AI Transparency Notes

FortiAI on FortiManager 1.1

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

May 27, 2026
FortiAI on FortiManager 1.1 AI Transparency Notes
02-110-1297460-20260527

TABLE OF CONTENTS

Change Log **5**
1. Purpose and Scope **6**
1.1 Document Purpose 6
1.2 What is FortiAI on FortiManager? 6
1.3 In-Scope Features 7
1.4 Out-of-Scope 7
1.5 Target Audience 8
2. Model(s) Used and Hosting **9**
2.1 Large Language Model 9
2.2 System Components 9
2.3 Vector Database Contents (RAG Knowledge Base) 10
2.4 Azure OpenAI Default Safety Policies 10
3. Design Methodology, Inputs and Outputs **11**
3.1 Design Philosophy 11
3.2 Interaction Model 11
3.3 Inputs 11
3.4 Outputs 12
4. Data Flows, Protection, and Retention **14**
4.1 Data Flow Overview 14
4.2 Data Masking and Obfuscation 15
4.3 Data Protection Layers 15
4.4 AI Proxy Logging 16
4.5 Client-Side Logging 16
4.6 Feedback Data Collection 17
5. Testing and Validation **18**
5.1 QA Testing 18
5.2 Logging Infrastructure Supporting Validation 18
5.3 Hallucination Reduction 18
5.4 Azure OpenAI Default Guardrails 19
6. Performance Metrics **20**
6.1 Metrics to Measure and Report 20
7. Known Limitations **21**
7.1 LLM Hallucination 21
7.2 Image Upload - No Data Masking Applied 21
7.3 User-Typed Sensitive Data - Partial Masking Gap 21
7.4 Scope of Data Submitted to the LLM 21
7.5 Masking Field Coverage 22
7.6 No Session Memory Across Conversations 22
7.7 FortiAI Feature Availability Depends on AI Service Connectivity 22
7.8 MCP Server Public API - Not Yet Available 22

8. Risk Analysis and Categorization **23**
Appendix A: Component Summary **24**
Appendix B: Glossary **25**

Change Log

Date	Change Description
2026-05-27	Initial release.

1. Purpose and Scope

1.1 Document Purpose

This AI Transparency Notes document provides a factual disclosure of the FortiAI feature embedded within FortiManager. It describes how the AI assistant works, what data it processes, how it is designed, and what safeguards are in place - enabling administrators, compliance teams, and auditors to make informed decisions about its deployment and use.

All content in this document is based directly on the FortiAI FortiManager product specification.

1.2 What is FortiAI on FortiManager?

FortiManager delivers a single-pane-of-glass management experience spanning connectivity, resource utilization, device settings, policy status, and alerts. It orchestrates FortiGates across varied form factors - VM, CNF, standard and ruggedized hardware - and extends management to the broader FortiOS-based networking ecosystem including FortiSwitch, FortiAP, and FortiExtender, ensuring consistent enforcement of unified security policies across the organization.

While FortiManager is a powerful platform, it carries a learning curve and some operations can become complex. FortiAI on FortiManager is an intelligent assistant designed to reduce that complexity: to simplify operations and to bring features that might otherwise be multi-step, complicated, or impossible to accomplish via the regular UI, while keeping security as a top priority built in from day one.

FortiAI provides a chat window within the FortiManager UI through which administrators can ask questions and request tasks in natural language. Answers and results are returned in the same chat window. Specific features can also be triggered directly from the UI via dedicated buttons - for example, explaining the difference between two configuration revisions, or explaining a planned installation.

FortiAI specifically harnesses Generative AI to automate routine tasks, optimize configurations, accelerate diagnostics, and streamline policy management. It integrates with FortiAIOps to support predictive network analysis and automated troubleshooting, enabling administrators to proactively identify and resolve potential issues before they impact network performance.

1.3 In-Scope Features

Agent / Feature	Description
CLI and Jinja Script Agent	Turns plain-language intent into CLI or Jinja scripts for FortiGate and FortiManager. Supports creating, modifying, and explaining scripts - whether starting from scratch or refining existing ones - and guides toward best practices.
Device Management Agent	Provides real-time device health, licensing, firmware status, upgrade readiness, and security risk assessment. Explains configuration changes in plain language - including the CLI commands used and their intent. Supports searching devices by interface configuration. Gives visibility into recent changes, installation activity, and devices needing attention.
Firewall Policy Agent	Enables natural language search, inspection, creation, editing, cloning, version comparison, and rollback of firewall policies and related objects across the FortiManager environment.
VPN and SD-WAN Network Diagnostic Agent	Diagnoses and resolves VPN, SD-WAN, performance, and device issues through natural conversation. Analyses symptoms, walks through root cause step by step, and guides administrators to restore network stability - without requiring deep technical knowledge upfront.
Provisioning Template Agent	Guides SD-WAN overlay design, site connectivity, and traffic flow optimization. Creates and manages reusable templates for BGP, IPsec VPN, and system settings. Highlights configuration differences and conflicts before changes are applied. Supports device onboarding, variable management, and consistent deployment at scale.
General Knowledge Agent	Answers product questions, compares features and specifications, navigates to specific UI pages within FortiManager, and locates relevant Fortinet documentation. Serves as a guided entry point for FortiManager and Fortinet product knowledge.

The agents can either work alone to solve problems in particular areas or team up to tackle more difficult ones.

1.4 Out-of-Scope

- Autonomous configuration changes without explicit administrator confirmation - all data mutations require user approval
- Processing of end-user personal data or enterprise business content unrelated to network management
- FortiAI does not modify user input or raw data beyond the masking and unmasking of sensitive fields described in the following topic: 4. Data Flows, Protection, and Retention on page 14
- MCP Server for public or third-party use - not yet released; Fortinet reserves the right to release it in a future version

1.5 Target Audience

Audience	Relevance
Network Security Administrators	Primary users; interact with FortiAI daily for device management, policy work, and diagnostics.
Security Operations Teams	Consumers of AI-generated diagnostics, reports, and remediation guidance.
IT Compliance and Audit Teams	Review of AI governance, data handling, and acceptable use boundaries.
CISOs / IT Leadership	Strategic oversight of AI deployment and associated risk posture.
Data Protection Officers	Privacy impact review; data flows, masking practices, and retention policies.

2. Model(s) Used and Hosting

2.1 Large Language Model

FortiAI on FortiManager currently uses Azure OpenAI GPT-4.1 as its large language model (LLM). Fortinet retains the right to update the LLM provider and model as necessary, and will document such changes in the product release notes to represent the current state of the production system.

The LLM is used to: generate text responses; call tools from the MCP server or UI; read and edit data in FortiManager's database (based on RBAC); display UI components to present results; obtain user confirmation before taking actions; summarize Fortinet documentation retrieved from the vector database via file search; and transcribe images to text.

Customer data is not used to train the LLM.

2.2 System Components

Component	Role	Hosting Location
FortiAI UI	Collects user queries and tasks. Displays results and answers. Obtains user permission before any data in FortiManager is updated.	On-premises (FortiManager appliance)
MCP Client (on FortiManager)	Assembles system prompts. Designs and executes plans for tasks. Calls tools from the MCP server or UI tools (to display information or retrieve input from UI). Queries the vector DB for documentation search. Performs data masking before sending to LLM and unmasking on receiving the response. Uses the LLM chat completion API.	On-premises (FortiManager appliance)
MCP Server (on FortiManager)	Provides tools to get data from FortiManager, update data to FortiManager, and run services on FortiManager and managed FortiGates. Implements authentication - only UI login users can access it. Implements RBAC permission checks per tool call, consistent with FortiManager's existing RBAC. Fortinet reserves the right to implement other authentication methods if releasing the FortiManager MCP server for public use.	On-premises (FortiManager appliance)

Component	Role	Hosting Location
AI Proxy	Fortinet's internal cloud service through which all Fortinet products access the LLM via a centralized hub. Logs all traffic to and from the LLM. Controls AI license and token usage.	Fortinet Internal cloud
Vector DB (RAG)	Contains curated Fortinet documentation used for retrieval-augmented generation. See Section 2.3 for content list.	Fortinet Internal cloud
Azure OpenAI GPT-4.1	The large language model. Performs text generation, tool calling, image transcription, and documentation summarization.	Microsoft Azure (accessed via Fortinet AI Proxy)

2.3 Vector Database Contents (RAG Knowledge Base)

The vector database used for Retrieval-Augmented Generation currently contains the following Fortinet documentation. Content is refreshed with each product release cycle.

- Latest release of the FortiGate Administrator Guide.
- Latest release of the FortiManager Administrator Guide.
- Latest release of the FortiManager CLI Reference.
- Latest FortiGate Product Matrix and Purchase Guide.
- Latest release of the FortiManager Best Practices guide.
- Latest release of the FortiManager Examples guide.
- Latest release of the FortiManager 4D Resources for SD-WAN.

2.4 Azure OpenAI Default Safety Policies

Azure OpenAI default guardrail and content control policies are enabled for the GPT-4.1 deployment used by FortiAI. The details of these policies are published by Microsoft and can be found at: Default Guardrails and controls policies for Microsoft Foundry Models (classic)

3. Design Methodology, Inputs and Outputs

3.1 Design Philosophy

FortiAI on FortiManager is designed with two primary objectives: (1) to simplify FortiManager operations and unlock capabilities that would otherwise require multiple steps, deep expertise, or are impossible via the standard UI; and (2) to keep user data safe, with security built into the design from day one.

FortiAI is explicitly designed as an advisor and assistant - not an autonomous agent. Any action that writes or modifies data in FortiManager requires explicit user confirmation through the UI before it is executed. FortiAI will not modify user input or data beyond the defined masking and unmasking of sensitive fields.

3.2 Interaction Model

User input may be processed in several ways depending on the task:

- Simple query: user input is sent directly to the LLM once and a response is returned
- Search and match: input is interpreted by the LLM and used to search or match against FortiManager data
- Agentic multi-step task: the LLM calls MCP server tools, collects returned data, and may send multiple rounds to the LLM before producing a final answer - with masking applied to sensitive data at each round

Importantly, some output may also be delivered directly from MCP tool results to the UI using standard UI components, without passing through the LLM for summarization. This path is used when structured data can be displayed directly without requiring language generation.

3.3 Inputs

Input Category	Examples	Sensitivity	Handling Before LLM Transmission
Natural language queries typed by user	"What devices have firmware vulnerabilities?", "Create a policy for HTTPS traffic"	Low to variable	Not modified. Note: sensitive values typed directly by the user (e.g. a device name or IP) may not be masked if not already present in session context - see Section 7.3

Input Category	Examples	Sensitivity	Handling Before LLM Transmission
User-uploaded images	Screenshots of CLI output or configuration pages submitted for transcription	Variable - potentially High	The entire image is submitted to the LLM without any data masking applied - see Section 7.2
FortiManager device configurations	Interface configs, policy packages, routing tables, VDOM settings, policy objects	High	Retrieved via MCP server tools. Sensitive fields masked by MCP Client before LLM transmission.
Device telemetry and monitoring data	CPU/memory usage, SD-WAN health checks, session data, bandwidth statistics	Medium	Retrieved via MCP server tools. Sensitive fields masked before LLM transmission.
Audit and change logs	Configuration change history, installation activity, revision differences	Medium	Retrieved via MCP server tools. Sensitive fields masked before LLM transmission.
Device inventory and topology	Device names, ADOM/VDOM structure, firmware versions, licensing status	Medium-High	Retrieved via MCP server tools. Sensitive fields masked before LLM transmission.
User feedback (opt-in)	Quality rating submitted via the Send Feedback feature	Low	Contains platform name (FortiManager), version number, user rating, user feedback, and conversation context immediately prior to submission only. No other user information is collected.

3.4 Outputs

Output Type	Description	Delivery Method	User Action Required
Natural language responses	Plain-text answers, explanations, and summaries generated by the LLM	Chat window	None (informational)
CLI / Jinja scripts	Generated or modified automation scripts	Chat window	User review before executing

Output Type	Description	Delivery Method	User Action Required
Device summary reports	Structured device health, firmware, licensing, and risk assessments	Chat window or PDF download	Review; act as appropriate
Policy drafts	New or edited firewall policy objects	Chat window + FortiManager UI	Explicit user confirmation required before writing to FortiManager DB
Provisioning templates	BGP, IPsec VPN, and system configuration templates	Chat window + FortiManager UI	Explicit user confirmation required before applying
Diagnostic findings	Root cause analysis and step-by-step remediation guidance for VPN, SD-WAN, and device issues	Chat window	Manual or user-confirmed remediation
Direct tool data via UI components	Structured data from MCP tool responses rendered directly in the UI without LLM summarization	Standard FortiManager UI components	None; displayed as-is from FortiManager data
UI navigation	FortiAI navigates the user to specific FortiManager UI pages in response to natural language requests	FortiManager UI	None (navigational)
Revision / installation explanations	Plain-language explanation of what changed between configuration revisions, or what a planned installation will do	Chat window (button-triggered from UI)	None (informational)
RAG-based documentation answers	Summaries drawn from the vector database (Fortinet admin guides, best practices, CLI reference, SD-WAN resources)	Chat window	None (informational)

4. Data Flows, Protection, and Retention

4.1 Data Flow Overview

When a user interacts with FortiAI on FortiManager, data flows as follows: it is read from FortiManager local databases via MCP server tools, then passes through the Fortinet AI Proxy, then to the LLM, then back to the AI Proxy, then to the FortiAI MCP Client, and finally to the FortiAI UI. The AI Proxy may save the LLM chat completion requests and responses, but the data at that point is already masked.

The following table describes each stage in detail:

#	Stage	Description
1	User Input (UI)	Administrator enters a natural language query or uploads an image in the FortiAI chat window. Alternatively, the user clicks a button in the FortiManager UI to trigger a specific feature (e.g. explain revision difference, explain planned installation).
2	MCP Client - Plan	The MCP Client assembles the system prompt and determines what MCP server tools or UI interactions are needed to fulfil the task.
3	MCP Server - Tool Calls	MCP server tools retrieve data from FortiManager local databases, or run services on FortiManager and managed FortiGates. Authentication is enforced: only users with a valid GUI session can invoke tools. RBAC permission checks are applied per tool call, consistent with FortiManager's role-based access control.
4	MCP Client - Masking	The MCP Client masking service replaces sensitive fields in the tool response data (IPs, MAC addresses, passwords, secret keys, certificates, names) with masked tokens. A local key-value mapping is maintained for unmasking later. Note: masking applies to MCP tool response data. If the user typed sensitive values directly into the chat, those may not be masked - see Section 7.3.
5	LLM Request via AI Proxy	The masked prompt is transmitted to the Fortinet AI Proxy, which routes it to Azure OpenAI GPT-4.1. All traffic to and from the LLM passes through the AI Proxy. The AI Proxy may log these requests and responses; since data is masked before this point, the logs contain masked data.

#	Stage	Description
6	LLM Inference	Azure OpenAI GPT-4.1 processes the masked prompt. The response may include generated text, tool call instructions, or structured output. For complex tasks, multiple rounds of LLM calls may occur (loop back to Step 3).
7	MCP Client - Unmasking	The MCP Client replaces masked tokens in the LLM response with their original values before the result is passed to the UI.
8	Output to UI	The response is delivered to the FortiAI UI. Output may be: (a) LLM-generated text displayed in the chat window; (b) data from MCP tool results rendered directly via standard UI components, without LLM summarization; (c) a downloadable PDF report; or (d) UI navigation to a specific FortiManager page.
9	User Confirmation (if data mutation)	If the task requires writing to or modifying the FortiManager database, the UI presents the proposed change to the user and requires explicit confirmation before any data is written. Without confirmation, no mutation occurs.

4.2 Data Masking and Obfuscation

Before data retrieved from FortiManager via MCP server tools is sent to the LLM, the MCP Client masking service identifies and replaces sensitive fields with masked tokens. The following field types are masked:

- IP addresses
- MAC addresses
- Passwords
- Secret keys
- Certificates
- Names (device names, usernames, and similar identifiers)

The masking service maintains a local key-value mapping of original-to-masked values for the session. After the LLM returns a response, the service replaces all masked tokens with their original values before displaying them to the user.

4.3 Data Protection Layers

The following access and data protection controls are implemented in FortiAI on FortiManager:

Layer	Control	Description
1	AI License Requirement	Users must hold a valid AI license to use any FortiAI feature, including access to the FortiManager MCP server.
2	Feature Toggle	Administrators can hide all FortiAI features entirely via a CLI configuration option, if desired.
3	FortiAI User Designation	Up to 3 administrators can be designated as FortiAI users. Only those designated administrators can use FortiAI features; other admins cannot access them.
4	Authentication	To use FortiAI from the UI, a user must have a valid, active FortiManager GUI login session.
5	RBAC Enforcement (MCP Server)	Each MCP server tool automatically checks the calling user's role-based permissions before executing, consistent with FortiManager's existing RBAC model.
6	Data Masking (MCP Client)	Sensitive fields in MCP tool response data are masked by the MCP Client before transmission to the LLM. Data is unmasked from the LLM response before display to the user.
7	Future MCP Public API (Planned)	If the MCP server is released for public use in a future version, strong authentication and same RBAC permission checks will be enforced for external callers.

4.4 AI Proxy Logging

All LLM traffic - chat completion requests and responses - passes through the Fortinet AI Proxy. The AI Proxy logs all this traffic. Because data is masked by the MCP Client before transmission to the AI Proxy and LLM, the logs contain masked data rather than original sensitive values.

4.5 Client-Side Logging

FortiAI maintains two dedicated log files on the FortiManager appliance for debugging purposes:

- MCP Client log: records failed LLM chat completion requests and responses (client-side communication)
- MCP Server log: records failed MCP tool requests and responses (server-side communication)

Additionally, the FortiAI UI provides a button that displays the current session's data masking key-value pairs in a table, allowing administrators to inspect what data has been masked within a session.

4.6 Feedback Data Collection

FortiAI includes an optional "Send Feedback" feature for users to rate the quality and performance of FortiAI responses. Feedback submissions contain only the following: the platform name (FortiManager), the FortiManager version number, the user's rating, the user’s feedback and the context of the conversation immediately prior to submission. No other user information is collected through this mechanism.

5. Testing and Validation

5.1 QA Testing

A dedicated QA team tests and validates FortiAI outputs. Testing covers the full pipeline: user input, MCP tool execution, LLM inference, data masking and unmasking, and final response quality.

5.2 Logging Infrastructure Supporting Validation

Log / Tool	What It Captures	Purpose
MCP Client log file	All LLM chat completion requests and responses (client-side)	Audit of AI interaction history; debugging response quality issues
MCP Server log file	All MCP tool requests and responses (server-side)	Audit of tool calls, data retrieval, and RBAC enforcement
Data Masking UI view	Current session masking key-value pairs, displayed in a table via a UI button	Allows administrators to verify what data has been masked in a session

5.3 Hallucination Reduction

GPT-4.1 may hallucinate in some outputs. The FortiAI system prompts are specifically designed to limit LLM output to content relevant to FortiManager and FortiGate. Testing has been conducted to reduce hallucination rates. A persistent disclaimer - "FortiAI can make mistakes" - is displayed in the FortiAI chat window at all times, ensuring users maintain appropriate critical review of AI-generated content.

5.4 Azure OpenAI Default Guardrails

Azure OpenAI default guardrail and content control policies are enabled on the GPT-4.1 deployment. These provide a baseline layer of safety filtering at the model level. Full details are published at: Default Guardrails and controls policies for Microsoft Foundry Models (classic)

6. Performance Metrics

6.1 Metrics to Measure and Report

Category	Metric	Measurement Method	Value
Accuracy	Response accuracy rate - percentage of LLM responses judged factually correct for the FortiManager/FortiGate context	Human expert evaluation on a holdout test set of representative queries	Under testing plan
Accuracy	All support agents correctness rate - percentage of generated scripts that execute on a target FortiGate without error	Cross-check of AI responses against FortiManager ground truth state	Under testing plan
Accuracy	Diagnostic accuracy - percentage of root causes correctly identified, validated against actual resolution outcomes	Post-incident correlation review	Under testing plan
Safety	Masking coverage rate - percentage of defined sensitive field types correctly masked before LLM transmission, across a representative dataset	Automated masking validation on masking UI	Under testing
Reliability	AI Proxy / LLM service availability (uptime percentage)	AI Proxy monitoring and incident log	Cluster high availability, geo redundancy
Reliability	Graceful degradation - FortiManager core functionality operates correctly when FortiAI service is unavailable	Failure injection testing (AI Proxy unreachable scenarios)	Under testing
User Quality	User feedback rating distribution (from in-product Send Feedback feature)	Aggregated from AI Proxy or FortiManager feedback data	Rating and Feedback is periodically reviewed by Product Management Team
Token Usage	FMG is only responsible for display, the token allocation and accounting is the FortiAI Proxy Server responsibility	UI has Token usage	Current monthly token usage/Total Monthly Entitled Tokens

7. Known Limitations

7.1 LLM Hallucination

GPT-4.1 may hallucinate - producing plausible but incorrect answers. While FortiAI system prompts are designed to constrain LLM output to FortiManager and FortiGate topics, and testing has been conducted to reduce this, hallucination cannot be fully eliminated. A disclaimer ("FortiAI can make mistakes") is permanently displayed in the chat window. Users should always validate AI-generated scripts, policy recommendations, and diagnostic conclusions before acting on them. For any action that modifies FortiManager data, a user confirmation step is required - but this does not substitute for careful human review of the proposed change.

7.2 Image Upload - No Data Masking Applied

When a user uploads an image for transcription, the entire image is submitted to Azure OpenAI GPT-4.1 without any data masking applied. If the image contains sensitive information - such as IP addresses, credentials, configuration data, usernames, or any other sensitive content - that information will be transmitted to the LLM unmasked. Users should be aware of this limitation before uploading screenshots or images containing sensitive data.

7.3 User-Typed Sensitive Data - Partial Masking Gap

The data masking service operates on data returned from FortiManager via MCP server tool responses. If a user types sensitive values directly into the chat window - for example, a device name that is not already present in the current session context - those values may not be masked before being sent to the LLM. Users should avoid typing sensitive credentials, keys, or other high-sensitivity values directly into the FortiAI chat window.

7.4 Scope of Data Submitted to the LLM

FortiAI may submit data collected from FortiManager to the LLM. This data may include device configurations, policy packages, and monitoring data from managed devices. While defined sensitive fields within this data are masked (Section 4.2), the overall volume and breadth of configuration data transmitted to the LLM is by design

- the AI requires this context to answer management questions accurately. Administrators should factor this into their assessment when deploying FortiAI in highly sensitive or restricted environments.

7.5 Masking Field Coverage

The masking service covers the field types identified as sensitive by the product team: IP addresses, MAC addresses, passwords, secret keys, certificates, and names. There may be additional fields that specific customers or environments consider sensitive which are not currently included in the masking scope. Requests for expanded masking coverage can be submitted through Fortinet support and will be evaluated for future releases.

7.6 No Session Memory Across Conversations

FortiAI does not retain context between separate user sessions. Each new session begins without knowledge of prior conversations. Users must re-provide relevant context when starting a new session.

7.7 FortiAI Feature Availability Depends on AI Service Connectivity

FortiAI features depend on connectivity to the Fortinet AI Proxy and Azure OpenAI. If these services are unavailable, FortiAI features will not function. FortiManager's core management functionality is unaffected by AI service interruptions - FortiAI is an enhancement layer and is not required for core operations.

7.8 MCP Server Public API - Not Yet Available

The FortiManager MCP Server has not yet been released for public or third-party use. Fortinet reserves the right to release it in a future version. When released, it will enforce strong authentication and the same RBAC-based permission checks as the current internal implementation.

8. Risk Analysis and Categorization

FortiAI on FortiManager is best categorized as a non-high-risk AI system under Regulation (EU) 2024/1689. Based on its documented intended purpose, FortiAI assists authorized administrators with network and security management tasks, including diagnostics, policy drafting, script generation, configuration explanation, and documentation retrieval. It is not intended to make, recommend, or materially influence decisions in any Annex III high-risk domain, including employment, education, law enforcement, access to essential services, migration, critical infrastructure safety decisioning, or the administration of justice.

FortiAI is also not designed or intended for any prohibited AI practice under Article 5, such as manipulative, exploitative, social scoring, certain biometric, or impermissible law-enforcement uses. Although FortiAI is not currently classified as high-risk, it is designed to support applicable transparency obligations where users interact with an AI system or receive AI-generated outputs. The AI Act transparency framework includes disclosure obligations for certain AI interactions and AI-generated content. As an additional safeguard, FortiAI does not autonomously execute proposed configuration changes; any such change requires review and affirmative confirmation by an authorized administrator before implementation.

Appendix A: Component Summary

Component	Technology / Protocol	Hosting Location
FortiAI UI	Web-based chat interface within FortiManager; standard FortiManager UI components for direct tool output	On-premises (FortiManager appliance)
MCP Client	Prompt assembly, task planning, tool orchestration, vector DB queries, data masking/unmasking, LLM chat completion API calls	On-premises (FortiManager appliance)
MCP Server	MCP protocol; provides tools for FortiManager DB read/write and FortiGate service execution; authentication and RBAC enforcement	On-premises (FortiManager appliance)
AI Proxy	Fortinet internal cloud service; LLM traffic routing, logging (masked data), AI license and token control	Fortinet internal cloud
Vector DB	RAG knowledge base containing Fortinet FortiGate and FortiManager documentation	Fortinet cloud / FortiManager
Azure OpenAI GPT-4.1	Large language model for text generation, tool calling, image transcription, documentation summarization	Microsoft Azure (via Fortinet AI Proxy)

Appendix B: Glossary

Term	Definition
ADOM	Administrative Domain - a logical partition in FortiManager for managing a subset of devices
AI Proxy	Fortinet's internal cloud service that routes LLM traffic for Fortinet products, logs all LLM interactions (in masked form), and controls AI licensing and token usage
FortiAI	The branded name for FortiAI's generative AI capabilities embedded within FortiManager
FortiAIOps	Fortinet's AI-driven operations product for predictive network analysis; integrates with FortiAI in FortiManager
GPT-4.1	The Azure OpenAI large language model currently used by FortiAI on FortiManager
Hallucination	A phenomenon where an LLM generates plausible-sounding but factually incorrect content
Jinja Script	A templating language used by FortiManager for dynamic, variable-driven configuration generation
MCP (Model Context Protocol)	A protocol enabling standardized communication between AI models and tool/data providers
MCP Client	The FortiManager component that orchestrates AI tasks: assembles prompts, plans and executes tool calls, queries the vector DB, and manages data masking/unmasking
MCP Server	The FortiManager component that exposes tools to read/write FortiManager data and run services on FortiManager and FortiGate devices; enforces authentication and RBAC
RAG	Retrieval-Augmented Generation - an AI architecture that retrieves relevant reference documents and includes them as context before generating a response
RBAC	Role-Based Access Control - restricts access to features and data based on the authenticated user's assigned role and its associated permissions
System API User	A FortiManager account type used for programmatic/API access, as distinct from an interactive GUI session user. System API users can access the MCP server subject to RBAC checks.
VDOM	Virtual Domain - a virtual instance within a FortiGate device providing logical network and policy segmentation

www.fortinet.com