length of passwords

When a password is encrypted, it uses more characters than the original plain text version of the password. For each type of password, the maximum number of plain text characters that can be entered and the maximum number of encrypted characters that are displayed through the CLI are shown in the table below.

Password type	Maximum plain text characters	Maximum displayed encrypted characters
Admin account	1024	65
Outbound connection credentials	128	215
LDAP server	60	122

Flash status word reference table

The flash status word is used in diagnosing NTP server synchronization issues.

It is displayed by the **`ntpq`** program **`rv`** command. It comprises a number of bits, coded in hexadecimal as follows:

Code	Tag	Message	Description
0001	TEST1	pkt_dup	duplicate packet
0002	TEST2	pkt_bogus	bogus packet
0004	TEST3	pkt_unsync	server not synchronized
0008	TEST4	pkt_denied	access denied
0010	TEST5	pkt_auth	authentication failure
0020	TEST6	pkt_stratum	invalid leap or stratum
0040	TEST7	pkt_header	header distance exceeded
0080	TEST8	pkt_autokey	Autokey sequence error
0100	TEST9	pkt_crypto	Autokey protocol error
0200	TEST10	peer_stratum	invalid header or stratum
0400	TEST11	peer_dist	distance threshold exceeded
0800	TEST12	peer_loop	synchronization loop
1000	TEST13	peer_unreach	unreachable or nonselect

Alarm categories

The table below lists the possible alarm categories that can be raised on the TelePresence Conductor. Each alarm is identified by a 5-digit **Alarm ID**. The first 2 digits of the **Alarm ID** categorize the alarm as follows:

Alarm ID prefix	Category
10nnn	Hardware issues
15nnn	Software issues
20nnn	Cluster-related issues
25nnn	Network and network services settings
30nnn	Licensing / resources / option keys
35nnn	External applications and services (such as policy services or LDAP/AD configuration)
40nnn	Security issues (such as certificates, passwords or insecure configuration)
50nnn	General TelePresence Conductor configuration issues
55nnn	B2BUA issues

Alarms list

The following table lists the alarms that can be raised on the TelePresence Conductor.

ID	Description	Solution
15001	The application watchdog has recently attempted to recover the system from a system hang - The application watchdog has recently attempted to recover the system from a system hang by restarting this peer. If the problem persists contact your Cisco support representative	Check the Event Log for more details
15002	Application watchdog has given up trying to bring service back - The application watchdog has restarted the system several times to attempt to automatically recover a hung service, and the system has failed to recover. Please contact your Cisco support representative	Reboot the system, check network and clustering configuration and fix any alarms
15003	Application watchdog is disabled - The application watchdog has been disabled. Recovery from a serious error is disabled	Enable the watchdog or reboot
20001	Recovering from network partition - A network partition or system restart has occurred and the number of reachable peers has changed. Recovering data - no service will be available while this is happening	Please wait - this process should take no more than 5 minutes
20002	Peers out of reach - One or more peers are out of reach	Check the address of the peers are correct and ensure your network is fully functional
30008	Invalid release key - The release key is not valid; if you do not have a valid key, contact your Cisco support representative	Add a release key
40003	Conferencing functionality disabled - The root user has the default password set; conferencing functionality is disabled	View instructions on changing the root password
40005	Conferencing functionality disabled - The admin user has the default password set; conferencing functionality is disabled	Change the admin password to enable conferencing functionality
50001	Incompatible role type - A conference alias or auto-dialed participant has a role type that is incompatible with the conference template with which it is associated	Ensure that all conference aliases and auto-dialed participants have a valid role type and template. Check logs to identify records with invalid roles
50002	No conference bridges configured - No conference bridges have been defined. The TelePresence Conductor will not work	Ensure you have created a conference bridge pool, and that it contains at least one conference bridge
50004	No conference templates - No conference templates have been configured. The TelePresence Conductor will not work	Create a conference template

ID	Description	Solution
50005	A conference bridge appears to be in a network partition - One or more of the conference bridges configured is reachable by some, but not all, of the peers. A network partition is probably stopping communication to this conference bridge	Investigate the network fault
50006	One or more conference bridges unusable - One or more conference bridges have a status of 'Unusable'	Check that the address, username and password are correct and that the conference bridge is reachable
50007	Invalid regex found - A conference alias or auto-dialed participant has invalid regex	Ensure that all conference aliases and auto-dialed participants have valid regex. Check logs to identify records with invalid regex
50008	Prohibited TelePresence MCU parameter set - A conference template has an advanced parameter configured that, when sent to the TelePresence MCU, will cause the conference to fail.	Remove the following conference bridge parameters from all conference templates: conferenceName, maximumVideoPorts, reservedVideoPorts, numericId, guestNumericId
50010	Invalid JSON found - A conference template for a TelePresence MCU has invalid JSON	Ensure that all conference templates have valid JSON. Check logs to identify records with invalid JSON
50011	All conference bridges unusable or in 'Busy out' state - All conference bridges either have a status of 'Unusable' or 'Busy out'	Check that the conference bridges' configuration, addresses, usernames and passwords are correct and that the conference bridges are reachable. Check that the conference bridges are administratively enabled.
50012	Discouraged TelePresence MCU parameter set - A conference template has an advanced parameter configured that, when sent to the TelePresence MCU, may cause the conference to fail.	Remove discouraged parameters from all conference templates
50013	Empty conference bridge pool - An active conference bridge pool does not contain any conference bridges	Add at least one conference bridge to each active conference bridge pool
50016	Conflict between conference alias and Call Policy prefix - A conference alias conflicts with the Call Policy prefix	Review your dial plan and change either the conflicting conference alias or the Call Policy prefix
50017	TelePresence MCU resource warning - TelePresence MCU port usage is approaching or has reached full capacity	Add an additional TelePresence MCU
50018	Conference bridge type mismatch - A conference bridge has a conference bridge type that is different from the conference bridge type of the pool to which it belongs	Ensure that all conference bridges in a pool are of the same conference bridge type as the pool itself
50019	Service Preference mismatch - A Service Preference contains conference bridge pools of more than one conference bridge type	Ensure that all pools within a Service Preference are of the same conference bridge type
50020	Unlinked conference bridge address - A conference bridge address is associated with a non-existent conference bridge	Ensure that the conference bridge address is associated with a valid conference bridge

ID	Description	Solution
50021	Unlinked auto-dialed participant - An auto-dialed participant has no conference template associated with it	Add a valid conference template to the unlinked auto-dialed participant
50022	Unlinked alias - A conference alias has no conference template associated with it	Add a valid conference template to the unlinked alias
50023	Invalid regex in conference alias - A conference alias contains invalid regex	Ensure the regex for the conference alias is valid
50024	Invalid Service Preference - A Service Preference has no conference bridge pools associated with it	Add at least one conference bridge pool to the Service Preference
50025	Unlinked conference bridge - A conference bridge is not in any pool	Delete unlinked conference bridge from list of conference bridges
50026	Invalid regex in auto-dialed participant - An auto-dialed participant contains invalid regex	Ensure the regex for the auto-dialed participant is valid
50027	Unlinked Service Preference - A Service Preference has a non-existent conference bridge pool associated with it	Add a valid conference bridge pool to the unlinked Service Preference
50028	Conference bridge has no address linked to it - There has been an error when mapping the conference bridge to its address	Ensure that the conference bridge has an address linked to it
50029	Conflict between conference alias and conference bridge dial plan prefix - A conference alias conflicts with a conference bridge's dial plan prefix	Review your dial plan and change either the conflicting conference alias or the conference bridge's dial plan prefix.
50030	No Service Preferences - No Service Preferences have been defined. The TelePresence Conductor will not work	Create a Service Preference
50031	A pool doesn't belong to a Service Preference - A pool doesn't belong to a Service Preference	Ensure that all pools belong to at least one Service Preference
50032	Too many auto-dialed/reserved participants for the conference template - Too many auto-dialed/reserved participants for the conference template. This conference cannot be created.	Increase the maximum number of participants allowed in this conference or remove some of the auto-dialed/reserved participants associated with this template.
50034	Duplicate conference bridge dial plan prefix - Conference bridge dial plan prefixes must be unique	Ensure that all conference bridge dial plan prefixes are unique
50035	TelePresence MCUs exceed capacity - Unable to create conferences as insufficient TelePresence MCU resource is available	Add an additional TelePresence MCU or reconfigure your conference templates to use less resource
50036	Duplicate conference template - There are conference templates that have a duplicate name	Ensure that all conference templates have unique names
50037	Duplicate conference bridge pool - There are conference bridge pools that have a duplicate name	Ensure that all conference bridge pools have unique names
50038	Duplicate conference alias - There are conference aliases that have a duplicate name	Ensure that all conference aliases have unique names
50039	Duplicate conference bridge Service Preference - There are conference bridge Service Preferences that have a duplicate name	Ensure that all conference bridge Service Preferences have unique names
50041	Duplicate auto dialed participant - There are auto dialed participants that have a duplicate name	Ensure that all auto dialed participants have unique names

ID	Description	Solution
50042	TelePresence Server device resource warning - An individual TelePresence Server's device usage is approaching or has reached full capacity	Add an additional TelePresence Server
50043	TelePresence Servers exceed capacity - Unable to create conferences as insufficient TelePresence Server resource is available	Reconfigure your conference templates to use less resource or add an additional TelePresence Server
50044	Duplicate quality setting - All quality settings must have a unique description	Ensure that all quality settings have a unique description
50045	Unlinked conference template quality - A conference template references an unknown quality	Reconfigure all conference templates to reference an existing quality
50046	Invalid conference template configuration - Provision for multiscreen must be set to 'No' for TelePresence MCUs	Set Provision for multiscreen to 'No' for all conference templates intended for TelePresence MCUs
50047	Unlinked auto-dialed participant quality - An auto-dialed participant references an unknown quality	Reconfigure all auto-dialed participants to reference an existing quality
50049	Invalid auto-dialed participant configuration - Provision for multiscreen must be set to 'No' for TelePresence MCUs	Set Provision for multiscreen to 'No' for all auto-dialed participants intended for TelePresence MCUs
50050	No qualities - No qualities are configured	Configure at least one quality
50051	Invalid auto-dialed participant configuration - TelePresence MCU auto-dialed participants must have a maximum quality of 1 token	Ensure that all TelePresence MCU auto-dialed participants have a maximum quality of 1 token
50052	Invalid conference template configuration - TelePresence MCU conference templates must have a maximum quality of 1 token	Ensure that all TelePresence MCU conference templates have a maximum quality of 1 token
50053	Invalid conference template configuration - All TelePresence Server conference templates must have a valid content quality	Ensure that all TelePresence Server conference templates have a valid content quality
50054	Invalid conference template configuration - All conference templates must have both a host and guest quality set	Ensure that all conference templates have both a host and guest quality set
50055	Invalid auto-dialed participant configuration - All auto-dialed participants must have a maximum quality set	Ensure that all auto-dialed participants have a maximum quality set
50056	Invalid conference template configuration - All conference templates must have a content quality set	Ensure that all conference templates have a content quality set
50057	Invalid conference template configuration - All TelePresence MCU conference templates must have content quality set to 'Off'	Ensure that all TelePresence MCU conference templates have content quality set to 'Off'
50058	Duplicate conference bridge name - All conference bridges must have a unique name	Ensure that all conference bridges have a unique name
50061	Endpoint has no associated codec - An endpoint has been configured without any codecs	Ensure that all endpoints have at least one codec associated with them
50062	Endpoint name not unique - An endpoint has been configured with a name that is not unique	Configure all endpoints have unique names

ID	Description	Solution
50063	Missing endpoint for endpoint codec - An endpoint codec has been configured without an associated endpoint	Ensure that all endpoint codecs are configured against a known endpoint
50064	Endpoint codec order is not unique - An endpoint codec has been configured with an order that is not unique	Ensure that all endpoint codecs associated with an endpoint have a unique order
50065	Unlinked endpoint codec - An endpoint codec has been configured to an unknown endpoint	Ensure that all endpoint codecs are configured against a known endpoint
50066	Conductor does not have exclusive access to conference bridge - A conference bridge is being accessed directly, bypassing Conductor	If the alarm 'Peers out of reach' is raised, address this first. Otherwise, ensure that all users are accessing conferences via Conductor and not directly via the conference bridge.
50067	Conference template is not linked to a Service Preference - A conference template's Service Preference is not found	Configure the conference template to use an existing Service Preference or add a new Service Preference for this conference template
50068	TelePresence MCU pool resource warning - TelePresence MCU pool port usage is approaching or has reached full capacity	Add an additional TelePresence MCU to the pool
50069	'Keep conference alive' setting will be ignored - If the parameter 'lastChairmanLeavesDisconnect' is set to 'Yes' on a conference template, then 'Keep conference alive' will be ignored for any auto-dialed participants	Ensure that 'Keep conference alive' is set to 'No' for any conference templates where the parameter 'lastChairmanLeavesDisconnect' is set to 'true'
50070	TelePresence MCU pool exceeds capacity - Unable to create conferences as insufficient TelePresence MCU resource is available in pool	Add an additional TelePresence MCU to the pool, or reconfigure your conference templates to use less resource
50071	TelePresence Server pool resource warning - TelePresence Server pool resource usage is approaching or has reached full capacity	Add an additional TelePresence Server to the pool
50072	TelePresence Server pool exceeds device capacity - Unable to create conferences as insufficient TelePresence Server device resource is available in pool	Add an additional TelePresence Server to the pool or reconfigure your conference templates associated with the pool to use less resource
50073	Conference bridge address not set - A conference bridge has been configured without an address.	Configure the address of the conference bridge or remove it from the Conductor.
50074	Invalid JSON found - An auto-dialed participant has invalid JSON set	Ensure that all auto-dialed participants have valid JSON. Check logs to identify records with invalid JSON
50075	Invalid auto-dialed participant parameter set, which may cause auto-dialed participant to fail - An auto-dialed participant has an auto-dialed participant parameter configured that, when sent to the conference bridge, may cause the auto-dialed participant to fail	Remove the following parameters from all auto-dialed participants: layoutControlEnabled, linkType, participantProtocol, password

ID	Description	Solution
50076	Invalid auto-dialed participant parameter set, which will cause auto-dialed participant to fail - An auto-dialed participant has an auto-dialed participant parameter configured that, when sent to the conference bridge, will cause the auto-dialed participant to fail	Remove the following parameters from all auto-dialed participants: autoConnect, deferConnection, participantName, conferenceName, address, addAsGuest, autoDisconnect, dtmfSequence
50077	A conference bridge is reporting zero available resource - A conference bridge is reporting zero available resource. This means the conference bridge is not available for use.	Check that the conference bridge is correctly licensed
50078	Quality setting has no audio quality level - A quality setting cannot be used, because there is no audio quality level selected	Ensure that all quality settings have an audio quality level configured
50080	Invalid Location reference - A Location references an unknown conference template	Ensure that all Locations reference a conference template that exists
50082	Conference bridge type mismatch - A pool with the wrong conference bridge type is in a Service Preference	Remove the pool from the Service Preference or change the conference bridge type of the pool and its conference bridges.
50084	Quality setting has no video quality - A quality setting cannot be used, because there is no video quality level selected	Ensure that all quality settings have a video quality level configured
50085	Invalid quality level for audio or video quality - The quality level that has been selected for either audio or video quality is no longer valid	Reconfigure the quality setting to have valid audio and video quality levels
50088	Invalid Location reference - A Location references an unknown IP address	Ensure that all Locations reference known IP addresses
50089	Incomplete Location - A Location with an ad hoc IP address does not have an ad hoc template associated with it	Ensure that all Locations with an ad hoc IP address also have an associated ad hoc template
50090	A Location without a rendezvous IP address - A Location with a conference type of 'Rendezvous' has been configured without a rendezvous IP address	Ensure that all Locations with a conference type of 'Rendezvous' are configured with a rendezvous IP address
50091	A Location without an ad hoc IP address - A Location with a conference type of 'ad hoc' has been configured without an ad hoc IP address	Ensure that all Locations with a conference type of 'ad hoc' are configured with an ad hoc IP address
50092	Conference template with invalid primary advanced template parameters - A conference template has been configured with primary advanced template parameters that relate to an incorrect conference bridge type	Ensure that all conference templates have primary advanced template parameters that match the conference bridge type for the template
50094	A Location with an incorrect template - A Location with a conference type of 'ad hoc' has been configured with a lecture-type ad hoc template	Ensure that all Locations with a conference type of 'ad hoc' are configured with meeting-type ad hoc templates
50095	Conference template with invalid cascade advanced template parameters - A conference template has been configured with cascade advanced template parameters that relate to an incorrect conference bridge type	Ensure that all conference templates have cascade advanced template parameters that match the conference bridge type for the template

ID	Description	Solution
50096	Conference bridge without a conference bridge pool - A conference bridge exists that does not belong to a conference bridge pool	Ensure that all conference bridges belong to a conference bridge pool
50097	Duplicate Location name - Two or more Locations have been configured with the same name	Ensure that all Locations have unique names
50098	Auto-dialed participant with invalid cascade advanced template parameters - An auto-dialed participant references advanced template parameters that relate to an incorrect conference bridge type	Ensure that all auto-dialed participants reference advanced template parameters that match the conference bridge type for the template associated with the auto-dialed participant
50099	TelePresence Server pool exceeds license capacity - Unable to create conferences as insufficient TelePresence Server license resource is available in pool	Add an additional TelePresence Server to the pool, add licenses to TelePresence Servers in the pool, or reconfigure your conference templates associated with the pool to use less resource
50100	Unreferenced conference template - A conference template is not referenced by either an alias or a Location.	Ensure that all conference templates are referenced by either a alias or a Location
50101	TelePresence Server license resource warning - An individual TelePresence Server's license usage is approaching or has reached full capacity	Add an additional TelePresence Server or add licenses to TelePresence Servers in the pool
50102	TelePresence Server's license capacity exceeded - Ad individual TelePresence Server's license resource is insufficient to create a conference	Add an additional TelePresence Server, add licenses to TelePresence Servers in the pool, or reconfigure your conference templates associated with the pool to use less resource
50103	Corrupted internal field - An internal field has become corrupt in its use of letter case. This may prevent joining to conferences.	Check the logs to identify the records where the field(s) have become letter case corrupt. Delete and re-add these records.
50104	TelePresence Server pool license resource warning - TelePresence Server pool license usage is approaching or has reached full capacity	Add an additional TelePresence Server to the pool or add licenses to TelePresence Servers in the pool
50106	Invalid type of Location for conference bridge pool - A conference bridge pool should reference a Location of type 'Rendezvous' or 'Both', not 'Ad hoc'	Change or remove the Location for the associated conference bridge pool
50107	Conference bridge pool has invalid Location - A conference bridge pool has a non-existent Location associated with it	Change or remove the Location for the associated conference bridge pool
50108	Conference template of type 'lecture' is missing an alias - A conference template of type 'lecture' is not referenced by either a 'host' or 'guest' role alias	Add an alias with a 'host' or 'guest' role for the corresponding conference template or remove the template
50109	Incorrectly configured Location - A conference bridge pool references a Location that does not have a trunk IP address	Ensure that all Locations that are referenced by a conference bridge pool have a trunk IP address
50110	Conflicting codec addresses - Two or more pre-configured endpoint codecs have the same IP address or FQDN.	Ensure that the codec IP addresses or FQDN values for all pre-configured endpoint codecs are unique.

ID	Description	Solution
50111	Old TelePresence MCU version configured - A TelePresence MCU is configured with a software version 4.3 or lower; some features are not supported.	Upgrade your TelePresence MCU to the latest software version
50112	Old TelePresence Server version configured - A TelePresence Server is configured with a software version lower than 4.1; some features are not supported.	Upgrade your TelePresence Server to the latest software version
50113	Invalid JSON found - A conference template for a TelePresence Server has invalid JSON.	Ensure that all conference templates have valid JSON. Check the logs to identify records with invalid JSON.
50114	Discouraged TelePresence Server parameter set - A conference template has an advanced parameter configured that, when sent to the TelePresence Server, may cause the conference to fail.	Remove discouraged parameters from all conference templates.
50115	Prohibited TelePresence Server parameter set - A conference template has an advanced parameter configured that, when sent to the TelePresence Server, will cause the conference to fail.	Remove prohibited parameters from all conference templates.
50116	Duplicate conference bridge - The same conference bridge has been added to TelePresence Conductor more than once	Ensure that all conference bridges are unique
50117	Conference bridge missing Encryption feature key - One or more conference bridges are missing the Encryption feature key. This is required for all TelePresence Conductor B2BUA links.	Contact your Cisco support representative.
50118	Call control destination not reachable - At least one call control trunk destination is unreachable.	Consult the Event Log for the call control destination(s) that failed and check the configuration for it on the TelePresence Conductor's Location page.

Related documentation

All documentation for the latest version of the TelePresence Conductor can be found at www.cisco.com.

Title	Link
Cisco TelePresence Conductor Clustering with Cisco Unified CM Deployment Guide	www.cisco.com
Cisco TelePresence Conductor Clustering with Cisco VCS (Policy Service) Deployment Guide	www.cisco.com
Cisco TelePresence Conductor Clustering with Cisco VCS (B2BUA) Deployment Guide	www.cisco.com
Cisco TelePresence Conductor with Cisco Unified Communications Manager Deployment Guide	www.cisco.com
Cisco TelePresence Conductor with Cisco VCS (Policy Service) Deployment Guide	www.cisco.com
Cisco TelePresence Conductor with Cisco VCS (B2BUA) Deployment Guide	www.cisco.com
Cisco TelePresence Conductor Getting Started Guide	www.cisco.com
Cisco TelePresence MCU Online Help	www.cisco.com
Cisco TelePresence MCU API Reference Guide	www.cisco.com
Cisco TelePresence Server API Reference Guide	www.cisco.com
Cisco TelePresence Server Online Help	www.cisco.com
Cisco TelePresence Video Communication Server Administrator Guide	www.cisco.com
Management Information Base for Network Management of TCP/IP-based internets: MIB-II	http://tools.ietf.org/html/rfc1213
Network Time Protocol website	www.ntp.org/
Regular Expression Pocket Reference, ISBN-10: 0596514271, ISBN-13: 978-0596514273	
RFC 3164: The BSD syslog Protocol	http://tools.ietf.org/html/rfc3164
RFC 3261: SIP: Session Initiation Protocol	http://tools.ietf.org/html/rfc3261
What warnings do I get on a Cisco TelePresence MCU that my conference is finishing? (knowledge base article)	www.cisco.com

Glossary

A glossary of TelePresence terms is available at: https://tp-tools-web01.cisco.com/start/glossary/.

Legal notices

Intellectual property rights

This Administrator Guide and the product to which it relates contain information that is proprietary to TANDBERG and its licensors. Information regarding the product is found below in the Copyright notice section.

TANDBERG® is a registered trademark belonging to Tandberg ASA. Other trademarks used in this document are the property of their respective holders. This Guide may be reproduced in its entirety, including all copyright and intellectual property notices, in limited quantities in connection with the use of this product. Except for the limited exception set forth in the previous sentence, no part of this Guide may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronically, mechanically, by photocopying, or otherwise, without the prior written permission of TANDBERG.

COPYRIGHT © TANDBERG

Copyright notice

The product that is covered by this Administrator Guide is protected under copyright, patent, and other intellectual property rights of various jurisdictions.

This product is Copyright © 2013, Tandberg Telecom UK Limited. All rights reserved.

TANDBERG is now part of Cisco. Tandberg Telecom UK Limited is a wholly owned subsidiary of Cisco Systems, Inc.

A list of the conditions of use can be found at:

http://www.cisco.com/en/US/docs/telepresence/infrastructure/conductor/license_info/Cisco_Conductor_EULA.pdf

This product includes copyrighted software licensed from others. A list of the licenses and notices for open source software used in this product can be found at:

http://www.cisco.com/en/US/products/ps11775/products_licensing_information_listing.html

This product includes software developed by Computing Services at Carnegie Mellon University (http://www.cmu.edu/computing/).

This product includes software developed by the University of California, Berkeley and its contributors.

IMPORTANT: USE OF THIS PRODUCT IS SUBJECT IN ALL CASES TO THE COPYRIGHT RIGHTS AND THE TERMS AND CONDITIONS OF USE REFERRED TO ABOVE. USE OF THIS PRODUCT CONSTITUTES AGREEMENT TO SUCH TERMS AND CONDITIONS.

Accessibility notice

Cisco is committed to designing and delivering accessible products and technologies.

The Voluntary Product Accessibility Template (VPAT) for Cisco TelePresence Conductor is available here:

http://www.cisco.com/web/about/responsibility/accessibility/legal_regulatory/vpats.html#telepresence

You can find more information about accessibility here:

www.cisco.com/web/about/responsibility/accessibility/index.html