



NICK'S SOFTWARE

Melbourne · 25+ years building custom software

NS DocNavigator

Intelligent Document Retrieval & Knowledge System

Turn thousands of pages of technical documentation into instant, AI-powered answers.

Executive Overview · May 2026

30+ FORMATS

PDF, Office, images,
HTML, web, SharePoint

100% ON-PREMISE

Your documents never leave
your infrastructure

EVERY ANSWER CITED

Page-level traceability
for every claim

The Challenge

Time lost

Engineers spend hours hunting through hundreds of pages of manuals for one specific procedure.

Knowledge silos

Critical information is locked in PDFs, Word docs and legacy formats — not searchable as a whole.

Inconsistent answers

Different people interpret the same documentation differently. No single source of truth.

No traceability

When answers are found, there's no record of which source was used or which version of the doc.

Onboarding gap

New staff take weeks to become productive with complex technical documentation.

The Solution: NS DocNavigator

An on-premise AI system that reads your documents and answers questions — with full source citations.

1

INGEST

Upload any document.

30+ file formats.

Automatic processing
and metadata extraction.

2

UNDERSTAND

AI builds a searchable
knowledge base with
vector + graph search.

Images, diagrams
and schematics captioned.

3

ANSWER

Ask questions in plain
language.

Get accurate answers
with full source citations
and follow-up suggestions.

How It Works

1

Upload Documents

PDF, Word, Excel,
PowerPoint, images,
HTML, and more

2

AI Processing

Text extraction, image
captioning, knowledge
graph building

3

Ask Questions

Natural language chat
with real-time streaming
answers

4

Get Answers + Sources

Every answer cites the
exact document, page,
and passage

Retrieval Intelligence

Eight retrieval and reasoning techniques layered for precision.

Hybrid AI Search

BM25 + semantic + RRF fusion.
BGE cross-encoder reranking.

Knowledge Graph

FalkorDB. Maps relationships between components, procedures, and specs.

Multi-Hop Reasoning

Decomposes complex queries into sub-questions, then synthesises.

CausalRAG

Cause-and-effect reasoning for troubleshooting and root-cause analysis.

Corrective RAG (CRAG)

Relevance check after rerank, self-reformulates and retries on miss.

Contextual Retrieval

Anthropic method — each chunk gets document-level context for accuracy.

Agentic Tool Calling

Model can call external tools mid-answer via MCP, then continue reasoning.

Semantic Chunking

Smarter document splitting that respects topic boundaries.

Document Handling & Workflow

Built for real documents, real teams, and real audit requirements.

30+ File Formats

PDF, DOCX, PPTX, XLSX, HTML, images, legacy Office, MD, CSV.

Image Understanding

Extracts and captions diagrams, schematics, and illustrations.

Web & SharePoint

Recursive crawler with NTLM auth, scheduled refresh, dead-link reports.

Source Citations

Page-level traceability — document, page, and exact passage.

Smart Conversation

Follow-up suggestions, intent-aware routing, full conversation history.

PDF Export

Export Q&A sessions to PDF for audit, hand-off, or sign-off.

Glossary

Domain-specific term resolution injected into every query.

RAGAS Evaluation

Built-in quality measurement — relevance, faithfulness, completeness.

Security & Governance

On-Premise Deployment

Your documents never leave your infrastructure. No cloud data transfer required.

Role-Based Access Control

Three roles (Admin, Editor, Viewer) with granular permissions on every endpoint.

JWT Authentication

Token-based auth with bcrypt password hashing and admin approval workflow.

API Key Masking

Credentials are masked in the UI and API responses — no accidental exposure.

Full Audit Trail

Every answer cites its sources. Real-time log monitoring for administrators.

SSRF & Injection Protection

Built-in safeguards against common web vulnerabilities (OWASP Top 10).

AI Provider Flexibility

Switch AI providers at runtime — no restart, no downtime, no lock-in.

CLOUD AI

Anthropic Claude
OpenAI GPT

Best quality answers.
Pay-per-use pricing.

LOCAL AI

Ollama (Llama, Mistral)
On your GPU

Zero external calls.
Full data privacy.

HYBRID

Mix cloud + local
per function

Cloud for chat, local
for ingestion. Best of both.

Business Value

Faster Answers

Seconds instead of hours
searching through manuals.

Full Traceability

Every answer cites its source
document and page number.

Reduced Onboarding

New staff productive from
day one
with AI-assisted access.

Data Sovereignty

Self-hosted. Your documents
stay
on your infrastructure.

Simple Deployment

```
$ docker compose up
```

One command. The entire system starts.

No External Databases

SQLite, LanceDB and FalkorDB run embedded.
No database servers to manage.

GPU Optional

Works on CPU. GPU accelerates processing
but is not required.

Web-Based UI

No client installation.
Any modern browser works.

Persistent Data

All data on Docker volumes —
survives container rebuilds.

Model Caching

AI models downloaded once and cached.
No re-downloads on updates.

Linux + Docker

Runs on any Linux server with Docker.
NVIDIA GPU recommended for production.



NICK'S SOFTWARE

Melbourne · 25+ years

Ready to see it?

Live demo available on request — invite-only access.

Self-hosted

30+ formats

Manufacturing-ready

Fully traceable

Get in touch

customdb@nicks-software.com

+61 416 563 915

www.nicks-software.com



SCAN TO LEARN MORE

www.nicks-software.com/ns-graphrag.html