

In mission-critical environments, every second matters. Intelligence officers rely on Mattermost to detect and counter evolving threats with sovereign, structured workflows purpose-built for disinformation response. Whether at headquarters or in the field, Mattermost ensures data control, operational agility, and secure cross-agency collaboration to keep the mission on track.

Core Capabilities

Mattermost Agents enable users and teams to interact with advanced large language models (LLMs) from within Mattermost. Key features include:

- **Conversational AI:** Interact with AI agents via direct messages or channels for answering questions, drafting content, or summarizing threads.
- **Image Analysis:** If using compatible LLM providers, agents can perform vision tasks alongside text-based queries.
- **Tool Integrations:** Built-in support for semantic server search, user lookup, Jira issue retrieval, and GitHub integration, all while respecting user permissions.
- **Multi-Agent Support:** Ability to create and configure multiple agents—each with distinct personalities, assigned LLMs, and fine-grained access permissions (team/channel/user level).
- **Chat Summarization:** Automated summarization of threads or large sets of messages, including meeting transcripts and unread channel messages.
- **Usage Analytics:** Detailed logging for all LLM usage, supporting cost tracking, resource planning, and billing.

Extensibility

LLM and Provider Flexibility

Agents natively support popular LLM providers (OpenAI, Anthropic, Cohere, Azure OpenAI, OpenAI-compatible APIs), each configurable with custom authentication, models, and advanced features such as web search. Custom instructions allow tailoring agent behavior (tone, workflow compliance, organizational jargon) for each agent setup. The service-based architecture means multiple bots can share LLM configurations or have fully distinct setups.





Retrieval-Augmented Generation (RAG)

Full support for RAG workflows is available via configurable “embedding search” (semantic search), powered by vector search in PostgreSQL with the pgvector extension. Integration with embedding providers (e.g., OpenAI) and flexible chunking/tokenization of source text ensures high-quality semantic retrieval. RAG enables AI agents to answer questions based on company knowledge (posts, docs, etc.) rather than just the base LLM training data.

Model Context Protocol (MCP) Integration

The Model Context Protocol (MCP) standard integration allows connecting agents to external tools and services outside of built-in Mattermost APIs. Agents can be configured to connect to multiple MCP servers, each managing its own connection/session context for isolation and security. This opens the door to advanced use cases, including corporate RAG chains, external data sources, and custom agent tools/extensions.

AI Guardrails & Security

Mattermost Agents incorporate multiple layers of security and governance:

- **Permission Controls:** Agent access can be restricted at the team, channel, or user level to manage sensitive information exposure.
- **Security-first Integrations:** Built-in tool integrations only work in direct messages and require explicit user approval before execution.
- **LLM Trace and Token Logs:** Administrators can audit all LLM interactions and trace usage, ensuring transparency and compliance.
- **API Call Hostname Allowlist:** Output token limits and context-specific prompt instructions help enforce organizational policies and reduce the risk of data leakage or AI misuse.
- **Information and Link Rendering Limitations:** Agent response limitations further reduce exfiltration risk (e.g., plain text URLs, not clickable links).

Feature Licensing Overview

Feature	License Required?
Basic agent config, single agent, DMs	No license needed
Multiple agents, RAG, fine-grained access	Entry, Enterprise, or Advanced
MCP, usage analytics, advanced summarization	Entry, Enterprise, or Advanced
Image analysis (vision capabilities)	No license needed
Basic tool integrations	No license needed

Extensibility & Guardrails Summary

Area	Capabilities	Licensing
LLM Provider Choice	OpenAI, Anthropic, Cohere, Azure, OpenAI-compatible	Some features require paid license
RAG/Semantic Search	Embedding + pgvector; chunking & custom models	Entry/Enterprise+
MCP Integration	Connect with external tools/services via open protocol	Entry/Enterprise+
Custom Agent Behavior	Instructions, context, tool use, per-agent personality	Supported; advanced control needs license
Security/Guardrails	Permissions, logs, allowlists, logging, output limits	Logging & analytics need license

Unlock Sovereign AI for Your Organization with Mattermost Agents

Mattermost Agents are a robust, enterprise-ready solution for embedding advanced, extensible, and auditable AI workflows within the Mattermost platform—with a particular emphasis on security, compliance, and support for the latest agentic and RAG paradigms. Whether your organization needs simple conversational AI, sophisticated RAG-enabled automation, or tightly controlled enterprise AI guardrails, Mattermost Agents provide the flexibility and control necessary to meet diverse organizational requirements.



Learn more about Mattermost Agents in our Agents Admin Documentation.

Documentation

