Published August 2026

VisibleThread 7.x: Customer Hosted Deployment Overview

1. Introduction

VisibleThread 7.x is a secure, browser-based application designed for GovCon and regulated industries to manage the proposal lifecycle. It supports FIPS 140-2 approved encryption allowing you to store & process Controlled Unclassified Information (CUI) and meet your CMMC 2.0 obligations. The platform's modular architecture ensures that connectivity-dependent features, including generative AI and external repositories, are optional. This ensures the deployment aligns strictly with your security and operational constraints.

Note on Security:

- **100% Isolated:** VisibleThread 7.x can operate entirely behind your firewall with no external data transmission.
- **Air-Gapped Ready:** The platform is built for "air-gapped" environments, sensitive, compartmented information facilities (SCIFs), or secure private clouds (Azure GCC High, AWS FedRAMP) with no internet connectivity.
- **Zero Visibility:** All data remains within your network; VisibleThread staff have no visibility into your proprietary information or documents.
- **Data Privacy:** All user and application data resides in the PostgreSQL database managed by you. Your data never leaves your network.
- **SSO:** SAML v2 based SSO using your IDP as the identity provider.
- **User and Workspace Segregation:** Unlimited number of shared or private workspaces can be created to allow users work on private projects or share projects and documents with teammates. Your system admin has full control over user workspaces.

Key System Interactions

This section highlights how the platform manages data and how it can be configured to connect with other systems:

- **Internal Data Control:** All document analysis, shreds, and proposal data are stored in a database you manage (PostgreSQL v16).
- **Repository Integration:** The platform can connect to your SharePoint (Standard, GCC, or GCC-High) to pull source documents directly into your workspace.
- **AI Gateway:** For AI features, VisibleThread can securely connect to Large Language Models (LLMs) that you provide via Azure AI, AWS Bedrock, or Local (Ollama) hosting. Your data stays within your network and is not used to train these models.

- **External Connectivity:** Optional features like Opportunity Search require outbound access to http://vtfetcher.visiblethread.com to pull down latest opportunity details.

2. Deployment Options Overview

VisibleThread 7.x features a modular framework where the Core platform serves as the base for optional add-on modules. Use this table to identify which services are standard and which require additional configuration.

Option	Name	Key Features	Requirements
Core Deployment	VT Core	Folders area access, Opportunity & Proposal Management, and Amendment Handling.	**16GB RAM** (minimum requirement), **100GB Storage**, PostgreSQL v16.
Add-on 1	Opportunity Search	Integration with sam.gov for sourcing and filtering opportunities directly in the platform.	Network access to the VisibleThread opportunity API (vtfetcher.visiblethread.com).
Add-on 2	SharePoint Integration	Seamlessly pull source documents directly from your own SharePoint repositories	Integration with your SharePoint Online environment.
Add-on 3	AI Features **All of these features can be turned on/off at the customers discretion*	**Visi Chat:** AI chatbot that is used to generate content.	Access to a Chat LLM (Large Language Model)
		VT RAG: Use your own documents to ground the Visi Chat AI.	Access to an embedding model
		Technical Alignment and Proposal Outlines	Access to the Document Analysis LLM

3. Infrastructure & Environment

Customer Hosted Options

- *Customer-Hosted On-Premises*: Full physical control for classified operations. Installs on Windows Server or Red Hat Enterprise Linux (RHEL).
- *Customer-Hosted Private Cloud:* Dedicated resources and isolation in environments like Azure GCC High or AWS FedRAMP.

System Requirements

VisibleThread requires:

- *RAM:* 16GB (Minimum requirement—this is an increase from previous versions).
- *CPU:* 2 CPU Cores.
- *Disk Space:* 100GB.
- *Database:* PostgreSQL v16 (standard local install or external like AWS RDS).
- *Database:* The pgvector extension is required.

4. Generative AI & LLM Strategy

VisibleThread 7.x can be configured to leverage LLM's to provide advanced AI capabilities, such as drafting new content. You can ground this AI context using documents from your own repositories. To support this, the platform can integrate with SharePoint Online, SharePoint GCC and SharePoint GCC-High. VisibleThread does not include an LLM as part of the deployment. The LLM must be provided and managed by your internal IT team.

Customization & Flexibility

VisibleThread acts as a secure gateway, allowing for high levels of administrative control:

- **Model Selection:** Customers may choose to use a single high-performance model for all tasks or use different models (e.g., a "lighter" model for chat and a more robust model for document analysis).
- **Feature Control:** AI capabilities are modular. For example, organizations can choose to enable Visi Chat while keeping RAG disabled.
- **Data Sovereignty:** Because the models are hosted within your own Azure, AWS, or local environment, all proprietary data stays within your network and is not used to train the LLMs.

Supported Model Interfaces

Models can be accessed via:

- Azure AI (GCC and non-GCC).
- AWS Bedrock
- Local Hosting: Customer-managed models (e.g., via Ollama).

While VisibleThread 7.x allows for customization, the following classes of models are recommended:

AI Component	Purpose	Recommended Model Class
Chat LLM	Visi Chat bot for content generation.	Mistral-Nemo, Anthropic Sonnet 4.5, AWS Nova Nova Pro 1, GPT 5.1. GPT 5.5.

AI Component	Purpose	Recommended Model Class
Embedding Model	Use Visi Chat with your documents, cites exact source information.	Mixedbread mxbai-embed-large or OpenAI text-embedding-3-small
Document Analysis LLM	AI-generated Technical Alignment and Proposal Outlines.	Anthropic Sonnet 4.5, AWS Nova Pro 1, GPT 5.1, 5.5.

**Note: model recommendations are subject to change.*

5. Architecture Overview

Single Node Deployment

The diagram below provides an overview of the VisibleThread platform and its interaction with other systems on your network.

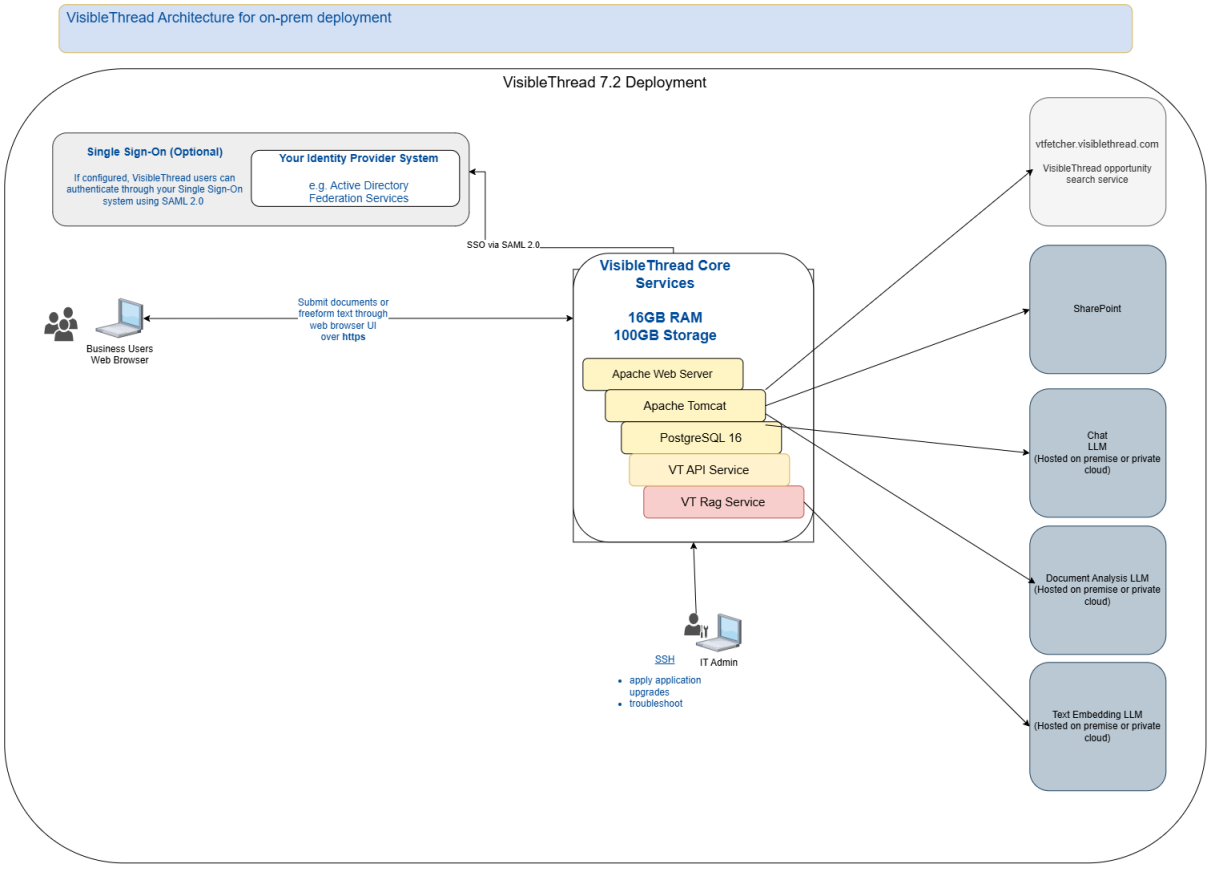

Published August 2026

Load Balanced Environment

VisibleThread can be deployed in a load balanced environment, allowing for horizontal scaling to support high availability. In this scenario multiple application servers can be deployed sharing an external database.