WEBINAR RECORDED

PECB

AI for NIST, ISO 27001, and CMMC Compliance: Opportunities and Pitfalls

JULY 30 | 03:00 PM CEST

Dr. Jorge Carillo
Managing Director at
PrivacyTech

Dr. Oz Erdem
Linqs Assurance,
Cyber, Compliance, AI & 3rd Party Risk Leader

#GlobalLeadingVoices

Agenda

- Understanding NIST, ISO/IEC 27001 and CMMC v2, and Overlap
- AI in Discovery and Control
- Risk Prioritization
- Use Cases
- Advantages
- Fundamental Limits and Caveats & Critical Risks
- Guardrails and Design Principles & Adoption Roadmap
- Cross-Border Data Transfer & Sovereignty
- Lawful Basis for Processing
- Purpose Limitation & AI Repurposing
- Data Subject Rights in AI Pipelines

AI Improves Speed and Coverage — It Cannot Certify Compliance

The Governance Rule

Speed and coverage improve. Accountability does not move.

1 Speed & Scale, Not Certification

AI accelerates evidence collection, control mapping, and monitoring, but it cannot certify compliance.

2 Non-Delegable Duties

Risk acceptance, control design, and formal assertions to auditors remain squarely human responsibilities.

3 Mandatory Review Gates

Every AI-suggested mapping, remediation action, or audit narrative requires control-owner approval before use.

4 Governance Comes First

RACI structures and sign-off gates must exist before any AI tool is deployed, not retrofitted afterward.

5 Defensibility Is the Metric

Success is measured by how findings hold up under assessor questioning, not by tool sophistication.

6 Organizational Accountability

The organization — not the vendor or the model — bears legal and contractual responsibility for compliance claims.

Understanding ISO/IEC 27001:2022

STANDARD OVERVIEW

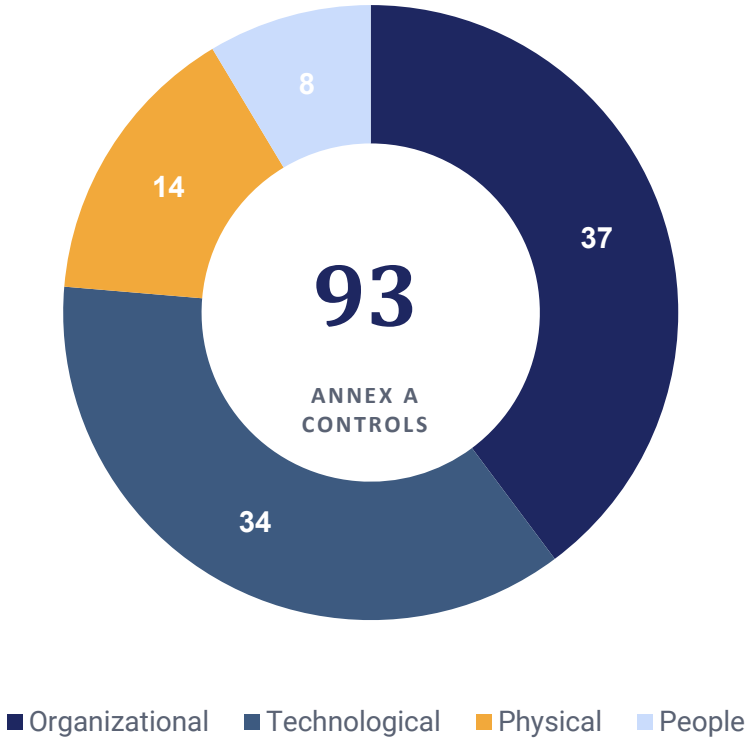

ISMS Foundation

Clauses 4–10 define management-system requirements: context, leadership, planning, and improvement.

Risk-Based Approach

Risk assessment and treatment plans drive control selection — not a static checklist.

Statement of Applicability

Documents justification for including or excluding each Annex A control.

Two-Stage Certification

Stage 1 documentation review, then Stage 2 operational effectiveness testing.

Continual Improvement

Management review and internal audits drive ongoing ISMS maturity.

Understanding CMMC v2 Level 2

STANDARD OVERVIEW

110	14	CUI
Security Requirements	Requirement Families	Primary Protection Scope

Protects Controlled Unclassified Information across the Defense Industrial Base, aligned to NIST SP 800-171 Rev 2 requirement families — from Access Control to System Integrity.

1 SSP and POA&M

A System Security Plan documents implementation; a POA&M tracks unmet requirements to closure.

2 Third-Party Assessment

Most Level 2 programs require a C3PAO assessment; some lower-risk programs allow self-assessment.

3 Boundary Scoping

CUI Assets, Security Protection Assets, Risk Managed Assets, and Out-of-Scope Assets must be delineated.

4 Contractual Flow-Down

Primes must flow CMMC obligations down to subcontractors handling CUI under 32 CFR Part 170 and DFARS (48 CFR Part 252.204.7012)

Compliance Overlap: The ISO ↔ CMMC Crosswalk

Where They Align

Shared Control Intent

Heavy overlap in access control, audit logging, configuration management, and incident response.

Unified Control Library

A converged control repository lets one evidence artifact satisfy multiple frameworks at once.

Many-to-Many Mapping

A single 800-171 requirement often maps to multiple Annex A controls, and vice versa — rarely 1:1.

SME REVIEW

Where AI Adds Value

Where They Diverge

800-171 is more prescriptive on media/physical protection; ISO emphasizes documented risk rationale.

Semantic Similarity Scoring

AI-assisted crosswalking uses embedding-based similarity to propose candidate requirement mappings.

SME Validation Required

Every AI-suggested mapping needs compliance SME sign-off before becoming an authoritative crosswalk.

AI in Discovery and Control Mapping

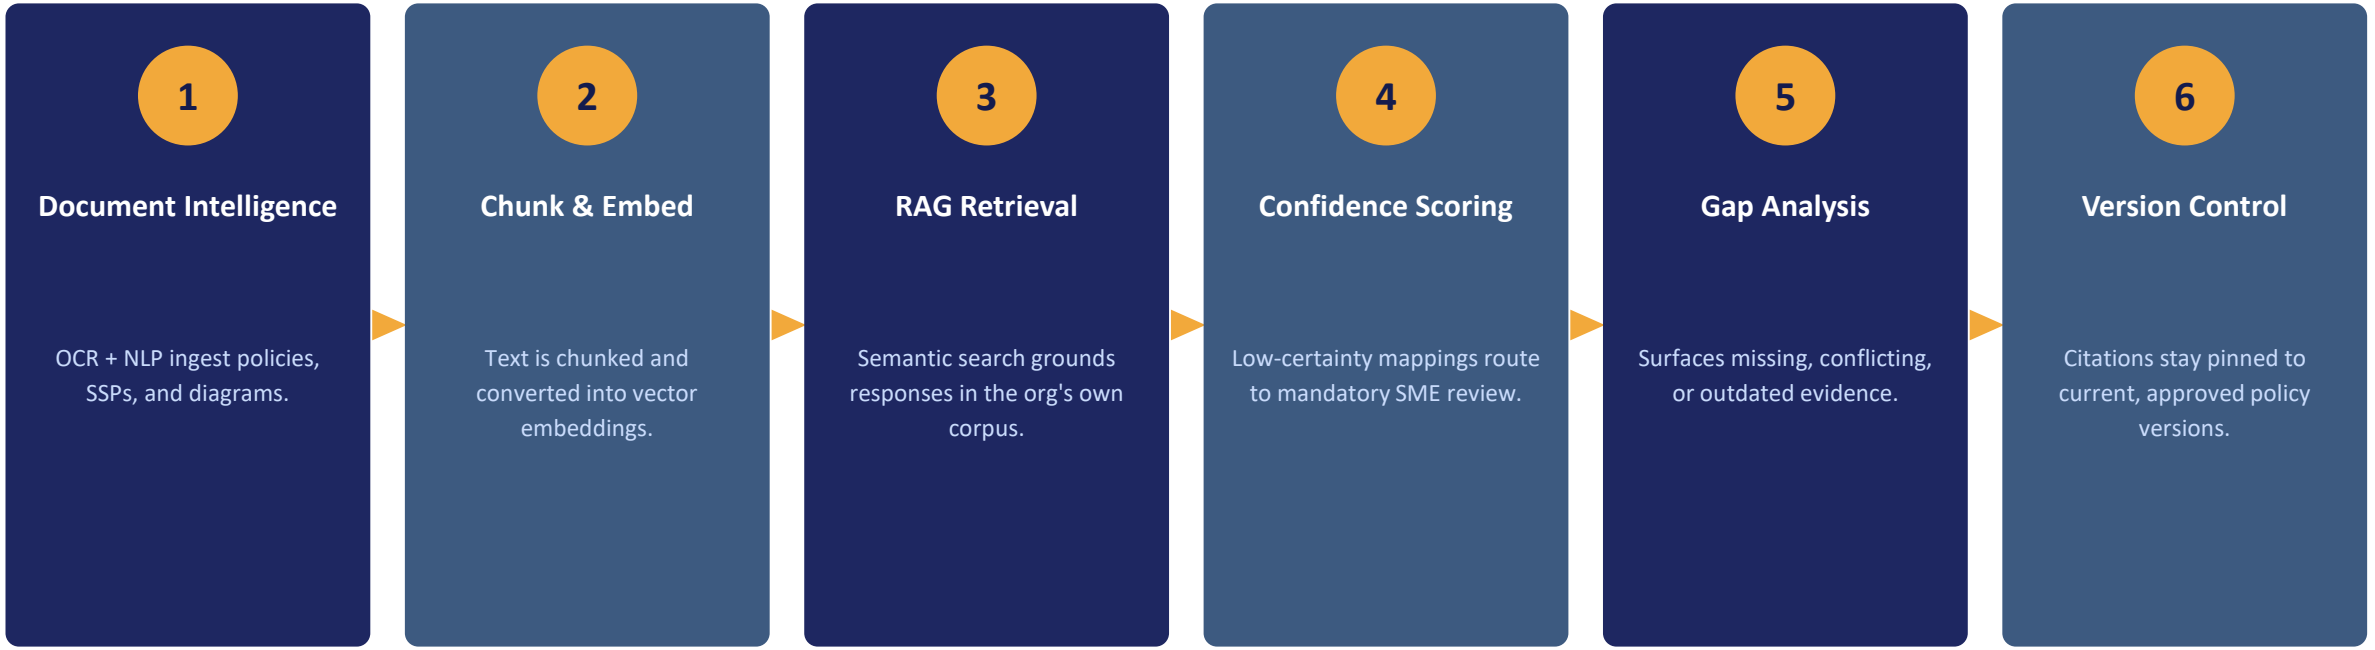

Every retrieval-augmented step above is grounded in the organization's own indexed corpus — never open internet knowledge — and every low-confidence output routes to mandatory SME review before publication.

Evidence Collection and Continuous Monitoring

API-BASED EVIDENCE SOURCES

1	**Continuous Control Monitoring**	Near-real-time dashboards replace point-in-time snapshots for current control health.
2	**ML-Based Drift Detection**	Unsupervised clustering and statistical baselining flag configuration drift and unusual access.
3	**Chain-of-Custody Evidence**	Automated packages are timestamped and hashed to preserve integrity for auditor review.
4	**Rolled-Up Health Scoring**	Technical telemetry aggregates into clause-level or requirement-family status indicators.
5	**Analyst Triage Still Required**	Anomaly models tuned to organizational baselines still generate false positives needing review.

Risk Prioritization and Incident Operations

AI-Assisted Automation

Risk-Weighted Remediation Queues

Blends CVSS severity, EPSS exploit likelihood, and asset criticality into a ranked queue.

Faster Incident Triage

Natural-language summarization drafts initial incident timelines.

Bounded Automated Response

Auto-isolate a host or disable an account for pre-approved, low-risk cases.

Requires Human Approval

High-Impact Actions

Data deletion, production shutdown, or external notification always need explicit sign-off.

POA&M Draft Review

AI-drafted remediation narratives need control-owner review before assessor submission.

Scoped Agent Permissions

Read-only defaults and narrow tokens limit the blast radius of automation error.

Practical Use-Case Alignment

USE CASES

Use Case	What AI Does	Framework Fit
Evidence Orchestration	Auto-collects and packages audit evidence mapped to specific citations.	ISO + CMMC
Security Analytics	UEBA and anomaly detection surface insider-threat and access indicators.	ISO + CMMC
Vulnerability Triage	AI-ranked remediation queues blend exploitability with asset criticality.	CMMC (RA / CA)
Governed Compliance Copilot	Answers policy questions grounded strictly in an approved internal corpus.	ISO + CMMC
Policy Drafting Assistance	Produces first-draft SSP / POA&M language for human editing.	CMMC (SSP)
Vendor Risk Review	Reviews vendor questionnaires and SOC 2 reports to flag gaps.	ISO (A.5.19–23)

Advantages of AI-Enabled Compliance

ADVANTAGES

Speed

Compressed Prep Time

Evidence gathering and first-draft mapping cycles shrink from weeks to days in mature deployments.

Coverage

Broader Coverage

Continuous scanning reaches across large, distributed environments beyond manual sampling capacity.

Consistency

Greater Consistency

Standardized mapping logic reduces variance in interpretation across business units and analysts.

Visibility

Real-Time Visibility

Live control-health dashboards directly support management review and continual improvement.

Cost

Lower Prep Cost

Reduced manual audit-prep labor frees subject-matter experts for judgment-intensive decisions.

Leverage

Cross-Framework Leverage

A single evidence base can serve ISO, CMMC, and SOC 2 simultaneously, cutting duplicate effort.

Fundamental Limits and Caveats

LIMITS

01	**Probabilistic, Not Deterministic**	Identical inputs can yield different phrasing or mapping confidence across separate model runs.
02	**Cannot Establish an ISMS**	A management system is an organizational commitment, not a generated document.
03	**Assessors Judge Effectiveness**	Auditors evaluate operational judgment, not the sophistication of the tooling used.
04	**No Delegated Risk Acceptance**	A named, accountable risk owner must exercise judgment; AI cannot formally accept residual risk.
05	**Generic Training, Specific Context**	Models trained on broad data can misapply requirement language to a unique environment.
06	**Over-Reliance Erodes Expertise**	Teams that stop understanding requirements cannot defend findings under assessor questioning.

Critical Risks: Data and Accuracy

CRITICAL RISK

CUI Leakage Exposure	Sending CUI to unapproved public LLM endpoints can violate CMMC boundary and DFARS flow-down rules.	CRITICAL
Hallucinated Evidence	AI asserting a control is implemented without a verifiable citation is an audit failure waiting to happen.	CRITICAL
Source Provenance	Every AI-generated statement in an audit package must trace to an authoritative, timestamped source.	HIGH
Training-Data Contamination	Using regulated environment data to train third-party models is frequently contractually disallowed.	HIGH
Stale Knowledge Bases	An unsynced RAG index can cite outdated policy versions, creating mapping errors auditors will catch.	MEDIUM
Mandatory Validation Gate	No AI-generated content enters an audit-facing package without a documented human review step.	CONTROL

Critical Risks: Security and Auditability

CRITICAL RISK

1	**Prompt Injection**	Malicious content embedded in ingested documents can hijack agent behavior or expose data.
2	**Excessive Agentic Permissions**	Overly broad API scopes let a compromised or misconfigured agent act beyond its intended purpose.
3	**Explainability Gap**	Black-box outputs without citations cannot be defended when an assessor asks 'how do you know?'
4	**Vendor Assurance Gaps**	Many AI tool vendors lack their own certifications or clear subprocessor disclosures.
5	**Silent Model Drift**	Undisclosed vendor model updates can change output behavior outside change-control review.
6	**Immutable Audit Logging**	Every AI-assisted decision needs a tamper-evident log of who, what, when, and source.

Guardrails and Design Principles

GUARDRAILS

✓ Human-in-the-Loop Gates	Mandatory approval checkpoints before anything reaches an audit package or attestation.
✓ Least-Privilege Access	Scoped, read-mostly credentials for agents; write or execute actions require explicit approval.
✓ Trusted Knowledge Boundary	RAG restricted to vetted, access-controlled internal corpora — never open internet for CUI work.
✓ Evidence Integrity Controls	Cryptographic hashing and immutable logging protect AI-generated artifacts and their sources.
✓ Change Control for Models	Model and version updates pass through the same change-management process as any system.
✓ Vendor Due Diligence	Data-use terms, subprocessor disclosure, training opt-out, and certifications reviewed pre-adoption.

Phased Adoption Roadmap

ROADMAP

01

Governance Foundation

Define AI use policy, RACI, data classification rules, and approval workflow before any tool is deployed.

02

Low-Risk Pilots

Document discovery, gap-analysis drafting, and an internal Q&A copilot on non-sensitive corpora.

03

Controlled Evidence Integration

Connect read-only technical sources; enable monitoring dashboards with human sign-off gates.

04

Assisted Remediation

AI-suggested remediation queues and playbook drafts, still requiring approval before execution.

Ongoing — Continuous Improvement: Quarterly review of AI-assisted finding accuracy, false-positive rates, and audit outcomes. No phase advances without a defined owner and a documented rollback plan.

Adopt AI as a Governed Tool — Not a Shortcut

KEY TAKEAWAYS

01 Adopt AI only where clear ownership exists for every AI-touched control or decision.

02 Ensure you can reconstruct who approved what, and from which source, at any time.

03 Keep AI scoped to acceleration and pattern recognition; keep attestation human-owned.

04 Judge vendors by data-handling maturity and certifications, not capability demos alone.

05 Document AI usage itself as an auditable control within the ISMS or SSP.

Bottom Line: AI is a force multiplier for a mature compliance program — not a substitute for one.

Where Does AI Actually Sit in Your Organization?

Three roles — each one changes how every risk in today's conversation lands.

A

AI as the compliance tool

The audit helper. Draft controls, controls, map evidence, run checks.

B

AI inside the certified scope

In the product. Part of the system system the certificate covers.

C

AI as shadow tooling

Informal use. Outside policy, inside the workflow.

Every risk we discuss today lands differently depending on which role the AI is playing.

Four Lenses

01 Cross-Border Data Transfer & Sovereignty

Where the data actually goes when the model answers.

02 Lawful Basis for Processing

Secure processing is not the same as lawful processing.

03 Purpose Limitation & AI Repurposing

The second life of data collected for something else.

04 Data Subject Rights in AI Pipelines

Do You Know Where Your LLM's Human Reviewers Sit?

(A) KEY CONCEPT

GDPR Chapter V — International Data Transfers

Personal data leaving the EEA needs an needs an adequacy decision, SCCs, or or BCRs — plus a Transfer Impact Assessment (post-Schrems II).

(B) WHY IT'S RELEVANT

Complements the CUI boundary discussion

CMMC protects US-boundary data. data. GDPR the EU direction of the same pipeline — running in parallel. parallel.

(C) IMPACT & ETHICS

Three jurisdictions in one call

A single LLM request may cross three borders before returning.

Individuals rarely know their data crossed a border to train train someone's compliance workflow.

CMMC/NIST-Compliant and GDPR Non-Compliant at the Same Time

(A) KEY CONCEPT

Dual-regime exposure

CMMC's US-jurisdictional CUI rules and GDPR's EU personal-data rules run in parallel. One compliant, one not — is possible.

(B) WHY IT'S RELEVANT

Two axes, one pipeline

Boundary controls cover one axis. The The personal-data axis runs alongside alongside — worth naming out loud. loud.

(C) IMPACT & ETHICS

Certificate ≠ lawful processing

A stamp on the wall doesn't answer the personal-data question.

Two frameworks, two duties, one responsibility.

What Happens to Your Data During Abuse Monitoring?

(A) KEY CONCEPT

Sub-processor transparency — GDPR Art. 28

Human reviewers used for safety and abuse detection are processing personal data — and must be disclosed.

(B) WHY IT'S RELEVANT

Extending the vendor assurance lens

Sub-processor chains often go three or three or four layers deep — worth mapping end-to-end.

(C) IMPACT & ETHICS

Contractors abroad may read prompts flagged for review.

Sensitive prompts, unseen readers

Transparency to the data subject subject is often absent from the the contract.

Is Your "US-region" Endpoint a US-only Path?

(A) KEY CONCEPT

Inference region ≠ full data lifecycle

Regional routing covers where the model answers — not where training data, safety review, and logs live.

(B) WHY IT'S RELEVANT

Boundary claims, sharpened

A concrete example of how a boundary boundary claim can be technically true true and legally insufficient.

(C) IMPACT & ETHICS

False confidence in isolation

"It stays in the US" can be geographically partial.

Precise language matters when when explaining data flows to to customers.

Have You Run a Transfer Impact Assessment on Your Compliance AI?

(A) KEY CONCEPT

TIA — a Schrems II requirement

A documented assessment of risks to data subjects' rights in the destination country — before the transfer.

(B) WHY IT'S RELEVANT

Building on vendor due diligence

The EU procedural step that gives vendor diligence a written form regulators can inspect.

(C) IMPACT & ETHICS

Evidence of thinking, not just doing

Regulators expect documented assessment, not verbal reassurance.

The TIA forces us to consider the consider the human whose data data is transferred.

On What Legal Basis Does Your AI Read That HR File?

(A) KEY CONCEPT

GDPR Art. 6 — six lawful bases

Every processing of personal data needs one: consent, contract, legal obligation, vital interests, public task, or legitimate interests.

(B) WHY IT'S RELEVANT

ISO 27001 approval, side by side with GDPR

An ISO-approved control can still lack a lack a GDPR lawful basis — the two two frameworks live in parallel.

(C) IMPACT & ETHICS

Secure ≠ lawful

Processing without a basis is unlawful regardless of security posture.

Security and lawfulness are not the same virtue.

Have You Documented Your Legitimate Interest Assessment?

(A) KEY CONCEPT

LIA — a three-part balancing test

Purpose, necessity, balancing — required when relying on Art. 6(1)(f) legitimate interests.

(B) WHY IT'S RELEVANT

The default basis for compliance AI

Most compliance-monitoring AI is defended under "legitimate interest" — interest" — rarely with a written LIA. LIA.

(C) IMPACT & ETHICS

Undocumented reasoning collapses fast

Without evidence, the basis dissolves under regulator questioning.

The balancing test weighs the individual's rights rights against organizational convenience.

When Does Monitoring Become Surveillance?

(A) KEY CONCEPT

Proportionality principle

Even lawful processing must be necessary and proportionate to the purpose it serves.

(B) WHY IT'S RELEVANT

Capability outruns intent

AI dramatically lowers the cost of monitoring — proportionality has to to scale with capability, not budget. budget.

(C) IMPACT & ETHICS

Quiet drift into workforce workforce surveillance

Continuous compliance monitoring can shift silently past its original scope.

What we can observe is not what we should observe.

Is Your AI Making "solely automated decisions" About People?

(A) KEY CONCEPT

GDPR Art. 22 — automated decisions

Individuals have the right not to be subject to decisions based solely on automated processing that produce legal or similarly significant effects.

(B) WHY IT'S RELEVANT

Legal weight for human-in-the-loop

Adds legal ballast to the human-in-the-in-the-loop guardrail — in some cases cases the human isn't optional.

(C) IMPACT & ETHICS

A reviewer who only confirms the model has not satisfied Art. 22.

The rubber stamp doesn't count

Meaningful oversight requires requires competence and authority to override.

What is the Lawful Basis for the Basis Itself?

(A) KEY CONCEPT

Layered processing

Training, fine-tuning, inference, and logging are each distinct processing acts — each needs its own lawful basis.

(B) WHY IT'S RELEVANT

One system, several processing acts

A single AI system may involve four or four or five processing acts stacked stacked together.

(C) IMPACT & ETHICS

Deployment approval ≠ lifecycle lifecycle approval

One "yes" at go-live doesn't cover the whole lifecycle.

Consent given once
is not consent given forever.

Was that Data Really Collected for this Purpose?

(A) KEY CONCEPT

GDPR Art. 5(1)(b) — purpose limitation

Personal data must be collected for specified, explicit, legitimate purposes — and not further processed in incompatible ways.

(B) WHY IT'S RELEVANT

SIEM logs feeding compliance AI

Logs collected for security monitoring, monitoring, reused for compliance training, are a secondary purpose worth naming.

(C) IMPACT & ETHICS

Purpose creep is quiet but compounding

Small reuses accumulate into a very different use.

Every reuse is a new promise to the individual.

Does Your Vector Database Remember What You Asked it to Forget?

(A) KEY CONCEPT

Right to erasure — meets embeddings

GDPR Art. 17. Deleting a source document does not remove its statistical footprint inside a vector index.

(B) WHY IT'S RELEVANT

A technical wrinkle on the RAG guardrail

RAG architecture and the right to erasure meet at the embedding layer. layer.

(C) IMPACT & ETHICS

Erasure may be technically partial

Some requests are hard to honor fully at the embedding level.

We should not promise what our what our architecture cannot cannot deliver.

Who Approved the Second Life of This Data?

(A) KEY CONCEPT

Compatibility assessment — Art. 6(4)

Before a new use, an evaluation is required: is the new purpose compatible with the original?

(B) WHY IT'S RELEVANT

Security logs into compliance training

Turning security telemetry into AI training data is a textbook Art. 6(4) case.

(C) IMPACT & ETHICS

Repurposing changes the original promise

Without an assessment, the second use breaks the first.

Consent is a contract, not a starting position.

Is Your RAG Retrieving Data that Was Never Meant to be Retrievable?

(A) KEY CONCEPT

Purpose creep by architecture

Once embedded, documents become retrievable by any authorized user, for any prompt.

(B) WHY IT'S RELEVANT

A concrete mechanism of drift

A specific way purpose limitation quietly breaks — visible only when you you inspect the index.

(C) IMPACT & ETHICS

Architecture becomes policy

And that architecture is often more permissive than intended.

Design choices are governance choices.

Would Your Training Pipeline Survive a Purpose Limitation Audit?

(A) KEY CONCEPT

Training as its own processing activity

Model training is a distinct processing act — often the most opaque and hardest to unwind after the fact.

(B) WHY IT'S RELEVANT

Complementing the contamination point

Adds the individual-rights dimension to dimension to the training-data conversation.

(C) IMPACT & ETHICS

Weights don't unlearn cleanly

Training data cannot be surgically removed from a trained model.

Some processing decisions are effectively permanent.

Why did the System Decide this About Me?

(A) KEY CONCEPT

Right to explanation — Art. 22 & Recital 71

Meaningful information about the logic, significance, and consequences of automated processing.

(B) WHY IT'S RELEVANT

Explainability, reframed

Explainability is not only an auditor auditor concern — it is a data subject subject right.

(C) IMPACT & ETHICS

A black box that satisfies an auditor may still fail a data subject.

Auditor-friendly ≠ subject-friendly

Explanation is dignity, not decoration.

Can You Show Me Everything Your AI Knows About Me?

(A) KEY CONCEPT

Right of access — GDPR Art. 15

Access spans vector databases, prompt caches, inference logs, and fine-tuning datasets — not just the CRM.

(B) WHY IT'S RELEVANT

Extending the immutable-logging lens

Immutable logs are an assurance mechanism — and also a data-subject subject surface.

(C) IMPACT & ETHICS

Most stacks can't fully answer

A comprehensive Art. 15 request against an AI stack is a stress test.

If we cannot show it, we should question keeping it.

Can Your AI Actually Forget Me, or Pretend to?

(A) KEY CONCEPT

Right to erasure — Art. 17 in AI stacks

Applies to embeddings, caches, and fine-tuning artifacts — not only the source record.

(B) WHY IT'S RELEVANT

Architecture must match the promise

The technical counterpart to the legal legal right — design has to support support what the policy claims.

(C) IMPACT & ETHICS

Deleting the row while leaving the embedding is legally exposed.

Partial deletion is a real risk

Honesty about limits is better better than performative compliance.

What if Someone Objects to Being Analyzed by Your Compliance AI?

(A) KEY CONCEPT

Right to object — GDPR Art. 21

Where processing rests on legitimate interest, individuals may object — and the objection must be considered.

(B) WHY IT'S RELEVANT

Individual-level limits on the pipeline

AI-driven compliance monitoring meets meets a per-person constraint.

(C) IMPACT & ETHICS

Objections must be honored

Unless compelling legitimate grounds are demonstrated and documented.

Consent-adjacent bases require ongoing negotiation.

Is Your Human-in-the-loop a Reviewer, or a Rubber Stamp?

(A) KEY CONCEPT

Meaningful human oversight

Requires three things at once: competence to judge, authority to override, and time to consider.

(B) WHY IT'S RELEVANT

Adding substance to the guardrail

The presence of a human is not the the same as meaningful review.

(C) IMPACT & ETHICS

Automation bias is the default

Under time pressure, most reviewers accept what the model suggests.

Oversight without authority is theater.

Does the Person Whose Data Trained Your Model Know It?

(A) KEY CONCEPT

Transparency — Arts. 13 and 14

Information duties at the point of collection — and again when personal data is obtained indirectly.

(B) WHY IT'S RELEVANT

Provenance is a compliance question

Where the training data came from is a legal legal question, not only an engineering one. one.

(C) IMPACT & ETHICS

Silent training erodes trust

Learning from personal data without notice is a trust cost that compounds.

Transparency is a precondition for legitimacy.

Compliance ≠ Ethics.
Which One are We Optimizing for?

ISO 27001 will certify that you have a process.

GDPR will judge whether that process respects a human.

CMMC will judge whether you kept the information in the right system & away from the unauthorized persons/entities.

AI can help you pass all three — but only after you decide which role your AI is playing.

Would you let an AI write the breach notification to the person whose data it leaked?

If the answer is “NO," the boundary of AI in your compliance program is somewhere before that point.

Find it — deliberately.

Two Lenses, One Governance

Speed & coverage

AI improves the speed and coverage of compliance work — it — it cannot certify compliance on on its own.

Lawfulness & the person

AI must also respect lawfulness, lawfulness, purpose, and the rights rights of the individual behind the the data.

COMBINED

Control, context, person

Governance that covers the control, the context, and the person the data belongs to.

THANK YOU

Dr. Ozkan Erdem

Linqs Assurance

✉ oz@linqs.net

✉ jorge@carrillo.cz

Dr. Jorge Carrillo

Managing Director, PrivacyTech

in https://www.linkedin.com/in/drozkanerdem/

in https://www.linkedin.com/in/privacytech/