White Label Consultancy

Intra-Group Data Transfers: Closing the Gap Between Data Protection Law and Operational Reality

September 2026

Table of Contents

Executive Summary

Privacy teams inside multinational groups are increasingly asked to answer a deceptively simple question: is our internal data sharing actually compliant? The honest answer, in most organisations, is that nobody has a complete picture. Data moves between entities for HR, IT, customer support, and product delivery in ways that were never formally mapped, let alone documented under a coherent legal framework.

An Intra-Group Data Transfer Agreement (IGDTA) is the instrument that closes this gap: a single contractual and governance framework that regulates how personal data flows between group entities, allocates privacy roles consistently, and centralises the mechanisms required for cross-border transfers. Done well, it replaces a patchwork of ad-hoc arrangements with one document the organisation can actually stand behind under regulatory scrutiny.

This article does not attempt to be a step-by-step drafting manual. Instead, it sets out the level of rigour a defensible IGDTA demands, from the structural decisions it requires to the multi-jurisdictional complexity it must absorb and the pitfalls that make so many existing agreements fragile in practice. The aim is to give Privacy leadership a clear bar against which to assess whether their current framework holds up, and what it would take to close the gap.

1. The Governance Gap: Why Ad-Hoc Intra-Group Sharing Doesn't Scale

Data protection law assigns obligations by role. A controller determines the purposes and means of processing and carries the primary duty set. A processor acts only on the controller's instructions and is bound by a defined set of contractual terms. Two entities that decide together are joint controllers. Two entities that each decide for themselves share data as independent controllers. The framework handles each of these cleanly enough when there are two parties and one relationship to document. It breaks down almost immediately inside a group of any real size.

A multinational group operates as a dense, constantly shifting network: dozens or hundreds of entities exchanging data across HR systems, IT infrastructure, customer platforms, and shared services, often across a dozen or more jurisdictions at once. Those exchanges do not share a single legal shape. A parent receiving workforce data for group reporting is a controller in its own right. A shared services centre running payroll is a processor. Product and marketing functions that genuinely decide together are joint controllers. The same entity occupies a different position in each, and each position brings a different set of contractual terms it must be under.

A single business process compounds this further. Onboarding an employee, resolving a support ticket, or running a group-wide security tool may touch several entities simultaneously, each with a different role, and may cross several jurisdictions with different transfer requirements on the way.

Without a unifying framework, organisations tend to default to one of two failure modes. The **first is silence**: intra-group flows are simply not documented at all, on the assumption that data "staying within the family" carries no legal consequence. An assumption most data protection laws do not support. The **second is fragmentation**: a sprawl of bilateral Data Processing Agreements (DPAs), side letters, and email chains, each drafted at a different time, by different people, under different assumptions, none of which reflect how the group actually operates today.

Both failure modes carry the same underlying risk: in a regulatory inquiry, an audit, or a data subject complaint, the organisation cannot produce a coherent account of who is processing what, on what legal basis, and under what safeguards. An IGDTA exists to prevent exactly that outcome, not as a paperwork exercise, but as the operational backbone that makes the rest of the privacy programme defensible.

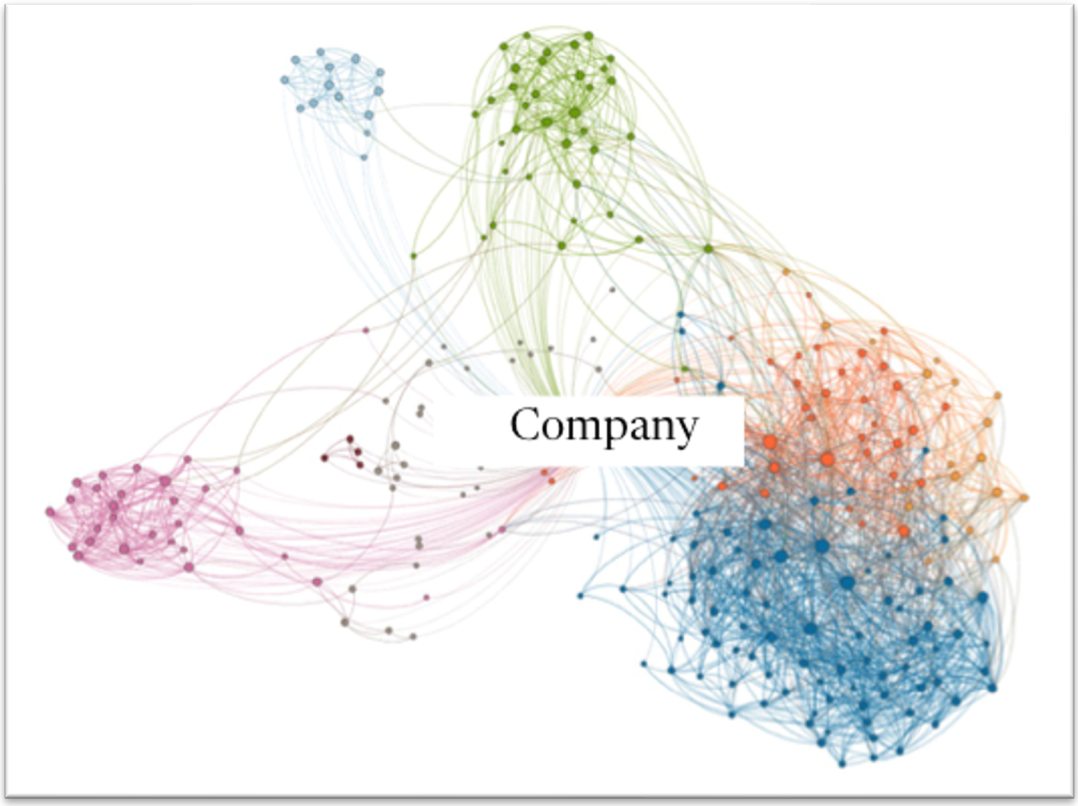

Figure 1: the operational reality inside a multinational group is a dense, many-to-many network - not the simple one-to-one relationship data protection law was built around.

The Cost of Getting This Wrong

The gaps described above surface quickly once a regulator, auditor, or data subject complaint asks a specific question: **who processed this data, on what basis, and under what safeguard?** Across data protection frameworks, whether the GDPR, the Saudi PDPL, the Brazilian LGPD, or comparable regimes elsewhere, supervisory authorities increasingly expect organisations to answer that question at the level of individual processing activities and transfers, not merely at the level of a general privacy policy. Since the invalidation of earlier transfer frameworks and the heightened scrutiny that

followed, supervisory authorities have shown a clear willingness to test whether an organisation's paper trail (its Records of Processing Activities (RoPA), its Transfer Impact Assessments (TIA), its intra-group contracts, actually matches operational reality.

For Privacy teams, this changes what "good enough" looks like. A fragmented set of legacy agreements may have been tolerable when regulatory attention was lighter; it is a considerably harder position to defend today, particularly for groups whose entities span jurisdictions with materially different transfer regimes. An outdated or incomplete IGDTA can quickly become a point of scrutiny during a regulatory inquiry or internal audit, particularly where it reveals inconsistencies with the RoPA or the group's public DPA.

There is also a less obvious cost: internal friction. Without a clear, current record of who is controller and who is processor for a given activity, every new project touching intra-group data, a shared HR platform, a group-wide analytics tool, a security monitoring system, starts by re-litigating the same questions from scratch. A framework that settles those questions once, and keeps the answers current, removes that cost from every project that follows.

What an IGDTA Actually Solves

An IGDTA is a contractual and governance framework used by multinational groups to regulate internal data sharing, allocate processing roles, and centralise the mechanisms needed for international transfers. It is not mandated by name in most data protection laws - the GDPR, for instance, does not prescribe an "IGDTA" as such - but it has become the practical standard for organisations that need a scalable answer to intra-group compliance.

A well-built IGDTA typically centralises four things at once:

1. a current and defensible allocation of controller, joint controller, and processor roles across entities;
2. the transfer mechanisms (such as Standard Contractual Clauses) needed to legitimise cross-border flows;
3. a description of processing activities consistent with the group's RoPA; and
4. a mechanism for new entities to formally adhere to the framework as the group changes.

It is worth being precise about what an IGDTA replaces, and what it does not. The table below situates it against adjacent instruments Privacy teams are likely to already have in place.

Instrument	Typical Purpose	Relationship to the IGDTA
Bilateral intra-group DPA	Governs a single controller-processor pair within the group	Superseded or absorbed by the IGDTA once a group-wide framework is in place
Joint Controller Agreement	Allocates responsibilities where two or more entities jointly determine purposes and means	Can sit alongside or be incorporated into the IGDTA for specific joint-controller scenarios

Vendor / third-party DPA	Governs data sharing with an external processor outside the group	Out of scope for the IGDTA, which is limited to intra-group flows
Standard Contractual Clauses (SCCs)	Provide the legal basis for a specific cross-border transfer	Incorporated into the IGDTA as schedules, rather than negotiated separately for each flow

The distinction matters because organisations sometimes attempt to stretch an existing bilateral DPA template across the whole group, or conversely, assume the IGDTA needs to absorb every vendor relationship. Neither works: the IGDTA's value lies precisely in being scoped to intra-group flows, with clean interfaces to the instruments that govern everything else.

2. How a Defensible IGDTA Is Structured

In practice, an IGDTA is rarely one flat document. A structure that scales well tends to separate a stable core - definitions, general obligations, liability principles, and the adherence mechanism - from a set of schedules designed to be updated independently: an entity and role schedule, a processing activity annex mapped to the RoPA, one or more transfer schedules organised by jurisdiction, and a Technical and Organisational Measures (TOMs) annex referencing the group's actual security measures rather than a generic template.

Four structural questions determine whether an IGDTA will hold up in practice, rather than becoming a document that is signed once and never revisited. Each targets a different dimension of the framework: who is actually in scope and how that is kept current; how entities are grouped into legal roles tied to real processing activity; whether the drafting reflects how the organisation actually operates; and how liability and incident response are allocated. Privacy teams evaluating their own framework (or assessing a proposal from external counsel to build one) should expect all four to be addressed explicitly.

1. Which entities are actually in scope, and how is that kept current?

The starting point is rarely the org chart. Large groups typically hold a raw list of active legal entities that runs into the hundreds, most of which turn out to have no relevant processing activity at all such as holding companies, dormant entities, special-purpose vehicles. Scoping an IGDTA properly means systematically narrowing that list to entities genuinely engaged in personal data processing, validated with stakeholders who understand what each entity does in practice (Company Secretary functions, regional operational leads), rather than inferred from a corporate register.

The harder problem is keeping that list current. Groups restructure constantly through acquisitions, divestments, mergers, new market entities and an IGDTA that cannot absorb those changes without a full renegotiation becomes stale within a year. A defensible framework builds in an adherence mechanism: a defined, low-friction process by which a new entity formally joins the

agreement, rather than relying on someone remembering to update a schedule.

Scoping exercise on a recent multi-jurisdictional engagement

Active entities on the corporate register	**150**
Formally excluded - no relevant processing activity	**85**
Confirmed in scope for the IGDTA	**65**

Figure 2: scoping exercise on a recent multi-jurisdictional engagement - narrowing the raw register down to a confirmed in-scope population.

2. **How are entities grouped into roles, and how is that tied to real processing activity?**

Legal roles are rarely obvious from an entity's position in the corporate structure. Two entities that look identical on an org chart may play entirely different roles: one might operate its own HR systems and act as a controller for employee data, while another exists purely to provide shared IT infrastructure and acts as a processor for the rest of the group.

The practical answer is a role-based grouping model rather than an entity-by-entity one: **entities are assessed against a consistent set of operational indicators: does it have its own employees, does it operate systems independently, does it deliver services to other group entities, does it process customer data in its own capacity?** Based on these indicators, entities are grouped into defensible categories such as controllers, processors for internal operations, and processors or sub-processors for products or services. This is what allows an agreement covering a hundred-plus entities to remain a coherent document rather than a hundred-plus individual annexes.

That grouping should stay anchored to the organisation's RoPA rather than developed in parallel. Where the RoPA already exists in a GRC or privacy management platform, the processing activity list in the IGDTA should be traceable back to it (the same categories of data subjects, the same categories of personal data, the same purposes) so the two documents never drift apart, and one cannot be produced to a regulator while contradicting the other.

Illustrative role-grouping model used to move from an entity-by-entity list to a scalable, defensible structure

G1 Controllers	→	G2 Processors - internal operations	→	G3 Processors - products	→	G4 Sub-processors - products

Figure 3: an illustrative four-group role-allocation model moving from an entity-by-entity list to a structure that scales.

3. **Does the drafting reflect how the organisation actually operates?**

The single most common reason IGDTAs fail in practice is that they are drafted as if the organisation were simpler than it is: liability provisions that assume a clean controller-processor split when the reality is more layered, sub-processing chains that are not reflected in the text, or Technical and Organisational Measures that were copied from a template and never reconciled with what the group's public-facing DPA actually commits to.

A defensible agreement is built from an honest assessment of the previous version, if one exists: what provisions caused friction in practice, where the onboarding mechanism for new entities broke down, and what internal positions on liability or transfer risk need to be reflected rather than assumed. The goal is not contractual elegance for its own sake. It is a document the organisation can actually operate under, year after year, as entities and activities change.

4. **How are liability and incident response actually allocated?**

Liability allocation deserves separate attention because it is where the abstraction of "intra-group cooperation" meets the reality of separate legal entities, often in separate jurisdictions with their own boards and own financial exposure. A single group-wide indemnity cap, applied uniformly regardless of an entity's role or the sensitivity of what it processes, tends to satisfy nobody in practice: it either under-protects entities handling higher-risk data or becomes commercially unworkable for smaller entities that never agreed to carry group-wide exposure.

The same care applies to breach notification. An IGDTA should specify, concretely, how an incident discovered by one entity is escalated to the entities actually responsible for regulatory notification and to data subjects, on what timeline, and who bears the cost of the response. Left unaddressed, this is consistently one of the first gaps to surface during an actual incident, precisely the moment an organisation can least afford to be improvising.

3. Who Needs to Be at the Table

An IGDTA that is built by Privacy and Legal in isolation, without input from the functions that actually generate the data flows, tends to be theoretically sound and practically wrong. The people who understand how data actually moves within the group sit in Legal, HR, IT, and Finance (not exclusively in the Privacy

function) and each brings a different, necessary piece of the picture.

- **Legal and Corporate/Company Secretary functions** typically hold the clearest picture of the group's actual entity structure, including which entities are active, dormant, or in the process of being wound down. They are often the only function positioned to own the entity list on an ongoing basis.

- **HR** is usually the function that can speak concretely to how employee data moves across entities: shared HR information systems, cross-border assignments, group-wide payroll or benefits platforms, all of which generate intra-group flows that Privacy would otherwise have to reconstruct from scratch.

- **IT** holds the systems inventory, including which entity operates a given platform, where it is hosted, and which other entities consume it as a shared service. This information directly determines whether an entity is a controller, a processor, or a sub-processor for that activity.

- **Finance and shared-services functions** can identify centralised processing hubs, such as shared service centres handling group-wide functions like billing, procurement, or expense management, that are easy to miss if the mapping exercise starts purely from a privacy or legal register.

In practice, the mapping exercise behind an IGDTA works best as a structured series of targeted conversations with each of these functions, rather than a single questionnaire sent out broadly and hoping for complete answers. The quality of the resulting agreement tends to track directly with how well this cross-functional input was gathered, not with how polished the final drafting is.

4. The Multi-Jurisdictional Layer: Why One Template Rarely Works Globally

Everything above becomes materially harder once a group operates across more than a handful of jurisdictions - the norm, not the exception, for organisations that need an IGDTA in the first place. Across frameworks, transfer mechanisms, role terminology, and even the definition of a "transfer" can vary, so a document built around one regulatory model tends to fail quietly elsewhere.

Region / Framework	What Changes for an IGDTA	Practical Implication
EU/EEA & UK GDPR	Mature SCC modules, adequacy decisions, and an established regulatory posture on intra-group transfers	Often the reference model, but UK and EU SCC mechanics diverge and must be tracked as separate schedules
Switzerland FADP	EU SCCs are recognised under the Swiss FADP, subject to the	EU SCCs can generally be used for Swiss transfers, but the agreement

	adaptations necessary for Swiss law.	should include the necessary Swiss-law adaptations
Saudi Arabia, DIFC, ADGM PDPL	Each has published its own formal SCCs (Saudi SCCs, DIFC SCCs, ADGM SCCs), structurally similar to the EU model but legally distinct instruments, not interchangeable with EU SCCs (ADGM being the exception, via its EU SCC Addendum)	EU SCCs cannot simply be reused; the transfer schedule needs the jurisdiction-specific SCC template, typically paired with a local transfer risk assessment
UAE Federal & Bahrain PDPL	Both permit contractual safeguards, but neither has published its own standalone SCC template; transfers often rely on adequacy lists or case-by-case regulator approval	The transfer schedule needs a bespoke contractual safeguard or BCR route, rather than a ready-made local SCC to plug in
LatAm (e.g., Brazil, Argentina, Colombia, Peru, Uruguay)	Several jurisdictions (e.g., Colombia, Peru, Uruguay) have formally adopted the Ibero-American Data Protection Network (RIPD) model clauses, while Brazil and Argentina maintain their own distinct SCC frameworks alongside or instead of the RIPD model	RIPD-based clauses can often serve as a common regional baseline (including in Argentina, which recognises them alongside its own 2016 clauses); Brazil is an exception, requiring its own ANPD template rather than the RIPD model
APAC (e.g., Singapore PDPA, China PIPL, Australia Privacy Act)	Transfer regimes range widely in formality: from Singapore's consent/contractual approach, to China's PIPL requiring a government-filed standard contract (or security assessment above certain thresholds) plus data-localisation defaults, to Australia's accountability-based model with no prescribed transfer instrument at all	A single global transfer schedule will not travel across APAC as-is; jurisdiction-specific carve-outs are needed, and China in particular may require a standalone filing rather than a contractual add-on

The practical consequence is that a genuinely global IGDTA is rarely a single uniform schedule. It is usually a common core with shared definitions, role allocation logic, governance and adherence mechanics paired with jurisdiction-specific transfer schedules that are built, and revisited, on their own timeline as laws in each region evolve. Treating every jurisdiction as a GDPR variant is one of the more common ways an otherwise well-structured IGDTA becomes indefensible in a specific market.

5. Lessons from the Field: A Recent Multi-Jurisdictional Scoping Exercise

The structural questions above are not abstract. They mirror, closely, a scoping and drafting exercise recently run for a multinational client operating across more than a dozen jurisdictions, and the shape that exercise took is a useful illustration of how the theory plays out in practice.

The starting point was a roughly 150 active legal entities pulled from the corporate register (the figure behind Figure 2 above). Very few of the client's own stakeholders could say, without checking, which of those entities actually processed personal data in any meaningful sense. A series of structured conversations - not a survey sent to the whole group - narrowed that list considerably: 85 entities were formally excluded for lack of relevant processing activity, a handful were held as "maybe" pending further validation, and the rest were confirmed in scope. The previous IGDTA was reviewed at this stage too; its party list was too out of date to rely on, but it was still useful for assessing whether the existing onboarding mechanism for new entities had actually worked in practice.

From there, the work split into three parallel workstreams:

- **Entities and roles.** Each in-scope entity was assessed against a consistent set of operational indicators - did it have its own employees, did it operate systems independently, did it deliver services to other group entities, did it process customer data in its own capacity. The answers grouped the entities into four defensible categories: controllers, processors for internal operations, processors for products, and sub-processors for products (the G1-G4 structure shown in Figure 3) - replacing what would otherwise have been well over a hundred individual entity-by-entity annexes.

- **Processing activities.** Rather than build a parallel description of what each entity did, the processing-activity list was anchored directly to the client's existing Record of Processing Activities (RoPA), exported from its GRC platform and restructured into a usable activity-mapping framework. Gaps and inconsistencies in that export still needed validating with the client but starting from the RoPA rather than beside it kept the two documents aligned from day one.

- **Cross-border transfers.** No formal transfer impact assessment framework existed at the outset. Building on prior comparative research, the team ran a country-by-country legal analysis and, working through a risk-based lens with the client, confirmed which jurisdictions needed additional transfer mechanisms. The result:

updated existing transfer mechanics, and new schedules for Latin America.

By the end of the project, the IGDTA had been comprehensively revised: an updated entity list and role allocation, processing activity descriptions traceable to the RoPA, TOMs aligned with the group's public-facing DPA, and a transfer framework built to absorb the next jurisdiction rather than requiring a fresh renegotiation. None of this came from more elegant drafting - it came from doing the mapping work in the right order before a single clause was written.

6. Common Pitfalls That Undermine an Existing IGDTA

Privacy teams reviewing an agreement already in place, whether inherited from a prior compliance exercise or drafted internally under time pressure, tend to encounter the same handful of weaknesses.

- **Outdated transfer mechanics.** Schedules still built on pre-2021 Standard Contractual Clauses, or referencing frameworks that have since been invalidated, rather than the current SCC modules and the supplementary transfer impact assessments regulators now expect as standard practice.

- **A stale entity list with no adherence process.** The schedule of parties reflects the group as it existed at signing, with no defined mechanism for new entities to join - so every acquisition or restructuring quietly falls outside scope, often unnoticed until an audit asks the question directly.

- **Role allocation that no longer matches operational reality.** Entities recorded as processors take on independent processing activities, or vice versa, without the agreement being updated - typically because the change happened at business level with no process for flagging it back to Legal or Privacy.

- **TOMs and RoPA references that have drifted from the documents they were meant to mirror.** The IGDTA becomes internally inconsistent with the group's own privacy documentation. This is precisely what a regulator or auditor is likely to test first, since it is the fastest way to establish whether the organisation's paper trail can be trusted.

- **Liability and breach-response provisions that were never stress-tested.** Clauses that read reasonably in the abstract but have not been checked against how an actual incident, or an actual dispute between two group entities, would unfold in practice.

- **Treating the IGDTA as a one-time deliverable rather than a living framework, with no defined trigger for revisiting it as the group changes.** This is addressed in the next section.

7. When an IGDTA Needs to Be Revisited

An IGDTA that was fit for purpose at signing does not stay that way indefinitely. A framework built on the core-and-schedules architecture described earlier should absorb

most of these events as a contained update to a single schedule. Only a narrow set of triggers justifies reopening the core agreement itself.

- **Mergers, acquisitions, or divestments** change the population of in-scope entities. This is the most common trigger, and the one most likely to be missed if ownership of the entity list sits with a corporate function that has no visibility into privacy obligations. Provided the adherence mechanism is well designed, this should be resolved entirely within the entity and role schedule, with no need to touch the core agreement.

- **Entry into a new jurisdiction** without an existing transfer schedule calls for a new schedule to be added, not a renegotiation of the agreement's general terms. The same applies when an existing jurisdiction's transfer regime changes materially, for example through a new adequacy decision, a new set of standard clauses, or an invalidated transfer mechanism: the fix belongs in that jurisdiction's transfer schedule.

- **A shift in how a shared service is delivered internally,** such as a function moving from a regional entity to a centralised shared-services entity, changes the underlying role allocation even though nothing about the service itself has changed from the end user's perspective. This is again a schedule-level update: the entity and role schedule needs to reflect the new allocation, while the core agreement's role definitions remain unaffected.

- **Findings from an internal audit, a regulatory inquiry, or an actual security incident** are the one category of trigger that can go either way. Most findings point to a schedule that has fallen out of sync with the RoPA or the TOMs annex, which is a contained fix. Occasionally, though, they expose a weakness in the core agreement itself, most often in the liability or breach-notification provisions discussed earlier. Distinguishing which of the two is the case is the first step in scoping any review triggered this way.

- **A scheduled review, run on a fixed cycle independent of any trigger**, is what catches the drift the four triggers above don't: changes that accumulate gradually rather than arriving as a single identifiable event – an entity whose activity shifted without a formal restructuring, a TOMs annex that fell behind an actual change in security controls, a transfer schedule that no longer matches current data flows. It is a backstop, not a replacement for the trigger-based monitoring.

Organisations that track this distinction, rather than treating every trigger as grounds for reopening the whole document, and that pair it with a review on a fixed cycle rather than waiting on triggers alone, are the ones able to keep an IGDTA current without turning every change in the group into a full renegotiation.

8. Closing the Gap: How White Label Consultancy Supports This Work

The value of an IGDTA lies less in how well it is drafted than in whether it is built to keep pace with a group that never stops changing. The frameworks that hold up are not the most elegantly drafted; they are the ones built to absorb a new entity, a new jurisdiction, or a new regulatory requirement without falling apart in the process.

Most groups already have the raw material: a RoPA, a set of legacy contracts, a rough sense of who processes what. What is usually missing is the discipline to turn that material into a framework that a regulator, an auditor, or a new business partner can actually be shown. That gap is where an experienced partner earns its place, not by producing a longer document, but by asking the right questions before drafting begins.

White Label Consultancy supports multinational groups end-to-end on exactly this work: scoping and mapping the corporate group, structuring defensible role allocations, running the cross-border transfer analysis across European, Gulf, Latin American and Asia-Pacific regimes, and drafting an agreement built to be maintained, not just signed. Our team combines GDPR and GCC-region data protection expertise with hands-on experience running exactly this kind of multi-jurisdictional mapping exercise for complex groups.

Authors

Dr. Magdalena Góralczyk

Data Protection Partner

White Label Consultancy

np@whitelabelconsultancy.com

Nauani Benevides

Consultant

White Label Consultancy

nbe@whitelabelconsultancy.com

Dr. Arina Kostina

Consultant

White Label Consultancy

ako@whitelabelconsultancy.com

Contact

hello@whitelabelconsultancy.com
+45 71 74 74 54

NORWAY

Main Office Oslo

+47 4141 2168

Fjordalléen 16
0250 Oslo

DENMARK

Copenhagen Office

+45 71 74 74 65

Dampfærgevej 27
2100 København Ø

POLAND

Warsaw Office

+48 515 07 99 77

Ul. Marszałkowska 58/15
00-545 Warszawa

UNITED ARAB EMIRATES

Dubai Office

+971 50 4536616

Dubai World Trade Center
Dubai

WHITELABELCONSULTANCY.COM