PEMBROKE PRIVACY

Records of Processing Activities

A Practical Guide to Creating, Structuring, Maintaining and Assuring an Accurate ROPA

Prepared by Daniel Whooley, Associate Consultant at Pembroke Privacy

Table of Contents

2. Executive Summary 3
3. Purpose and Regulatory Context 4
3.1. Who Must Maintain a ROPA 5
3.2. What Regulators and Stakeholders Expect 6
4. Controller and Processor ROPAs 6
5. Minimum ROPA Requirements 8
5.1. Identity and Contact Details 8
5.2. Purposes of Processing 8
5.3. Categories of Data Subjects and Personal Data 8
5.4. Categories of Recipients 9
5.5. International Transfers 9
5.6. Retention Periods 9
5.7. Security Measures 10
5.8. Example of a ROPA Entry 11
6. Additional Fields for Privacy Governance 12
7. Choosing the Right Level of Detail 14
8. A Nested ROPA Structure 16
9. Creating the Initial ROPA 18
9.1. Scope and Govern 18
9.2. Assemble Existing Evidence 18
9.3. Validate with the Business 19
9.4. Reconcile, Assure and Approve 19
10. ROPA Accuracy Through Embedded Workflows 20
11. How the UK GDPR differs from the EU GDPR 21
12. Conclusion 22
13. Appendix A: Embedding ROPA Updates in Business Workflows 24
14. Appendix B: ROPA Quality-Assurance Checklist 26

© 2026 Pembroke Privacy

2.

Executive Summary

A Record of Processing Activities (ROPA) is the written record that Article 30 of the General Data Protection Regulation (GDPR) requires most controllers and processors to maintain.[1] For a controller, it sets out, for each processing activity, why personal data is used, whose data and which categories are involved, who receives it, whether it leaves the European Economic Area (EEA), how long it is kept and, in general terms, how it is protected. A processor's record is narrower and describes the processing it carries out on behalf of each controller it serves. Both must be made available to the supervisory authority on request.

Most organisations have a ROPA in some form, with many built in 2018 when GDPR came into effect. Far fewer have one they would be comfortable producing to a supervisory authority, relying on to scope a data breach, or sharing with a client conducting due diligence. The difference is rarely the template. More often, the document is outdated, having been left unchanged since its creation; it is too vague to meet regulatory expectations; or it describes processing activities that exist in theory but do not match the processing carried out in practice. The obligation to have a complete and accurate ROPA is not just legal theory. Supervisory authorities regularly issue fines on this matter, with the Irish Data Protection Commission finding in June 2026 that the Health Service Executive (HSE) had infringed Article 30 GDPR by failing to maintain a complete and compliant record of processing activities at the time of a 2018 ransomware attack,[2] a reminder that a ROPA is judged at the moment it is needed, not the moment it was written.

Creating a ROPA and keeping it accurate are very different undertakings. Creation is a project with a scope, a timetable, and an end date. Accuracy is an operating discipline, activated whenever a system is procured, a contract is renewed, a product is launched or a team is reorganised. Organisations that treat the ROPA as a project tend to produce a careful snapshot that begins to drift on the day it is approved.

This paper argues that the accuracy of a ROPA is a question of operating model rather than effort. Pembroke Privacy recommends a nested structure that mirrors how the organisation is run, moving from department to business process to processing activity, with named owners at each level. The initial build should start from existing documentation and be validated with the people who perform the work. Maintenance should then be embedded in workflows the organisation already operates, such as procurement, change management, project initiation, data protection impact assessment (DPIA) screening, strategy changes, team overhauls, and contract renewal, so that updates are small, frequent and evidenced as a by-product of decisions already being made. Risk-tiered owner attestation and

1 GDPR, Article 30.

2 DPC, Inquiry into Midlands Regional Hospital Tullamore (IN-19-9-4) (10 June 2026).

© 2026 Pembroke Privacy

independent assurance complete the model, providing confidence that the record reflects reality and not merely that every field has been completed.

In summary, Pembroke Privacy recommends that organisations: (1) structure the ROPA by department, business process and processing activity, with a named owner at each level; (2) build the first version from existing evidence and validate it with the people who perform the work; (3) attach ROPA updates to workflows the organisation already operates, such as procurement, change management and DPIA screening; and (4) confirm accuracy through risk-tiered owner attestation and independent assurance.

The paper distinguishes between what Article 30 expressly requires, what other provisions of the GDPR require, what regulators recommend and what Pembroke Privacy recommends. Where the paper states that the GDPR "requires" something, it describes a legal obligation, where it states that Pembroke Privacy "recommends" something, it describes practical judgement. The paper leads with the EU GDPR as supervised in Ireland, and the section 'How the UK GDPR differs from the EU GDPR' summarises the material differences under the UK GDPR.

3.

Purpose and Regulatory Context

Article 30 sits within Chapter IV of the GDPR, alongside the controller's obligation to implement appropriate measures and be able to demonstrate compliance (Article 24) and the obligation of data protection by design and by default (Article 25). Recital 82 explains that controllers and processors should maintain records of processing activities to demonstrate compliance and make them available to the supervisory authority so that they can be used to monitor processing operations.[3] Supervisory authorities describe the ROPA as one of the means by which controllers demonstrate and implement the accountability principle in Article 5(2), and state that a well-drafted record shows that a controller is aware of, and has considered the purpose of, all processing taking place within the organisation.[4]

The ROPA does not, on its own, prove that processing is lawful. Its value lies in being the authoritative index of an organisation's processing, the point from which lawful basis analyses, DPIAs, processor contracts, transfer assessments, retention rules, and security controls can be located and tested. It is the document that is pulled when a breach occurs, when an access request or complaint arrives, when a supplier is reviewed, when a privacy notice is updated and when a regulator asks questions. A record that cannot support those uses is not performing its accountability function, however complete its columns appear. For the same reason, the ROPA should not be treated as a static inventory of the data an organisation holds. It describes what the organisation does with personal data and why.

3 GDPR, Recital 82 and Articles 24, 25 and 30.

4 DPC, Guidance Note: Records of Processing Activities (April 2023).

© 2026 Pembroke Privacy

3.1. Who Must Maintain a ROPA

Article 30(1) applies to controllers and Article 30(2) to processors, together with their representatives where applicable. The record must be in writing, including in electronic form, and made available to the supervisory authority on request. Following a 2022 sweep of thirty organisations' records, in which some organisations struggled to respond within the time requested, a supervisory authority indicated that ten days should be sufficient notice for any organisation to produce its ROPA.[5]

Article 30(5) disapplies the obligation for an enterprise or organisation employing fewer than 250 persons, unless the processing it carries out is likely to result in a risk to the rights and freedoms of data subjects, is not occasional, or includes special category data or personal data relating to criminal convictions and offences. The Article 29 Working Party, whose position paper the European Data Protection Board (EDPB) has endorsed, confirmed that these conditions are alternatives, so that any one of them triggers the obligation. It also confirmed that a smaller organisation need only record the activities that meet one of the conditions, and that an organisation's regular processing of employee data cannot be regarded as occasional.[6] In practice, few employers can rely on the derogation for their core activities.

In May 2025 the European Commission proposed, as part of its fourth simplification omnibus ("Omnibus IV"), to replace Article 30(5) so that the obligation would not apply to organisations employing fewer than 750 persons unless their processing is likely to result in a high risk within the meaning of Article 35.[7] The EDPB and the European Data Protection Supervisor supported the objective while requesting clarifications, noting that the record remains a useful tool for transparency and data subject rights.[8] Even where a derogation applies, a structured record remains the most efficient way to meet transparency, access, and breach notification obligations, none of which is affected by the proposal.

5 Ibid.

6 Article 29 Working Party, Position Paper on the Article 30(5) derogations (19 April 2018), endorsed by the EDPB.

7 European Commission, Proposal for a Regulation amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2019/1020 and (EU) 2023/2854 as regards the extension of certain mitigating measures available for small and medium-sized enterprises to small mid-cap enterprises and further simplification measures (21 May 2025).

8 EDPB-EDPS Joint Opinion 01/2025 on the Proposal for a Regulation on simplification measures for SMEs and SMCs, in particular the record-keeping obligation under Art. 30(5) GDPR (adopted 8 July 2025).

© 2026 Pembroke Privacy

3.2. What Regulators and Stakeholders Expect

Authorities have recommended that larger organisations divide their ROPA by business function so that processing activities are not missed. They have also cautioned against entries that lack sufficient detail and described the ROPA as a living document that should be reviewed and updated on a continuous basis.[9] Beyond the regulator, the record is relied upon by the data protection officer (DPO) in monitoring compliance, by internal audit, by boards seeking assurance and, increasingly, by clients and acquirers conducting due diligence.

An incomplete record leaves processing that has never been assessed for lawful basis, DPIA requirements, transfers, or retention, and that will be missed when a breach or access request needs to be scoped, and incompleteness can itself be an infringement. A vague record cannot support a privacy notice, a breach assessment, or a regulator's questions. An overly detailed record becomes unmaintainable and obscures the processing that carries real risk. An outdated record is arguably the most dangerous because it is relied upon with a confidence it no longer deserves.

4 .

Controller and Processor ROPAs

A controller determines the purposes and means of processing, and a processor processes personal data on the controller's behalf (Article 4(7) and (8)). The EDPB describes these as functional concepts, which must follow what the parties actually do.[10] The same organisation will often be a controller for its own HR, finance and marketing activities and a processor for services it provides to clients. Each role carries its own Article 30 obligation, and Table 1 compares what each record must contain.

Element	Controllers	Processors
Identity and contact details of controllers, joint controllers, processors, representatives, and DPOs	✓	✓
Purposes of processing	✓	x
Categories of processing carried out for each controller	x	✓

9 DPC, Guidance Note (n 4).

10 EDPB, Guidelines 07/2020 on the concepts of controller and processor, version 2.1.

© 2026 Pembroke Privacy

Element	Controllers	Processors
Categories of data subjects and personal data	✓	x
Categories of recipients	✓	x
International transfers and safeguards (where applicable)	✓	✓
Retention periods (where possible)	✓	x
General description of security measures (where possible)	✓	✓

Table 1: Controller and processor records compared under Article 30 GDPR

The differences reflect the role the organisation plays in relation to each activity. A processor does not decide why personal data is processed, so purposes, recipients and retention are decisions recorded in the controller's record. The processor's record instead evidences what it does, for whom, where and how securely. The two records therefore describe the same processing from different vantage points. The controller's record should show the processor as a recipient and should reflect any transfers made by the processor or its sub-processors. The processor's record should be traceable to its contract with each controller.

That contract is the processor's primary source. Article 28(3) requires a Data Processing Agreement (DPA) to outline the subject matter and duration of the processing, its nature and purpose, the type of personal data and the categories of data subjects, and to bind the processor to act only on documented instructions. A processor whose Article 30(2) record is aligned with those contractual descriptions has a record it can defend. Where a processor goes beyond its instructions and determines purposes and means for itself, Article 28(10) treats the processor as a controller for that processing, because it now determines the purposes and means.[11] That processing therefore belongs in its controller record.

Pembroke Privacy's ROPA collection sheets for controllers and processors illustrate how the two records can be aligned without being treated as identical. The processor sheet's scope note records that it covers processing carried out on behalf of controllers and intentionally excludes controller-only fields such as purpose, lawful basis, categories of data subjects and personal data, recipients and retention periods, while using the same department, processing activity and sub-processing activity hierarchy as the controller sheet. Where an organisation acts in both roles, Pembroke Privacy recommends separate registers, or clearly segregated views, with common identifiers for shared systems and the contract reference recorded in both. Ownership of the controller record should sit with business process owners, and ownership of the processor record with the relevant service or client delivery owner.

1 EDPB, Guidelines 07/2020 (n 10), noting that a processor that goes beyond the controller's instructions and starts to determine its
1 own purposes and means of processing will be considered a controller in respect of that processing.

© 2026 Pembroke Privacy

5 .

Minimum ROPA Requirements

5.1. Identity and Contact Details

Article 30(1)(a) requires the name and contact details of the controller and, where applicable, any joint controller, the controller's representative, and the DPO. In practice, the record should identify the legal entity rather than a brand, because each entity in a corporate group is a separate controller. Organisation-level details can sit in a header, but joint controllers vary by activity and should be recorded against the activities concerned. Common errors include trading names, contact details for individuals who have since left and the omission of an Article 27 representative where one is required.

Where the list of contacts is long, such as a group with many entities or a processor acting for many controllers/clients, it may be more practical to maintain those details in a separate contacts register referenced from the ROPA. That register then forms part of the record and must be kept to the same standard of accuracy and be capable of being produced alongside the ROPA on request. A ROPA that points to a contacts register which is not kept up to date carries a greater risk than an incomplete entry, because the reference gives an appearance of completeness that the underlying document does not support.

5.2. Purposes of Processing

Article 30(1)(b) requires the purposes of the processing. This is the field on which every other field depends because lawful basis, necessity, data minimisation, and retention can only be assessed against a stated purpose. The expected level of detail is a short statement of why the processing happens, specific enough for a reader to judge whether each category of data is needed. A department name is not a purpose and nor is a system. The most common mistakes are describing the technology rather than the purpose, bundling unrelated purposes into one entry and adopting privacy notice wording without checking it against practice.

5.3. Categories of Data Subjects and Personal Data

Article 30(1)(c) requires a description of the categories of data subjects and of personal data. Entries should identify every group whose data is processed, including indirect data subjects such as emergency contacts, witnesses, dependants, and named care recipients, who are frequently overlooked, and should note where children or other vulnerable individuals are involved. Field-by-field listings of personal data are unnecessary, but "various" is not a category. Special category data and criminal offence data should be identified explicitly, as should free text, attachments, call recordings, and paper files, which routinely contain more than the structured fields suggest.

© 2026 Pembroke Privacy

5.4. Categories of Recipients

Article 30(1)(d) requires the categories of recipients to whom personal data have been or will be disclosed, including recipients in third countries or international organisations. Processors are recipients.

Public authorities that receive data in the framework of a particular inquiry are not treated as recipients under Article 4(9). This exception is narrow and covers only non-routine requests, such as a one-off written request from a tax authority during a specific tax audit. Authorities that receive personal data routinely, such as a tax authority receiving payroll submissions, or a police authority where an organisation regularly releases CCTV footage on request, should be recorded as recipients. In practice, an organisation cannot list every regulator or authority that might one day request personal data, but it should record those that have requested data before and those with which contact is expected.

A useful entry covers internal recipients on a need-to-know basis, processors, and independent controllers. Pembroke Privacy recommends recording each recipient's role, because confusing processors with independent controllers is one of the most frequent errors and has consequences for contracts, transparency, and transfers.

5.5. International Transfers

Article 30(1)(e) requires identification of transfers to a third country or international organisation and, for transfers under the second subparagraph of Article 49(1), documentation of suitable safeguards. Recording the mechanism relied on for all third-country transfers, such as an adequacy decision, standard contractual clauses (SCCs) or binding corporate rules, together with any transfer impact assessment (TIA), is good practice that supports compliance with Chapter V along with meeting an Article 30 requirement. The transfers most often missed are remote support access from outside the EEA and sub-processor hosting locations.

5.6. Retention Periods

Article 30(1)(f) requires the time limits for erasure of the different categories of data. An entry should detail how long the data is kept, what triggers that period to start, and which rule in the retention schedule the period comes from. Where a fixed period cannot be set, the entry should explain how the period will be decided instead.

Common mistakes include applying one period to every type of data, stating a period without saying when it starts, and recording periods that differ from the retention schedule or are never actually enforced. Where data is kept longer in practice than the documented period, the gap should be recorded and fixed.

© 2026 Pembroke Privacy

5.7. Security Measures

Article 30(1)(g) requires a general description of the technical and organisational measures referred to in Article 32(1). The word "general" matters. The description should be specific to the activity but should not reproduce configuration detail that would itself create a risk if the record were circulated widely.

© 2026 Pembroke Privacy

5.8. Example of a ROPA Entry

Requirement	Inadequate entry	Proportionate entry
Identity	Group plc, privacy@group.com for every entity	Legal entity name Registered address DPO contact and representative if applicable Joint controllers recorded per activity
Purpose	HR administration	Administer sickness absence, statutory sick pay, medical certification, and return-to-work records
Data subjects	Staff	Employees directors contractors emergency contacts
Personal data	Various	Absence dates duration medical certificates statutory sick-pay eligibility return-to-work notes limited functional information
Recipients	Third parties	HR managers on a need-to-know basis payroll provider medical practitioner employment or labour regulator where required
Transfers	Cloud: SCCs	Within the EEA limited US HR-system support, relying on approved SCCs transfer impact assessment in-region hosting restricted support access
Retention	As required	Employment plus 6 years emergency contacts deleted on replacement or termination identity copies deleted after verification unless legally required

© 2026 Pembroke Privacy

Requirement	Inadequate entry	Proportionate entry
Security	Password protected	MFA role-based HR access encryption field masking audit logs quarterly access review secure deletion

Table 2: Inadequate and proportionate entries for each Article 30(1) requirement

6.

Additional Fields for Privacy Governance

Several matters that are central to compliance are required by other provisions of the GDPR but are not listed in Article 30 and recording them within the ROPA framework is what turns the record from a description into a governance tool. Pembroke Privacy recommends separating the fields Article 30 requires from those that other provisions or good practice add, by grouping the columns under "Mandatory documentation requirements" and "Additional accountability fields" and labelling each column with its legal source.

Supervisory authorities expect the ROPA to be capable of being read as a standalone document,[12] yet much of the supporting analysis, such as DPIAs, Legitimate Interests Assessments (LIAs) and TIAs, is too substantial to reproduce in a register. Pembroke Privacy's recommendation is to record the conclusion in the ROPA, such as the lawful basis relied on or the outcome of DPIA screening, and to link to the specialist record for the reasoning. Table 3 shows, for each field, whether it is legally required elsewhere and how Pembroke Privacy recommends it is maintained.

Field	Legal source	Recommended treatment
Article 6 lawful basis	Article 6; Art. 13(1)(c)	ROPA field
Article 9 condition	Article 9(2)	ROPA field

1 DPC, Guidance Note (n 4).
2

© 2026 Pembroke Privacy

Field	Legal source	Recommended treatment
Criminal offence data condition	Article 10	ROPA field
Privacy notice reference	Articles 12-14	ROPA field
Source of personal data	Article 14	ROPA field
Legitimate Interests Assessments (LIAs)	Accountability evidence	Linked record
Consent records	Article 7(1)	Linked record
Data subject rights handling	Articles 15-22	Link to Process for handling SARs
DPIA	Article 35	ROPA field + linked DPIA
Processors and contracts	Article 28	ROPA field + contract reference
Retention schedule	Article 5(1)(e)	Linked record
Transfer assessments	Articles 44-49	Linked record
Security controls	Article 32	Linked record / security register
Breach history	Article 33(5)	Linked record
Automated decision-making / profiling	Articles 13(2)(f), 22	ROPA field
Ownership, approval, and review	Accountability (Arts. 5(2), 24)	ROPA field

Table 3: Governance fields beyond Article 30 and their recommended treatment

© 2026 Pembroke Privacy

7.

Choosing the Right Level of Detail

The question Pembroke Privacy is asked most often about ROPAs is how much detail is enough. Regulator guidance does not prescribe a single level of granularity, and different approaches can be defensible. What matters is that each entry is specific enough to support the legal analysis and the uses described in 'Purpose and Regulatory Context', and general enough to be maintained.

Getting this distinction right determines whether the ROPA is usable, because each concept answers a different question and only one of them is the natural unit of a ROPA entry.

Much of the confusion arises from treating five different concepts as interchangeable:

- A business process is an end-to-end sequence of work the business recognises, such as recruitment or accounts payable.
- A processing activity is a set of operations serving a defined purpose with a coherent legal profile and is the natural unit of a controller ROPA entry.
- A purpose is the reason the processing takes place.
- A system is the technology used, and it is a means of processing and may support several activities.
- A data flow is the movement of personal data between parties, systems, or locations, which informs the recipient and transfer fields but is usually better captured in data-flow diagrams than as ROPA rows.

A record organised by system tends to become a technical inventory without purposes, while a record organised only by department tends to collapse distinct purposes into a single vague entry. Entries that are too broad can be difficult to understand without detailed knowledge of how the organisation operates. For example, a heading such as "customer data" says little about what processing is actually taking place, why it occurs, who receives the data, or how long it is retained. Likewise, "HR administration" may encompass recruitment, payroll, performance management, sickness absence, employee benefits, disciplinary processes, and many other activities, each of which may have a different purpose, lawful basis, recipient profile, retention period, or Article 9 condition.

Broad headings can still be useful as parent records under which more specific processing activities are grouped. However, they should not be the only level at which processing is recorded. The ROPA should contain activities at a level of detail that allows a reader, including someone with no prior knowledge of the organisation, to understand what processing is taking place and assess the associated compliance requirements.

© 2026 Pembroke Privacy

Excessively granular records create the opposite problem. Treating every system action, screen, or workflow step as a separate activity increases the maintenance burden, duplicates legal analysis across multiple entries, and makes it harder to identify the processing activities that carry genuine compliance or operational risk.

Pembroke Privacy recommends combining operations into a single entry where they share the same purpose, lawful basis, categories of data subjects and personal data, recipients, retention rule, risk profile and owner, and separating them where any of these differs materially. The strongest reasons to separate are:

- A different Article 6 basis or Article 9 or 10 condition.
- Special category or criminal offence data arising in only part of a process.
- Different recipients or international transfers.
- A different retention period or trigger.
- A materially higher risk, including where only part of a process requires a DPIA.
- A different accountable owner.

Documenting employee attendance and leave provides a useful illustration. A single "leave" entry may appear convenient, but it often combines activities with materially different legal and operational characteristics. Paid time off and annual leave may not involve special category data at all, whereas statutory sick leave and absence typically involve health data and may rely on Articles 9(2)(b) and 9(2)(h). Domestic violence leave may require restricted access and a different retention approach. Recording all of these activities under a single heading obscures the differences in purpose, lawful basis, special category processing, access controls, and retention that the ROPA is intended to capture. Table 4 illustrates the spectrum.

Area	Too broad	Too detailed	Proportionate
Sickness absence	HR administration	Separate entries for uploading a medical certificate, approving an absence, and generating an absence report	"Statutory sick leave and absence", with its own purpose, Article 9 condition, recipients, and retention period
Supplier management	Finance: supplier information	Separate entries for supplier creation, bank-detail verification, payment runs, and callback checks	"Supplier onboarding and supplier bank-detail management"
Customer onboarding	Customer data	One entry for each form, workflow step, or portal screen	"Verify identity, complete required checks, and establish the customer relationship"

© 2026 Pembroke Privacy

Area	Too broad	Too detailed	Proportionate
Recruitment	HR administration	Separate entries for receiving CVs, arranging interviews, recording interview notes, and issuing offer letters	"Candidate recruitment and selection"

Table 4: Entries that are too broad, too detailed, and proportionate

Consistency across departments requires more than a shared template. Controlled lists for data subject categories, personal data categories, recipient types, lawful bases, and retention triggers help ensure that similar processing is described consistently. A naming convention for processing activities further improves consistency across the record.

8.

A Nested ROPA Structure

A flat list of processing activities becomes difficult to use once it extends beyond a few dozen rows. Pembroke Privacy therefore recommends organising the record hierarchically. The first level is the branch, business unit or department. The second is the high-level business process. The third is the processing activity or subprocess, which is the level at which the Article 30 fields are completed. A fourth level of subprocess detail should be added only where it improves clarity, legal analysis, ownership, or risk management.

Nesting improves usability because readers navigate a structure that mirrors the organisation, rather than filtering hundreds of rows, and it is consistent with the regulatory recommendation to break the record down by business function.[13] It improves maintenance because a change points to a defined branch, so that a new payroll provider leads the owner to a handful of related entries rather than to everything the HR function does. It also makes ownership visible since each level can carry a named owner.

Organising primarily by department works well where the organisation chart is stable and ownership is clear. Where processes routinely cross departments, as customer onboarding often does across sales, compliance, and operations, organising by end-to-end process may be more natural. In either case, each activity should have a single owning node, with contributing functions recorded as a field and other branches referring to the activity identifier rather than duplicating the entry. Common services such as

1 Ibid.
3

© 2026 Pembroke Privacy

email, document management, identity and access management, the IT service desk and CCTV are best represented in a dedicated enterprise services branch. A business process should be divided into subprocesses only where the separation tests in 'Choosing the Right Level of Detail' are met, to avoid excessive detail.

Unique identifiers make the structure work. A hierarchical scheme, such as HR, HR.03 and HR.03.02, allows parent-child relationships to be read from the identifier itself, while a parent identifier field and owner fields at each level allow the record to be filtered, reported, and reconstructed. The template should also state whether child entries inherit attributes from their parent, so that exported versions are unambiguous. Table 5 shows how the sheet's Human Resources branch could be adapted.

Identifier	Level	Node	Why it is a separate node
HR	1	Human Resources	Organisational branch with an accountable owner
HR.01 to HR.04	2	Employee file Pay and benefits Attendance and leave Mobility & temporary assignments	Business processes recognised by the HR function
HR.03.01	3	Paid time off and annual leave	No Article 9 condition recorded for working-time retention
HR.03.02	3	Statutory sick leave and absence	Health data under Articles 9(2)(b) and 9(2)(h)
HR.03.03	3	Occupational health and reasonable accommodation	External occupational health provider where clinical records held separately by the provider
HR.03.04 to HR.03.07	3	Maternity Paternity, parental, and parent's leave Carers and medical care Force majeure leave	Distinct statutory leave regimes, evidence, and retention
HR.03.08	3	Domestic violence leave	Highly restricted case access due to the sensitive nature of the activity
ENT.01	Reference	HR information system (enterprise service)	Referenced by HR entries rather than restated in each

Table 5: Nested structure for the Human Resources branch of the controller collection sheet

© 2026 Pembroke Privacy

9.

Creating the Initial ROPA

Few organisations need to start from a blank spreadsheet, and most of the required information already exists in fragments across the organisation, either through data flows, vendor procurement due diligence, or other risk management regimes. The difficult task is to assemble it, validate it with the people who do the work and reconcile the inconsistencies that emerge. The Information Commissioner's Office (ICO) suggests beginning with an information audit, questionnaires, meetings with key business functions and a review of existing policies, procedures, contracts, and agreements.[14] A proportionate build typically moves through four phases.

9.1. Scope and Govern

Before gathering information, the organisation should define the scope of the exercise. This should establish which legal entities, functions, and processing activities are in scope, whether the ROPA will cover controller processing, processor processing, or both, the level of detail required, the template and tooling to be used, and the timetable for completion. Responsibility for collecting, compiling, and verifying the data should also be assigned at the outset. The organisation should identify an accountable executive, a privacy lead, and named owners for each department and business process, as ownership agreed early is far easier to maintain than ownership introduced after the record has been completed. Once scope and ownership have been agreed, the initial structure can be mapped from the organisation chart, process maps, and service catalogue, and the relevant stakeholders identified, including IT, procurement, legal, records management, and information security.

9.2. Assemble Existing Evidence

Draft entries should be pre-populated from privacy notices, policies, processor contracts, and the contracts register, DPIAs and LIAs, the retention schedule, system and application inventories, data-flow documentation, the breach log and any earlier ROPA. The organisation should then look for processing that nobody has declared. Reconciling accounts payable records against known processors, reviewing third-party applications connected to the organisation's identity provider, tracing a test access request through the organisation's systems and inspecting physical and archived records each surface processing that self-reporting tends to miss.

1 ICO, Documentation.
4

© 2026 Pembroke Privacy

The organisation should also engage with IT to obtain a list of the applications connected to its identity provider or Single Sign-On (SSO) system. This is the same check described above, and it is the most reliable way to confirm that all vendors are correctly recorded.

9.3. Validate with the Business

Workshops or structured interviews with the people who perform each process should walk through it as it is actually performed: what data comes in and from where, what happens to it, who sees it, where it goes and when it is deleted. Asking "why", until the answer becomes a reason rather than a department, is the most reliable way to arrive at a genuine purpose. The privacy or legal team should then confirm the Article 6 basis for each purpose, any Article 9 or 10 condition and national law provision, and whether an LIA, consent mechanism or DPIA is required. Processors, recipients, transfers, and retention periods should be checked against contracts, sub-processor lists, transfer assessments, and the retention schedule rather than accepted from recollection.

9.4. Reconcile, Assure and Approve

Where sources conflict, for example where a privacy notice does not reflect operational practice or a retention period differs from the retention schedule, the ROPA should record the actual position, and the issue should be documented for remediation. The purpose of the ROPA is to describe what the organisation does, not what it intends to do.

The build process often identifies gaps, inconsistencies, and compliance issues that were not previously visible, and these findings should be recorded in a risk register, with each action assigned an owner, risk rating, priority, and target completion date. The register then becomes the organisation's roadmap for improving its privacy framework.

A risk-based sample of entries should be reviewed against the checklist in Appendix B, with particular attention given to higher-risk processing activities. Owners should confirm the accuracy of their entries, accountable executives should approve their areas of responsibility, and the privacy team should publish the first version of the ROPA with appropriate access controls and change history enabled. A documented issue with a clear owner and remediation plan is more defensible than an entry that conceals or ignores the problem.

© 2026 Pembroke Privacy

10.

ROPA Accuracy Through Embedded Workflows

An annual review asks owners to recall twelve months of change in a single sitting, competes with every other priority in the business and is generally viewed as additional bureaucracy. At best, it records the true position at one point in time. The central recommendation of this paper is to attach ROPA updates to workflows the organisation already operates. Vendor procurement, contract signature, and internal restructuring all pass through a documented workflow with an approval step, so the ROPA can either be an input to the change or be updated as a by-product of the decision, rather than being a separate compliance task. This is data protection by design in practice.

The workflows that generate the most change, and therefore offer the most value from integration, are procurement and supplier onboarding, contract creation and renewal, IT change management and new system approval, project initiation, and DPIA screening. Each can carry a short set of ROPA questions, typically no more than five, asking whether the change introduces a new purpose, new categories of data or data subjects, new recipients or processors, a new location or transfer, or a change to retention. A "no" to every question closes the loop, a "yes" routes the change to the relevant activity owner. Appendix A lists each triggering event in summary form, together with the individual responsible for initiating the update and the evidence to retain.

An example of embedding the ROPA into a workflow is the replacement of a payroll provider. An internal procurement process captures the decision and identifies that a new processor will receive employee data. That single change may affect only a small number of ROPA entries, but those entries can then be reviewed in a targeted and proportionate way. Recipient details may need to be updated, contract references amended, international transfer information reviewed, and any related DPIA reconsidered if the change is material. Where the processor receives special category data, the organisation should also confirm exactly what information will be shared and whether any supporting privacy documentation requires updating. The process owner initiates the changes, the privacy team validates the compliance elements, and the procurement records, signed contract, and any transfer assessment are retained as evidence.

Rather than reviewing an entire business function, the organisation updates only the affected activities, and aligns with the "little and often" maintenance model. Process owners update individual entries when a triggering change occurs, using a short form that captures what changed, why and from when. The privacy team reviews the change and checks any legal analysis. Periodic owner reviews, in which each owner attests that their entries remain accurate, with the frequency set according to the risk of the processing, confirm that nothing has been missed. A risk-based assurance programme, carried out by

© 2026 Pembroke Privacy

someone independent of the process owner, such as the privacy team, internal audit or an external reviewer, tests whether the record matches reality. The annual review becomes a backstop rather than the primary maintenance mechanism, and because most changes will already have been captured, it becomes an exercise in confirmation rather than reconstruction.

11.

How the UK GDPR differs from the EU GDPR

Organisations subject to the UK GDPR face a broadly similar Article 30 obligation. The ICO's guidance sets out substantially the same categories of information for controllers and processors as the EU GDPR, with the main drafting difference being that a third country is any country outside the UK[15]. There are, however, several important differences that organisations should consider when designing and maintaining a UK ROPA.

First, the UK Data Protection Act 2018 requires additional information to be recorded in some circumstances. Where special category or criminal offence data is processed under a Schedule 1 condition that requires an appropriate policy document, paragraph 41 of Schedule 1 requires the organisation to record the condition relied upon, how the processing satisfies Article 6, and whether retention and erasure take place in accordance with that policy document[16]. An appropriate policy document is a UK-specific document that explains the Schedule 1 condition relied upon and the organisation's retention and deletion arrangements for the data concerned. This is a legal requirement with no direct equivalent under the EU GDPR.

15 ICO, What do we need to document under Article 30 of the UK GDPR?

16 Data Protection Act 2018 (UK), Schedule 1, paragraph 41.

17 ICO, Recognised legitimate interest (23 March 2026).

18 DSIT, Data (Use and Access) Act 2025: data protection and privacy changes (27 June 2025).

19 ICO (n 17), noting that the guidance is under review following the Act.

© 2026 Pembroke Privacy

Second, the Data (Use and Access) Act 2025 amends the UK GDPR and the Data Protection Act 2018. The changes most relevant to the ROPA framework are a new "recognised legitimate interest" lawful basis in Article 6(1)(ea), which covers five pre-approved purposes and removes the balancing test for them,[17] a more permissive framework for solely automated decision-making subject to safeguards, revised rules for international transfers, and a new requirement to have a process for handling data protection complaints.[18] These changes affect the lawful basis, automated decision-making, transfer and rights fields of a UK ROPA framework, even where the Article 30 content itself is unchanged. The ICO has indicated that its documentation guidance is under review as a result of the Act,[19] and organisations should monitor it for changes.

Third, organisations operating in both jurisdictions should not assume that a single analysis serves both regimes. Transfers should be assessed separately under each regime even though the European Commission renewed its adequacy decision for the UK in December 2025.[20] Organisations should record the applicable regime for each activity and, where processing is subject to both, record the lawful basis and transfer mechanism separately for each.

12.

Conclusion

Article 30 sets out a short list of information, but the value of a ROPA depends less on its columns than on whether it remains true. An accurate record is proportionate in its detail, structured so that people can find and own their part of it, and maintained as a by-product of decisions the organisation is already making.

The approach set out in this paper combines four elements: a clear distinction between what Article 30 requires and what good governance adds; a nested structure with identifiers and named owners; an initial build grounded in existing evidence and validated with the people who do the work; and a "little and often" maintenance model supported by risk-tiered attestation and independent assurance. None of these elements is prescribed by the GDPR, and regulator guidance leaves room for different approaches. Together, however, they give an organisation a defensible basis for demonstrating accountability, and a record it can rely on when it matters.

17 ICO, Recognised legitimate interest (23 March 2026).

18 DSIT, Data (Use and Access) Act 2025: data protection and privacy changes (27 June 2025).

19 ICO (n 17), noting that the guidance is under review following the Act.

20 European Commission, Adequacy decisions.

© 2026 Pembroke Privacy

For organisations ready to act, the practical starting points are to confirm an accountable executive and named process owners, to map the nested structure from the organisation chart and service catalogue, to attach a short set of ROPA questions to the procurement, change management and DPIA screening workflows, and to schedule a risk-based sample review of existing entries against the checklist in Appendix B.

The ultimate test of a ROPA is whether the organisation can rely on it when it matters. The ROPA should not stand alone. It is most effective as an index to the organisation's wider privacy and information governance framework, connecting privacy notices, lawful basis analyses, DPIAs, contracts, transfer assessments, retention schedules, security controls and breach response, so that each can be found, tested and kept consistent with the others. Organisations that build their record in this way will find that the effort of maintaining it falls while its value, to the privacy team, to the business and to the regulator, rises.

© 2026 Pembroke Privacy

13.

Appendix A: Embedding ROPA Updates in Business Workflows

Workflow	Trigger Event	Responsible Individual	Evidence Retained
Project initiation	New project involving personal data	Project manager	Business case
DPIA screening and assessment	DPIA screening completed or DPIA approved	DPIA author	DPIA screening outcome, DPIA reference
Procurement and supplier onboarding	New supplier with access to personal data	Procurement owner	Due diligence records, supplier Data Processing Agreement
Contract creation and renewal	New agreement, renewal, or sub-processor change	Contract owner	Data Processing Agreement, sub-processor list
New system or technology approval	New system approved or implemented	IT owner	Data Processing Agreement, System approval record
Change management	Change affecting personal data processing	Change requester	Change ticket
Product and service development	New feature, product, or use of personal data	Product owner	Product approval or release record
Data migration	Approved migration or data transfer project	Migration lead	Data mapping, migration records
Records retention review	Retention schedule change or disposal exercise	Records manager	Updated retention schedule
Security risk assessment	Assessment of a system processing personal data	Security owner	Risk assessment

© 2026 Pembroke Privacy

Workflow	Trigger Event	Responsible Individual	Evidence Retained
International transfer assessment	New transfer or change in transfer arrangements	Privacy or legal lead	Transfer Impact Assessment, transfer mechanism
Business restructuring	Reorganisation or ownership change	HR or operations lead	Organisation change record
Mergers and acquisitions	Acquisition, divestment, or integration activity	Corporate development lead	Data Protection Gap Assessment, Due diligence findings
Incident and breach management	Personal data breach or near miss	Incident manager	Incident record, breach log
Process closure or decommissioning	Retirement of a process or system	Process or system owner	Disposal or migration records
Privacy notice changes	Privacy notice updated	Privacy lead	Updated notice

Table 6: Organisational workflows that trigger ROPA reviews

© 2026 Pembroke Privacy

14.

Appendix B: ROPA Quality-Assurance Checklist

#	Check	Result
1	The entry describes a processing activity with a specific purpose.	
2	All applicable mandatory Article 30 fields are completed for the relevant role.	
3	Data subject and data categories are specific, and special category and criminal offence data are flagged.	
4	The Article 6 basis is mapped to each purpose, with any Article 9 or 10 condition and national law provision.	
5	Reliance on legitimate interests or consent is supported by a linked LIA or consent record.	
6	Recipients are complete, processors are named and each recipient's role is identified.	
7	Third countries are named, and the transfer mechanism and any TIA are referenced.	
8	Retention states a period, a trigger, and a schedule reference, and matches actual practice.	
9	The security description is specific to the activity and pitched at a general level.	
10	DPIA screening status is recorded, and high-risk processing has a current DPIA reference.	
11	The entry is consistent with the privacy notice, the contracts register and the system inventory.	
12	An owner, approver, last and next review dates, and review trigger are recorded.	
13	The entry has a unique identifier and parent and is not duplicated elsewhere.	
14	Known gaps are recorded in the gap log with an owner and due date.	

Table 7: ROPA quality-assurance checklist

© 2026 Pembroke Privacy