Moving from Readiness to Responsible Implementation

AIIM 2026 WORKSHOP

AI+IM GLOBAL SUMMIT

opentext™

Moving from readiness to responsible implementation

AIIM 2026 Workshop

YOUR FACILITATORS

Alan Pelz-Sharpe
Founder & Principal Analyst, Deep Analysis
Author of the white paper underlying this workshop

Tracy Caughell
OpenText
Sr. Director, Product Management, OpenText Content Cloud

Workshop agenda

Introductions

Why This Conversation, Why Now
Why pilots stall · The IM discipline as competitive advantage

Key Concepts
White paper highlights · The data challenge

Workshop 1: Target with precision

Workshop 2: Understand the data

Workshop 3: Build in protection from the start

Your Action Plan
90-Day roadmap · How to get a seat at the table

Takeaways & Next Steps
Key points recap · Resources to take home

Show of hands….

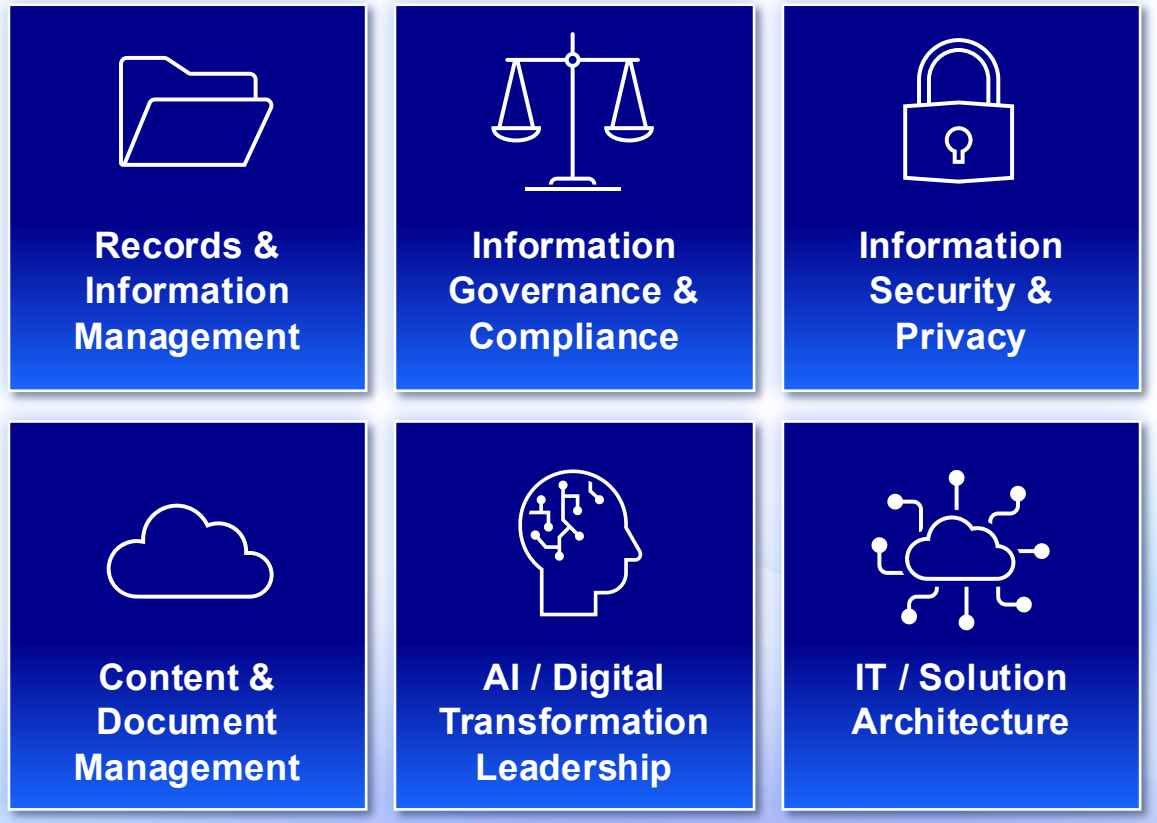

Which best describes your primary focus area?

True/False – AI for Enterprise Content…

1. The more information I give to the AI the better the answers will be

- *False. AI can provide better answers when focused on the right content!*
- *The likeliness for hallucination and distraction raises with the amount of content that is unrelated to the ask.*

2. AI will bring order in my information chaos

- *False. AI does not overcome the "garbage in garbage out" principle – it just amplifies it!*
- *Properly structured and governed information is not just good for humans but also for AI!*

3. Let's just connect all information repositories to a central super AI

- *False. A central AI does not understand how information is structured and which permissions and controls apply.*
- *The central AI will be crawling information, create performance problems and still will never be up-to-date in real-time.*
- *You must bring AI to content – not content to AI!*
- *Let the agents work together and specialize in ecosystems*

Never have I ever…

…seen AI expose content that it shouldn’t

Never have I ever …seen AI expose content that it shouldn't

Tips: Curated Content, Information management, **Security**

- "Good AI requires Good Content Management"
- Ensure users are unable to access content using AI that they wouldn't otherwise have access to.
- Utilize your robust content management security and access controls
- Lifecycle management will increasingly become important to make sure only the most up to date information is provided

Never have I ever…

…used AI for a critical decision

Never have I ever … used AI for a critical decision

- First determine if this is in fact a critical decision
  - Do you have boundaries for what decisions AI can make by itself and what requires HITL?
- Use AI to inform, but not make critical decisions
- Important skill going forward, knowing when to use AI
  - Improve clarity, speed or coverage
  - Avoid when empathy, legal interpretation, or nuanced judgement needed

Never have I ever…

…seen a really bad AI hallucination

“AI ruined my weekend”

Never have I ever … seen a really bad AI hallucination

- Use AI in support of work, not as an authoritative source
- Use curated data because ‘garbage in leads to garbage out’
- Hallucinations are particularly noticeable:
  - Misinterprets regulations or procedures
  - Misleading summaries
  - Fabricated details

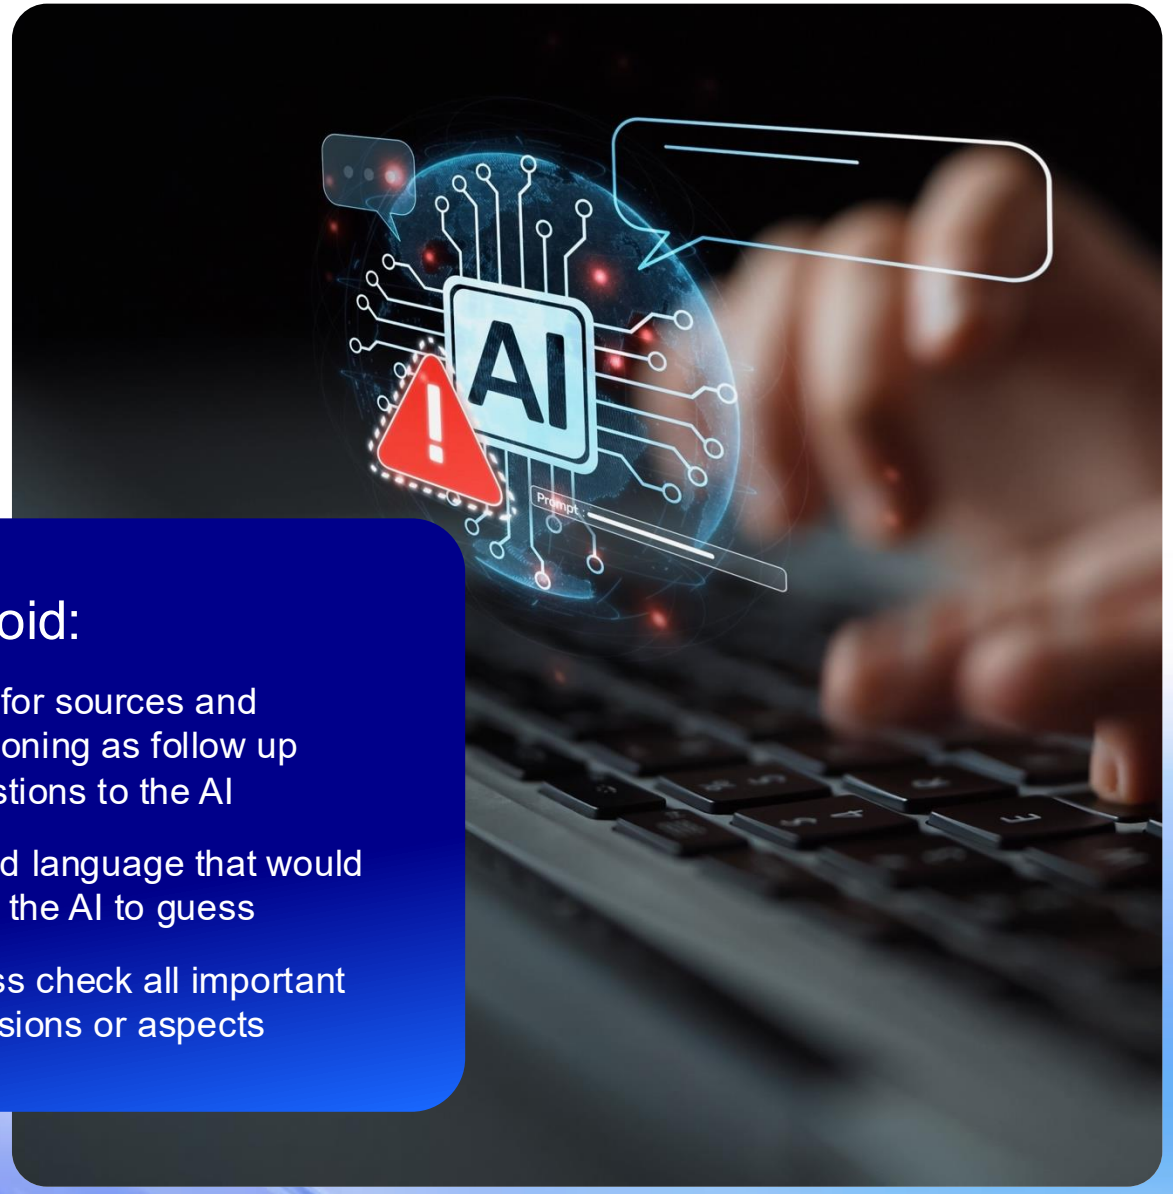

To avoid:

- Ask for sources and reasoning as follow up questions to the AI
- Avoid language that would lead the AI to guess
- Cross check all important decisions or aspects

What you'll get from this workshop

A clear framework

Three practical steps to move from planning to action: target precisely, understand your data, build in protection.

Hands-on practice

Four workshop segments where you apply concepts to your own organization's challenges and use cases.

Reusable tools

Worksheets, checklists, and a 90-day action plan you can put to work on Monday morning.

A seat at the table

The language, evidence, and strategy to position IM expertise as essential to every AI initiative in your organization.

Why this conversation? Why now?

"The era of general AI readiness planning is over — it's time for focused action."

Pilots are stalling

Organizations have been "getting ready" for 4 years. Yet most AI projects never leave the proof-of-concept stage — or fail at scale.

The data gap is the culprit

The #1 barrier isn't the AI model — it's the unstructured data that feeds it. Documents, emails, manuals: uncontrolled, unvetted, ungoverned.

Information management is the missing piece

Information managers hold the keys. Data quality, governance, and lifecycle management are prerequisites for AI success — not afterthoughts.

Information managers are instrumental to AI success

The discipline of IM is just as important as the AI technology itself.

What AI Needs to Succeed

- Accurate, current, trusted source content
- Structured understanding of business data models and lifecycles
- Clear states: draft vs. approved vs. archived
- Knowledge of what is redundant, obsolete, or transitory (ROT)
- Governed access — right data to right audience
- Provenance: traceable, citable sources

What IM Professionals Bring

- Classification taxonomies and metadata standards
- Content lifecycle governance frameworks
- Records retention and disposition expertise
- Compliance, privacy, and security knowledge
- Institutional knowledge of where data lives and why
- Policies that translate to AI guardrails

4 years of AI evolution — where are we now?

2022	2023	2024	2025	2026→
LLM Discovery	**Productivity assistants**	**Agentic workflows**	**Readiness fatigue**	**Time to act**
ChatGPT moment. Organizations start asking: 'What is this?'	Copilot, assistants. 'How do we use this at work?'	Embedding AI into operations. Pilots launched across enterprises.	Pilots stall. Data gaps exposed. 'Why isn't this working?'	The readiness phase is over. Responsible activation begins.

The shift from chat → copilot → agentic has one common thread: it all depends on your data.

Three steps from readiness to responsible implementation

The answer is not more planning — it's intelligent activation of the right data.

01

Target precisely

Identify and select only the unstructured data collections essential to your priority AI initiatives. You don't need all your data — just the right subsets.

- Focus on specific, bounded collections
- Map data to use cases first
- Resist the 'boil the ocean' instinct

02

Understand the data

Assess the suitability of individual items and entire collections. Understand shape, lifecycle, volatility, and state before activating anything.

- Shape: content, length, format
- Lifecycle: creation → use → archival
- Volatility & growth rate matter

03

Build in protection

Design for veracity, privacy, and security from day one. These three filters must be built into the data pipeline before AI ever touches the content.

- Veracity: state-aware, ROT-free retrieval
- Privacy: PII classification & redaction
- Security: RBAC, audit trails, guardrails

Unstructured data: Shape vs. lifecycle

"Unstructured" does not mean "uniform." Each collection differs in format, content, growth, and management needs.

SHAPE · What it looks like

Content type	Policy doc, email, manual, contract, schematic
Format	PDF, Word, HTML, image, video, structured XML
Length	500-word KB article vs. 10,000-word operating manual
Headings & structure	Internal consistency of a collection

LIFECYCLE · How it lives & dies

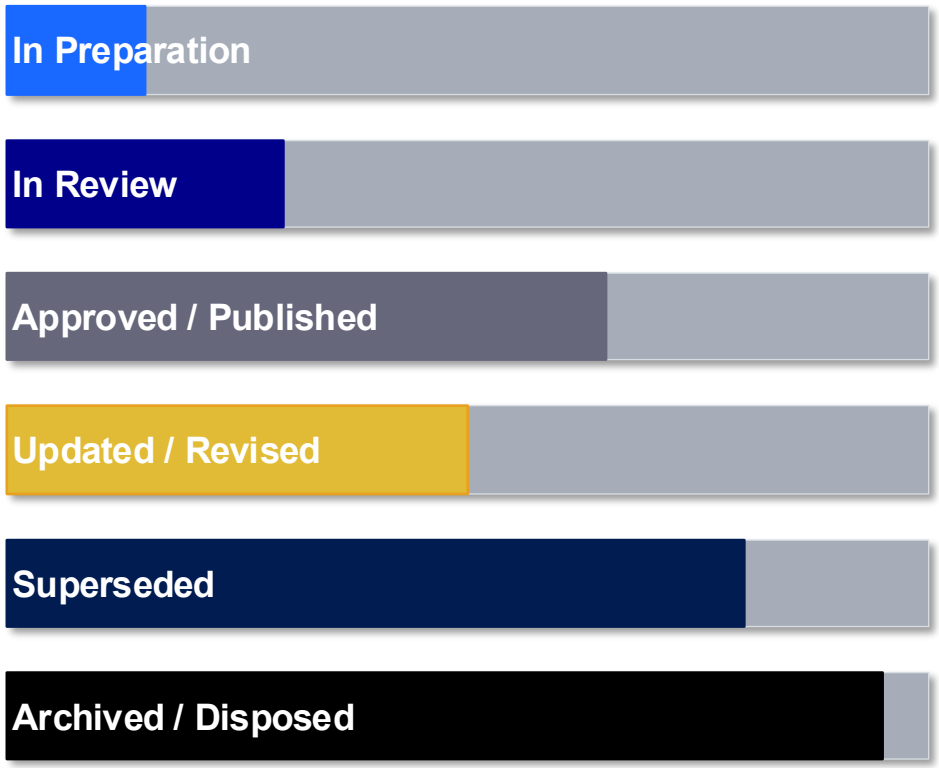

% of typical knowledge base in this state

The three protective gates: Veracity · Privacy · Security

Protection is not a post-deployment add-on. These three gates must be designed in from day one.

VERACITY

Trust in what AI knows

- State-aware retrieval: approved/published only
- Exclude drafts, archived & deprecated content
- Provenance & citation: every output traceable
- ROT awareness: flag redundant, obsolete, transitory data

Greatest risk is not technical failure but propagation of misinformation.

PRIVACY

Control over sensitive data

- Automated PII classification & redaction
- Audience-based data segregation (internal vs. external)
- Consent & purpose limitation mapped to legal basis
- "Clean" versions created for AI contexts

Broadly ingesting unstructured data without controls creates massive compliance and liability risk.

SECURITY

Protection of the pipeline

- Role-based access control (RBAC) for AI agents
- Input/output sanitization and monitoring
- Immutable audit trails: prompt → data → response
- Guard against prompt injection & data poisoning

Activating data for AI creates new attack surfaces and high-value extraction targets.

Context is everything

Context without governance = Risk

AI systems need to know not just the content of a document — but who created it, when, for what purpose, and for which audience. Without a modern CMS, that context simply doesn't exist.

Fragmented data = Failed AI

When content lives in legacy file shares, disconnected repositories, and email archives, AI can't reliably access or assess it. Modernization connects the content layer to the intelligence layer.

OpenText Content Cloud

Provides intelligent content management from capture and IDP through information governance and data archiving — fueling AI-driven work at scale with built-in governance.

Great GenAI requires great content management

- Unified metadata across all content types
- Automated classification and tagging
- Lifecycle state management at scale
- Version control and provenance tracking
- Fine-grained access controls & audit logs
- Retention and disposition automation
- Integration with AI and RAG pipelines

Automating curation of context assures the highest accuracy and most useful GenAI responses

Target with precision

Workshop 1

✦ WORKSHOP ✦

Part 1: Define the right use case
Precision before technology

15 min

Discussion: Work in groups of 3–4. Use the Handout #1 worksheet provided.

Debrief: Each group shares their use case and #1 data gap with the room. (2 min per group)

5 min

Pick your use case

Each group selects ONE specific AI use case from your organization (or a hypothetical one from the list on your handout). Be specific — 'AI for customer service' is too broad. 'AI to answer product warranty questions from the support knowledge base' is specific.

5 min

Map the data required

For that use case: what specific collections of unstructured data would the AI need? Name them. Where do they live today? Who owns them? Use your handout to capture this.

5 min

Identify the gaps

What is missing, unknown, or ungoverned about that data? What would have to be true for it to be AI-ready? Capture 2–3 specific barriers on your handout.

HANDOUT / WORKSHEET

Handout #1: Use case definition worksheet

GROUP MEMBERS: _____

STEP 1: Define your AI use case

Use case (be specific — what will AI do, for whom, using what data?):

Use case category:

- ☐ Customer service / support
- ☐ HR / internal helpdesk
- ☐ Contracts / legal review
- ☐ Knowledge management
- ☐ Compliance / audit
- ☐ Other: _____

STEP 2: Map the required data collections

Data Collection Name	Where It Lives Today	Owner / Steward	Estimated Size

STEP 3: Top 2 data barriers / gaps

1. _____

Step 2: Understand the terrain

Workshop

✦ WORKSHOP ✦

Part 2: Activate the right data
Not all data

15 min

Discussion: Continue in the same groups. Use Handout #2.

Debrief: Share one insight per group — what's the most surprising thing you discovered about your data?

Shape & lifecycle

- For your primary data collection: what is its shape? (format, typical length, heading structure)
- What is its lifecycle? Who creates it, who reviews it, how is it updated, how is it retired?
- Is it well-managed in a system today, or scattered across file shares/email?

Volatility & growth

- How often is this collection updated? (daily, weekly, monthly, rarely?)
- Is it growing fast? Is growth driven by new content or accumulation of outdated content?
- High volatility + high growth = strong candidate for AI; also highest maintenance requirement.

The three gates

- VERACITY: Which documents could be outdated, draft, or ROT? What policy would exclude them?
- PRIVACY: Does this collection contain PII, confidential info, or privileged content?
- SECURITY: Who should this AI be able to serve? Internal only? Customers? Both?

HANDOUT / WORKSHEET

Handout #2: Data activation assessment

Data Collection Being Assessed: _____ Use Case: _____

SHAPE & LIFECYCLE

Format(s):

Created by / Reviewed by:

How is it retired / archived?

VOLATILITY & GROWTH

Update frequency:

☐ Daily ☐ Weekly ☐ Monthly ☐ Rarely

Estimated # of documents / items:

AI suitability (circle): High Medium Low

THE THREE PROTECTIVE GATES

VERACITY

Which documents could be outdated or ROT?

What policy would exclude them?

PRIVACY

Does this collection contain PII or confidential content?

How would you classify & redact?

SECURITY

Who should this AI serve? (internal / external / both)

What access controls are needed?

Step 3: Build in protection

Workshop

✦ WORKSHOP ✦

Part 3: Embed AI where work happens
OpenText leads

10 min

AI creates value only when embedded in the workflows where decisions are made.

Content capture & IDP

Intelligent document processing at ingestion — classify, extract, enrich, and route content automatically as it enters your organization.

Governance & compliance

Automated retention decisions, disposition triggers, and compliance monitoring powered by AI that understands the content lifecycle.

Search & discovery

AI-augmented enterprise search that surfaces the right content in the right context — not just keyword matching but semantic understanding.

Analytics & insights

AI that identifies patterns across unstructured data — surfacing risk signals, customer trends, and operational intelligence from your content.

Workflow automation

AI agents embedded in approval chains, review workflows, and case management — reducing manual handling of routine information tasks.

Discussion: Where in your organization's workflows would AI-embedded content management have the most immediate impact?

Integrated AI is a secure and sustainable AI

Unlock ready-to-use AI without the risk of a fragmented infrastructure

Uncompromised security

- Follows existing security policy with no extra management
- Honors existing permissions through in-place indexing

Deep business grounding

- Native metadata and knowledge graph semantics for superior AI outcomes
- Understand document context tied to your core business processes

A connected AI-to-AI ecosystem

- Bring intelligent content assistance to other AI systems
- Collaborate seamlessly with:
  - Microsoft Copilot
  - Salesforce Agentforce
  - Gemini Enterprise
  - SAP Joule (coming soon)
- Cross-system tasks, agents, and blended grounding while maintaining strict data ownership

Always current data

- Data remains fresh without extra effort – answers are real-time based on the ground-truth in your content repository
- A full-stack AI deployed how you need it included with Content Aviator

Seamless, native user experience

- Keep your teams working within the familiar UI of choice, including Microsoft 365 applications
- Put AI at their fingertips without switching applications

Responsible activation: 90-day plan

✦ WORKSHOP ✦

Part 4: Your 90-day action plan
Getting a seat at the table

10 min

Individual exercise: Use Handout #3. Then share one commitment with the room.

Weeks 1–4	Weeks 5–8	Weeks 9–12
CONVENE & sELECT	**ANALYZE & dESIGN**	**PROTOTYPE & gOVERN**
• Identify the owners of your top AI use case • Map the specific bounded data collections involved • Hold a 1-hour 'Data Activation Kickoff' meeting • Deliverable: One-page Data Activation Charter	• Profile selected collections: shape, lifecycle, volatility • Run a veracity/privacy/security assessment (Handout #2) • Work with IT to design the retrieval pipeline • Deliverable: Technical design doc for secure AI data feed	• Build minimal secure pipeline to a test environment • Draft AI Data Supply Chain Policy • Draft Context & Audience Policy for AI access • Deliverable: Working prototype + 2 draft policies

Getting a seat at the table: Lead with business value. Quantify the cost of AI failure from bad data. Own the governance solution.

HANDOUT / WORKSHEET

Handout #3: My Personal 90-Day Action Plan

Weeks 1–4: CONVENE & SELECT

The AI use case I will focus on:

The data collection I will assess first:

The person I need to bring to the table:

My specific action this week:

Weeks 5–8: ANALYZE & DESIGN

Volatility / growth rate of my target collection:

Top veracity concern:

Top privacy concern:

Who I need in the room for the design session:

Weeks 9–12: PROTOTYPE & GOVERN

The policy we are missing today:

Who needs to approve that policy:

What success looks like at 90 days:

How I'll demonstrate IM value to AI leadership:

My one public commitment from today's workshop: _____

HANDOUT / WORKSHEET

Handout #4: Key Questions by Common AI Use Case

Use this as a reference when evaluating AI use cases in your organization.

Customer-Facing AI / Chatbots

- Is the knowledge base current and approved?
- Are there PII risks in support ticket histories?
- Who governs what the AI is allowed to say?
- How are outdated product docs excluded?

HR / Internal Helpdesk AI

- Are HR policies version-controlled and approved?
- What employee data could be exposed in RAG?
- How is RBAC configured — manager vs. employee?
- How do you handle multi-jurisdiction policy?

Contract / Legal Review AI

- Are contracts in a governed repository?
- Is there a retention policy for expired contracts?
- How do you handle legal privilege and confidentiality?
- Can AI access superseded contract templates?

Compliance / Audit AI

- Is regulatory content updated after amendments?
- How do you ensure AI only sees current guidance?
- Are audit trails required for every AI output?
- Who validates AI-generated compliance summaries?

Knowledge Management AI

- How do you measure the KB's volatility rate?
- What % of your KB is likely ROT right now?
- Is there a formal review cycle for KB articles?
- How do you handle knowledge from departing SMEs?

Operations / Field Service AI

- Are technical manuals version-controlled centrally?
- Who updates operating procedures and how fast?
- Can AI differentiate in-service vs. EOL products?
- What happens when field data conflicts with KB?

Closing

Key Takeaways

1

The readiness phase is over.

Organizations that keep planning instead of acting will fall behind. The time for focused, responsible action is now.

2

You don't need all your data.

Success comes from activating only the precise subsets of unstructured data that serve specific, high-value AI use cases.

3

Shape + Lifecycle + Volatility = AI Readiness.

Understanding these three properties of your data collections is the foundation of every responsible AI implementation.

4

Protection must be designed in, not bolted on.

Veracity, privacy, and security are non-negotiable filters — build them into your data pipeline before AI ever touches the content.

5

Information managers hold the keys.

Your expertise in governance, lifecycle, and data quality is exactly what AI initiatives need. Use that to earn a seat at every AI decision table.

AI: Moving from Readiness to Responsible Implementation

Whitepaper
By Deep Analysis | OpenText 2026

AI-readiness starts with trusted, governed content

AI-assessment
and
Best Practice Guide

Modernize content management for the AI era

IDC Modernization Assessment
and
Best Practice Guide

The readiness phase is over; time for responsible action

Alan Pelz-Sharpe
Founder & Principal Analyst
Deep Analysis

alanpelzsharpe@deep-analysis.net
linkedin.com/in/alan-pelz-sharpe-876b2

Tracy Caughell
Sr. Director, Product Management
OpenText Content Cloud

tcaughel@opentext.com
linkedin.com/in/tracycaughell

opentext™