

Specimen: five rows of the obligation register

ISO/IEC DIS 27045.2, Edition 1, stage 40.60 Close of voting reached 7 August 2026. The publication estimate of 15 September 2027 is our arithmetic from the stage sequence, not an ISO date. The document is guidance: it is not certifiable and creates no transition deadline.

What this is. This is a labelled specimen: the first five rows of the register, the column order and the plan. It is not a client record and it is not the standard. 26 units and 8 characteristics, with the evidence an auditor samples and the change table from the first draft. It does not reproduce the ISO text and it does not certify.

The obligation register

Ref	Clause	Title	Text	Requirement in our words	What to do	Evidence to keep	Owner
ISO27045-1	1	Scope	Public	The scope states what the document is for and who it is for. It gives guidance on managing the risks and navigating the threats that arise across the life cycle of a big data system, from the characteristics that make it different: volume and velocity, variety and variability, volatility and veracity, validity and value. It says the guidance applies when big data is used in the design and rollout of AI systems. It closes by saying that every organisation which develops or uses big data systems is covered, whatever its type, size or purpose.	Read the scope as an applicability test, because that is what it is. The document does not ask whether an organisation is large, regulated or in a particular sector. It asks whether the organisation develops or uses big data systems. If the answer is yes, the guidance is written for you, and the practical question becomes which of the eight characteristics your systems actually show. The second thing the scope settles is that this is not a technology-only document. It names AI systems in the same sentence as the big data life cycle, so a data platform that feeds a model is inside the guidance, and so is a model that produces data the platform then holds. Write the applicability note against the scope's own words, not against a vendor's summary.	A one-page applicability note naming each big data system, the business owner and the date.; The data inventory extract that shows those systems, with the data types and volumes.; The risk assessment or method that already covers them, so the guidance can attach to it.; The board or risk committee minute that records the decision to prepare now or wait.	
ISO27045-2	2	Normative references	Public	The normative references are the documents whose content is required to apply this one. This draft names two: ISO/IEC 22989:2022 for artificial intelligence concepts and terminology, and ISO/IEC 27005:2022 for the management of information security risk. Dated references apply as the edition cited. Undated references apply as the latest edition, including amendments.	Two references is a short list and it tells you how the document is meant to be used. The risk process comes from 27005:2022, so clause 7 is a big data lens on a process an ISMS manager already runs, not a second risk method. The AI vocabulary comes from 22989:2022, so the document does not define AI for itself and does not need the reader to agree a new glossary. Check the reference list in your own method and supplier documents. A list that still shows 27005 alone is out of date against this draft, and a list that cites an older 22989 edition will not match the terms this document uses. When the standard is published, re-read the list before you re-baseline: a normative reference can be added or dated between the draft and the published edition.	The current reference list from the risk method, with the edition and year for each entry.; The controlled copies of ISO/IEC 27005:2022 and ISO/IEC 22989:2022, or the subscription that gives access to them.; A note showing which clauses of your method depend on each reference.	

Ref	Clause	Title	Text	Requirement in our words	What to do	Evidence to keep	Owner
ISO27045-3	3	Terms, definitions and abbreviated terms	Group	This clause holds the vocabulary the rest of the document uses. It has two parts: the defined terms and the abbreviated terms. The definitions come from the two normative references and from the wording this document itself adds, and the clause says which source applies to each term.	Treat this clause as the contract between the document and your existing glossary. Twenty-two terms are defined, and several of them are the words a risk workshop argues about: asset, threat, vulnerability, risk, validity, veracity. If your risk register uses its own definitions, the gap between the two is where a finding will come from. The abbreviated terms matter less, but they are a quick check on whether a supplier's questionnaire is using the same vocabulary. The abbreviations are data protection officer, identity and access management, large language model, multi-party computation and personally identifiable information.	The organisation's glossary, with the source and edition for each term it defines.; A mapping from the organisation's terms to the document's terms, with the differences named.; The review record showing the risk and privacy owners agreed the mapping.	
ISO27045-3.1	3.1	Terms and definitions	Public	Twenty-two terms are defined for the document, drawn from the two normative references and from the document itself. They run from artificial intelligence system and asset through big data, data classification and data life cycle to validity, value, variability, variety, velocity, veracity, volatility, volume and vulnerability. Each entry names its source, and several are marked as modified from that source.	The terms that matter for the work are the ones this draft changed. Big data is defined by the characteristics that require a scalable architecture, which is why the scope and the risk lens are built on characteristics rather than on a size threshold. Data life cycle now runs from the first idea to discontinuation, which gives retention and deletion a place in the life cycle instead of leaving them outside it. Risk is the effect of uncertainty on objectives, so a big data risk is still an effect on an objective, not a data quality defect on its own. Read the definitions in this order: big data, data life cycle, risk, threat, vulnerability, then the eight characteristics. That order lets each later term rest on the one before it. The changed terms are listed with their old and new wording on the terms page.	The glossary mapping, with the eight characteristic definitions quoted from the published edition once it exists.; The risk register's term set, checked against the document's definitions.; The training or briefing record that shows the risk team has seen the changed terms.	
ISO27045-3.2	3.2	Abbreviated terms	Public	Five abbreviations are defined: DPO for data protection officer, IAM for identity and access management, LLM for large language model, MPC for multi-party computation and PII for personally identifiable information.	The list is short and each abbreviation points at a control area the document expects the reader to know. IAM signals that access control is part of the answer to big data risk. MPC signals that privacy-enhancing computation is in scope as a control option, not only as a vendor feature. PII and DPO signal that privacy is treated as a first-class half of the document, which the title also says. Use the list to check a supplier questionnaire. If a supplier answers a big data security question without naming who holds access, how identity is managed or where personal data is processed, the answer is thin against this vocabulary.	The supplier questionnaire, with an IAM and a privacy section that use the same abbreviations.; The access review record for the big data platform.; The privacy impact assessment for any system that processes PII at scale.	

The dated plan

Date	Action	What it covers	Evidence to keep
2026-10-28	Record the applicability decision	Decide whether the organisation develops or uses a big data system, name the systems, and record the decision with its owner and date. Clause 1 is the test.	A one-page applicability note citing clause 1.
2026-11-11	List the characteristics present	For each of the eight characteristics, write present, absent and the reason. Keep veracity and validity apart.	The eight-row characteristic sheet.
2026-11-25	Add the lens to the risk method	Write the characteristic lens into the existing risk method at identification, analysis, treatment and review. Clause 7 follows ISO/IEC 27005:2022.	A versioned method with the lens added.
2026-12-09	Walk one risk end to end	Take the characteristic most present and walk it from identification to a treated risk with an owner.	One completed risk row an auditor can sample.
2027-01-20	Name the evidence	Use the evidence list to name the artefact for each characteristic in scope, and where in the estate it lives.	An evidence list with no blank owner.

Date	Action	What it covers	Evidence to keep
2027-03-15	Brief the control owners	Brief the people who supply the evidence in plain words, and add the big data questions to the internal audit checklist.	A briefing record and a checklist version.
2027-05-20	Rehearse the audit question	Run one internal review of the characteristic sheet against the evidence. Keep the findings apart from the incident register.	A review note and any corrective action.
2027-07-15	Prepare the board paper	Write the one-page decision: what the guidance is, what it is not, what it changes for us, and the cost of preparing now against waiting.	The board paper and its date.
2027-09-15	Check the stage record	This is the publication estimate, not an ISO date. Check the ISO record for a stage change and update the plan if the date moved.	A dated stage note.
After publication	Re-read the withheld clauses	Clauses 4, 6, 7 and Annex A are not public in the draft. Read them on publication and record what they add to the register.	A reading note and a new register version.

The characteristic worksheet

Characteristic	Term	Definition in our words	Present?	Reason if absent	Threat the draft names	Evidence	Owner
Volume	3.1.21	How much data there is, measured by what it costs to compute, store and manage while it is processed. It is a workload question as much as a size question.			Layers of storage and transmission, different networks and protocols, and a mix that makes normalisation hard. AI systems can produce data faster than the estate can trace or store it.		
Velocity	3.1.18	How fast data moves through the estate: created, then transmitted, stored, analysed and visualised. The rate covers the whole life cycle, not only the moment of creation.			A breach that is over before a batch control notices, a stream that bypasses the control in front of the warehouse, and a model whose data changed before its review.		
Variety	3.1.17	How many shapes the data takes: formats, logical models, timescales, meanings and sources. Unstructured and semi-structured data sit here alongside tables.			Sensitive data that classification never reached, and re-identification by correlating an anonymised set with a public one. Text, images and sensor data need handling rules that a table policy does not give.		
Variability	3.1.16	The data keeps its name while its speed, shape, meaning or quality moves under it. A schema, a unit or a source can change without the data set being renamed.			A schema change that stops a masking rule matching, a semantic change that makes an old classification wrong, and a quality drop that makes a report misleading without making it fail.		
Volatility	3.1.20	How quickly the data goes stale, which is why a data set stops suiting the current analysis and why an audit trail can lose its integrity.			Data kept because storage is cheap, an audit trail whose timestamps no longer line up with the decision it explains, and a retention rule that has drifted from the analysis it serves.		
Veracity	3.1.19	How complete and accurate the data is, read together with where it came from and what it is being used to support. Provenance is part of the answer.			Risks that aggregate and magnify, decisions built on data whose source is uncertain, and AI-generated data that is unreliable in a way the reader cannot see.		
Validity	3.1.14	Whether the data is fit for the purpose it is put to, judged on relevance, consistency, precision and correctness. Accurate is not the same as valid: a record can be right and still incomplete.			A model that loses validity when the data it was trained on changes, analysis that answers a question nobody asked, and a data set that is accurate but not fit for the decision.		
Value	3.1.15	What the organisation learns from finding patterns in the data, and the business benefit that follows. The document calls it the most important characteristic.			Insight worth stealing and worth manipulating, derived data sets that inherit weak access rules from the raw data, and a board metric that can be changed without a trace.		

Not legal advice; does not certify. This is our reading of the public draft, not the official text. Record ee9973f021b64da4, retrieved 2026-10-21. Source: <https://www.iso.org/standard/88970.html>. The evidence list and change table are in the CSV and XLSX.