Cisco Meeting Server

Cisco Meeting Server Release 2.5

Single Combined Server Deployment Guide

November 22, 2019

Contents

What's new

Version	Change
November 22, 2019	Moved Dial Transform overview from API guide to this guide.
November 13, 2019	New version for 2.8. Call capacity information updated.
November 13, 2019	Title updated to add 2.7.
September 30, 2019	Minor correction.
August 15, 2019	Minor correction.
July 19, 2019	Minor correction.
June 24, 2019	Minor corrections.
June 03, 2019	Minor corrections.
March 19, 2019	Added note on need for each participant in dual homed conference to have unique "from:" SIP address.
February 18, 2019	Minor corrections to WebRTC app and Web Bridge section.
January 31, 2019	Clarification added to Streamer component support information.
January 29, 2019	Recording indicator for dual homed confererences section updated.
January 08, 2019	Minor corrections to appendix on scaling deployments
January 07, 2019	Miscellaneous correction
January 02, 2019	New version for 2.5. Added information on new browser support for WebRTC app, and hosting branding files on a local server.
November 21, 2018	Miscellaneous corrections
October 30, 2018	Miscellaneous corrections
October 02, 2018	Added sections on setting recorder resolution and recording indicator for dual homed conferences.

Version	Change
October 01, 2018	New version for 2.4. Announcing removal of H.323 Gateway, SIP Edge, TURN Server, XMPP Server components in a future version of the Meeting Server software.

1 Introduction

The Cisco Meeting Server software can be hosted on specific servers based on Cisco Unified Computing Server (UCS) technology as well as on the Acano X-Series hardware, or on a specification-based VM server. The term Meeting Server is used throughout this document as a generic term to refer to the Cisco Meeting Server 2000, Cisco Meeting Server 1000, specification-based VM hosts and Acano X-series servers .

This guide covers the Meeting Server deployed as a single combined server deployment, the deployment has no scalability or resilience. The server comprises a number of components, see Figure 1.

Note: Although this guides covers the SIP edge and TURN server components within the Meeting Server, customers are encouraged to start planning their transition to using Cisco Expressway at the edge of their network and the Meeting Server in the core of their network. The SIP edge, TURN server, internal Firewall and H.323 gateway components will be removed from the Meeting Server software at some point in the future.

In addition, in the future, the Cisco Meeting WebRTC App and Cisco Jabber will be the supported apps to join Meeting Server hosted conferences, in addition to SIP endpoints, and Lync/Skype for Business clients in dual homed conferences. On withdrawal of the native Cisco Meeting Apps, the XMPP server component will be removed from the Cisco Meeting Server software.

Note: The components greyed out in the figure below will be removed from the Meeting Server software at some point in the future.

Figure 1: Single combined server deployment

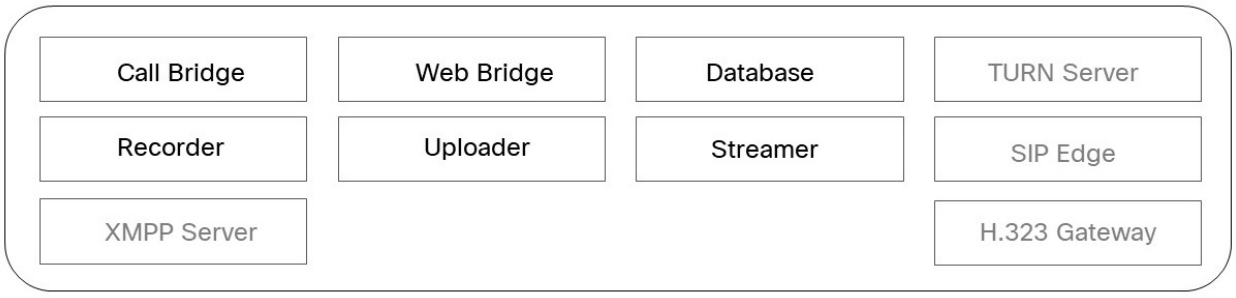

Not all of these components need to be configured, you only need to configure the components that are appropriate to your deployment. This is discussed in Chapter 2.

1.1 Using the Cisco Expressway-E as the edge device in Meeting Server deployments

Over the previous few releases of Cisco Expressway software, edge features have been developed to enable the Cisco Expressway-E to be used as the edge device in Meeting Server deployments. Use the TURN server capabilities in Cisco Expressway-E to connect:

- participants using the WebRTC app to conferences hosted on the Meeting Server,
- remote Lync and Skype for Business clients to conferences hosted on the Meeting Server.

In addition, the Cisco Expressway-E can be used as a SIP Registrar to register SIP endpoints or to proxy registrations to the internal call control platform (Cisco Unified Communications Manager or Cisco Expressway-C).

Table 1 below indicates the configuration documentation that covers setting up Cisco Expressway-E to perform these functions. Table 2 below shows the introduction of the features by release.

Note: Cisco Expressway-E can not be used to connect remote Cisco Meeting App thick clients (Windows/Mac desktop or iOS) to conferences hosted on the Meeting Server. Nor can the Cisco Expressway-E be used between on premises Microsoft infrastructure and the Meeting Server. In deployments with on-premises Microsoft infrastructure and the Meeting Server, the Meeting Server must use the Microsoft Edge server to traverse Microsoft calls into and out of the organization.

Note: If you are configuring dual homed conferencing between on-premises Meeting Server and on-premises Microsoft Skype for Business infrastructure, then the Meeting Server automatically uses the TURN services of the Skype for Business Edge.

Table 1: Documentation covering Cisco Expressway as the edge device for the Meeting Server

Edge feature	Configuration covered in this guide
Connect remote WebRTC apps	Cisco Expressway Web Proxy for Cisco Meeting Server Deployment Guide
Connect remote Lync and Skype for Business clients	Cisco Meeting Server with Cisco Expressway Deployment Guide
SIP Registrar or to proxy registrations to the internal call control platform	Cisco Expressway-E and Expressway-C Basic Configuration (X8.11)

Table 2: Expressway edge support for the Meeting Server

Cisco Expressway-E version	Edge feature	Meeting Server version
X8.11	Supported: - load balancing of clustered Meeting Servers, - Microsoft clients on Lync or Skype for Business infrastructure in other organizations, or Skype for Business clients on Office 365 (not "consumer" versions of Skype). - interoperability between on-premise Microsoft infrastructure and on-premise Meeting Server, **where no Microsoft calls traverse into or out of the organization.** - standards based SIP endpoints. - standards based H.323 endpoints. - Cisco Meeting App thin client (Web RTC app) using TCP port 443. Not supported: - off premise Cisco Meeting App thick clients (Windows/Mac desktop or iOS). - interoperability between on-premise Microsoft infrastructure and on-premise Meeting Server **where Microsoft calls traverse into or out of the organization**, in this scenario, the Meeting Server must use the Microsoft Edge server to traverse Microsoft calls into and out of the organization. See Cisco Meeting Server with Cisco Expressway Deployment Guide (2.4/X8.11.4).	2.4
X8.10	Supported: - Microsoft clients on Lync or Skype for Business infrastructure in other organizations, or Skype for Business clients on Office 365 (not "consumer" versions of Skype), - standards based SIP endpoints, - Cisco Meeting App thin client (Web RTC app) using UDP port 3478 to connect to the Meeting Server via the Expressway reverse web proxy. Not supported: - load balancing of clustered Meeting Servers, - off premise Cisco Meeting App thick clients (Windows/Mac desktop or iOS) or Cisco Meeting App thin client (Web RTC app) using TCP port 443, - interoperability between on premises Microsoft infrastructure and Meeting Server; in this scenario, the Meeting Server must use the Microsoft Edge server to traverse Microsoft calls into and out of the organization. See Cisco Expressway Web Proxy for Cisco Meeting Server	2.3

Cisco Expressway-E version	Edge feature	Meeting Server version
X8.9	Supported: - Microsoft clients on Lync or Skype for Business infrastructure in other organizations, or Skype for Business clients on Office 365 (not "consumer" versions of Skype), - standards based SIP endpoints. Not supported: - load balancing of clustered Meeting Servers,, - off-premise Cisco Meeting App thick clients (Windows/Mac desktop or iOS) and Cisco Meeting App thin client (WebRTC app), - interoperability between on premises Microsoft infrastructure and Meeting Server; in this scenario, the Meeting Server must use the Microsoft Edge server to traverse Microsoft calls into and out of the organization See Cisco Expressway Options with Meeting Server and/or Microsoft Infrastructure	2.2

You are encouraged to migrate your Meeting Server deployments from using the Meeting Server edge components to using the Expressway X8.11 (or later) TURN server. The SIP edge, TURN server, internal Firewall and H.323 gateway components will be removed from the Meeting Server software at some point in the future

1.2 Using the Cisco Expressway-C with the Meeting Server in the core network

In addition to deploying Cisco Expressway-E at the edge of the network, Cisco Expressway-C can be deployed in the core network with the Meeting Server. If deployed between the Meeting Server and an on-premises Microsoft Skype for Business infrastructure, the Cisco Expressway-C can provide IM&P and video integration. In addition the Cisco Expressway-C can provide the following functionality:

- a SIP Registrar,
- an H.323 Gatekeeper,
- call control in Meeting Server deployments with Call Bridge groups configured to load balance conferences across Meeting Server nodes.

Table 3:Additional documentation covering Cisco Expressway-C and the Meeting Server

Feature	Configuration covered in this guide
Call control device to load balance clustered Meeting Servers	Cisco Meeting Server 2.4+, Load Balancing Calls Across Cisco Meeting Servers
SIP Registrar	Cisco Expressway-E and Expressway-C Basic Configuration (X8.11)

Feature	Configuration covered in this guide
H.323 Gatekeeper	Cisco Expressway-E and Expressway-C Basic Configuration (X8.11)

1.2.1 Supported deployments

Figure 2 and Figure 3 illustrate recommended Meeting Server deployments.

Both deployments show an Expressway pair (Expressway-C and Expressway-E) being used as the edge device for the Meeting Server. The Expressway-E is located in the DMZ, while the Expressway-C is located in the internal network between the Meeting Server and Cisco Unified Communications Manager.

The Cisco Meeting WebRTC App can connect via the TURN server on the Expressway-E , or via the TURN server on the Meeting Server if already configured for native Cisco Meeting Apps (Windows, Mac and iOS). Native Cisco Meeting Apps need to connect via the XMPP/Load Balancer components of the Meeting Server, as the Expressway does not support XMPP.

Figure 3 illustrates Microsoft infrastructure added to the deployment to support dual homed conferencing.

Figure 2: Cisco Unified Communications Manager-centric deployment example

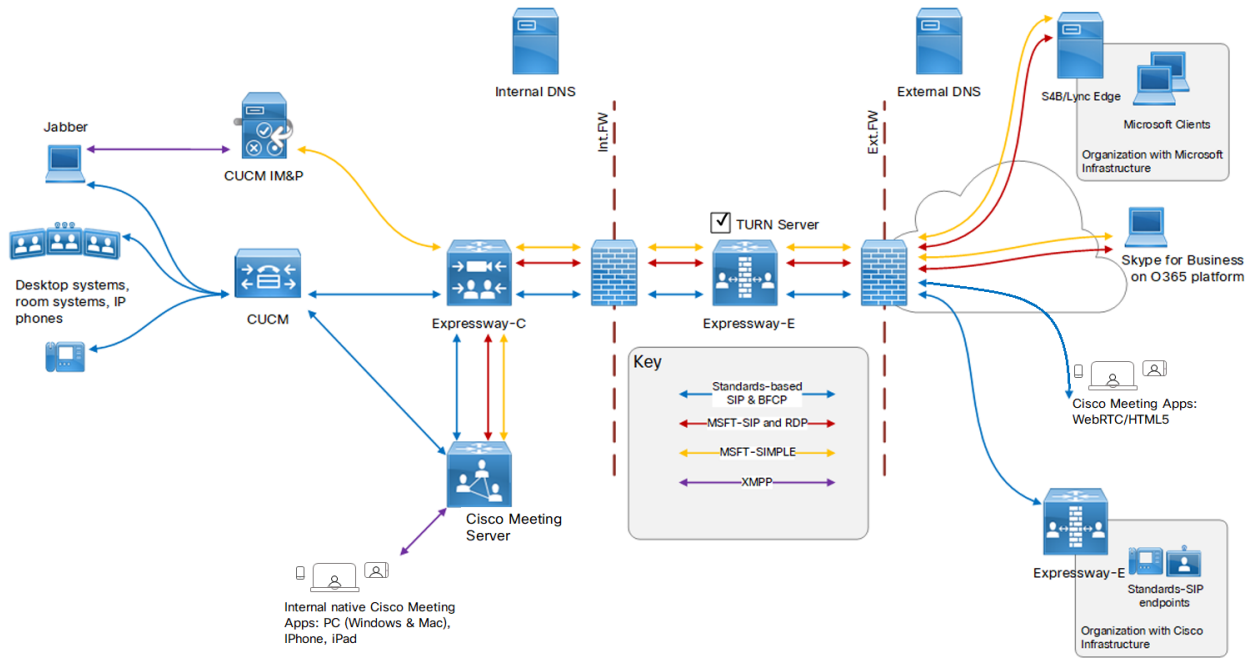

Figure 3: Cisco and Microsoft Infrastructure on-premises deployment example

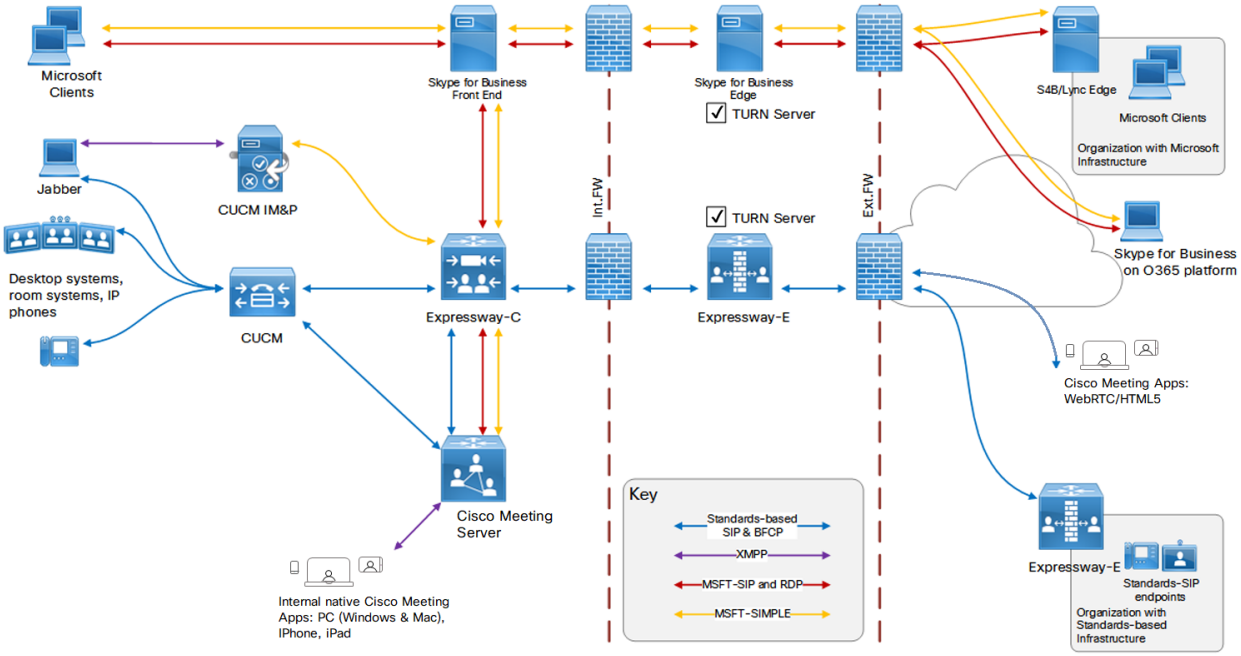

1.2.2 Using the Cisco Expressway H.323 gateway component

In line with Cisco's goal of a single Edge solution across the Cisco Meeting Server and Cisco Expressway, Cisco plans to end of life the Meeting Server H.323 Gateway component. From version 2.4 of the Meeting Server software, there will be no further bug fixes for the H.323 Gateway component. The H.323 component will be removed from the Meeting Server software in a future release. Customers are encouraged to start evaluation of the more mature H.323 Gateway component in the Cisco Expressway, and plan their migration over.

Any H.323 endpoints registered to Expressway-E or Expressway-C will not consume Rich Media Session (RMS) licenses when calling into the Cisco Meeting Server from Expressway version X8.10 onwards.

1.3 How to use this guide

This deployment guide follows on from the appropriate Installation Guide for your server, and assumes that you have completed the installation instructions already. This guide should be read and used in conjunction with the appropriate Certificate Guidelines.

In addition to this deployment guide and the Certificate Guidelines, the reference material shown in the figure below can be found on the Cisco Meeting Server documentation page.

Note: Throughout this guide, the term coSpace has been renamed space.

Figure 4: Overview of guides covering the Meeting Server

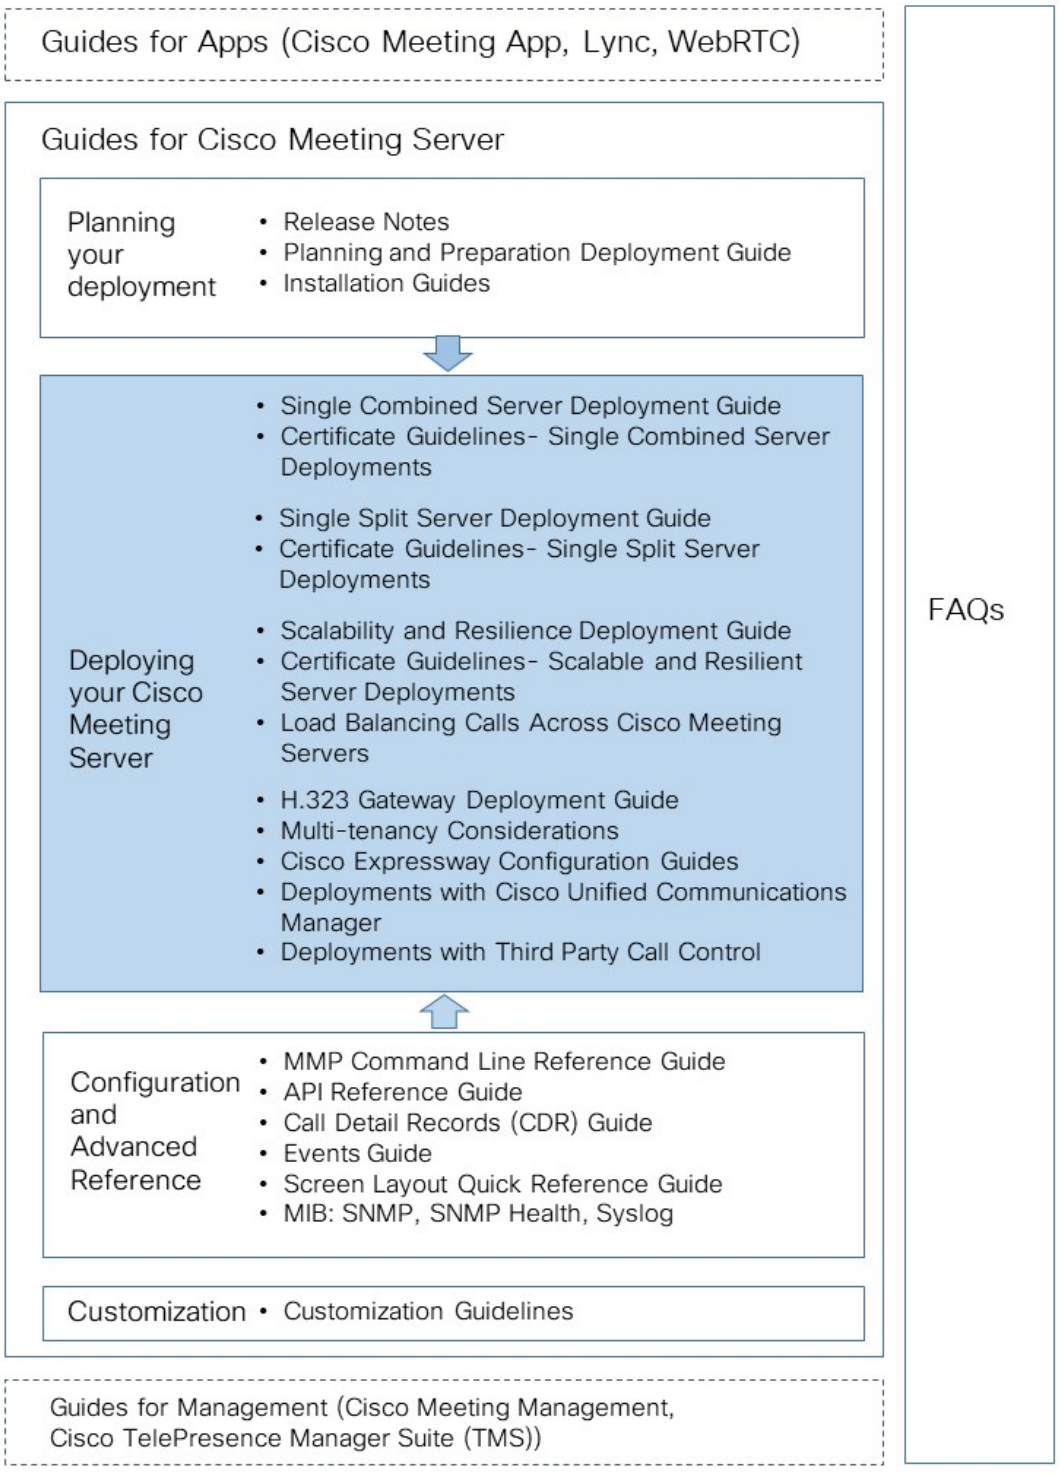

Note: The address ranges we use in Cisco user documentation are those defined in RFC 5737 which are explicitly reserved for documentation purposes. IP addresses in Meeting Server user documentation should be replaced with correct IP addresses routable in your network, unless otherwise stated.

1.3.1 Commands

In this document, commands are shown in black and must be entered as given–replacing any parameters in <> brackets with your appropriate values. Examples are shown in blue and must be adapted to your deployment.

1.4 Management and network interfaces

There are two layers to the Meeting Server: a Platform and an Application.

- The Platform is configured through the Mainboard Management Processor (MMP). The MMP is used for low level bootstrapping, and configuration via its command line interface, see the Cisco Meeting Server MMP Command Line Reference for details..

  Note: On the Acano X-series servers the MMP can be accessed via the serial Console port or using SSH on the Ethernet interface labeled Admin. In virtualized deployments the MMP is accessed on virtual interface A.

- The Application runs on this managed platform with configuration interfaces of its own. The application level administration (call and media management) is done via either the Call Bridge's Web Admin Interface, or through the API. Either can be configured to run on any one of the A-D Ethernet interfaces.

  In virtualized deployments one Ethernet interface (A) is created, but up to three more can be added (B, C and D).

  On the Acano X-series servers there are five physical Ethernet interfaces labeled Admin, A, B C and D. There is no physical separation between the media interfaces A-D on an X-series server, but the Admin interface is physically separate. Each interface is configured independently at the IP level. IP forwarding is not enabled in either the Admin or host IP stack.

  CAUTION: Do not set the Web Admin to listen on the Admin interface. If you do, it may cause out of memory problems if there are a lot of web/API requests, resulting in various processes such as syslog or web proxy being killed to maintain core functionality. If this happens you will see a loss of Syslog messages and Web Admin access. See this FAQ.

See the appropriate Installation Guide for details.

1.4.1 Application Programming Interface

The Meeting Server supports an Application Programming Interface (API). The API uses HTTPS as a transport mechanism and is designed to be scalable in order to manage the potentially very large numbers of active calls and spaces available in a deployment.

The API includes LDAP server access methods for adding, configuring and modifying LDAP servers, and support for multi-tenancy for searching calls through an additional Tenant ID. Other additions include posting to space message boards, the ability to filter the set of active call legs to just those experiencing "alarm" conditions (for example, packet loss or excessive jitter) and the ability to retrieve system-wide status values.

Multi-tenancy means that groups of users can be entirely segmented within the solution as required by service provider deployments e.g. users will only be able to meet, assign users to spaces, and search in the directory within the same configured customer groups.

Refer to the Cisco Meeting Server API Reference guide for more details.

1.5 Obtaining information on hosted conferences

There are two mechanisms for obtaining information on conferences hosted on the Meeting Server which remove the need to constantly poll the API: Call Detail Records and Events.

1.5.1 Call Detail Records (CDRs)

The Meeting Server generates Call Detail Records (CDRs) internally for key call-related events, such as a new SIP connection arriving at the server, or a call being activated or deactivated.

The server can be configured to send these records to a remote system to be collected and analyzed. There is no provision for records to be stored on a long-term basis on the Meeting Server, nor any way to browse CDRs on the Meeting Server itself.

The CDR system can be used in conjunction with the Meeting Server API, with the call ID and call leg IDs values being consistent between the two systems to allow cross referencing of events and diagnostics.

The Meeting Server supports up to four CDR receivers, enabling you to deploy different management tools or multiple instances of the same management tool. For more information, see the Cisco Meeting Server Call Detail Records Guide.

1.5.2 Events

From version 2.4, the Meeting Server can notify an "events client" in real-time of changes that are occurring on the Meeting Server. The Meeting Server acts as a server for the events, and the events client could be for example, a web-based management application. Cisco Meeting Management acts as an events client.

Note: You can construct your own events client, which is similar to constructing an API client. The events client needs to support HTTP and WebSocket libraries, both are available in common scripting languages like Python. The events port on the Meeting Server is the same port as you configured for the Web Admin, typically TCP port 443 on interface A.

Rather than continually poll an API resource on the Meeting Server, an events client can subscribe to an event resource to receive updates. For example, after establishing a WebSocket connection between the events client and the Meeting Server, the events client can subscribe to the event resource `callRoster` and receive updates on the participant list of an active conference to find out when a new participant joins, or an existing participant changes layout etc.

For more information, see the Cisco Meeting Server Events Guide.

1.6 Cisco licensing

You will need licenses for the Cisco Meeting Server. For information on purchasing and assigning licenses, see Section 1.6.3 and Section 1.6.4.

1.6.1 Cisco Meeting Server licensing

The following features require a license installed on the Meeting Server before they can be used:

- Call Bridge
- Recording
- Streaming

From version 2.4, you no longer need to purchase a branding license to apply single or multiple branding to the WebRTC app login page, IVR messages, SIP or Lync call messages or invitation text.

The XMPP activation key is included in the Cisco Meeting Server software.

In addition to feature licenses, user licenses also need to be purchased, there are 3 different types of user licenses:

- PMP Plus,
- SMP Plus,
- ACU

For information on user licensing, see Section 1.6.2.

Note: From version 2.4, you have the choice of purchasing an activation key with SIP media encryption enabled or SIP media encryption disabled (unencrypted SIP media) for the Cisco

Meeting Server 1000, Cisco Meeting Server and the VM software image. For more information on the unencrypted SIP media mode and activation key see Appendix D.

1.6.2 Cisco user licensing

Cisco Multiparty licensing is the primary licensing model used for Cisco Meeting Server; Acano Capacity Units (ACUs) can still be purchased, but cannot be used on the same Call Bridge as Multiparty licenses. Contact your Cisco sales representative if you need to migrate ACUs to Multiparty licenses.

Multiparty licensing is available in two variations: Personal Multiparty Plus (PMP Plus) licensing, which offers a named host license, and Shared Multiparty Plus (SMP Plus) licensing, which offers a shared host license. Both Personal Multiparty Plus and Shared Multiparty Plus licenses can be used on the same server.

1.6.2.1 Personal Multiparty plus licensing

Personal Multiparty Plus (PMP Plus) provides a named host license assigned to each specific user who frequently hosts video meetings. This can be purchased through Cisco UWL Meeting (which includes PMP Plus). Personal Multiparty Plus is an all-in-one licensing offer for video conferencing. It allows users to host conferences of any size (within the limits of the Cisco Meeting Server hardware deployed). Anyone can join a meeting from any endpoint, and the license supports up to full HD 1080p60 quality video, audio, and content sharing.

Note: The initiator of an Ad Hoc conference can be identified and if they have been assigned a PMP Plus license then that is used for the conference.

Note: To determine the number of active calls using the PMP Plus licence of an individual, use the parameter **`callsActive`** on API object **`/system/multipartyLicensing/activePersonalLicenses`**. We generally allow 2 calls to be active allowing for one starting and other finishing. If the call is on a cluster of Call Bridges then use the parameter **`weightedCallsActive`** on API object **`/system/multipartyLicensing/activePersonalLicenses`** for each Call Bridge in the cluster. The sum of **`weightedCallsActive`** across the cluster matches the number of distinct calls on the cluster using the individual's PMP Plus license. If a PMP Plus licence is exceeded, then SMP Plus licences are assigned, see Section 1.6.5.

1.6.2.2 Shared Multiparty plus licensing

Shared Multiparty Plus (SMP Plus) provides a concurrent license that is shared by multiple users who host video meetings infrequently. It can be purchased at a reduced price with a UCM TP Room Registration license included when purchasing room endpoints, or it can be purchased separately. Shared Multiparty Plus enables all employees who do not have Cisco UWL Meeting licenses to access video conferencing. It is ideal for customers that have room systems deployed that are shared among many employees. All employees, with or without a Cisco UWL

Meeting license have the same great experience, they can host a meeting with their space, initiate an ad-hoc meeting or schedule a future one. Each shared host license supports one concurrent video meeting of any size (within the limits of the hardware deployed). Each Shared Multiparty Plus license includes one Rich Media Session (RMS) license for the Cisco Expressway, which can be used to enable business-to-business (B2B) video conferencing.

Note: To determine the number of SMP Plus licences required, use the parameter **`callsWithoutPersonalLicense`** on API object **`/system/multipartyLicensing`**. If the calls are on a cluster of Call Bridges then use the parameter **`weightedCallsWithoutPersonalLicense`** on API object **`/system/multipartyLicensing`** for each Call Bridge in the cluster. The sum of **`weightedCallsWithoutPersonalLicense`** across the cluster matches the number of distinct calls on the cluster which require an SMP Plus license.

1.6.2.3 Cisco Meeting Server Capacity Units

Acano Capacity Units (ACUs) have been renamed Cisco Meeting Server Capacity Units. Each Capacity Unit (CU) supports 12 audio ports or the quantity of concurrent media streams to the Cisco Meeting Server software shown in Table 4.

Table 4: Capacity Unit Licensing

Media Stream	Number of licenses per Capacity Unit	Number of licenses required per call leg
1080p30	0.5	2
720p30	1	1
480p30	2	0.5

Each CU also entitles the Licensee to content sharing in each meeting containing at least one video participant. For more information refer to the terms and conditions of the CU license.

1.6.3 Obtaining Cisco user licenses

This section assumes that you have already purchased the licenses that will be required for your Meeting Server from your Cisco Partner and you have received your PAK code(s).

Follow these steps to register the PAK code with the MAC address of your Meeting Server using the Cisco License Registration Portal.

1. Obtain the MAC address of your Meeting Server by logging in to the MMP of your server, and enter the following command: **`iface a`**

   Note: This is the MAC address of your VM, not the MAC address of the server platform that the VM is installed on.

2. Open the Cisco License Registration Portal and register the PAK code and the MAC address of your Meeting Server.
3. The license portal will provide a zipped copy of the license file. Extract the zip file and rename the resulting .lic file to **`cms.lic`**.
4. Using your SFTP client, log into Meeting Server and copy the **`cms.lic`** file to the Meeting Server file system.
5. Restart the Call Bridge using the command **`callbridge restart`**
6. After restarting the Call Bridge, the license status can be checked by typing **`license`**

   The activated features and expirations will be displayed.

1.6.4 Assigning Personal Multiparty licenses to users

Follow these steps to apply Multiparty licensing to the Meeting Server.

Note: This procedure requires that users imported from a single LDAP source are either all licensed or all not licensed.

1. Create a userProfile (POST /userProfiles) or update an existing one (PUT to /userProfiles/<user profile id>) with the hasLicence field set to “true” to indicate users associated with this userProfile have a Cisco user license.

   Or create a userProfile or update an existing one with the hasLicence field set to “false” to indicate users associated with this userProfile do not have a Multiparty license. Alternatively, leaving the hasLicense field unset will select the default setting of false.
2. Create an ldapSource (POST /ldapSources) or update an existing one (PUT to /ldapSources/<ldap source id>) with the userProfile id parameter. This associates the userProfile created in step 1 with the appropriate LDAP source.
3. POST /ldapSyncs with ldapSource id parameter to sync the LDAP source. All imported users will be associated with the given userProfile

To determine whether a specific user has as a license, use GET /users/<user id> to retrieve the userProfile associated with this user.

Note: If the userProfile is deleted, then the userProfile is unset for the ldapSource and the imported users.

1.6.5 How Cisco Multiparty licenses are assigned

When a meeting starts in a space, a Cisco license is assigned to the space. Which license is assigned by the Cisco Meeting Server is determined by the following rules:

- if one or more members with a Cisco PMP Plus license has joined a space, then one of their licenses will be used, if not, then
- if the person that created the space (the owner) has a Cisco PMP Plus license, then the license of that owner is assigned, if not, then
- if the meeting was created via ad hoc escalation from Cisco Unified Communications Manager, then Cisco Unified Communications Manager provides the GUID of the user escalating the meeting. If that GUID corresponds to a user with a Cisco PMP Plus license, the license of that user is assigned, if not, then
- if present a Cisco SMP Plus license is assigned.

1.6.6 Determining Cisco Multiparty licensing usage

The following objects and fields have been added to the API to enable Admins to determine the consumption of Multiparty licenses:

- /system/licensing object, enables an Admin to determine whether components of the Cisco Meeting Server have a license and are activated,
- /system/multipartyLicensing object returns the number of licenses available and in use,
- /system/multipartyLicensing/activePersonalLicenses object indicates the number of active calls that are using a Personal Multiparty Plus user license,
- userProfile field as part of LDAP Sync
- hasLicense field to the userProfile, this indicates if a user has a license
- ownerId and ownerJid fields per /coSpace object. If present, the ownerId field holds the GUID of the user that owns this coSpace, and ownerJid holds the JID of the user.

Note: The owner is set using the field ownerJid when POSTing or PUTing a /coSpace object. When GETing the /coSpace both the ownerJid and ownerId are returned for the user.

For more information on these additional object and fields to support Cisco Multiparty licensing, refer to the Cisco Meeting Server API Reference Guide.

2 General concepts for deployment

This chapter provides an overview of the general concepts for deploying the Meeting Server in a single combined server deployment. Figure 5 illustrates a typical deployment.

Figure 5: Example of a Meeting Server deployment using an Acano X-series server in a single combined server deployment

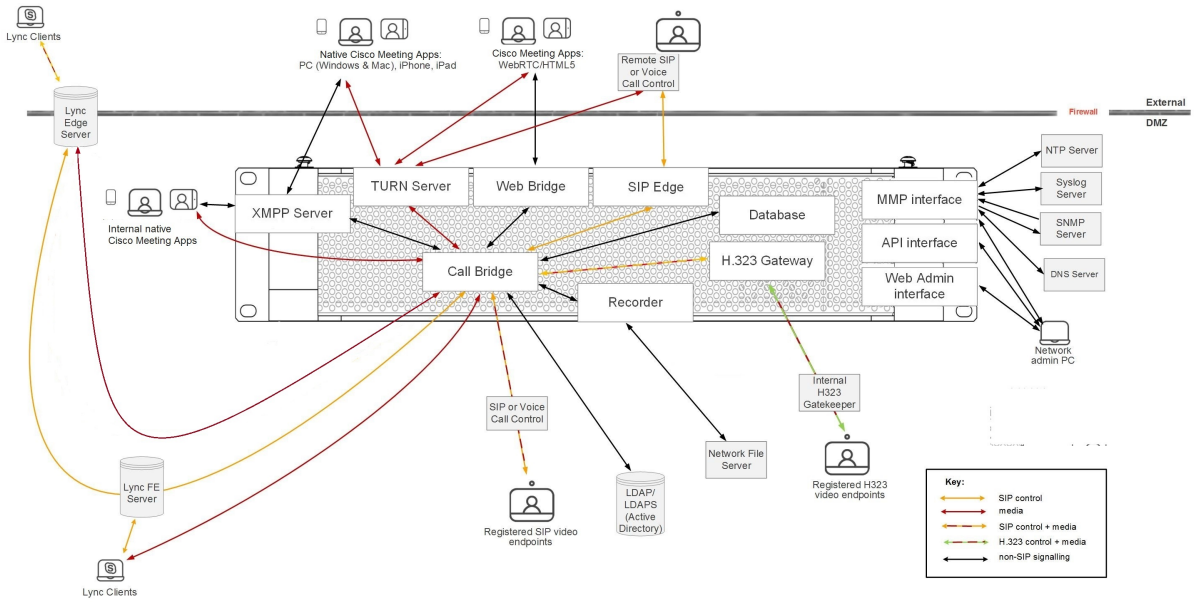

Note:

- The Meeting Server includes a Recording facility and a Streaming facility. Only enable the Recorder/Streamer on the same server as the Call Bridge if you are simply evaluating the features. For normal deployment enable the Recorder/Streamer on a different server to the Call Bridge. If you intend to deploy the Recorder and Streamer on the same Meeting Server, you will need to size the server appropriately for both uses. See Chapter 14 for more information on recording and Chapter 15 for more information on streaming.
- The SIP Edge component is still a beta feature in version 2.4, and should not be deployed in a production network. It will be removed in a future version of the Cisco Meeting Server software. You are advised to start migrating over to the SIP Edge component in Cisco Expressway X8.11.
- The Meeting Server includes an H.323 Gateway. The gateway is designed to be used only with the Call Bridge. Other than a brief summary in Section 2.6 this guide does not cover the H.323 Gateway, instead see the H.323 Gateway Deployment Guide for more information. Note that Cisco plans to end of life the Cisco Meeting Server H.323 Gateway component in November 2018, after which there will be no further development or feature releases related

to the H.323 Gateway. Customers are encouraged to start evaluation of the more mature H.323 Gateway component in the Cisco Expressway, and plan their migration over.

2.1 Web Admin

The Web Admin is a web based interface to configure the Meeting Server.

After configuring the Web Admin Interface for HTTPS access, as described in the Meeting Server installation guide, type the hostname or IP address of the server in a web browser to reach the login screen of the Web Admin Interface. See Web Admin Interface – Configuration menu options for details of the configuration accessible through the Web Admin Interface. The Web Admin Interface can only be used to configure the Meeting Server currently logged into.

An alternative to using the Web Admin Interface, is to use a REST API tool for example Postman or Chrome Poster, to access the Meeting Server's API. The Meeting Server API is routed through the Web Admin Interface, so an HTTPS connection is setup between the browser and the Meeting Server. The API Reference Guide is available here.

2.2 Call Bridge

The Call Bridge is the component on the Meeting Server that bridges the conference connections, enabling multiple participants to join meetings hosted on the Meeting Server or Lync AVMCUs. The Call Bridge exchanges audio and video streams so that participants can see and hear each other.

2.2.1 Call Bridge license

The Call Bridge license allows the Call Bridge to be used for media calls. The license needs to be installed on:

- the Cisco Meeting Server 1000,
- VM servers with Cisco Meeting Server software installed and configured as a combined server deployment (all components are on the same server),

You need to have the Call Bridge activated to create any calls, if you require demo licenses to evaluate the product then contact your Cisco sales representative or Cisco partner.

Acano X-Series Servers do not require a license. VMs configured as Edge servers do not require a license for the Call Bridge.

To apply the license after uploading the license file, you need to restart the Call Bridge. However, you must configure the Call Bridge certificates and a port on which the Call Bridge listens before you can do this. These steps are part of the Meeting Server configuration and described in Section 4.3 and the Certificate Guidelines for Single Combined Server Deployments.

The banner “This CMS is running in evaluation mode; no calls will be possible until it is licensed.” is displayed in the Web Admin interface until a valid cms.lic file is uploaded. After you upload the license file, the banner is removed.

2.3 Database

The Call Bridge reads from and writes to the database storing the space information, for example the members of spaces, chat messages occuring between members of a space, recent activity within a space.

In a single combined deployment the database is created and managed automatically by the Call Bridge and does not require a license or being enabled.

2.4 New WebRTC App and Web Bridge

Version 2.3 of the Meeting Server introduced the new WebRTC app which receives and transmits higher quality video using H.264, and has an improved user interface, similar to the Cisco Meeting App version 1.10 for Windows, Mac and iOS.

In version 2.4, the browsers supporting the WebRTC app are extended to:

- Google Chrome for Windows, macOS and Android. Use Chrome version 66 or later. We strongly recommend using the most recent version of Chrome.
- Mozilla Firefox for Windows and macOS. Use Firefox is 59.0.2 or later. We strongly recommend using the most recent version of Firefox.
- Apple Safari for macOS. Use Safari 11.1 or later. We strongly recommend using the most recent version of Safari.

Note: Content cannot be sent from Safari on macOS, iOS or from Chrome on Android, this is a browser limitation.

For more information on browser support and supported devices, see the Cisco Meeting App WebRTC Important Information.

Version 2.5 supports additional browsers, these are:

- Safari on iOS for iPads, running the latest version of iOS (recommended). iOS 11.0 is the minimum supported release.
- Safari on iOS for iPhones, running the latest version of iOS (recommended). iOS 11.0 is the minimum supported release. (This is beta quality in version 2.5.x).

  Note: We have tested the WebRTC app using the Safari browser on iPad Air 2 and iPad Pro 12.9 inch (2nd generation) with iOS 11.4.1, iPad (6th generation) with iOS 12.0.1, iPhone 6 on iOS12, iPhone 7 on iOS 12 and 12.1, iPhone 8 Plus on iOS12 and 12.1, and iPhone X on iOS 11.4.1.

- the latest version of Microsoft Edge (Microsoft Edge 42/Microsoft EdgeHTML 17) on Microsoft Windows 10 (this is beta quality in version 2.5.x).

  Note: There are limitations using the WebRTC app with Microsoft Edge and Mozilla FireFox browsers:

  - Using the WebRTC app with Microsoft Edge will not work if using the TURN server in Cisco Expressway or using the Meeting Server TURN with TCP.
  - Using the WebRTC app with Firefox will not work if using the TURN server in Cisco Expressway with TCP, but will work with the Meeting Server TURN with TCP.

  See Cisco Meeting App WebRTC Important Information for further details on these and other limitations.

Note: There are differences between the new WebRTC app and the new Cisco Meeting App version 1.10 for Windows and Mac. Refer to the Feature Comparison Matrix that accompanies the user documentation for these differences.

Behind the WebRTC app is a new Web Bridge, there is a minor change to the functionality and configuration of the new Web Bridge. This change is:

- the legacy mode for guest access on the Web Admin interface (**Configuration** > **General** > **Guest access via ID and passcode**) has no effect. From version 2.3, if passcodes are required for guest, then the passcode needs to be supplied at the same time as the guest id.

Note: If you have a single combined Meeting Server deployment then the Web Bridge will be upgraded to the new version when you upgrade the Meeting Server software to version 2.3. For deployments involving multiple Meeting Servers, we recommend that you upgrade all Meeting Servers to the same version to avoid the risk of any incompatibilities between versions.

If you are using the WebRTC app you will need to enable and configure the Web Bridge, refer to the sections Configuring the Web Bridge and Web Admin interface settings for the Web Bridge . The WebRTC app works on HTML5-compliant browsers and uses the WebRTC standard for

video and audio. For a list of tested browsers see the Meeting Server FAQ here. Using the Web Bridge does not require a license, but it does require an enabled Call Bridge.

2.4.1 Customizing the WebRTC sign in page

From version 2.3, the redesigned Web Bridge can only be customized through the API; it is no longer possible to upload a new background image and logo for the WebRTC app using the Web Admin interface.

The new look and feel for the WebRTC app, has resulted in changes to the design elements that can be rebranded. From 2.3, only these elements can be rebranded via the API:

- sign in background image for WebRTC app,
- sign in logo,
- text below sign in logo,
- text on browser tab.

Note: Customers who have previously used the API for branding archive application should review the 2.4 customization guidelines to confirm that their existing archive is still compatible. Incompatible archives will result in the Web Bridge failing to start correctly.

Figure 6: WebRTC app assets

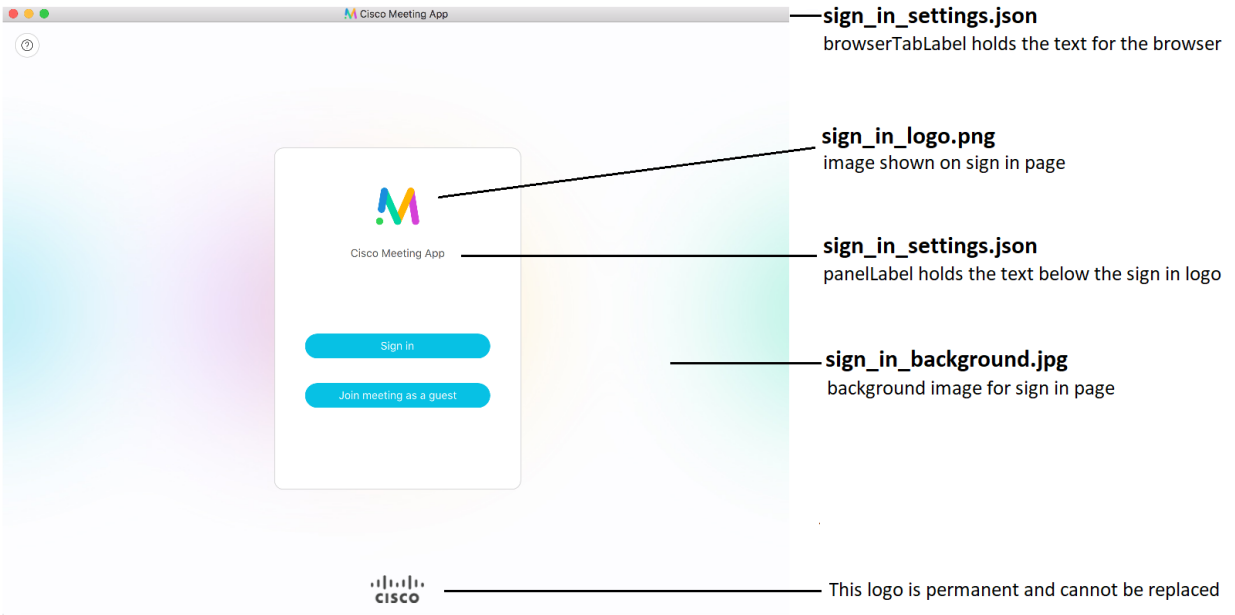

In addition to the design elements listed above which can be rebranded, IVR messages, SIP/Lync call messages and the text shown in invitations to join conferences, can be customized via the API.

There are two levels of branding:

- Single brand via API: only a single set of resources can be specified (1 WebRTC page, 1 set of voice prompts etc). These resources are used for all spaces, IVRs and Web Bridges.
- Multiple brand via API: different resources can be used for different spaces, IVRs and Web Bridges. These resources can be assigned at the system, tenant or space/IVR level.

Note: From version 2.4, you no longer need to purchase a branding license to apply single or multiple branding to the WebRTC app login page, IVR messages, SIP or Lync call messages or invitation text.

See the Cisco Meeting Server 2.4 Customization Guidelines for examples on using the API to undertake this level of customization.

2.4.2 Hosting branding files locally

Note: Hosting branding files locally on Acano X Series servers is beta quality in version 2.5.x.

Prior to version 2.5, using branding files for the Meeting Server required you to configure a separate web server to hold the branding files (voice prompts and lobby screen branding assets). From version 2.5 one set of branding files can be held locally on the Meeting Server. These locally hosted branding files are available to the Call Bridge and Web Bridge once the Meeting Server is operational, removing the risk of delays in applying customization due to problems with the web server. The images and audio prompts replace the equivalent files built into the Meeting Server software; during start up, these branding files are detected and used instead of the default files. Locally hosted branding files are overridden by any remote branding from a web server.

You can change these locally hosted files simply by uploading a newer version of the files and restarting the Call Bridge and Web Bridge. If you remove the locally hosted files, the Meeting Server will revert to using the built-in (US English) branding files after the Call Bridge and Web Bridge have been restarted, providing a web server has not been set up to provide the branding files.

Note: To use multiple sets of branding files, you still need to use an external web server.

For more information on hosting branding files locally, see the Cisco Meeting Server 2.5 Customization Guidelines.

2.5 On screen messaging

The Meeting Server provides the ability to display an on-screen text message to participants in a meeting hosted on the Meeting Server; only one message can be shown at a time. Using the

API, the duration that the message is displayed can be set, or made permanent until a new message is configured. Use the **`messageText`**, **`messagePosition`** and **`messageDuration`** parameters for API object **`/calls`**.

For users of SIP endpoints and Lync/Skype for Business clients, the on-screen text message is displayed in the video pane. The position of the message in the video pane can be selected from top, middle or bottom.

On screen messaging is also sent to other devices that are using ActiveControl in the deployment, for instance CE8.3 endpoints, and individual Meeting Servers not in a cluster but with the in-call message feature enabled. Meeting Servers in a cluster also support on screen messaging through a proprietary mechanism.

2.6 TURN server

Note: Cisco plans to remove the TURN server component from the Cisco Meeting Server software in a future version. Customers are encouraged to plan their migration over to using Cisco Expressway for TURN, see the Cisco Expressway Traffic Classification Deployment guide for deployment information.

The TURN server provides firewall traversal technology, allowing the Meeting Server to be deployed behind a Firewall or NAT. To connect to the deployment from external Cisco Meeting Apps, Lync clients or SIP endpoints registered to a SIP or voice call control device, you need to enable the TURN server, refer to the sections on Configuring the TURN server and Web Admin interface settings for the TURN server. If you are using Cisco Meeting Apps you also need to configure the Web Admin interface to allow the Call Bridge and external clients to access the TURN server. Using the TURN server does not require a license.

The TURN server listens on both ports 443 and 3478 for both UDP and TCP connections. Media sent over TCP is encrypted using TLS. The TURN server supports TCP to UDP interworking (see Figure 7). A browser can send TCP media to the TURN server which converts it to standard UDP media. This is useful when UDP traffic from browsers is blocked.

Figure 7: TURN server supporting TCP and UDP

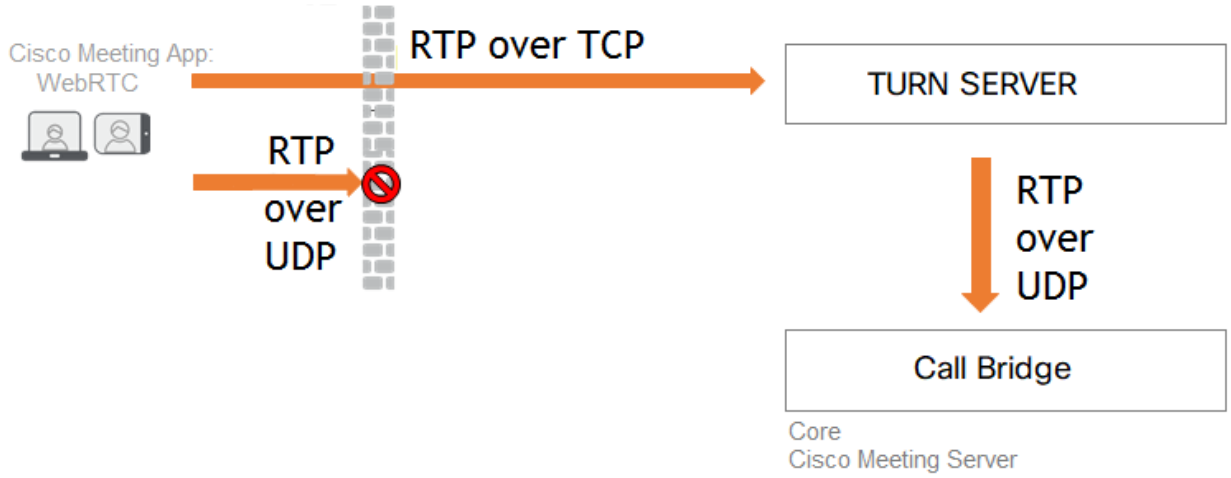

From version 2.0.4, the TURN server in a combined server deployment must be configured to listen on the loopback interface. See Section 4.7 for details.

Note: The Web Bridge sends STUN traffic to the TURN server in order to determine round trip time. For scalable deployments with multiple Web Bridges and TURN servers, the round trip time enables the Web Bridge to select the best TURN server for the session. From a network and firewall perspective, this will appear as though the Meeting Server is sending STUN traffic to it's own public IP address, network tools may flag this as an attack. This traffic can either be allowed or blocked, if this traffic is blocked the Web Bridge will choose one of the TURN servers for the WebRTC client but it might not be the best one for the WebRTC client in question. However, it should not have any other impact on the Meeting Server.

2.6.1 Enabling and disabling UDP signaling for SIP

The "UDP signaling for SIP" setting allows you to completely disable SIP over UDP, or to enable "single address" or "multi address" mode. Single address mode corresponds to the SIP over UDP behavior in versions prior to 2.2 and is the default, multi address mode allows SIP over UDP on multiple interfaces.

Use multi address mode if the Call Bridge is configured to listen on more than one interface for SIP over UDP traffic. Disable "UDP signaling for SIP" if you use SIP over TCP, or require that all of your network traffic is encrypted .

The "UDP signaling for SIP" mode is set through the Web Admin interface of the Call Bridge. Log into the Web Admin interface and select **Configuration>Call settings**, see Figure 8.

Figure 8: Settings for UDP signaling for SIP

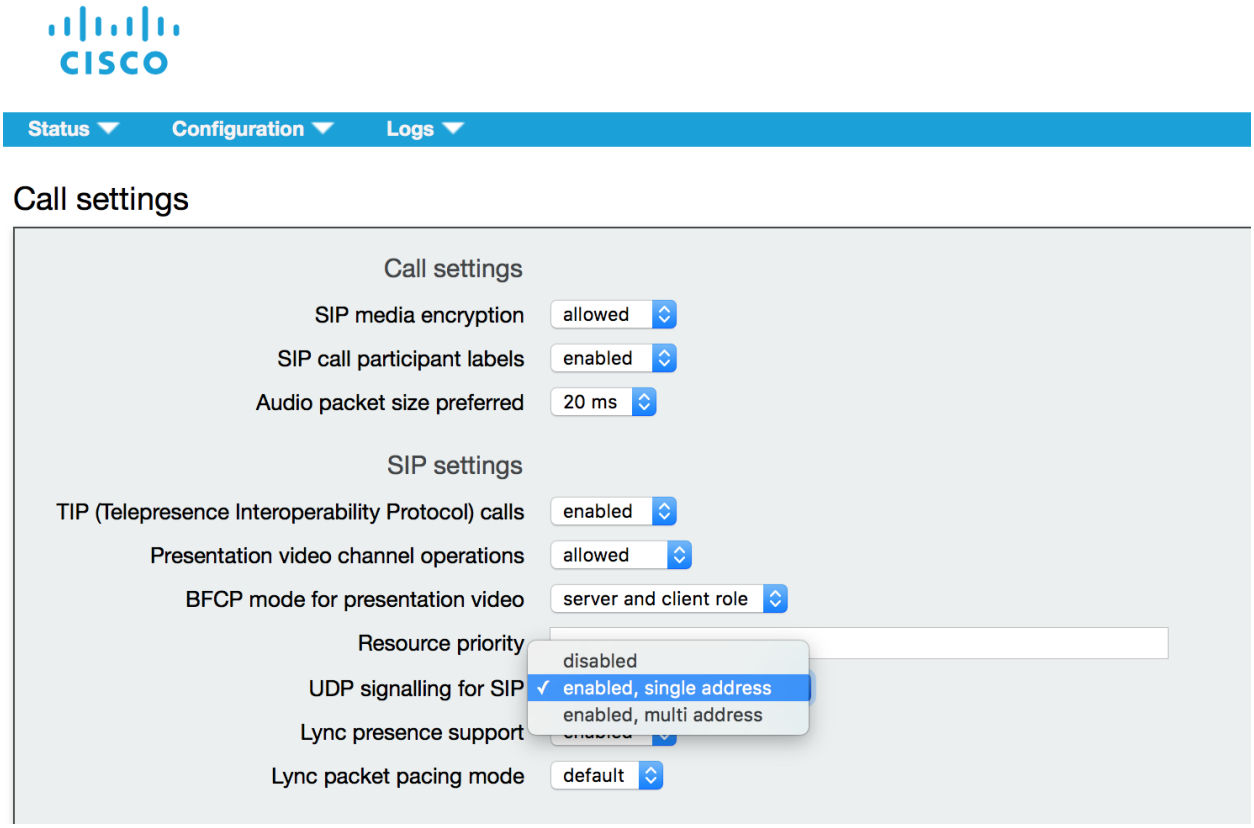

2.7 XMPP server

Note: Cisco is simplifying the Cisco Meeting Server and Cisco Meeting App interaction, and as a result the app dependence on XMPP will be removed. Once this development is complete, Cisco will remove XMPP from the Cisco Meeting Server product line. Customers are encouraged to start planning the migration to the Cisco Meeting WebRTC app rather than using the Cisco Meeting App thick clients (Windows, Mac and iOS).

Customers who are using Cisco Meeting Apps require an XMPP license installed on the server(s) running the XMPP server application. The XMPP license is included in the Cisco Meeting Server software. You will also need a Call Bridge activated on the same Cisco Meeting Server as the XMPP server.

The XMPP server handles the signaling to and from Cisco Meeting Apps, including the WebRTC app. If you are NOT planning to use the Cisco Meeting Apps for PC, iOS (iPhone and iPad), Mac or WebRTC Client you do not need to enable the XMPP server, disregard all sections referring to the XMPP server.

2.7.1 Deploying Cisco Meeting Apps

Note: Acano clients are now referred to as Cisco Meeting Apps in the Meeting Server documentation.

If you are using any of the Cisco Meeting Apps you need to enable the XMPP server, refer to the sections on Configuring the XMPP server and Web Admin interface settings for XMPP.

CAUTION: The maximum number of concurrent XMPP clients supported by the current Meeting Server software is 500.This maximum is a total number of all different clients (Cisco Meeting App, WebRTC Sign-in and WebRTC Guest clients) registered at the same time to clustered Meeting Servers. If the number of concurrent XMPP registrations exceeds 500 sessions, some unexpected problems with sign in may occur or it may lead to a situation where all currently registered users need to re-sign in, this can cause a denial of service when all users try to sign in at the same time.

The following diagram shows example control and media flows during a Cisco Meeting App call.

Figure 9: Example call flow diagram

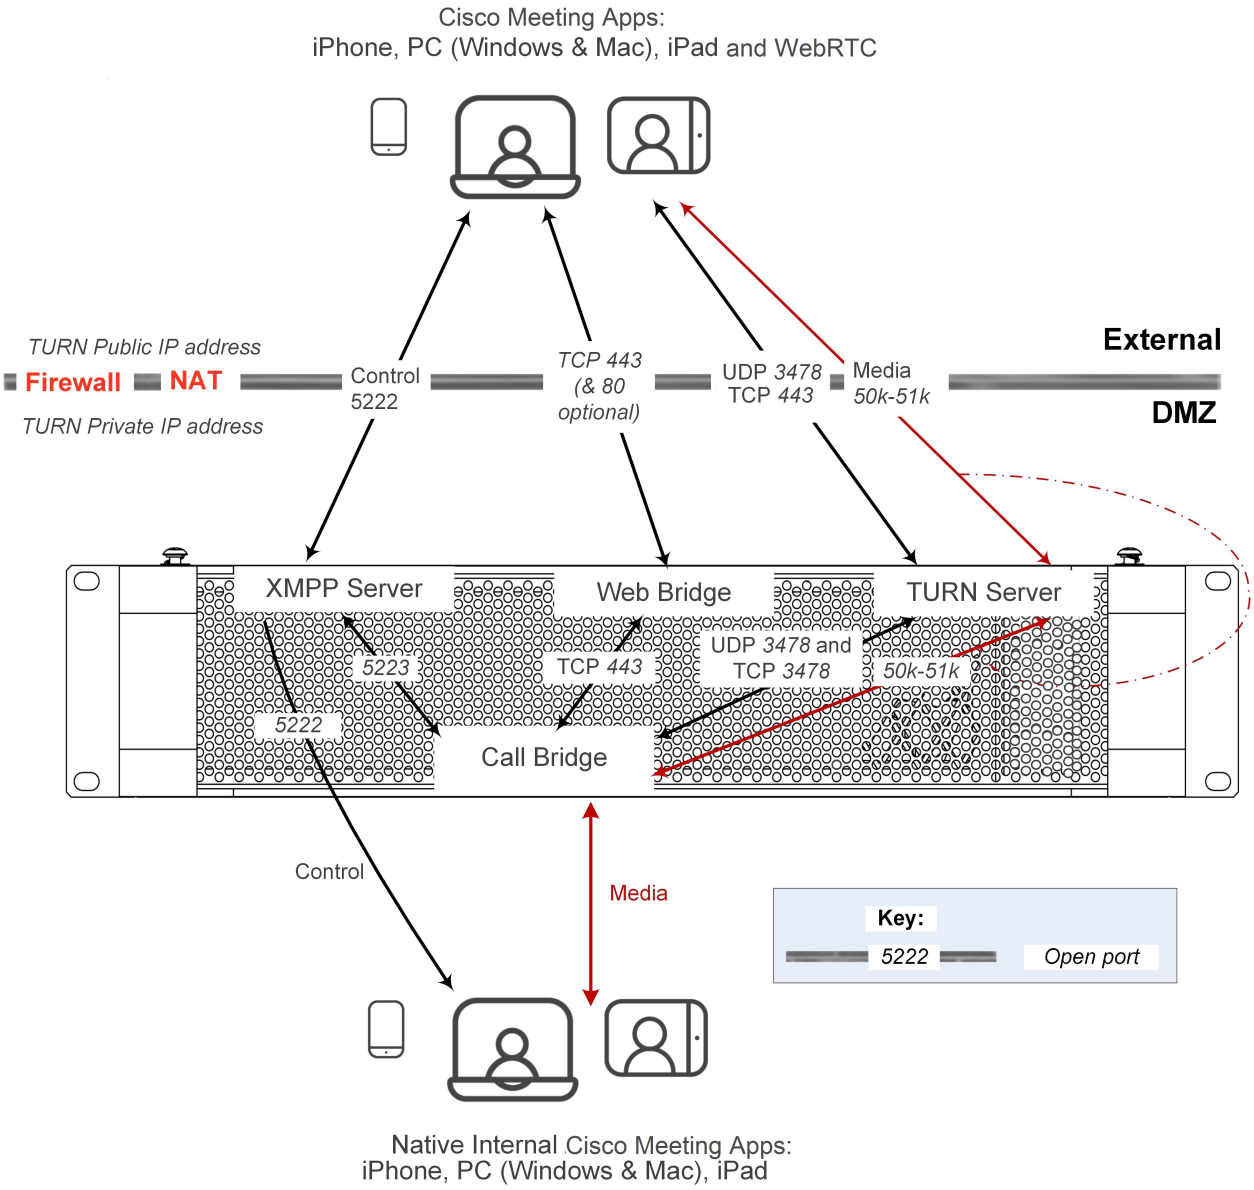

Points to note on the figure above:

- The following ports must be open:
  - UDP Port 3478 between the TURN server and remote Cisco Meeting Apps for PC, Mac, iOS and the WebRTC client. Note: if Port 3478 is blocked, the WebRTC client will fallback to Port 443. The remote Cisco Meeting Apps for PC, Mac and iOS will not fallback to Port 443.

Note: If you use Port 443 for both the Web Bridge and the TURN server then they must be on different interfaces of the Meeting Server. Alternatively chose a different port for the TURN server.

  - UDP Port 50000-51000 from Call Bridge to TURN server (for media). Although the range between the TURN server and the external Cisco Meeting Apps is shown as 50000-51000, future releases may require a wider range of 32768-65535.
  - TCP Port 443 (HTTPS) from Call Bridge to Web Bridge (for guest login). The Web Bridge is required for Cisco Meeting Apps to look up guest login when using a web browser that does not support WebRTC (for example Internet Explorer).

- The TURN server listens on TCP port 3478 for communication from the Call Bridge. You do not need to open UDP port 3478 or TCP port 3478 as they are internal to the Meeting Server.
- Internal clients connect directly to the XMPP server on port 5222 and media connects directly between the Cisco Meeting App and the Call Bridge.
- External Cisco Meeting Apps establish a control connection to the XMPP server (black line). Media can go directly from the Cisco Meeting App to the Call Bridge (dashed red line) or be relayed via the TURN server if required (red line).
- Another deployment option is to enable the XMPP server on a second interface and connect that interface to the private network. Then internal clients can connect directly to the XMPP server. Separate internal and external SRV records for the XMPP service need to be configured, directed to the two interfaces that the XMPP server is listening on.
- Both internal and external Cisco Meeting Apps use ICE/TURN to find suitable candidates for connectivity and choose the best: in the case of internal clients this will always be the local host candidates on the internal network.

2.8 H.323 Gateway

Note: Cisco plans to remove the H.323 Gateway component from the Cisco Meeting Server software in a future version. Customers are encouraged to start evaluation of the more mature H.323 Gateway component in the Cisco Expressway, and plan their migration over.

The H.323 Gateway enables an H.323 call to connect to the Call Bridge. The H.323 Gateway does not provide firewall traversal or call control, you are recommended to deploy an H.323 Gatekeeper to perform these functions. With the H.323 Gateway enabled, you can make the following calls:

H.323 call > H.323 GW > space

H.323 call > H.323 GW > Call Bridge->Lync

H.323 call > H.323 GW > Call Bridge->SIP device

H.323 call > H.323 GW > Call Bridge->Cisco Meeting App

The H323 Gateway can be enabled on the same server as the Call Bridge or on a separate one. By default the H.323 Gateway uses port 6061.

Refer to the H.323 Gateway deployment guide for more information.

2.9 SIP trunks and routing

The Meeting Server requires SIP trunks to be set up from one or more of the following: SIP Call Control, Voice Call Control and Lync Front End (FE) server. Changes to the call routing configuration on these devices are required to route calls to the Meeting Server that require the XMPP service or Web Bridge service for interoperability.

2.10 Support for Lync and Skype for Business

2.10.1 Support for Lync and Skype for Business clients

You can use Skype for Business clients, and Lync 2010 and Lync 2013 clients connected to a Skype for Business server, Lync 2010 or 2013 server .

The Meeting Server uses:

- the RTV codec transcoding up to 1080p with the 2010 Lync Windows client and 2011 Lync Mac clients,
- the H.264 codec with the 2013 Lync Windows client and Skype for Business client.

The Meeting Server will provide both RTV and H.264 streams when a mixture of clients versions are connected.

Lync 2010 and 2013 clients and Skype for Business clients can share content. The Meeting Server transcodes the content from native Lync RDP into the video format used by other participants in the meeting and sends it as a separate stream. Lync and Skype for Business clients also receive content over a RDP stream and can display it separately from the main video.

The Lync FE Server will need a Trusted SIP Trunk configured to route calls originating from Lync endpoints through to the SIP video endpoints i.e. to route calls with destination in the SIP video endpoint domain through to the Call Bridge.

The SIP Call Control will require configuration changes to route calls destined to the Lync/Skype for Business client domain to the Call Bridge so that SIP video endpoints can call Lync/Skype for Business clients.

The dial plan routes Lync/Skype for Business calls between these two domains in both directions.

The Meeting Server includes support for Lync Edge to enable Lync/Skype for Business clients outside of your firewall to join spaces.

Dual homed conferencing functionality improves how the Meeting Server communicates with the Lync AVMCU, resulting in a richer meeting experience for both Lync/Skype for Business and Cisco Meeting App users. Appendix E describes the dual homed conference experience.

2.10.2 Support for Dual Homed Conferencing

Dual homed conferencing requires the Lync Edge settings to be configured on the Lync Edge server settings on the Meeting Server for conference lookup. If you already have an on-prem Lync deployment or Lync Federation deployment working with the Meeting Server deployment, then no additional configuration is required on the Meeting Server. If this is a new deployment, then you need to setup the Meeting Server to use the Lync Edge server, see Chapter 8.

For information on the features which improves the experience of participants in Lync/Skype for Business meetings, see:

- FAQ on the improvements in meeting experience for Lync participants,
- FAQ on RDP support,
- FAQ on multiple video encoder support.

2.11 Recording meetings

The Recorder component on the Meeting Server adds the capability of recording meetings and saving the recordings to a document storage such as a network file system (NFS).

2.11.1 License keys for recording

Recording is controlled by license keys, where one license allows one simultaneous recording. The license is applied to the server hosting the Call Bridge (core server) which connects to the Recorder, not the server hosting the Recorder.

Note: The recommended deployment for production usage of the Recorder is to run it on a dedicated VM with a minimum of 4 physical cores and 4GB . In such a deployment, the Recorder should support 2 simultaneous recordings per physical core, so a maximum of 8 simultaneous recordings.

To purchase recording license keys, you will need the following information:

- number of simultaneous recordings,
- MAC address of interface A on the servers hosting the Call Bridges.

You can purchase recording license keys through Cisco's ecommerce tool.

2.12 Streaming meetings

The Streamer component adds the capability of streaming meetings held in a space to the URI configured on the space.

An external streaming server needs to be configured to be listening on this URI. The external streaming server can then offer live streaming to users, or it can record the live stream for later playback.

Note: The Streamer component supports the RTMP standard in order to work with third party streaming servers that also support the RTMP standard. However, we have only tested against Vbrick as an external streaming server.

2.12.1 License keys for streaming

You will need one or more licenses for streaming which is loaded on the Meeting Server hosting the Call Bridge, not the server hosting the Streamer. One 'recording' license supports 1 concurrent streaming or 1 recording, existing recording licences will allow streaming. Contact your Cisco sales representative or partner to discuss your licensing requirements.

2.13 Diagnostics and troubleshooting

In addition to using Syslog records (see Section 3.1.4) to help diagnose deployment issues, the following features are available on the Meeting Server:

- SIP tracing
- log bundle
- generate keyframe for specific call leg
- regular reporting of registered media modules
- retrieving diagnostics on Recorder/Sreamer/Web Bridge

2.13.1 SIP Tracing

You can enable additional SIP tracing using the **Logs > Detailed tracing** page in the Web Admin Interface. These logs may be useful when investigating call setup failure issues for SIP endpoints and should be disabled at all other times. To prevent the verbose logging being enabled for longer than necessary, it automatically shuts off after a choice of 1 minute, 10 minutes, 30 minutes or 24 hours. Refer to the Meeting Server Support FAQs on the Cisco website for more troubleshooting information.

Diagnostics for failed login attempts include:

- the IP address of the far end included in event log messages relating to logins
- audit messages generated for unsuccessful logins (minus the user name) and log in session timeouts. They are also generated for successful logins.

2.13.2 Log bundle

Meeting Server can produce a log bundle containing the configuration and state of various components in the Meeting Server. This log bundle will help Cisco Support speed up their analysis of your issue.

If you need to contact Cisco support with an issue, follow these steps to download the log bundle from the Meeting Server.

1. Connect your SFTP client to the IP address of the MMP.
2. Log in using the credentials of an MMP admin user.
3. Copy the file logbundle.tar.gz to a local folder.
4. Rename the file, changing the logbundle part of the filename to identify which server produced the file. This is important in a multi-server deployment.
5. Send the renamed file to your Cisco Support contact for analysis.

2.13.3 Ability to generate a keyframe for a specific call leg

The **`generateKeyframe`** object is added to **`/callLegs/<call leg id>`**. POST to **`/callLegs/<call leg id>/generateKeyframe`** to trigger the generation of a new keyframe in outgoing video streams for the call leg in question. This is a debug facility, and Cisco Support may ask you to use the feature when diagnosing an issue.

2.13.4 Reporting registered media modules in syslog

Syslog can print a message every 15 minutes to allow people to monitor whether all media modules are alive and well.

An example from an Acano X3 server:

```
Apr 21 09:53:50 user.info cms-emea-01 host: server: INFO : media module status
11111111111
```

2.13.5 Retrieving diagnostics on a Recorder/Streamer/Web Bridge

There are API objects that enable the retrieval of:

- the number of **`activeRecordings`** on **`/recorders/<recorder id>`**
- the number of **`activeStreams`** on **`/streamers/<streamer id>`**:

and to retrieve the **`status`** on **`/recorders/<recorder id>`**, **`/streamers/<streamer id>`**, **`/webBridges/<web bridge id>`**. The table below shows the status settings for the components.

Status	Component	Recorder	Streamer	Web Bridge
unused	component is unused	✓	✓	✓
success	connected to the queried Call Bridge	✓	✓	✓
connectionFailure	could not connect to the queried Call Bridge	✓	✓	✓
invalidAddress	the configured URL is invalid	✓	✓	
dnsFailure	the configured URLcannot be resolved by the DNS server	✓	✓	
remoteFailure	a connection was established with the component but the Call Bridge received a failure response	✓	✓	
unknownFailure	an unknown failure occured	✓	✓	
lowDiskSpace	has limited disk space available	✓		

3 Prerequisites

3.1 Prerequisites

This chapter describes the changes to your network configuration that you need to consider before installing and configuring the Meeting Server; some of these items can be configured beforehand.

3.1.1 DNS configuration

The Meeting Server needs a number of DNS SRV and A records. See Appendix A for a full list, but specific records are also mentioned elsewhere.

3.1.2 Security certificates

You will need to generate and install X.509 certificates and keys for services which use TLS; for example, Call Bridge, Web Admin Interface (the Call Bridge's interface), Web Bridge, TURN server, and the XMPP server.

The Certificates Guidelines for single combined deployments contains both background information on certificates and instructions, including how to generate self-signed certificates using the Meeting Server's MMP commands. These certificates are useful for testing your configuration in the lab. However, in a production environment we **strongly recommend** using certificates signed by a Certificate Authority (CA).

Instructions that were previously in this guide concerning certificates have been removed and replaced by a single step referencing the Certificate Guidelines.

Note: If you self-sign a certificate, and use it, you may see a warning message that the service is untrusted. To avoid these messages re-issue the certificate and have it signed by a trusted CA: this can be an internal CA unless you want public access to this component.

3.1.3 Firewall configuration

See Appendix B for the list of ports which need to be opened on your firewall, and Section 17.6 for advice on creating Firewall rules.

3.1.4 Syslog server

The Meeting Server creates Syslog records which are stored locally and can also be sent to a remote location. These records are useful when troubleshooting because they contain more detailed logging than is available on a Meeting Server's own internal log page. Internal syslog

messages can be downloaded over SFTP, however Cisco recommends that the host server is configured to send debug information to a remote Syslog server.

Note: The Syslog server must use TCP, not UDP. Check that your Syslog server is configured to use TCP.

Follow the instructions below to define a Syslog server.

1. SSH into the MMP and log in.
2. Enter the following command, **`syslog server add <server address> [port]`**

   Examples:

   ```
   syslog server add syslog01.example.com 514
   syslog server add 192.168.3.4 514
   ```
3. Enable the Syslog server by entering:

   ```
   syslog enable
   ```
4. Optionally, if you want to send the audit log to a Syslog server follow these steps.

   (The audit log facility records configuration changes and significant low-level events. For example, changes made to the dial plan or configuration of a space via the Web Admin Interface or the API, are tracked in this log file, and tagged with the name of the user that made the change. The file is also available via SFTP.)

   a. Create a user with the audit role.

      ```
      user add <username> (admin|crypto|audit|appadmin)
      user add audituser audit
      ```
   b. Log out of the MMP and log back in with the newly created user account.
   c. Enter the command (this command can only be run by a user with the audit role):
      ```
      syslog audit add <servername>
      syslog audit add audit-server.example.org
      ```

Note: Normally local Syslog files are overwritten in time, but you can permanently store system and audit log files using the **`syslog rotate <filename>`** and **`syslog audit rotate <filename>`** commands. These files can also be downloaded over SFTP. See the MMP Command Reference.

3.1.5 Network Time Protocol server

Configure a Network Time Protocol (NTP) server to synchronize time between the Meeting Server components.

Note: Sharing a common view of time is important for multiple reasons, it is necessary when checking for certificate validity and to prevent replay attacks.

1. If necessary, SSH into the MMP and log in.
2. To set up an NTP server, type:

   ntp server add <domain name or IP address of NTP server>

To find the status of configured NTP servers, type **ntp status**

See the MMP Command Reference for a full list of **ntp** commands.

3.1.6 Call Detail Record support

The Meeting Server generates Call Detail Records (CDRs) internally for key call-related events, such as a new SIP connection arriving at the server, or a call being activated or deactivated. It can be configured to send these CDRs to a remote system to be collected and analyzed. There is no provision for records to be stored on a long-term basis on the Meeting Server, nor any way to browse CDRs on the Meeting Server.

The Meeting Server supports up to two CDR receivers, enabling you to deploy two different management tools or two instances of the same management tool for resiliency. If you are using Acano Manager, the Acano Manager server **must** be one of your CDR receivers, you can either add a second Acano Manager server or add a different management platform.

You can use either the Web Admin Interface or the API to configure the Meeting Server with the URI of the CDR receivers. If you are using the Web Admin interface go to **Configuration > CDR settings** and enter the URI of the CDR receivers. Refer to the Call Detail Records Guide or the API Reference guide for details on using the API to configure the Meeting Server with the URIs of the CDR receivers.

3.1.7 Host name

Cisco recommends that the Meeting Server is given its own hostname.

1. If necessary, SSH into the MMP and log in.
2. Type:

   hostname <name>
   hostname london1
   hostname mybox.example.com

3. Type:

   reboot

Note: A reboot is required after issuing this command.

3.1.8 Other requirements

- Access to an LDAP server to import users. This can be a Microsoft Active Directory (AD) server or an OpenLDAP server.

  If you plan for users to utilise the Cisco Meeting Apps to connect to the Meeting Server, then you must have an LDAP server. User accounts are imported from the LDAP server. You can create user names by importing fields from LDAP as described in LDAP configuration. The passwords are not cached on the Meeting Server, they are managed centrally and securely on the LDAP server. When a Cisco Meeting App authenticates, a call is made to the LDAP server.

- Decision on a dial plan to use to reach calls hosted on the Call Bridge. The dial plan will depend on your environment; that is whether you are making one or more of the following types of call: Lync, SIP (including voice) or Cisco Meeting App calls. Instructions for deploying this dial plan are given in Chapter 6.

- Access to one or more of the following to test the solution: Lync clients, SIP endpoints, SIP phones and/or Cisco Meeting Apps as appropriate.

- Access to a SIP Call Control platform if you intend to make SIP calls. Chapter 7 and Chapter explain how to set up a SIP trunk to the Cisco VCS and summarizes the required dial plan configuration changes. Information on setting up the SIP Trunk to a Cisco Unified Communications Manager (CUCM), the Avaya CM and Polycom DMA is provided in the Cisco Meeting Server Deployments with Call Control guide; you can use other call control devices not listed in the guide.

- If you intend to integrate the Meeting Server with an audio deployment, the Meeting Server must connect to a Voice Call Control device attached to a PBX; it is not possible to connect a Meeting Server directly to a PBX.

- If deploying in a Lync environment, access to the Lync Front End (FE) server to make dial plan configuration changes there. The changes required are given in this document.

3.1.9 Specific prerequisites for a virtualized deployment

- A host server that complies with the resources specified in the Installation Guide for Cisco Meeting Server Virtualized Deployments.

3.1.10 Specific prerequisites for Acano X-series server

- A suitable environment: refer to the Acano Hardware/Environmental Data Sheet for details on the required power and cooling

- The Acano X- series server has two power modules, and country-specific power cables are supplied for the AC power supplies. The server will work with just a single power unit connected. To implement power supply redundancy you must connect both modules to

power supplies. Connecting the modules to independent power supplies allows for the greatest resiliency.

- 2U of rack space if using the rack mounting kit; 3U of rack space if installing on a shelf
- A minimum of two Ethernet links:
  - One for the MMP (labeled Admin on the back of the Acano X-series server). The speed can be 100M or 1G.
  - One for a media interface (there are four labeled A to D). The speed can be 1G or10G.

  IP addresses can be configured statically or automatically via DHCP or SLAAC/DHCPv6. Ethernet links will operate at the speed of the network switch; the switch port should be set to auto negotiate speed. If you are using a speed of 10G be sure to use the appropriate cable.

See the Installation Guide for the Acano X-series server for full details.

4 Configuring the MMP

The Meeting Server components are configured using the MMP.

4.1 Creating and managing MMP and Web Admin interface user accounts

You should have created an MMP administrator user account by following the Cisco Meeting Server Installation Guide ; if so, go on to the next section. The same account is used to access the Web Admin Interface.

(If you do not have these MMP administrator user accounts, you will have to use the emergency admin recovery procedure detailed in the Installation Guide appropriate to your deployment.)

Note: See the MMP Command Reference Guide for the full range of MMP commands, including setting up additional administrator user accounts and user accounts with other roles.

4.2 Upgrading software

The Cisco Meeting Server 2000, Cisco Meeting Server 1000 and Acano X-series servers ship with the latest software release available at the time of shipment, but may not be up-to-date. Equally, if you downloaded the software some days ago, we advise you to check on the Cisco website in case a later version is available, and if so, upgrade to the latest version.

The following instructions apply to all types of deployment:

1. To find out which software version is running on the Meeting Server, SSH into the MMP of the server, log in and type:

   ```
   version
   ```

2. Before upgrading your Meeting Server:
   a. take a backup of the current configuration on the server . Save the backup safely to a local server. See the MMP Command Reference guide for full details. Do NOT use the automatic backup file that is created during the upgrade process.
   b. save the cms.lic and certificate files to the local server.
   c. using the Web Admin interface, check that all calls (SIP and clients) are working and no fault conditions are listed.
3. To upgrade, first download the appropriate software file from the Cisco website. Click on this link, then click on the appropriate Meeting Server type listed in the right hand column of the web page and follow any instructions displayed with the download link.
4. Use an SFTP client to upload the new software image to the MMP of the Meeting Server. For example:

```
sftp admin@10.1.124.10
put upgrade.img
```

where 10.1.x.y is an IP address or domain name.

5. To upgrade the server, connect via SSH to the MMP and type:

   ```
   upgrade
   ```

   wait approximately 10 to 12 minutes for the server to restart, and for the Web Admin interface to be available.

6. To verify that the upgrade was successful, SSH into the MMP, log in and type the following command:

   ```
   version
   ```

This completes the upgrading of the Meeting Server deployment. Now verify that:

- dial plans are intact,
- XMPP service is connected (if used),
- and no fault conditions are reported on the Web Admin interface and log files.

Check that you are able to connect using SIP and Cisco Meeting Apps (as well as Web Bridge if that is supported).

Note on rollback procedure: If anything unexpected happens after you upgrade the server and you decide to downgrade, simply upload the software release for the previous version, and type **`upgrade`**. Then use the MMP command **`factory_reset app`** on the server. Once the server has rebooted from a factory reset, use the **`backup rollback <name>`** command to restore the backup configuration files on the server. Providing you restore the backup file that was created from the server, the license file and certificate files will match the server.

4.3 Configuring the Call Bridge

The Call Bridge needs a key and certificate pair that is used to establish TLS connections with SIP Call Control devices and with the Lync Front End (FE) server. If you are using Lync, this certificate will need to be trusted by the Lync FE server.

Note: SIP and Lync calls can traverse local firewalls using the SIP Edge component, this is a beta feature and should not be used in production environments. If you plan to evaluate this feature, note that you need to configure trust between the Call Bridge and the SIP Edge, for more information see Chapter 13.

Note: SIP and Lync calls can traverse local firewalls using the Cisco Expressway, you will need to configure trust between the Call Bridge and the Cisco Expressway. Cisco Expressway must be running X8.9 or later. For more information, see Cisco Expressway Options with Cisco Meeting Server and/or Microsoft Infrastructure (Expressway X8.9.2) or if running X8.10 see Cisco Expressway Web Proxy for Cisco Meeting Server (X8.10) and Cisco Expressway Session Classification Deployment Guide (X8.10).

The command **`callbridge listen <interface>`** allows you to configure a listening interface (chosen from A, B, C or D). By default the Call Bridge listens on no interfaces.

1. Create and upload the certificate as described in the Certificate Guidelines.
2. Sign into the MMP and configure the Call Bridge to listen on interface A.

   ```
   callbridge listen a
   ```

   Note: the Call Bridge must be listening on a network interface that is not NAT'd to another IP address. This is because the Call Bridge is required to convey the same IP that is configured on the interface in SIP messages when talking to a remote site.

3. Configure the Call Bridge to use the certificates by using the following command so that a TLS connection can be established between the Lync FE server and the Call Bridge, for example:

   ```
   callbridge certs callbridge.key callbridge.crt
   ```

   The full command and using a certificate bundle as provided by your CA, is described in the Certificate Guidelines.
4. Restart the Call Bridge interface to apply the changes.

   ```
   callbridge restart
   ```

4.4 Configuring the Web Admin interface for HTTPS access

The Web Admin Interface is the Call Bridge's user interface. You should have set up the certificate for the Web Admin Interface (by following one of the Installation Guides). If you have not, do so now.

1. The installation automatically set up the Web Admin Interface to use port 443 on interface A. However, the Web Bridge also uses TCP port 443. If both the Web Admin Interface and the Web Bridge use the same interface, then you need to change the port for the Web Admin Interface to a non-standard port such as 445, use the MMP command **`webadmin listen <interface> <port>`**. For example:

   ```
   webadmin listen a 445
   ```
2. To test that you can access the Web Admin Interface, type your equivalent into your web browser: **`https://meetingserver.example.com:445`**

If it works, proceed to the next section.

3. If you cannot reach the Web Admin Interface:

   a. Sign into the MMP, type the following and look at the output:

      webadmin

      The last line of the output should say **"webadmin running".**

   b. If it does not there is a configuration problem with your Web Admin Interface. Check that you have enabled it by typing:

      webadmin enable

   c. The output of the **webadmin** command should also tell you the names of the certificates you have installed, e.g. webadmin.key and webadmin.crt.

      Note: They should be the same names of the certificates you uploaded previously.

      Assuming these are the names then type:

      pki match webadmin.key webadmin.crt

      This will check that the key and certificate match.

   d. If you are still experiencing issues, troubleshoot the problem as explained in the Certificates Guidelines.

4.5 Configuring the XMPP server

If you are using the Recorder or Streamer components or any of the Cisco Meeting Apps including the WebRTC Client you now need to configure the XMPP server and then enable it. Otherwise, skip this section.

Note: The Recorder and Streamer components behave as XMPP clients, so the XMPP server needs to be enabled on the Meeting Server hosting the Call Bridge.

From Cisco Meeting Server 2.0, the XMPP license is included in the Cisco Meeting Server software.

Note: You will need a Call Bridge activated on the same Meeting Server as the XMPP server.

1. To create DNS A and SRV records for the Meeting Server

   a. Create DNS A record for the fully qualified domain name (FQDN) of the server that will be used to host the XMPP Server and set it to the IP address of the interface that the XMPP server is listening on.

b. Create DNS SRV record for _xmpp-server._tcp for port 5269 resolving to the DNS A record created in step a above.

c. Create DNS SRV record for _xmpp-client._tcp for port 5222 resolving to the DNS A record created in step a above.

d. Test the above by running the following commands from a PC. They should return the correct IP addresses for these domains:

```
nslookup -querytype=srv _xmpp-server._tcp.example.com
nslookup -querytype=srv _xmpp-client._tcp.example.com
```

2. Sign into the MMP and generate the private key and certificate using the information in the Certificate Guidelines.

The XMPP server can be configured to listen on any subset of the four media interfaces and ignore connections from any interface in the complement.

3. Establish an SSH connection to the MMP and log in.
4. To configure the XMPP server to use one or more interfaces enter the following command:

   `xmpp listen <interface whitelist>`

   The following is an example where interface is set to interface A and B.

   `xmpp listen a b`
5. Assign the certificate and private key files that were uploaded earlier, using the command:

   `xmpp certs <keyfile> <certificatefile> [<cert-bundle>]`

   where keyfile and certificatefile are the filenames of the matching private key and certificate . If your CA provides a certificate bundle then also include the bundle as a separate file to the certificate. See the Certificate Guidelines for further information
6. Configure the XMPP server with the following command:

   `xmpp domain <domain name>`

   The following is an example where the domain name is example.com.

   `xmpp domain example.com`
7. Enable the XMPP service:

   `xmpp enable`
8. To allow a Call Bridge to access the XMPP server securely (after configuration), provide a component name for the Call Bridge to use to authenticate e.g. **`cb_london`**:

   `xmpp callbridge add <component name>`

   for example

   `xmpp callbridge add cb_london`

   A secret is generated; for example, you see:

```
cms>xmpp callbridge add cb_london
Added callbridge: Secret: aB45d98asdf9gabgAb1
```

9. Make a note of the domain, component and secret generated in the previous steps because they are required when you use the Web Admin interface to configure the Call Bridge access to the XMPP server (so that the Call Bridge will present the authentication details to the XMPP server).

   Note: If you lose the details, use the MMP command **`xmpp callbridge list`** to display them.

4.5.1 Configuring XMPP multi-domains

A single XMPP server can host multiple XMPP domains. For example, both example.com and example.org can exist on the same Meeting Server. It is possible to configure multiple tenants with the same XMPP domain (as in previous releases), or each tenant with their own domain, or mix these schemes.

Note: It is strongly recommended that multiple XMPP domains are not used for a single tenant, or in cases where tenants are not used.

To configure multiple domains for the XMPP server to listen to, use the MMP command:

`xmpp multi_domain add <domain name> <keyfile> <certificatefile> [<crt-bundle>]`

where:

`<keyfile>` is the private key that you created for the XMPP server

`<certificatefile>` is the signed certificate file for the XMPP server

`[<crt-bundle>]` is the optional certificate bundle as provided by the CA

Note: You also need to add a DNS SRV record for each additional XMPP domain, and to add the domain to the Incoming Calls page on the Web Admin interface (Configuration>Incoming calls).

Note: Restart the XMPP server for the configured multiple domains to take effect.

Note: The XMPP server will not start if the private key or certificate files are missing or invalid.

To list the domains that the XMPP server is listening to, use the command:

`xmpp multi_domain list`

To delete a domain that the XMPP server is listening to, use the command:

`xmpp multi_domain del <domain name>`

4.6 Configuring the Web Bridge

The Web Bridge is used by the WebRTC app. If you are deploying the WebRTC app you need to set the network interface for the Web Bridge and then enable it. Otherwise, skip this section.

1. SSH into the MMP.
2. Configure the Web Bridge to listen on the interface(s) of your choice with the following command:

   webbridge listen <interface[:port] whitelist>

   The Web Bridge can listen on multiple interfaces, e.g. one on public IP and one on the internal network. (However, it cannot listen on more than one port on the same interface.)

   The following is an example where interfaces are set to interface A and B, both using port 443.

   webbridge listen a:443 b:443

3. Create DNS A record for the Web Bridge and set it to resolve to the IP address of the Ethernet interface you want the Web Bridge to listen on.
4. Create a certificate and private key for the Web Bridge as described in the Certificates Guidelines. Upload the certificate file to the MMP via SFTP.
5. Add the Call Bridge certificate to the Web Bridge trust store as described in the Certificates Guidelines document.
6. The Web Bridge supports HTTPS. It will forward HTTP to HTTPS if configured to use "http-redirect". To do so:

   a. Enable HTTP redirect with the following command:

      webbridge http-redirect enable

   b. If required (see the note below), set the Windows MSI, Mac OSX DMG and iOS installers that are presented to WebRTC users:

      webbridge msi <url>

      webbridge dmg <url>

      webbridge ios <url>

   Note: If you only use browsers that support WebRTC (e.g. Chrome) you do not need to set these download locations because browser functionality will be used for guest access to space. However, if you use browsers that do not (e.g. IE, Safari) then configure these locations so that when the Meeting Server detects the device being used (iOS device, Mac, or PC), it can redirect the user to the configured client download link for that device, and prompt the user to install the correct Cisco Meeting App so that they can join the meeting. After installation, the user is connected to the space as a Guest.

7. Enable the Web Bridge with the following command:

   ```
   webbridge enable
   ```

8. Use the Web Admin interface to configure the settings through which the Call Bridge communicates with the Web Bridge, see Chapter 11

4.7 Configuring the TURN server

CAUTION: Your TURN server password and credentials must be unique. Do not reuse your admin username or password.

1. SSH into the MMP.
2. Configure the TURN server with the following command:

   ```
   turn credentials <username> <password> <realm>
   ```

   The following is an example where username is `myusername`, the password is `mypassword` and it uses the realm `example.com`.

   ```
   turn credentials myTurnUsername myTurnPassword example.com
   ```

3. If the TURN server has a public IP address rather than being NAT'ed (see Figure 2), this step is not required, go on to step 4. If the TURN server is located behind a NAT, set the public IP Address that the TURN Server will advertise using:

   ```
   turn public-ip <ip address>
   ```

   The following is an example where a public IP address is set to 5.10.20.99

   ```
   turn public-ip 5.10.20.99
   ```

   CAUTION: Locating the TURN server behing a NAT requires careful configuration of the NAT, to ensure connectivity always works. This is due to how Interactive Connectivity Establishment (ICE) works, and is not a problem specific to the TURN deployment within the Meeting Server. For information on deploying a TURN server behind NAT, see Appendix G.

   Note: The IP address set here should not be confused with the IP addresses set in the Web Admin Interface **Configuration > General** page. The MMP commands configure the TURN server itself, while the **Configuration > General** page settings allow the Call Bridge and external clients to access the TURN server, and are explained in Web Admin interface settings for the TURN server.

Figure 10: TURN server public IP address (not NAT'ed) - Combined server deployment

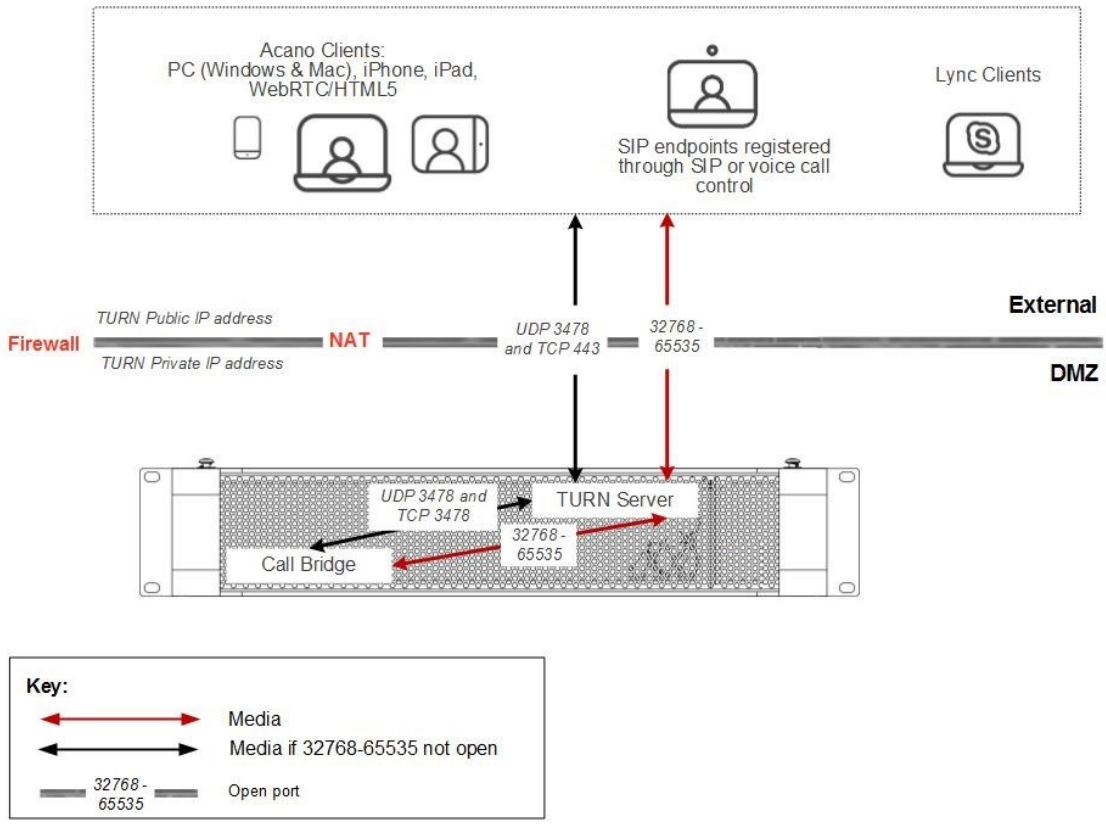

Note: Although the port range between the TURN server and the external clients is shown as 32768-65535, currently only 50000-51000 is used. The required range is likely to be larger in future releases.

4. Configure the TURN Server to listen on a specific interface using:

   `turn listen <interface whitelist>`

   In a single combined server deployment, the TURN server must be configured to listen on the loopback interface. Ensure that the whitelist of interfaces to listen on contains at least one interface, and specify the loopback interface. The loopback interface must not be the first interface in the whitelist.

   For example:

   `turn listen c lo` when a Call Bridge or Web Bridge is colocated on the same server as this TURN Server

   Note: You can specify more than one interface for the TURN server to listen on. If specifying multiple interfaces for the TURN server, the first one must be the public interface, i.e. the one

on the public network, or the one that a NAT forwards to. For example, **`turn listen b a`** where **`b`** is the NAT'd interface and **`a`** is the private internal interface.

5. Select the port for the TURN server to listen on using:

   `turn tls <port|none>`

   for example:

   `turn tls 443`

   Note: For maximum connectivity from external locations, Cisco recommends that port 443 is used for both the TURN Server and the Web Bridge. However, to set up TCP to UDP interworking on a TURN server, the Web Bridge and TURN Server must listen on different interface:port combinations.

   To run both the TURN server and the Web Bridge on port 443 requires them to be run on separate servers/VMs, or if on the same server/VM they need to be on different interfaces and different subnets.

   If this is not possible then select a non-standard port for the TURN server, for example: **`turn tls 447`** and use the **`tcpPortNumberOverride`** parameter to configure the port on the Call Bridge (see step 7).

6. Since media sent over TCP is encrypted using TLS, a certificate is required on each TURN server that carries out TCP to UDP interworking. The certificate should be signed by the same CA as that used for the Web Bridge.

   a. Generate a private key and the Certificate Signing Request (.csr) file for the TURN server. For information on how to generate a private key and .csr file, refer to the Certificate Guidelines.

      Note: The public key is created and held within the .csr file.

   b. Submit the .csr file to the CA for signing.

   c. SSH into the MMP

   d. Disable the TURN server interface before assigning the certificate

      `turn disable`

   e. Upload the signed certificate and intermediate CA bundle (if any) to the Meeting Server using SFTP.

   f. Check that the certificate (and certificate bundle) and the private key match

      `pki match <certicatefile> <cert bundle/CA cert> [<CA cert>]`

g. Check that the specified certificate is signed by the root CA using the certificate bundle to determine the chain of trust

```
pki verify <certicatefile> <cert bundle/CA cert> [<CA cert>]
```

h. Assign the certificate (and certificate bundle) and private key pair to the TURN server

```
turn certs <keyfile> <certificatefile> [<cert-bundle>]
```

i. Re-enable the TURN server

```
turn enable
```

7. If in step 5 you set a non-standard port for TCP on the TURN Server, use the API parameter **tcpPortNumberOverride** on object /turnServers/<turn Server id> to configure this value on the Call Bridge.

   For example, for the TURN server which will interwork the media, POST to the Call Bridge's /turnServers node the following parameter values replaced by your values:

   ```
   tcpPortNumberOverride = 447
   ```

   Note: This parameter is not required for configured Lync Edge servers, where the TCP port number can always be determined automatically.

8. Use the Web Admin interface to configure the settings through which the Call Bridge communicates with the TURN server, see Chapter 12.

5 LDAP configuration

If you plan for users to utilize the Cisco Meeting Apps to connect to the Meeting Server, then you must have an LDAP server (currently Microsoft Active Directory, OpenLDAP or Oracle Internet Directory LDAP3, see note below). The Meeting Server imports the User accounts from the LDAP server.

You can create user names by importing fields from LDAP, as described in this section. The passwords are not cached on the Meeting Server, a call is made to the LDAP server when a Cisco Meeting App authenticates, and therefore passwords are managed centrally and securely on the LDAP server.

Note: From version 2.1, the Meeting Server supports Oracle Internet Directory (LDAP version 3). This must be configured through the API, not the Web Admin interface.

To configure the Meeting Server to support Oracle Internet Directory, the Meeting Server should not use the LDAP paged results control in search operations during LDAP sync. POST to `/ldapServers` or PUT to `/ldapServers/<ldap server id>` the request parameter `usePagedResults` set to false.

5.1 Why use LDAP?

Using LDAP to configure the Meeting Server is a powerful and scalable way to set up your environment: defining your organization's calling requirements within the LDAP structure minimizes the amount of configuration required on the Meeting Server.

The server uses the concept of filters, rules and templates, which allow you to separate users into groups, for example:

- Everyone in the HR department
- Staff at grade 11 and above
- Job title = 'director'
- People whose surname starts with 'B'

5.2 Meeting Server settings

The examples in this section explain how to configure a single LDAP server (in this case Active Directory), using the Web Admin interface on the Meeting Server. However, the Meeting Server supports multiple LDAP servers which can be configured via the API, see the LDAP Methods section in the API Reference guide.

When configuring a cluster of Call Bridges, the simplest method is to use the API. If configuring multiple Call Bridges via the Web Admin interface, each must have identical configuration.

Note: The Web Admin Interface only allows you to configure one LDAP server.

To set up the Meeting Server to work with Active Directory, follow these steps:

1. Sign in to the Web Admin Interface and go to **Configuration > Active Directory**.
2. Configure the connection to the LDAP server in the first section with the following:
   - Address = this is the hostname or IP address of your LDAP server
   - Port = usually 636
   - Username = the Distinguished Name (DN) of a registered user. You may want to create a user specifically for this purpose.
   - Password = the password for the user name you are using
   - Secure Connection = tick this box for a secure connection

     For example:

     ```
     Address:  ldap.example.com
     Port:     636
     Username: cn=Fred Bloggs,cn=Users,OU=Sales,dc=YourCompany,dc=com
     Password: password
     ```

     Note: For further details of the permissions required by the user name and password credentials, see Appendix F.

     Note: The Meeting Server supports secure LDAP. By default the LDAP server runs on port 636 for secure communications and port 389 for insecure communications. The Meeting Server supports both, but we recommend using 636. Note that you must select Secure Connection (see above) for communications to be secure: using port 636 alone is not enough.

     Note: When LDAP servers are configured with secure connection, connections are not fully secure until TLS certificate verification has been configured using the **`tls ldap`** command on the MMP.
3. Type the Import Settings which will be used to control which users will be imported.
   - Base Distinguished Name = the node in the LDAP tree from which to import users. The following is a sensible choice for base DN to import users

     ```
     cn=Users,dc=sales,dc=YourCompany,dc=com
     ```

- Filter = a filter expression that must be satisfied by the attribute values in a user's LDAP record. The syntax for the Filter field is described in rfc4515.

  A rule for importing people into the main database might reasonably be 'import anyone with an email address', and this is expressed by the following filter:

  ```
  mail=*
  ```

  For testing purposes you may want to import a named user (e.g. fred.bloggs)and a group of test users whose mail address starts with “test”; for example:

  ```
  (|(mail=fred.bloggs*)(mail=test*))
  ```

  If you wanted to import everyone apart from one named user (e..g. fred.bloggs), use this format:

  ```
  (!(mail=fred.bloggs*))
  ```

  To import users that belong to a specific group, you can filter on the memberOf attribute. For example:

  ```
  memberOf=cn=apac,cn=Users,dc=Example,dc=com
  ```

  This imports both groups and people that are members of the APAC group.

  To restrict to people (and omit groups), use:

  ```
  (&(memberOf=cn=apac,cn=Users,dc=Example,dc=com)(objectClass=person))
  ```

  Using an extensible matching rule (LDAP_MATCHING_RULE_IN_CHAIN / 1.2.840.113556.1.4.1941), it is possible to filter on membership of any group in a membership hierarchy (below the specified group); for example:

  ```
  (&(memberOf:1.2.840.113556.1.4.1941:=cn=apac,cn=Users,dc=Example,
  dc=com)(objectClass=person))
  ```

  Other good examples which you can adapt to your LDAP setup include:

  Filter that adds all Person and User except the ones defined with a !

  ```
  (&(objectCategory=person)(objectClass=user)(!(cn=Administrator))(!
  (cn=Guest))(!(cn=krbtgt)))
  ```

  Filter that adds same as above (minus krbtgt user) and only adds if they have a sAMAccountName

  ```
  (&(objectCategory=person)(objectClass=user)(!(cn=Administrator))(!
  (cn=Guest))(sAMAccountName=*))
  ```

  Filter that adds same as above (Including krbtgt user) and only adds if they have a sAMAccountName

  ```
  (&(objectCategory=person)(objectClass=user)(!(cn=Administrator))(!
  (cn=Guest))(!(cn=krbtgt))(sAMAccountName=*))
  ```

This filter only imports specified users within (|(tree

```
(&(objectCategory=person)(objectClass=user)(|(cn=accountname)
(cn=anotheraccountname)))
```

Global Catalog query to import only members of specified security group (signified with =cn=xxxxx

```
(&(memberOf:1.2.840.113556.1.4.1941:=cn=groupname,cn=Users,
dc=example,dc=com)(objectClass=person))
```

4. Set up the Field Mapping Expressions

   The field mapping expressions control how the field values in the Meeting Server's user records are constructed from those in the corresponding LDAP records. Currently, the following fields are populated in this way:

   - Display Name
   - User name
   - space Name
   - space URI user part (i.e. the URI minus the domain name)
   - space Secondary URI user part (optional alternate URI for space)
   - space call id (unique ID for space for use by WebRTC client guest calls)

   Field mapping expressions can contain a mixture of literal text and LDAP field values, as follows:

   ```
   $<LDAP field name>$
   ```

   As an example, the expression

   ```
   $sAMAccountName$@example.com
   ```

   Generates:

   ```
   fred@example.com
   ```

For more information see More Information on LDAP Field Mappings.

Note: Each imported user must have a unique XMPP user ID (JID), constructed using the JID field in the Field Mapping Expressions section of the **Configuration > Active Directory**. In order to construct a valid JID, any LDAP attribute used in the JID field mapping expression must be present in each LDAP record that is to be imported. To ensure that only records that have these attributes present are imported, we recommend that you include presence filters (i.e. those of the form (<attribute name>=*)) using a '&' (AND) in the Filter field under Import Settings for each attribute used in the JID field mapping expression.

For example, suppose your JID field mapping expression is `$sAMAccountName$@company.com`, and you wish to import users who are members of the group cn=Sales,cn=Users,dc=company,dc=com, an appropriate import filter would be:

```
(&(memberOf=cn=Sales,cn=Users,dc=company,dc=com)(sAMAccountName=*))
```

5. To synchronize with Active Directory, select **Sync now** or activate the synchronization by using the appropriate API call (see the Cisco Meeting Server API Reference Guide).

   Note: that you must manually resynchronize whenever entries in the LDAP server change.

6. View the result of the synchronization by going to **Status > Users**.

It is possible to choose whether to use OU separation when importing from the LDAP server. In the Web Admin Interface, go to **Configuration > Active Directory** and in the **Corporate Directory Settings** section select **Restrict Search to Searcher OU** to enable the search only within the OU of the user account.

5.3 Example

This example assigns a space to a particular group of users and a Call ID for this space using an 88 prefix in front of the regular telephone number.

1. Create the group in the LDAP structure called "space" and assign the required members to that group.

2. Use the following filter which uses the extensible matching rule (LDAP_MATCHING_RULE_IN_CHAIN / 1.2.840.113556.1.4.1941) to find all the users that are a member of the "space" group:

```
(&(memberOf:1.2.840.113556.1.4.1941:=cn=space,cn=Users,dc=lync,
dc=example,dc=com)(objectClass=person))
```

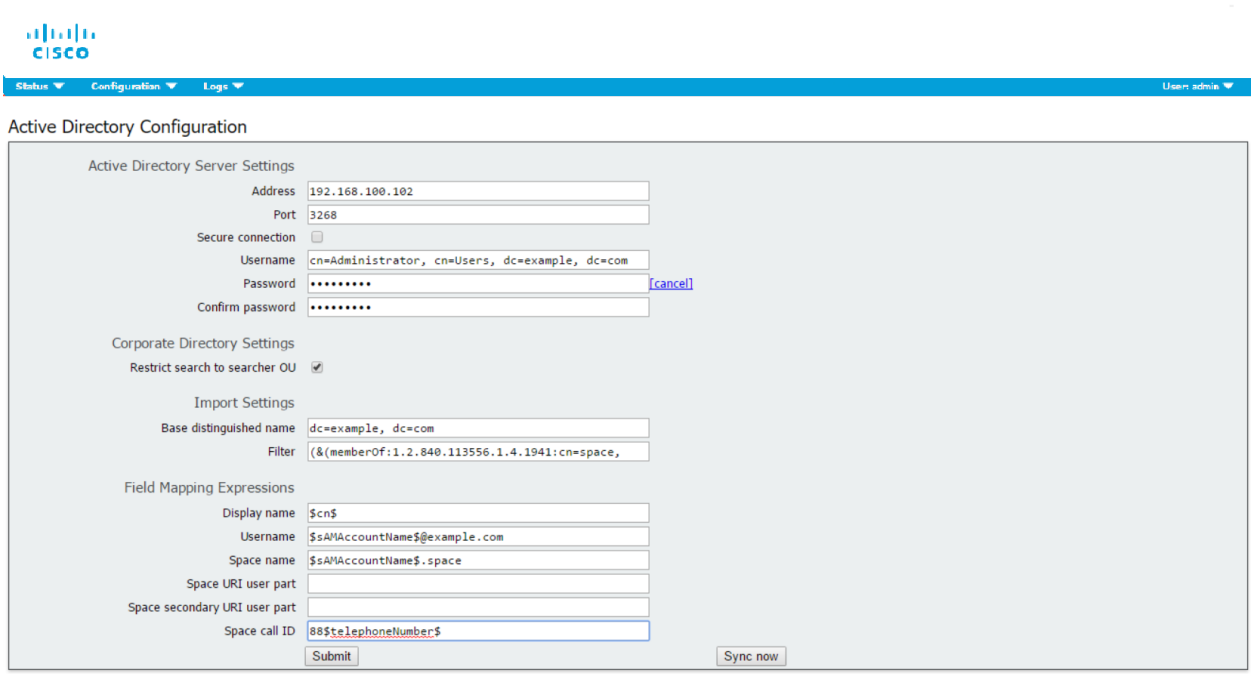

3. Then synchronizing a particular user in the directory called:

   ```
   cn = Fred Blogs
   TelePhoneNumber = 7655
   sAMAccountName = fred.blogs
   ```

   creates the following space which can be viewed on the **Status > Users** page.

Name	XMPP id
Fred Blogs	`fred.blogs@xmpp.example.com`

   And the following space that can be viewed on the **Configuration > space** page.

Name	URI user part
fred.blogs	`fred.blogs.space`

5.4 Enforcing passcode protection for non-member access to all user spaces

When spaces are auto-generated via an LDAP sync, they are all created without a passcode. By default **`nonMemberAccess`** is set to **`true`** so that the existing behavior remains unchanged, no passcode is required to access the space and non-members are able to access the created spaces.

Setting **`nonMemberAccess`** to **`false`** allows a company to enforce passcode protection for non-member access to all user spaces.

To ensure the member must configure non-member access and set a passcode as part of the LDAP sync:

- Either POST to **`/ldapSources`** or PUT to **`/ldapSources/<ldap source id>`** the request parameter **`nonMemberAccess`** set to **`false`**.
- To retrieve the **`nonMemberAccess`** setting, use GET on **`/ldapSources/<ldap source id>`**.

Note: Spaces created before version 2.4 (when this parameter was introduced) are unaffected by any LDAP syncs.

6 Dial plan configuration – overview

6.1 Introduction

For the Meeting Server to be integrated in a SIP, Lync and voice environment, connections need to be set up from the SIP Call Control, Lync FE Server and Voice Call Control to the Meeting Server. Changes to the call routing configuration on these devices is required in order to correctly route the calls that require the Meeting Server.

Figure 11 assumes a company deployment which has a mix of SIP video endpoints, Lync clients and IP phones: the Meeting Server enables connectivity between Lync clients and SIP video endpoints, and between Lync clients and IP phones.

The SIP video endpoints are configured on a domain called vc.example.com and the Lync clients on example.com. You will need to adapt the example, as appropriate.

Note: Although this figure and subsequent diagrams in this Deployment Guide use an Acano X series deployment as the example, the instructions apply equally to virtualized deployments.

Figure 11: Example deployment for dial plan configuration

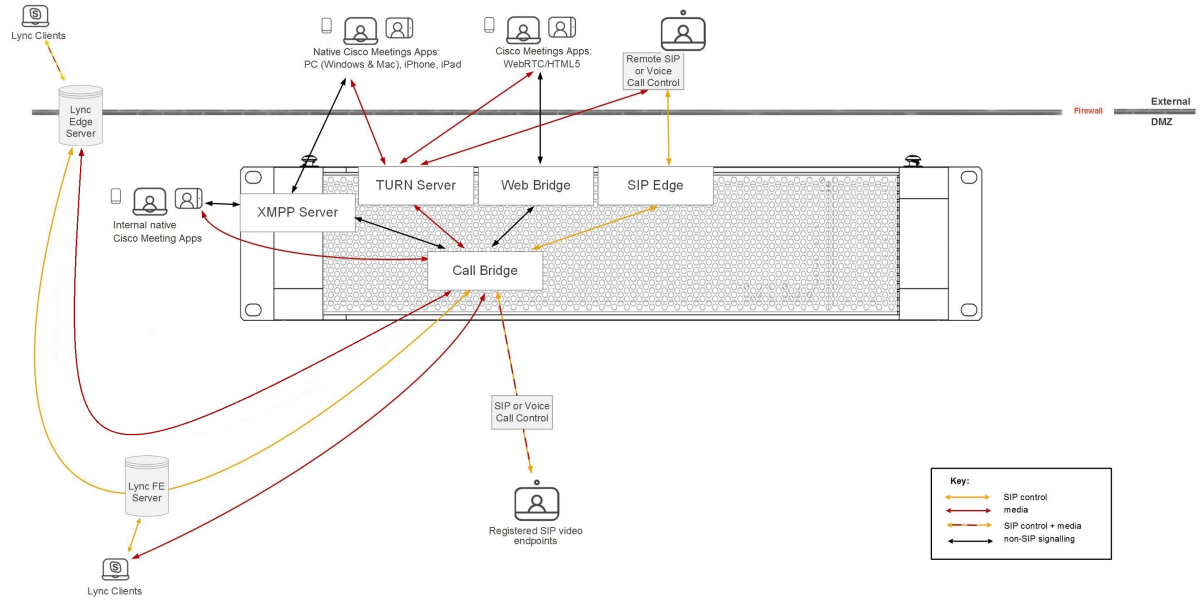

As shown in the figure above, the Lync FE server needs a trusted SIP Trunk to the Meeting Server, configured to route calls originating from Lync clients through to Meeting Server spaces, Cisco Meeting App users (native and WebRTC) and also SIP video endpoints. The subdomains vc.example.com (for SIP video endpoints) and meetingserver.example.com (for

spaces and Cisco Meeting Apps) should be routed through this trunk from the Lync FE server to the Meeting Server.

Note: Connections to Office 365 or on-premise Lync deployments in another organization, should route to a Cisco Expressway. See the Expressway deployment guides for more information.

The SIP Call Control platform needs a SIP trunk set up to route calls to the example.com domain (for Lync Clients) and meetingserver.example.com (for spaces and Cisco Meeting Apps) to the Meeting Server.

The Meeting Server requires a dial plan to route calls with domain example.com to the Lync FE server and subdomain vc.example.com to the SIP Call Control platform.

The next section discusses the two configuration pages in the Web Admin interface of the Meeting Server that determine how the Meeting Server handles incoming calls and outbound calls.

Following this chapter, Chapter 7 and Chapter 8 provide step-by-step instructions on configuring the total solution.

6.2 Web Admin Interface configuration pages that handle calls

This section explains the configuration pages in the Web Admin interface that the Meeting Server uses to determine how to handle each call.

Two configuration pages in the Web Admin Interface control how the Meeting Server behaves for incoming and outgoing calls: **Outbound calls** and **Incoming calls**. The Outbound Calls page controls how outbound calls are handled; the Incoming calls page determines whether incoming calls are rejected. If they are not rejected, but matched and forwarded, then information about how to forward them is required and the Incoming Calls page has two tables – one to configure matching/rejection and the other to configure the forwarding behavior.

6.2.1 Outbound calls page

The Outbound Calls page allows you to configure appropriate dial plans comprising a number of dial plan rules. A dial transform can be applied to Outbound calls to control the routing of the outbound calls, see Dial Transforms.

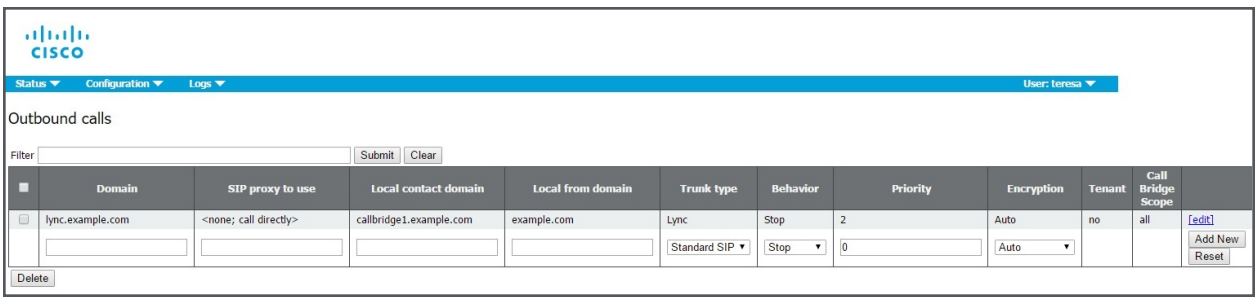

Domain: the domain to match in order to apply the dial plan rule; either a complete value (e.g. "example.com") or a "wildcarded" one (e.g. "*.com").

SIP proxy to use: each entry/rule in a dial plan matches on the Domain of the outgoing call (see below) and determines which SIP proxy to use (or whether it is a direct call).

Local contact domain: is the domain that will be used for the contact URI for calls using this dial plan rule.

CAUTION: If you are using Lync, we suggest that you use the **Local contact domain**. If you are not using Lync we recommend that the **Local contact domain** field is left blank to avoid unexpected issues with the SIP call flow.

CAUTION: For each Lync domain you need to create an outbound rule – follow the procedures described in this section. If you have many Lync domains you can consider creating an outbound rule with a wildcard domain.

Local from domain: is the domain the call uses as its originator ID/Caller ID.

Trunk type: usually, you set up rules to route calls out to third party SIP control devices such as CiscoExpressway, Avaya Manager or Lync servers. Therefore, there are currently three types of SIP trunks you can configure: Standard SIP, Avaya and Lync.

Note: A common use of the Meeting Server is with an Avaya PBX; these calls will be audio-only. However, the Meeting Server does not impose this restriction on interoperability with Avaya products (some of which support video also): therefore a call of type of 'avaya' does not imply that the call is audio-only.

Behavior and **Priority**: Dial plan rules are tried in the order of the Priority values. If a rule is matched, but the call cannot be made, then other lower priority rules may be tried. If a rule has a behavior of STOP, then no further rules are used.

Encryption: select from **Auto**, **Encrypted**, **Unencrypted**.

CAUTION: The default **Encryption** behavior mode is **Auto**. Ensure all "Lync" outbound dialing rules are explicitly set to **Encrypted** mode to prevent the Call Bridge attempting to use unencrypted TCP for these connections in the event of the TLS connection attempt failing.

6.2.2 Incoming call page: call matching

The top table in the Incoming Call page is the Call Matching table. The rules defined in the Call Matching table govern how the Meeting Server handles incoming SIP calls. Any call routed to the Meeting Server on any domain can be tested for a match for IVRs, Cisco Meeting App users or for preconfigured spaces on that server.

The example Call matching rule below seeks to match all calls coming in on the **meetingserver.example.com** domain to both Cisco Meeting App users and spaces.

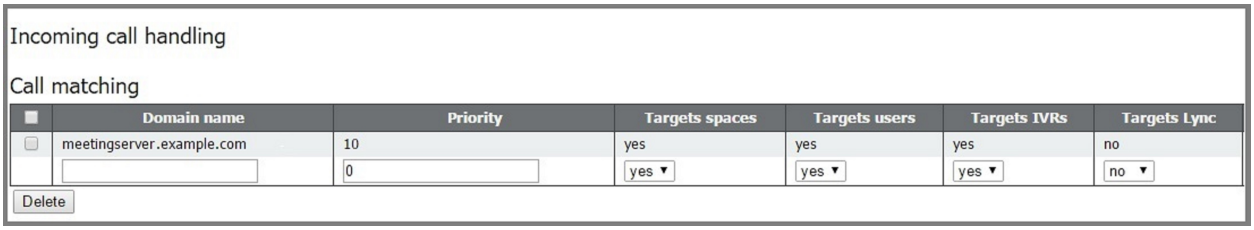

For example, if the incoming call was to **name.space@meetingserver.example.com** and there was a configured space called **name.space** the call would be routed to the space with that name. If the incoming call was to **firstname.lastname@meetingserver.example.com** the call would be routed to that user with that first and last name.

Alternatively, you can choose not to route calls to users or spaces on a per domain basis; that is, you can use one incoming domain for spaces and another for users.

It is recommended that rules are created for every domain expected for incoming calls. With some call control solutions the domain may be the IP address or hostname of the server. In these cases the highest priority domain is expected to be the main domain, with IP address and hostname rules having lower priority.

Rules with a higher priority value are matched first. In cases where multiple rules have the same priority then matching occurs based on alphabetical order of the domain.

After a rule is executed rules further down the list are ignored for the call.

If all Call Matching rules fail, the next table (Call Forwarding) is used as described in the next section.

Points to note:

- Matching for space and/or users is only done on the part of the URI before the @.
- The highest priority rule that matches a space is used to form the URI in the invitation text. It is expected that the highest priority rules are for the deployment as a whole rather than for individual IP addesses or hostnames.
- Do not leave the Domain field blank in a rule, otherwise the Call Bridge will refuse the call.
- No rules in the Call matching table will result in all domains being matched.

6.2.3 Call forwarding

If an incoming call fails to match any of the rules in the Call Matching table, the call will be handled according to the Call Forwarding table. In this table you can have rules to decide whether to reject the call outright or to forward the call in bridge mode, for example resolving to a Lync conference. By defining rules, you decide whether to forward the call or not. It might be appropriate to “catch” certain calls and reject them.

Rules can overlap, and the **Domain matching pattern** can include wildcards, for example: exa*.com; but do not use "*" as a match all, otherwise you will create call loops. Order rules using the **Priority** value; higher numbered rules are tried first.

For calls that will be forwarded, you can rewrite the destination domain using the **Forwarding domain**. A new call is created to the specified domain. The **Caller ID** setting allows the forwarded call to either preserve the original calling party's ID or to generate a new one. Select **pass through** to preserve the calling party's ID or **use dial plan** to generate a new calling party ID according to your call routing configuration.

The example Call Forwarding rule below forwards calls for the domain lync.example.com and the routing is determined by the call routing rules.

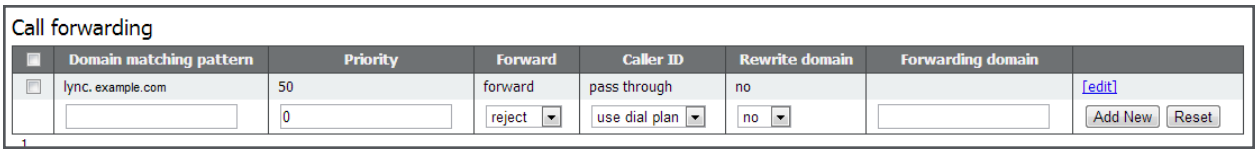

An incoming call is terminated if does not match any of the rules in the Call Matching table and does not match any of the **Domain matching patterns** in the Call Forwarding table.

6.3 Dial Transforms

Dial Transforms are applied to outgoing calls prior to the Outbound rules taking effect. When dial transforms are applied, the outbound dial plan rules are applied to the transformed number. Dial Transforms only affect Outbound calls, they do NOT affect gateway calls.

There are three stages to the transform:

- A "type" is applied, which defines the type of preprocessing to apply to the transform.
  - Raw: produces one component – $1
  - Strip: removes dots, dashes, spaces and produces one component – $1
  - Phone: use to transform to an international phone number – produces two components $1county code and $2number

  Note: A phone URI is recognized as a purely numeric string (optionally prefixed by a '+') when it begins with a valid international dial code (e.g. 44 for UK or 1 for US) followed by the correct number of digits for a phone number for that region.
- The components are matched using regular expressions to see if the rule is valid
- An output string is created from the components according to the defined transform

Examples

Example	Type	Match	Transform
For US numbers, use 'vcs1' directly	Phone	($1/01/)	$2@vcs1
For UK numbers, add a prefix and use 'vcs2'	Phone	($1/44/)	90044$2@vcs2
For UK numbers starting with a 7, add '90044' as a prefix, add '123@mobilevcs' as a suffix	Phone	($1/44/)($2/^7/)	90044$2{}123@mobilevcs
For unrecognized all-digit strings, use '@vcs3' as a suffix	Strip	($1/(\d){6,}/)	$1@vcs3
Replace + with 00	Strip	($1/\+(\d)+/)	$1{/\+/00/}
Replace an alphanumeric regex e.g. (.*)@example.com and replace with \1.endpoint@vc.example.com	Raw	($1/(.*) @example.com/)	${/@example.com$/ .endpoint@vc.example.com/}

For a single Meeting Server, use the **Configuration > Outbound Calls** page in the Web Admin Interface to control how dialed numbers are transformed. If a match expression is provided, the regular expression determines whether the specified transform expression is applied

For example, the dial plan in the screen shot below ensures that outbound "+1" (US) calls use one Call Bridge and +44 (UK) calls use another.

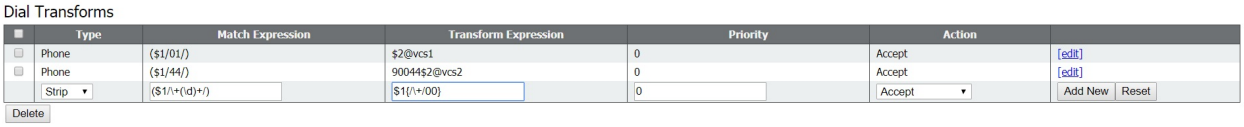

7 Dial plan configuration – SIP endpoints

7.1 Introduction

This chapter describes the configuration to enable SIP video endpoints to dial into a meeting hosted on the Meeting Server. Work through the steps in the order provided, adapting the example as appropriate.

7.2 SIP video endpoints dialing a meeting hosted on the Meeting Server

This first step considers the configuration required on the call control device and on the Meeting Server to direct SIP video endpoints to meetings hosted on the Meeting Server.

Figure 12: Example of SIP video endpoints calling into Meeting Server hosted calls

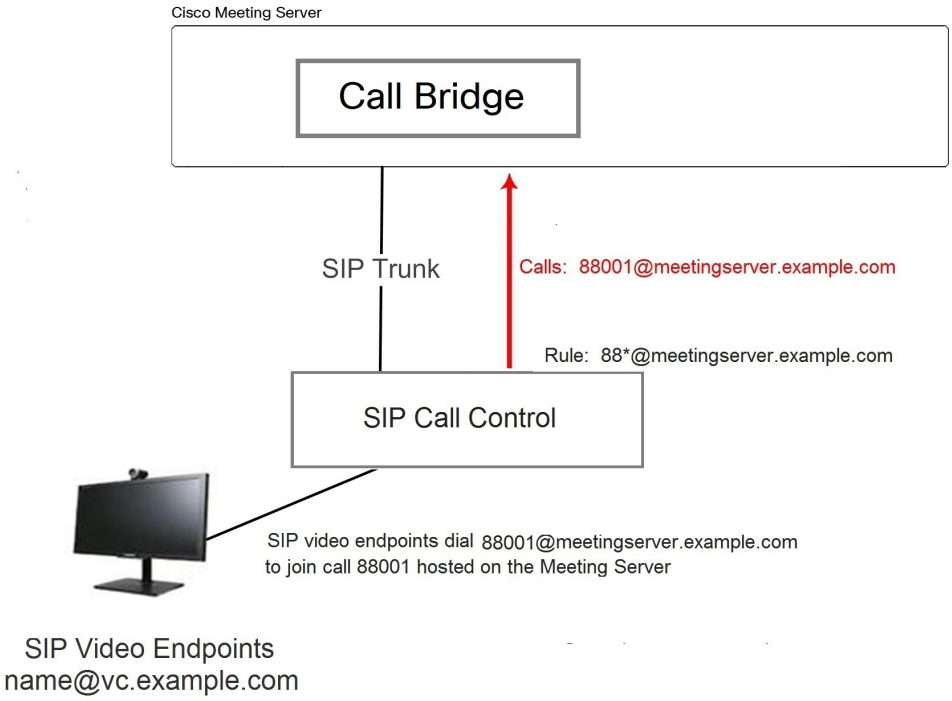

7.2.1 SIP call control configuration

This example assumes the SIP Call Control is a Cisco VCS, but similar steps are required on other call control devices, for example using the Cisco Unified Communications Manager, see the Cisco Meeting Server with Cisco Unified Communications Manager Deployment Guide.

1. Sign in to the VCS as an administrator.
2. Set up a zone to route calls to the Meeting Server
   a. Go to **VCS Configuration > Zones > New**.
   b. Create the zone with the following:
      - H.323 Mode = Off.
      - SIP Mode = On
      - SIP Port = 5060 (5061 if using TLS)
      - SIP Transport = TCP or TLS, as appropriate
      - SIP Accept Proxied Registrations = Allow
      - Authentication Policy = Treat as authenticated
      - SIP Authentication Trust Mode = Off
      - Peer 1 Address = the IP address of the Call Bridge
3. Add a search rule to route calls to the Meeting Server. For example to route any calls on SIP endpoints to a meeting on the Meeting Server using the domain **`meetingserver.example.com`**.
   a. Go to **VCS Configuration > Dial Plan > Search rules**
   b. Give the rule a suitable name, e.g. **`Route EPs to Meeting Server`**.
   c. Set the following:
      - Source = Any
      - Request Must Be Authenticated = No
      - Mode = Alias pattern match
      - Pattern Type = Regex
      - Pattern String = **`.*@meetingserver.example.com`**
      - Pattern Behavior = Leave
      - On Successful Match = Stop
      - Target = the zone you created for the Meeting Server.

7.2.2 Meeting Server configuration

1. Sign in to the Web Admin Interface on the Meeting Server.
2. Either create a space on the Meeting Server for endpoints to dial into:
   a. Go to **Configuration >space**
   b. Add a space with:

- **Name** =<string>, for example. `Call 001`
- **URI** =<user part of the URI>, for example. `88001`

or use an already existing space.

Note: spaces can also be created or modified from the API. See the API Reference guide.

3. Add an inbound dial plan rule for incoming calls to the Meeting Server.
   a. Go to **Configuration > Inbound Calls** and add a dial plan rule with the following details:
      - **Domain name** = <FQDN of the Meeting Server>, for example `meetingserver.example.com`
      - **Targets spaces** = `yes`
      - **Targets IVRs** = `yes`
      - optional **Targets users** = `yes`
      - **Targets Lync** = `yes` Note: this is required later in Section 8.1.2

      Note: See Section 6.2.2 for more information on the **Inbound calls** page of the Web Admin interface.

4. Add an outbound dial plan rule for outbound calls to SIP endpoints via the VCS.
   a. Go to **Configuration > Outbound Calls** and add a dial plan rule with the following details:
      - **Domain** =<domain to match> such as `example.com` or `*.com`
      - **SIP Proxy to use** = <the IP address or FQDN of your VCS>
      - **Local Contact Domain** =

        Note: The local contact domain field should be left blank unless setting up a trunk to Lync (as in Section 8.1.2).

      - **Local From Domain** = <FQDN of the Meeting Server>
      - **Trunk type**=`Standard SIP`.

      Note: See Section 6.2.1 for more information on the **Outbound calls** page of the Web Admin interface.

SIP video endpoints can now dial into a call 88001 hosted on the Meeting Server by dialing `88001@meetingserver.example.com`, and the Meeting Server can call out to SIP endpoints.

Before moving onto creating dial plans for Lync in Chapter 8, consider whether to:

- configure the media encryption setting, see Section 7.3,
- enable TIP support for Cisco CTS endpoints, see Section 7.4,
- configure an Interactive Voice Response (IVR), see Section 7.5.

7.3 Media encryption for SIP calls

The Meeting Server supports media encryption for SIP connections, including Lync calls, made to or from the Meeting Server. This is configured in the **Configuration > Call settings** page in the Web Admin Interface.

1. Sign in to the Web Admin Interface and go to **Configuration > Call settings**
2. Select the appropriate **SIP media encryption** setting (**allowed**, **required** or **disabled**).
3. Change the bandwidth settings for SIP, CMA (Cisco Meeting App) or Server reflexive.
4. To select applying these changes to SIP calls already in progress, click the **Apply to Active Calls** button at the end of the page, or to select applying these changes to future SIP calls click the **Submit** button.

Note: The SIP Encryption field in the Web Admin Interface **Configuration > Outbound Calls** page allows you to set the SIP control encryption behavior for each Outbound Calls rule. This separates the control and media encryption behavior, allowing a TLS control connection to be used in the absence of media encryption; you can also set the bahavior via the API.

7.4 Enabling TIP support

If you use endpoints such as the Cisco CTS range, you need to select TIP protocol support. Enable it as follows:

1. In the Web Admin Interface go to **Configuration > Call settings** and in the SIP Settings section, set TIP (Telepresence Interoperability Protocol) to **enabled**.

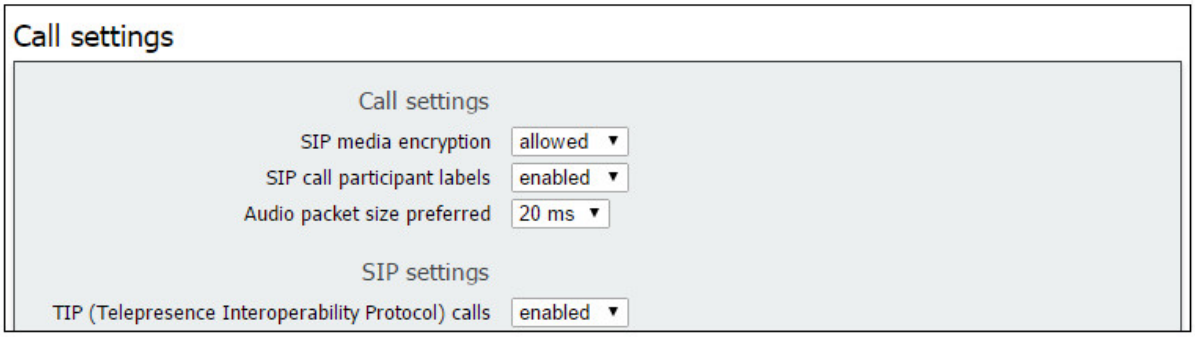

2. Set both SIP Bandwidth Settings to at least 4000000.

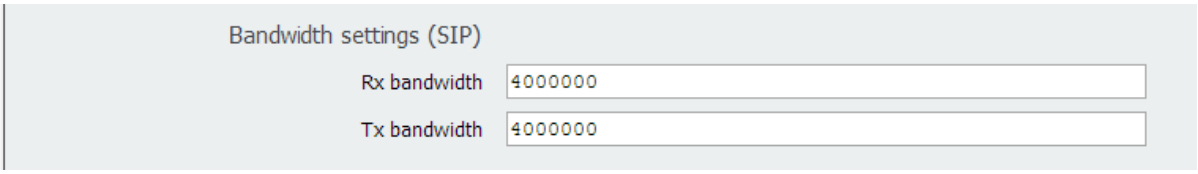

3. Click **Submit**.

7.5 IVR configuration

You can configure an Interactive Voice Response (IVR) to manually route to pre-configured calls. Incoming calls can be routed to the IVR where callers are greeted by a prerecorded voice message inviting them to enter the ID number of the call or space that they want to join. Video participants will see a welcome splash screen. After entering the ID, users are routed to the appropriate call or space, or prompted to enter a PIN if the call or space has one. (Callers are disconnected after the third incorrect call ID.)

If you intend to use an IVR follow these instructions:

1. Sign into the Web Admin Interface and go to **Configuration > General**.
2. In the **IVR** section, configure the following:
   - **IVR numeric ID** = <numeric call ID that users call to reach the IVR>
   - **Joining scheduled Lync conferences by ID**= “not allowed” or “allowed” depending on your policy.
3. On **Configuration > Incoming Calls** set **Target IVRs** = "yes" to match incoming calls to the IVR.
4. Configure the appropriate routing on your SIP Call Control to ensure that calls to the numbers set in the previous step are routed to the Meeting Server.

7.6 Next steps

Now follow the steps in Chapter 8 to configure dial plans to integrate Meeting Server with Lync deployments.

8 Dial plan configuration – integrating Lync/Skype for Business

Throughout this chapter, references to Microsoft Lync also mean Microsoft Skype for Business.

Note: For Call Bridge integration with Lync Edge, the Call Bridge needs its own login account. For each Lync call to or from the Call Bridge, the server requests TURN resources from the Lync Edge using that account. Until that call is disconnected, that resource is considered "Used" from a Lync point of view. Lync will only allow up to 12 TURN allocations per user account; therefore, with 1 registration, only 12 calls are possible.

8.1 Lync clients dialing into a call on the Meeting Server

This section details the configuration required to enable Lync endpoints to join a meeting hosted on the Meeting Server. It uses the same call number/URI as used in Section 7.2; adapt the example as appropriate.

Figure 13: Example Lync clients calling into Meeting Server hosted meetings

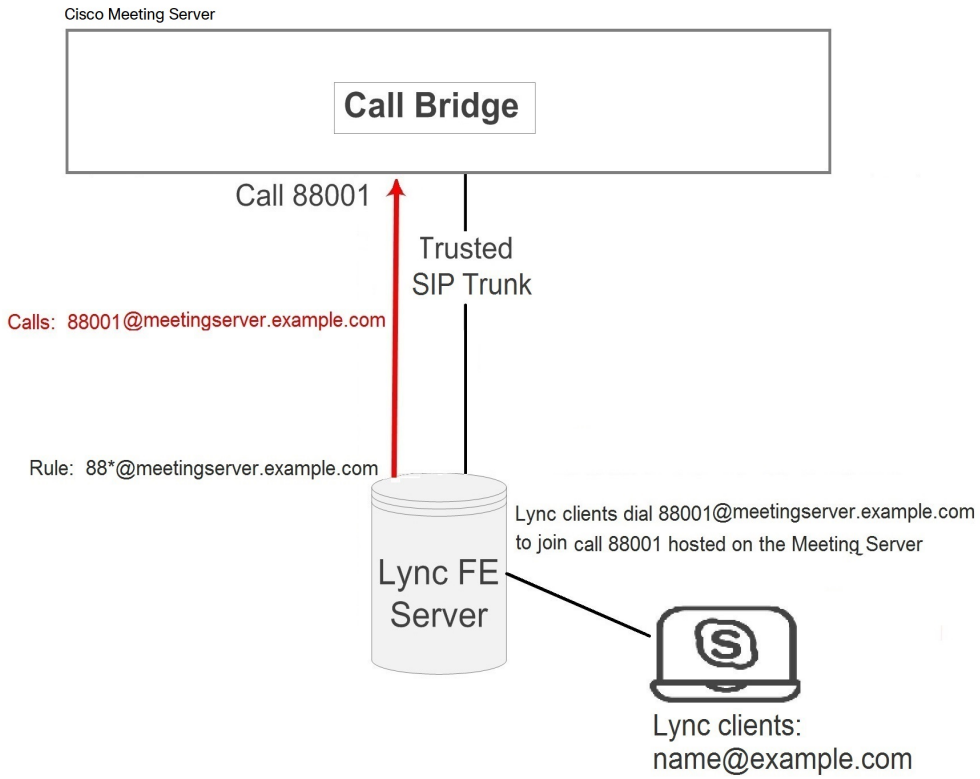

8.1.1 Lync Front End (FE) server configuration

CAUTION: This section provides an example for configuring a static route between a Lync FE server and the Meeting Server, it is only a guideline and is not meant to be an explicit set of instructions for you to follow. Cisco strongly advises you to seek the advice of your local Lync server administrator on the best way to implement the equivalent on your server's configuration.

Note: Before configuring a static route from the Lync FE server, ensure that you have installed certificates on the Meeting Server which will be trusted by the Lync FE server – as described in the Certificate Guidelines.

To route calls originating from Lync clients to the Meeting Server, add a Lync static route pointing to the Meeting Server. This involves setting the Meeting Server as a trusted application for the Lync FE server and adding the static route.

1. Open the Lync Server Management Shell.
2. Create a new application pool that will contain the Meeting Server as a trusted application.

   ```
   New-CsTrustedApplicationPool -Identity fqdn.meetingserver.com -ComputerFqdn fqdn.meetingserver.com -Registrar fqdn.lyncserver.com -site 1 -RequiresReplication $false -ThrottleAsServer $true -TreatAsAuthenticated $true
   ```

   Replacing

   - `fqdn.meetingserver.com` with the FQDN of the Meeting Server, the identity MUST be the CN specified in the Call Bridge's certificate.
   - `fqdn.lyncserver.com` with your Lync FE Server or FE Pool FQDN
3. Add the Meeting Server as a trusted application to the application pool.

   ```
   New-CsTrustedApplication -ApplicationId meetingserver-application -TrustedApplicationPoolFqdn fqdn.meetingserver.com -Port 5061
   ```

   Replacing

   - `meetingserver-application` with name of your choice
   - `fqdn.meetingserver.com` with the FQDN of the Meeting Server
4. Create the static route between the Meeting Server and the Lync FE server.

   ```
   $x=New-CsStaticRoute -TLSRoute -Destination "fqdn.meetingserver.com" -MatchUri "meetingserver.example.com" -Port 5061 -UseDefaultCertificate $true
   ```

   Replacing

   - `fqdn.meetingserver.com` with your FQDN of the Meeting Server
   - `meetingserver.example.com` with the URI matching the domain used for all of your Meeting Server calls.

5. Add the new static route to the existing collection of static routes
   `Set-CsStaticRoutingConfiguration -Identity global -Route @{Add=$x}`
6. Optional. Before enabling the static route, consider changing the default screen resolution for Lync calls from the default of VGA to HD720p. To enable HD720p on Lync:
   `Set-CsMediaConfiguration -MaxVideoRateAllowed Hd720p15M`
7. Enable the new static route.
   `Enable-CsTopology`

Note: Users may have to logout and login again to update to the new HD720p setting, all other settings are automatic and should work within a few minutes.

8.1.2 Adding a dial plan rule on the Meeting Server

1. Sign in to the Web Admin Interface of the Meeting Server, go to **Configuration > Outbound Calls**
2. At the bottom of the Outbound calls table, create a new dial plan rule
   a. In the **Domain** field, enter the Lync domain that will be matched for calls that need to be sent to Lync. For example, **`example.com`**
   b. **SIP Proxy to Use** field, enter the address (IP address or FQDN) of the proxy device through which to make the call.
      - Either leave this field blank and the server will perform a DNS SRV lookup for the called domain using **`_sipinternaltls._tcp.<yourlyncdomain>.com`**
      - or enter the IP address or FQDN of the Front End Pool (or Lync sip domain) and the server will first perform a DNS SRV lookup for that defined domain using **`_sipinternaltls._tcp.<Server address>.com`** and then perform a DNS A record lookup for the Host entered if the SRV lookup fails to resolve
      - or enter the IP address or FQDN of your Lync FE server
   c. **Local Contact Domain** field, enter the FQDN of your Meeting Server. For example: **`meetingserver.example.com`**

      Note: The only case in which this field should be set is when setting up a trunk to Lync; otherwise it should be left blank.

   d. **Local From Domain** field, enter the domain that you want the call to be seen as coming from (the Caller ID) e.g. **`meetingserver.example.com`**

      Note: If you leave **Local From Domain** blank, the domain used for the Caller ID defaults to that entered as the Local Contact Domain.

e. **Trunk Type** field, select `Lync`

f. In the **Behavior** field, select **stop** or **continue** depending on whether the next outbound dial plan rule is tried if this rule fails to result in a connected call.

g. **Priority** field, assign a Priority level to determine the order in which dial plan rules will be applied. Rules with higer prioity vales are applied first.

h. **Encryption** field, select **Auto**, **Encrypted** or **Unencrypted** according to whether encrypted SIP control traffic on calls made via this rule, is enforced.

i. Select **Add New**.

Note: Tenant and Call Bridge scope can only be set through the API.

After completion you should be able to call from the Lync environment to the Meeting Server and from the Meeting Server to Lync.

In the example, the Lync clients can now dial into a call 88001 hosted on the Meeting Server by dialing 88001@example.com.

8.2 Integrating SIP endpoints and Lync clients

To allow SIP endpoints to dial a Meeting Server space, implement the steps in Section 7.2; to allow Lync clients to dial a Meeting Server space, implement Section 8.1.

Then both SIP video endpoint users and Lync client users can enter the same call by dialing `<call_id>@meetingserver.example.com`

Figure 14: Example of SIP video endpoints and Lync clients calling into Meeting Server hosted meetings

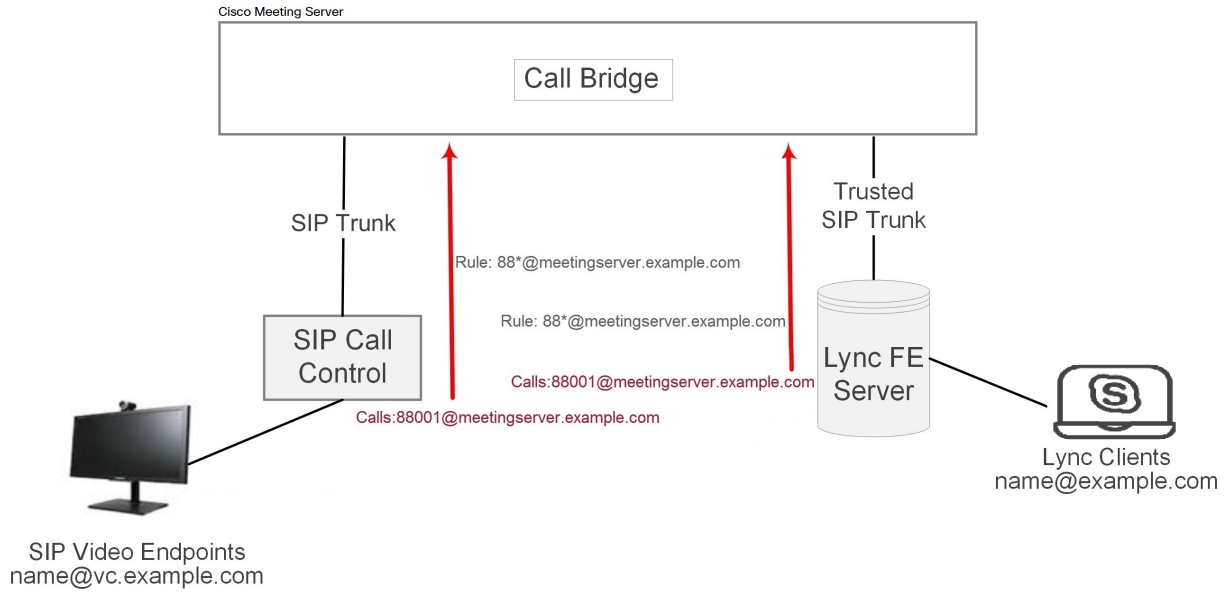

8.3 Adding calls between Lync clients and SIP video endpoints

This section assumes the completion of the configuration described in the two dial plan configuration sections (Section 7.2 and Section 8.1). It expands the example to allow Lync and SIP video endpoints to call each other in a call using the Meeting Server as a gateway to transcode the video and audio (see the figure below).

Note: The Outbound Calls page was used previously to set up a SIP trunk from the Meeting Server to the Cisco VCS. In order to configure the Meeting Server to act as a "point-to-point bridge" between Lync and SIP environments, you need to configure call forwarding as described in this section and also set up a SIP trunk from the Meeting Server to other SIP call control devices you are using such as the Lync FE server, Cisco VCS, CUCM, Avaya CM or Polycom DMA.

Figure 15: Example of SIP video endpoints and Lync clients in calls

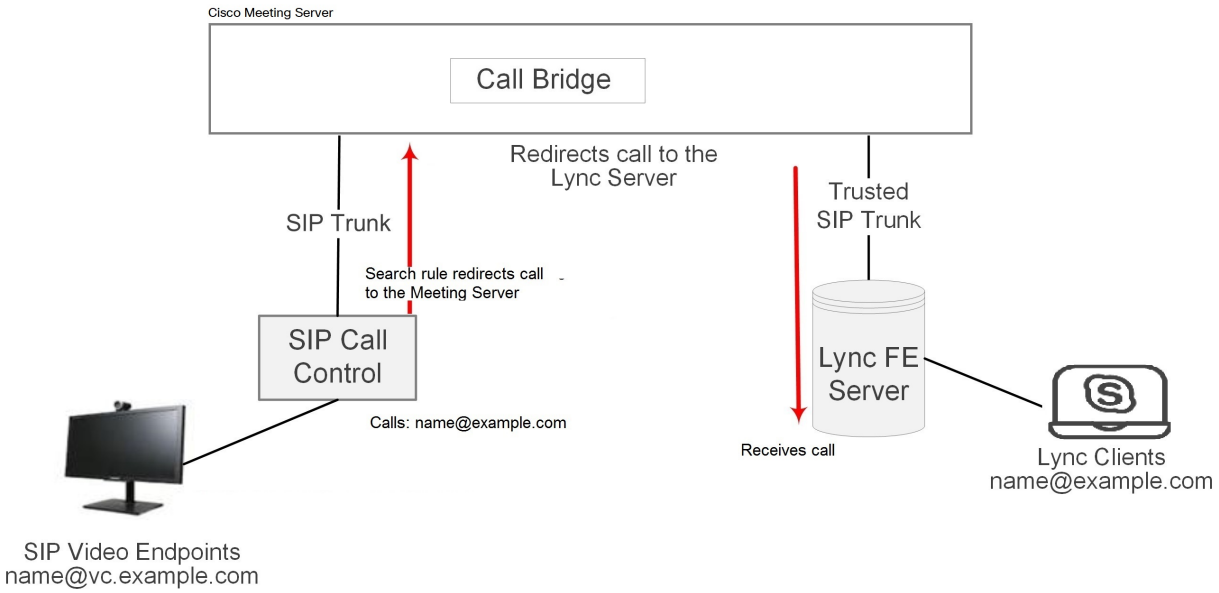

In this example:

- A Lync user can dial **<name>@vc.example.com** to set up a call with a SIP video endpoint, for example **meetingroom1@vc.example.com**.
- A SIP video endpoint can dial **<name>@example.com** to set up a call with a Lync endpoint, for example **roberta.smith@example.com**.

Adapt the example as appropriate.

8.3.1 Lync Front End server configuration

To allow Lync clients to call SIP video endpoints:

- Add a Lync static route pointing to the Meeting Server that will redirect calls for `@vc.example.com`. Follow the steps on creating a Lync static route given in Section 8.1

this will route Lync client calls to SIP video endpoints.

8.3.2 VCS configuration

To allow SIP video endpoint to call Lync clients:

- Add a search rule on the VCS (SIP call control device) to route calls with the suffix `@example.com` to the Meeting Server.

this will route SIP video endpoint calls to Lync clients.

8.3.3 Meeting Server configuration

Create two forwarding rules on the Meeting Server, one to forward calls to SIP endpoints, and the other to forward calls to Lync clients. Then create two outbound dial plan rules one to route outbound calls to SIP endpoints, and the other to route outbound calls to Lync clients.

1. Sign in to the Web Admin Interface and go to **Configuration > Incoming Calls**.
2. In the **Call forwarding** section, create two new rules:
    a. Create a call forwarding rule for calls to vc.example.com
        - **Domain matching pattern** = `vc.exa*.com`
          Wildcards are permitted in any part of a domain matching pattern, but do not use "*" as a match all, otherwise you will create call loops.
        - **Priority** = <number> any value is acceptable, including 0 if there are no other forwarding rules configured. To ensure a rule is always used, set its priority as the highest of any rules configured.

          (Rules are checked in order of priority; highest priority first. If two Domain Matching Patterns match a destination domain, the rule with the higher priority is used.)
        - **Forward** = `forward`

          (If you select "reject", calls that matched the Domain Matching Pattern are not forwarded but terminate.)
        - **Caller ID** = `use dial plan` this will use the domain from the outbound dial plan.
        - **Rewrite Domain** = `no`
          The call will be forwarded using the domain that was called.

          (If you select yes here, you must then complete the **Forwarding domain** field. The original domain will be replaced with the one you enter in **Forwarding domain** before the call is forwarded.)
        - Click **Add new**

b. Create a call forwarding rule for calls to example.com
   - Domain matching pattern = **`exa*.com`**
   - Priority: <number>
   - Forward = **`forward`**
   - Caller ID = **`use dial plan`**
   - Rewrite Domain = **`no`**
   - Click Add new.

3. Go to Configuration>Outbound calls page, create two new rules:
   a. Create a dial plan for calls to domain vc.example.com for SIP endpoints, this is a repeat of step 4 in Section 7.2.2.
      - In the Domain field, enter the SIP domain that will be matched for calls that need to be sent to SIP endpoints. For example, **`vc.example.com`**
      - SIP Proxy to use= <the IP address or FQDN of your VCS>
      - Local Contact Domain =

        Note: The local contact domain field should be left blank.

      - Local From Domain = <FQDN of the Meeting Server>
      - Trunk type=**`Standard SIP`**.
      - Select Add New.
   b. Create a dial plan rule for calls to domain example.com for Lync clients, this is a repeat of Section 8.1.2.
      - In the Domain field, enter the Lync domain that will be matched for calls that need to be sent to Lync. For example, **`example.com`**
      - SIP Proxy to Use field, enter the address (IP address or FQDN) of the proxy device through which to make the call.
        - Either leave this field blank and the server will perform a DNS SRV lookup for the called domain using **`_sipinternaltls._tcp.<yourlyncdomain>.com`**
        - or enter the IP address or FQDN of the Front End Pool (or Lync sip domain) and the server will first perform a DNS SRV lookup for that defined domain using **`_sipinternaltls._tcp.<yourlyncdomain>.com`** and then perform a DNS A record lookup for the Host entered if the SRV lookup fails to resolve
        - or enter the IP address or FQDN of your Lync FE server
      - Local Contact Domain field, enter the FQDN of your Meeting Server. For example: **`meetingserver.example.com`**

Note: The only case in which this field should be set is when setting up a trunk to Lync; otherwise it should be left blank.

- **Local From Domain** field, enter the domain that you want the call to be seen as coming from (the Caller ID), this will be the FQDN of the Call Bridge, e.g. `meetingserver.example.com`

  Note: If you leave **Local From Domain** blank, the domain used for the Caller ID defaults to that entered as the Local Contact Domain.

- **Trunk Type** field, select `Lync`
- In the **Behavior** field, select **stop** or **continue** depending on whether the next outbound dial plan rule is tried if this rule fails to result in a connected call.
- **Priority** field, assign a Priority level to determine the order in which dial plan rules will be applied. Rules with higer prioity vales are applied first.
- **Encryption** field, select **Auto**, **Encrypted** or **Unencrypted** according to whether encrypted SIP control traffic on calls made via this rule, is enforced.
- Select **Add New**.

SIP video endpoints can now call Lync clients by dialing `<name>@example.com`, and Lync clients can call SIP video endpoints by dialing `<endpoint>@vc.example.com`.

8.4 Integrating Cisco Meeting App with SIP and Lync clients

Note: Cisco Meeting App users are not permitted to call out to Lync meetings.

Refer to the sections on LDAP Configuration and Configuring the XMPP server for instructions on configuring your Meeting Server to use the Cisco Meeting App.

If you are using the same LDAP configuration to create both Lync accounts and Cisco Meeting App accounts, and using the Meeting Server as a Lync gateway, then problems can occur with users calling Cisco Meeting App clients rather than the intended Lync client. To prevent this happening set up rules for Call matching and Call forwarding, this is explained below.

For example, assume there is an account `fred@example.com` on the Meeting Server and a `fred@lync.example.com` account on the Lync FE server. If a call arrives at the Meeting Server and no Call matching rules are configured, the Meeting Server will ignore the domain and the call will go to the Meeting Server's `fred@example.com` account. The Meeting Server check's whether there is a user "fred" locally, ignoring the `xxxx` in `fred@xxxx`.

The solution is to configure a **Call matching** rule on the **Incoming Calls** page to match the domain for local Cisco Meeting App users and a **Call forwarding** rule to forward calls to Lync clients. For the **Call matching** rule, set the **Domain name** field to something distinct from the

domain that the Lync FE server uses, for example **example.com**. In the **Call forwarding** section create a rule specifying the Lync domain in the **Domain matching pattern** field, for example **lync.example.com**. A call to **fred@example.com** will reach the Cisco Meeting App user but a call to **fred@lync.example.com** will be forwarded to Fred's Lync client.

8.5 Integrating Lync using Lync Edge service

For NAT traversal using the Lync Edge server, follow the configuration steps in this section to configure Lync Edge settings on the Meeting Server. This is required to support Dual Homed Conferencing or if the Lync Edge performs the TURN/ICE role for Lync calls, rather than the Meeting Server.

8.5.1 Lync Edge call flow

To establish a call from the Meeting Server to the Lync Edge server (see Figure 16 below):

1. The Call Bridge makes a "register" SIP call to the Lync FE server.
2. The "register" is acknowledged.
3. The Call Bridge sends a "service" to the Lync FE server.
4. The FE server returns the URI of the media relay authentication server (MRAS). (The Lync Edge Server acts as a MRAS.)
5. The Lync client initiates an incoming call.
6. The Call Bridge sends "service" messages to the Lync FE server to request MRAS credentials to use the Lync Edge MRAS service
7. The Lync FE server returns the credentials for the Call Bridge to use, as well as the UDP and TCP ports, and the MRAS URI once again
8. The Call Bridge resolves this MRAS URI using DNS and starts sending STUN messages directly to the Lync Edge server
9. The call media then flows directly between the Call Bridge and Lync Edge's TURN server on UDP port 3478 and returns from the Lync Edge server to the Call Bridge on a port in the ephemeral range above.

Therefore the following ports need to be opened in the firewall for the media between Call Bridge and the Lync Edge server: UDP 3478 outgoing and 32768-65535 incoming.

Figure 16: Call Bridge to Lync Edge server call flow

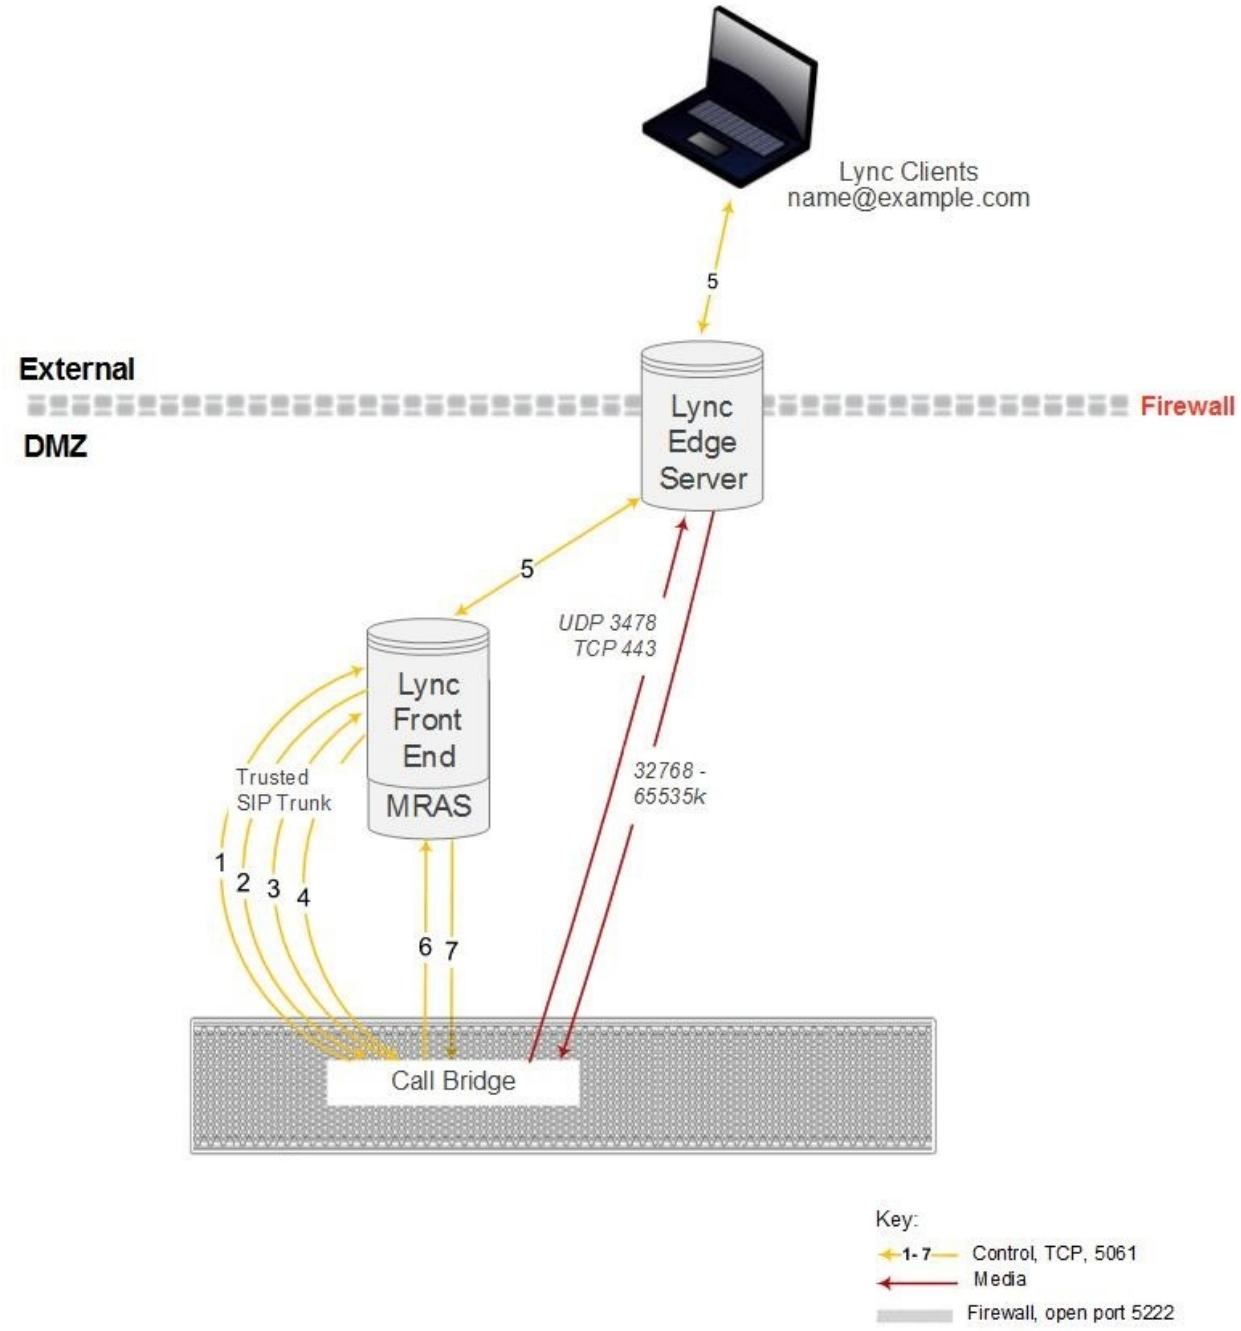

8.5.2 Configuration on Meeting Server to use Lync Edge

To use a Lync Edge server, log into the Web Admin Interface of the Meeting Server, go to **Configuration > General** and configure the Lync Edge Settings. (When a Lync Edge server is configured, it takes the TURN / ICE role for Lync calls, and so at some level is an alternative to the TURN server settings above).

You also need to create a Lync user client account to set up the Meeting Server- Lync Server Edge configuration.

Follow these steps to set up the Meeting Server to use the Lync Edge server:

1. Ensure that you have the appropriate DNS records in place; see Appendix A for a list of DNS records needed for the single combined server type deployment.

2. Create a new user in your LDAP directory, just as you would any other user in your directory, for example, firstname="edge", second name = "user".

3. Log into the user manager on your Lync FE Server and create a Lync Client user from the user you created in the previous step. Do this in the same way as you would any other user to enable them to use Lync. Using the example name above creates a Lync client user called **edge.user@lync.example.com**

4. Sign in to the Web Admin Interface of the Meeting Server, and go to **Configuration > General**. Configure the Lync Edge Settings by entering the Lync FE Server Address (or a host name that resolves to this). For Username enter the Lync client user name created in the previous step.

5. Complete the Number of Registrations field, if necessary.

   This field overcomes a feature of the Lync Edge server that limits the number of simultaneous calls that it will run for one registered device. By entering a number greater than 1, the Call Bridge will make that number of registrations, thereby increasing the number of simultaneous calls that the Meeting Server can make out through the Lync Edge Server.

   Entering a number greater than 1 adds a number to the end of your Lync Edge username and registers with the resulting username. For example, if you configured Username as edge.user@lync.example.com and set Number of Registrations to 3, you will need to create the following users in your Lync environment so that they can be used with the Edge server:
   edge.user1@lync.example.com
   edge.user2@lync.example.com
   edge.user3@lync.example.com

   We recognize that this requires some administrative overhead; however it is due to a limitation of the Lync Edge server as explained above.

   Leave the Number of Registrations blank to only make a single registration as **edge.user@lync.example.com**.

   Note: There is no need to enter the password for the Lync users because the Lync FE server trusts the Call Bridge.

Points to note about configuring the Lync Edge:

- The Meeting Server supports Lync content (presentations contributed over RDP) from external Lync clients whose media arrives via the Lync Edge server. In addition, space (URIs) now report back as busy or available based on how many participants are currently in the space so that Lync clients that have spaces in their favorites can see the space status.
- If you are using a Lync AVMCU, you need to configure the Lync edge settings in order to register with the Lync FE server.
- Cisco Meeting Apps continue to use the Meeting Server TURN server even if a Lync Edge server is configured.
- If you have a Lync Edge server configured, all Lync calls will use that server for ICE candidate gathering and external media connectivity. If you do not have a Lync Edge server configured, but have configured a Cisco Expressway in your deployment, then the Lync calls will be handled by the configured TURN server in the Expressway.
- In a typical Lync Edge deployment, the internal interface of the Lync Edge server will not have a default gateway defined; only the external interface has a default gateway defined. If the Call Bridge interface is not on the same local subnet as the internal interface of the Lync Edge server, then you must define a static and persistent network route to the Lync Edge server so it can route packets to the Meeting Server correctly, using the internal interface. To add a static and persistent network route to the Lync Edge Server, open CMD and issue the command below , replacing the example data with your own IP information.

  Example Command:

  ```
  route add –p 10.255.200.0 mask 255.255.255.0 10.255.106.1
  ```

  In this example a network route is added that allows the entire subnet of 10.255.200.0 to route through the gateway of 10.255.106.1; 10.255.106.1 is the gateway of the subnet for the internal interface on the Lync Edge server.

  Failure to add this route will result in all STUN packets sent by the Meeting Server to the Lync Edge server to go unanswered, which can result in call failures.

8.6 Direct Lync federation

The Meeting Server supports direct federation with Microsoft Lync, by putting the Call Bridge on a public IP address with no involvement from NAT. This allows calls to be made from the Meeting Server direct to any Lync domain and vice versa.

To allow inbound calls you must:

1. Create the DNS SRV record _sipfederationtls._tcp.domain.com that points to the FQDN of the Meeting Server. This step is required as Call Bridge will need to have a public IP, and NAT is not supported in this scenario.

2. Add a DNS A record that resolves the FQDN of the Meeting Server to a public IP address.
3. Upload a certificate and certificate bundle to the Meeting Server that complies with the following:
   a. The certificate must have the FQDN as the CN, or if using a certificate with a SAN list then ensure that the FQDN is also in the SAN list. Note: if the certificate contains a SAN list, then Lync will ignore the CN field and only use the SAN list.
   b. The certificate must be signed by a public CA.

      Note: you are advised to use the same Certificate Authority (CA) that is trusted by Lync FE servers. Contact your Lync adviser for details of the CA and for support on the Meeting Server-Lync integration.

   c. The certificate bundle must contain the Root CA's certificate and all intermediate certificates in the chain in sequence, so that a chain of trust can be established.

      Note: for more information on certificates refer to the Introduction in the Cisco Meeting Server Certificate Guidelines.

   d. Open the appropriate Firewall ports as stated in Appendix B for example: TCP 5061, UDP 3478, UDP 32768-65535, TCP 32768-65535

For outbound calls from the Meeting Server:

1. Create an outbound dial rule, leave the **Domain** and **SIP proxy** fields blank, and set **Trunk** type as Lync. Also set the appropriate **Local contact domain** and the **Local from domain** fields.

   If specifying individual domains in outbound dial plan rules, ensure that all domains configured on the Lync side have been added. The domains in use can be read from the Lync Server Topology Builder. Note that if additional domains are later added to Lync, then these should also be added to the outbound dial plan rules.

8.7 Calling into scheduled Lync meetings directly and via IVR

Pre-requisite on Lync deployment: This feature requires a working Lync deployment with telephone dial-in capabilities already enabled. The Lync deployment requires one or more on-prem Lync FE servers to be configured.

Note: The on-prem Lync FE servers need to be configured even if your Lync deployment does not support external Lync or Skype for Business clients.

The Meeting Server supports calling into a scheduled Lync meeting from WebRTC or SIP endpoint, using the Lync call ID to join the call; Cisco Meeting App users can only be added to a Lync meeting by a Lync client. This feature requires one or more Lync FE servers to be configured on the Meeting Server for conference lookup. You can configure one via the Web Admin interface under the Lync Edge settings from **Configuration > General**, and one or more via the API (create them as TURN servers with type "lyncEdge"). Refer to Configuration on Meeting Server to use Lync Edge for instructions on how to do this. If there are multiple FE servers in a Pool, use the Pool FQDN as the Server Address.

Note: For Lync meeting resolution, the Meeting Server uses the Lync meeting ID and DNS lookup of _sipinternaltls._tcp.lync-domain, rather than outbound rules. Set DNS SRV record _ sipinternaltls._tcp.lync-domain on your DNS server or if you do not want to use a DNS SRV record then setup a record on the Meeting Server with the command **`dns app add rr <DNS RR>`**. For more information on using the dns app command see the MMP Command Line Reference; for a list of DNS records needed for the single combined type deployment see Appendix A.

Configure the Lync FE servers, then follow the task sequence in Table 5 below:

Table 5: Task sequence to configure Lync FE servers

Sequence	Task	On the Web Admin Interface	Via the API
1	Configure the Call Bridge IVR(s) to allow entry of Lync conference IDs	If you have set up an IVR via the Web Admin Interface: Go to **Configuration > General** in the IVR section, set **Joining scheduled Lync conferences by ID** to **allowed**	If you have set up IVRs through the API: Set **`resolveLyncConferenceIds`** to **`true`** for the configured IVR
2	Allow direct dialing to Lync conference IDs from standard SIP systems. Note: you may choose to extend an existing configured domain to allow Lync conference access, or to create a new one for this purpose.	Go to **Configuration > Incoming calls**, and for one or more configured call matching domains, set **Targets Lync** to **yes**	Set **`resolveToLyncConferences`** t to **`true`** on the incoming dial plan rule
3	Allow Lync conference ID entry via the Web Bridge call join interface	If you have set up the Web Bridge via the Web Admin Inteface: Go to **Configuration > General** in the Web bridge settings section ensure that **Joining scheduled Lync conferences by ID** is set to **allowed**	If you have set up Web Bridges through the API: Set **resolveLyncConferenceIds** to **`true`** on the Web Bridge

If a call is being matched against Lync conference IDs, the Call Bridge first checks that the call ID does not apply to a space, if it does not then the Call Bridge identifies a Lync FE server that it has been configured with, that has advertised itself as having the capability to resolve IDs. The Call Bridge queries the Lync FE server to determine whether the call ID in question corresponds to a Lync conference – if it does, the look up is deemed to have been successful and the call is joined to the Lync call. If the call ID is not recognized as corresponding to a Lync conference then no further Lync FE servers will be queried.

Note: You may get unexpected results if you add the settings of multiple Lync FE servers that are in different Lync deployments. For instance, if multiple Lync conferences in different Lync deployments use the same call ID, then more than one Lync FE server may respond positively to the lookup, in which case the "first" successful Lync resolution is used.

Note: Each participant connecting through a Meeting Server to a Lync meeting is required to have a unique "from:" SIP address to avoid participant conflicts in the Lync AVMCU. Telephone participants connecting through a PSTN gateway are at a high risk of encountering participant conflicts due to the generic outgoing callerID information. It is recommended that all telephone participants connect to Lync meetings through the Lync PSTN Conferencing/Mediation Server rather than through the Meeting Server Dual Home gateway.

The text in the invitations sent for scheduled Lync meetings can be customized to include the necessary details to allow users to join via the Meeting Server. These details should be placed in the custom footer section. For example **'For SIP/H.323 endpoints, join by calling join@example.com and entering the conference ID above. For WebRTC go to join.example.com and enter the conference ID above.'** The URIs in this must match those configured above. Please see the Microsoft documentation https://technet.microsoft.com/en-us/library/gg398638.aspx for more details.

8.8 Choosing Call Bridge mode to connect participants to Lync conferences

Version 2.3 allows you to choose the behavior of the Call Bridge when connecting participants to Lync conferences, using the Meeting Server API. A request parameter **`lyncConferenceMode`** has been added when POSTing to **`/callProfiles`** or PUTing to **`/callProfile/<call profile id>`**.

Set to **`dualHomeCallBridge`** if you want the calls on the same Call Bridge to be combined into one conference. This will result in a single conference on the Call Bridge, the Call Bridge will call out to the AVMCU meeting.

Set to **`gateway`** if you do not want the calls to be combined into one conference. Each SIP participant will be in their own conference with an associated call out to the AVMCU meeting.

Note: Set **`lyncConferenceMode`** to **`gateway`** to disable dual home conferencing.

9 Office 365 Dual Homed Experience with OBTP Scheduling

9.1 Overview

"Office 365 Dual Homed Experience with OBTP (One Button To Push) Scheduling" allows participants to join Office 365 meetings using Cisco endpoints that support OBTP.

The host schedules a meeting using Microsoft Outlook with Skype for Business plugin, and adds participants and conference rooms (including OBTP-enabled endpoints) and a location to meet in.

To join the meeting, participants using a OBTP-enabled endpoint simply push the OBTP button on the endpoint or touchscreen. Skype for Business clients click a link to join the meeting as normal.

Note: If using Office 365, only invited OBTP-enabled endpoints or Skype for Business clients with Office 365 can join the Lync meeting; Cisco endpoints cannot join the meeting manually, via the Meeting Server IVR. This is a key difference to an on-premise Lync deployment,which allows any Cisco endpoint to join manually via the Meeting Server IVR.

Note: "Office 365 Dual Homed Experience with OBTP (One Button To Push) Scheduling" is supported from Version 2.2, and requires Cisco TMS 15.5, and Cisco TMS XE 5.5 or later.

9.2 Configuration

Note: This feature requires the Call Bridge to connect to the public internet in order to contact Office 365. You will need to open TCP port 443 on your firewall for outgoing traffic.

To set up this method of joining Office 365 meetings, sign into the Web Admin interface of the Meeting Server, navigate to **Configuration>Incoming calls** and configure a **Call matching** rule for incoming calls with the **Targets Lync Simplejoin** field set to `true` . This tells the Meeting Server how to resolve the Lync Simple Meet URL sent in the Office 365 invite.

To have the ability to call participants as well as meetings, use an existing outbound dial plan rule to route the outbound calls, or create a new outbound dial plan rule.

9.3 In-conference experience

"Office 365 Dual Homed Experience with OBTP Scheduling" provides the "dual homed experience" with 2-way audio, video and content sharing. Office 365 clients have the familiar in-conference experience determined by the Lync AVMCU, and participants using OBTP enabled endpoints have a video conferencing experience determined by the Meeting Server. All see the combined participants lists.

Note: Controls on clients do not work conference wide, and can give rise to some strange behavior. For example, if a Skype for Business client mutes an endpoint connected to the Meeting Server then the endpoint will mute, but no notification is sent to the endpoint to say it has been muted; the endpoint cannot unmute itself. If a Skype for Business client mutes all endpoints connected to the Meeting Server and then unmutes them, all the endpoints will remain muted.

Note: ActiveControl functionality such as muting and dropping participants only affect participants on the local Call Bridge and not on the Lync AVMCU.

10 Web Admin interface settings for XMPP

Note: Cisco plans to remove the XMPP server component from the Cisco Meeting Server software in a future version. Customers are encouraged to start planing their migration to using the Cisco Meeting WebRTC App rather than using the Cisco Meeting App thick clients (Windows/Mac desktop and iOS).

This section explains how to configure the settings through which the Call Bridge communicates with XMPP server.

Note: If you are not using the Cisco Meeting Apps including the WebRTC app, skip this chapter.

10.1 XMPP server connections

Figure 17 and Table 6 show the ports used for the native Cisco Meeting App connections.

Figure 17: Native Cisco Meeting App port usage

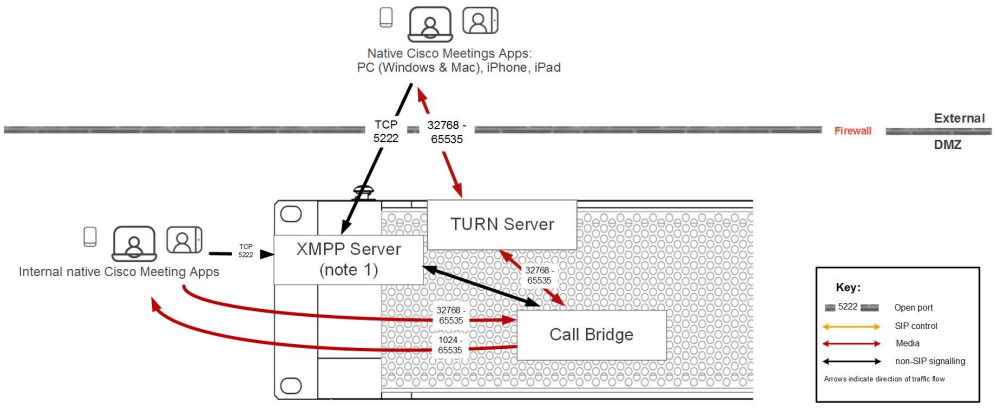

Table 6: Ports required for Native Cisco Meeting App connections

Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
XMPP server	Internal or external native Cisco Meeting Apps	5222	TCP	Incoming	For both internal and external nativeCisco Meeting Apps

Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
XMPP server	External nativeCisco Meeting Apps	3478 (note 1)	UDP (STUN)	Incoming	
Call Bridge	Internal native Cisco Meeting Apps	32768-65535	UDP (STUN RTP)	Incoming	
Call Bridge	Internal native Cisco Meeting Apps	1024-65535 (note 2)	UDP (STUN RTP)	Outgoing	
TURN server	External native Cisco Meeting Apps	32768-65535 (note 3)	Media UDP (STUN RTP)	Incoming	
TURN server	External native Cisco Meeting Apps	32768-65535 (note 3)	Media TCP (RTP)	Incoming	
Call Bridge	XMPP server				Internal to Meeting Server, does not require open ports
Call Bridge	TURN server	32768-65535 (note 3)	Media UDP (STUN RTP)	Incoming and outgoing	
Call Bridge	TURN server	32768-65535 (note 3)	Media TCP (RTP)	Incoming and outgoing	
Call Bridge	TURN server	3478 (note 1)	UDP (STUN)	Incoming	

Note:
1) If the media ports (32768-65535) are not open then port 3478 will be used.
2) Exact range depends on far end
3) Although the range is shown as 32768-65535, currently only 50000-51000 is used. A wider range is likely to be required in future releases.

10.2 XMPP settings

Follow the steps in order.

1. Ensure that you have installed a security certificate for the XMPP server.
2. Ensure that you have configured the XMPP server using the MMP.
3. If you are using a virtual host, ensure that you have uploaded the license key file.
4. Log in to the Web Admin Interface and configure the XMPP server settings as follows:

a. Go to **Configuration > General**

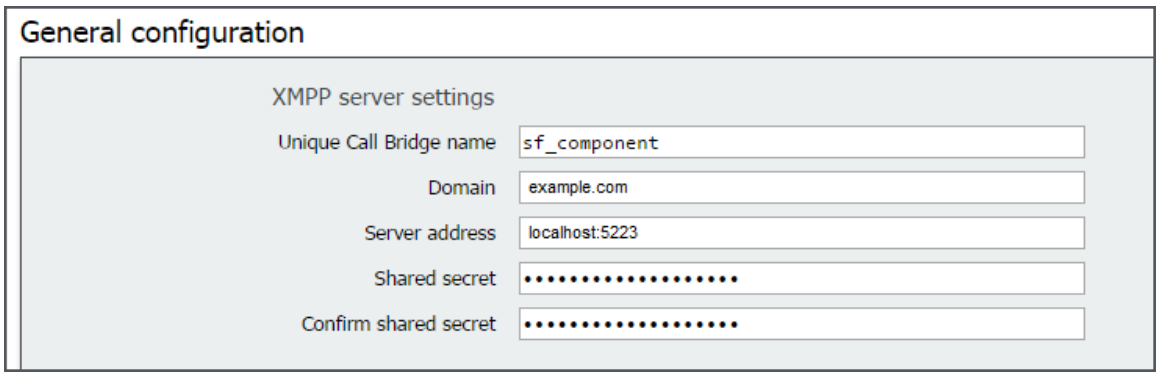

b. Complete the fields in the XMPP Server Settings section.

- Unique Call Bridge name (this is the component name set up previously, no domain part is required, as shown):

  cb_london

- Domain (this is the XMPP server domain set up previously:

  example.com

- Server Address is the IP address or hostname of the XMPP server, with an optional <port> (default is 5223):

  localhost:5223

  Note: If you are using DNS to locate the XMPP server it is not necessary to configure the server address.

- Shared secret: as generated during the XMPP server configuration (see step 9 in Section 4.5).

c. Save your configuration by selecting **Submit** at the bottom of this page.

5. Go to **Status > General** and verify the server connection.
   You should see details similar to the following in the XMPP Connection field:

System status

Uptime	8 hours, 4 minutes, 18 seconds
Build version	1.8.3
XMPP connection	connected to localhost (secure) for 8 hours, 4 minutes, 16 seconds
Authentication service	registered for 8 hours, 4 minutes, 16 seconds

6. On a PC, download the Cisco Meeting App software from the Cisco web site and install it.
7. Log in to the Cisco Meeting App using one of the newly created user accounts. Then check

that you can make calls as expected.

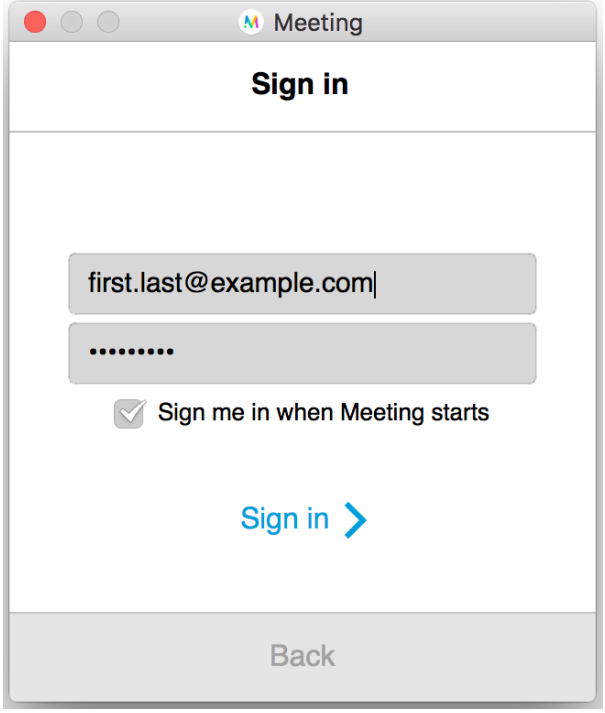

10.3 Client-based space creation and editing

Cisco Meeting App users can create spaces. These spaces have URIs and IDs by default, allowing them to be easily dialed by SIP endpoints. The SIP dial-in URI is automatically created; however, you can enter a preferred left-hand side of the SIP URI. The Meeting Server will automatically ensure that it is a unique URI. To create the full URI it combines the user entered piece with the highest priority domain in the incoming dial plan that resolves to that space. This means users can now create spaces and email the SIP URI so that others can join. This makes it straightforward to bring SIP endpoints into a space.

Refer to the Cisco Meeting App User Guide for details on creating spaces, and inviting guests and other members to meetings using the space.

Note: spaces can also be created using the Meeting Server API (see the API Reference guide), or by using the Web Admin Interface **Configuration > Spaces** page.

---

11 Web Admin interface settings for the Web Bridge

This section explains how to configure the settings through which the Call Bridge communicates with the Web Bridge server. This allows you to use WebRTC video calls and meetings.

If you are testing the WebRTC app, follow the instructions in Section 11.2 in the order provided at any time after the initial Meeting Server configuration has been completed. If you are not using the WebRTC app, skip this chapter.

11.1 Web Bridge connections

Table 7 show the ports used for WebRTC app connections. Section 11.1.1 describes the call flow between the WebRTC app and components in the Meeting Server.

Figure 18: WebRTC Client port usage

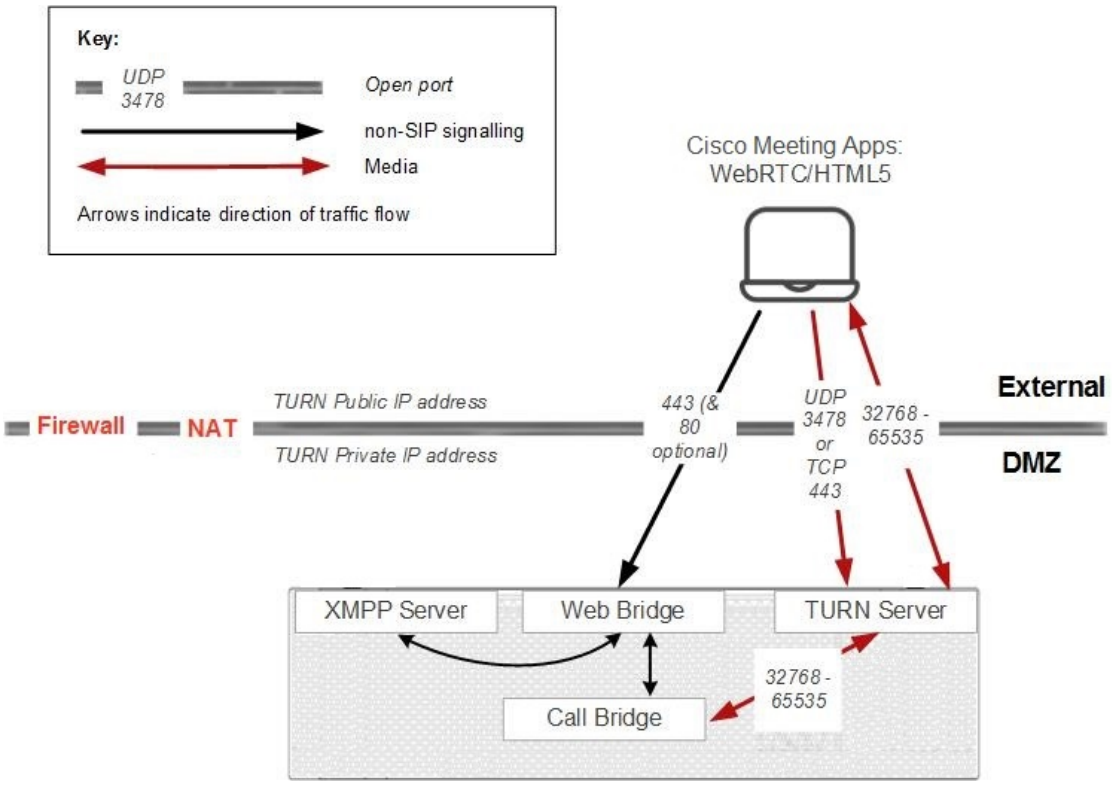

Table 7: Ports required for WebRTC app connections

Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
Web Bridge	WebRTC apps	443 (note 1)	TCP (HTTPS)	Incoming	
Web Bridge	WebRTC apps	80	TCP (HTTP)	Incoming	
TURN server	WebRTC apps	32768-65535 (note 4)	Media UDP (STUN RTP)	Incoming	
TURN server	WebRTC apps	32768-65535 (note 4)	Media TCP (RTP)	Incoming	
TURN server	WebRTC apps	3478 (note 3)	Media UDP (STUN RTP)	Incoming	
TURN server	WebRTC apps	443 (notes 1 and 3)	Media TCP (RTP)	Incoming	
Call Bridge	Web Bridge				Internal to Meeting Server, does not require open ports
XMPP server	Web Bridge				Internal to Meeting Server, does not require open ports
Call Bridge	TURN server	32768-65535 (note 4)	Media UDP (STUN RTP)	Incoming and outgoing	
Call Bridge	TURN server	32768-65535 (note 4)	Media TCP (RTP)	Incoming and outgoing	

Note:
1) To run both the Web Bridge and the TURN server on port 443 requires the two components to be run on a different interface and port combination, if this is not possible then use port 447 for the TURN server.
2) If you have setup TURN TLS (see Section 4.7) to use a different port to 443, for example 447, then that port will be used.
3) If the media ports (32768-65535) are not open then port 3478 will be used and if that is not open, then the port setup for TURN TLS will be used, if setup.
4) Although the port range between the TURN server and the External clients is shown as 32768-65535, currently only 50000-51000 is used. The required range is likely to be larger in future releases.

11.1.1 Web Bridge call flow

1. PC web browser opens HTTPS connection to Web Bridge
2. User is prompted to **Join Call** (see step 3) or **Sign In** (see step 4)
3. If **Join Call** is selected, user is prompted to enter the Call ID and Passcode (if required)
    a. Web Bridge queries Call Bridge to validate Call ID and Passcode
    b. If successful, the User is prompted to enter a Name to be displayed in the call
    c. Upon completing these steps and clicking Join Call, the WebRTC app sends an http message to the Web Bridge on port 443, which requests temporary credentials from the Call Bridge over port 443
    d. Web Bridge then connects to the XMPP Server on port 5222, using the above temporary credentials, and the Call Bridge validates the credentials
    e. Call Bridge requests allocations from the TURN Server to use for this call on UDP 3478
    f. WebRTC app requests allocations from the TURN Server to use for this call on UDP 3478 (or TCP 443)
    g. If the UDP STUN packets sent by the WebRTC app to the TURN server are successful, the WebRTC app will send media from the TURN server, with a Media Port range of 32768-65535
    h. If the UDP STUN messages are un-successful, the WebRTC app will fall back and send messages to the TURN Server on TCP Port 443
    i. If the TCP connection is successful, Media will also be sent to the TURN Server on TCP Port 443
    j. The TURN Server will then relay this WebRTC Media to Call Bridge, converting to UDP if received as TCP from WebRTC app
4. If **Sign In** is selected, user is prompted to enter Username and Password
    a. Web Bridge will do DNS Lookup for the SRV record of _xmpp-client._tcp for the domain entered in the Username field
    b. Web Bridge connects to the XMPP Server returned in the DNS lookup and sends the Credentials as supplied for verification
    c. If Login is successful, the User is logged into the WebRTC app and is shown a Client view similar to the PC XMPP app
    d. Upon attempting a New Call or joining a Meeting, the app will connect as follows
    e. Web Bridge signals Call Bridge for the call request over XMPP
    f. Call Bridge opens connections to the TURN Server to request allocations for ports to use for this call on UDP 3478

g. Once TURN allocations have succeeded, Call Bridge answers the call and sends the address and ports to use back to Web Bridge to be relayed to the WebRTC app

h. WebRTC app requests allocations from the TURN Server to use for this call on UDP 3478 (or TCP 443)

i. If the UDP messages are successful, the WebRTC app will send messages to the media port range of 32768-65535, using the specific ports relayed to it from Call Bridge

j. If the UDP messages are un-successful, the WebRTC app will fall back and send messages to the TURN Server on TCP Port 443

k. If the TCP connection is successful, media will also be sent to the TURN Server on TCP Port 443

l. The TURN Server will then relay this WebRTC media to Call Bridge, converting to UDP if received as TCP from WebRTC app

11.2 Web Bridge settings

Follow these steps:

1. Ensure that you have installed the Web Bridge certificate.

   Note: If you are intending to use branding you need to apply the license key to the server.

2. Ensure that you have configured the Web Bridge.
3. Sign in to the Web Admin Interface and configure the Meeting Server as follows:
   a. Go to **Configuration > General**.
   b. Set the following where:
      - **Guest account client URI** = The URI that the Call Bridge uses to talk to the Web Bridge, this could be the same URI to reach the guest account. This does not need to be configured if set in the API. If set here, you need to include https:// for example **`https://join.example.com`**.

        Note: You will need to set a DNS A record so that **`https://join.example.com`** points at the Web Bridge, see Appendix A.

- **Guest account JID domain** = guest account JID, e.g. example.com

- **Guest access via ID and passcode**: if set to **secure** then guests will be able to join a space by entering the Call ID and passcode on the Web Bridge. If set to **not allowed** then guests will not be able to join via entering the Call ID and passcode. If set to **legacy**, and using a 2.3 release or later, then the behavior is the same as **secure**.
- **Guest access via hyperlinks**: a guest can join a space by clicking a hyperlink included in a meeting invite. If this is set to **allowed**, a link that includes a secret and an id will immediately resolve to a conference that might be passcode protected. The complexity of the secret prevents brute force attacks. If set to **not allowed**, a guest user will need to enter the conference details before the conference lookup is performed.
- **User sign in**: if set to **allowed**, registered users and guests can join spaces using the WebRTC app. If set to **not allowed**, registered users will not be able to sign in using the WebRTC app.
- **Joining scheduled Lync conferences by ID**: if set to **allowed**, WebRTC apps can join scheduled Lync conferences by entering the Lync meeting id from the WebRTC landing page.
- **External access > Web Bridge URI**: enter the URL used to access the Web Bridge for your Meeting Server. For example: **https://join.example.com**. This is used to generate meeting invites and the cross-launch URL for the Cisco Meeting App.

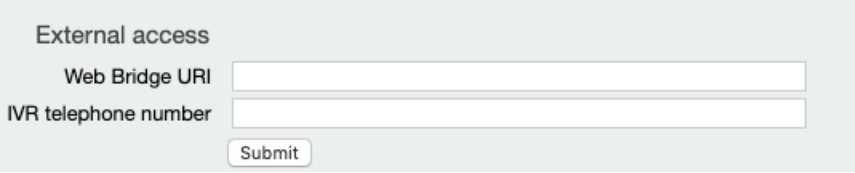

4. Open a web browser and go to the URI for the Guest account client, for example:**https://join.example.com** to test the configuration.

   Guest users selecting the web link will see a landing page, click on **Join call** then enter the Call ID and passcode, if one is set up for the call.

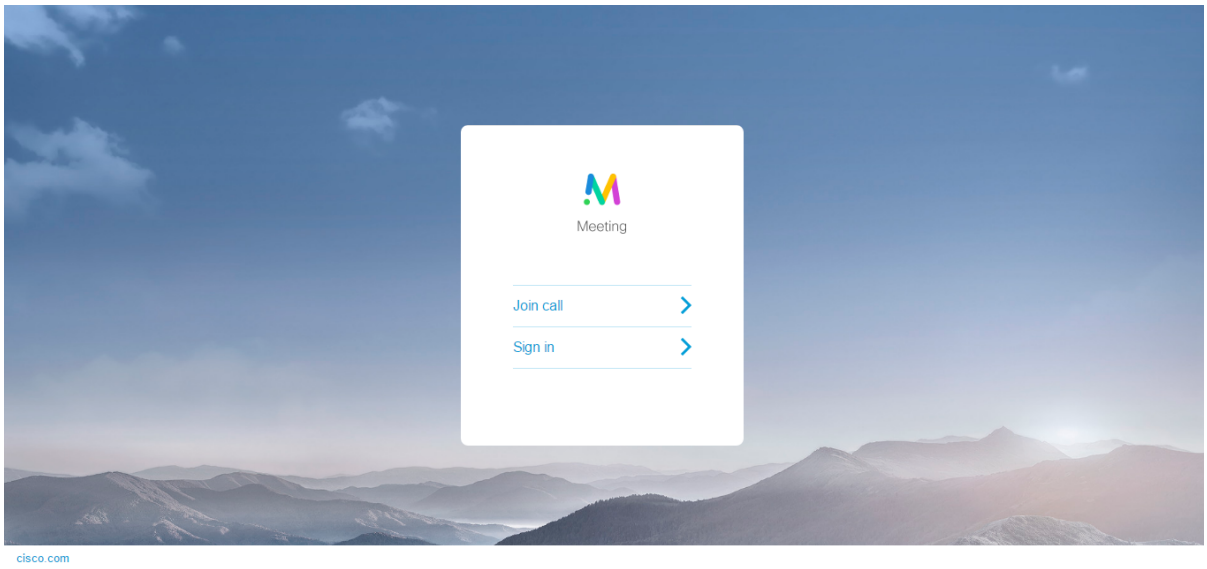

If **User sign in** via the Web Bridge is set to **allowed** in the Web Admin interface (**General configuration > Web bridge settings**), Cisco Meeting App users who do not have access to a native Cisco Meeting App but have an account, can click on **Sign in** and then enter their username and password (see note below). After signing in they see their spaces, participate in meetings and can invite participants to meetings; all from the WebRTC app.

Note: To ensure users with a Cisco Meeting App account see their XMPP/URI domain displayed when they sign in to their account using the WebRTC app, give the Call matching rule for the XMPP domain the highest priority on the **Incoming calls** page of the Web Admin interface. Otherwise, the user may see the FQDN of the Web Bridge that they landed on when they signed in.

11.3 Web browsers supporting the WebRTC app

Prior to version 2.4 of the Meeting Server software, the WebRTC version of the Cisco Meeting App was only supported on Chrome. From version 2.4 the WebRTC app is supported on the following browsers:

- Google Chrome for Windows, macOS and Android. Use Chrome version 66 or later. We strongly recommend using the most recent version of Chrome.
- Mozilla Firefox for Windows and macOS. Use Firefox is 59.0.2 or later. We strongly recommend using the most recent version of Firefox.
- Apple Safari for macOS. Use Safari 11.1 or later. We strongly recommend using the most recent version of Safari.

Note: Content cannot be sent from Safari on macOS, iOS or from Chrome on Android, this is a browser limitation.

For more information on browser support and supported devices, see the Cisco Meeting App WebRTC Important Information.

Version 2.5 supports additional browsers, these are:

- Safari on iOS for iPads, running the latest version of iOS (recommended). iOS 11.0 is the minimum supported release.
- Safari on iOS for iPhones, running the latest version of iOS (recommended). iOS 11.0 is the minimum supported release. (This is beta quality in version 2.5.x).

  Note: We have tested the WebRTC app using the Safari browser on iPad Air 2 and iPad Pro 12.9 inch (2nd generation) with iOS 11.4.1, iPad (6th generation) with iOS 12.0.1, iPhone 6 on iOS12, iPhone 7 on iOS 12 and 12.1, iPhone 8 Plus on iOS12 and 12.1, and iPhone X on iOS 11.4.1.

- the latest version of Microsoft Edge (Microsoft Edge 42/Microsoft EdgeHTML 17) on Microsoft Windows 10 (this is beta quality in version 2.5.x).

  Note: There are limitations using the WebRTC app with Microsoft Edge and Mozilla FireFox browsers:

  - Using the WebRTC app with Microsoft Edge will not work if using the TURN server in Cisco Expressway or using the Meeting Server TURN with TCP.
  - Using the WebRTC app with Firefox will not work if using the TURN server in Cisco Expressway with TCP, but will work with the Meeting Server TURN with TCP.

  See Cisco Meeting App WebRTC Important Informationfor further details on these and other limitations.

12 Web Admin interface settings for the TURN server

Note: Cisco plans to remove the TURN server component from the Cisco Meeting Server software in a future version. Customers are encouraged to plan their migration over to using Cisco Expressway for TURN, see the Cisco Expressway Traffic Classification Deployment guide for deployment information.

This section explains how to configure the settings through which the Call Bridge communicates with the TURN server. The TURN server allows you to use the built-in firewall traversal technology when traversing a firewall or NAT.

Follow the instructions in Section 12.2 in the order provided at any time after the initial Meeting Server configuration has been completed.

12.1 TURN server connections

The TURN server listens on both ports 443 and 3478 for both UDP and TCP connections. From version 2.0.4, the TURN server will never listen on port 443 on the loopback interface and by default, the Call Bridge tries to contact the TURN server using TCP port 3478 rather than TCP port 443 as in previous releases.

Figure 19 and Table 8 show the ports used by the TURN server.

Figure 19: Ports used by TURN Server

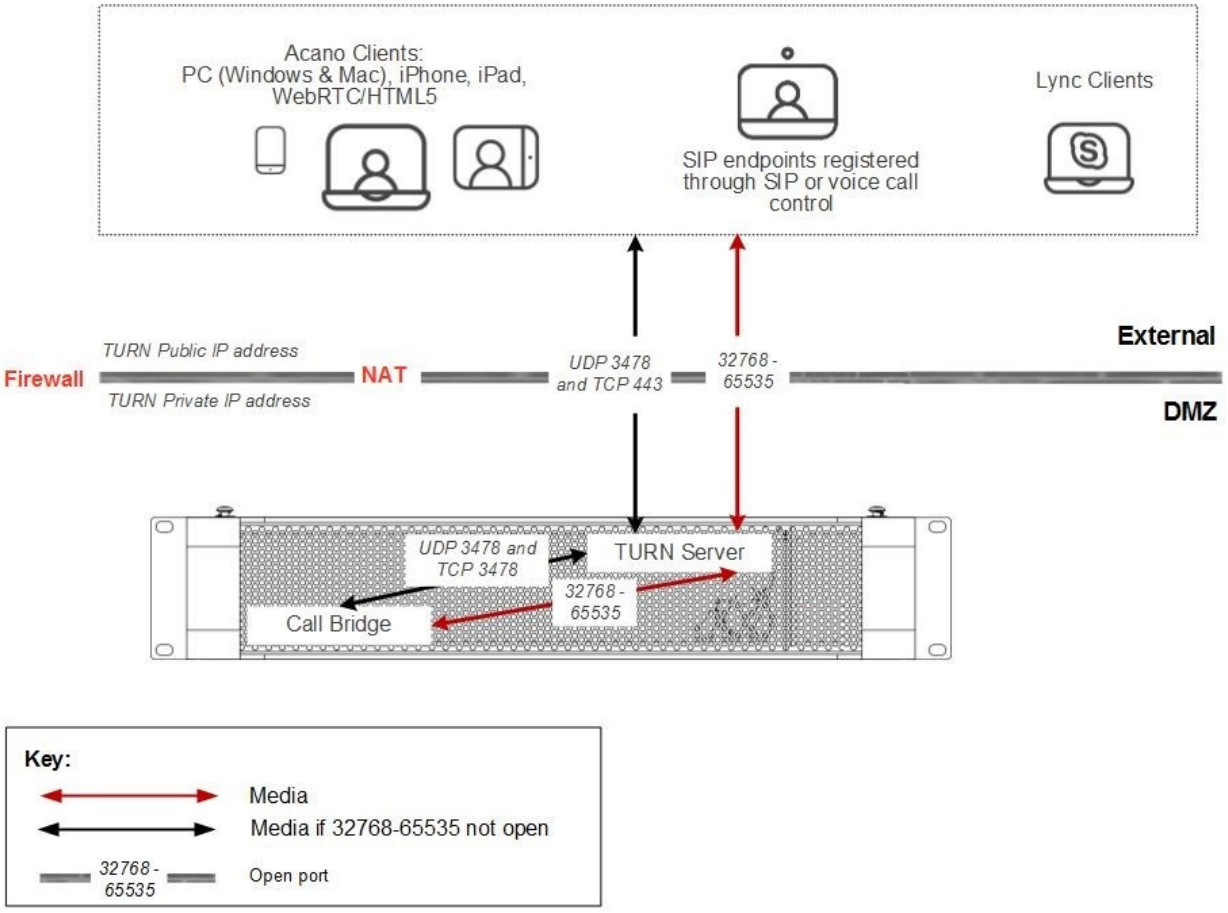

Table 8: Ports required for TURN server connections

Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
TURN server	Call Bridge and remote devices (note 1).	32768-65535 (note 2)	Media TCP (RTP)	Incoming and outgoing	
TURN server	Call Bridge and remote devices.	32768-65535 (notes 2 and 3)	Media UDP (STUN RTP)	Incoming and outgoing	
TURN server	Call Bridge and remote devices.	3478 (note 3)	UDP (STUN)	Incoming	
TURN server	Call Bridge and remote devices.	3478 (note 3)	TCP (STUN)	Incoming	Typically won't be used by remote devices and doesn't need opening in external fire-wall.
TURN server	Call Bridge and remote devices.	443 (see notes 3,4,5)	UDP (STUN)	Incoming	Typically won't be used by remote devices and doesn't need opening in external fire-wall.
TURN server	Call Bridge and remote devices.	443 (see notes 3,4,5)	TCP (STUN)		

Note:
1) Remote devices includeCisco Meeting Apps, WebRTC clients and SIP endpoints or voice control.
2) Although the range is shown as 32768-65535, currently only 50000-51000 is used. A wider range is likely to be required in future releases.
3) If the media ports (32768-65535) are not open then TCP/UDP port 3478/443 used to connect to the TURN server will be used to relay media
4) UDP/TCP port /443 can be changed. Using the MMP command **`turn tls <port>`** will change the UDP/TCP port that the TURN server listens.
5) The TURN server will not listen on port 443 on the loopback interface. This is to avoid port clashes with other services that may be running on port 443 on the loopback interface.

12.2 TURN server settings

Follow the steps in order.

1. Ensure that you have configured the TURN server.
2. Log into the Web Admin Interface and configure the Meeting Server as follows:
   a. Go to **Configuration > General**.
   b. Set the following:

   - TURN Server Address (Server) = internal server IP address that the Call Bridge will use to access the TURN server to avoid firewall traversal for internal call control
   - TURN Server Address (Clients) = public IP address assigned to the TURN server that external clients will use to access the TURN server. This will be the IP address entered in Section 4.7 when you configured the TURN server.

   Note:
   For example, if the interface of the TURN Server is on IP address XX.XX.XX.XX and NAT'ed to an external IP address YY.YY.YY.YY then enter XX.XX.XX.XX as the TURN Server Address (Server) and YY.YY.YY.YY as TURN Server Address (Client). If the interface is on the external IP then no need to enter a client address.

   You can enter a DNS name instead of an IP address in both fields, if the DNS name resolves to the appropriate IP address.

   If you are using a public IP address, leave TURN Server Address (Clients) address blank and set TURN Server Address (Server) to the public IP address or DNS name used

   - Username and Password = your information

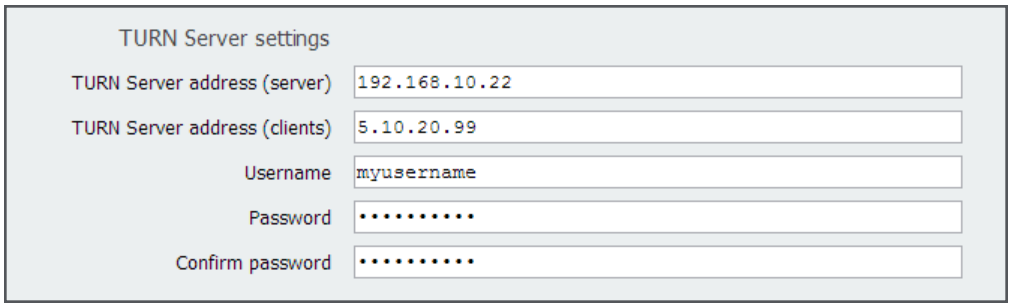

13 SIP and Lync call traversal of local firewalls (BETA)

Note: SIP and Lync call traversal of local firewalls is a beta feature and should not be used in production environments. This edge feature will be removed from the Cisco Meeting Server software in a future version.

The Meeting Server supports traversal of local firewalls for SIP endpoints and Lync calls. The Call Bridge uses the TURN server component within the Meeting Server to traverse the local firewall and sends the SIP signal via a new SIP Edge component. A third party SIP firewall traversal device is not required.

You need to set up an Outbound dial plan rule, even if you don't plan on making outgoing calls via the SIP edge. This is because any new transaction within the call or sending content to a Lync device uses the outbound rule. Once the SIP Edge has been configured and enabled, incoming calls from SIP devices are automatically traversed across the firewall.

Note: this feature assumes that the remote SIP device can see the TURN server. It does not require the remote SIP device to be ICE aware. It also requires the remote SIP device to be able to contact the SIP Edge, the SIP Edge can either have a public IP address or sit behind a NAT with appropriate forwarding of traffic.

Figure 20 shows a schematic for SIP call traversal using the TURN server and SIP Edge component in the DMZ network which is accessible to remote SIP devices via the public IP address 203.0.113.1. The Call Bridge is deployed in the private network and accesses the TURN server and SIP Edge via the internal IP address 198.51.100.1. The table below Figure 20 lists the ports required to be open for SIP/Lync call traversal.

Figure 20: Firewall traversal for remote SIP devices using the SIP Edge component

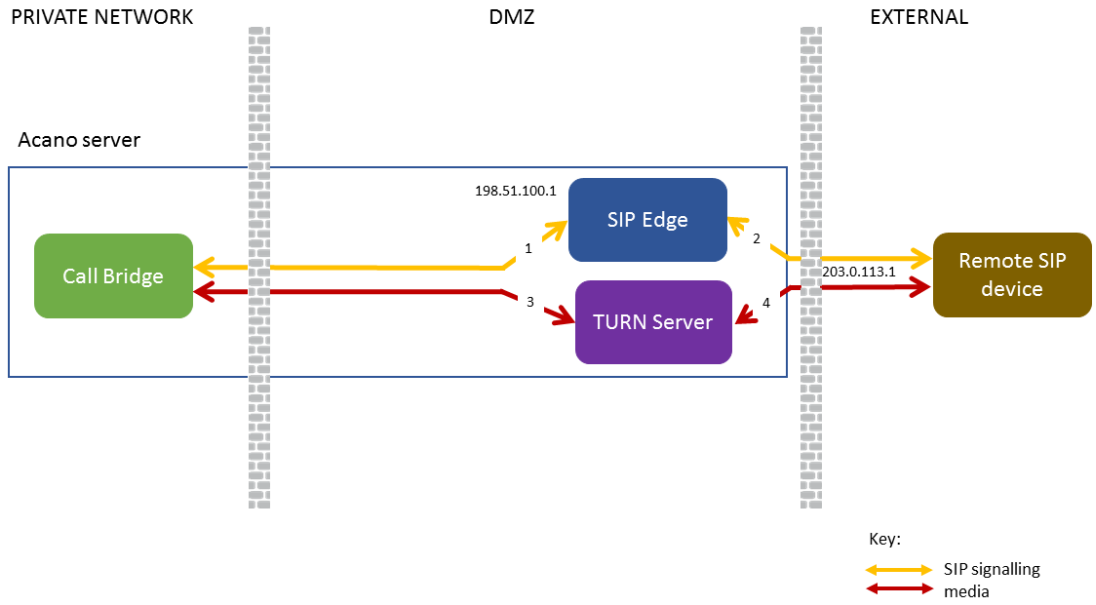

Note: SIP endpoints or Lync clients external to the network may still need a third party device to route SIP signaling across their firewall.

Table 9 below lists the ports required to be open for SIP/Lync call traversal.

Table 9: Ports to open for SIP/Lync call traversal

Link	Component	Connecting to	Destination port to open port	Traffic type	Traffic direction with respect to component	Configurable
1	Edge server (private interface)	Call Bridge	Any unused port, e.g. 3061	SIP, TLS	Incoming from Call Bridge	Yes
2	Edge server (public interface)	Remote SIP devices	5061 (Note 1)	SIP TLS (Note 2)	Incoming/Outgoing	Yes
3	Call Bridge	TURN server	32768-65535	Media UDP (STUN, RTP)	Incoming/Outgoing	No

Link	Component	Connecting to	Destination port to open port	Traffic type	Traffic direction with respect to component	Configurable
3	Call Bridge	TURN server	32768-65535	Media TCP (RDP)	Incoming/Outgoing	No
3	Call Bridge	TURN server	3478	Media UDP	Incoming	No. Internal to Meeting Server.
3	Call Bridge	TURN server	3478	Media TCP	Incoming	Yes, see note 4. Internal to Meeting Server.
4	TURN server	Remote SIP devices	32768-65535	Media UDP (STUN, RTP)	Incoming/Outgoing	No
4	TURN server	Remote SIP devices	32768-65535	Media TCP (RDP)	Incoming/Outgoing	No
3,4	TURN server	Remote SIP devices	3478	UDP	Incoming/Outgoing	Yes
3,4	TURN server	Remote SIP devices	443 (Note 3)	TCP	Incoming/Outgoing	Yes

Note 1: Port 5061 is normally used for SIP TLS

Note 2: Only SIP TLS is supported, there is no support for UDP or TCP on port 5060

Note 3: Using the MMP command **`turn tls <port>`** will change the TCP port that the TURN server listens on for both Call Bridge and App connections..

13.1 Configuring SIP/Lync call traversal

1. Set up a SIP Edge on the Edge server.

    a. Configure the internal interface and port for communication to the Call Bridge. If the Call Bridge is using port 5061, then use a different port. In a single combined deployment, the SIP Edge needs to listen to the local interface.
    `sipedge private <local interface>:port`

    for example: **`sipedge private lo:3061`**

    b. Configure the external interface and port on the SIP Edge
    `sipedge public b:5061`

    for example: **`sipedge public b:5061`**

Note: The SIP Edge always uses TLS for communication. Typically SIP TLS uses port 5061.

c. If the SIP Edge is behind a NAT, then configure the public address of the NAT.
sipedge public-ip <address>

for example: **sipedge public-ip 203.0.113.0**

Note: DNS records used for external connections must match the public address

d. Set up a certificate, key file and trust bundle on the SIP Edge. These files are used to communicate with the internal Call Bridge and the external SIP server. If you have previously assigned a public CA signed certificate to the Call Bridge, then you can use the same certificate on the SIP Edge. If the SIP Edge has direct federation with a Lync Edge server, then the certificate file must be signed by a public CA trusted by the Lync deployment (as was previously required for the Call Bridge).

Combine the Call Bridge certificate and the chain of CA certs into one file and use this as the SIP Edge certificate <certificatefile>. To enable the SIP Edge to trust the Call Bridge for the TLS trunk, use the Call Bridge certificate as the <trust-bundle>.
sipedge certs <keyfile> <certificatefile> <trust-bundle>

for example:

sipedge certs sipedge.key sipedge.crt callbridge.crt

Note: SIP Edge certificates need to be signed by a public CA and trusted by the third party SIP server. Apply the certificate to the SIP Edge and the Call Bridge. For more information on certificates refer to the Certificate Guidelines.

e. Enable the SIP Edge
sipedge enable

2. Set up a trunk from the Call Bridge to the SIP Edge. Note that you need to configure trust between the Call Bridge and the SIP Edge before creating the trunk.

Note: Currently, only one trunk can be made to any SIP Edge, and only one trunk can be made from the same Call Bridge.

a. Set up certificates for the connection to the SIP Edge. This uses the certificate from the SIP Edge as used above.
callbridge trust edge <certificate file>

for example: **callbridge trust edge sipedge.crt**

b. Create the trunk using the IP address and port of the private interface.
callbridge add edge <ip address>:<port>

for example: **`callbridge add edge 198.51.100.0:3061`**

3. Set DNS records to point to SIP Edge(s) for SIP and/or Lync. DNS can point to multiple SIP Edges for resilience. Use the **`_sips._tcp<domain>`** SRV record for the external TLS connection.

4. Configure the TURN server as explained in Chapter 12. Note that TURN TLS is required for Lync content data packets as they use TCP and not UDP.

5. Create the outbound dial plan rule. Use the API to PUT to the relevant outbound dial plan rule '/api/v1/outboundDialPlanRules/ <outbound dial plan rule id>', with 'callRouting=traversal'

   Points to note:

   - You cannot use the Web Admin interface to select the call routing.
   - Outgoing calls use the certificates that you setup for incoming calls, see step 1d above.
   - The SIP Edge only supports TLS. All dial plan rules targeting the SIP Edge must be set to 'sipControlEncryption=encrypted'.
   - The Call Bridge determines the next hop of the signaling by doing a DNS lookup. It then sends this information to the SIP Edge using the outbound rules.

Table 10 below outlines the call flow to establish an outgoing call from the Meeting Server to a remote SIP device via the SIP Edge server.

Table 10: Call flow from the Meeting Server to Remote SIP device via SIP Edge

See signals in Figure 20 and Figure 4	Call Flow
1	Call Bridge uses an outbound dial plan rule to route SIP signaling via the SIP Edge server
1	Call Bridge sends a DNS request to resolve the next hop to send the request to
1	Call Bridge sends the requests to the SIP Edge with both local address and the TURN address for receiving media
2	The SIP Edge server makes the outgoing call to the remote SIP device
	The remote SIP device answers the call
3,4	Media flows between the TURN server and the remote SIP device

14 Recording meetings

14.1 Overview

The Recorder component on the Meeting Server adds the capability of recording meetings and saving the recordings to a document storage such as a network file system (NFS).

The Recorder should be enabled on a different Meeting Server to the server hosting the conferences, see Figure 21. Only locate the Recorder on the same Meeting Server as the Call Bridge which is hosting the conferences (local) for the purposes of testing the deployment.

The recommended deployment for production usage of the Recorder is to run it on either a dedicated VM, or as part of a combined Recorder/Streamer VM. This VM should be sized with 1 vCPU and 0.5 GB of memory per concurrent 720p30 recording, with a minimum of 4 vCPU and a maximum of 24 vCPU. If combined with a Streamer, then the sum of both components memory and RAM needs to be allocated, subject to the same minimum and maximum vCPU value.

Where possible it is recommended that the Recorder is deployed in the same physical locality as the target file system to ensure low latency and high network bandwidth. It is expected that the NFS is located within a secure network.

The recorder uses variable bit rate, so it is not possible to accurately predict how much storage a recording will take. Our testing has shown that the size of 720p30 recordings ranges between 300MB to 800MB for 1 hour. In terms of budgeting it would be safe to assume 1GB per hour.

Note: Depending on the mechanism you use to store the recordings you may need to open external firewall ports so that the recorder, uploader and storage system can communicate. For example: NFS running version 2 or 3 of the port mapper protocol uses TCP or UDP ports 2049 and 111.

Note: Do not use the Firewall component on the Meeting Server if using either the Recorder or Uploader.

Note: At the end of recording a meeting, the recording is automatically converted to MP4. The converted file is suitable for placing within a document storage/distribution system, for example, in a network file system (NFS) they are stored in the NFS folder spaces/<space ID; tenant spaces are stored in tenants/<tenant ID>/spaces/<space ID>.

Figure 21: Permitted deployment for recording: remote mode

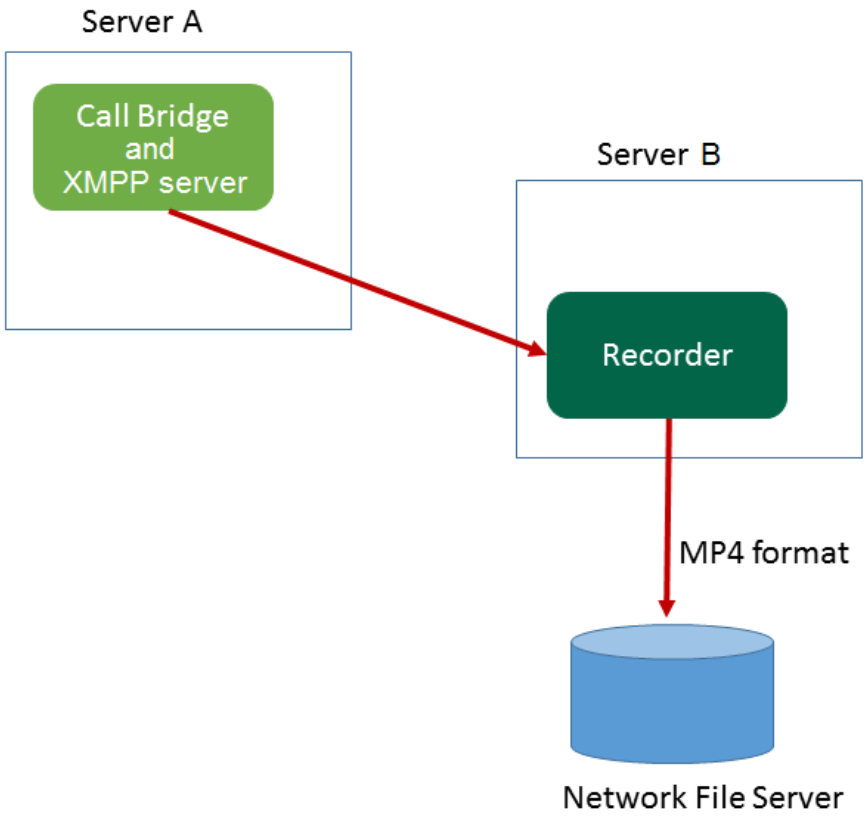

The Recorder also supports redundant configurations, see Figure 22. If you use multiple recorders then the solution load balances recordings between all recording devices and no knowledge of the physical location of recording devices is known.

Figure 22: Permitted deployments for recording: multiple recorders

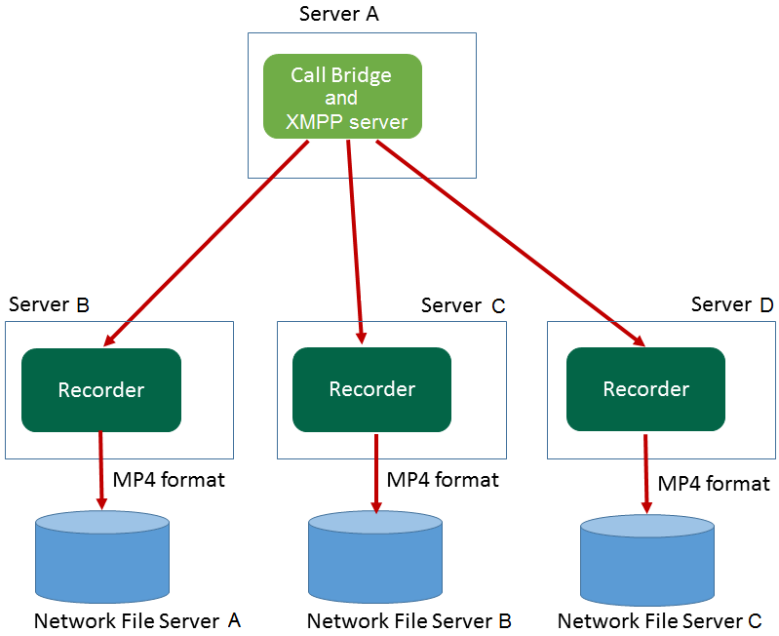

For testing purposes the Recorder can be co-located on the same server as the Call Bridge hosting the conferences. This may support between 1 to 2 simultaneous recordings.

Note: Acano X-series servers used in the single combined deployment mode should only be used for testing the Recorder, they should not be used in production networks to host the Recorder.

Figure 23: Permitted deployment for testing : local mode

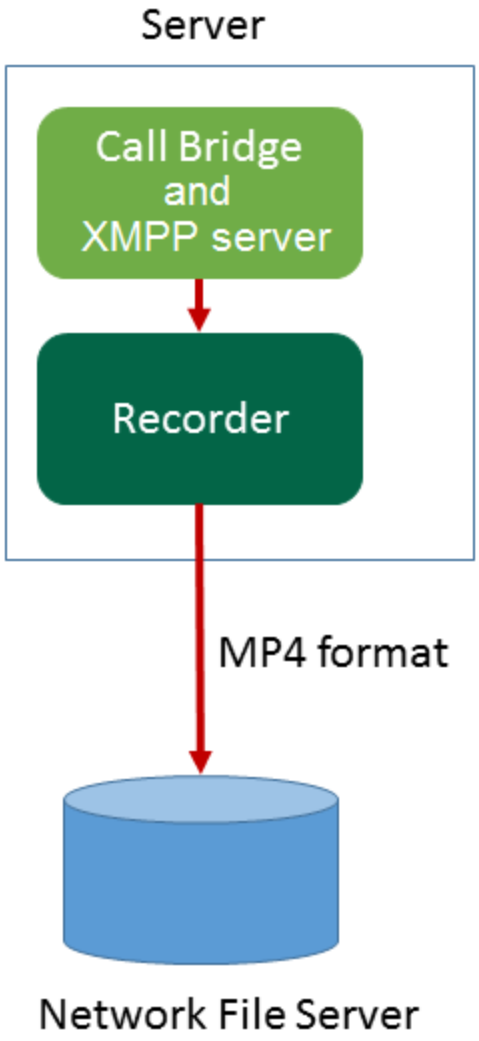

14.2 Configuring the Recorder

- Use MMP commands to enable the Recorder on a Meeting Server, specify which Call Bridges within the deployment will work with the Recorder and where to save the recordings. The additional MMP commands are given in the MMP Command Line Reference guide.
- Specify the HTTPS URL address that the Call Bridge will use to reach this recorder. Either POST the URL to the /recorders object or PUT to the /recorders/<recorder id> object.

- Use the new recordingMode parameter on the API object /callProfiles or /callProfiles/<call profile id> to select whether a meeting can be recorded or not. Options for this are:

  Automatic - recording occurs without any user intervention, if recording cannot occur the meeting still occurs.

  Manual - users can manually start and stop the recording using DTMF.

  Disabled - no users can record.

- Control which users have permission to start and stop recording by setting the recordingControlAllowed parameter on callLegProfiles.

- Use the new startRecording and stopRecording parameters for /dtmfProfiles and /dtmfProfiles/<dtmf profile id> to map the DTMF tones for starting and stopping recording.

  Note: The additional API objects are given in the Cisco Meeting Server API Reference guide.

- At the end of recording a meeting, the recording is automatically converted to MP4. The converted file is suitable for placing within a document storage/distribution system, for example, in a network file system (NFS) they are stored in the NFS folder spaces/<space ID; tenant spaces are stored in tenants/<tenant ID>/spaces/<space ID>

Note: For the first 5 seconds after recording is started, the call will not be recorded. This is intentional and is to aid lipsync between video and audio in the recording. If you record for less than 5 seconds a small file will be saved on the NFS, but it will not play back.

14.3 Example of deploying recording

Note: If you plan to save the recordings on a NFS server running Windows 2008 R2 SP1, there is a windows hotfix required to fix permission issues: https://support.microsoft.com/en-us/kb/2485529. Consult your Microsoft Windows Administrator before applying this fix.

Note: The Recorder behaves as an XMPP client, so the XMPP server needs to be enabled on the Meeting Server hosting the Call Bridge.

This example gives the steps to deploy a recorder remote to the Call Bridge. It assumes that you already have a working Call Bridge and XMPP server.

1. Create a certificate and private key for the Recorder, following the steps described in the Certificates Guidelines for an internal CA signed certificate.

2. SSH into the MMP of the Meeting Server hosting the Recorder.

3. Configure the Recorder to listen on the interface(s) of your choice with the following command:
   recorder listen <interface[:port] whitelist>

   The Recorder can listen on multiple interfaces, e.g. one on public IP and one on the internal network. (However, it cannot listen on more than one port on the same interface.)

   The following is an example where interfaces are set to interface A and B, both using port 8443.
   recorder listen a:8443 b:8443

   To use a local Recorder, the Recorder must listen on the loopback interface lo:8443, for example
   recorder listen lo:8443 b:8443

4. Upload the certificate file, key file and certificate bundle to the MMP via SFTP.
   recorder certs <keyfile> <certificatefile> [<crt-bundle>]
5. Add the Call Bridge certificate to the Recorder trust store using the command:
   recorder trust <crt-bundle>
6. Specify the hostname or IP address of the NFS, and the directory on the NFS to store the recordings
   recorder nfs <hostname/IP>:<directory>

   Note: The Recorder does not authenticate to the NFS.

7. Use the recorder command to list the details for the recorder, for example:
```
cms1> recorder
Enabled                 : true
Interface whitelist     : a:8443 b:8443
Key file                : recorder0.key
Certificate file        : recorder0.cer
CA Bundle file          : recorder.crt
Trust bundle            : callbridge.crt
NFS domain name         : examplecompany_nfs
NFS directory           : /home/examplecompany/nfs
```

8. Enable the Recorder:
   recorder enable
9. Create DNS A record for the Recorder and set it to resolve to the IP Address of the Ethernet interface you want the Recorder to listen on.
10. Use the API of the Meeting Server hosting the Call Bridge to configure the settings through which the Call Bridge will communicate with the Recorder.

    a. Specify the HTTPS URL address that the Call Bridge will use to reach this recorder. Either POST the URL to the /recorders object or PUT to the /recorders/<recorder id> object

Note: If using a local Recorder, the URL must be the loopback interface, for example `https://127.0.0.1:8443`

b. Select whether a meeting can be recorded or not and whether the recording will start without any user intervention. Use the recordingMode parameter on the API object /callProfiles or /callProfiles/<call profile id>

c. Control which users have permission to start and stop recording. Use the recordingControlAllowed parameter on /callLegProfiles

d. Use the startRecording and stopRecording parameters for /dtmfProfiles and /dtmfProfiles/<dtmf profile id> to map the DTMF tones for starting and stopping recording. For example: **7 to start and **8 to stop recording.

11. Remember to set the permissions on your NFS to rw and change the chown and chmod permissions on the directory. For example:

```
sudo chown nobody:nogroup /record
sudo chmod -R 777 /record
```

14.4 Recorder licensing

14.4.1 Recorder licensing

You will need a license for each Call Bridge, and one or more licenses for recording which is loaded on the Call Bridge server, not the Recorder server. A recording license supports 1 concurrent recording. Contact your Cisco sales representative to discuss your licensing requirements.

14.5 Setting the resolution of the Recorder

From version 2.4, you can configure the recording resolution of the Meeting Server Recorder. The resolution is configured on the Recorder component itself and is not passed to the Recorder by the Call Bridge. To configure the resolution use the MMP command: `recorder resolution <audio|720p|1080p>`.

If no resolution is configured, the default setting is 720p30. Audio recordings are stored in .mp4 file format.

Table 11 provides typical specifications for the different recorder settings, the recommendations are based on our internal testing.

Table 11: Recorder resolution specifications

Recorder setting	Percentage of physical core per recording	RAM required per recording	Typical disk usage per hour	Planned disk usage per hour (recommended)
1080p	100%	1GB	1GB to 1.6GB	2GB
720p	50%	0.5GB	300MB to 800MB	1GB
audio	20%	125MB	70MB	100MB

14.5.1 Example of setting the recording resolution

This example assumes you already have a working Recorder.

1. SSH into the MMP of the Meeting Server hosting the Recorder.
2. Disable the Recorder to change the configuration.
   ```
   recorder disable
   ```
3. Configure the Recorder to record meetings at the specified resolution using the MMP command:
   ```
   recorder resolution <audio|720p|1080p>
   ```
   For example:
   ```
   recorder resolution 1080p
   ```
4. Re-enable the Recorder so that it picks up the new configuration.
   ```
   recorder enable
   ```

14.6 Recording indicator for dual homed conferences

For dual homed conferences, recording should be done using the Microsoft recording method on the Lync/Skype endpoint. We do not recommend using Cisco Meeting Server to record dual homed conferences.

From version 2.4, a recording icon indicates to SIP participants connected to the Meeting Server that a Lync/Skype endpoint is recording the conference on the Lync/Skype side.

Meeting Server adds a recording icon to the video pane composed for non-ActiveControl endpoints. Table 12 below shows the icons that Meeting Server will display to indicate that a dual homed conference is being recorded.

Table 12: Recording indicators

Icon displayed	Description
	Meeting is being recorded via the Meeting Server.
	Meeting is being recorded by a Lync/Skype endpoint
	Meeting is being recorded via the Meeting Server and by a Lync/Skype endpoint.
	The meeting is not being recorded (no icon displayed).

Note: The Cisco Meeting App shows the recording state using its own icons, they do not distinguish between local and remote recording. Meeting Server icons are not overlaid on the Cisco Meeting App video pane.

14.7 Recording with Vbrick

Note: The Uploader component is a fully released feature in version 2.4.0.

The Uploader component simplifies the work flow for uploading Meeting Server recordings to the video content manager, Vbrick, from a configured NFS connected to a Meeting Server. No manual importing of recordings is required.

Once the Uploader component is configured and enabled, recordings are pushed from the NFS to Vbrick, and an owner is assigned to the recording. The Rev portal applies security configured by your administrator to your video content, only allowing a user to access the content that they are permitted to access. Vbrick emails the owner when the recording is available in the owner's Rev portal. Owners of a recording access the video content through their Rev portal, and can edit and distribute as necessary.

Note: If a file is added to the NFS share within a space directory, the file will be uploaded to Vbrick as though it were a valid recording. Take care to apply permissions to your NFS share so that only the Recorder can write to it.

Note: Depending on the mechanism you use to store the recordings you may need to open external firewall ports so that the recorder, uploader and storage system can communicate. For

example: NFS running version 2 or 3 of the port mapper protocol uses TCP or UDP ports 2049 and 111.

Note: Do not use the Firewall component on the Meeting Server if using either the Recorder or Uploader.

14.7.1 Prerequisites for the Meeting Server

Uploader installation. The Uploader component can be installed on the same server as the Recorder component, or on a separate server. If installed on the same server as the Recorder, then add a couple of vCPUs for it to use. If run on a different server, then use the same server specification as for the Recorder: dedicated VM with a minimum of 4 physical cores and 4GB.

CAUTION: The Uploader must run on a different Meeting Server to the Call Bridge hosting the conferences.

Read and Write access to the NFS share. The Meeting Server running the Uploader will require Read and Write permissions for the NFS. Write permission is required to allow the Uploader to re-write the name of the mp4 file when upload is completed.

Note: If the NFS is set or becomes Read Only, then the Uploader component will continuously upload the same video recording to Vbrick. This is a result of the Uploader being unable to mark the file as upload complete. To avoid this, ensure that the NFS provides read/write access.

API Access to Vbrick Rev. Configure API access for a user on Vbrick Rev.

API Access to Call Bridge. Configure API access for a user on the Meeting Server running the Call Bridge.

Trust Store Store the certificate chains from the Vbrick Rev server, and the Meeting Server running the Web Admin interface for the Call Bridge. The Uploader needs to trust both the Vbrick Rev and the Call Bridge.

Decide who has access to the video recordings. Access to uploaded video recordings can be set to: All Users, Private, and for only space owners and members.

Default state of video recordings. Decide whether the video recordings are immediately available after upload (Active), or that the owner of the video recording needs to publish it to make the recording available (Inactive).

Table 13: Port Requirements

Component	Connecting to	Destination port to open
Call Bridge	NFS (version 3)	2049
Uploader	Web Admin of Call Bridge	443 or port specified in Uploader configuration

Component	Connecting to	Destination port to open
Uploader	Vbrick Rev server	443 for video uploads and API access to Vbrick Rev server

14.7.2 Configuring the Meeting Server to work with Vbrick

These steps assume that you have already setup the NFS to store recordings.

1. Establish an SSH connection to the MMP of the Meeting Server where you want to run the Uploader. Log in.
2. For new Vbrick installations, ignore this step. If you are reconfiguring a Vbrick installation then first disable Vbrick access to the Meeting Server.
   `uploader disable`
3. Specify the NFS that the Uploader will monitor.
   `uploader nfs <hostname/IP>:<directory>`
4. Specify the Meeting Server that the Uploader will query for recording information, for example the name of the Meeting Server hosting the space associated with the recording.
   `uploader cms host <hostname>`
5. Specify the Web Admin port on the Meeting Server running the Call Bridge. If a port is not specified, it defaults to port 443.
   `uploader cms port <port>`
6. Specify the user with API access on the Meeting Server running the Call Bridge. The password is entered separately.
   `uploader cms user <username>`
7. Set the password for the user specified in step 6. Type
   `uploader cms password`
   you will be prompted for the password.
8. Create a certificate bundle (crt-bundle) holding a copy of the Root CA's certificate and all intermediate certificates in the chain for the Web Admin on the Meeting Server running the Call Bridge.
9. Add the certificate bundle created in step 8 to the Meeting Server trust store.
   `uploader cms trust <crt-bundle>`
10. Configure the Vbrick host and the port to which the Uploader will connect.
    `uploader rev host <hostname>`
    `uploader rev port <port>`

    Note: The port defaults to 443 unless otherwise specified.

11. Add a Vbrick Rev user who has API permission to upload video recordings.
    `uploader rev user <username>`
12. Set the password for the user specified in step 11. Type
    `uploader rev password`
    you will be prompted for the password.

13. Create a certificate bundle (crt-bundle) holding a copy of the Root CA's certificate and all intermediate certificates in the chain for the Vbrick Rev server.
14. Add the certificate bundle created in step 13 to the Vbrick Rev trust store.
    uploader rev trust <crt-bundle>
15. Set access to the video recording.
    uploader access <Private|Public|AllUsers>
16. Give members of the space the ability to view or edit the recordings.
    uploader cospace_member_access <view|edit|none>

    Note: This step requires the listed members to have valid email addresses which are associated with accounts on Vbrick. For example **user1@example.com**

17. Decide whether the owner of the space is the single owner of the video recordings.
    uploader recording_owned_by_cospace_owner <true|false>

    Note: This step also requires the owner of the video recordings to have a valid email address which is associated with an account on Vbrick.

18. If the owner of the space is not listed in Vbrick Rev, then set the username of the fallback owner. If the fallback owner is not specified, then owner will default to the user configured on the MMP.
    uploader fallback_owner <vbrick-user>
19. Enable comments to the video recordings.
    uploader comments enable
20. Enable ratings for the video recordings.
    uploader ratings enable
21. Set the download permission for the video recordings.
    uploader downloads enable
22. Set the default state of the video recording when first uploaded to Vbrick Rev.
    uploader initial_state <active|inactive>
23. Decide whether to delete the video recording from the NFS after upload is complete
    uploader delete_after_upload <true|false>
24. Enable the Uploader to access the Meeting Server
    uploader enable

Note: Set **messageBoardEnabled** to **true** to see the messages being posted in the space indicating that the recording is available.

15 Streaming meetings

The Streamer component adds the capability of streaming meetings held in a space to the URI configured on the space.

An external streaming server needs to be configured to be listening on this URI. The external streaming server can then offer live streaming to users, or it can record the live stream for later playback.

Note: The Streamer component supports the RTMP standard in order to work with third party streaming servers that also support the RTMP standard. However, we have only tested against Vbrick as an external streaming server.

The Streamer connects to an external server using RTMP with an overall bitrate of 2Mbps. The video is encoded using H.264 at 720p30, while the audio is 64kbps AAC-LC. All traffic between the Streamer and the external streaming server is unencrypted.

The Streamer should be hosted on another Meeting Server instance than the server hosting the Call Bridge, see Figure 24. If the Streamer is hosted on the same server as the Call Bridge (local), then it should only be used for testing purposes.

The recommended deployment for production usage of the Streamer is to run it on a separate VM. This VM should be sized with 1 vCPU and 1GB of memory per 6 concurrent streams, with a minimum of 4 vCPUs and a maximum of 32vCPUs.

Note: These VM specifications are currently being evaluated, and the sizes are likely to be reduced.

For more details on VM specification see Unified Communications in a Virtualized Environment – Cisco (https://www.cisco.com/c/dam/en/us/td/docs/voice_ip_comm/uc_system/virtualization/cisco-collaboration-virtualization.html) and the Cisco white paper on Load Balancing Calls Across Cisco Meeting Server.

Where possible, it is recommended that the Streamer is deployed in the same physical locality as the Call Bridge to ensure low latency and high network bandwidth. If there are network connection issues between the Call Bridge and the Streamer, then the resultant stream could be affected.

Note: you may need to open firewall ports if the streaming destination URIs are on the external side of a firewall.

Figure 24: Permitted deployment for streaming: remote mode

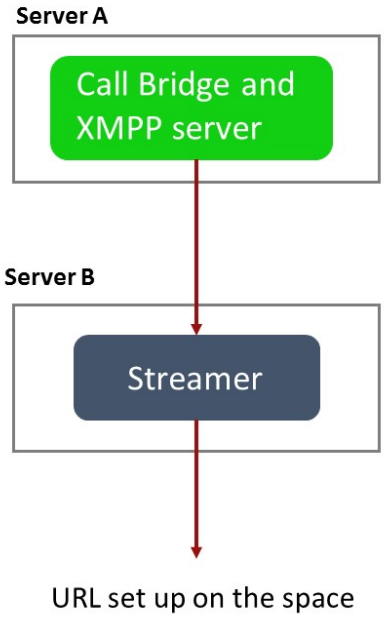

The Streamer also supports redundant configurations, see Figure 25, 15.1, 15.1 and 15.1. If you use multiple streamers then the solution load balances between available streaming devices. To restrict the use of specific Streamers to specific Call Bridges use the Call Bridge Group functionality introduced in version 2.1.

Figure 25: Permitted deployments for streaming: multiple streamers

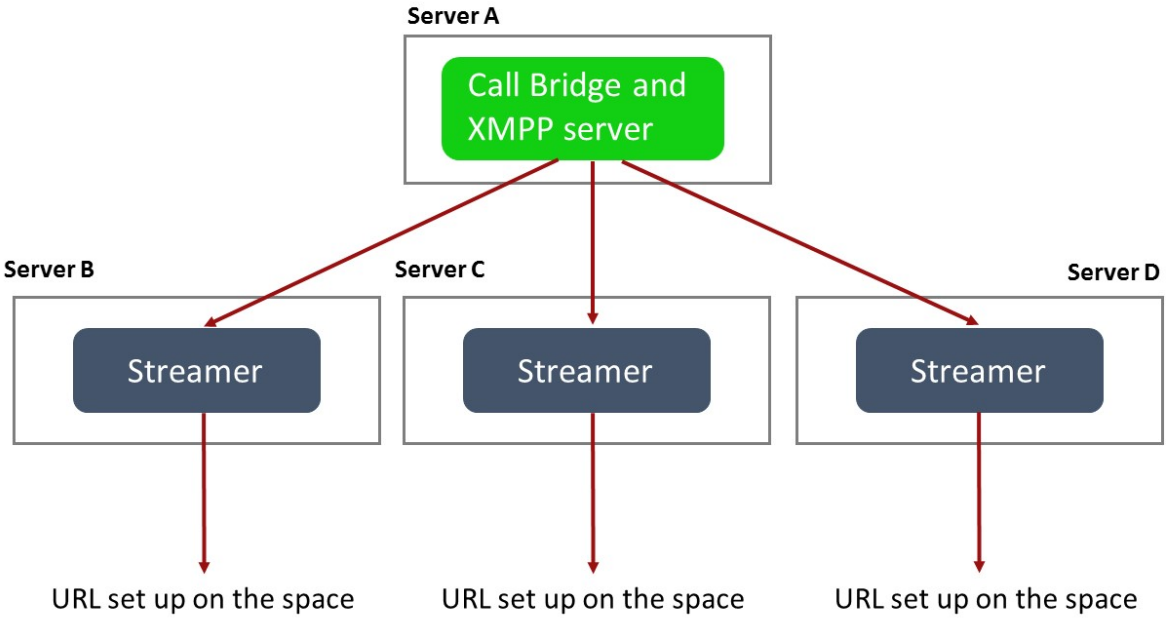

For testing purposes the Streamer can be co-located on the same server as the Call Bridge. This may support between 1 to 2 simultaneous streamings.

Note: Acano X series servers used in the single combined deployment mode should only be used for testing the Streamer, they should not be used in production networks to host the Streamer.

15.1 Overview of steps to configuring the Streamer

Use MMP commands to enable the Streamer on a Meeting Server, specify which Call Bridges within the deployment will work with the Streamer and where to save the streamings.

- Specify the HTTPS URL address that the Call Bridge will use to reach this streamer. Either POST the URL to the `/streamers` object or PUT to the `/streamers/<streamer id>` object.
- Use the new `streamingMode` parameter on the API object `/callProfiles` or `/callProfiles/<call profile id>` to select whether a meeting can be streamed or not. Options for this are:

  Automatic - streaming occurs without any user intervention, if streaming cannot occur the meeting still occurs.

  Manual - users can manually start and stop the streaming using DTMF.

  Disabled - no users can stream.

- Control which users have permission to start and stop streaming by setting the **`streamingControlAllowed`** parameter on **`/callLegProfiles`**.
- For each space that a user would like to stream, POSTor PUT to **`/coSpaces`** the **`streamURL`** parameter specifying the streaming destination URL.
- Use the new startStreaming and stopStreaming parameters for **`/dtmfProfiles`** and **`/dtmfProfiles/<dtmf profile id>`** to map the DTMF tones for starting and stopping streaming.

15.2 Example of deploying streaming

Note: The Streamer behaves as an XMPP client, so the XMPP server needs to be enabled on the Meeting Server hosting the Call Bridge.

This example gives the steps to deploy a streamer remote to the Call Bridge. It assumes that you already have a working Call Bridge and XMPP server.

1. Create a certificate and private key for the Streamer, following the steps described in the Certificates Guidelines for an internal CA signed certificate.
2. SSH into the MMP of the Meeting Server hosting the Streamer.
3. Configure the Streamer to listen on the interface(s) of your choice with the following command:
   `streamer listen <interface[:port] whitelist>`
   The Streamer can listen on multiple interfaces, e.g. one on public IP and one on the internal network. (However, it cannot listen on more than one port on the same interface.)
   The following is an example where interfaces are set to interface A and B, both using port 8445.
   `streamer listen a:8445 b:8445`
   To use a local Streamer, the Streamer must listen on the loopback interface lo:8445, for example
   `streamer  listen lo:8445 b:8445`
4. Upload the certificate file, key file and certificate bundle to the MMP via SFTP.
   `streamer certs <keyfile> <certificatefile> [<crt-bundle>]`
5. Add the Call Bridge certificate to the Streamer trust store using the command:
   `streamer trust <crt-bundle>`
6. Use the streamer command to list the details for the streamer, for example:

```
cms1> streamer
Enabled                : true
Interface whitelist    : a:8445 b:8445
Key file               : streamer0.key
```

```
Certificate file        : streamer0.cer
CA Bundle file          : streamer.crt
Trust bundle            : callbridge.crt
```

7. Enable the Streamer:
   streamer enable
8. Create DNS A record for the Streamer and set it to resolve to the IP Address of the Ethernet interface you want the Streamer to listen on.
9. Use the API of the Meeting Server hosting the Call Bridge to configure the settings through which the Call Bridge will communicate with the Streamer.

   a. Specify the HTTPS URL address that the Call Bridge will use to reach this streamer. Either POST the URL to the **/streamers** object or PUT to the **/streamers/<streamer id>** object

   Note: If using a local Streamer, the URL must be the loopback interface, for example **https://127.0.0.1:8445**

   b. POST to **/coSpaces** or PUT to **/coSpaces/<coSpace id>** the **streamUrl** which determines where streaming is streamed to, if streaming is initiated

   c. Select whether a meeting can be streamed or not and whether the streaming will start without any user intervention. Use the **streamingMode** parameter on the API object **/callProfiles** or **/callProfiles/<call profile id>**

   d. Control which users have permission to start and stop streaming. Use the **streamingControlAllowed** parameter on **/callLegProfiles**

   e. For each space that a user would like to stream, POSTor PUT to **/coSpaces** the **streamURL** parameter specifying the destination URL to stream to.

   Note: some streaming services require username and password, others provide a unique stream key. For example, for vBrick:
   streamUrl=rtmp://<username>:<password>@<vbrick IP/FQDN>/live/PullStream1

   and for YouTube:
   streamUrl=rtmp://a.rtmp.youtube.com/live2/<stream key>

   f. Use the **startStreaming** and **stopStreaming** parameters for **/dtmfProfiles** and **/dtmfProfiles/<dtmf profile id>** to map the DTMF tones for starting and stopping streaming. For example: **7 to start and **8 to stop streaming.

15.3 Streamer licensing

You will need one or more licenses for streaming which is loaded on the Meeting Server hosting the Call Bridge, not the server hosting the Streamer. One 'recording' license supports 1 concurrent streaming or 1 recording, existing recording licences will allow streaming. Contact your Cisco sales representative or partner to discuss your licensing requirements.

16 Support for ActiveControl

From version 2.1, the Meeting Server supports ActiveControl for hosted calls. For participants using a Cisco SX, MX or DX endpoint with CE 8.3+ software installed, ActiveControl allows the meeting participant to receive details of the meeting and perform a few administrative tasks during the meeting, using the endpoint interface.

16.1 ActiveControl on the Meeting Server

The Meeting Server supports sending the following meeting information to ActiveControl enabled endpoints:

- Participant list (also known as the roster list) so that you can see the names of the other people in the call and the total number of participants,
- indicator of audio activity for the currently speaking participant,
- indicator of which participant is currently presenting,
- Indicators telling whether the meeting is being recorded or streamed, and if there are any non-secure endpoints in the call,
- on screen message which will be displayed to all participants, see Section 2.5.

In addition, the Meeting Server can control the following features on ActiveControl enabled endpoints:

- select the layout to be used for the endpoint,
- disconnect other participants in the meeting, see Section 16.4

Note: These features are configured using the API of the Meeting Server, see defaultLayout parameter on the API objects: /calls, /callLegProfile and /coSpace.

16.2 Limitations

- If an ActiveControl enabled call traverses a Unified CM trunk with a Unified CM version lower than 9.1(2), the call may fail. ActiveControl should not be enabled on older Unified CM trunks (Unified CM 8.x or earlier).
- ActiveControl is a SIP only feature. H.323 interworking scenarios are not supported.

Note: ActiveControl uses UDT transport for certain features, for example sending roster lists to endpoints and allowing users to disconnect other participants while in a call. See Section 16.4 for the steps to follow on the Meeting Server.

16.3 Overview on ActiveControl and the iX protocol

ActiveControl uses the iX protocol, which is advertised as an application line in the SIP Session Description Protocol (SDP). The Meeting Server automatically supports ActiveControl, and the feature cannot be disabled. In situations where the far end network is not known or is known to have devices that do not support iX, it may be safest to disable iX on SIP trunks between the Meeting Server and the other call control or Video Conferencing devices. For instance:

- for connections to Unified CM 8.x or earlier systems the older Unified CM systems will reject calls from ActiveControl-enabled devices. To avoid these calls failing, leave iX disabled on any trunk towards the Unified CM 8.x device in the network. In cases where the 8.x device is reached via a SIP proxy, ensure that iX is disabled on the trunk towards that proxy.
- for connections to third-party networks. In these cases there is no way to know how the third-party network will handle calls from ActiveControl-enabled devices, the handling mechanism may reject them. To avoid such calls failing, leave iX disabled on all trunks to third-party networks.
- for Cisco VCS-centric deployments which connect to external networks or connect internally to older Unified CM versions. From Cisco VCS X8.1, you can turn on a zone filter to disable iX for INVITE requests sent to external networks or older Unified CM systems. (By default, the filter is off.)

16.4 Disable UDT within SIP calls

ActiveControl uses the UDT transport protocol for certain features, for example sending roster lists to endpoints, allowing users to disconnect other participants while in a call, and inter-deployment participation lists. UDT is enabled by default. You can disable UDT for diagnostic purposes, for example if your call control does not use UDT, and you believe this is the reason the call control does not receive calls from the Meeting Server.

Using the Meeting Server API:

1. Create a compatibility profile with the sipUdt parameter set to “false”. Either POST sipUdt=false to the **`/compatibilityProfiles`** object or PUT to **`/compatibilityProfiles/<compatibility profile id>`** object
2. Disable the use of UDT at the system level, by adding the compatibilityProfile parameter and id (from step 1) to the system profile. PUT compatibityProfile=<compatibility profile id> to the **`/system/profiles/`** object.

16.5 Enabling iX support in Cisco Unified Communications Manager

Support for the iX protocol is disabled by default in Cisco Unified Communications Manager. To enable iX support, you must first configure support in the SIP profile and then apply that SIP profile to the SIP trunk.

Configuring iX support in a SIP profile

1. Choose **Device > Device Settings > SIP Profile**. The Find and List SIP Profiles window displays.
2. Do one of the following:
   a. To add a new SIP profile, click **Add New**.
   b. To modify an existing SIP profile, enter the search criteria and click **Find**. Click the name of the SIP profile that you want to update.

   The SIP Profile Configuration window displays.
3. Check the box for **Allow iX Application Media**
4. Make any additional configuration changes.
5. Click **Save**

Applying the SIP profile to a SIP trunk

1. Choose **Device > Trunk**.

   The Find and List Trunks window displays.
2. Do one of the following:
   a. To add a new trunk, click **Add New**.
   b. To modify a trunk, enter the search criteria and click **Find**. Click the name of the trunk that you want to update.

   The Trunk Configuration window displays.
3. From the SIP Profile drop-down list, choose the appropriate SIP profile.
4. Click **Save**.
5. To update an existing trunk, click **Apply Config** to apply the new settings.

16.6 Filtering iX in Cisco VCS

To configure the Cisco VCS to filter out the iX application line for a neighbor zone that does not support the protocol, the zone must be configured with a custom zone profile that has the SIP UDP/IX filter mode advanced configuration option set to On.

To update advanced zone profile option settings:

1. Create a new neighbor zone or select an existing zone (**Configuration > Zones > Zones**).
2. In the Advanced parameters section, for **Zone profile**, choose *Custom* if it is not already selected. The zone profile advanced configuration options display.

3. From the **SIP UDP/IX filter mode** drop-down list, choose **On**.
4. Click **Save**.

16.7 iX troubleshooting

Table 14: Call handling summary for calls that contain an iX header

Scenario	Outcome
Unified CM 8.x or earlier	Calls fail
Unified CM 9.x earlier than 9.1(2)	Calls handled normally but no ActiveControl
Unified CM 9.1(2)	Calls handled normally plus ActiveControl
Endpoint - no support for iX and no SDP implementation	Endpoint may reboot or calls may fail

17 Additional security considerations & QoS

This chapter discusses other security features available on the Meeting Server that are in addition to authentication provided through X.509 certificates and public keys.

Note: The commands listed in this chapter are also listed in the MMP Command Reference guide.

17.1 Common Access Card (CAC) integration

The Common Access Card (CAC) is used as an authentication token to access computer facilities. The CAC contains a private key which cannot be extracted but can be used by on-card cryptographic hardware to prove the identity of the card holder.

The Meeting Server supports administrative logins to the SSH and Web Admin Interface using CAC. Use the MMP commands in Table 15 below to configure CAC for your deployment.

Table 15: MMP commands to configure CAC logins

MMP commands	Description
`cac enable|disable [strict]`	Enables/disables CAC mode with optional strict mode removing all password-based logins
`cac issuer <ca cert-bundle>`	Identifies trusted certificate bundle to verify CAC certificates
`cac ocsp certs <keyfile> <certificatefile>`	Identifies certificate and private key for TLS communications with OCSP server, if used
`cac ocsp responder <URL>`	Identifies URL of OCSP server
`cac ocsp enable|disable`	Enables/disables CAC OCSP verification

17.2 Online Certificate Status Protocol (OCSP)

OCSP is a mechanism for checking the validity and revocation status of certificates. The MMP can use OCSP to work out whether the CAC used for a login is valid and, in particular, has not been revoked.

17.3 FIPS

You can enable a FIPS 140-2 level 1 certified software cryptographic module, then cryptographic operations are carried out using this module and cryptographic operations are restricted to the FIPS approved cryptographic algorithms.

Table 16: MMP commands to configure FIPS

MMP commands	Description
fips enable|disable	Enables/disables the FIPS-140-2 mode cryptography for all cryptographic operations for network traffic. After enabling or disabling FIPS mode, a reboot is required
fips	Displays whether FIPS mode is enabled
fips test	Runs the built-in FIPS test

17.4 TLS certificate verification

You can enable Mutual Authentication for SIP and LDAP in order to validate that the remote certificate is trusted. When enabled, the Call Bridge will always ask for the remote certificate (irrespective of which side initiated the connection) and compare the presented certificate to a trust store that has been uploaded and defined on the server.

Table 17: MMP commands to configure TLS

MMP commands	Description
tls <sip|ldap> trust <crt bundle>	Defines Certificate Authorities to be trusted
tls <sip|ldap> verify enable|disable|ocsp	Enables/disables certificate verification or whether OCSP is to be used for verification
tls <sip|ldap>	displays current configuration

17.5 User controls

MMP admin users can:

- Reset another admin user's password
- Set the maximum number of characters that can be repeated in a user's password – and there are a number of other user password rule additions
- Limit MMP access by IP address
- Disable MMP accounts after configurable idle period

17.6 Firewall rules

The MMP supports the creation of simple firewall rules for both the media and admin interfaces. Note that this is not intended to be a substitute for a full standalone firewall solution and therefore is not detailed here.

Firewall rules must be specified separately for each interface. After setting up a firewall rule on an interface, remember to enable the firewall on that interface. See the MMP Command Reference for full details and examples.

CAUTION: We recommend using the serial console to configure the firewall, because using SSH means that an error in the rules would make the SSH port inaccessible. If you must use SSH then ensure that an allow **`ssh rule`** is created for the ADMIN interface before enabling the firewall.

17.7 DSCP

You can enable DSCP tagging for the different traffic types on the Meeting Server (see the MMP Command Reference).

1. Sign in to the MMP.
2. Use **`dscp (4|6) <traffic type> (<DSCP value>|none)`** to set the DSCP values as required. For example: **`dscp 4 oa&m 0x22`** which sets operations, administration and management for IPv4.
3. Alternatively, use the **`dscp assured (true|false)`** command to force the use of the assured or non-assured DSCP values for the "voice" and "multimedia" traffic types. For example: **`dscp assured true`**

Note: DSCP tagging is for all packets being sent from the Meeting Server only. For PC Client DSCP tagging, Group Policy must be used to define desired DSCP values because Windows controls this, and normal user accounts have no permissions to set DSCP.

18 Diagnostic tools to help Cisco Support troubleshoot issues

18.1 Log bundle

From version 2.2, the Meeting Server can produce a log bundle containing the configuration and state of various components in the Meeting Server. This log bundle will aid Cisco Support speed up their analysis of your issue. It will include some of the following files:

- syslog
- live.json
- dumps
- db

If you need to contact Cisco support with an issue, follow these steps to download the log bundle from the Meeting Server.

1. Connect your SFTP client to the IP address of the MMP.
2. Log in using the credentials of an MMP admin user.
3. Copy the file logbundle.tar.gz to a local folder.
4. Rename the file, changing the logbundle part of the filename to identify which server produced the file. This is important in a multi-server deployment.
5. Send the renamed file to your Cisco Support contact for analysis.

Initial file size of the log bundle.tar.gz is 1 Kb, after transfer via SFTP the size will increase depending on the number of files and their size.

18.2 Ability to generate a keyframe for a specific call leg

A new **`generateKeyframe`** object has been added to **`/callLegs/<call leg id>`**. POST to **`/callLegs/<call leg id>/generateKeyframe`** to trigger the generation of a new keyframe in outgoing video streams for the call leg in question. This is a debug facility, and Cisco Support may ask you to use the feature when diagnosing an issue.

18.3 Reporting registered media modules in syslog

From version 2.2, syslog will now print a message every 15 minutes to allow people to monitor whether all media modules are alive and well.

An example from an Acano X3 server:

```
Apr 21 09:53:50 user.info cms-emea-01 host: server: INFO : media module status
11111111111
```

Appendix A DNS records needed for the deployment

Note: You can configure the DNS resolver(s) to return values which are not configured in external DNS servers or which need to be overridden; custom Resource Records (RRs) can be configured which will be returned instead of querying external DNS servers. (The RR is not available to clients.) See the MMP Command Reference for details.

Note: Verify that no A or SRV records already exist for any Meeting Servers before defining the records below.

Table 18: DNS records required for deployment

Type	Example and Description
SRV(*)	**_xmpp-client._tcp.example.com** Resolves to: The A record **xmpp.example.com** below. Usually this is port 5222. Description: Used by clients to login. The SRV record must correspond to the domain used in your XMPP usernames.
SRV(*)	**_xmpp-server._tcp.example.com** Resolves to: The A record xmpp.example.com below. Usually this is port 5269. Description: Used to federate between XMPP servers. The SRV record must correspond to the domain used in your XMPP usernames.
A	**xmpp.example.com** Resolves to: IP address of the XMPP server. Description: Used by clients to login.
A / AAAA	**join.example.com** Resolves to: IP address of Web Bridge. Description: This record is not used by the Meeting Server directly; however, it is common practice to provide an end user with an FQDN to type into the browser which resolves to the Web Bridge. There is no restriction or requirement on the format of this record.

Type	Example and Description
A / AAAA	**uk.example.com** Resolves to: IP address of the Call Bridge. Description: Used by the Lync FE server to contact the Call Bridge.
A / AAAA	**ukadmin.example.com** Resolves to: IP address of the MMP Interface IP address of the Web Admin Interface. Description: This record is used purely for admin purposes; when system administrators prefer a FQDN to remember for each MMP interface.
SRV(*)	**_sipinternaltls._tcp.<yourLyncdomain>** Resolves to: The A record of the Lync FE server or FE Pool. Description: If you have an FE pool, you can have multiple FE records pointing to individual FE servers within the pool. You also need this record if you want Meeting Server to resolve Lync meetings by Lync meeting IDs.
A / AAAA	**fe.<yourLyncdomain>** Resolves to: IP address of the Lync FE server. Description: You will need one record for each individual FE server.
SRV(*)	**_sipfederationtls._tcp.<yourSIPdomain>** Resolves to: The FQDN of the Call Bridge. Description: This record is required for Lync federation.
A	**callbridge.example.com** Resolves to: IP address of the Call Bridge. Description: Required for Lync federation as the Call Bridge will need to have a public IP address, and NAT is not supported in this scenario.

(*) SRV records do not resolve directly to IP addresses. You need to create associated A or AAAA name records in order to satisfy the SRV requirements.

Appendix B Ports required for the deployment

The following diagram shows the connections to the Meeting Server and location of the firewall in a single combined server deployment. Use the tables below the diagram to identify which ports to open

Figure 26: Ports that must be open in a single combined server deployment

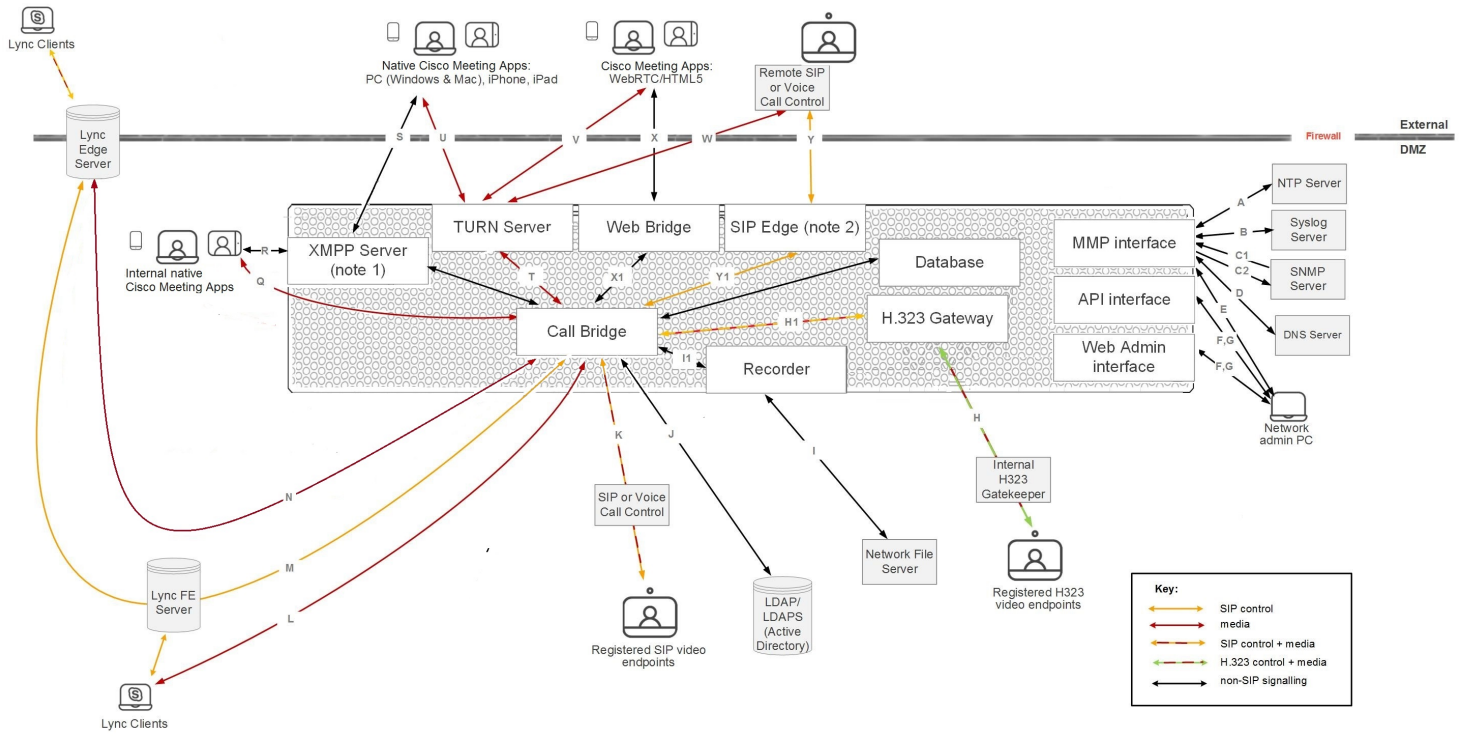

Note:
1) The figure above shows the XMPP server listening on an external port. If you prefer the XMPP server to not listen on one of the interface ports (A-D), then instead configure the Load Balancer to listen on the external port and have the Load Balancer relaying the information to the XMPP server.

B.1 Configuring the Meeting Server

Table 19 lists the ports to use to configure the Meeting Server.

Table 19: Ports for administration of the Meeting Server

Code	Connect to	Destination port to open	Method	Traffic type	Traffic direction with respect to Meeting Server	Additional information
E	MMP	22	SSH	TCP	Incoming	Secure login to MMP
F	API or Web Admin	80	HTTP	TCP	Incoming	Port configurable through MMP
G	API or Web Admin	443	HTTPS	TCP	Incoming	Port configurable through MMP

B.2 Connecting services

Use Table 20 to identify which ports are used to connect different services to the Cisco Meeting App .

Table 20: Ports to open to connect services

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
A	MMP	NTP server	123	TCP or UDP	Outgoing	
B	MMP	Syslog server	514	TCP	Outgoing	Default port, different port configurable through MMP
C1	MMP	SNMP server	161	UDP	Incoming	
C2	MMP	SNMP TRAP	162	TCP or UDP	Outgoing	

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
D	MMP/Call Bridge/Web Bridge	DNS server	53	TCP or UDP	Outgoing	On X series servers both the Admin port and the interface being used (A to D) need to be able to access the DNS server on port 53.
	Call Bridge	CDR recipient device		TCP	Outgoing	set URI of CDR recipient in Web Admin interface, or API using API object /system/cdrReceivers/

B.3 Using Meeting Server components

Use Table 21 to identify which ports are used to connect to the components in the Cisco Meeting App.

Table 21: Ports to open on the Meeting Server to use components

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
H1	Call Bridge	H.323 Gateway	6061	TCP (SIP)	Outgoing	Port configurable through MMP
H	H.323 Gateway	H.323 Gatekeeper	1720	TCP (H.225)	Incoming	Port not configurable
			port on H.323 Gatekeeper for next hop	TCP (H.225)	Outgoing	
H	H.323 Gateway	H.323 Gatekeeper	1024-65535 (note 1)	TCP (H.245)	Incoming	Port not configurable
			port on H.323 Gatekeeper for next hop	TCP (H.245)	Outgoing	

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
H	H.323 Gateway	H.323 Gatekeeper	32768-65535 (note 2)	UDP media	Incoming and outgoing	
I1	Call Bridge	Recorder	8443		Outgoing	Port configurable through MMP. For a local recorder use the loopback interface, eg lo:8443
I	Recorder	Network File Server (NFS)				Use the MMP command **`recorder nfs <hostname/IP<directory>`** to specify where to store the recordings on the NFS
J	Call Bridge	LDAP/LDAPS (Active Directory)	389/636 (note 3)	TCP/TCP (SIP TLS)	Outgoing	Port configurable through Web Admin interface
K	Call Bridge	Internal registered SIP endpoint or voice call control	5060	SIP UDP	Incoming and outgoing	
K	Call Bridge	Internal registered SIP endpoint or voice call control	5060	TCP (SIP)	Incoming and outgoing	
K	Call Bridge	Internal registered SIP endpoint or voice call control	5061	TCP (SIP TLS)	Incoming and outgoing	
K	Call Bridge	Internal registered SIP endpoint or voice call control	32768-65535	UDP (STUN RTP, BFCP)	Incoming	
L	Call Bridge	Lync client, AVMCU	32768-65535	UDP (STUN RTP)	Incoming	

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
L	Call Bridge	Lync client, AVMCU	1024-65535 (note 1)	UDP (STUN RTP)	Outgoing	
L	Call Bridge	Lync client, AVMCU	32768-65535	TCP (RDP)	Incoming	
L	Call Bridge	Lync client, AVMCU	1024-65535 (note 7)	TCP (RDP)	Outgoing	
M	Call Bridge	Lync FE server	5061	TCP (SIP TLS)	Incoming and outgoing	
N	Call Bridge	Lync edge server	3478	UDP	Outgoing	
N	Call Bridge	Lync edge server	443	TCP	Outgoing	
N	Call Bridge	Lync edge server	32768-65535 (note 2)	UDP (STUN RTP)	Incoming	
	Call Bridge	XMPP server				Internal to Meeting Server, does not require open ports
Q	Call Bridge	Internal native Cisco Meeting Apps	32768-65535	UDP (STUN RTP)	Incoming	
Q	Call Bridge	Internal native Cisco Meeting Apps	1024-65535 (note 1)	UDP (STUN RTP)	Outgoing	
R, S	XMPP server	Internal or External native Cisco Meeting Apps	5222	TCP	Incoming	For both internal and external native Cisco Meeting Apps

Code	Component	Connecting to	Destination port to open	Traffic type	Traffic direction with respect to component	Additional information
T, U,V,-W	TURN server	Call Bridge and remote devices	32768-65535 (notes 2 and 4)	Media UDP (STUN RTP)	Incoming and outgoing	
T, U,V,-W	TURN server	Call Bridge and remote devices	32768-65535 (notes 2 and 4)	Media TCP (STUN RTP)	Incoming and outgoing	
T, U,V,-W	TURN server	Call Bridge and remote devices	3478 (note 4)	UDP (STUN)	Incoming	
T, U,V,-W	TURN server	Call Bridge and remote devices	3478 (note 4)	TCP (STUN)	Incoming	
T, U,V,-W	TURN server	Call Bridge and remote devices	443 (notes 4, 5, 6)	UDP (STUN)	Incoming	
T, U,V,-W	TURN server	Call Bridge and remote devices	443 (notes 4, 5, 6)	TCP (STUN)	Incoming	
X	Web Bridge	WebRTC clients	80	TCP (HTTP)	Incoming	
X	Web Bridge	WebRTC clients	443 (notes 6 and 8)	TCP (HTTPS)	Incoming	
X1	Call Bridge	Web Bridge		TCP		Internal to Meeting Server, does not require open ports
Y	SIP Edge server (public interface)	Remote SIP devices	5061	TCP (SIP TLS) (note 9)	Incoming and outgoing	Port configurable through MMP
Y1	SIP Edge server (private interface)	Call Bridge	any unused port e.g. 3061	TCP (SIP TLS)	Incoming	Port configurable through MMP
	Call Bridge	Database				Internal to Meeting Server, does not require open ports

Note:
1) Exact range depends on far end.
2) Although the range is shown as 32768-65535, currently only 50000-51000 is used. A wider range is likely to be required in future releases.
3) Port 636 (secure) and 389 (non-secure) are commonly used for this function but the port is configurable through the Web Admin interface. The same applies to 3268 and 3269 (non-secure and secure) global catalog LDAP requests.
4) If the media ports (32768-65535) are not open then TCP/UDP port 3478/443, used to connect to the TURN server, will be used to relay media.
5) UDP/TCP port 443 can be changed. Using the MMP command turn tls <port> will change the second UDP/TCP port that the TURN server listens on.
6) The TURN server will not listen on port 443 on the loopback interface. This is to avoid port clashes with other services that may be running on port 443 on the loopback interface.
7) Exact range depends on configuration of Lync server.
8) To run both the TURN server and the Web Bridge on port 443 requires the two components to be run on different interface: port combination, if this is not possible then use port 447 for the TURN server.
9) Port 5061 only supports SIP TLS, there is no support for UDP or TCP.
10) Lync clients includes AVMCU—the same ports need opening.

Appendix C Growth in scaling deployments

Table 22 below demonstrates the expansion in maximum call capacities on Meeting Servers by upgrading to later software versions. Bold indicates a new feature in that software version. Note that there are different capacities for a single or cluster of Meeting Servers compared to load balancing calls within a Call Bridge Group.

Table 22:Evolution in Meeting Server call capacity

Software version		2.0		2.1		2.2		2.3		2.4 and 2.5		2.8		
Cisco Meeting Server platform		1000	2000	1000	2000	1000	2000	1000	2000	1000	2000	1000 M4	1000 M5	2000
Individual Meeting Servers or Meeting Servers in a cluster (notes 1,2 3 and 4)	1080p30 720p30 SD Audio	**48** **96** **192** **3000**	NA NA NA NA	48 96 192 3000	NA NA NA NA	48 96 192 3000	**250** **500** **1000** **3000**	48 96 192 3000	250 500 1000 3000	48 96 192 3000	**350** **700** 1000 3000	48 96 192 **1700**	48 96 192 **2200**	350 700 1000 3000
	HD par-ticipants per conference per server	96	NA	96	NA	96	**100**	96	100	96	**450**	96	96	450
	WebRTC connections per Web Bridge	100	NA	100	NA	100	100	100	100	100	100	100	100	100

Table 22:Evolution in Meeting Server call capacity (....continued)

Software version Cisco Meeting Server platform		2.0 1000	2.0 2000	2.1 1000	2.1 2000	2.2 1000	2.2 2000	2.3 1000	2.3 2000	2.4 and 2.5 1000	2.4 and 2.5 2000	2.8 1000 M4	2.8 1000 M5	2.8 2000
Meeting Servers in a Call Bridge Group	Call type supported	NA	NA	**Inbound SIP**	NA	Inbound SIP **Outbound SIP**		Inbound SIP Outbound SIP **Cisco Meeting App**		Inbound SIP Outbound SIP Cisco Meeting App		Inbound SIP Outbound SIP Cisco Meeting App	Inbound SIP Outbound SIP Cisco Meeting App	Inbound SIP Outbound SIP Cisco Meeting App
	1080p30 720p30 SD Audio Load limit	NA	NA	**48** **96** **192** **3000** **96,000**	NA	48 96 192 3000 96,000	**250** **500** **1000** **3000** **500,000**	48 96 192 3000 96,000	250 500 1000 3000 500,000	48 96 192 3000 96,000	250 500 1000 3000 500,000 (note 5)	48 96 192 1700 96,000	48 96 192 2200 96,000	250 700 1000 3000 700,000 (note 5)
	Number of HD participants per conference per server	NA	NA	96	NA	96	**100**	96	100	96	**450**	96	96	450
	WebRTC connections per Web Bridge	NA	NA	100	NA	100	100	100	100	100	100	100	100	100

Note 1: Maximum of 24 Call Bridge nodes per cluster; cluster designs of 8 or more nodes need to be approved by Cisco, contact Cisco Support for more information.

Note 2: Clustered Cisco Meeting Server 2000's without Call Bridge Groups configured, support integer multiples of maximum calls, for example integer multiples of 700 HD calls.

Note 3: Up to 16,800 HD concurrent calls per cluster (24 nodes x 700 HD calls).

Note 4: A maximum of 2600 participants per conference per cluster depending on the Meeting Servers platforms within the cluster.

Note 5: For versions 2.4 and 2.5, a Cisco Meeting Server 2000 with Call Bridge Groups enabled only supports 500 HD calls. The **loadLimit** for a Cisco Meeting Server 2000 is set at 500,000 due to the total number of video calls supported remaining at 1,000. The Meeting Server does not offer a way to limit the minimum resolution of calls, so it will always be possible to have a mix of HD and lower resolution calls. There is no mechanism to make use of the increased scale in HD or full HD in a load balanced configuration.

Note 6: VMWare have made changes in ESXi versions 6.0 update 3, 6.5 update 2 and 6.7 in order to overcome security concerns with Intel chipsets. As a consequence of these changes the throughput of audio calls on Cisco Meeting Server has been reduced in version 2.8 (video capacity is unaffected).

Appendix D Activation key for unencrypted SIP media

Prior to version 2.4, you could only purchase an activation key with SIP media encryption enabled. Media includes audio, video, content video and ActiveControl data. From version 2.4, you have the choice of purchasing an activation key with SIP media encryption enabled or SIP media encryption disabled (unencrypted SIP media) for the Cisco Meeting Server 1000, Cisco Meeting Server 2000 and the VM software image. Choose either encrypted or unencrypted options under the software pids R-CMS-K9 and R-CMS-2K-K9.

Note: Current Call Bridge activations are unaffected, unless an activation key is uploaded with SIP media encryption disabled.

D.1 Unencrypted SIP media mode

If the activation key for SIP media encryption disabled is uploaded to the Meeting Server, then the following occurs:

- media sent between the Meeting Server and SIP devices is unencrypted,
- media sent over distribution links between clustered Call Bridges is unencrypted,
- call signalling remains encrypted,
- media in calls between the Meeting Server and Cisco Meeting App, on any platform, remains encrypted,
- an error message is returned if the **`sipMediaEncryption`** parameter is set to anything other than **`prohibited`** on the following API objects:

  `/calls/<call id>/participants`

  `/calls/<call id>/callLegs`

  `/callLegs/<call leg id>`

  `/callLegProfiles` and **`/callLegProfiles/<call leg profile id>`**

  `/callLegs/<call leg id>/callLegProfileTrace`
- an error message is displayed if the **SIP media encryption** field on the the **Configuration**>**Call settings** web page of the Web Admin interface is set to anything other than **disabled**.

Note: If SIP media encryption is disabled, call signaling can still be encrypted on outbound calls, if required, by setting the **`sipControlEncryption`** parameter on **`/outboundDialPlanRules`**.

D.2 Determining the Call Bridge media mode

To determine whether the Call Bridge uses encrypted or unencrypted SIP media use a GET on API object **`/system/licensing`**. If the **`features`** response value has the **`status`** of **`callBridgeNoEncryption`** set to **`activated`** then an activation key for unencrypted media is loaded on theCall Bridge. Other valid settings for the **`status`** of **`callBridgeNoEncryption`** are **`noLicense`**, **`grace`** or **`expired`**.

`callBridgeNoEncryption` also has an **`expiry`** field in the form of a string.

Appendix E Dual Homed Conferencing

E.1 Overview

Dual homed conferencing also improves the user experience for both Lync client users and Cisco Meeting App users in Lync scheduled meetings and in Lync drag and drop style meetings (also known as ad hoc calls). Lync participants can use drag and drop to add Cisco Meeting App users to a Lync meeting, and can use conference controls to mute Cisco Meeting App users or disconnect them. For Cisco Meeting App users joining a Lync scheduled conference, they will see the video from up to five Lync participants, as well as video from the Cisco Meeting App users. Lync users see video in a gallery format from all of the Cisco Meeting App users, as well as the Lync users in the meeting. Both Lync users and Cisco Meeting App users receive a full combined list of participants in the meeting.

Note: The "Add Participant" button on the Lync/Skype for Business client does not work in ad hoc dual homed conferences. Do not use the "Meet Now" button as a workaround, as this will this will leave an active call between the Meeting Server and the AVMCU.

Lync participants can also directly dial into a Meeting Server space or use drag and drop to add a Meeting Server space to a Lync meeting. These are useful if a large meeting is being held in a Cisco Meeting Server space which the Lync user wants to join. In the first case they will receive a composed layout of multiple participants. When adding a complete space to a Lync meeting, the Lync user will receive only one video stream from the space (the main speaker) and will not receive a full combined participant list. They can continue to add additional Lync participants as normal.

Note: Dual-homed conferences with a Meeting Server cluster are not currently supported with Expressway X8.11 as the edge for the Meeting Server, unless at least some of the Microsoft traffic flows directly between one of the Meeting Servers in the cluster and the Microsoft infrastructure (and not through Expressway). Dual-homing is supported with Expressway X8.11 as the edge for standalone Meeting Servers.

E.2 Consistent meeting experience in dual homed conferences

The Meeting Server sends two H.264 video streams stream per video participant to the AVMCU, a high resolution video stream and a low resolution video stream, see Figure 27. Lync, Skype for Business and O365 clients that support the high resolution, subscribe to and receive the high quality video stream. Clients that select a lower quality, because of bandwidth restrictions, window size, layout, CPU power or being on a mobile device, subscribe to and

receive the lower quality streams, and do not reduce the video quality nor degrade the video experience for other participants.

Note: Ensure that the bandwidth of the SIP trunk is set sufficiently high to accommodate the two video streams. We recommend 8MB for LANs and 2.5MB for WANs.

Figure 27: Dual media streams to AVMCU

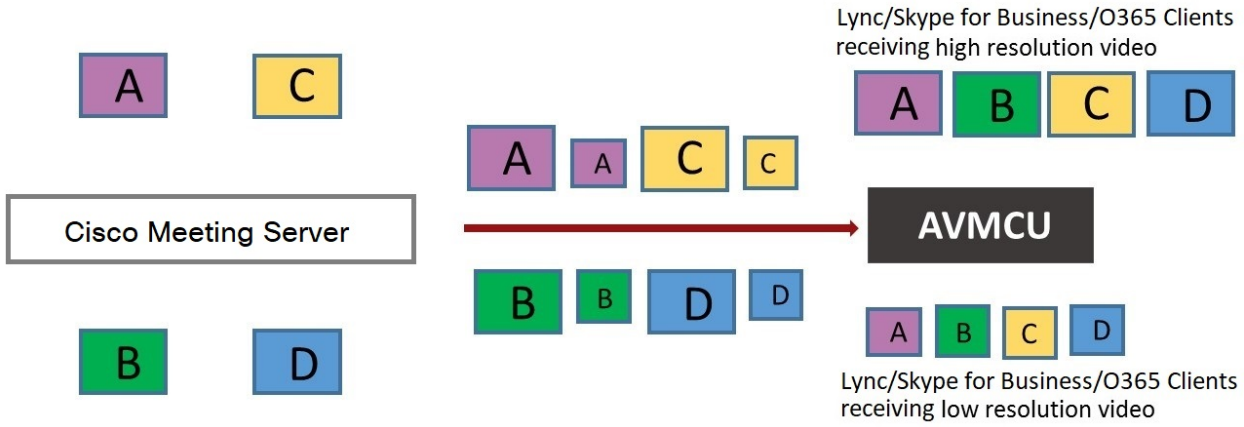

Note: Any devices using Microsoft RTVideo will not benefit from this feature.

E.2.1 Summary of user experiences

Dual homed conferencing combined with support for RDP and multiple video encoders, results in a richer meeting experience for both Lync and Cisco Meeting app users.

- Both Lync client users and Cisco Meeting App users see familiar screen layouts.
- Both Lync client users and Cisco Meeting App users receive a full combined list of all participants in the meeting, regardless of where they are connected.
- Lync client users see a non-square aspect ratio for video from SIP endpoints and Cisco Meeting Apps.
- Lync client users see content in a separate area of their screen rather than in the main video area.

- The Meeting Server sends video using the best quality codec supported by each participant in Lync meetings. This optimizes the experience for all Lync client users in a meeting, when a mixture of Lync client versions are used by participants.
- The Meeting Server sends two H.264 video streams stream per video participant to the AVMCU, a high resolution video stream and a low resolution video stream, to preserve the high resolution experience for clients that support it, when clients that can only support low resolution join the meeting.
- Chat works in Lync AVMCU conferences with Cisco Meeting App users in spaces. and in direct calls between a Cisco Meeting App user and a Lync client.

Note: For the best user experience during meetings, use Lync 2013, Skype for Business 2015 or later, which allow multiple video streams to be transmitted to the Meeting Server. This enables an endpoint or Cisco Meeting App user connecting to the Meeting Server to view multiple Lync participants. Lync 2010 only provides a single loudest speaker stream, if the loudest speaker is on the Meeting Server side of the conference already, then Cisco Meeting App users and SIP endpoint users will not view the Lync participants.

For more information on RDP and multiple video encoder support, see these FAQs:

- RDP support,
- multiple video encoder support.

E.3 Mute/unmute meeting controls in dual homed conferences

Version 2.4 of the Meeting Server software introduced improved mute/unmute meeting controls in dual homed conferences for:

- on-premise and Office 365 Lync/Skype for Business clients,
- end point users,
- Cisco Meeting App users.

Note: This section assumes that muting and unmuting is enabled using the API of the Meeting Server.

Muting/unmuting:

- Lync clients can mute and unmute anyone in the dual homed conference, this means themselves and others, and they can mute and unmute the audience too.
- All endpoint users can now mute Lync clients,

- Endpoint users on the Lync side of the AVMCU can now mute and unmute themselves (self) and other endpoints (either on the Lync clients/endpoints connected to the AVMCU or on the Meeting Server side). Prior to version 2.4, only endpoint users on the Meeting Server side of the AVMCU could mute and unmute themselves (self) and others.
  - For non-ActiveControl endpoints, the Meeting Server sends DTMF key sequences for each mute and unmute, and overlays an icon on the media stream to the endpoint to indicate whether the endpoint is muted or unmuted.
  - For ActiveControl endpoints running CE 9.2.1or later software, the endpoint handles the icons and messages (the Meeting Server does not overlay icons).
- Once an ActiveControl endpoint is muted it has to be unmuted locally so as to ensure the privacy of any local conversation. For example, when a remote participant mutes an ActiveControl endpoint and then tries to unmute it, the ActiveControl endpoint will mute itself again until it is locally unmuted.
- When a remote participant tries to unmute a non-ActiveControl endpoint, the non-ActiveControl endpoint will be unmuted.
- Cisco Meeting App users and Cisco Meeting Management users can mute and unmute Lync clients. They also see the correct mute state of all participants in the meeting.

Muting/unmuting Cisco Meeting App users:

- Information on local muting and unmuting of a Cisco Meeting App user is not passed to Lync clients in dual homed conferences. However, if a Lync client remotely mutes a Cisco Meeting App user and the Cisco Meeting App unmutes itself, the Meeting Server tells the Lync clients about the unmuting.
- When a remote participant tries to unmute a Cisco Meeting App user, the Cisco Meeting App user will remain locally muted. Note: other participants will still see them as unmuted, although they are actually muted.
- The Cisco Meeting App shows the mute/unmute state using its own icons. Meeting Server icons are not overlaid on the Cisco Meeting App video pane.

E.4 Configuring the Dual Homed Lync functionality

If you already have an on-prem Lync deployment or Lync Federation deployment working with the Meeting Server deployment, then no additional configuration is required on the Meeting server.

If this is a new deployment, then make sure that you configure the Lync Edge settings on the Meeting Server, see Section 8.5.

E.4.1 Troubleshooting

If users are unable to join a Lync conference via the IVR or using a dial plan rule that resolves to "Lync", the first thing to do is to verify that the "Lync Edge" settings have been set up - the same mechanism is used to resolve Lync conferences as is used to find the Edge server. The Meeting Server must query the Lync FE server to find both of these.

If this fails, a message will be logged in the event log to say that the conference ID cannot be found:

```
lync conference resolution: conference "1234" not found
```

This may mean that the conference does not exist, but there are also other possible causes.

If SIP traffic tracing is enabled, there should be a 'SERVICE' message sent to the Lync FE server just before the above message is logged, which should be replied to with a 200 OK. Check that this message is sent to the correct IP, which should be that of a Lync FE server.

If this message is not sent (it does not show up in the logs), then it is possible that the Call Bridge is unable to find the Lync server using a DNS SRV lookup for the **`_sipinternaltls._tcp.lyncdomain`** record, and so does not know where to send it. Enabling DNS tracing and retrying should confirm this. However this can also happen if the Lync Edge settings have not been configured on the Meeting Server.

If the Service message is sent but the Lync server replies with "403 unauthorized", then the most likely cause of this is that the local contact domain in the outbound dial plan rule for this Lync domain is not set correctly. It should be set to the FQDN of the Meeting Server, which should be the same as the FQDN supplied in the CN of the Call Bridge's certificate.

Appendix F More information on LDAP field mappings

This section provides additional information for LDAP field mappings that you set up for the Meeting Server.

Parts of an LDAP field value can be substituted by means of a sed-like construction, as follows:

$<LDAP field name>|'/<regex>/<replacement format>/<option>'$

where:

- **<option>** can be g, to replace every match of **<regex>** with **<replacement format>**, or blank to match only the first
- parts of **<regex>** can be tagged for use in **<replacement format>** by enclosing them in round brackets
- tagged matches can be referenced in **<replacement format>** as \x where x is a digit from 0 to 9. Match 0 corresponds to the entire match, and matches 1-9 the 1st to 9th tagged sub-expressions
- single quotes inside the substitution expression must be escaped with a backslash, as must backslash characters themselves
- any character other than a single quote, a backslash, or the digits 0-9 can be used in place of the forward slash that separates the components of the substitution expression
- if the separating character is to be used as a literal within the expression, it must be escaped with a backslash.

As an example, the following would convert addresses in the format:

firstname.lastname@test.example.com

into the format:

firstname.lastname@xmpp.example.com JIDs
$mail|'/@test/@xmpp/'$

and the following would remove every lower case 'a' from the user's full name:

$cn|'/a//g'$

A sensible set of expressions for use might be:

```
Full name:              $cn$
JID:                    $mail|'/@test/@xmpp/'$
space URI:            $mail|'/@.*//'$.space
space dial-in number: $ipPhone$
```

Note: The LDAP server credentials are used to read the following fields (for security reasons you may want to restrict the fields and permissions available using those credentials):

- mail
- objectGUID
- entryUUID
- nsuniqueid
- telephoneNumber
- mobile
- sn
- givenName

Appendix G Using TURN servers behind NAT

The TURN server can be deployed behind a NAT, and the NAT address specified using the MMP command **`turn public-ip`**. However, due to how Interactive Connectivity Establishment (ICE) works, careful configuration of the NAT is required to ensure connectivity always works.

This appendix provides an overview of how ICE works. It explains:

- how candidates are identified,
- how connectivity is checked,
- the effect of NAT in front of the TURN server,
- how NAT affects external Cisco Meeting App users.

Note: Issues can arise when the only available path includes both relay candidates. This requires the firewall to be correctly configured, so that all clients are able to send and receive video and audio.

G.1 Identifying candidates

ICE works by gathering a list of candidate addresses and ports, and then finding which pairs of these candidates allow media to be exchanged. When multiple candidate pairs are available then a priority scheme is used to determine which pair is used.

Typically, three candidates might exist:

1. Host candidate
2. Server Reflexive candidate
3. Relay candidate

G.1.1 Host candidate

The most simple candidate is the host candidate. This is the address used by the host interface. This is often on a local network and not routable.

G.1.2 Server Reflexive candidate

The server reflexive candidate is the address that the TURN server sees incoming packets coming from. To determine this, the host sends packets to a defined port on the TURN server (normally port 3478) and the TURN server replies with information about where the packets came from.

Figure 28: Server Reflexive candidate

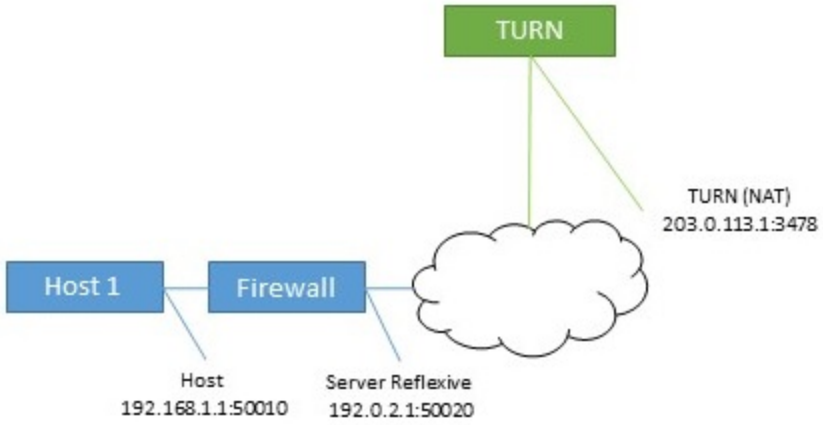

In cases where the host is behind a firewall carrying out NAT, then this is different to the host candidate. In many cases, packets sent to this port and address will be forwarded back to the host.

Figure 29: Effect of a host behind a firewall carrying out NAT

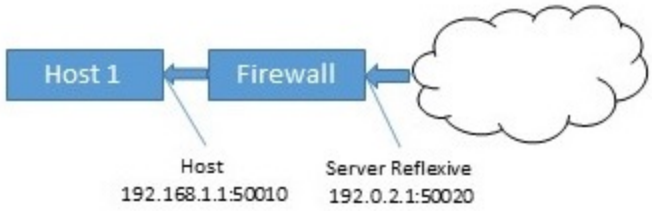

G.1.3 Relay candidate

The final candidate is the relay candidate. This candidate is created by the TURN server in response to requests from the host. The relay address of this candidate is the TURN server interface address, when NAT is used the relay address is changed to an address from NAT.

Figure 30: Relay candidate

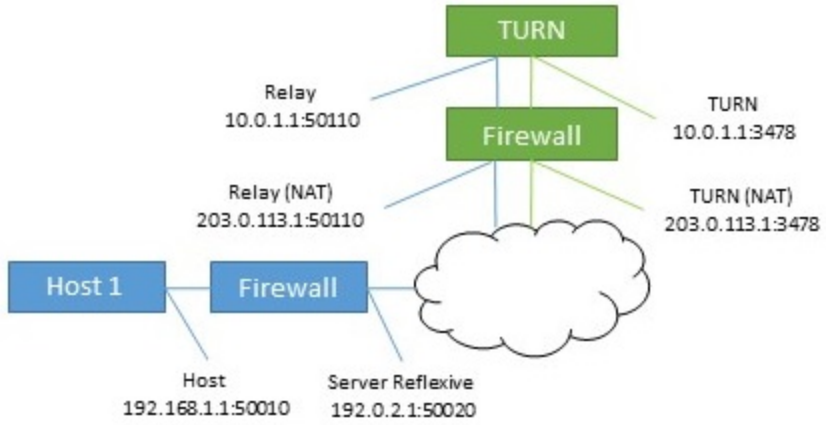

Data sent to this relay address is then sent back to the host via the TURN server.

Figure 31: TURN server returns relay address to host

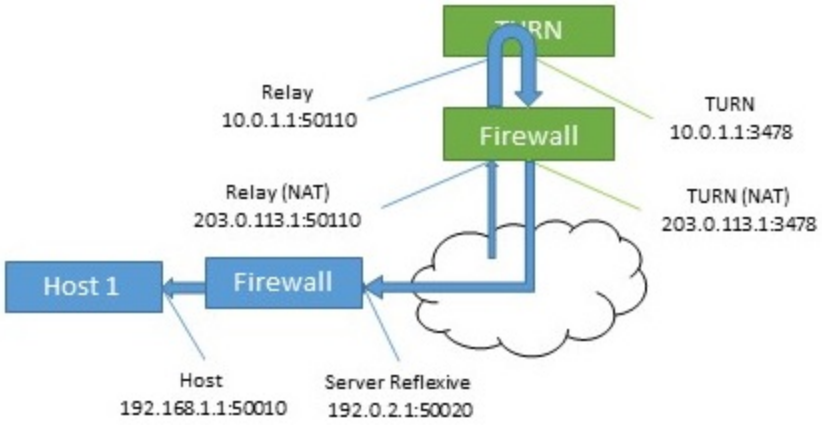

This relay candidate has a second use. It can also be used by the host to send packets to the far end. This occurs when there is no other path possible. Note that these packets come from the TURN server itself, so will only get their NAT address when rewritten by the firewall.

Figure 32: Host sending packets to the far end

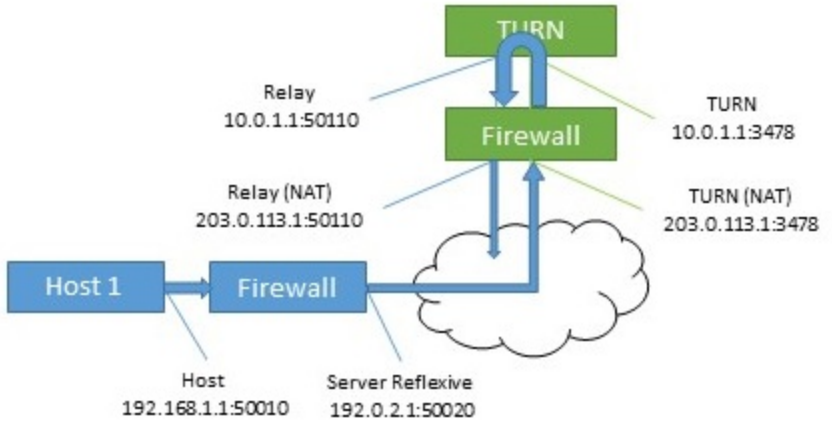

G.2 Checking connectivity

Once candidates are known then connectivity checks are undertaken. Each host tries to contact the far end host, server reflexive and relay addresses directly. It then also uses its relay to attempt connections to the same far end candidates.

Table 23: Candidates for two hosts (using same TURN server)

Host	Type	Address:port
1	Host	192.168.1.1:50010
1	Server Reflexive	192.0.2.1:50020
1	Relay	203.0.113.1:50110
2	Host	172.16.1.1:50100
2	Server Reflexive	198.51.100.1:50040
2	Relay	203.0.113.1:50510

Table 24: Candidate pairs formed by host 1

Source	Destination Type	Destination address
Host (192.168.1.1:50010)	Host	172.16.1.1:50100
Host (192.168.1.1:50010)	Server Reflexive	198.51.100.1:50040
Host (192.168.1.1:50010)	Relay	203.0.113.1:50510
Relay (10.0.1.1:50110)	Host	172.16.1.1:50100

Source	Destination Type	Destination address
Relay (10.0.1.1:50110)	Server Reflexive	198.51.100.1:50040
Relay (10.0.1.1:50110)	Relay	203.0.113.1:50510

Typically, the relay addresses are only required when the hosts have limited network access. For example, a user in a coffee shop or hotel may not be able to access any higher numbered ports.

When both hosts have restricted access then a path that involves both relay candidates can be formed. In this case, the traffic flows out of one relay candidate and into the other before being forwarded on to the far end.

Figure 33: Host to host media path using relay to relay path (no NAT)

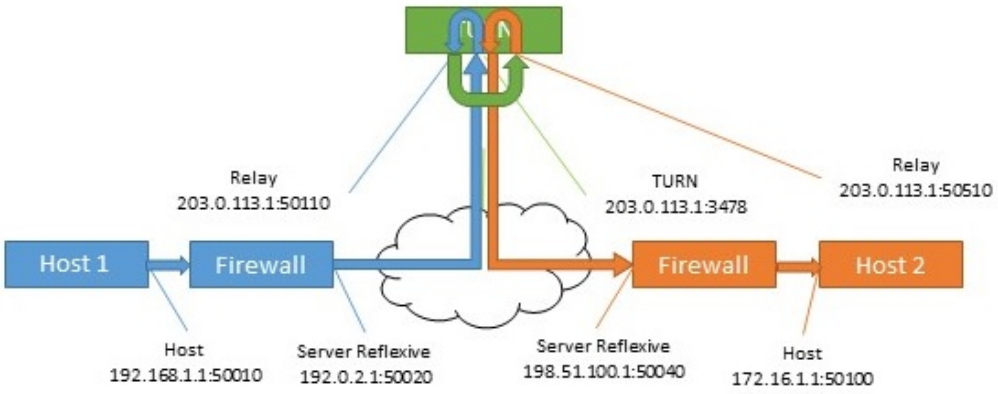

G.3 NAT in front of the TURN server

When NAT is present in front of the TURN server, the flow becomes more complicated. The relay candidates are expecting to receive traffic from one of the other hosts candidates. If the packets are sent from the TURN server's interface, and are not rewritten by the firewall, then they will appear to be coming from an unknown address. This prevents a succesful connectivity check and in cases where the other paths are not available, there are no routes for media to take.

Figure 34: Host to host media path using relay to relay path (with NAT)

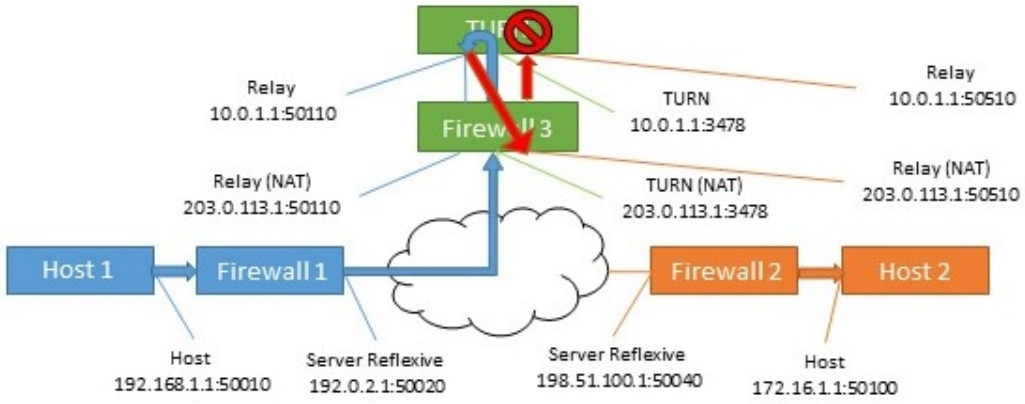

Table 25: Host to host media path using relay to relay path (with NAT)

Source address (in packets)	Destination	Action at destination
192.168.1.1:50010	203.0.113.1:3478 via Firewall	Firewall 1 rewrites source address
192.0.2.1:50020	203.0.113.1:3478	Firewall 3 rewrites destination address and forwards to the TURN server
192.0.2.1:50020	10.0.1.1:3478	TURN serevr internally maps this to the relay address for this source, and sends to far end's relay.
10.0.1.1:50110	203.0.113.1:50510 via Firewall	Firewall 3 rewrites destination address
10.0.1.1:50110	10.0.1.1:50510	TURN server sees unexpected source address and drops traffic.

The solution for this is known as hairpin NAT, loopback NAT or NAT reflection. In this the source address of the traffic is rewritten as well as the destination. The source address is then the address of the firewall, which means it matches one of the candidates.

Table 26: Host to host media path using relay to relay path (with hairpin NAT)

Source address (in packets)	Destination	Action at destination
192.168.1.1:50010	203.0.113.1:3478 via Firewall	Firewall 1 rewrites source address
192.0.2.1:50020	203.0.113.1:3478	Firewall 3 rewrites destination address and forwards to the TURN server.

Source address (in packets)	Destination	Action at destination
192.0.2.1:50020	10.0.1.1:3478	TURN server internally maps this to the relay address for this source, and sends to far end's relay.
10.0.1.1:50110	203.0.113.1:50510 via Firewall	Firewall 3 rewrites both source and destination addresses.
203.0.113.1:50110	10.0.1.1:50510	TURN server internally maps traffic from relay to assigned host.
10.0.1.1:3478	198.51.100.1:50040 via Firewall	Firewall 3 rewrites source address.
203.0.113.1:3478	198.51.100.1:50040	Firewall 2 rewrites destination address.
203.0.113.1:3478	172.16.1.1:50100	Arrives at final destination.

For details on how to enable this functionality, refer to your firewall documentation.

G.4 TURN server, NAT and the Cisco Meeting App

The effect of NAT on external Cisco Meeting App users needs to be considered in deployments where one Meeting Server is configured as a Core server with an internal interface, while another Meeting Server is configured as an Edge server set up on with two interfaces (internal and external). For Cisco Meeting App users working remotely, the Cisco Meeting App may be unable to see any ephemeral UDP ports.

In this case there is no server reflexive candidate for the Call Bridge, since the address seen by the TURN server is the same as the host candidate.

Figure 35: Split Meeting Server deployment with external Cisco Meeting App users (no NAT)

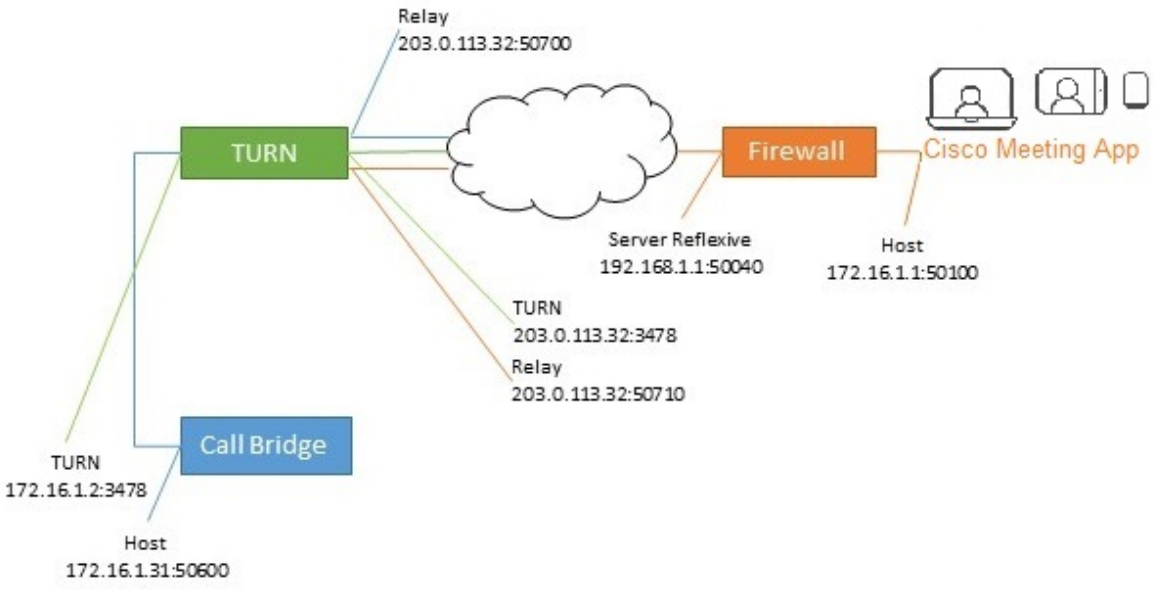

Since the Call Bridge running on the Core server is only on the internal network it has no route to the Cisco Meeting App's host address, server reflexive or the relay address. Likewise the Cisco Meeting App cannot see the Call Bridge's host, or its relay address.

However, the relay ports can see each other, and therefore a path for media can be established.

Figure 36: Relay ports establishing the media path

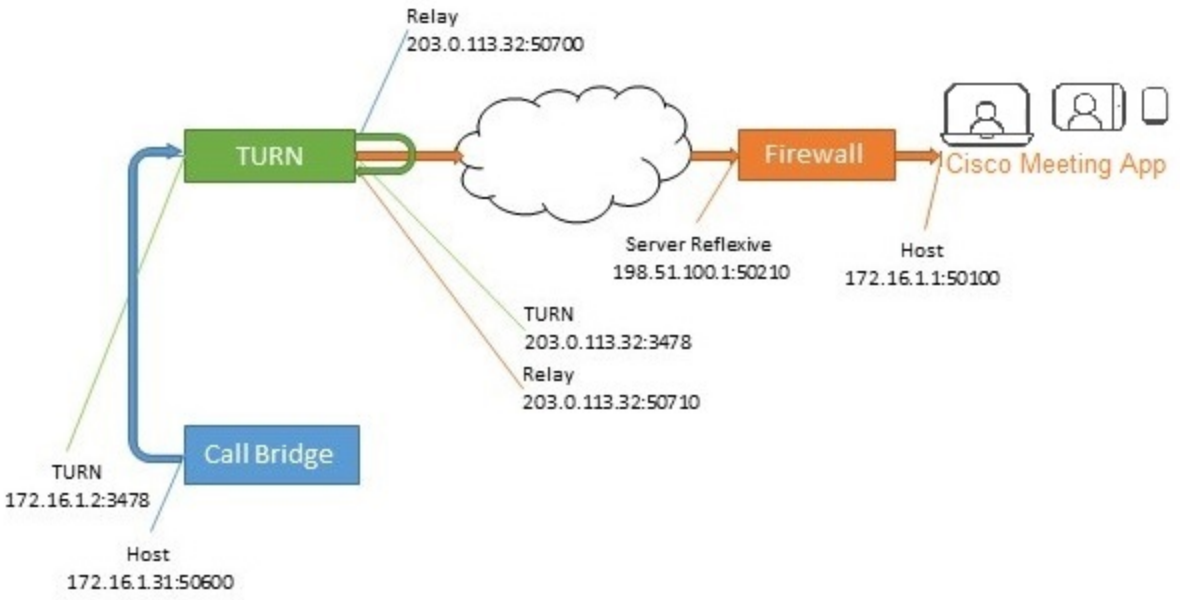

As in the general case, when the TURN server is behind a NAT this picture is further complicated.

Figure 37: Split Meeting Server deployment with external Cisco Meeting App users (with NAT)

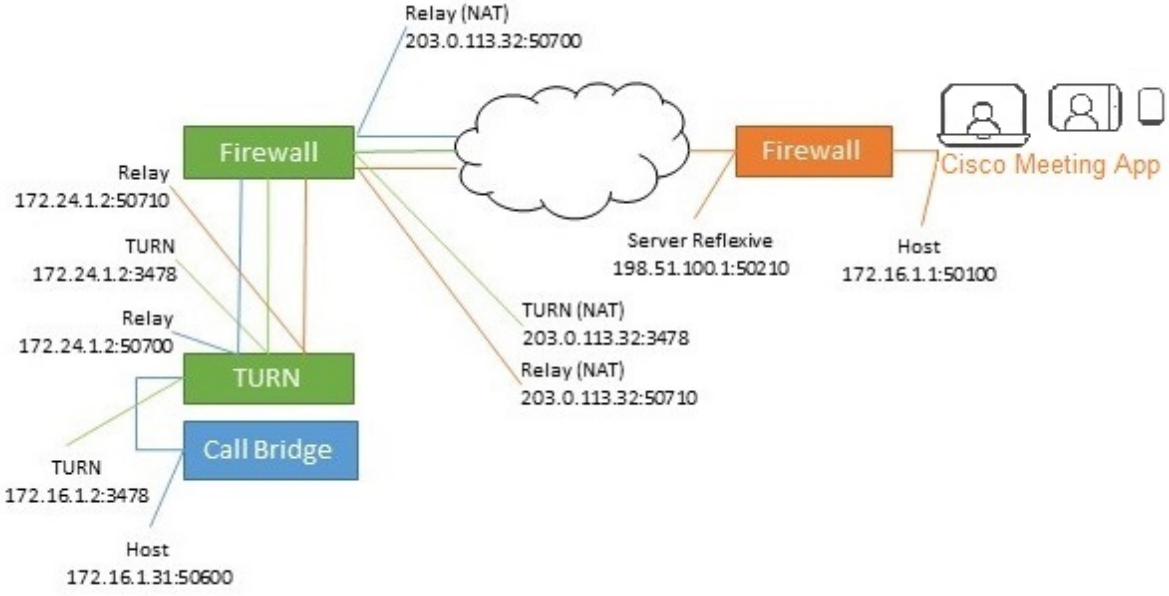

The solution for this is identical to the general case. The source address of traffic needs to be rewritten by the firewall so that it appears as coming from the correct address.

Figure 38: Relay ports establishing the media path

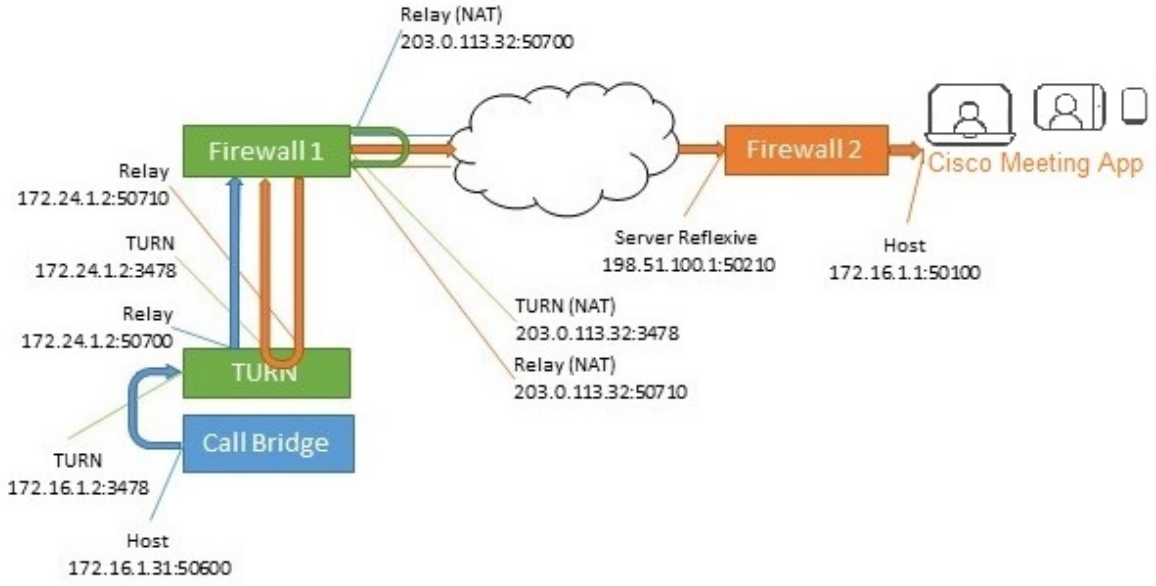

Table 27: Host to host media path using relay to relay path (with hairpin NAT)

Source address (in packets)	Destination	Action at destination
172.16.1.31:50600	172.16.1.2:3478	TURN internally maps this to the relay address for this source, and sends to the far end's relay.
172.24.1.2:50700	203.0.113.32:50710 via Firewall	Firewall 1 rewrites both source and destination addresses.
203.0.113.32:50700	172.24.1.2:50710	TURN server internally maps traffic from relay to assigned host.
172.24.1.2:3478	198.51.100.1:50510 via Firewall	Firewall 1 rewrites source address.
203.0.113.32:3478	198.51.100.1:50510	Firewall 2 rewrites destination address.
203.0.113.32:3478	172.16.1.1:50100	Arrives at final destination.

Appendix H Using a standby Meeting Server

The instructions in this appendix apply to:

- virtualized deployments (including the Cisco Meeting Server 1000
- the Acano X-series servers.

H.1 Backing up the currently used configuration

1. Establish an SSH connection to the currently used Meeting Server using an SSH utility such as OpenSSH or PuTTY.
2. Issue the command:

   backup snapshot <name>

   This backup includes IP addresses, passwords and certificates into a file called <name>.bak. We recommend using a name in the format servername_date (for example, test_server_ 2014_09_04).

   A successful backup creation returns:

   cms> backup snapshot test_server_2014_09_04.bak ready for download
3. Download the backup file using an SFTP client (e.g. WinSCP).

Note: We recommend backing up your Meeting Server regularly, e.g. once a day and that you store copies of the backup externally to the Meeting Server and the standby server.

H.2 Transferring a backup to the standby server

We recommend that you keep the standby sever running at all times.

1. Copy all the certificates and the cms.lic file from the standby server in case they differ from the original server that the backup was created on. Store them somewhere safe.
2. Establish an SFTP connection with the standby server.
3. Upload the previously saved backup file on to the standby server.
4. Issue the MMP backup list command to confirm that the backup file was successfully uploaded. This should return something similar to:

   cms> backup list test_server_2014_09_
5. Enter the following command and confirm to restore from the backup file:

   backup rollback <name>

This overwrites the existing configuration and reboots the Meeting Server. Therefore a warning message is displayed. The confirmation is case sensitive and you must press upper case **Y**, otherwise the operation will be aborted.

Note: It is not possible to create a backup from one type of deployment (virtualized or Acano X-series server) and roll it back on the other type.

A successful operation returns:

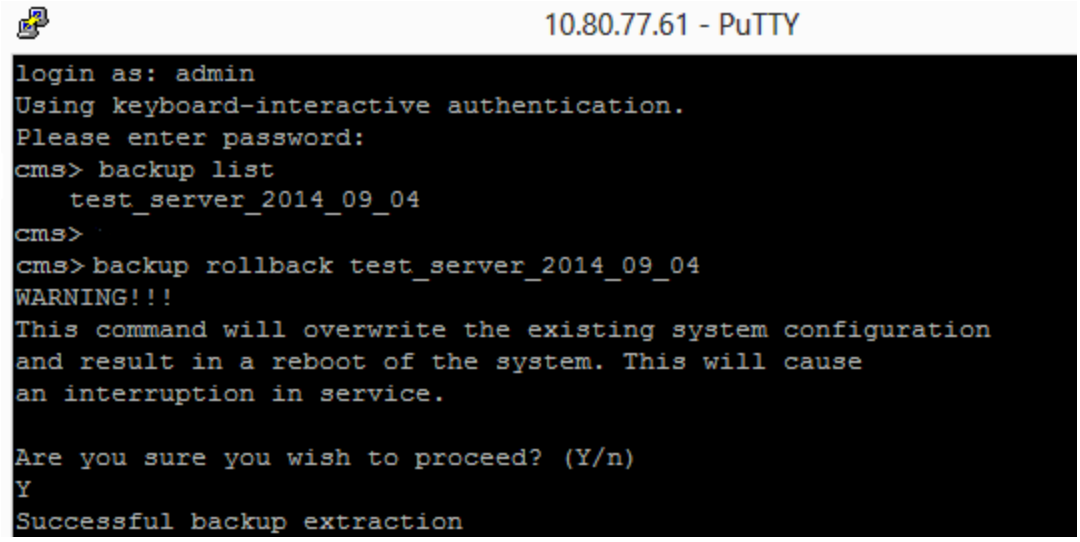

When you restore from the backup, everything is overwritten including the IP address, certificates and the cms.lic file. Therefore if you are restoring onto a different server from the one that the backup was made on, you must manually copy the original cms.lic file and any certificates that are not valid on the new server. Note that the cms.lic file is tied to the MAC address of the server; therefore after the backup has been restored to the new server, the license from one server will be invalid on another one.

6. Establish an SFTP connection with the standby server
7. Upload the previously saved original cms.lic file back on to this server
8. If necessary:
   a. Put back any certificates and private keys (if the restored versions are not valid on the standby server).
   b. Assign these certificates to their corresponding services using the following commands:

```
callbridge certs nameofkey nameofcertificate
webbridge certs nameofkey nameofcertificate
webadmin certs nameofkey nameofcertificate
xmpp certs nameofkey nameofcertificate
webbridge trust nameofcallbridgecertificate
```

c. Restart any service for which you changed the certificate

```
xmpp restart
callbridge restart
webbridge restart
webadmin restart
```

After the new server has fully booted up, it will be fully operational, and will take over the services of the original server.

H.3 Time for swapping servers

If the standby server is kept powered on, typical restore times for virtualized Meeting Servers is 2-4 minutes (and forAcano X-Series Servers this is 6-8 minutes) to restore the configuration, copy the cms.lic file and restart the XMPP server. If certificate files also need to be restored, additional time may be required.

Appendix I Web Admin Interface – Configuration menu options

The **Configuration** tab on the Call Bridge's Web Admin interface allows you to configure the following options:

- General
- Active Directory
- Call settings
- Outbound calls and Incoming calls
- Outbound calls and Incoming calls
- CDR settings
- Spaces
- Cluster
- CMA user settings

I.1 General

Use the **Configuration > General** page to set up and configure:

- **XMPP server settings**. Use these fields to configure the settings through which the Call Bridge communicates with the XMPP server. See Web Admin interface settings for XMPP. Use MMP commands to configure the XMPP server itself. See Configuring the XMPP server. Note that you only need to configure and enable the XMPP server if you are using the Recorder or Streamer components or any Cisco Meeting App (including the WebRTC Client).
- **TURN server settings**. Use these settings to allow the Call Bridge and external clients to access the TURN server. See Web Admin interface settings for the TURN server. Use MMP commands to configure the TURN server itself. See Configuring the TURN server.
- **Lync Edge settings**. Use these settings if you are integrating your Call Bridge with Lync Edge. See Configuration on Meeting Server to use Lync Edge.
- **Web bridge settings**. Use these settings if you are using WebRTC video calls and meetings. These settings allow the Call Bridge to communicate with the Web Bridge server. See Web Admin interface settings for the Web Bridge .
- **IVR**. Use these settings if you are using an Interactive Voice Response (IVR) to manually route to pre-configured calls, so callers are greeted by a prerecorded voice message

inviting them to enter the ID number of the call or space that they want to join. See IVR configuration.

- **External access**.Use these settings to enter the URI for the Web bridge that WebRTC clients will use to access it. This field must be filled in manually on every Call Bridge in the cluster for Meeting App clients to generate WebRTC URLs.
  Enter the IVR telephone number to add an extra option in the invite list, and to provide a call in "Phone" line to the meeting email and invitation templates.

I.2 Active Directory

If you want users to use Cisco Meeting Apps to connect to the Meeting Server, then you must have an LDAP server. The Meeting Server imports the User accounts from the LDAP server.

Note: You can use OpenLDAP and Oracle Internet Directory (LDAP version 3), however, this needs to be configured via the API–it cannot be configured through the Web Admin interface.

Use the **Configuration > Active Directory** page to set up the Meeting Server to work with Active Directory. See LDAP configuration.

I.3 Call settings

Use the **Configuration > Call settings** page to:

- Allow media encryption for SIP calls (including Lync).
- Specify whether participant label overlays are shown on SIP calls.
- Specify the preferred size (in milliseconds) for outgoing audio packets; 10ms, 20ms, or 40ms.
- Enable TIP support. (You need to enable TIP support if you use endpoints such as the Cisco CTS range.)
- Allow presentation video channel operations–if this is set to **prohibited** then no content channel video or BFCP capability will be advertised to the far end.
- If presentation video channel operations are allowed for SIP calls, this setting determines the Call Bridge's BFCP behavior, one of:
  - **server role only**–this is the normal option for a conferencing device, and is intended for use with BFCP client mode devices (for instance, SIP endpoints).
    or
  - **server and client role**–this option allows the Call Bridge to operate in either BFCP client or BFCP server mode in calls with remote devices.

This setting allows improved presentation video sharing with a remote conference-hosting device.

- Set the value for the Resource-Priority header field in outgoing SIP calls. This setting tells the Meeting Server how much priority you will allow the bandwidth to allocate for presenting. This depends on the bandwidth capability of the network environment and other factors such as if there are any immersive systems that push HD, for example.
- Enable and disable UDP signaling for SIP. Set to one of:
  - **disabled|enabled**: disable if you use SIP over TCP, or require that all of your network traffic is encrypted.
  - **enabled, single address** mode corresponds to the SIP over UDP behavior in versions prior to 2.2 and is the default.
  - **enabled, multi address** if the Call Bridge is configured to listen on more than one interface.
- Enable Lync presence support. This setting determines whether or not this Call Bridge should supply information on destination URIs it serves to Lync presence subscribers.
- Leave the Lync packet pacing mode set to **default**. Do not change the setting to **delay** unless instructed to do so by Cisco Support.

Note: For more information on each field, you can use the hover-over text that displays for each individual field, or see Dial plan configuration – SIP endpoints.

The **Call settings** page also allows you to change the bandwidth settings for SIP, Cisco Meeting Server (CMA), Server reflexive, Relay, VPN, and Lync content. The settings are measured in bits-per-second, for example, 2000000 is 2Mbps. We dedicate at least 64kbps for audio, and recommend 2Mbps for a 720p30 call, or around 3.5Mbps for a 1080p30 call. More bandwidth would be required for 60fps.

You may need to change some of the bandwidth settings if you allow SIP media encryption, or enable TIP support, for example. In the case of 3 screen TIP calls, the bandwidth numbers seen on the **Call settings** page get automatically tripled, so you do not need to manually set them to 6Mbps for example. However, we would normally recommend (3x) 4Mbps for most CTS calls.

I.4 Outbound calls and Incoming calls

Use the **Configuration > Outbound calls / Incoming calls** pages to determine how the Meeting Server handles each call.

The **Outbound calls** page controls how outbound calls are handled; the **Incoming calls** page determines whether incoming calls are rejected, or matched and forwarded. If they are matched and forwarded, then information about how to forward them is required. The **Incoming calls**

page has two tables–one to configure matching/rejection and the other to configure forwarding behavior.

For more information on completing these fields, see Web Admin Interface configuration pages that handle calls.

I.5 CDR settings

Use the **Configuration > CDR settings** page to enter the URI of the CDR receivers.

The Meeting Server generates Call Detail Records (CDRs) internally for key call-related events, such as a new SIP connection arriving at the server, or a call being activated or deactivated. It can be configured to send these CDRs to a remote system to be collected and analyzed. You can not store records on a long-term basis on the Meeting Server, or browse CDRs on the Meeting Server.

For more information on completing these fields, see Call Detail Record support and the Call Details Record Guide.

You can also use the API to configure the Meeting Server with the URI of the CDR receivers. See the API Reference guide.

I.6 Spaces

Use the **Configuration > Spaces** page to create a space on the Meeting Server to dial into. This allows, for example, endpoints and Meeting App to dial in.

Add a space with:

- **Name** for example. `Call 001`
- **URI** for example. `88001`

On this page you can also optionally specify Secondary URI user part, Call ID, Passcode, and Default Layout.

You can also use the API to create spaces. See the API Reference guide.

I.7 Cluster

Note: The **Configuration > Cluster** page only appears in the Web Admin interface if all the databases are running as a cluster and all the Call Bridges have been connected to the database cluster.

Within your Meeting Server deployment, you can enable Call Bridge clustering which will allow multiple Call Bridges to operate as a single entity and scale beyond the capacity of any single Call Bridge.

You have a choice whether to setup the Call Bridges in the cluster to link peer-to-peer, or for calls to route via call control devices between the clustered Call Bridges.

For more information, see the Clustering Call Bridges section in the Cisco Meeting Server 2.3, Scalability and Resilience Deployment Guide.

I.8 CMA user settings

Use the **Configuration > CMA user settings** page to allow or not allow incoming calls to Cisco Meeting App users.

By default incoming calls to Cisco Meeting Apps are allowed, however this behavior can be changed so that incoming calls are not allowed to users of the Cisco Meeting App.

Cisco Legal Information

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco Trademark