

BEYOND ONE-SIZE COMPLIANCE

Delivering Across Borders and Regulators

A Practitioner's Assessment of Evidence, Approaches, and Recommendations for Cross-Border Regulatory Complexity

Giovanni Leonardi

May 2018



Contents

1. Executive Summary

2. The Scale of the Problem

3. Why the Global-First Model Fails

3.1 The Specification Problem

3.2 The Sequencing Problem

3.3 The Governance Problem

4. The Evidence from Practice

4.1 Pattern 1: Regulatory-Led Architecture

4.2 Pattern 2: Jurisdictional Autonomy Within Portfolio Guardrails

4.3 Pattern 3: Regulatory Relationship as a Delivery Dependency

4.4 Pattern 4: Cross-Jurisdictional Conflict Resolution

5. Recommendations

5.1 Recommendation 1: Conduct a Full Regulatory Landscape Assessment Before Design

5.2 Recommendation 2: Design from the Constraint Inward

5.3 Recommendation 3: Establish a Jurisdictional Delivery Model

5.4 Recommendation 4: Integrate Regulatory Engagement into the Programme Plan . .

5.5 Recommendation 5: Build Portfolio-Level Conflict Resolution

5.6 Recommendation 6: Plan for Ongoing Regulatory Divergence

6. Conclusion

3

3

4

4

4

5

5

5

6

6

6

7

7

7

7

7

8

Delivering Across Borders and Regulators

Portfolio · No. 87

Originally written May 2018 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The portfolio that treats jurisdictional variation as a late-stage complication will discover that compliance is not a layer you add — it is a geometry that shapes everything beneath it.”

1. Executive Summary

Organisations operating across multiple countries face a challenge that is widely acknowledged but poorly addressed: every jurisdiction has its own regulatory framework, and those frameworks rarely align. A programme that is compliant in one country may be non-compliant in the next. A design decision that satisfies one regulator may directly contradict the requirements of another. A timeline that accommodates one jurisdiction's consultation cycle may collide with another's enforcement deadline.

This white paper examines the structural forces that make multi-jurisdictional programme delivery so consistently difficult, draws on the observable patterns across sectors that have grappled with this challenge — financial services, energy, telecommunications, and pharmaceuticals — and sets out a defensible approach to portfolio-level governance that treats jurisdictional variation as a design parameter rather than a late-stage complication.

The central argument is that the dominant model — a single global programme with local compliance bolt-ons — is structurally inadequate. The organisations that deliver successfully across jurisdictions do so by inverting the relationship between global design and local regulation: they design from the regulatory constraint inward, not from the global standard outward.

2. The Scale of the Problem

The challenge is not theoretical. Consider the regulatory landscape facing a financial services organisation operating across the European Union, the United States, and the Asia-Pacific region in 2018. In data protection alone, the EU's General Data Protection Regulation has just come into force, imposing a harmonised but demanding framework across twenty-eight member states. The United States has no equivalent federal framework, relying instead on a patchwork of sector-specific and state-level legislation. Australia has recently strengthened its Privacy Act. Singapore's Personal Data Protection Act operates under different principles again. Japan, South Korea, and Hong Kong each

have their own regimes with their own enforcement philosophies.

This is one regulatory domain in one sector. Multiply it across prudential regulation, conduct regulation, anti-money laundering, sanctions, operational resilience, outsourcing, and consumer protection, and the combinatorial complexity becomes clear. A large multinational programme may face hundreds of distinct regulatory requirements across dozens of jurisdictions, with no central authority to arbitrate conflicts between them.

The organisations that underestimate this complexity share a recognisable pattern of failure: they design a global solution, discover jurisdictional conflicts during implementation, attempt to resolve them through exceptions and workarounds, and end up with a fragmented landscape that is neither globally consistent nor locally compliant.

3. Why the Global-First Model Fails

The instinct to design globally and adapt locally is understandable. It promises efficiency, consistency, and economies of scale. It aligns with how most multinational organisations think about programme delivery: define the target state centrally, then roll it out across geographies with local adjustments.

The problem is that regulatory requirements are not adjustments. They are constraints. And constraints do not layer neatly on top of a design — they shape the design itself.

3.1 The Specification Problem

A global programme typically produces a single set of requirements, a single architecture, and a single target operating model. Local regulatory requirements are then mapped against this global specification, and gaps are identified for local remediation.

This approach fails because it treats regulation as a gap to be closed rather than a parameter to be designed for. By the time the global specification exists, hundreds of design decisions have already been made — about data flows, about process sequences, about technology platforms, about organisational structures — and many of those decisions implicitly assume a regulatory environment that does not exist uniformly across the delivery geography.

3.2 The Sequencing Problem

Global programmes typically sequence their delivery by capability or by technology component: build the platform first, then deploy the processes, then train the people. Multi-jurisdictional delivery often requires the opposite: sequence by jurisdiction, because the regulatory approval process in each country has its own timeline, its own dependencies, and its own prerequisites that may have nothing to do with the programme's technical readiness.

- A regulator in one jurisdiction may require a formal impact assessment before any technology change is deployed.

- Another may require a period of parallel running that overlaps with the planned cutover window.
- A third may require regulatory approval of the target operating model before implementation begins – a process that can take six to twelve months.

Sequencing by capability in a multi-jurisdictional context produces a delivery plan that looks efficient on paper and disintegrates on contact with regulatory reality.

3.3 The Governance Problem

Global programme governance typically reports through a single Programme Board or Steering Committee. Multi-jurisdictional delivery requires governance that can accommodate fundamentally different risk appetites, different regulatory relationships, and different definitions of what "done" means in each country.

A decision that is routine in one jurisdiction may be a regulatory event in another. A risk that is accepted at the global level may be unacceptable to a local regulator. A timeline commitment made to the global steering committee may be impossible to honour in a jurisdiction where the regulator has not yet approved the approach.

The portfolio that treats jurisdictional variation as a late-stage complication will discover that compliance is not a layer you add – it is a geometry that shapes everything beneath it.

4. The Evidence from Practice

Across the sectors that routinely deliver multi-jurisdictional programmes, a set of patterns has emerged that distinguishes successful delivery from the more common pattern of delay, rework, and fragmentation.

4.1 Pattern 1: Regulatory-Led Architecture

The organisations that deliver successfully start with the regulatory constraints, not with the target state. They map the full regulatory landscape across all delivery jurisdictions before a single design decision is made. They identify the binding constraints – the requirements that are non-negotiable in at least one jurisdiction – and design the global architecture to accommodate them from the outset.

This approach is more expensive in the design phase. It requires regulatory expertise earlier in the programme lifecycle, and it often produces a global architecture that is more complex than a jurisdiction-blind design would be. But it avoids the far greater cost of discovering architectural conflicts during implementation, which is the point at which remediation is most expensive and most disruptive.

4.2 Pattern 2: Jurisdictional Autonomy Within Portfolio Guardrails

Successful multi-jurisdictional portfolios give local delivery teams genuine autonomy over how they meet regulatory requirements within their jurisdiction, while maintaining portfolio-level guardrails on architecture, data standards, and interoperability.

This is a harder governance model than either full centralisation or full federation. It requires clear boundaries: what is mandated globally (architecture principles, data taxonomy, security standards) and what is delegated locally (process design, regulatory engagement, implementation sequencing). The boundaries must be explicit, documented, and enforced — not left to negotiation on a case-by-case basis.

4.3 Pattern 3: Regulatory Relationship as a Delivery Dependency

In every jurisdiction, the programme's relationship with the local regulator is a delivery dependency, not an administrative formality. The programmes that manage this well treat regulator engagement as a workstream with its own plan, its own milestones, and its own risks — integrated into the programme plan at the same level as technology delivery or organisational change.

The programmes that manage it poorly treat regulatory approval as a gate that will open when the programme is ready. This assumption is almost always wrong: regulators have their own timelines, their own review processes, and their own priorities, and they are under no obligation to align them with the programme's delivery schedule.

4.4 Pattern 4: Cross-Jurisdictional Conflict Resolution

Conflicts between jurisdictional requirements are inevitable. Two regulators may impose contradictory requirements on the same data flow. Two jurisdictions may have incompatible definitions of a key concept. Two regulatory timelines may be physically impossible to satisfy simultaneously.

The organisations that handle this well have an explicit conflict resolution mechanism at the portfolio level. This is not a matter of choosing which regulator to satisfy — that is rarely an option. It is a matter of identifying design solutions that satisfy both requirements, or of structuring the delivery so that each jurisdiction receives a variant that meets its specific requirements while maintaining the maximum possible consistency with the global design.

“The real skill in multi-jurisdictional delivery is not managing complexity — it is deciding which complexities to absorb into the design and which to resolve through structural separation.”

5. Recommendations

Based on the evidence from practice, this paper recommends the following approach for organisations facing multi-jurisdictional programme delivery.

5.1 Recommendation 1: Conduct a Full Regulatory Landscape Assessment Before Design

Before any design work begins, commission a comprehensive mapping of the regulatory requirements across all delivery jurisdictions. This assessment should identify not only the current requirements but also the regulatory pipeline — forthcoming changes that will take effect during the programme's delivery window. The output should include an explicit identification of binding constraints and cross-jurisdictional conflicts.

5.2 Recommendation 2: Design from the Constraint Inward

Use the binding constraints identified in the landscape assessment as the starting parameters for global architecture and design. The global design should accommodate the most demanding jurisdictional requirement in each domain, not the average. Where this produces unacceptable complexity or cost, the decision to accept a lower standard in specific jurisdictions should be explicit, risk-assessed, and approved at the portfolio level.

5.3 Recommendation 3: Establish a Jurisdictional Delivery Model

Structure the portfolio with explicit jurisdictional delivery teams that have authority over local implementation within portfolio guardrails. Define the guardrails clearly: what is globally mandated and what is locally delegated. Ensure each jurisdictional team includes regulatory expertise as a core capability, not an advisory function.

5.4 Recommendation 4: Integrate Regulatory Engagement into the Programme Plan

Treat regulator engagement in each jurisdiction as a first-class delivery workstream. Map the regulatory approval process, identify the dependencies between regulatory milestones and delivery milestones, and plan accordingly. Do not assume that regulatory approval will align with the programme's preferred timeline.

5.5 Recommendation 5: Build Portfolio-Level Conflict Resolution

Establish an explicit mechanism for resolving cross-jurisdictional conflicts at the portfolio level. This mechanism should have the authority to make binding design decisions, the regulatory expertise to assess compliance implications, and the delivery perspective to evaluate implementation feasibility. It should produce documented decisions with clear rationale, because regulators in affected jurisdictions may ask to see them.

5.6 Recommendation 6: Plan for Ongoing Regulatory Divergence

The regulatory landscape will continue to evolve after the programme delivers. Design the target state to accommodate ongoing jurisdictional variation, not just the variation that exists at the point of delivery. This means building configurability into systems, modularity into processes, and regulatory intelligence into the operating model that will sustain the delivered capability.

6. Conclusion

Multi-jurisdictional programme delivery is not a variant of single-jurisdiction delivery with extra compliance work. It is a fundamentally different problem that requires a fundamentally different approach to architecture, governance, sequencing, and stakeholder management.

The organisations that succeed treat jurisdictional variation as a first-order design parameter — as fundamental to the programme's architecture as the technology platform or the operating model. The organisations that struggle treat it as a local adaptation problem to be solved after the global design is complete.

The evidence from practice is clear: the cost of designing for jurisdictional variation from the outset is a fraction of the cost of discovering it during implementation. The recommendations in this paper are not novel — experienced practitioners in multinational organisations will recognise every one of them. But they are consistently underweighted in programme planning, and the result is a pattern of delay, rework, and regulatory friction that is as predictable as it is avoidable.