

The Autonomous Transfer Problem: When AI Agents Send Data Across Sovereign Data Boundaries.

Márton Domokos, Senior Counsel, CMS Hungary
Andrea Orbán, CIPP/E, CIPM, Senior Corporate Counsel, Privacy and Compliance, Iron Mountain

October 12, 2026

Key Points Outline

Core Dilemma: Organizations must keep sensitive data inside sovereign boundaries. However, autonomous AI agents process text via cloud tools, risking cross-border data transfer.

The Key Challenge: How can organizations harness autonomous AI efficiency without allowing sensitive text fragments to escape required sovereign jurisdictions?

Purpose: Share practical architectural guardrails using Iron Mountain’s product as a reference case study—replacing policy promises with infrastructure-enforced controls.

What is InSight DXP & Why It’s Our Case Study

What is Iron Mountain InSight® DXP?

An enterprise **Intelligent Content Platform** that unifies physical and digital information, converting unstructured documents (paper, PDFs, records) into structured, AI-ready intelligence.

Core Capabilities: Agentic Workflows, Semantic Search & Knowledge Graphs, Automated Classification & Extraction, Human-in-the-Loop (HITL) Validation

🦻 *In Plain English: It turns your mountain of messy PDFs into smart data—without losing anything in translation!*

⚖️ Lawyer's Reality Check

“Policy promises are like pinky-swears. Technical locks are like a steel vault with lasers.”

Why Compliance Loves It:

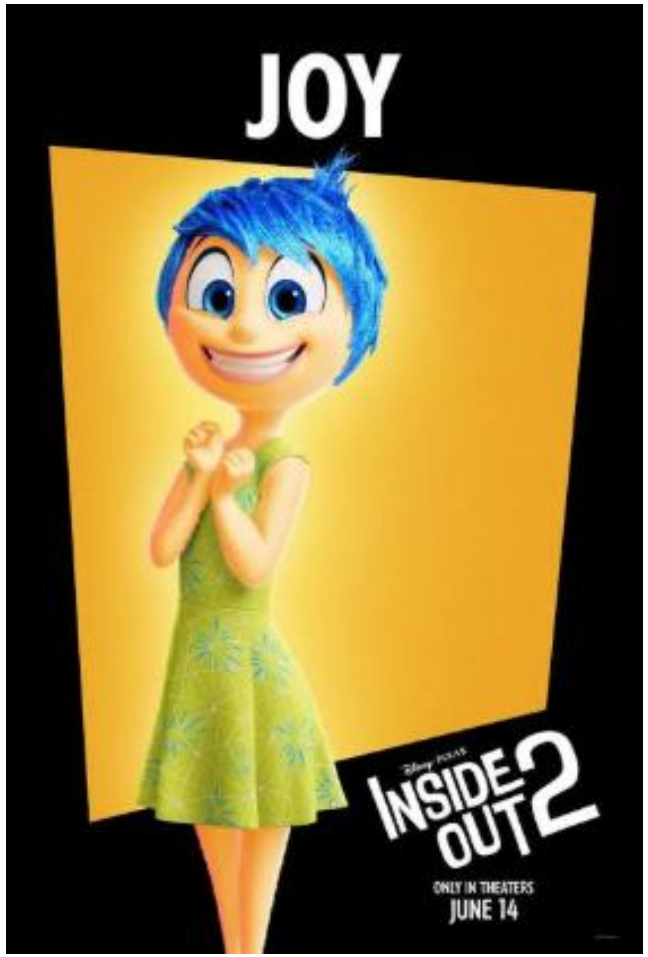
- Zero unapproved cross-border data field trips 🧳
- GDPR auditors leave with a smile 😊
- Legal team finally gets to sleep at night 😴

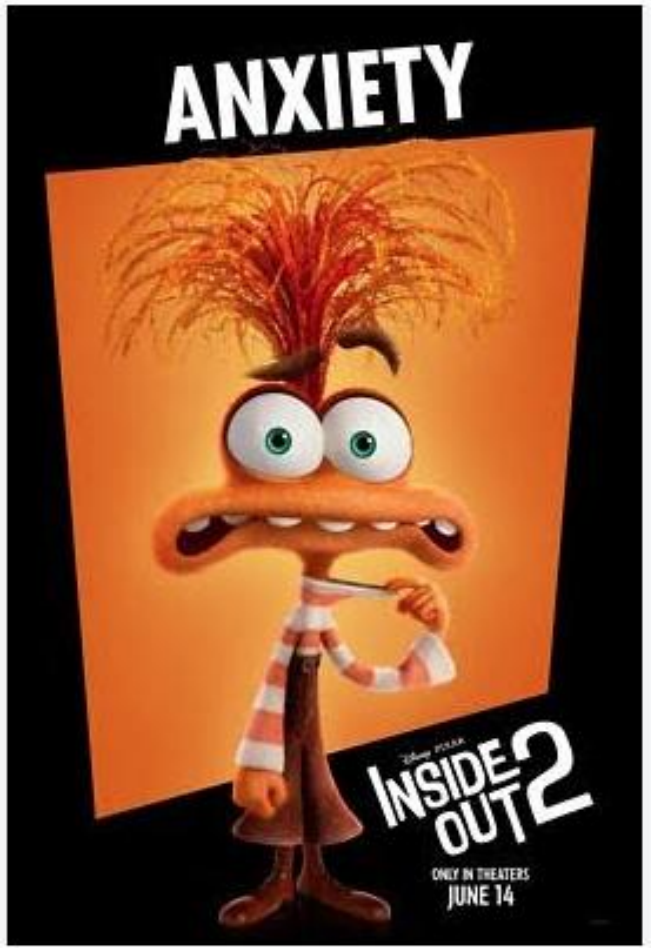
Why Use DXP as Today's Case Study?

- **Infrastructure-Enforced Controls:** Demonstrates how real software architecture replaces unenforceable policy promises with technical software locks.
- **Sovereignty & Privacy by Design:** Illustrates how autonomous AI can operate safely within strict EU Data Boundaries and GDPR compliance frameworks.
- **Proven Operational Impact:** Provides validated real-world metrics across government, judicial, and enterprise legal deployments.

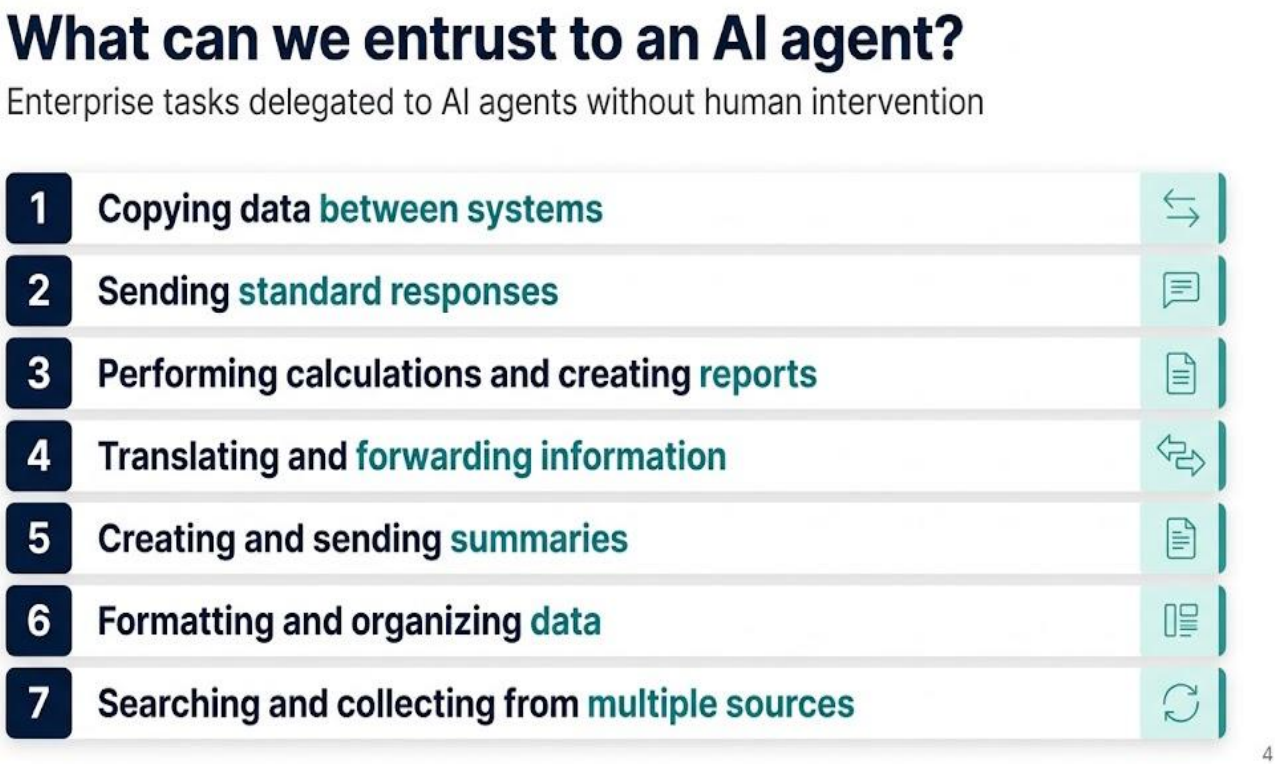


Key Presenters



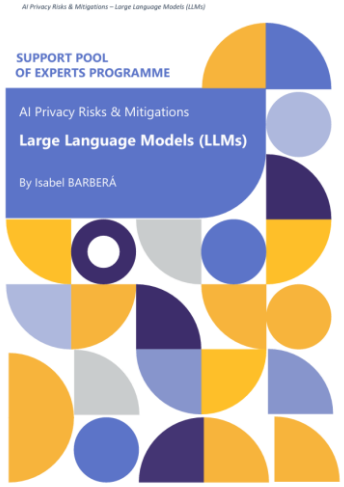


We heard in a presentation on “Agentic AI Beyond the Pilot”…



“Copying data between systems” ranks first among tasks that can be delegated to an AI agent without human intervention.

Key Authorities



Agentic AI

EDPS
EUROPEAN DATA PROTECTION SUPERVISOR

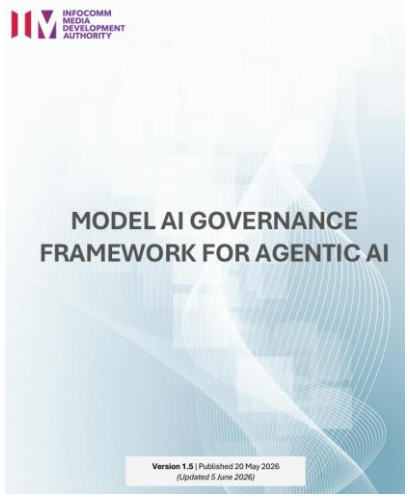
IA agentique et données personnelles : la CNIL et le Conseil de l'IA et du Numérique publient une note exploratoire

20 juillet 2026

CNIL. PROTÉGER les données personnelles
ACCOMPAGNER l'innovation
PRÉSERVER les libertés individuelles

ICO tech futures: Agentic AI

ico.



Securing Agentic AI – An Addendum to the Guidelines and Companion Guide on Securing AI Systems

CSA SINGAPORE
Cyber Security Agency of Singapore

Communiqué: G7 Data Protection and privacy Authorities Roundtable

Office of the Privacy Commissioner of Canada

Key Topics (Agenda)

1.	The "Transfer Problem": Control Over Cross-Border Data Flows	→	Guardrail 1: Containing Data Flows via Local Infrastructure & Tenancy
2.	The "Privacy Risk": Raw Personal Data in the Agent	→	Guardrail 2: Sanitizing Text Fragments via Automated Privacy Gate
3.	The "Subprocessor ~~Risk~~ Sprawl"	→	Guardrails 3 & 4: Closed Toolkits & Model Training Isolation
4.	The "Autonomy Risk": Unbounded Agents	→	Guardrail 5: Controlled Autonomy & Human "AI Shepherds"
5.	The "Auditability Risk": Proving Where the Data Went	→	Guardrail 6: Tamper-Proof Audit Records & AI Risk Registry

The "Transfer Problem": Control Over Cross-Border Data Flows

The Problem: An AI agent splits documents into text fragments and passes them to cloud tools, LLMs, search services, APIs or external servers to perform its tasks – potentially outside the EU or subject to "third-country" jurisdiction. The content may leave the organization's controlled environment without the user realizing it.

Why It Matters Legally: Fragments may contain personal data. Each such call may be a restricted transfer under GDPR Chapter V (Art. 44–49) of the GDPR, needing an adequacy decision or SCCs plus a transfer impact assessment – yet the flow is often invisible to the controller.

What Regulators Expect:	• Give the agent access only to the information and tools needed for a defined purpose (ICO). • Verify where third-party services process data and perform a transfer impact assessment (EDPB). • Whitelist outbound network and API access (CSA).

Guardrail 1: Containing Data Flows via Local Infrastructure & Tenancy

The Transfer Problem

An AI agent chunking text fragments and passing them to unvetted cloud tools or external servers outside local jurisdiction.

How DXP Architecture Solves It

- **Regional EU Data Boundaries**: In European deployments, the architecture enforces strict EU Data Boundary controls. All prompt data, context windows, and generated outputs physically remain within designated local European cloud data centers.
- **Physical vs. Logical Tenancy Separation**: Organizations can choose single-tenant physical separation (dedicated compute, storage, and networking resources where no infrastructure is shared) or multi-tenant logical separation (where data is segregated cryptographically using tenant-specific encryption keys and isolated storage buckets)

The Legal Takeaway

Physical and legal custody of data remains anchored inside local sovereign boundaries, fulfilling data transfer requirements and security mandates.

The "Privacy Risk": Raw Personal Data in the Agent

The Risk: Chunked text carries names, addresses and ID numbers straight into three possible destinations: the model's context window (the information the model receives to generate response), persistent memory (for use in later interactions) and onward tool calls (to other services, such as search tools, APIs or external systems).

Why It Matters Legally: Chunking does not automatically remove or anonymise personal data. Once unmasked personal data enters the agent chain, every downstream call is a transfer of personal data; persistent memory must comply with retention and profiling rules.

What Regulators Expect:	• Restrict access to personal data unless strictly required, using anonymisation instead where possible (CSA). • Remove sensitive data from inputs before they reach the model (EDPB). • Partition agent memory per processing purpose (no shared memory that collects everything from every task) and let stored data expire automatically (CNIL).

Guardrail 2: Sanitizing Text Fragments via Automated Privacy Gate

The Privacy Risk

Chunking text risks passing raw personally identifiable information (PII) directly into an AI model's reasoning engine.

How DXP Architecture Solves It

- **Automated Ingestion DLP Gate:** A Privacy microservice positions an automated Data Loss Prevention (DLP) filter at document ingestion as a mandatory hard gate.
- **Local Sanitization & Tokenization:** Personal names, addresses, and ID numbers are masked locally before model processing, allowing AI to read only tokenized structural patterns.

The Legal Takeaway

Embeds Privacy by Design to sanitize raw identity markers locally, preventing unlawful international transfers of personal data.

The "Subprocessor ~~Risk~~ Sprawl"

The Problem: An agent can decide, based on the task and its instructions, to call external tools or APIs to retrieve information or perform an action. An employee asks an agent to summarise a confidential contract. The agent uses an external API to extract the text or a web tool to look up related information, potentially transmitting parts of the contract to that service.

Why It Matters Legally: unvetted parties, no contractual safeguards, lack of transparency. Art. 28(2) of the GDPR requires prior authorisation of sub-processors and Article 28(4) requires flow-down of obligations; unknown recipients also undermine transparency transfer mapping (Article 30).

What Regulators Expect:	• Vet, contract with and monitor every third-party service the agent uses (EDPB). • Grant agents only the minimum tools and data access needed for each task (IMDA). • Let agents connect only to whitelisted, trusted servers and test third-party tools in sandboxes first (CSA).

Guardrails 3 & 4: Closed Toolkits & Model Training Isolation

The Subprocessor Risk

Autonomous AI agents dynamically calling web tools can lead to dynamic data egress and unvetted subprocessor sprawl.

How DXP Architecture Solves It

- **Closed Capability** or shared across tenants. Accuracy is enhanced using
- **Manifest:** Agents operate strictly within pre-screened internal tools, blocking arbitrary outbound web calls or uninspected APIs.
- **Zero Model Retraining Commitment:** Customer data is never used to train foundation models synthetic data templates.

The Legal Takeaway

Locks down the subprocessor ecosystem to ensure GDPR Article 28 compliance and guarantees complete contractual and technical data isolation.

The "Autonomy Risk": Unbounded Agents

The Risk: an AI agent with a high degree of autonomy can make decisions and take actions without asking a human to approve each step. For example, an agent asked to analyse a contract might independently decide to send parts to an external API, use a web tool to retrieve additional information, incorporate the results into its analysis or pass them to another tool, and take further actions based on the results.

Why It Matters Legally: autonomous actions may trigger Article 22 of the GDPR where decisions are based solely on automated processing and have legal or similarly significant effects, or Article 14 of the AI Act where the system is classified as high-risk and requires human oversight. Unauthorised data disclosure may also be difficult to reverse once the information has reached an external party.

What Regulators Expect:	• Set the agent's level of autonomy according to data sensitivity and task criticality (IMDA). • Have the agent ask a human before accessing personal information (ICO). • Require human approval for high-risk or irreversible actions (CSA).

Guardrail 5: Controlled Autonomy & Human "AI Shepherds"

The Autonomy Risk

Uncontrolled, default agent autonomy making consequential decisions without oversight creates massive legal and operational compliance risks.

How DXP Architecture Solves It

- **Static Permission Boundaries & Circuit Breakers:** Workflows operate within fixed boundaries and pause automatically on illegible input or confidence exceptions.
- **Operational "AI Shepherds":** Tasks encountering exceptions are securely routed to human reviewers for validation before execution continues.

The Legal Takeaway

Complies with automated decision-making regulations by maintaining human control over high-stakes outputs while automating routine tasks.

The "Auditability Risk": Proving Where the Data Went

The Risk: An AI agent may interact with several tools and services during a single task. If the organisation does not keep adequate logs, it may be difficult to establish exactly what happened + without complete records of prompts, model responses and tool calls, the organisation cannot reconstruct which data left, to whom, and on what legal basis.

Why It Matters Legally: Organisations need reliable evidence and logs of agent activity (accountability), maintain appropriate records of processing (Article 30), assess and notify personal data breaches + "AI incidents" (AI Act), and conduct DPIAS.

What Regulators Expect:	• Carry out a DPIA where agentic processing is likely to be high risk (ICO). • Make each task traceable: the personal data used, the agents involved and the third-party services contacted (CNIL). • Keep immutable, tamper-evident logs of prompts, responses and tool calls (CSA).

Guardrail 6: Tamper-Proof Audit Records & AI Risk Registry

The Auditability Risk

Inability to prove exact prompt, model response, and execution histories to regulators or Data Protection Officers during audits.

How DXP Architecture Solves It

- **Trace V2 Audit Schema in WORM Vaults:** Records system prompts, model responses, and tool calls in immutable Write-Once, Read-Many vaults.
- **Internal AI Risk Registry:** Catalogs platform features by risk level to automatically enforce stricter DLP masking and oversight controls.

The Legal Takeaway

Provides tamper-proof evidence for Data Protection Impact Assessments (DPIA) and regulatory compliance audits.

Real-World Proof Points: Productivity & Sovereignty

Public Sector Impact

EU Government Agency: Automated 5M+ handwritten public records, reducing processing costs by 50%+ with 99.5% accuracy.

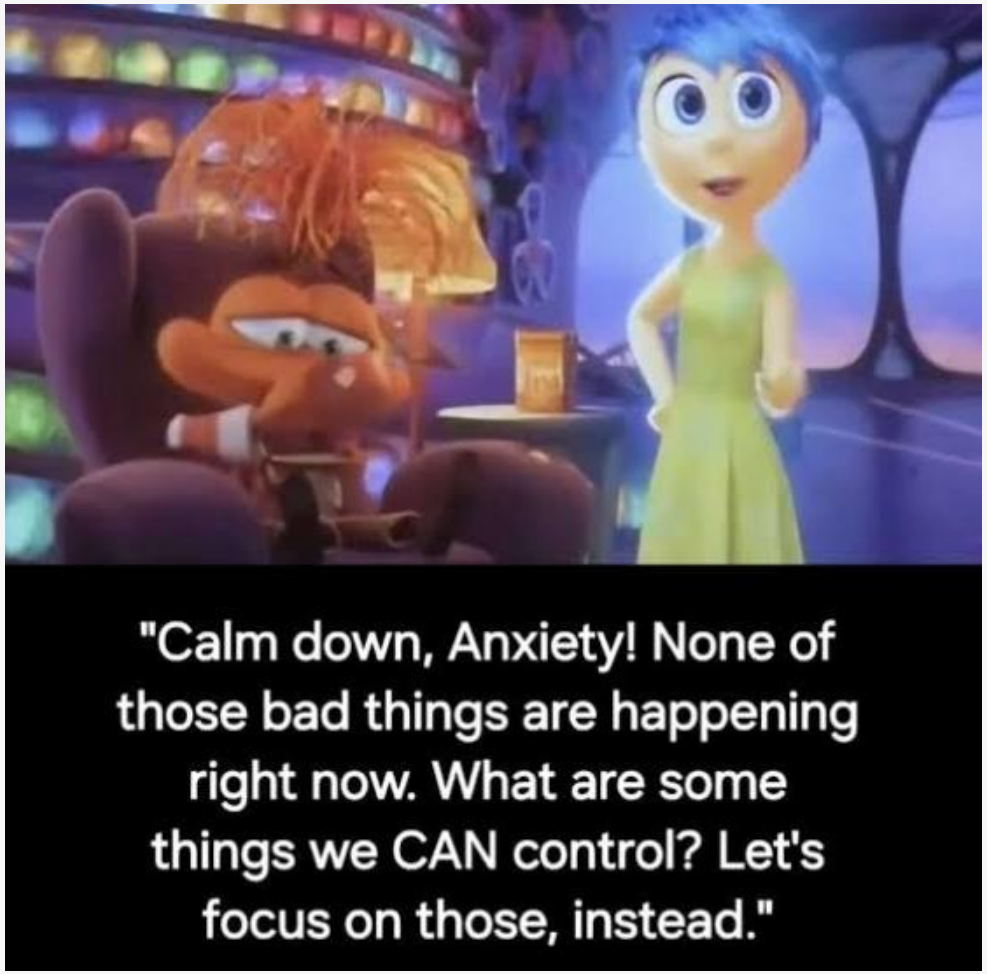
Judicial & Legal Ops

- **UK HM Courts & Tribunals:** Converted 870k pages across 41M files, accelerating document delivery from 5 to 3 days under statutory compliance.
- **Internal Legal Ops:** Reduced contract setup overhead by 25 mins per document, saving 5,500 hours annually within local sovereign boundaries.

Core Conclusion

Architecture-enforced data governance allows legal teams to confidently deploy AI while honoring every sovereign data requirement.

Key Takeaway



Thank you for your attention!



Márton Domokos
Senior Counsel

T +36 1 483 4824
E marton.domokos@cms-cmno.com



Andrea Orbán
Senior Privacy Counsel

T +36 70 4303107
E andrea.orban@emea.ironmountain.com