DEN AGENTIC AI · TECHNICAL READINESS CHECKLIST

Research, RAG & Knowledge Readiness Checklist

Assess whether your organisation's knowledge infrastructure is ready to support grounded AI — including retrieval-augmented generation (RAG), AI-assisted research, and organisational knowledge systems.

IT & Data Leaders | Knowledge Managers

THREE CONCEPTS THIS CHECKLIST COVERS

AI-Assisted Research

Using AI tools to discover, synthesise, and extract information from internal and external sources — faster and more consistently than manual research.

e.g. Perplexity for market research, NotebookLM for document synthesis, ChatGPT for competitive analysis

RAG (Retrieval-Augmented Generation)

A technique where AI retrieves relevant content from a defined knowledge source before generating its response — so answers are grounded in your actual documents, not general training data.

e.g. AI that answers questions using your company policies, contracts, or product documentation

Organisational Knowledge Systems

The infrastructure for capturing, organising, maintaining, and accessing institutional knowledge — including SOPs, policies, training materials, and operational guides.

e.g. SharePoint, Notion, Confluence — or the AKE (Agentic Knowledge Engine)

1 Document & Knowledge Infrastructure

Score: / 5

- [ ] **Core internal documents exist in a digital, findable format** CRITICAL
  SOPs, policies, procedures, and guides are written down — not just in people's heads — and stored somewhere accessible
- [ ] **Documents are stored in a central location that staff can access** CRITICAL
  Not spread across personal drives, email threads, or scattered SharePoint folders with no clear hierarchy
- [ ] **Documents are version-controlled — old versions are clearly marked or archived** IMPORTANT
  AI retrieving outdated policy versions is a governance and accuracy risk — version clarity is non-negotiable for RAG
- [ ] **Someone owns document maintenance — updates are made when processes change**
  Knowledge that is not maintained becomes a liability when AI retrieves it as current information
- [ ] **Documents are written clearly enough for an AI to extract accurate answers**
  Vague, jargon-heavy, or inconsistently structured documents produce poor RAG results regardless of system quality

2 Data Quality for RAG

Score: / 5

- [ ] **Key terms and definitions are consistent across all documents** CRITICAL
  "Revenue", "customer", "region", "approved" must mean the same thing everywhere a RAG system will search
- [ ] **Documents are regularly reviewed for accuracy — not just created once** IMPORTANT
  A quarterly review cycle minimum — AI retrieving outdated regulatory or product information causes real harm
- [ ] **Sensitive documents are clearly labelled with access classifications**
  Not all knowledge should be equally accessible. Clear classification prevents AI from surfacing confidential content inappropriately

- [ ] **Structured data (tables, databases) has documented field definitions**
  When AI accesses structured data, ambiguous column names or undocumented calculations produce unreliable answers

- [ ] **You can identify which knowledge sources are "trusted" vs "informal"**
  Official policy documents vs Slack conversations vs personal notes — RAG systems must be fed from trusted sources only

3 AI-Assisted Research Readiness

Score: / 5

- [ ] **Staff know the difference between general AI answers and grounded AI answers** CRITICAL
  ChatGPT answering from training data vs a RAG system answering from your documents are fundamentally different in reliability
- [ ] **Research tasks have a defined verification step before conclusions are used** IMPORTANT
  AI research output — from any tool — must be reviewed and fact-checked before being included in reports or decisions
- [ ] **Staff can evaluate the reliability of AI research sources**
  They understand that Perplexity citing a low-quality blog is different from NotebookLM citing your own internal report
- [ ] **Research prompts are structured with clear scope and source constraints**
  e.g. "Summarise only from the uploaded documents" vs open-ended research — staff know which to use for which task
- [ ] **Research outputs are saved and attributed — not lost after the session ends**
  AI research sessions have no memory by default. Important findings should be documented, not relied upon in the same chat thread days later

4 RAG System Readiness

Score: / 5

- [ ] **You can identify the 5–10 knowledge sources the AI must access to be useful** CRITICAL
  Core policies, product guides, pricing documents, FAQs — you know specifically what the AI needs to retrieve to answer real questions
- [ ] **Access controls are in place — not everyone should see everything** CRITICAL
  A RAG system must respect existing organisational permissions. HR documents should not be surfaced to all staff via an AI assistant
- [ ] **A process exists for updating the knowledge base when content changes** IMPORTANT
  When a policy changes, who updates the RAG knowledge source? Without this process, the AI answers from outdated information
- [ ] **You can test and validate that the AI is retrieving from the right sources**
  You have (or will establish) a way to audit what the AI actually retrieved when it generated a response — traceability is a governance requirement
- [ ] **A named person or team is responsible for the knowledge pipeline**
  Without ownership, knowledge bases degrade. The RAG system is only as reliable as the content fed into it

SCORE SUMMARY

/ 5	/ 5	/ 5	/ 5	TOTAL / 20
Knowledge Infrastructure	Data Quality for RAG	Research Readiness	RAG System Readiness	

Score 0 – 8
Not RAG-Ready
Focus on document consolidation, version control, and knowledge ownership before implementing any RAG system. A knowledge consolidation sprint is the right first step.

Score 9 – 14
Partially Ready
You can begin a limited RAG pilot on a specific, well-documented knowledge domain. Identify your most reliable document set and start there. Build governance alongside the pilot.

Score 15 – 20
RAG-Ready
Strong foundations. You are ready to design and deploy a governed RAG system or evaluate the Agentic Knowledge Engine (AKE) for your organisation.

SCORED 15+ AND READY TO BUILD?

Agentic Knowledge Engine (AKE) — Limited Beta

A governed RAG-based intelligence layer that connects LLMs to your organisation's proprietary documents, policies, and structured data — with access controls, source traceability, and human oversight built in.

denagenticai.com
/advisory