Workday Submission on the EDPB's Guidelines 02/2026 on Anonymisation

Response to the Public Consultation — 30 October 2026

Workday appreciates the opportunity to comment on the EDPB's Guidelines 02/2026 on Anonymisation (the "Guidelines"), open for public consultation following their adoption on 7 July 2026. Workday is the enterprise AI platform for managing people, money, and agents. Workday unifies HR and Finance on one intelligent platform with AI at the core to empower people at every level with the clarity, confidence, and insights they need to adapt quickly, make better decisions, and deliver outcomes that matter. Workday is used by more than 11,000 organizations around the world and across industries – from medium-sized businesses to more than 65% of the Fortune 500.

Please do not hesitate to contact Barbara Cosgrove, Chief Privacy and Digital Trust Officer, at barbara.cosgrove@workday.com if you have any questions or would like further information.

Workday provides enterprise AI applications for human capital management, financial management and planning to thousands of customers across the EU. In that role, we process data on behalf of our customers and work closely with them to apply privacy-protective techniques. We are glad to see the EDPB take on the difficult task of giving organisations a workable, up-to-date test for when data falls outside the scope of the GDPR.

Good anonymisation guidance gives organisations the confidence to put real investment behind privacy-enhancing technologies. The Guidelines, on the whole, take important steps to deliver that benefit. That said, we think a handful of elements in the current Guidelines need further work before the Guidelines are finalised, namely their:

- treatment of processors, which as drafted could make controllers less willing to share anonymised or minimised data with the vendors that process it on their behalf;
- one-sided build-out of the "reasonably likely to be used" test, where the factors pointing toward risk are spelled out far more fully than the factors pointing away from it;
- framing of the "simplified" approach, which could end up functioning as the default option rather than one of two equally valid paths;
- grouping of "anonymous," "de-identified" and "de-personalised" data under a single warning, as though the three carried the same meaning;
- light coverage of how the No Inference and singling-out criteria apply once AI models and synthetic data enter the picture; and
- silence on how specific de-identification tools map onto the three criteria in practice.

Each point is set out in more detail below, and we would be glad to walk through any of them with the EDPB.

The Guidelines Should Reflect the Realities of Processing Arrangements

The Guidelines provide that, where an entity processes data on behalf of another, whether that data is personal data for the processing entity depends on whether it is personal data for the controller. In practice, this means that information remains personal data for a processor whenever it is personal data for the controller, regardless of the safeguards the processor applies or the limited nature of the data it actually receives.

This does not match how a great many processing relationships actually work, and it removes an incentive the GDPR is otherwise meant to create. Controllers often strip or minimize a dataset before it ever reaches a processor, specifically so that individuals are protected and exposure is limited if something goes wrong. Automatically imputing the controller's status under paragraph 15 is awkward next to the Court of Justice's approach in the EDPS v SRB case, which looked at anonymity from the recipient's own vantage point. In other arrangements, a controller instructs its processor to apply de-identification and transformation techniques itself and isolate the resulting output for narrower purposes such as analytics, benchmarking, or model training. Leaning on the relative approach in paragraph 89 and Example 21 of the Guidelines, the EDPB notes that when a company holds both an aggregate dataset and an internal report that would allow inference, anonymity is only broken from the perspective of those who actually have access to the internal report. We urge the EDPB to explicitly recognise that implementing strict technical isolation and organisational/contractual safeguards (i.e., context controls) means that original production data is not a "means reasonably likely to be used" inside an isolated processing environment. These technical and organisational boundaries protect the irreversibility of the transformation pipeline rather than simply attempting to replace it.

Workday recommends the EDPB:

- Assess anonymity from the processor's own perspective, rather than deriving it automatically from the controller's status, consistent with the Guidelines otherwise contextual, entity-relative approach.
- Confirm that an entity processing data on behalf of another may hold anonymous data where the dataset it actually receives, the instructions it operates under, and the technical and contractual safeguards in place, leave it without a realistic way to identify anyone.
- Give meaningful credit to strong, enforceable contractual and organisational safeguards as one part of a technically sound anonymisation outcome, not a replacement for one.
- Confirm that under the relative approach in paragraph 89 and Example 21, strict technical isolation and access controls ensure that original production data is not "reasonably likely to be used" to re-identify data within an isolated environment.

Clarify "Reasonably Likely to Be Used" – in Both Directions

Data is identifiable, and thus not anonymous, if an individual can be distinguished from others "using means reasonably likely to be used" in a way that allows them to be treated differently. The Guidelines work

through a detailed list of factors that can raise the likelihood of identification, a wide range of entities whose perspective might matter, and the conditions needed to rebut the presumption that people generally follow the law. By contrast, the factors that lower the likelihood of identification are mentioned only in passing, scattered through the text rather than organised into anything comparable. If that imbalance carries through to the final text, many organisations will simply default to treating data as personal, with little gain for individuals and even less reason to keep investing in genuine technical de-identification.

Three specific areas would benefit from clarification. First, the Guidelines should make clear in paragraph 31 that unauthorised or accidental access only counts in the analysis when the facts actually point to a realistic chance of both that access occurring and re-identification succeeding, rather than simply because such a thing, such as an innocent mistake, could happen. Second, a legal prohibition on identification should be presumed effective unless there is concrete evidence that the entities involved, or others like them, have broken a similar rule in similar circumstances. Third, once technical measures have already done the heavy lifting of reducing identifiability, specific, verifiable and enforceable contracts should count for something in the analysis, rather than being treated as largely beside the point.

Workday recommends the EDPB:

- Build out a set of risk-reducing factors to sit alongside the existing risk-increasing ones, covering things like how current and accessible auxiliary data actually is, what an identification attempt would cost in time and expertise, and what access controls are in place and give guidance on how the two sides should be weighed against each other.
- Limit the relevance of unauthorised or accidental access to cases where the circumstances genuinely point toward both access occurring and re-identification succeeding.
- Treat legal prohibitions as carrying real weight, displaceable only by concrete evidence tied to the specific entities involved.

Keep the Simplified Approach Clearly Optional

The Guidelines offer both a contextual approach, which captures the full nuance of the legal test, and a simplified approach, which the Guidelines acknowledge can set a higher bar than the law requires. Workday is glad to see that there is a choice. Our concern is with how the simplified approach might be used in practice: because it is quicker to apply and comes stamped with EDPB guidance, it could end up as the path most organisations default to, rather than one of two equally available options.

Workday recommends the EDPB:

- Spell out that the simplified approach is optional, does not make data personal just by using it in the body of the text (not just footnote 36), and that an organization can switch over to a contextual assessment whenever it wants without being penalised for the switch.

Distinguish "Anonymous" from "De-Identified" and "De-Personalised"

Paragraph 40 of the Guidelines warns against describing data as anonymous, de-identified, or de-personalised when people remain identifiable — but by naming all three terms together, it reads as though they mean the same thing. They do not, in common usage: "anonymous" expresses a legal conclusion that data sits outside the GDPR altogether, while "de-identified" and "de-personalised" are the terms organisations reach for when de-identifying measures have been applied but the resulting data is still personal data, much closer in substance to pseudonymisation. Folding all three into one warning would muddy a distinction that organisations currently use to communicate accurately about how identifiable their data still is.

Workday recommends the EDPB:

- Narrow the warning in paragraph 40 to describing data as "anonymous" or "anonymised" while people remain identifiable, rather than sweeping in every term used to describe a de-identification effort.
- Confirm that "de-identified" and "de-personalised" remain appropriate terms for data that has been through de-identification measures but may still be within the scope of the GDPR, ensuring organisations can accurately describe data subjected to strong technical de-identification and context controls.

Sharpen the No Inference and Singling Out Criteria, Especially for AI Use Cases

Workday welcomes inference as one of the three criteria, and the idea that an inference only counts against anonymity when it is specific and meaningful. Where that line actually falls, though, is not yet clear. The discussion of de-aggregation, AI models and synthetic data is open-ended enough that an organisation could reasonably wonder whether routine analytics or a published statistic trips the criterion any time it might, alongside some other piece of information out there, say something about a person — even in cases where no particular individual is ever realistically at risk of being identified. In addition, singling out is not explained in the Guidelines, yet it is a concept that helps decide whether a dataset that fails No Record Isolation can still count as anonymous. Finally, the Guidelines say little about how the three-criteria test applies once data is being used to train or evaluate an AI model, or to generate synthetic data — exactly the kind of use case where companies are now trying to find privacy-protective paths forward.

Additionally, we welcome the clarification in paragraph 10 that anonymisation does not require deleting the original source data. Building on this, and the standard for statistical generalisation in footnote 43, the Guidelines should clarify the operational reality of AI development. Specifically, if a transformation pipeline

feeding into an isolated machine learning development environment breaks the 1:1 record link—whether by breaking records into narrow, non-unique feature pairs using advanced data transformation, generating synthetic data, or training model weights under the test in footnote 43—and transient pipeline logs (the "keys") are subsequently purged, that output can be considered fully anonymised under EU law. This remains true even while the controller or processor continues to host the original source data in a separate technically isolated production environment.

Workday recommends the EDPB:

- Confirm that a finding which holds true across an entire population — and would apply just as well to someone who was never in the dataset — is a generalisation, not an inference about a specific person, and so does not trip the No Inference criterion.
- Bring the standard for AI models in line with Opinion 28/2024: the test should be whether a model can be made to output someone's personal data using means reasonably likely to be used, not simply whether the model is theoretically capable of saying something about people.
- Explain what singling out actually means, and confirm that it only makes someone identifiable where an entity can both link an isolated record to that specific person and use it to treat them differently.
- Walk through with an example how the three criteria apply to training an AI model, evaluating one, and generating synthetic data.
- Confirm that pursuant to paragraph 10 and footnote 43, the output of a machine learning transformation pipeline (such as synthetic data or model weights) is considered fully anonymised when the 1:1 record link is broken and transient keys are purged, even if the controller or processor retains the original source data in a separate, technically isolated production environment.

Give Practical, Technique-Level Guidance on De-Identification Tools

The three criteria in the new Guidelines give analysts a cleaner conceptual tool, but that hands-on, technique-level detail has not carried over. Companies that have already built de-identification tools into their products are left guessing at how those investments will hold up against the new test, at exactly the moment when regulatory clarity would do the most to encourage further adoption, including in AI-related processing.

Workday recommends the EDPB:

- Publish a reference spanning both the methods covered in Opinion 05/2014 and newer tools like differential privacy, federated learning, secure multi-party computation and synthetic data generation.
- For each one, note in plain terms what it tends to achieve, where it tends to come up short, and which design choices change the outcome, without endorsing any particular technology.

Calibrate Documentation and Reassessment Expectations to Risk

Workday supports documenting how an anonymisation decision was reached and revisiting that decision periodically as techniques move forward and new auxiliary data turns up. What the Guidelines do not do is say how much documentation or review is actually "appropriate" as there is no sense of scale built in. Taken together with the Guidelines' broader expectations around security incidents and safety margins, that open-endedness could tie up significant resources reviewing large or long-retained datasets for little real gain, especially once an original dataset has already been deleted and there is not much left to do in response to a review besides noting that it happened.

Workday recommends the EDPB:

- Scale documentation to risk instead of applying one standard everywhere, weighing how sensitive the data is, how many people are affected, and how broadly it is accessed, and offer a sample template organisations can adapt rather than build from scratch.
- Tie reassessment to specific events rather than an open-ended schedule, such as the data being shared more widely, new auxiliary data surfacing that changes the risk picture, a security incident, or a genuine breakthrough in re-identification methods.

Support a Single, Consistent Standard for Anonymisation Across EU Law

The meaning of "anonymised" is being worked out in several places at the same time: in the European Commission's specification proceedings under Article 6(11) of the Digital Markets Act, in these Guidelines, and in the Digital Omnibus negotiations. For a platform operating across the EU's digital rulebook, it would be hard to manage a world where the same dataset counts as anonymous under one instrument and personal data under another. Finalising these Guidelines without an eye on those other processes also risks the guidance itself needing a rewrite soon after it is adopted, which would recreate the very uncertainty it is meant to resolve.

Workday recommends the EDPB:

- Work with the European Commission and national authorities toward one GDPR-consistent meaning of "anonymised" across EU digital law, including the DMA and the Data Act.
- Plan to revisit these Guidelines once the Digital Omnibus negotiations wrap up and the DMA standard is settled.

Conclusion

Workday appreciates the effort behind modernising the EDPB's guidance on anonymisation and is glad to contribute to that process. The changes outlined above are meant to make the final Guidelines easier to

apply in practice while still holding the line on what it actually means for data to sit outside the scope of the GDPR. We are happy to discuss any of these points further at the EDPB's convenience.