



Making Enterprises AI-native

An Enterprise solution that serves as a full AI-powered assistant integrating with almost all sources of Enterprise Data plus a multi-modal inference engine used per purpose and usage patterns

PILOTS LIVE

4

Fortune 500 enterprises

1

QBA PIPELINE

\$5M

pre-signed via existing channel

POSITIONING

Sovereign AI

EU + US regulated

Smart ● Secure ● Sovereign

Executive Summary

Q-GPT is an Enterprise AI knowledge platform enabling secure conversational AI across the entire enterprise ecosystem. Built on open-source technologies advanced Retrieval Augmented Generation (**RAG**) delivering accurate, cited responses with full data sovereignty.

Q-GPT powers business-critical intelligence, governance, and automation. It is designed to work with multiple models and internal vector database. Its innovation combines hybrid search, contextual retrieval, **vLLM** inferenc.



40+ Data Sources

Connects to SharePoint, SAP S4/HANA, Microsoft D365, Microsoft 365, Slack, Outlook, G-Suite and more.

Hosting can be On-prem/ Hybrid/ Private Cloud

Enterprise data is kept secured and sovereign by the choice of models and deployment.

Data Security Compliant

Built for EU AI Act, DORA and GDPR regulations. DPIA compliant for ISO 27001 and SOC2.

Open Source

No vendor lock-in. MIT Licensed core platform. No separate subscription costs other than token fees for chosen models.

White Labeling

Q-GPT can be white labeled by enterprises with their own branding and appearance.


Problem(s)

- ✗ Training Public LLMs
- ✗ Unassembled Enterprise Knowledge
- ✗ Tribal dependency
- ✗ Inefficient audit & compliance
- ✗ Uncontrolled AI usage risk
- ✗ Frenzied investments on AI initiatives
- ✗ Risk concentration with vendor lock-in
- ✗ Data sovereignty not respected
- ✗ Hallucinations used for enterprise decisions

Solutions

- ✓ Unified Enterprise Knowledge Graph
- ✓ Governed AI with source level ACLs
- ✓ Indexed, retrievable, reasoning-ready knowledge
- ✓ Reduced dependency on tribal knowledge
- ✓ Controlled innovation without regulatory risk
- ✓ Air-gapped LLM access to preserve data sovereignty
- ✓ Inhouse RAG and Prompt Orchestration
- ✓ Reinforced knowledge graph to reduce hallucinations
- ✓ Controlled TCO for AI investments

Q-GPT Introduction



Unified Knowledge Access

Native connectors to **40+ enterprise systems** (*Microsoft 365, Google Workspace, SAP S4/HANA, Microsoft D365, databases, and more*)

Self-Hosted Deployment

Complete **On-premises** or **Private Cloud** deployment-data never leaves the customer environment

RAG-Powered Accuracy

Every response is grounded in actual documents with **full citation to sources**

Agentic Capabilities

Deep research, code execution, web search, and custom tool integration

Permission Aware

Automatically syncs document permissions from source systems. Users only see what they're authorized to access

Q-GPT Built on Core AI Technologies



Q-GPT Leveraging Open-Source

Current AI Landscape



AI Learning Platforms



- o AI-enhanced LMS
- o Training-focused
- o SaaS-first deployment
- o Limited cross-system intelligence

Enterprise AI Search



- o Strong knowledge retrieval
- o Cloud-native
- o Closed model ecosystems
- o Not infrastructure-level platforms

AI Automation Assistants



- o Workflow automation
- o IT-centric use cases
- o SaaS dependent
- o Not learning-centric



Q-GPT Advantage



Sovereign by Design

- o Fully self-hosted
- o EU data compliance native
- o Zero data externalization

Model Independence

- o LLM-agnostic architecture
- o No vendor lock-in
- o Future-proof AI strategy

Enterprise-Wide Intelligence

- o 40+ system connectors
- o Hybrid + contextual retrieval
- o Cross-application reasoning

AI Learning Everywhere

- o Continuous learning across tools
- o Not confined to LMS
- o Embedded enterprise intelligence

Open-Source Community Impact

Active contributor community enabling rapid feature development | Transparency in AI implementation—organizations can audit the code | Ecosystem of extensions, connectors, and integrations

Q-GPT Product Roadmap


Our roadmap includes two breakthrough technologies that will further differentiate the platform

Q4 2025

Q1 2026

Q2 2026

Current Interfaces



StructRAG with Knowledge Graph Integration

Q-GPT will go beyond traditional RAG by implementing structured knowledge representation based on the StructRAG research. The Knowledge Graph module will extract entities and relationships from documents, enabling:

- **Entity-Centric Queries**
- **Graph-Enhanced Retrieval**
- **Relationship Mapping**
- **Multi-Hop Reasoning**


Recursive Language Models - RLM

A breakthrough inference strategy for handling unbounded context lengths. Unlike traditional approaches that struggle with "context rot" in long documents, RLM enables:

- **Context-Centric Decomposition**
- **Recursive Sub-Queries**
- **Unbounded Context Handling**
- **Smaller Models, Better Results**

Q-GPT Role of AI

AI is not an add-on feature in Q-GPT—it is the core product. Every aspect of the platform is designed around AI capabilities



Q-GPT Connectors

Connector Management

 Q-GPT

Exit Admin

Connectors

Existing Connectors

Add Connector

Document Management

Document Sets

Explorer

Feedback

Custom Assistants

Assistants

Slack Bots

MCP Actions

OpenAPI Actions

Standard Answers

Q-GPT version: 0.0.0-dev

Arbab

Add Connector

Search Connectors

Popular

 Confluence

 Sharepoint

 Google Drive

 Salesforce

 HubSpot

 Gong

 Frappe HRMS

Knowledge Base & Wikis

 Confluence

 Sharepoint

 Coda







Knowledge & Wikis


Confluence


Notion


SharePoint


Coda


GitBook

Cloud Storage


Google Drive


Dropbox


S3


Azure


Oracle

Collaboration


Slack


Teams


Gmail


Discord


Email

CRM & Sales


Salesforce


HubSpot


Gong Yoo


Fireflies


Highspot

Operations


Jira


Zendesk


Asana


Productboard


Freshdesk

Code


GitHub


GitLab


Bitbucket

Q-GPT MCP Actions

MCP Management



MCP Actions

Connect MCP (Model Context Protocol) servers to add custom actions and tools for your assistants.

Q Search servers...

Outlook
Arcade MCP Server

Teams
Arcade MCP Server

Google Search
Arcade MCP Server

Categories

- All Tools**
- Productivity & Docs**
- Communication**
- Developer Tools**
- Search**
- Sales & CRM**
- Entertainment**
- Finance & Payments**
- Customer Support**

MCP Action Hub



Confluence



SharePoint



Notion



Zendesk



Slack



Dropbox



Google Drive



Asana



Gmail



Teams



HubSpot




Salesforce



Asana



Outlook



Google Docs



Google Sheets

Q-GPT Multi Modal Options

LLM Setup

 **LLM Setup**

Enabled LLM Providers

If multiple LLM providers are enabled, the default provider will be used for case!

OpenRouter

(OpenRouter)

Default

Edit

Azure OpenAI

(Microsoft Azure Cloud)

Enabled

Edit

Set as default

Add LLM Provider

Add a new LLM provider by either selecting from one of the default provide

OpenAI

Set up


Anthropic

Set up


Ollama

Set up

Available LLMs

 Open AI	 Claude	 Gemini	 Mistral
 AWS Bedrock	 Azure OpenAI	 Llama	 Nemotron
 Moonshoot	 ZAI Zhipu	 Qwen	 DeepSeek
 xAI	 IBM Granite	 Cohere	 Llama CPP

Q-GPT Metrics



Ethical Design & Guardrails within Q-GPT



Personal Data Handling

- **Data Sovereignty by Design**
 - Fully self-hosted — all data remains within customer infrastructure
 - No external AI providers or third-party data sharing
 - GDPR-compliant and aligned with EU data localization (ideal for regulated industries)
- **Enterprise-Grade Encryption**
 - AES-256 encryption at rest
 - TLS 1.3 encryption in transit
 - Encrypted backups and secure credential storage
- **Data Retention Controls**
 - Configurable TTL policies for chat history
 - Automatic expiration and cleanup
 - User-controlled conversation history



Transparency to Users

- **Citation System**
 - Every AI response includes source document citations
 - Transparent visibility into referenced documents
 - Enables verification and builds trust
- **Audit Logging**
 - Complete query and response history with timestamps
 - User attribution for all interactions
 - Admin action audit trails for compliance
- **Explainable Retrieval**
 - Visibility into retrieved documents
 - Insight into document ranking logic
 - Clear understanding of why specific answers were generated



Bias & Fairness Management

- **User Feedback Integration**
 - Users can upvote/downvote responses and documents
 - Feedback adjusts document ranking scores over time
 - Continuous improvement based on user signals
- **Source Diversity**
 - Configurable source filtering prevents over-reliance on single systems
 - Time filtering reduces bias toward stale information
 - Balanced retrieval across document types
- **Permission-Based Filtering**
 - Results are filtered based on user permissions
 - Prevents information asymmetry from AI access
 - Maintains existing organizational security boundaries



Risk & Incident Handling

- **Role-Based Access Control (RBAC)**
 - Granular roles
 - SSO integration
 - Multi-factor authentication support
- **Permission Synchronization**
 - Automatic sync from multiple source systems
 - Post-query filtering to prevent unauthorized data exposure
 - Conservative defaults -restrict access on sync failure
- **Incident Response**
 - Comprehensive logging for investigation
 - Rate limiting to prevent abuse
 - Query validation for security protection
- **Responsible Deployment**
 - Clear security best-practice documentation
 - Regular open-source security updates
 - Enterprise-grade support for critical deployments

Thank You

Your Right-Sourcing **IT** Partner.

