



On-Premise Knowledge Agent

Enterprise AI Knowledge Management System

For Banking, Financial Services & Government

On-Premise Knowledge Agent
Enterprise AI Knowledge Management

Data Stays On-Premise | Private Network | System Active

AI Assistant | Retrieval Flow | Analytics | Configuration | Audit & Security

Document Repository

- Financial_Report_Q4.pdf (2.4 MB, 1247 embeddings)
- Policy_Document_2024.docx (856 KB, 634 embeddings)
- Technical_Specifications.pdf (4.1 MB, 2103 embeddings)
- Compliance_Guidelines.pdf (1.8 MB, 892 embeddings)
- Internal_Memo_Jan.docx (423 KB)

AI Knowledge Assistant
Ask questions about your internal documents

What are the key compliance requirements mentioned in the Q4 financial report?

10:34 AM

Based on the Q4 Financial Report (pages 12-15), the key compliance requirements include:

- SOX compliance for financial reporting and internal controls
- GDPR data protection standards for customer information
- ISO 27001 information security management
- Regular quarterly audits by external auditors
- Board-level oversight of risk management

These requirements are being actively monitored with a 98.5% compliance rate as of Q4 2024.

Sources Referenced:

- Financial_Report_Q4.pdf (p.12-15)
- Compliance_Guidelines.pdf (p.8)

94.0% confidence

10:34 AM

Ask a question about your documents... [Send]

247 Total Docs | 52.3K Embeddings | 8.2 GB Indexed

96.8% Accuracy Rate

1.2s Response Time

52.3K Embeddings

AI-Powered Knowledge Intelligence

GPT-4 Turbo • NLP • Vector Embeddings • Machine Learning • Data Analytics

Sector-Specific AI Solutions

Banking & Financial Services | Ministry of Corporate Affairs

Enterprise-grade conversational AI with LLM inference, embedding-based search, and zero data leakage. All processing on-premise for complete data sovereignty and regulatory compliance.

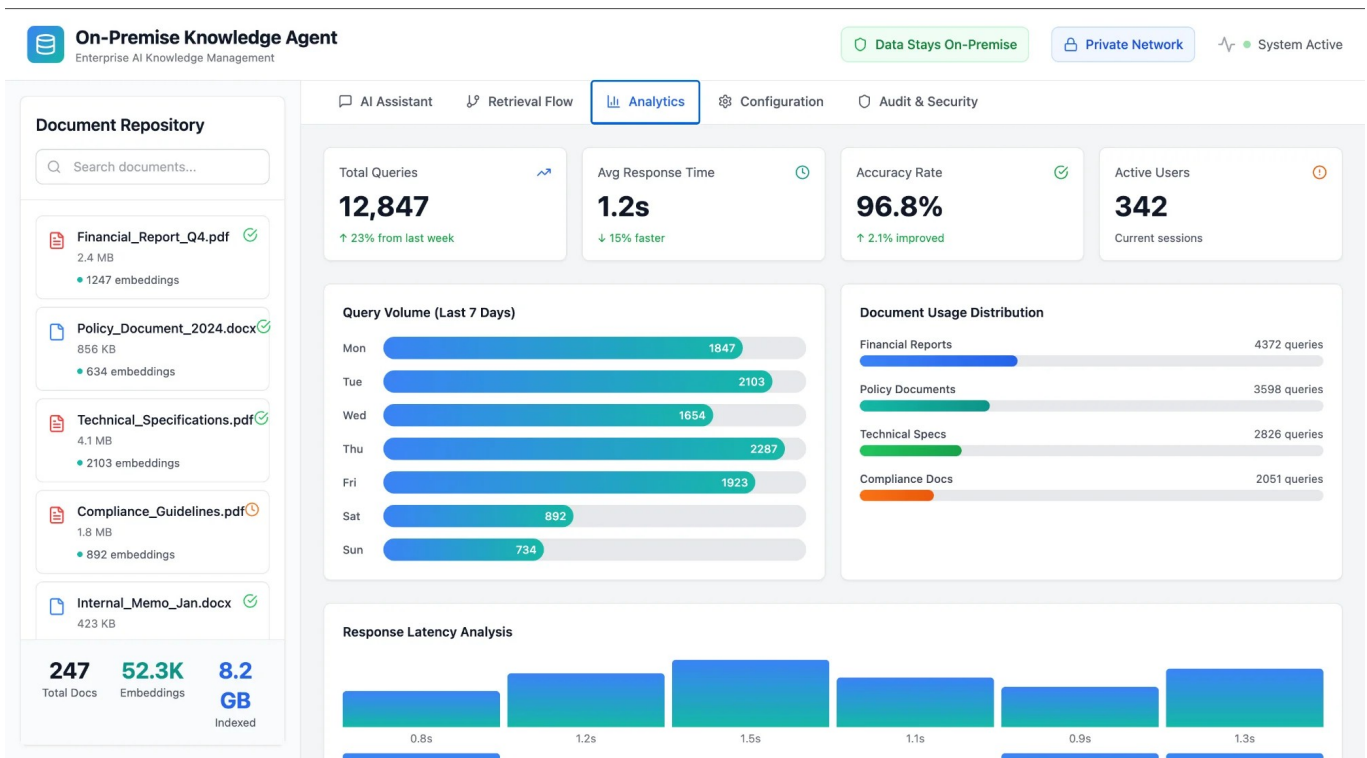
Sector-Specific AI Applications

Banking & Financial Services

- **Compliance Document Retrieval:** Instant AI-powered access to RBI circulars, SEBI guidelines, KYC norms, AML policies, and regulatory frameworks using NLP-based semantic search
- **Financial Policy Intelligence:** Query internal banking policies, loan documentation, credit risk guidelines, and audit reports with contextual understanding
- **Customer Query Resolution:** Staff can instantly retrieve product information, interest rates, service charges, and procedural guidelines to serve customers efficiently
- **Risk Management:** Analyze historical fraud patterns, credit risk assessments, and compliance violations using ML-powered analytics

Ministry of Corporate Affairs

- **Corporate Law Repository:** Instant NLP-based retrieval of Companies Act provisions, SEBI regulations, MCA circulars, and corporate governance guidelines
- **Compliance & Filing Intelligence:** Query MCA-21 filing requirements, ROC forms, director KYC norms, and annual return procedures with semantic understanding
- **Policy Analysis:** Analyze amendments, notifications, and policy changes using ML algorithms to identify impacts on corporate entities
- **Audit & Investigation Support:** Access historical case laws, inspection reports, and enforcement actions with vector-based similarity search



Real-time analytics dashboard: Monitor query volume, response times, accuracy metrics, and document usage patterns with comprehensive business intelligence insights

AI & Analytics Technical Specifications

LLM Engine

GPT-4 Turbo

Embedding Model

text-embedding-ada-002

Vector Dimensions

1536 dimensions

NLP Processing

Tokenization, NER

Search Type

Hybrid Vector+Keyword

Data Capacity

52.3K embeddings

The screenshot shows the 'On-Premise Knowledge Agent' configuration interface. The top navigation bar includes 'AI Assistant', 'Retrieval Flow', 'Analytics', 'Configuration' (highlighted), and 'Audit & Security'. The 'Configuration' section is divided into three main areas:

- Language Model Configuration:** LLM Selection is set to 'GPT-4 Turbo (On-Premise)'. Temperature is 0.7. Max Tokens is 2048. Top-P Sampling is 0.9.
- Embedding Model Configuration:** Embedding Model is 'text-embedding-ada-002 (Local)'. Vector Dimensions are 1536. Retrieval Depth is 5 documents. Similarity Threshold is 0.75.
- Advanced Retrieval Settings:** Hybrid Search (Vector + Keyword) and Re-ranking Enabled are checked. Query Expansion is unchecked. Chunk Size is 512 tokens and Chunk Overlap is 100 tokens.

On the left, a 'Document Repository' sidebar shows a search bar and a list of documents: 'Financial_Report_Q4.pdf' (2.4 MB, 1247 embeddings), 'Policy_Document_2024.docx' (856 KB, 634 embeddings), 'Technical_Specifications.pdf' (4.1 MB, 2103 embeddings), 'Compliance_Guidelines.pdf' (1.8 MB, 892 embeddings), and 'Internal_Memo_Jan.docx' (423 KB). A summary at the bottom of the sidebar indicates 247 Total Docs, 52.3K Embeddings, and 8.2 GB Indexed.

✓ AI & NLP

GPT-4, embeddings

✓ Machine Learning

Re-ranking, similarity

✓ Data Analytics

Real-time dashboards

✓ Enterprise Security

Zero data leakage, RBAC

✓ Make in India

DPIIT recognized startup

Contact Us

<https://carnotresearch.com/>

contact@carnotresearch.com

+91 8449710411

Office 2B-4A, Block B, Research and Innovation Park
IIT Delhi, Hauz Khas, New Delhi

Artificial Intelligence + Big Data Analytics + Enterprise Security