

DocsChatV2 Web Application – Product User Guide

Table of Contents

- DocsChatV2 Web Application – Product User Guide 1
- Table of Contents..... 1
- 1. End User Guide 1
 - 1.1 Introduction 1
 - 1.2 Getting Started 1
 - 1.3 Understanding the DocsChatV2 User Interface 1
 - 1.4 Managing Conversations 7
 - 1.5 Attaching and Managing Files 8
 - 1.6 Chatting with the AI 9
 - 1.7 Notifications and Messages..... 10
 - 1.8 Best Practices 10
- 2. Administrator Guide..... 11
 - 2.1 Administrator Overview..... 11
 - 2.2 Accessing the Admin Panel 11
 - 2.3 Managing AI Service Settings 11
 - 2.4 File and Conversation Limits..... 11
 - 2.5 Configuration Propagation..... 12
 - 2.6 Best Practices for Administrators..... 12
- 3. Frequently Asked Questions (FAQ) 12
- 4. Glossary 13

1. End User Guide

1.1 Introduction

DocsChatV2 is an AI-powered chatbot that enables users to ask questions about documents stored in their OneDrive account. The application allows users to attach files and receive answers grounded strictly in the context of those files.

The application is designed to provide a continuous conversational experience where users can return to previous conversations, retain attached documents, and continue discussions without losing context.

1.2 Getting Started

1.2.1 Accessing the Application

1. Open the DocsChatV2 Web Application using the Web Application URL provided by your organization.
2. Sign in using your Microsoft Entra ID account.
3. After successful authentication, the main DocsChatV2 workspace is displayed.

1.3 Understanding the DocsChatV2 User Interface

After signing in, users are presented with the main DocsChatV2 workspace. The interface is designed to create a smooth and continuous workflow where conversations, documents, and AI responses remain connected throughout the user journey.

The application interface is organized into three primary sections:

Section	Purpose
Left Sidebar	Manage and switch between conversations
Central Chat Area	View questions, AI responses, and interact with the chatbot
File Picker Panel	Browse, select, and manage OneDrive documents

Each section works together to help users maintain conversational continuity and efficiently interact with document-based AI responses.

1.3.1 Left Sidebar – Managing Conversations

The left side of the application contains the **Conversation Sidebar**, which acts as the navigation area for all user conversations.

The sidebar allows users to maintain continuity between conversations by preserving chat history, attached files, and conversation-specific context.

Features Available in the Left Sidebar

New Conversation Button

Located at the top of the sidebar, this option allows users to start a new conversation.

Selecting this option creates a fresh conversation workspace with no previously attached files or messages.

Conversation List

The sidebar displays all previously created conversations for the logged-in user.

Each conversation remains accessible until deleted by the user or removed due to configured conversation limits.

Conversation Switching

Clicking on any conversation in the sidebar automatically restores:

- Previous questions and AI responses
- Attached documents associated with the conversation
- Conversation-specific context

This allows users to continue conversations naturally without needing to repeat earlier actions.

Persistent Conversation Context

Files attached to a conversation remain linked to that specific conversation even after the user leaves and returns later.

When reopening an existing conversation, DocsChatV2 restores:

- Chat history
- Previously selected files
- AI context associated with the conversation

This continuity helps users resume work seamlessly without re-uploading documents or re-establishing context.

1.3.2 Central Chat Area – Viewing Questions and Responses

The center section of the application contains the **Chat Area**, where users interact directly with the AI assistant.

This section displays the full conversation thread in chronological order, making it easy to follow the flow of questions and responses.

The Chat Area also contains the **Input Bar**, which is used to submit questions and attach files.

Components of the Central Chat Area

Component	Description
Chat Window	Displays conversation history between the user and the AI
Input Bar	Used to enter questions and send prompts
Attachment Icon	Opens the File Picker Panel
Send Button	Submits the user query
Stop Button	Cancels an in-progress AI response

User Questions and AI Responses

Questions entered by the user appear directly within the conversation thread.

AI-generated responses appear below the related question, maintaining a continuous conversational structure that is easy to follow.

Streaming Responses

AI responses are streamed in real time as content is generated.

While a response is being generated:

- The **Send Button** automatically changes into a **Stop Button**
- Users can click the **Stop Button** to cancel the response immediately

Once the response completes or is stopped, the button changes back to the **Send Button**.

This behavior provides users with better control over long-running responses.

Input Bar and Message Submission

The **Input Bar** is located at the bottom of the Chat Area and contains three primary elements:

Position	Component
Left Side	Attachment Icon
Center	Message Input Field
Right Side	Send Button / Stop Button

Attachment Icon

The attachment icon opens the File Picker Panel, allowing users to select OneDrive documents for the active conversation.

Message Input Field

Users type questions into the input field.

For best results, questions should be:

- Clear
- Specific
- Related to the attached documents

Send Button

The Send Button submits the user query to the AI service.

Stop Button

While the AI is generating a response, the Send Button changes into a Stop Button, allowing users to stop generation instantly.

Response Formatting and Citations

AI responses support markdown formatting, including:

- Headings
- Bullet lists
- Numbered lists
- Bold text
- Code blocks

Responses may also include references to the source document names used to generate the answer.

AI Response Behavior

DocsChatV2 is designed to answer questions strictly from the documents attached to the active conversation.

If relevant information is not found within the attached documents, the AI responds with:

“This information is not in the provided documents.”

This behavior ensures that responses remain grounded in user-provided content.

1.3.3 File Picker Panel – Selecting and Managing Documents

The **File Picker Panel** allows users to browse, select, and manage OneDrive documents for the current conversation.

The panel is opened by clicking the **Attachment Icon** located within the Input Bar.

At the top of the File Picker Panel are four tabs that help users locate and manage files efficiently:

Tab	Purpose
Browse	Navigate OneDrive folders and files
Recent	Access recently used files
Shared	Access files shared through OneDrive
Selected	View files attached to the current conversation

Browse Tab – Navigating OneDrive Files and Folders

The **Browse** tab displays the user's OneDrive folder structure starting from the root directory.

Users can:

- Navigate through folders and subfolders
- Open directories
- Browse available documents
- Select files required for the conversation

This tab is useful when users know where their files are stored within OneDrive.

Recent Tab – Accessing Recently Used Files

The Recent tab displays documents recently accessed by the user through Microsoft Graph integration.

This allows users to quickly locate recently used files without manually browsing folders.

Typical use cases include:

- Continuing ongoing work
- Reusing recently opened documents
- Accessing frequently referenced files

Shared Tab – Viewing Shared Documents

The Shared tab displays files that have been directly shared with the user through OneDrive.

Users can select shared files in the same way as personal OneDrive files.

Selected Tab – Managing Attached Documents

The **Selected** tab displays all documents currently attached to the active conversation.

Users can:

- Review attached files
- Confirm correct document selection
- Remove unwanted files before asking questions

Because files remain associated with the conversation, reopening a previous conversation automatically restores the selected documents.

This continuity allows users to continue conversations without reattaching files.

1.4 Managing Conversations

1.4.1 Creating a New Conversation

1. Click **New Conversation** in the sidebar.
2. A new conversation is initialized.

1.4.2 Switching Conversations

1. Click any conversation in the sidebar.
2. The selected conversation's history and attached files are loaded automatically.

1.4.3 Conversation Limits

- Each user can maintain a limited number of conversations.

- Default limit: 20 conversations per user.
- Older conversations may need to be deleted if the configured limit is reached.

1.5 Attaching and Managing Files

DocsChatV2 allows users to attach OneDrive documents to conversations so the AI can generate context-aware responses using the content of those files.

Attached files remain associated with the active conversation to maintain continuity across user sessions.

1.5.1 File Limits

Setting	Default Value	Description
Maximum Files	5	Maximum files allowed per conversation
Maximum File Content Size	50 KB	Maximum combined content size of all selected files in a conversation

1.5.2 Supported File Types

DocsChatV2 supports the following document formats:

- PDF (.pdf)
- Word Documents (.docx)
- Excel Files (.xlsx, .xls)
- PowerPoint Files (.pptx)
- Text Files (.txt)
- Markdown Files (.md)
- CSV Files (.csv)
- JSON Files (.json)
- XML Files (.xml)
- HTML Files (.html, .htm)
- Rich Text Files (.rtf)

1.6 Chatting with the AI

1.6.1 Asking Questions

1. Ensure relevant files are attached to the active conversation.
2. Enter a question into the Input Field.
3. Click the **Send Button**.

1.6.2 Response Behavior

- Responses are generated only from attached documents.
- If relevant information is unavailable, the AI responds:

“This information is not in the provided documents.”

1.6.3 Streaming Responses

- Responses are streamed in real time.
- Users may stop generation using the **Stop Button**.

1.6.4 Source Citations

Responses may include references to the source file names used to generate the answer.

1.7 Notifications and Messages

1.7.1 Toast Notifications

Toast notifications are displayed for:

- Errors
- Warnings
- Success messages

1.7.2 Configuration Warning

Users may see the following message:

“Configuration may be out of date”

To refresh configuration:

1. Click **Refresh**
2. Updated settings are reloaded automatically

1.7.3 Banner Notifications

Banner notifications are displayed when administrators update user-impacting settings such as:

- File limits
- Conversation limits

1.8 Best Practices

- Attach only relevant documents
- Use clear and specific queries
- Avoid mixing unrelated topics in one conversation
- Review attached files before sending queries

2. Administrator Guide

2.1 Administrator Overview

Administrators have access to all end-user features, plus the **Admin Panel**, where application configurations are managed.

2.2 Accessing the Admin Panel

1. Log in as a user with admin privileges.

2. Click **Admin Panel**.

2.3 Managing AI Service Settings

2.3.1 Configuration Fields

- LLM Provider (Azure OpenAI / Google Gemini)
- Endpoint URL
- API Key

2.3.2 Steps

1. Navigate to **LLM Provider**
2. Enter required values
3. Save configuration

2.4 File and Conversation Limits

2.4.1 File Limits

Setting	Range
Total File Content Size	1–100 KB
File Count	1–10

2.4.2 Conversation Limits

Setting	Range
Conversations per User	1–100

2.5 Configuration Propagation

- Configuration changes are stored in Azure Cosmos DB.
- Connected client applications poll for configuration updates every 30 minutes.
- When configuration changes are detected, users are notified through an in-application banner prompting them to refresh the configuration.

2.6 Best Practices for Administrators

- Apply configuration changes during periods of low usage.
- Validate configuration changes in a non-production environment before applying them in production.
- Set file and conversation limits based on observed usage patterns.

3. Frequently Asked Questions (FAQ)

Q1. How are responses generated?

Responses are generated based only on attached document content.

Q2. Why is information missing?

The AI only uses attached files.

Q3. Can others see my chats?

No, data is fully isolated per user.

Q4. What happens when a conversation is deleted?

All associated messages and files are deleted and removed from Azure Cosmos DB.

4. Glossary

Term	Definition
LLM (Large Language Model)	AI model used for generating responses
RAG (Retrieval-Augmented Generation)	Combines retrieval with AI generation
Azure Cosmos DB	Microsoft's globally distributed, multi-model NoSQL database used to store the application's data

Microsoft Entra ID (MSAL)	Microsoft's cloud identity service; MSAL (Microsoft Authentication Library) signs in users and acquires access tokens
Azure Function	Serverless compute service
Microsoft Graph API	REST API used to access Microsoft 365 services, including OneDrive files