



WHITEPAPER

From scattered documents to AI you can trust.

How to turn the knowledge your organization already holds into governed, AI-ready data, and keep control of it.

SDEN WHITEPAPER 2026



sdn.ai

01 / 18

Folders were never enough.

For decades, documents were filed by project, team or date, but the knowledge inside was rarely structured.

Wasted expert time

Duplicate work and rework, slower decisions, and knowledge rebuilt from scratch on every project.

Compliance and risk exposure

Version conflicts, permissions that do not follow the data, and no complete inventory of what AI receives.

AI without grounding

Inconsistent answers, sensitive information surfacing, and responses nobody can trace back to a source.



What AI-ready knowledge means.

Before an assistant or an agent reads a document, four things have to be true about it.

01

Safe

Sensitive content is classified, and access follows the permissions you already have.

02

Relevant

Each set holds only what a given use, team or assistant needs.

03

High quality

Outdated, duplicate and superseded documents are flagged before AI relies on them.

04

Enriched


Every file carries consistent metadata, entities and a link to its source.



A chain of dependency, from documents to confident action.

Enterprise knowledge only becomes usable when every link holds.

- 1 Documents and systems hold enterprise knowledge
- 2 Repositories store them
- 3 Organized for storage, not retrieval
- 4 sDEN Foundation connects them into a governed knowledge graph
- 5 Knowledge becomes consistently discoverable
- 6 Private AI retrieves it with its sources
- 7 Teams act with confidence



Six steps from scattered sources to governed knowledge.

Each step builds on the previous one, and the last keeps the whole chain current.

01 Connect


02 Extract and understand

03 Define your taxonomy

04 Link the knowledge graph

05 Curate and publish

06 Monitor and refresh



Connect, then read everything.

01 Connect

Connect to file shares, document repositories and business systems such as PLM, QMS and ERP, honoring existing permissions and controls. sDEN reads, chunks and normalizes the content.


02 Extract and understand

Read text, tables and drawings, and describe every file with consistent metadata: document type, topics, dates, parts and suppliers, sensitivity and quality signals.



Your taxonomy, linked into one graph.

- 03 **Define your taxonomy**
Start from your own business taxonomy, or let sDEN propose one from your data for your experts to validate.
- 04 **Link the knowledge graph**
Resolve entities across systems and link them with their sources, lineage and access rules.




Publish governed sets. Keep them current.

05 Curate and publish

Slice the knowledge by relevance, topic, time, quality and sensitivity into governed sets for search, RAG and agents.


06 Monitor and refresh

Watch source systems for changes, process new content and keep every set current, so your AI runs on up-to-date, trusted information.




Better data underneath everything your documents touch.

sDEN Foundation connects the systems and documents your teams already depend on into one governed knowledge graph, enriching data where it lives, without replacing repositories or migrating a single folder.




Find by meaning

Discover and filter knowledge by real business entities and relationships, not folders and file names.




Govern at scale

Lineage, classification and existing permissions travel with the data, so the graph never shows more than a user may see.



Ground your AI

Give assistants and agents source-linked context, so every answer cites the path it took.



No rip-and-replace

Your repositories, identities, workflows and controls stay in place. The graph remains when models or vendors change.


Each document becomes data your business can query.

The knowledge that runs your operations is already captured, spread across drawings, dense tables, scans and contracts.

Decisions about assets, quality and investments run on data and dashboards, yet much of what they need stays in documents that cannot be searched, compared or analyzed together.

ILLUSTRATIVE EXAMPLE

ASSEMBLY DRAWING
DWG-4471-210, Rev C



Text

Tables

Drawings

Scans

Spreadsheets

Forms

SCHEMA: DRAWING REGISTER

Drawing number

DWG-4471-210, Rev C

Verified

Title block

Part numbers

P-20418 and 36 more

Verified

Bill of materials

Material

EN AW-6082 T6

Verified

Notes

Tolerances

ISO 2768-mK

Verified

Title block

Latest revision

Rev C, released

Verified

Revision table

Notes and holds

12 notes captured

Human reviewed

Notes column


Every record carries its source location and its review status, so a person can verify any field against the document it came from.

Manufacturing has the data. The challenge is putting it to work.

Document information is fragmented across file shares, PLM and project systems, supplier packages and departmental tools.

MANUFACTURING USE CASES

- Engineering
- Quality
- Supply
- Maintenance
- Enterprise AI



The same foundation, wherever knowledge is the asset.

INSTITUTIONS & RESEARCH

Put institutional knowledge to work, under institutional control.

- Connect institutional knowledge
- Keep sensitive data inside
- Respect roles and permissions
- Build internal capability



PROFESSIONAL SERVICES

Reuse what your firm knows, without exposing what it must protect.

- Reuse project knowledge
- Protect client confidentiality
- Access by practice and matter
- Answers you can cite



Five decisions you keep. Four questions to ask.

Sovereignty applies to the full stack, not only to where a model runs.

01

Where data resides

Where it is stored and in which jurisdiction it is processed.

02

Who operates

Administrators, cloud providers and subcontractors.

03

Which models

Host or fine-tune them, and switch vendors.

04

How AI is governed

Security, audit, access controls, IP and compliance.

05

Whether service continues

Through new prices or terms, outages or foreign law.

Territorial

Are data and compute in an approved jurisdiction?

Operational

Can we, or a trusted local provider, administer, secure and keep it running?

Technological

Do we control models, code, interfaces and IP, or can we replace suppliers?


Legal

Which laws apply, and can a foreign authority gain access or interrupt service?

Sources: IBM (decisions), McKinsey (dimensions)

Retrieve, contextualize, act. Through one gateway.

Every stage passes through the same controls: policy, private inference and audit.



Illustrative architecture. Designed to your requirements, not a live product interface.

Seven domains to control.

01

Data

Use, flows, retention, third-party training.

02

Identity and access

Least privilege, SSO and MFA, secrets, logging.

03

Model

Replace providers; protect weights and fine-tuning.

04

Infrastructure

Clouds, regions, backups, subprocessors.

05

Reversibility

Export data, prompts, agents, evaluations, logs.

06

Governance

Named owner, approvals, oversight, incidents.

07

Supplier resilience

A plan for new prices, terms, API cuts or outages.

The goal is not to own the whole stack, but to keep the ability to control, audit, switch and continue.

Pay for your capacity. Not each API call.

Hosted AI bills every request. Private AI runs on capacity you size.

\$0

external model API fees

FOR THIS FULLY LOCAL DEPLOYMENT

Hosted AI

Usage-based API fees

 Your teams

 Provider API



Every request adds to the bill

Private AI

Shared private capacity

 Your teams

 Private capacity



Compute + setup + operations

Seven questions. A sovereign organization answers yes to most.

- | | | | | | |
|----|--|--------------------------|----|--|--------------------------|
| 01 | Do we know what data our AI tools receive, and where it is stored? | ☐ | 02 | Are confidential data excluded from vendor training unless authorized? | ☐ |
| 03 | Can we export our data, prompts, workflows and model assets? | ☐ | 04 | Are security, subprocessors, data access and applicable law clear in contract? | ☐ |
| 05 | Can we switch model or hosting provider without rebuilding from zero? | ☐ | 06 | Do we have oversight, logging, access rights and an AI governance policy? | ☐ |
| 07 | Can sensitive workloads run in-jurisdiction, in private cloud or on premise? | ☐ | | | |

Ready to build AI you can govern, audit and own?

Discuss your data estate, institutional priorities and infrastructure requirements.

THE STRATEGIC IMPERATIVE



Authority over data

Define where information resides, who can access it, and how it is used by AI.



Freedom to adapt

Evaluate model, platform, and provider choices against your long-term requirements.



Capacity to operate

Give internal teams the knowledge, processes, and visibility to own the system.

YOUR CONTACTS AT sDEN



Sami Dzogang

Founder & CEO



John Constantine

Head of Sales

Request a Demo

sden.ai/demo