THE DMS, COMPLIANCE &
AI READINESS CHECK 2026

7 checks before scaling AI across legal work

Includes a practical DMS, Compliance & AI checklist

epona

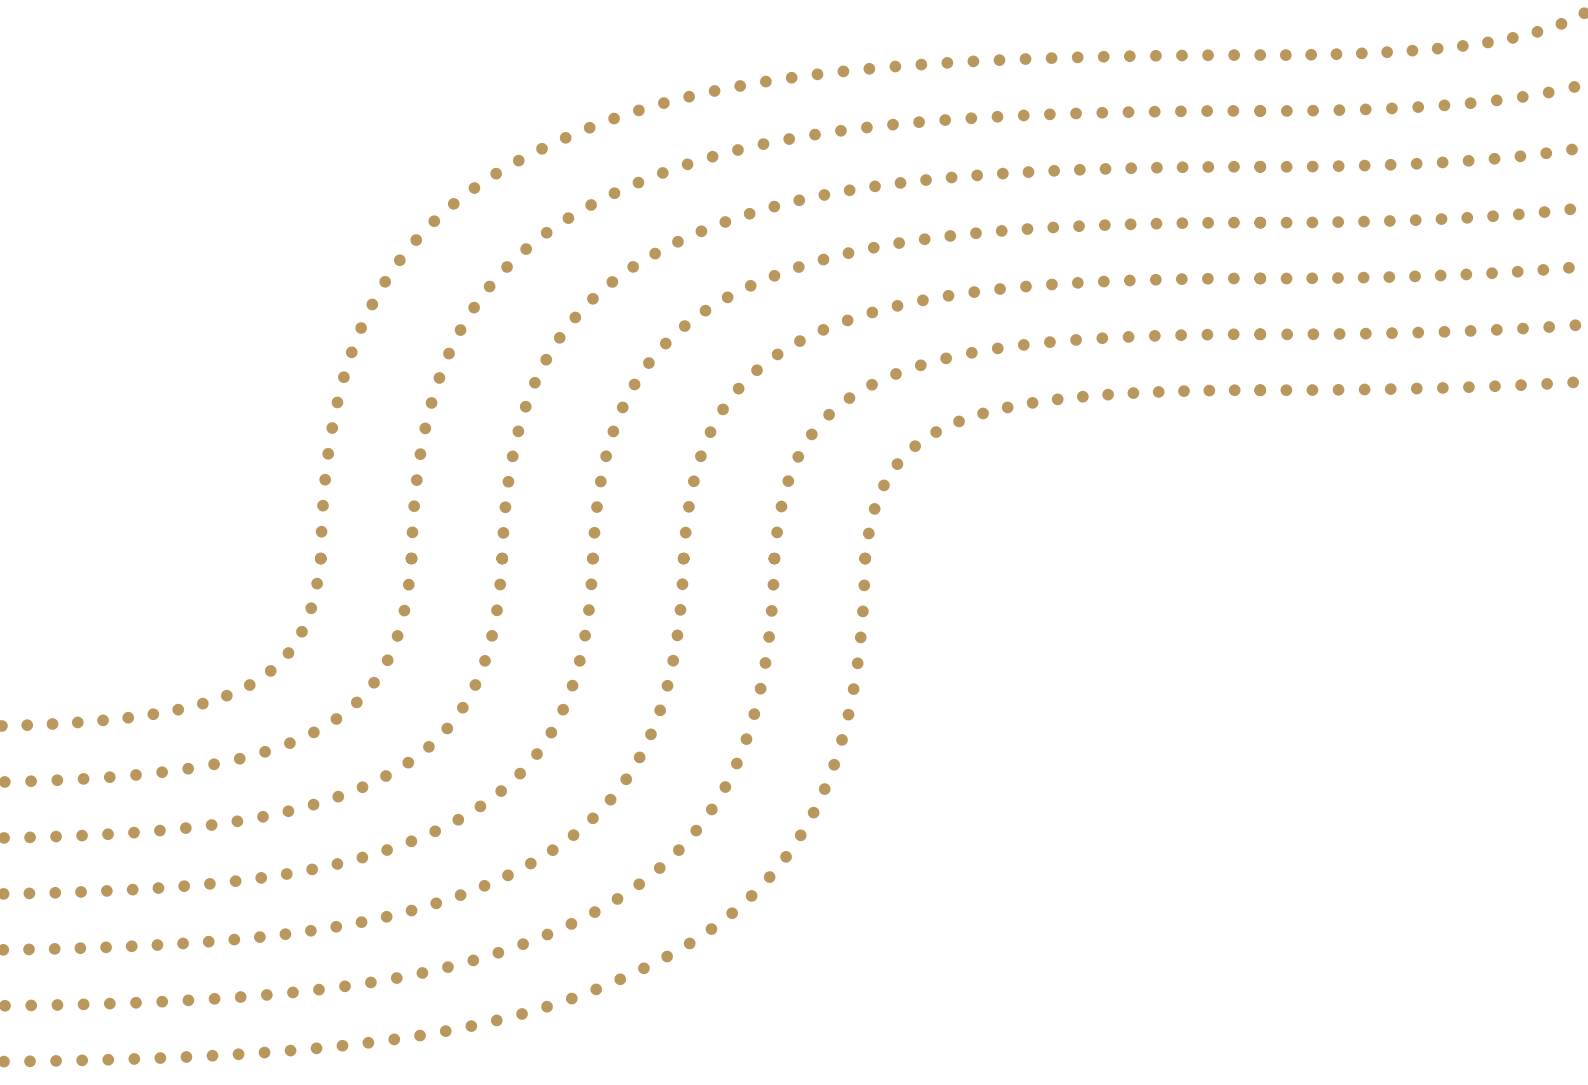
AI practical guide for legal professionals

AI in legal is moving into the systems where legal work happens every day: documents, emails, matters, templates, knowledge and workflows.

That makes one question more urgent than ever:

Is your DMS ready to support AI without increasing compliance risk?

This practical guide gives legal leaders, IT teams and compliance stakeholders seven practical checks to assess whether their document environment can support AI safely, defensibly and at scale.



Introduction:
AI makes the DMS a compliance issue

Document Management Systems were usually judged mostly on usability. Can lawyers file emails quickly? Can they retrieve the right version? Does the system work inside Outlook? Is the search good enough? Those questions still matter, but we can tell you: in 2026, they are no longer enough.

As AI becomes part of legal work, the DMS becomes part of the firm's risk architecture. AI does not operate in a vacuum. It draws on documents, emails, matter history, metadata, permissions and knowledge sources. If that environment is structured, governed and secure, AI can become a serious productivity layer. If that environment is messy, overexposed or poorly controlled, AI can amplify the problem.

This is the uncomfortable reality for law firms and legal departments: AI risk is often not caused by the AI tool itself. It starts with:

- Access rights
- Oversharing
- Poor matter separation
- Weak metadata
- Unclear data flows
- Systems that were never designed for AI scale

That is why DMS, compliance and AI now belong in the same conversation.

In this guide, you will find seven practical checks to assess whether your DMS is ready for secure AI adoption. You will learn where confidentiality risks often appear and which governance questions to ask before scaling AI. Next to that, we will talk some more about how to connect permissions, metadata, auditability and daily legal workflows into one clearer operating model.
At the end, you will also find the DMS, Compliance & AI Checklist, a practical one-page tool to use in a DMS selection meeting, AI governance workshop or internal compliance review.

Team Epona

Check 1:
Review permissions before Copilot or agents do

The first question is simple: who can see what? In many legal environments, permissions have grown organically over years. Teams were added to matters, folders were copied, external sharing links were created, inheritance was left untouched and access exceptions became normal. People usually notice this only when a client asks difficult questions, or when AI starts surfacing information that was technically accessible but practically invisible before.

When users can ask natural language questions across content, hidden permission problems become visible. A document that nobody searched for yesterday may become an answer tomorrow.

Your first DMS compliance check should therefore focus on oversharing.
Review matter workspaces, SharePoint sites, Teams channels, legacy libraries and “everyone” links. Identify places where access is broader than the matter requires. Pay extra attention to sensitive work such as M&A, investigations, employment matters, public sector work, financial institutions and litigation strategy.

A good AI-ready DMS does not only store documents. It enforces matter-level separation, supports clear access models and makes it possible to explain why a user can or cannot access certain information.



Check 2:
Validate whether confidentiality follows the document

Legal confidentiality is not a folder setting. It must follow the document through daily work.

That means labels, classification, encryption and access controls need to be connected. If a confidential document is downloaded, shared, emailed or used in an AI-enabled workflow, the governance model should still make sense.

Sensitivity labels are useful only when people understand when to apply them, when they are automated where possible and when they connect to real policy. The goal is not to create more administrative work. The goal is to make confidentiality visible and enforceable.

For legal teams, this matters because documents rarely stay still. They move between Outlook, Word, Teams, SharePoint, client portals and matter workspaces. AI adds another layer, because content may be summarised, compared, classified or used to support a workflow.

Ask this: if a document is client confidential inside the DMS, does that status remain meaningful across the wider Microsoft 365 environment?



Check 3:
How about your data boundaries

One of the biggest compliance questions around legal AI is also one of the most practical: where does the data go?

A DMS selection process often focuses on features, user experience and search. But the deeper question is control. Under whose legal reach does your data fall? Who controls the environment? Where does AI processing happen? Is data duplicated across multiple systems? Can you explain the architecture to clients and auditors?

This is especially important when AI is added to a platform. A system may be hosted in Europe, but that does not automatically answer every question about legal control, vendor dependency or data access. Legal teams increasingly need to understand not only where content is stored, but also which platform layer handles identity, permissions, retention, audit and AI interaction.

A Microsoft-native approach can simplify that discussion for organisations already working in Microsoft 365. Documents, identity, permissions, retention, audit and compliance controls can remain part of the same governed environment. Epona builds on that foundation by adding the legal structure around matters, documents and email workflows.



Check 4:
Is the auditability good enough for legal work

AI adoption becomes much more serious when someone asks: how was this output created?

Legal teams need to know which sources were used, which permissions applied, which user initiated the action and whether sensitive information was involved. Without traceability, speed can become a liability.

This is why the DMS is evolving from storage into part of a provenance layer. It should help preserve relationships between matters, documents, emails, queries, sources and outputs. For legal teams, that is not a technical luxury. It supports defensibility.

Auditability should cover several questions.

- Who accessed the document?
- Was it shared externally?
- Which version was used?
- Was the document classified?
- Did the action happen inside the approved environment?
- Can the firm reconstruct what happened if a client, auditor or regulator asks?

"Without governance and architecture, AI amplifies risk instead of reducing it."

- Bart Bogaerts, Partner Manager at Epona



Check 5:
Research metadata before relying on AI output

AI needs context. In legal, context means more than file content. It includes matter, client, jurisdiction, document type, version, status, confidentiality level, author, date and relationship to other documents.

Poor metadata creates weak AI. It also creates risk. If the system cannot distinguish between draft and final, internal and external, current and outdated, confidential and general, it becomes harder to trust what AI produces. This is where DMS maturity becomes directly connected to AI readiness. Metadata should not be treated as an administrative burden. It becomes part of the intelligence layer.

Modern legal platforms increasingly use automatic classification and metadata suggestions to reduce manual effort. That is important because lawyers want to work. The system should make good behaviour easy.



Check 6:
Assess your policy for prompt and agent behaviour

Most firms have some form of AI policy. Far fewer have a practical operating model for what users can do with legal content.

A useful policy should answer concrete questions. Which documents may be used with approved AI tools? Which content must never leave the governed environment? When does a user need approval? How should outputs be checked? Which sources should AI trust first? What happens when an agent acts across multiple systems?

This becomes more important as AI moves from prompting to agents. A prompt is usually a one-off instruction. An agent follows a repeatable workflow. It may retrieve documents, consult a knowledge base, apply a template, prepare a draft and route something for review. That is powerful, but it requires guardrails. The goal is to give people a safe route.



Check 7:
Review whether the DMS fits the way lawyers work

The strongest compliance model is the one people use. If filing is slow, lawyers work around it. If permissions are confusing, teams create shortcuts. If metadata takes too long, people skip it. If the DMS sits outside Outlook and Word, adoption drops.

This is why usability and compliance cannot be separated. A legal DMS must support the daily workflow: email filing, matter access, document creation, co-authoring, search, version control and mobile work. When that happens inside the Microsoft environment lawyers already use, governance becomes less intrusive.

Hidde Reitsma, lawyer and partner AMS Advocaten captured this well:

"The most intuitive way of working inside the Office environment remains the backbone. That may not sound exciting, but it is the real strength of the product."

That quote works because it gets to the heart of AI readiness. The best governance model is not the one that looks strongest in a policy document. It is the one that becomes part of the way work already happens.

The DMS, Compliance & AI Checklist

AI makes existing risks visible faster. A permission that is too broad, an old sharing link, an unclear data structure or a document that falls outside the right governance model can remain almost unnoticed for years. Once AI can summarise, combine and retrieve documents through natural language, that same mess suddenly becomes an operational risk.

This checklist helps legal, IT and compliance teams make the conversation concrete before AI is scaled further. The questions are not about the most polished demo or the smartest AI feature, but about control, explainability and trust. That is exactly where the difference lies between a DMS that only stores documents and a DMS that can safely move with AI.

Check	Question	Why this matters
☐ 1. Legal reach	Where does our legal data really live, and under whose legal reach does it fall?	Hosting location alone is not enough. A firm must understand which legal framework can affect access, control and disclosure.
☐ 2. Data control	Who controls the environment where our documents, emails and matter data live?	The more control sits outside your own governed environment, the harder it becomes to explain risk to clients, auditors and internal stakeholders.
☐ 3. AI processing	Where does AI processing take place, and does client data stay inside approved boundaries?	AI value depends on access to content, but legal teams need confidence that sensitive data is not moving through uncontrolled layers.
☐ 4. Permissions	Are permissions enforced consistently across matters, Teams, SharePoint and AI workflows?	AI will respect what the environment allows. If access is too broad, AI may surface information users should never practically see.
☐ 5. Oversharing	Do we know where broad links, inherited rights or “everyone” access still exist?	Oversharing is one of the most common hidden risks in AI adoption. It often starts as convenience and becomes a compliance issue later.
☐ 6. Confidentiality	Can confidential documents be labelled, protected and audited across daily workflows?	Confidentiality must follow the document through Outlook, Word, Teams, SharePoint and AI-supported work, not only inside one folder.
☐ 7. Governance model	Do we operate one consistent governance model, or several disconnected layers?	Fragmented governance creates legal fog. A consistent model makes policies easier to apply, monitor and defend.
☐ 8. Audit readiness	Can we explain our data flows, architecture and vendor choices to clients and auditors?	Compliance is not only about being secure. It is also about being able to prove why the architecture is defensible.
☐ 9. Vendor dependency	How dependent are we on the vendor’s roadmap, data structure and AI architecture?	A DMS decision is also a long-term platform decision. Dependency becomes more important once AI, governance and workflows are built on top.
☐ 10. Exit strategy	Can we move our data cleanly if our platform strategy changes?	Future flexibility matters. If leaving the platform is difficult, the firm may be accepting more strategic risk than it realises.

How to use this checklist
Use these questions in a DMS selection meeting, AI governance workshop or internal compliance review. Score each item as clear, unclear or risk area. The goal is not to solve everything in one session. The goal is to identify where the organisation needs a stronger answer before AI becomes part of everyday legal work.

A practical rule of thumb: if you cannot explain the answer simply to a client, auditor or managing partner, the architecture probably needs more attention before AI is scaled further.

What Epona adds

Epona365 is built on Microsoft 365 and SharePoint for law firms and corporate legal departments. It brings document and email management into the Microsoft environment, while adding the matter-centric structure, legal workflows, metadata logic and user experience that legal work requires.

That matters a lot regarding AI. A legal AI strategy is only as strong as the environment behind it. When documents, emails, matters, permissions and governance are structured inside the systems your organisation already controls, AI can be introduced with more confidence.

Epona helps legal teams create that foundation: practical document management, secure matter-based working, Microsoft-native governance and a platform that can support future AI-enabled workflows.

Read more on the Epona365 DMS

Ready for the next step?

Want to assess whether your DMS is ready for secure AI adoption? Book a 30-minute Epona demo and see how a Microsoft-native legal document environment can support compliance, confidentiality and AI readiness.

Contact us



AI Risk Heatmap: Where confidentiality risk usually starts

Low risk:	Structured matters, clear permissions, labels, audit trail
Medium risk:	Inconsistent metadata, manual filing, legacy folders
High risk:	Oversharing links, broad Teams access, unclear external sharing
Critical risk:	Sensitive client data used outside governed environments

Visit our website: epona.com

THE DMS, COMPLIANCE & AI READINESS CHECK 2026

7 checks before scaling AI across legal work

Includes a practical DMS, Compliance & AI checklist

epona

Empowering Legal Efficiency