



## CHAPTER 3

# Cisco IPICS UMS

---

A Cisco IPICS deployment includes one or more Cisco Unified Communications Server (UCS) components that host virtual instances of the Unified Media Service (UMS).

For information about UMS system requirements and capacity, see *Cisco IPICS Compatibility Matrix*.

For information about the ports and transport protocols that the UMS uses, see [Table 2-2 on page 2-24](#).

This chapter includes these topics:

- [UMS Overview, page 3-1](#)
- [When is a UMS Required?, page 3-2](#)
- [UMS Instances for Locations, page 3-2](#)
- [UMS Scaling, page 3-3](#)
- [UMS Resource Allocation, page 3-3](#)
- [Remote Users, page 3-4](#)
- [UMS Audio Mixing, page 3-4](#)

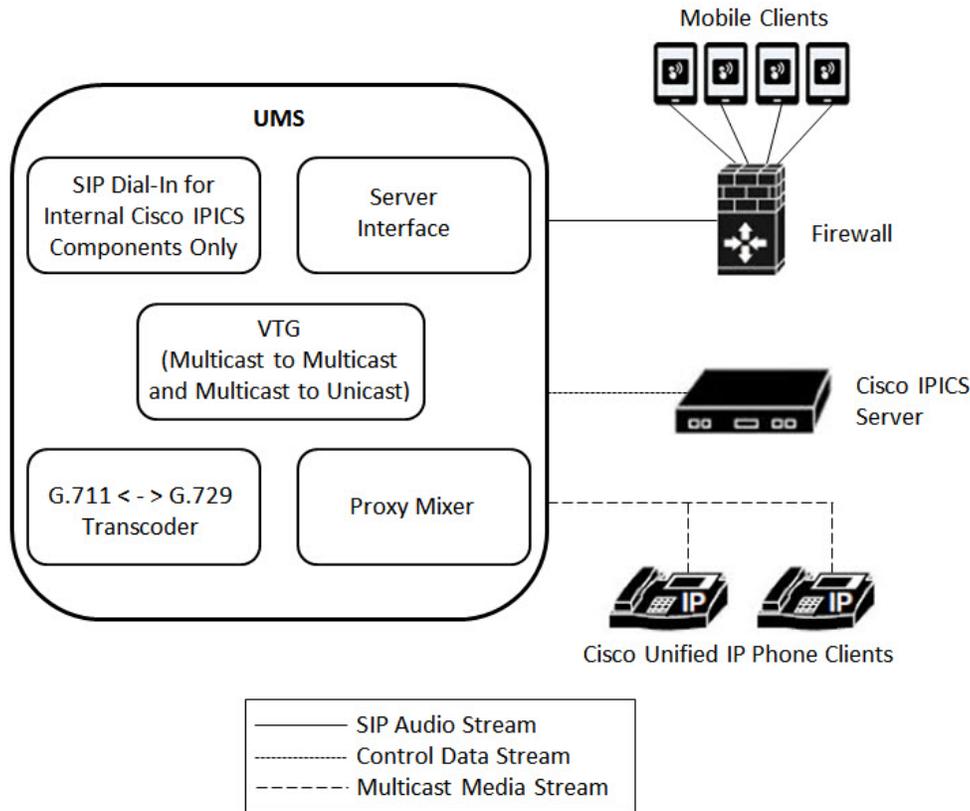
## UMS Overview

The UMS is a highly available, software-based media engine that performs several core functions in a Cisco IPICS deployment including:

- Media transcoding between G.711 and G.729 media streams
- Media stream mixing and floor control
- Talker ID for P25 and SIP based endpoints
- SIP termination
- SIP to multicast, SIP to SIP, multicast to SIP, and multicast to multicast media connectivity
- Floor control for mobile clients and Cisco Unified IP Phone clients

[Figure 3-1](#) illustrates the use of a UMS in a Cisco IPICS deployment.

Figure 3-1 UMS in a Cisco IPICS Deployment



## When is a UMS Required?

A UMS is required for any Cisco IPICS deployment in which unicast and multicast endpoints are joined or channels are combined.

Examples include:

- Using VTGs
- Using Cisco IPICS incidents
- Having dial-in users

Dial-in users joining VTGs or channels

## UMS Instances for Locations

Cisco IPICS relies upon the concept of resource *locations* to effectively manage network traffic. A location is a multicast enabled domain within a network. A domain can comprise multiple geographic locations. Some networks may be fully multicast enabled across the entire network while others may not route multicast traffic beyond a geographic boundary or a single subnet. Cisco IPICS considers each separate multicast domain to be a different location. Any resources (channels, radios, endpoints, and so on) that are not in the same multicast domain as the Cisco IPICS server are considered to be in the *Remote* location.

A UMS is assigned to a location when it is configured. All IPICS servers and components that are in the same multicast domain should be assigned to the same location. While a single IPICS server can administer UMS resources in multiple locations, multicast media is confined to its originating location. Media can move between locations via a Multicast–Unicast–Multicast (MUM) trunk configured on a T1 or E1 loopback or via a GRE tunnel. Only baseband audio data is carried across the MUM trunk; talker ID, control data and supplemental services are not distributed between locations.

The following guidelines apply to locations:

- A channel is associated to a location
- A VTG is a global resource and can be used to span a channel to another channel in another location
- Cisco Mobile Client users are always considered to be remote

## UMS Scaling

Each instance of the UMS supports up to 100 simultaneous audio streams. A Cisco IPICS deployment can be scaled to support many hundreds of users by adding UMS resources. The number of UMS resources that are required is based on sum of the streams. Cisco The IPICS server maintains a list of available UMSs and locations and distributes media streams on a round-robin basis.

Clients consume resources differently; the UMS sends one audio stream to a registered mobile client regardless of how many channels are displayed on the client. When deploying Cisco IPICS to support more than 1,000 simultaneous audio streams, server and media streams must be distributed across multiple physical servers. Additional virtual cores and memory may be required for optimized performance. Consult your Cisco representative for assistance with large scale installations.

## UMS Resource Allocation

UMS resources are dynamically allocated in the following situations:

- Activating or deactivating a VTG—When a VTG is activated by a Cisco IPICS user or triggered by a Cisco IPICS policy, one UMS resource is allocated for each channel in the VTG. The UMS resource is released when the user deactivates the VTG
- Activating or changing an incident—An incident is a special case of a VTG in which media other than voice can be associated with an event. Each channel or VTG in an incident consumes one UMS resource.
- Authenticating a mobile client—Mobile clients connect to Cisco IPICS via a SIP session between the client and the UMS, which consumes one UMS resource.
- Authenticating a Cisco video surveillance IP camera—Cisco video surveillance cameras that run the SIP client create a SIP session between the camera and the UMS, which consumes one UMS resource.
- A dial-in user joining a channel or VTG—Each Channel or VTG that the Cisco IPICS dial engine accesses consumes one UMS resource. (This UMS resource can service multiple users simultaneously.)

## Remote Users

A Cisco IPICS mobile client user is, by definition, a remote user. A remote user accesses Cisco IPICS through a SIP-based (unicast) connection and obtains a media connection to the Cisco IPICS server. When the user joins a channel or VTG, Cisco IPICS configures a resource on the UMS to enable a multicast connection from the UMS to the user.

This multicast connection is made one time for a channel or VTG, regardless of the number of users who select the channel or VTG. When the last remote user disconnects from the channel or VTG, the resource is released in the UMS and becomes available.

When a remote user obtains a media connection on the Cisco IPICS server, the UMS sends and receives multicast streams as follows:

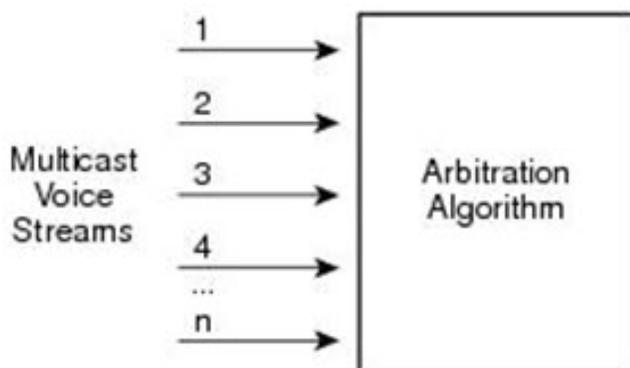
1. After the user selects a resource, Cisco IPICS allocates a UMS resource for the user and allocates a multicast address from the multicast pool. Cisco IPICS then performs an IGMP join operation on the multicast address so that when additional users select the same resource, the Cisco IPICS server can continue to use same the multicast address.
2. When the user begins to talk, Cisco IPICS transmits the audio to the multicast address of the selected resources.
3. When the UMS receives the multicast packets, it forwards the packets to the multicast address that has been allocated from the multicast pool. Cisco IPICS receives that multicast audio stream and forwards it as a unicast stream to all remote users who have selected that resource.

## UMS Audio Mixing

The UMS uses a communications system called *Hoot 'n' Holler* (or *hootie*) in which the three most recent talkers are mixed into one multicast output stream. The UMS provides “always on” multi-user conferences. The UMS version of Hoot 'n' Holler is based on the hardware implementation of audio mixing that available in Cisco devices that run UN-enabled versions of Cisco IOS.

In the Cisco Hoot 'n' Holler over IP implementation, all participants in a VTG can speak simultaneously. However, when voice packets from various sources arrive at the UMS, the arbitration algorithm selects only the three most active voice streams and mixes them. If other voice streams are present, the UMS drops the longest talker by using a round-robin arbitration algorithm. See [Figure 3-2](#).

**Figure 3-2** *Mixing Voice Streams*



[Table 3-1](#) shows an example of how mixing works in a VTG that has four active users on a channel.

**Table 3-1**      **Mixing Example**

<b>Event</b>	<b>Remarks</b>
User A starts speaking.	1 user speaking.
User B and User C join User A.	3 users speaking simultaneously. Cisco arbitration engine at the UMS receives 3 voice streams.
User D starts speaking while the other 3 users continue speaking.	Cisco arbitration engine at the UMS receives 4 voice streams. The algorithm can present up to 3 voice streams. It drops the voice stream from the longest talker, User A, and adds User D to the streams that it presents. The voice streams are now from User B, User C, and User D.
After 2 seconds, all 4 users are still speaking.	The current longest talker, User B, is dropped, and User A is added. Voice streams are now User C, User D, and User A.
After 2 seconds, all 4 users are still speaking.	The current longest talker, User C, is dropped, and User A is added. Voice streams are now User D, User A, and User B.
All users continue speaking.	The round-robin process of dropping the current longest talker and adding the other user every 2 seconds continues.

