

Microsoft Teams

- Overview
 - Microsoft Teams recording features
 - Voice, video, and screen/application window share recording:
 - Instant message, attachment capture and archiving:
 - Version support
- Deploying Microsoft Teams recording
 - Server sizing
 - Recording Server on Azure virtual machine for voice, video, and screen/application share recording
 - Recording Server for IM and attachment recording
 - Recording Director
 - Media Recorder and SQL Server
 - Load-balancing and Failover
 - Preparation
 - Installation
 - Configuration
 - Microsoft Teams metadata

Overview

Microsoft Teams recording features

- Voice, video, screen/application window share, instant messages, and attachment recording and archiving
- Integration with the official Teams Compliance Recording API and Graph API
- Support for load balancing and failover
- Support for Azure cloud and hybrid on-prem deployments
- Supports all type of Teams endpoints and devices

Voice, video, and screen/application window share recording:

- Bot based integration, where the Teams platform is able to automatically invite the recorder bot into P2P calls or meetings for the configured users
- Fail-close configuration option (in case of recorder failure, the recorded user cannot join the call)
- Supports always-on, selective, and on-demand recording
- Supports all call scenarios where the recorded user is a participant
- Built-in announcement and notification (provided by the Teams platform)

Instant message, attachment capture and archiving:

- Event/webhook, DLP API based integration, where the Teams platform sends all updates once the recorder subscribes
- Support for on-prem, cloud and hybrid deployments
- Supports always-on recording mode only
- Supports all instant message scenarios, teams, channels, P2P, and meeting chats
- Supports all formatting options, emojis, giphys, stickers, and other apps
- **Limitation: no selective capturing/archiving support, only tenant-wide capturing**
- **Limitation: chat sessions initiated from external domains won't be captures**

Version support

Switch Name & Model	Microsoft Teams
Supported Microsoft Teams Versions for voice, video, and screen/application window share recording	<p>Voice, video, and screen/application window share recording (Compliance Recording) is available to all:</p> <ul style="list-style-type: none"> • Microsoft 365 A3/A5/E3/E5/Business Premium and • Office 365 A3/A5/E3/E5 users • with no additional consumption charge. <p>For more information, see https://docs.microsoft.com/en-us/microsoftteams/teams-recording-policy</p>

Supported Microsoft Teams Versions for instant message, attachment capture and archiving	<p>Instant message, attachment capture and archiving require one of the following licenses:</p> <ul style="list-style-type: none"> • Office 365 A5/E5 • Microsoft 365 A5/E5 • Microsoft 365 Information Protection and Governance • Office 365 Advanced Compliance <p>For more information, see https://docs.microsoft.com/en-us/graph/teams-licenses</p>
Supported Endpoint / DeviceTypes	All

If you are on a different version, contact your Microsoft representative for more information.

Deploying Microsoft Teams recording

The following section contains all the necessary steps for setting up a Microsoft Teams recording infrastructure.

Server sizing

Recording Server on Azure virtual machine for voice, video, and screen/application share recording

According to the Microsoft requirements, the Microsoft Teams Bot service must run on an Azure Virtual Machine. Requirements for the Recording Server role which will host the Verba Microsoft Teams Bot service and the Unified Call Recorder service are:

- Azure Compute Unit (ACU) should be 200 or higher and 1:1 ratio for vCPU:Core
<https://docs.microsoft.com/en-us/azure/virtual-machines/acu>
- We recommend using the Dv2 series virtual machines, Standard_D3_v2 or above
<https://docs.microsoft.com/en-us/azure/virtual-machines/dv2-dsv2-series#dv2-series>
- ILPIP (Instance Level Public IP Address) for the Recording Servers
Private IP addresses are not recommended by Microsoft, due to potential performance limitation caused by the Azure NAT, although will work
- For resilient and/or high volume configurations, multiple virtual machines (running the Recording Server role) have to be deployed. In order to distribute the load across multiple Verba Microsoft Teams Bot services, an **Azure Application Gateway** has to be deployed in front of the VMs.
<https://docs.microsoft.com/en-us/azure/application-gateway/overview>

Due to performance limitations in the Microsoft Local Media SDK, the Recording Server sizing for Microsoft Teams recording is different from other integrations. The following table summarizes the server sizing for Microsoft Teams recording:

Modality	per vCPU	
	w 60% avg. load	w 80% avg. load
Voice	50	65
Voice and Screen & Application Share Up to x1 1080p stream / call	30	40
Voice and Video Up to x4 360p video streams / call	15	18


The values indicate the number of maximum simultaneous calls for the specified modalities per vCPU.

Testing was done with mixed audio on Azure D3v2 (4 cores), D4v2 (8 cores), and D5v2 (16 cores) virtual machines.

Recording Server for IM and attachment recording

The IM recording architecture includes two server roles: Recording Director and Media Recorder. These roles have different sizing numbers and different factors have to be taken into account.

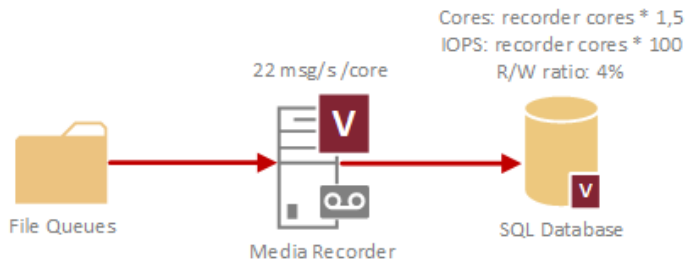
Recording Director



The Recording Director component **has to be sized based on the real-time incoming load**. The minimum CPU requirement is 4 CPU cores. It can process **1500 messages every second with a single CPU core**, and 6000 messages every second with 4 cores.

In the case of higher incoming loads, the network bandwidth also has to be considered. 1000 messages per second incoming load generate 48 Mbps traffic (or 6 MB/s) between the Teams side and the Recording Director, and 42 Mbps traffic (or 5.2 MB/s) between the Recording Director and the file queue storage.

Media Recorder and SQL Server



The Media Recorder component does not have to be sized for real-time processing, since the recorded data is stored already in the file queue storage. Instead, the **Media Recorder can be sized based on the overall message count a day**. If there are more incoming messages than the real-time processing capacity of the Media Recorder(s), then the messages will be inserted into the database later, so they will be also available for search and replay through the web interface later. However, sufficient processing capacity should be provided so it can process the daily message load at least within 16 hours.

The minimum CPU requirement is 4 CPU cores. It can process **22 messages every second with a single CPU core**. In the case of multiple Media Recorder servers, all servers have to have the same number of cores.

The Recorder Director and the Media Recorder components can be co-located on the same server. In this case, the resources will be shared between them.

The **SQL Server** has to be sized based on the fully utilized CPU cores of the Media Recorder server(s). The SQL Server needs to have **one and a half times more CPU cores than the Media Recorder server(s)**. On the SQL Server physical disk, **every fully utilized Media Recorder CPU cores generate 100 IOPS**.

The following table shows the expected incoming message rates at different user numbers:

	1K Users	10K Users	100K Users
Average during the day*	1.6 msg/s	16.6 msg/s	166.6 msg/s
Low message rate**	2.7 msg/s	27.7 msg/s	277.7 msg/s
Medium message rate**	4.1 msg/s	41.6 msg/s	416.6 msg/s
High message rate**	6.9 msg/s	69.4 msg/s	694.4 msg/s

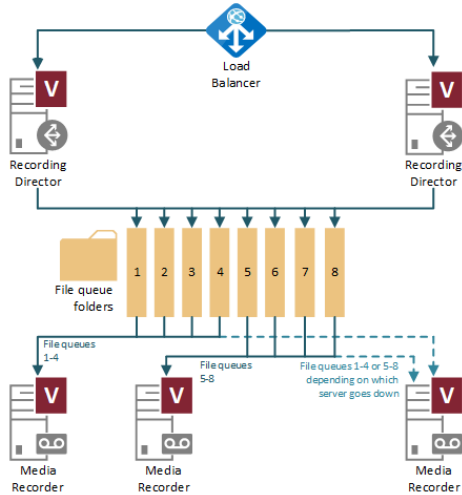
*Based on Slack usage statistics

**Based on Cisco IM/P sizing

For requirements for other components and server roles, see [Server sizing and requirements](#)

Load-balancing and Failover

Large deployments may require multiple VMs and other Azure components. In the case of IM and attachment recording, multiple Media Recorders can be deployed behind the Recording Director(s). If multiple Recording Directors are required, then a load-balancer has to be placed in front of them.



For the failover and load-balancing options for voice, video, and screen/application window share recording, see:

[Microsoft Teams Recording Failover and Load-balancing Design](#)

Preparation

The Microsoft Teams integration requires additional prerequisites and configuration on Azure and O365, see [Configuring Microsoft Teams Recording](#) for more information.

Make sure that all the required prerequisites are installed on each server prior to the installation.

- [Prerequisites](#)
- [Installing the required prerequisites](#)

For guidance on configuring the necessary firewall port, visit [Firewall configuration for Microsoft Teams recording deployments](#)

Installation

The following articles contain all the step for installing the various server roles:

- [Installing a Verba Single Server solution](#)
- [Installing a Verba Media Repository](#)
- [Installing a Verba Recording Server](#)

Configuration

For the configuration guide for voice / video / screen share recording, see [Configuring Microsoft Teams Recording](#).

For chat and attachment recording, see [Microsoft Teams Chat Recording](#).

Microsoft Teams metadata

The system captures the following metadata specific to Microsoft Teams voice/video/screen recordings.

Metadata Field	Description	Template	Available
Start Date	Start date of the conversation	Standard	Yes
Start Time	Start time on the conversation	Standard	Yes
End Date	End date of the conversation	Standard	Yes
End Time	End time of the conversation	Standard	Yes
Duration	Length of the conversation	Standard	Yes
User	Name of the recorded user	Standard	Yes
From	Phone number, Button name, User name	Standard	Yes
From Info	User / contact name	Standard	Yes
To	Phone number, Button name, User name	Standard	Yes
To Info	User / contact name	Standard	Yes
Direction	Direction of the call from the system perspective, requires configuring internal number/domain patterns	Standard	Yes
Direction (User)	Direction of the call from the recorded user perspective	Standard	Yes
From (Verba)	Name of the Verba user associated with the calling party	Standard	Yes
To (Verba)	Name of the Verba user associated with the called party	Standard	Yes
Location	Hostname of the recording server	Standard	Yes
End Cause	Normal, Hold, Transfer, Conference, Device Change, From Terminated, To Terminated	Standard	Yes
Audio Codec	Audio codec of the recorded streams	Standard	Yes
Video codec	Video codec of the recorded streams	Standard	Yes
Platform Call ID	Unique conversation identifier received from the recorded platform	Standard	Yes
Silence Ratio	Ratio of silence in the conversation	Standard	Yes
Talkover Ratio	Talkover ratio of the conversation	Standard	Yes
Longest Silence	Length of the longest silence present in the conversation	Standard	Yes
User ID / Agent ID	Azure AD Object ID for the recorded user	Standard	Yes
From Device	Device ID of the calling party	Standard	No
To Device	Device ID of the called party	Standard	No
Dialed Number	Original dialed number	Standard	No
From IP	IP address of the recording bot	Standard	Yes
To IP	IP address of the recording bot	Standard	Yes
From Proxy IP	IP address of the proxy server associated with the caller party	Standard	No
To Proxy IP	IP address of the proxy server associated with the called party	Standard	No
Source Platform	Microsoft Teams	Standard	Yes
Conversation Type	Voice, Video, Screen Share	Standard	Yes

Forward Reason	Forward reason for the conversation (e.g. forwarded, transferred, team call, delegated, etc.)	Standard	No
Recording failed	Indicates if the recording has failed and the metadata was imported during CDR reconciliation	Standard	No
Media Length	Length of the media file related to the conversation in hhh:mm:ss format	Standard	Yes
Media Error	Shows the media processing errors during recording	Standard	Yes
Voice Quality	Overall voice quality check score for the conversation	Standard	Yes
Record Type	Standard	Standard	Yes
2N Source	In case of duplicate (2N) recording, records are marked as primary or secondary	Standard	No