Isle of Man Government

DIGITAL ISLE OF MAN

CONSULTATION PAPER

Data Asset Register and Data Asset Registrar

Architecture, Governance, Classification, Operation and Technical Design

2026 | CONSULTATION PAPER

Digital Isle of Man | Department for Enterprise

CONSULTATION DRAFT

NOT FOR LEGAL RELIANCE

This White Paper is a consultation draft issued pursuant to the Foundations (Amendment) Bill 2025. It does not constitute legal advice and should not be relied upon as a statement of current law. All proposals are subject to change following the consultation process and the legislative and regulatory process.

Issued by:

Digital Isle of Man

Department for Enterprise

Isle of Man Government

Author:

Aga Strandskov

Head of Data Strategy

Digital Isle of Man

Date issued:

March 2026

Table of Contents

1. Executive Summary

Data has become one of the most economically significant resources of the twenty-first century, yet it remains one of the least governed. Where land, securities, intellectual property and other asset classes benefit from centuries of registry design, title systems and transactional frameworks, data operates with no formal registration mechanism, no standardised governance architecture, and no reliable means by which third parties can verify claims of stewardship, provenance or rights. Organisations hold vast quantities of data but typically cannot answer the most basic questions about it: what data assets do they hold, who has rights over them, what can lawfully be done with them and what are they worth.

The Foundations (Amendment) Bill 2025, now progressing through Tynwald, proposes to change this. The Bill introduces the Data Asset Foundation (DAF) regime — a pioneering statutory framework that establishes data as a registrable personal property right. At its operational heart sit two new institutions: the Data Asset Register (the Register) and the Data Asset Registrar (the Registrar).

This White Paper sets out the proposed architecture, legal foundations, operational mechanics, governance structure, classification system and policy rationale for the Register and Registrar. It is published for public consultation prior to the making of regulations and subsidiary instruments under the Bill.

The Core Proposition

The Regime does not claim property rights in raw data or pure information. Instead, it creates a new class of personal property — outside the traditional common law categories of things in possession and things in action — that attaches to a specific curated, structured, governed and classified instantiation of data, held within a foundation structure and recorded on a constitutive register. Under s.80(4) of the Bill, the personal property right vests in the Data Asset Foundation upon full registration. The Register is therefore constitutive: it does not merely record a pre-existing right, but creates one.

Registration follows a two-stage process. At Stage 1, provisional registration under s.79 records the asset on the Register and assigns a unique Data Asset Identifier, but confers no property right. At Stage 2, full registration under s.80 follows accreditation by an independent accredited assurance provider, at which point the property right vests and the asset becomes fully legally cognisable. This graduated model balances accessibility with assurance: organisations can enter the regime at low cost, while the property right is reserved for assets that have been independently verified against declared attributes.

What the Register Does — and Does Not Do

The Register records structured metadata, governance information, rights declarations, classification profiles and assurance status in respect of registered data assets. It does NOT hold underlying datasets. It does NOT validate data quality or accuracy. It does NOT automatically create proprietary rights: those arise only upon full registration following accreditation (Bill s.80). It does NOT override applicable data protection, intellectual property, or sector regulation.

What This Paper Covers

The Paper addresses the full scope of the proposed regime across twenty-three sections. It begins with the economic and policy context for a data asset registry (Section 2), the six policy objectives the Register is designed to achieve (Section 3), and the scope and definition of a "data asset" under the regime. It then sets out the Register's legal architecture: a five-group field structure, the Data Asset Identifier (DAI), and the Register Data Dictionary (Section 5); a two-axis classification system with four overlay attributes (Section 6); the two-stage registration process with a proposed micro-asset tier for SME accessibility (Section 7); the governance, powers and principles of the Registrar (Section 8); a three-stage review and appeals framework (Section 9); and a four-category enforcement and sanctions regime (Section 10).

The Paper then addresses lifecycle management, including three-tier version control and growth parameter declarations (Section 12); automated verification through the DAR Verification Agent (DAR-VA) operating within a confidential computing architecture using Trusted Execution Environments (Section 13); access, disclosure and confidentiality (Section 14); security interests and encumbrances (Section 15); the interface with Isle of Man data protection law (Section 16); AI governance, including AI-Use Tier requirements and DAR Provenance Certificates (Section 17); semantic interoperability through a machine-readable metadata profile built on DCAT v3, DPROD, ODRL and SHACL (Section 18); the DAR Asset Passport — a tokenised, W3C Verifiable Credential that travels with the asset through the ecosystem (Section 19); interaction with intellectual property, financial services and insolvency law (Section 20); cross-border recognition (Section 21); the five-tier hybrid technical architecture with KSI Blockchain integrity (Section 22); and the liability, indemnification and insurance framework (Section 23).

Subject to consultation: the proposals herein are not final policy. The Department for Enterprise is seeking views on the proposed Data Asset Register and Data Asset Registrar under the Data Asset Foundations framework.

2. The Data Economy: Context and Challenges

2.1 The Scale of the Data Economy

Data has become one of the most economically significant resources of the twenty-first century. The EU Data Economy is expected to be valued at €630 billion in 2025 and rise to €743 billion in 2030, reflecting a 3.3% CAGR between 2025 and 2030.[1] The OECD has identified data as a critical input across virtually every sector of economic activity, from financial services and healthcare to agriculture and public administration.[2] Widely cited IDC projections put the global datasphere at approximately 175 zettabytes by 2025, driven by the proliferation of connected devices, digital services, and artificial intelligence systems.

Yet despite this scale, the data economy operates with remarkably underdeveloped legal and institutional infrastructure. Where other asset classes - land, securities, intellectual property, ships, aircraft - benefit from centuries of registry design, title systems and transactional frameworks, data remains largely unregistered, untitled, and uncollateralised. This institutional gap represents both a market failure and an opportunity.

2.2 The Governance Gap

The central challenge of the contemporary data economy is not the absence of data, but the absence of trusted governance frameworks that enable data to be reliably identified, valued, transacted and protected. Organisations hold vast quantities of data but typically cannot answer fundamental questions: What data assets do we hold? Who has rights over them? What can we lawfully do with them? What are they worth? Can they serve as collateral?

This governance gap manifests in several ways. First, there is no universally accepted definition of a 'data asset' that distinguishes governed, curated datasets that can be trusted from raw, uncontrolled data accumulations. Second, existing legal frameworks - data protection law, intellectual property law, contract law - each address aspects of data governance but none provides a comprehensive framework for treating data as an asset class. Third, the absence of formal registration mechanisms means that third parties cannot reliably verify claims about data assets rights, creating information asymmetries that inhibit market formation.

[1] The European Data Market Study 2024-2026 (March, 2025)
[2] OECD, *Measuring the Economic Value of Data and Cross-Border Data Flows* (2020)

The consequences are material. Weak data governance increases duplication, inaccuracy, inconsistency, compliance risk, transaction friction, and missed opportunities for reuse, collaboration, derived products and financing. These costs are widely recognised, even where precise quantification varies by source and methodology.

2.3 International Responses

Jurisdictions worldwide are responding to these challenges from different philosophical starting points. The European Union has adopted an access-and-governance approach through the Data Governance Act and the Data Act, creating frameworks for data sharing, access and reuse without establishing a general property right in data. China has taken the most expansive approach, establishing dozens of data exchanges and adopting a 'bundle of rights' theory separating data resource holding rights, processing rights and product management rights. The UK's Data (Use and Access) Act 2025 creates Smart Data Schemes and Digital Verification Services but stops short of recognising data as property. Singapore has invested heavily in digital government, Smart Nation infrastructure and AI-governance frameworks, while Japan's Society 5.0 initiative prioritises data interoperability[3].

None of these approaches appears to have established a statutory constitutive register that creates a property right in a registered data asset in the manner proposed here. The Isle of Man's proposed Regime is intended to address that gap.

2.4 Why the Isle of Man

The Isle of Man (IoM) is uniquely positioned to pioneer this framework. As a Crown Dependency with its own parliament (Tynwald), legal system and regulatory infrastructure, the Island can legislate with agility while maintaining common law compatibility with the United Kingdom. The Island has a proven track record of regulatory innovation in financial services, igaming and digital economy governance. Its Foundations Act 2011 provides a tested legal vehicle for asset holding, and its compact scale enables the close collaboration between government, regulators and industry that novel frameworks require.

The proximity to the UK legal system is particularly significant following the UK Property (Digital Assets etc) Act 2025[4], which confirms that a thing is not prevented from being the object of personal property rights merely because it is neither a thing in possession nor a thing in action. The UK Act is deliberately principles-based and leaves the detailed boundaries of such rights to judicial development. The proposed IoM Regime would take a different route by creating a statutory registration mechanism for a defined category of data asset. The relationship is therefore conceptual rather than structural. The UK development may assist courts and counterparties in understanding the legal context in

[3] https://www8.cao.go.jp/cstp/english/society5_0
[4] https://www.legislation.gov.uk/ukpga/2025/29

which the IoM model operates, but recognition of any particular Manx-registered right in another jurisdiction would remain subject to that jurisdiction's private international law and domestic legal principles.

2.5 The Challenges This Regime Must Address

In designing the Register and Registrar, this paper confronts five structural challenges that any data asset regime must resolve:

Non-rivalrousness: unlike land or chattels, data can be copied and used simultaneously by unlimited parties. The regime must create meaningful property rights without claiming a monopoly over information content.

Dynamic character: data assets change continuously. A register designed for static assets will fail. The regime must accommodate continuous evolution while maintaining identity and integrity.

Data protection primacy: any regime that treats data as property must ensure absolute compatibility with data protection law. Registration must never override data subject rights.

Valuation uncertainty: data assets resist conventional valuation methodologies. The regime must create conditions that enable market-based valuation without mandating specific approaches.

AI governance: the rapid development of artificial intelligence creates urgent questions about training data rights, model provenance and derivative output ownership that the regime must address.

The sections that follow set out how each of these challenges is addressed.

3. Purpose and Policy Objectives

3.1 Policy Objectives

The Register is designed to achieve six policy objectives, each flowing directly from the regime established by the Foundations (Amendment) Bill 2025:

Objective	Description
Legal Certainty	Provide a formal mechanism to record claims of stewardship, rights assertions, licences, and security interests in defined data assets, and to confer a statutory personal property right upon full registration (Bill s.80(4)).
Legal Concept Recognition	Provide a mechanism to create a carefully delimited property right in a registered data asset: a curated, structured, governed and classified instantiation of data held within a foundation structure and recorded on a constitutive register, with legally enforceable exclusivity attaching to the registered instantiation rather than to the underlying information content.
Trust and Transparency	Enable third parties to verify the governance posture, assurance status, distribution classification, and declared use restrictions of a registered asset.
Market Enablement	Support financing, licensing, collateralisation, collaboration, and cross-border trade involving data assets by providing a reliable, public, and certified record.
Responsible Innovation	Encourage structured governance, risk classification, and proportionate assurance, calibrated to the nature, sensitivity, and distribution scope of the asset.
Jurisdictional Differentiation	Establish the Isle of Man as a jurisdiction of trust and leadership in data economy governance, attracting data intensive enterprise and data backed financial instruments.

3.2 Legislative Hierarchy

A key source of clarity in any new statutory regime is how obligations are allocated across the legislative and sub-legislative hierarchy. This Paper adopts that hierarchy throughout. Where a matter is stated to be established by the Bill, that statement is descriptive of the Bill. Where a matter is stated to be prescribed, set or specified by regulations, that statement describes a proposal for secondary legislation under the Bill. Where a matter is

described as guidance, standards, protocols, architecture or operational practice, that statement is a recommendation in this Paper for consultation and should not be read as enacted law unless and until adopted through an appropriate legal or administrative instrument. The following matrix sets out the intended allocation for the Register:

Instrument	Subject Matter
Primary Legislation (Foundations (Amendment) Bill 2025)	Establishes the Register and the Registrar (s.77). Defines the two-stage registration process: provisional (s.79) and full (s.80). Confers the personal property right (s.80(4)). Defines disposal on removal (s.81). Provides for equivalent data registers (s.82). Sets the framework for accreditation of assurance providers (s.84) and data enforcers (s.86). Provides High Court jurisdiction (s.90).
Regulations (made by the Department)	Including: Prescribe the form of the DDI (s.76). Prescribe particulars required for registration (s.78). Prescribe the form of the Register (s.77(5)(b)). Set the period for completing provisional registration (s.80(1)(a)). Prescribe the remediation period (s.85(6)). Set criteria for accrediting Assurance Providers (s.84(2)). Establish the data governance framework (s.69). Specify equivalent data registers (s.82(1)). Prescribe version control thresholds and growth parameters.
Guidance (issued by Registrar / Department)	Including: Data Asset Description Schema (DADS). Classification guidance including the two-axis matrix and overlay attributes. Material change trigger list. Version control tier guidance with worked examples. AI-Use Tier governance requirements. Data protection pre-registration verification protocol. Service standard targets. Publication matrix. Security interests operational guidance.
Technical Standards (maintained by Registrar)	Register Data Dictionary (versioned). Cryptographic integrity specifications (KSI Blockchain model).Machine-readable. API standards. Evidence package format. Provenance certificate format. Audit log format and retention schedule.

How to read this Paper
Unless the context otherwise requires, this Paper distinguishes between:
(a) Bill provisions, which describe matters provided for in the Foundations (Amendment) Bill 2025;
(b) proposed regulations, which describe matters that may be prescribed by regulations made under the Bill; and
(c) recommended operational design, which describes policy, guidance, technical standards or administrative practices proposed in this Paper for consultation.

References to what the Registrar “must” or “may” do should be read in that context. Nothing in this Paper should be taken to imply that a proposed operational or technical measure is already enacted law unless expressly stated.

4. Scope and Definition

4.1 What is a Data Asset Under the Regime?

The Bill defines a data asset for the purposes of the Regime as data that has been (a) dedicated to a data asset foundation, (b) the subject of a data asset dedication instrument, (c) registered as a data asset on the data asset register pursuant to section 80, and (d) described as a fully registered data asset on the register (s.67). The registration process — in particular the dedication instrument and the accreditation — is what transforms raw data into a data asset with legal status.

A 'provisional data asset' exists from acceptance of the application under s.79, but the full personal property right arises only upon full registration under s.80(2). This two-stage structure is central to the Regime and to the design of the Register.

Bill s.80(4) — The Property-Creating Provision

Upon the registration of a data asset in the data asset register, a personal property right in the data asset described therein is vested in the data asset foundation, notwithstanding that such data asset may be neither (a) a thing in possession; nor (b) a thing in action. This is the provision that makes the Isle of Man Regime unique: it creates a new class of property outside the traditional common law categories.

4.2 The Non-Rivalrous Nature of Data and the Property Right

Data is fundamentally non-rivalrous: it can be simultaneously used by unlimited parties without diminishment. This characteristic presents a foundational challenge for any property regime. The academic literature — most prominently Professor Lothar Determann's 'No One Owns Data' (2019, UC Hastings) — argues that data is unsuitable for property treatment precisely because it is non-rivalrous, non-excludable without technical measures, and contextually dependent.

This Regime does not claim property rights in raw data or pure information. It follows the position of the Law Commission of England and Wales in its 2023 Digital Assets Final Report: 'pure information' is explicitly excluded from the third category of personal property. Instead, the property right created by registration under s.80(4) attaches to the registered data asset — the specific curated, structured, governed, and classified instantiation of data held within a foundation structure and recorded on the constitutive Register.

Foundational Principle: What Registration Creates

The property right vests in the registered data asset as a governed, classified, curated instantiation — not in the underlying information content. Registration does not grant a monopoly over facts, ideas or raw data. Third parties remain free to independently compile equivalent datasets from their own sources. The property right gives the registrant enforceable legal claims to control access, license use, grant security interests and dispose of the registered instantiation through the Register.

4.2.1 Alignment with the UK Property (Digital Assets etc) Act 2025

The UK Property (Digital Assets etc) Act 2025 provides that a thing is not prevented from being capable of attracting personal property rights merely because it is neither a thing in possession nor a thing in action. The Act does not comprehensively define the outer boundaries of such rights; rather, it confirms doctrinal openness and leaves further development to the courts.

The proposed IoM Regime is not identical to the UK approach. Instead of relying primarily on case-by-case doctrinal development, it would create a statutory registration mechanism through which a defined personal property right vests upon full registration. The comparison is therefore best understood as conceptual support, not direct equivalence.

A registered data asset is intended to be structured so as to support arguments that it is sufficiently defined, identifiable, stable and legally controllable to justify recognition as a distinct registrable asset. It is definable through its Stable Attribute Profile, which provides precise boundaries including field architecture, classification coordinates, AI-Use Tier, schema version, and cryptographic hash. It is identifiable by third parties through the constitutive Register with its unique DAI reference, foundation identity, and publicly queryable status. It is capable of assumption through transfer via re-registration, licensing through recorded interests, and security interests through registrable charges. It has permanence through the cryptographic integrity chain and the Stable Attribute Profile that ensures essential characteristics persist. And it is rivalrous in the legally operative sense, because the registration itself — and the bundle of enforceable rights it creates — is exclusive to the registered foundation.

4.2.2 Constructing Rivalrousness

The rivalrousness of a registered data asset rests on three supporting pillars. First, intellectual property precedent demonstrates that the law routinely creates rivalrous rights in non-rivalrous information: copyright creates a rivalrous right (the exclusive right to reproduce) in a non-rivalrous subject (the creative expression); patents create a rivalrous right (the exclusive right to exploit) in a non-rivalrous subject (the technical information).

The Regime follows this established pattern. A necessary distinction should be noted: copyright and patents protect expression and inventions that meet statutory originality or novelty thresholds, whereas a curated dataset may not meet either threshold. The Regime addresses this by locating the property right not in the informational content (which may lack originality) but in the registered instantiation as a governed, classified, and accredited asset held within a foundation structure. The registration itself, combined with the governance framework and accreditation, is what constitutes the legally operative 'work' in which the property right vests. This means that if two parties independently compile equivalent datasets from the same underlying sources, both may register and both may hold valid property rights in their respective registered instantiations, just as two authors may independently hold copyright in similar works created without copying. The property rights are not in conflict because each attach to a distinct registered asset with its own DAI, governance structure, and integrity chain.

Second, international practice demonstrates that property-like rights in data can be constructed as a bundle of specific, enforceable claims rather than absolute dominion. China's data property framework, introduced through the 2022 Opinions on Building a Foundational System for Data, establishes three distinct rights for data handlers: the right to hold data resources, the right to process and use data, and the right to manage data products. It should be noted that the Chinese system deliberately avoids the concept of data ownership and focuses instead on use and transfer rights; it is best characterised as a 'system without ownership' that promotes data commercialisation through registered rights rather than through traditional property. The IoM Regime takes a different and more explicit approach by creating a statutory property right through registration, but the Chinese experience validates the core proposition that enforceable, registrable rights in data can support market formation, collateralisation, and commercial exchange.

Third, the Law Commission's identification of 'control' as the operative concept for digital assets provides the doctrinal foundation: a person in control must be able to exclude others from the benefit of the thing, put the thing to the uses of which it is capable, and identify themselves as the person with control. The foundation structure satisfies all three limbs.

4.2.3 Non-Exclusive Property and the Emerging Doctrinal Landscape

The preceding analysis suggests that rivalrousness can be accommodated within the existing doctrinal framework. The Regime may also be supported by a separate line of modern property theory which does not treat absolute exclusivity as an invariable condition of property. On that view, the Regime can be presented not as a departure from property doctrine, but as an instance of a broader and well-established understanding of property as a structured allocation of legal powers and controls.

The classical Blackstonian description of property as "sole and despotic dominion" is no longer taken in modern legal theory to be a complete account of property relations. A highly influential modern approach, associated with Hohfeldian analysis and with Honoré's

account of the standard incidents of ownership, treats property as a complex and variable set of legal incidents rather than a single absolute right. On that approach, rights of possession, use, management, income, capital, security, transmissibility and residuarity may be divided, limited, and shared without necessarily depriving an arrangement of its proprietary character. That description is consistent with familiar legal structures such as trusts, leases, easements, co-ownership, and licensed exploitation of intangible assets.

Applied to data, that analysis has important implications. Lametti's formulation of property as a social institution allocating to individuals "a measure of control" and "some degree of exclusivity," rather than absolute dominion, is especially relevant. A registered right under section 80(4) may plausibly be analysed in those terms: it allocates a defined measure of control through registration, governance, classification and enforcement, together with a degree of exclusivity in relation to the registered instantiation and the legal powers attached to it. On this analysis, the Regime need not claim exclusive control over information as such. Rather, the right attaches to the registered data object or instantiation recognised by the Regime and to the legal incidents that the Regime confers in respect of it.

Recent scholarship on data governance supports the proposition that data-related entitlements may be allocated as a differentiated bundle of rights rather than through a single exclusive ownership model. Eckardt and Kerber, applying economic property-rights theory to IoT data and the EU Data Act, analyse how legal and technical arrangements can distribute rights to control, access, use, share and monetise the same data among different actors. Their work shows that data governance can be structured through overlapping and functionally differentiated entitlements, even where no single actor enjoys simple, absolute dominion.

Other recent scholarship also cautions against assuming that data governance must be organised around a single exclusionary ownership model. Purtova and van Maanen examine both "data as an economic good" and "data as a commons," and show that commons-oriented thinking can illuminate shared and institutional forms of control over data, even though they also warn that purely data-centric governance models have significant limits. That literature does not establish a consensus in favour of non-exclusive property rights in data, but it does support the narrower proposition that data governance need not be conceptualised solely through the lens of absolute exclusivity.

Comparative regulatory developments point in a similar direction. The EU Data Act does not create a general ownership right in data. Instead, it establishes harmonised rules on fair access to and use of data, clarifies who may use certain categories of data and on what terms, and redistributes practical control by granting access and sharing rights to users and third parties in defined circumstances. The Data Governance Act likewise facilitates data sharing, intermediation and re-use through institutional mechanisms rather than by conferring comprehensive exclusive title in data. These measures are therefore better understood as allocating differentiated rights and obligations in relation to data than as adopting a classical exclusive-ownership model.

Chinese policy developments also indicate movement toward registrable and commercially exploitable interests in data, although the legal position remains evolving

and should not be overstated. Official statements confirm that China is accelerating the development of a data property-rights registration system. Academic and policy commentary further indicates that local pilots and exchange-based mechanisms have been used to support the registration, licensing and financing of processed data products, including data-backed lending transactions associated with data exchanges. The safer conclusion is therefore not that China has settled a general law of "data ownership," but that it is developing institutional mechanisms for recognising and commercialising data-related interests without relying on a simple model of exclusive ownership of information as such.

The Isle of Man Regime can therefore be situated at the intersection of two propositions. First, for compatibility with orthodox common-law reasoning, it may continue to rely on the analysis set out in section 4.2.2, including the continuing relevance of rivalrousness in the Law Commission's treatment of core digital assets. Secondly, it may also be explained by reference to the broader proposition that proprietary treatment can rest on a structured allocation of control, use, transfer and enforcement rights, and that exclusivity may be present by degree rather than as an all-or-nothing condition. That second proposition does not displace orthodox doctrine, but it provides an additional analytic basis on which the Regime may be defended.

On that basis, the Regime's classification system may be described as calibrating the degree of control and exclusivity associated with a registered data right, rather than as assuming that all proprietary rights must take the form of absolute exclusion from underlying information. To the extent that the Regime allows the same underlying information to be reflected in multiple registered instantiations, each with its own governance framework and legal incidents, that feature is capable of analysis as a deliberate allocation of differentiated rights in relation to data objects constituted by the Regime. The legal strength of that position, however, would depend on the legislation making clear that the proprietary claim attaches to the registered instantiation and its associated legal incidents, not to information in the abstract.

4.2.4 What Registration Does Not Grant

Clarity requires stating what the property right does not encompass. Registration grants no monopoly over underlying facts or information: third parties remain free to independently compile equivalent datasets. It does not restrict pre-existing third-party rights: intellectual property, contractual, and data protection rights held by others are unaffected. It does not expand the lawful basis for processing: registration never constitutes a lawful basis under Article 6 of the Applied GDPR. And it does not override data subject rights: GDPR rights continue in full and take priority over the property right.

4.3 What the Register Does Not Do

- Store underlying datasets, raw data files, or database contents.
- Validate the quality or accuracy of the data described in a registration.

- Automatically create proprietary rights: those arise only upon full registration following accreditation.
- Override applicable data protection legislation, intellectual property law, or sector-specific regulation.
- Adjudicate disputes over rights: it records declared rights and provides notice. Disputes are resolved through the s.90 mechanism or the courts.

The Register records structured metadata, governance information, certified assurance status, and rights declarations. A third party relying on the Register relies on the fact of registration and recorded status — not on a warranty from the Registrar of the underlying truth of declarations.

5. The Data Asset Register: Legal Architecture

Status: Bill framework for the Register and its legal effects; DAI design and field architecture proposed for implementation through regulations.

5.1 Legal Status

The Data Asset Register is the statutory public register established under the Foundations (Amendment) Bill 2025 (s.77(5)) for the purpose of recording, certifying, maintaining, and publishing information about data assets registered under the Data Asset Foundation Regime. The Register has the following legal characteristics:

- It is a statutory register: its existence, functions, and legal effects are defined by primary legislation.
- It is a public register: certain information is accessible to the public in accordance with the Access Framework (Section 14).
- It is constitutive: full registration creates the legal existence of a data asset as a distinct personal property right under the Regime (s.80(4)).
- It is authoritative: entries in the Register are presumed, in the absence of fraud or manifest error, to accurately reflect the legal position as at the time of entry.
- It is dynamic: it records not just the current state of a data asset's legal attributes but the history of changes to those attributes over time.

5.2 The Five-Group Field Architecture

The Register organises information into five functional groups, each serving a distinct legal and operational purpose. This architecture designed to support the full lifecycle of a registered data asset:

Field Group	Function
Group A: Foundation and Governance Fields	Records the identity, governance structure, and stewardship roles of the Data Asset Foundation holding the asset. Includes: foundation identifier, legal name, registration number, jurisdiction of formation, business and technical data stewards, data enforcer, charter hash and beneficiary details (where applicable).

Group B: Asset Core Fields	Records the defined characteristics of the data asset itself. Includes: asset identifier (DAI), asset name, asset type, classification profile (class code and overlay attributes), domain, owner role, lifecycle status, purpose, charges, restrictions, DDI hash, dedicator identity, and Stable Attribute Profile (SAP) including growth parameter declarations.
Group C: Rights and Encumbrances Fields	Records the allocation of rights in the data asset among multiple parties, including ownership, access licences, security interests, encumbrances, and AI-Use licensing terms. This group provides the notice function critical to market enablement.
Group D: Provenance and Composition Fields	Records the derivation history and lineage of the data asset, linking it to source assets, transformation events, and computational processes. Includes: source asset identifiers, source data hashes, derivation type, processing logic hash, processing attestation, and AI provenance metadata where applicable.
Group E: Assurance and Compliance Fields	Records the data asset's accreditation status, data quality grade, compliance posture, data protection registration, privacy classification, version history, and integrity chain status. This group is the primary indicator of governance health.

5.3 The Data Asset Identifier (DAI)

At the time of provisional registration, this Paper proposes that the Registrar assign each data asset a Data Asset Identifier (DAI). The DAI is a unique, immutable, jurisdiction-specific alphanumeric code that serves as the persistent legal reference for the asset throughout its lifecycle. The DAI has the following properties:

- It is generated algorithmically at registration, incorporating a jurisdiction prefix (IOM), a timestamp element, and a cryptographic component derived from the initial hash of the asset's schema descriptor.
- It is immutable: once assigned, the DAI does not change even if the asset's contents are amended, the holding entity changes, or rights are transferred.
- It is unique and collision-resistant: no two data assets share a DAI, and a DAI is never reused even after deregistration.
- It is persistent: a DAI remains the legal reference even when an asset enters remediation or is archived.

- It supports version suffixes: where a Tier 2 amendment is recorded, the DAI is extended with a version suffix (e.g. DAR-00147-v2) while retaining the original registration date and reference.
- It is machine-readable and structured to enable interoperability with other national and international data asset registries.

The DAI serves as the primary reference in all legal instruments, licences, security documents, and regulatory filings that relate to registered data assets. Under the proposed design, parties dealing in registered data assets would be expected to reference the DAI in transaction documents.

5.4 The Register Data Dictionary

The Register Data Dictionary is a versioned, change-controlled artefact maintained by the Registrar that defines all field specifications, controlled vocabularies, validation rules, and technical standards for the Register. It is the single authoritative source for the meaning, format, and permissible values of every data element in the Register.

Changes to the Dictionary follow a formal change control process, with material changes subject to public notice and a minimum 30-day comment period. Each published version of the Dictionary is assigned a version number and effective date, and the Register records which Dictionary version was operative at the time of each registration event.

6. Data Asset Classification

Status: proposed classification framework for implementation through regulations and the data governance framework; not itself enacted by the Bill.

6.1 Design Principles

Every registered data asset is assigned a classification profile that determines its assurance requirements, residency expectations, disclosure defaults, and renewal cadence. The classification system is built on a core design principle: each classification axis should be coherent, with every point on the axis distinguished from its neighbours on exactly the same logical dimension.

The classification separates into two independent axes — Distribution Scope and Data Sensitivity — with four independent overlay attributes. The class code describes what the asset is and how it is shared. The overlay attributes describe what is done with it and under what constraints.

6.2 Axis 1: Distribution Scope

Distribution Scope describes the boundary within which the data asset is shared or made available. It has four tiers:

Code	Tier	Definition
I	Internal	The data asset is held and used solely within the DAF's ecosystem. No access is granted to any external party. Includes internal R&D, analytics, risk management and custody-ready datasets.
C	Controlled-Share	The data asset is shared with a defined, contracted set of external partners. The Foundation can identify every party with access.
F	Federated	The data asset is pooled within a formal multi-party consortium, data trust, or federated research arrangement. Multiple parties contribute to and access a shared data pool under collective governance.
D	Distributed	The data asset is made available openly, via marketplace, public API, or to any qualified recipient without prior bilateral agreement. The Foundation does not control the recipient set at distribution.

6.3 Axis 2: Data Sensitivity (Risk Profile)

Data Sensitivity describes the nature of the data in the asset with respect to the individuals it concerns and the data protection obligations it attracts. It has three tiers:

Code	Tier	Definition
1	Non-personal / Anonymous / Synthetic	No personal data: either never personal, fully and irreversibly anonymised, or synthetically generated. A Privacy Non-Applicability Statement is required.
2	Pseudonymised / Aggregated	Personal data that has been pseudonymised or aggregated such that individuals are not directly identifiable but re-identification risk remains. Data protection obligations apply.
3	Identifiable / Special Category	Personal data from which individuals are directly identifiable, or special category data (health, biometric, genetic, racial/ethnic, political, religious, trade union, sexual orientation, criminal).

6.4 The Classification Matrix

Crossing the four Distribution Scope tiers with the three Data Sensitivity tiers produces twelve base class codes. Under the proposed classification framework, every registered data asset would be assigned exactly one class code:

Distribution \ Sensitivity	1: Non-personal	2: Pseudonymised	3: Identifiable
I: Internal	I1	I2	I3
C: Controlled-Share	C1	C2	C3
F: Federated	F1	F2	F3
D: Distributed	D1	D2	D3

6.5 The Four Overlay Attributes

This Paper recommends that every registered data asset be assigned four overlay attributes in addition to its class code. These attributes are independent of one another and

of the class code. The combination of class code and four overlay attributes constitutes the complete classification profile.

6.5.1 Commercial Intent (CI)

Value	Label	Description
CI-PG	Public Good	Non-commercial use for public benefit, open research, transparency, or regulatory compliance.
CI-PC	Pre-Commercial	Internal R&D, evaluation, or preparation for future commercial exploitation. Not yet actively commercialised.
CI-CO	Commercial	Active licensing, data products, B2B supply, or commercial API access. Revenue is or may be generated.
CI-MK	Marketplace	Active trading via a formal data marketplace, exchange, or broker.

6.5.2 AI Use Tier (AI)

Value	Label	Description
AI-0	No AI Use	Not used in training, fine-tuning, validation, or inference of any AI system. Technical measures preventing AI ingestion required. Full reservation against any AI training.
AI-1	Internal Analytics	Used only in internal AI analytics. AI processing log maintained; model outputs tagged with DAR provenance. Full reservation against external AI use.
AI-2	Licensed AI Training	Licensed for use in external AI training. Licensing agreement should specify model type, training methodology, output restrictions, and provenance attribution requirements.
AI-3	Open AI / GPAI	Available for general-purpose AI training. Published terms of use required. Machine-readable provenance metadata compatible with EU AI Act training data summary template.

6.5.3 Regulated Sector (RS)

Value	Label	Description
RS-0	General	No sector-specific regulatory overlay.

RS-F	Financial Services	Subject to FSA obligations, AML/CFT, financial instrument rules.
RS-H	Health	Subject to health data legislation, clinical governance requirements.
RS-L	Legal / Privilege	Subject to legal professional privilege or court confidentiality.
RS-C	Critical Infrastructure	Data relating to critical national infrastructure operations.
RS-X	Other Regulated	Other sector-specific regulation applies; must be specified.

6.5.4 Residency Tier (RT)

Value	Label	Description
RT-0	No Restriction	No residency constraint. Data may be stored and processed in any jurisdiction.
RT-1	Documented Compliance	No mandatory residency but documented compliance with transfer mechanisms is required.
RT-2	Approved Jurisdictions	Storage and processing restricted to jurisdictions on the Registrar's Recognised Jurisdiction list.
RT-3	Sovereign Gateway (at rest)	Under the proposed framework, data at rest should be within the IoM Sovereign Gateway. Processing may occur in approved jurisdictions.
RT-4	Full Sovereign	Under the proposed framework, data at rest and in use should be within the IoM Sovereign Gateway, with confidential computing requirements.

7. The Registration Process

Status: Bill framework, supplemented by proposed regulatory and operational design.

7.1 Two-Stage Registration

The Bill establishes a two-stage registration process. This section first describes that statutory framework. It then identifies proposed regulatory and operational features recommended in this Paper to support accessibility, proportionality and implementation.

Stage 1: Provisional Registration (s.79)

The council of a Data Asset Foundation applies to the Registrar to register data as a data asset by providing the prescribed particulars and an executed Data Asset Dedication Instrument (DDI). The Registrar accepts the application if satisfied that it complies with s.78. Upon acceptance, the data is recorded as a 'provisional data asset' on the Register. A DAI is assigned. The provisional data asset is publicly visible on the Register but no personal property right yet exists.

Stage 2: Full Registration (s.80)

The DAF must complete registration within the period prescribed by regulations by submitting confirmation of a data asset accreditation issued by an Accredited Assurance Provider (AAP) and meeting any outstanding requirements. When the Registrar is satisfied, the Register is updated to record the asset as 'fully registered' and the personal property right vests in the DAF under s.80(4). If requirements are not met within the prescribed period, the Registrar removes the provisional data asset and the registration is void.

The Bill does not itself create a self-attested micro-asset pathway. This Paper recommends that regulations may provide for a limited exception for specified low-risk assets, subject to conditions, liability allocation and AI-use restrictions.

Proposed regulatory and operational extension: Micro-Asset Tier

To support accessibility for startups and SMEs, this Paper recommends that regulations may create a Micro-Asset Tier. Under that proposal, assets classified as Internal/Non-personal (I1) with a declared valuation below a prescribed threshold could qualify for self-attested full registration for a limited period, subject to enhanced liability and restricted downstream use. This would bypasse mandatory third-party AAP accreditation for a period of 12 months, provided the foundation assumes legal liability for the accuracy of the Stable Attribute Profile (SAP). To mitigate the risk of the Micro-Asset Tier being used to obtain the property right and DAR Provenance Certificates for assets intended for AI training marketplaces without independent assurance, self-attested micro-assets should be

restricted to AI-Use Tier AI-0 (No AI Use) or AI-1 (Internal Analytics). Assets requiring AI-2 (Licensed AI Training) or AI-3 (Open AI / GPAI) classification would have to undergo full AAP accreditation regardless of valuation, reflecting the heightened governance requirements and third-party reliance implications of AI training data. The Registrar may not issue DAR Provenance Certificates for self-attested micro-assets.

7.2 The Application

A registration application under s.78 must include the following elements:

Element	Description
Data Asset Dedication Instrument (DDI)	An executed DDI effecting the dedication of the data to the Foundation and specifying the rights and permissions transferred.
Foundation Details (Group A)	Foundation identifier, legal name, registration number, stewardship roles, data enforcer, and charter hash.
Asset Description (Group B)	Completed Data Asset Description Schema (DADS) covering: asset name, type, proposed classification profile (class code and overlay attributes), domain, purpose, restrictions and Stable Attribute Profile with growth parameter declarations.
Provenance Declaration (Group D)	Description of the data's origin, references to any source data assets (by DAI), and description of any transformation processes applied. Where the data includes AI-generated content, AI provenance metadata must be provided.
Rights Declaration (Group C)	Current ownership details, any existing licences, any existing security interests or encumbrances.
Data Protection Declaration (Group E)	Whether the asset contains personal data, DPA registration number, Article 6 lawful basis declaration, and for Sensitivity Level 3 assets, DPIA evidence and IoM Information Commissioner consultation evidence.
Accreditation (for full registration)	Confirmation of data asset accreditation issued by an AAP (s.84), certifying the asset meets the data governance framework requirements.
Fee	Payment of the applicable registration fee as published in the Registrar's fee schedule.

7.3 Registrar Decision Framework

The Registrar makes decisions on registration applications using a structured decision framework. The possible outcomes are:

Status	Meaning and Effect
Provisionally Registered	Application accepted under s.79. DAI assigned. No personal property right. DAF must complete accreditation within prescribed period.
Fully Registered	Accreditation confirmed and all requirements met under s.80. Personal property right vested. Asset is active on the Register.
Under Review	Registrar is assessing the accuracy of the record (s.83(3)). Caution flag visible to public and professional users.
Remediation	Accreditation has expired or been revoked (s.85(4)). 90-day period to restore. Personal property right remains but asset cannot be utilised.
Removed	Asset removed from the Register. Personal property right extinguished (s.81(8)). Disposal under s.81(3). Record permanently archived.

7.4 Prohibited Registrations

The Registrar must refuse to register a data asset where: the data was collected unlawfully; the data is subject to a binding deletion order with no valid override; the data is classified under a national security framework prohibiting registration; a court order prohibits transfer or commercial exploitation; registration would facilitate money laundering, terrorist financing, or sanctions evasion; the applicant cannot demonstrate a valid lawful basis for processing any personal data within the asset; or the application describes something that does not qualify as a data asset under the statutory definition.

8. The Data Asset Registrar

Status: mixed — Bill framework for appointment and core functions (s.87–s.91), supplemented by proposed review, appeals and fee structure.

8.1 Appointment and Governance

The Data Asset Registrar is the competent authority established under the Foundations (Amendment) Bill 2025 (s.77) to administer the Register and exercise the regulatory, investigatory, and enforcement functions conferred by the Act. The Department for Enterprise appoints the Registrar (s.77(1)), who holds office under terms set by the Department and may be removed only for inability or unfitness (s.77(4)).

The Registrar's governance model balances operational independence with democratic accountability. The governance structure comprises:

- The Registrar, being a senior public official or body with appropriate legal, technical, and regulatory expertise.
- An Advisory Council of five to seven members including representatives of the legal profession, data technology, financial services, civil society, and academic expertise.
- An independent dispute resolution mechanism (s.90) for resolving registration disputes, with ultimate recourse to the High Court.

8.2 Functions of the Registrar

The Registrar exercises four categories of statutory function:

Registration Functions

Receiving, processing, and determining applications for registration. Assigning DAIs and updating the Register. Recording amendments to Asset Records upon application or Registrar-initiated action. Processing cessation of registration and archiving records. Issuing DAR Provenance Certificates for AI governance compliance.

Oversight Functions

Overseeing the process of verification of the accuracy and completeness of information submitted, including data protection pre-registration verification for assets containing personal data. Maintaining the Register in accordance with statutory and regulatory requirements. Updating the Register to reflect accreditation suspension, revocation, or expiry (s.85(3)). Marking assets as 'under review' during accuracy assessments (s.83(3)). Issuing Certified Extracts from the Register. Monitoring version control compliance and

growth parameter adherence. Receiving and acting on Data Enforcer reports, updating asset status upon Data Enforcer stop directions, and recording the Data Enforcer's annual opinion on the Register.

Regulatory and Enforcement Functions

Removing data assets following remediation failure (s.85(8)). Removing data assets from DAFs that fail to meet establishment requirements (s.68(5)). Cooperating with the Information Commissioner, the Financial Services Authority, and other relevant regulators. Referring matters of suspected criminal conduct to the appropriate authorities.

Policy and Development Functions

Publishing guidance on the operation of the Register and registration requirements. Maintaining the Register Data Dictionary and controlled vocabularies. Maintaining the DAR Metadata Profile (DCAT v3 / DPROD / ODRL) and the controlled vocabulary SKOS concept schemes. Issuing and revoking DAR Asset Passports. Advising the Department on regulations and subordinate legislation. Engaging with international counterparts on compatible frameworks and mutual recognition. Commissioning and publishing research on data economy development. Reporting annually on Register operations and data economy development.

8.3 Principles Governing the Registrar

This Paper recommends that the Registrar exercise all functions in accordance with six principles: (1) legality (all actions authorised by the Act); (2) proportionality (intervention proportionate to risk); (3) transparency (basis for significant decisions published); (4) consistency (like cases treated alike); (5) efficiency (processes completed within published service standards); and (6) independence (no instructions from Government or registrants on specific decisions).

8.4 Fees and Financial Sustainability

The Registrar may charge prescribed fees. The fee structure is designed to be cost-reflective, to incentivise early registration, to scale with complexity, and to remain competitive. The fee structure is set by regulations on the recommendation of the Registrar following consultation. Fee categories include: initial registration, annual maintenance, certification and extracts, transfer, security interest notation, version amendments, and provenance certificate issuance.

9. Review and Appeals Framework

This Part sets out the framework for reviewing and appealing from the Registrar's decisions

9.1 Scope of Reviewable Decisions

The Review and Appeals Framework applies to all significant decisions made by the Registrar in the exercise of its statutory functions. Reviewable decisions include:

1. **Registration decisions:** refusal of an application for provisional registration (s.79); refusal to record full registration (s.80); and conditions imposed on registration.
2. **Classification decisions:** the Registrar's determination of classification coordinates where the applicant disputes the assigned profile, and Registrar-initiated reclassification of an existing asset.
3. **Enforcement decisions:** the imposition of administrative penalties under the enforcement framework; the issuance of formal compliance notices; suspension of an asset's status; compulsory deregistration; and publication of enforcement notices.
4. **Amendment decisions:** Registrar initiated amendments to Asset Records under s.83(6); determination of whether a change constitutes a Tier 2 or Tier 3 event; and refusal to record a notified amendment.
5. **Accreditation-related decisions:** marking an asset as 'remediation' under s.85(4); removal of an asset following remediation failure under s.85(8); and refusal to accept a replacement AAP accreditation.
6. **Access and disclosure decisions:** refusal of an application for access to Confidential or Professional tier information; and the Registrar's determination to reclassify information between access tiers.
7. **Fee decisions:** where the Registrar's application of the fee schedule to a particular registration or amendment is disputed.

Decisions that are not reviewable under this framework include: the Registrar's referral of suspected criminal conduct to prosecuting authorities (which is a discretionary public interest function); the Registrar's policy guidance and published interpretative positions (which may be challenged by judicial review but not through this administrative mechanism); and interim measures taken under emergency powers (which are subject to expedited review rather than the standard process).

9.2 Three-Stage Review Process

The framework provides a three-stage process, designed to resolve disputes at the lowest appropriate level while preserving access to independent adjudication:

Stage 1: Internal Reconsideration

Under the proposed review framework, a party aggrieved by a Registrar decision may request internal reconsideration within 28 days of the decision being notified. The reconsideration is conducted by a senior officer of the Registrar who was not involved in the original decision. The reconsideration officer reviews the original decision on the papers, taking into account any additional representations submitted by the applicant. The reconsideration officer may: affirm the original decision; vary the decision (including by substituting a different sanction or amending conditions); or set aside the decision and remit the matter for fresh determination. The reconsideration must be completed within 42 days. The original decision remains in effect during reconsideration unless the reconsideration officer directs otherwise.

Stage 2: Independent Review Panel (s.90)

If the applicant is dissatisfied with the outcome of internal reconsideration, or if the matter involves a dispute between parties (such as a priority dispute between competing security interest holders), the matter may be referred to an Independent Review Panel under the s.90 mechanism. The Panel is constituted from a standing list of qualified reviewers maintained by the Department, drawn from persons with expertise in data governance, property law, regulatory law, and relevant technical domains. The Panel comprises one or three members depending on the complexity and value of the matter.

The Panel conducts a de novo review of the decision on the merits, not limited to the material before the Registrar. The Panel may receive written submissions, request oral representations, commission independent technical evidence (including from DAR-VA specialists where verification issues are in dispute), and inspect relevant Register entries and Enforcement Register records. The Panel may: affirm, vary, or set aside the Registrar's decision; substitute its own decision; award costs; and make recommendations to the Registrar on systemic issues identified during the review.

Under the proposed framework, the Panel would issue a reasoned written determination within 90 days of the reference. Determinations are published (with appropriate redaction of confidential information) and are binding on the Registrar and the parties, subject to appeal to the High Court. The Panel's published determinations form a body of precedent that guides the Registrar's future decision-making and provides transparency to market participants.

Stage 3: Appeal to the High Court

A party dissatisfied with the Panel's determination may appeal to the High Court of the Isle of Man. The appeal is on a point of law or on the ground that the Panel's decision was unreasonable in the Wednesbury sense (i.e. a decision that no reasonable Panel, properly directing itself, could have reached). The High Court may not substitute its own assessment of the merits for the Panel's except on a question of law. This limitation preserves the specialist expertise of the Panel while ensuring that legal errors and procedural unfairness are correctable.

The High Court may: affirm the Panel's determination; set it aside and remit the matter to the Panel for reconsideration in light of the Court's findings; or, in exceptional cases, substitute its own decision where remittal would serve no useful purpose.

9.3 Interim Relief

At any stage of the review process, the affected party may apply for interim relief to stay the effect of the Registrar's decision pending the outcome of the review. This Paper proposes that the test for interim relief mirror the American Cyanamid principles applied in Manx law: the applicant would need to demonstrate a serious question to be tried, that damages would not be an adequate remedy, and that the balance of convenience favours the grant of interim relief. In cases involving suspension or deregistration, the urgency of the matter and the potential for irreparable harm to the property right are relevant considerations.

For emergency enforcement measures (immediate suspension in response to a Critical DAR-VA alert or a Data Enforcer stop direction), this Paper recommends that the Registrar initiate an expedited internal review within 7 days. The affected party may apply to the Panel for emergency interim relief on 48 hours' notice.

9.4 Relationship to Existing Dispute Mechanisms

The Review and Appeals Framework does not displace existing legal remedies. A party may pursue judicial review of the Registrar's conduct in parallel with or instead of the statutory review process, although the Court will ordinarily expect the statutory mechanism to have been exhausted before entertaining a judicial review application. Disputes between private parties arising from transactions in registered data assets (e.g. contractual disputes between licensor and licensee, or priority disputes between secured creditors) are resolved through the s.90 Panel mechanism where the dispute relates to a matter recorded on the Register, and through the courts where it does not.

10. Enforcement and Sanctions

Status: proposed enforcement architecture for implementation through the Bill, regulations, Registrar procedures and guidance.

10.1 Proposed enforcement framework

The Bill provides the statutory basis for review, remediation and removal. This section sets out the sanctions architecture recommended in this Paper for implementation through regulations, Registrar procedures and published guidance.

Unless expressly stated otherwise, references in this section to penalties, notice categories, publication practices, tariff levels and offence structures should be read as proposals for consultation rather than as enacted law.

10.2 Design Principles

The sanctions framework should be governed by four principles, each flowing from the Registrar's statutory principles (Section 8.3):

1. **Proportionality.** Sanctions are calibrated to the severity of the breach, the culpability of the party, and the harm or risk of harm caused. Minor administrative failures attract different consequences from deliberate misrepresentation.
2. **Graduated escalation.** The framework provides a graduated scale of intervention, from advisory notices through to deregistration and criminal referral. The Registrar is expected to apply the least intrusive sanction capable of achieving compliance, escalating only where lesser measures have failed or where the severity of the breach warrants immediate action.
3. **Transparency.** The basis for significant enforcement decisions is published, consistent with the Registrar's transparency principle. Published enforcement action serves a dual purpose: it provides notice to third parties who may be relying on the Register, and it creates a deterrence effect that supports voluntary compliance.
4. **Due process.** No sanction is imposed without the affected party having had a reasonable opportunity to make representations. Emergency measures (such as immediate suspension in response to a Critical DAR-VA alert) take effect immediately but are subject to expedited review.

10.3 Categories of Breach

The framework should distinguish four categories of breach, each attracting a different range of sanctions:

Category 1: Administrative Non-Compliance

Failures of a procedural or administrative character that do not directly affect the integrity of the Register or the rights of third parties. Examples include: late filing of a Tier 1 administrative update; failure to pay annual maintenance fees within the prescribed

period; incomplete submissions that are capable of correction; and minor departures from the prescribed form of notification.

Proposed consequences: Advisory notice specifying the breach and requiring correction within 28 days. If uncorrected: formal warning recorded on the Register against the DAF's entry (visible at the Professional Access tier). Persistent administrative non-compliance across three or more consecutive reporting periods: administrative penalty not exceeding a prescribed amount (set by regulations, calibrated to the classification profile of the asset). The administrative penalty is a civil debt recoverable by the Registrar.

Category 2: Material Non-Compliance

Breaches that affect the accuracy of the Register, the integrity of a registered asset, or the governance framework, but fall short of deliberate misrepresentation. Examples include: failure to notify a Tier 2 amendment within 28 days of the triggering event; failure to respond to a DAR-VA Red alert within the prescribed period; continued utilisation of a data asset during remediation in breach of s.85(1); failure to comply with a remediation plan directed by the Data Enforcer; and failure to maintain a functioning DAR-VA access gateway as required by the registration conditions.

Proposed consequences: Under this Paper's proposals: formal compliance notice specifying the breach, the required corrective action, and the deadline for compliance (not less than 14 days, not more than 90 days, depending on severity). If uncorrected: the Registrar would be empowered to impose an administrative penalty (at a higher tariff than Category 1, prescribed by regulations), to mark the asset as 'Under Review' (s.83(3)), visible at the Public Access tier, alerting third parties, and to suspend the asset's status, preventing new licences or security interests from being registered. Persistent or repeated material non-compliance: the Registrar would be empowered to initiate remediation proceedings under s.85(4), triggering the 90-day remediation clock. If remediation fails: deregistration under s.85(8).

Category 3: Serious Non-Compliance

Breaches that undermine the integrity of the Register, cause or risk causing material harm to third parties, or involve a wilful disregard of the Regime's requirements. Examples include: deliberate misstatement of the Stable Attribute Profile, classification coordinates or growth parameters; deliberate obstruction of the DAR-VA's access to the data asset; making a false declaration to the Registrar in connection with an application for registration, amendment or transfer; utilising a data asset in a manner that the registrant knows to be inconsistent with the declared classification or AI-Use Tier; and failure to comply with a Data Enforcer stop direction.

Proposed consequences: Under this Paper's proposals, the Registrar would be empowered to impose an administrative penalty at the highest prescribed tariff. The Registrar would be empowered to immediately suspend the asset and mark it as 'Suspended - Enforcement Action' on the Register (visible at the Public Access tier). The Registrar would be empowered to initiate compulsory deregistration without a remediation period where the breach is sufficiently serious that remediation would not restore confidence in the registration. The

Registrar would be required to notify the Data Enforcer. Where the breach involves personal data, the Registrar would be required to notify the Information Commissioner. The Registrar would be empowered to publish a public enforcement notice setting out the nature of the breach and the sanctions imposed.

Category 4: Criminal Conduct

Conduct that may constitute a criminal offence under the Act, data protection legislation or general criminal law. This Paper further recommends that the Act or associated legislation create specific offences addressing conduct that undermines the integrity of the Register or the operation of the Regime:

1. **Knowingly or recklessly making a false declaration** to the Registrar in connection with a registration application, amendment notification or any other submission to the Register.
2. **Knowingly or recklessly providing false information** to the Data Enforcer or an Accredited Assurance Provider in connection with the performance of their statutory or professional functions.
3. **Tampering with or obstructing the DAR Verification Agent,** including interfering with the integrity of the TEE attestation process, corrupting verification data presented to the DAR-VA via the access gateway, or fabricating verification outputs.
4. **Utilising a data asset that has been suspended or is subject to a stop direction,** knowing that the suspension or direction is in effect.
5. **Fraudulent registration:** registering a data asset on the basis of information known to be materially false, for the purpose of obtaining the benefits of registration (including the property right, security interest capability or DAR Provenance Certificates).

Proposed consequences: Under this Paper's proposals, the Registrar would refer the matter to the appropriate prosecuting authority. On conviction: a fine not exceeding a prescribed maximum (to be set by the Act, with provision for the Department to increase by order). The Registrar would also be empowered to impose any of the administrative sanctions available for Category 3 breaches. Where a DAF council member is convicted of a Category 4 offence in connection with their duties, the Registrar would be empowered to direct the DAF to remove the individual from the Council.

10.4 Administrative Penalty Framework

It is proposed that administrative penalties should be civil in character and determined by the Registrar in accordance with a structured assessment prescribed by regulations and supported by published guidance. The factors to be considered include: the severity and duration of the breach; whether the breach was deliberate, reckless, or negligent; the extent of any harm caused or risk of harm to third parties, data subjects, or the integrity of the Register; the party's cooperation with the Registrar's investigation and willingness to remediate; any previous enforcement history; and any profit or benefit obtained as a result of the breach.

This Paper recommends that penalty levels be prescribed by regulations, expressed as fixed amounts, daily rates for continuing breaches, or percentages of the annual maintenance fee payable in respect of the asset. It is further proposed that the Registrar publish penalty guidelines setting out the indicative ranges for each category of breach. All penalties would be payable to the Registrar and applied to the operational costs of the Register.
A party subject to an administrative penalty may appeal through the Review and Appeals Framework (see 8a).

10.5 Enforcement Register

This Paper recommends that the Registrar maintain an Enforcement Register, accessible at the Professional Access tier (Tier 1), recording all formal enforcement actions taken. Each entry records: the DAI of the affected asset, the nature of the breach, the category, the sanction imposed, the date of the action, and the outcome (including any successful appeal). The Enforcement Register provides an auditable record of the Registrar's use of enforcement powers and supports the consistency principle by making precedent visible.

11. Key Roles in the Regime

The Regime creates a structured ecosystem of roles with defined responsibilities:

Role	Description and Statutory Basis
Data Asset Registrar	Appointed by the Department (s.77(1)). Establishes and maintains the Register. Makes registration decisions. Issues Certified Extracts and Provenance Certificates. May charge prescribed fees. Cannot be instructed by Government on specific decisions.
Council of the DAF	Governing body of the Data Asset Foundation. Makes registration applications (s.78(1)). Must comply with data governance framework (s.87). Must notify Registrar of Register mistakes without undue delay (s.83(4)). Offence for failure to notify (s.83(9)). Responsible for maintaining growth parameters within declared thresholds.
Data Enforcer	Mandatory appointment by every DAF (s.86(1)). Functions specified in the data governance framework. Must notify Registrar of Register mistakes (s.83(5)). Offence for failure to notify (s.83(10)). Has standing to bring proceedings (s.86(8)). Liable for own fraud, wilful misconduct, or gross negligence (s.86a). Must ensure data protection compliance is maintained throughout asset lifecycle.
Accredited Assurance Provider (AAP)	Person or body accredited by the Department (s.84(1)). Issues and renews data asset accreditations. Must be satisfied that asset meets data governance framework requirements (s.84(4)). Must verify the methodology supporting any claim that an asset is classified as Sensitivity Level 1 (non-personal) or Sensitivity Level 2 (pseudonymised), including any anonymisation, pseudonymisation and re-identification risk assessment required by the data governance framework. May disclose information to Registrar (s.84(5)). Listed on public register.
Business Data Steward	Named individual with accountability for the business governance of the asset. Identifier recorded publicly in Group A fields. Role defined in data governance framework.
Technical Data Steward	Named individual responsible for technical integrity, schema maintenance, and version control compliance. Identifier

	recorded publicly in Group A fields. Role defined in data governance framework.
Registered Agent	The authorised contact person or service provider designated to manage the foundation's account credentials, official communications and administrative interactions with the Register on behalf of the foundation.

12. Data Asset Lifecycle Management

Status: Bill framework for asset record amendment (s.83), supplemented by proposed lifecycle management design including version control tiers and growth parameters.

12.1 The Dynamic Data Asset and the Stable Attribute Profile (Structure vs Content)

Some data assets — live databases, real-time feeds, continuously updated AI training sets — change continuously. The Regime addresses this through the concept of the Stable Attribute Profile (SAP): the registration records the asset's structural attributes that define its identity, while permitting content to evolve within declared parameters. A registered asset whose SAP is maintained is treated as the same asset across all content updates.

The distinction between structural and content attributes is fundamental to the lifecycle framework:

Attribute Domain	Structural (identity defining) vs Content (variable within thresholds)
Schema	Structural: field names, data types, relational structure. Content: record count, field values, null rates.
Classification	All classification changes are always structural. Any change to Distribution Scope, Sensitivity Level, or AI-Use Tier triggers re-registration.
Source	Structural: named source systems, collection methodologies. Content: number of sources within declared scope, ingestion volume.
Curation	Structural: declared methodology, QA standards, anonymisation technique. Content: individual quality scores, error rates within tolerances.
Temporal	Structural: declared temporal scope, update frequency commitment (committed to in DDI). Content: actual date range, update intervals within declared frequency.

12.2 Three-Tier Version Control

The Regime establishes a three-tier version control system that calibrates notification and re-registration requirements to the significance of the change:

Tier	Description
Tier 1: Administrative Updates (No SAP amendment)	Content changes within declared thresholds: new records within growth parameters, routine quality corrections, scheduled refreshes. These are logged in the cryptographic audit trail but require no Registrar notification.
Tier 2: Material Amendments (SAP update required)	Structural changes that modify the asset's identity without fundamentally transforming it: new schema fields, anonymisation technique change, geographic scope expansion, new source system addition. Under the proposed framework, the foundation should notify the Registrar within 28 days and submit an amended SAP. The Registrar assigns a version suffix (e.g. DAR-00147-v2). The original registration date and reference are retained. Security interests and licences continue unless their terms specify otherwise.
Tier 3: Fundamental Change (New registration required)	Triggered when any of the following thresholds is crossed: any classification change (Distribution Scope, Sensitivity Level, or AI-Use Tier); schema transformation exceeding 40% field removal or renaming, or primary key relationship change; replacement of more than 60% of declared source systems; purpose transformation changing the asset's commercial character; or abandonment of declared curation methodology. The existing registration is marked 'superseded' and a fresh application is required. Security interests and licences do not automatically transfer to the new registration (subject to secured party's consent).

12.3 Growth Parameter Declarations

To provide objective criteria for distinguishing Tier 1 from Tier 2 changes, this Paper recommends that every SAP include quantitative growth parameter declarations specifying the expected operating envelope for the asset:

- Record volume: maximum growth rate per period (e.g. 'up to 30% annual growth').
- Source count: maximum number or range of contributing sources (e.g. '10–20 IoM-regulated banks').
- Temporal scope: rolling or fixed with end date (e.g. 'rolling 5-year window').
- Update frequency: minimum and maximum intervals (e.g. 'daily to weekly').

Actual growth exceeding declared parameters by more than 20% triggers a Tier 2 notification requirement within 28 days. Persistent exceedance across three consecutive

measurement periods triggers a mandatory AAP review. The baseline growth parameters are determined by data dedicator's commitment in DDI.

Growth parameter compliance is monitored continuously by the DAR Verification Agent (Section 13). The Agent compares actual asset metrics against declared parameters within a confidential computing enclave, producing TEE-attested Verification Reports that feed into the Register's integrity chain and trigger the graduated notification framework when thresholds are approached or breached.

12.4 Amendment of Asset Records

The Registrar updates Asset Records in two circumstances: on the application of the DAF council or any person with standing (s.83(1)), and on the Registrar's own initiative where accuracy requires it (s.83(6)). All amendments are recorded with full audit detail. The Register does not delete previous entries: all historical states are permanently accessible, creating an indelible title chain.

12.5 Accreditation and Remediation

A DAF must not utilise a data asset unless a valid accreditation is in place (s.85(1)). Accreditation may be suspended or revoked under the data governance framework (s.85(2)). When accreditation expires or is revoked, the Registrar marks the asset as 'remediation' (s.85(4)), triggering a 90-day period (extendable once by a further 90 days) in which the DAF must restore valid accreditation. During remediation, the personal property right remains in effect but the asset cannot be utilised. If accreditation is not restored, the asset is removed from the Register.

12.6 Transfer and Disposal

A registered data asset may be transferred between DAFs through a simplified or modified registration process (s.80(5)). When a DAF is wound up or a data asset is removed from the Register, disposal or reversion follows a statutory waterfall (s.81(3)).

12.7 Dormancy, Suspension and Deregistration

An asset that ceases updates for longer than twice its declared update frequency (or 12 months, whichever is shorter) is flagged as dormant by the Registrar. Security interests and licences are unaffected by dormancy status. Suspension may be triggered by non-renewal, material non-compliance reported by an AAP, or an enforcement notice from the Information Commissioner. Suspension prevents new interests or licences being registered but existing arrangements continue. Deregistration may be voluntary (where no outstanding security interests exist) or compulsory (where the foundation is wound up, suspension exceeds 12 months, or the asset no longer exists). Deregistered assets are

marked but not deleted from the Register: the historical record is preserved as evidence of prior registration.

12.8 Continuous Integrity Verification

The Register employs a checkpoint model based on the Estonian KSI Blockchain approach. Periodic hash checkpoints are generated at intervals no greater than the declared update frequency for each asset. Amendment hashes are generated immediately before and after any Tier 2 amendment, creating a cryptographic boundary between versions. For dynamic assets, continuous automated verification is performed by the DAR Verification Agent (Section 13), which monitors compliance against declared parameters and produces TEE-attested Verification Reports that are logged to the integrity chain. An annual AAP attestation confirms that the integrity chain is unbroken, all checkpoints and verification reports have been generated, and the current state of the asset is consistent with the registered SAP.

13. Automated Verification and Confidential Computing

Status: recommended technical architecture for automated verification; not itself a statement of enacted law.

13.1 The Verification Gap

For dynamic assets with continuous updates, declared growth parameters, and real-time AI-Use restrictions, periodic human review creates blind spots. The verification gap manifests in three ways: growth parameter exceedance may go undetected between annual AAP reviews; schema drift (gradual structural changes that individually fall below Tier 2 thresholds but cumulatively transform the asset) is invisible without continuous monitoring; and AI-Use compliance violations - such as data being ingested by a model at a tier above the declared AI-Use Tier - cannot be detected by the Registrar because the Register doesn't hold the underlying data.

13.2 The DAR Verification Agent (DAR-VA)

The DAR Verification Agent is a standardised DAR component specified by the Registrar and operated by the Registrar within the Register's infrastructure that performs continuous automated verification of a registered data asset against its Stable Attribute Profile and growth parameter declarations. The key design principle: the DAR-VA verifies *compliance against declared parameters* - it does not assess data quality (however, those checks may be added in later stages), validate business logic or perform content review.

The DAR-VA operates within the Register's confidential computing infrastructure.

It is a centralised service, operated by the Registrar as part of the Register's technical architecture (Section 22). The verification logic, TEE provisioning, attestation management, and report generation are all functions of the Register. This is consistent with the Registrar's oversight functions (Section 8.2) and with the principle that the Register is the authoritative source of integrity verification, not the foundation.

The foundation provides a secure access gateway.

To enable the DAR-VA to perform its verification checks, under the proposed architecture, the foundation would be required, as a condition of registration, to maintain a secure, standardised access gateway through which the DAR-VA can query the data asset's metadata, schema, record counts, and other parameters needed for verification. The gateway specification is published by the Registrar as part of the technical standards for the Register. The gateway provides read-only, metadata-level access; the DAR-VA does not ingest or store the underlying data content.

The TEE enclave is instantiated by the Register.

When the DAR-VA performs a verification cycle, it instantiates a TEE enclave within the Register's infrastructure, queries the foundation's access gateway from within the enclave, performs the parameter comparisons, and produces a TEE-attested Verification Report. The raw data accessed through the gateway is processed within the enclave and never persists outside it. The cryptographic attestation proves that the verification was performed by the Register's authorised code, in an unmodified enclave, and that the results were not tampered with.

The trust model is Register-centric.

Under the proposed architecture, the Registrar would control the verification logic, the TEE provisioning, and the attestation chain. The foundation controls access to its data asset through the gateway. Neither party trusts the other: the TEE ensures the Registrar cannot see the raw data, and the attestation ensures the foundation cannot manipulate the verification results. This mutual distrust model is the foundation of the confidential computing architecture described in Sections 13.3 and 20.4.

The DAR-VA checks seven dimensions:

Dimension	What the Agent Verifies
Record Volume	Current count vs declared growth parameters. Alerts at 80% threshold; mandatory notification at 100%.
Schema Compliance	Current schema structure vs registered SAP field definitions. Detects field additions, removals, type changes, and key modifications.
Source Integrity	Active source systems vs declared source list. Detects new undeclared sources or cessation of declared sources.
Temporal Compliance	Update frequency vs declared cadence. Flags dormancy (absence of updates beyond twice the declared frequency).
Sensitivity Boundary	For Sensitivity Level 1 (non-personal) assets: statistical re-identification risk assessment against declared anonymisation standard. Flags if risk exceeds declared threshold.
Data Protection	Performs rule-based and statistical checks designed to detect the likely presence of personal data or other protected data categories in circumstances where the registered classification or declared controls would make such presence material to compliance.
Integrity Chain	Verifies that hash checkpoints have been generated at required intervals and that the chain is unbroken.
Access Pattern	For assets with AI-Use restrictions (AI-0, AI-1): monitors access patterns for signatures consistent with AI training ingestion (bulk sequential reads, embedding generation patterns).

13.3 Confidential Computing and the Trust Model

Here's the critical design challenge: the verification agent inspects the actual data asset to perform its checks, but the Register doesn't hold the underlying data, and the Registrar should not have access to it. The solution is confidential computing.
The DAR-VA runs inside a Trusted Execution Environment (TEE) – an isolated enclave where code and data are encrypted in memory and inaccessible to the host operating system, hypervisor or system administrators. The verification logic is loaded into the TEE, the TEE accesses the asset's data store, performs the parameter comparisons, and produces a cryptographically attested Verification Report within an isolated execution environment. The design objective is that raw data is processed within the enclave and is not exposed to the Registrar or external operators, subject to the security assumptions and residual risks inherent in the approved TEE platform.

The trust model works as follows:

Party	Trusts	Does Not Need to Trust
Registrant	TEE hardware integrity (attested by manufacturer). Verification logic is the Registrar's published, auditable code.	The Registrar does not access raw data. No third party sees the underlying dataset.
Registrar	The signed Verification Report.	Does not need to access the raw data. Does not need to trust the registrant's self-reporting.
Third parties (secured parties, licensees)	The Registrar's Verification Report, which is anchored to the TEE attestation.	Do not need direct access to the asset or the registrant's infrastructure.

13.4 Verification Reports and Notification Framework

The DAR-VA produces periodic Verification Reports (at intervals matching the declared update frequency, minimum monthly) containing: a pass/fail against each of the seven dimensions; quantitative metrics (current record count vs declared threshold, current schema hash vs registered hash, etc.); the TEE attestation certificate; and a timestamp chain linking to the Register's integrity layer.

The notification framework operates on a graduated escalation:

1. **Green (compliant)**: All dimensions within parameters. Report logged to integrity chain. No external notification.
2. **Amber (approaching threshold)**: One or more dimensions at 80%+ of declared parameter. Foundation's Technical Data Steward notified automatically. Registrar notified if the condition persists for three consecutive reporting cycles.
3. **Red (parameter breach)**: One or more dimensions exceeded declared parameters. Foundation's Business and Technical Data Stewards and Data Enforcer notified. Registrar notified. 28-day Tier 2 notification clock starts automatically. If the breach is a Tier 3 trigger (classification change, schema transformation exceeding 40%),

the Registrar is immediately notified and the asset status is updated to "Under Review".

4. **Critical (integrity failure)**: Integrity chain broken, TEE attestation failure, or DAR-VA tampered with. Registrar immediately notified. Asset status updated to "Under Review". AAP notified for emergency review.

14. Access, Disclosure and Confidentiality

Status: proposed access and disclosure framework for implementation through regulations and Registrar guidance.

14.1 Tiered Access Model

The Register serves two competing interests: transparency and commercial confidentiality. These are reconciled through a Tiered Access Model:

Access Tier	Available To and Scope
Tier 0: Public Access	Any person, without registration or fee, via web portal and API. Can access: asset existence, DAI, asset name, asset type, classification profile (class code and overlay attributes), asset status, data quality grade, foundation name, stewardship role identifiers, Data Enforcer identifier and annual report and AI-Use Tier.
Tier 1: Professional Access	(If granted, not automatic) Regulated financial institutions, legal practitioners, accredited assurance providers, registered agents, its Data Enforcer and other approved professionals. Additional fields include: rights schedule summary, encumbrance notices, provenance summary, accreditation dates, provider identity, version history, and growth parameter declarations.
Tier 2: Confidential Access	The DAF council, the Registrar, and regulators with statutory access rights (limited to that right). Full access to all fields including: council members, detailed provenance records, compliance logs, processing logic identifiers, data protection declarations, and the full audit trail.

14.2 Safe Publication Principles

This Paper recommends that the Registrar apply safe publication principles to ensure that public disclosure does not inadvertently reveal commercially sensitive information, personal data, or information that could be used to reverse-engineer proprietary methods. The principles require that: public fields contain only metadata, never underlying data; classification profiles are published but detailed governance logs are not; cryptographic

hashes are published but the data they verify is not; and the Registrar conducts a disclosure impact assessment before adding any new field to the public tier.

15. Security Interests and Encumbrances

Status: Bill framework for security interests (s.81), supplemented by proposed interaction rules for version control.

A key objective of the Register is to enable registered data assets to serve as security for credit, thereby unlocking significant economic value. The Register includes a Security Interests Module for notice filing of security interests over registered data assets. The design aligns with the UNCITRAL Model Law on Secured Transactions (2016), designed to enhance the availability of credit for SMEs, which covers intangible movable property and uses registration as the primary method of making security interests effective against third parties.

15.1 Creation and Registration

A security interest over a registered data asset is created by agreement between the grantor (the asset holder) and the grantee (the secured party). It is perfected by registration of the interest in the Rights and Encumbrances fields of the Asset Record. An unperfected security interest is valid between the parties but is not enforceable against third parties who deal in good faith with the registered owner. The module records: the identity of the secured party, the nature and extent of the security interest, the date and time of registration, any conditions or limitations on the security, and the DAI of the secured asset.

15.2 Priority and Enforcement

Priority among competing security interests is determined by timestamp of registration: the first registered interest has priority, consistent with the UNCITRAL default rule. The secured party's remedies on default include: appointment of a receiver over the data asset, sale of the data asset through the Register, and licence of the data asset to generate revenue to discharge the secured obligation. The Registrar's role is limited to recording the interest and publishing it in accordance with the access framework — it does not adjudicate priority disputes, which are resolved through the s.90 mechanism or the courts.

Consideration: in addition to specific interests over a single DAI, the Register could also support portfolio wide Floating Charges. A grantee could register a charge over all current and future data assets held by a specific DAF within defined classification profiles (e.g. all "C1" and "C2" assets). This charge would crystallise into a fixed interest upon a default event or notice of insolvency.

15.3 Interaction with Version Control

Where a Tier 2 amendment is made to a secured asset, existing security interests automatically attach to the new version unless the security agreement specifies otherwise. Where a Tier 3 fundamental change triggers new registration, security interests do not automatically transfer. The secured party must consent to the new registration or enforce their security.

16. Data Protection Interface

Status: recommended data protection interface design for implementation through regulations and the data governance framework; not itself a statement of enacted law. Existing data protection obligations under the Applied GDPR and DPA 2002 continue to apply independently.

Foundational Principle

Registration on the Data Asset Register is NEVER a lawful basis for processing personal data under Article 6 of the Applied GDPR. This Paper recommends that this principle be stated on the face of the primary legislation. The Registrar is not a data protection regulator.

16.1 The IoM Data Protection Framework

The Isle of Man's data protection framework is established by the Data Protection Act 2018, which applies the EU General Data Protection Regulation as 'Applied GDPR' in Manx law. The IoM holds EU adequacy status (originally granted in 2004) and was granted UK law enforcement adequacy in February 2025. This dual adequacy status is a significant asset for the Regime, and this Paper recommends that the Department consult with the ICO before the Regime enters force to ensure that the data asset property framework does not inadvertently affect the IoM's adequacy position.

16.2 Mandatory Pre-Registration Verification

This Paper recommends that, before any data asset containing personal data can be registered, the following verification steps be completed:

Step	Requirement
1. DPA Registration	The controller should hold a current DPA registration. The DPA registration number should be recorded in the Group E fields.
2. Lawful Basis Declaration	For all assets at Sensitivity Level 2 or 3, the foundation should declare the Article 6 lawful basis for each category of processing. For special category data, the applicable Article 9 condition should also be declared.
3. Data Protection Impact Assessment	For assets at Sensitivity Level 3, the foundation should provide evidence that a DPIA has been completed (Article 35). Where

	the DPIA indicates high residual risk, evidence of prior consultation with the Information Commissioner (Article 36) should be provided.
4. Transparency Obligation	The foundation should confirm that data subjects have been informed, in accordance with Articles 13 and 14, that their data may be held within a registered data asset foundation.
5. Anonymisation Verification	For assets claiming Sensitivity Level 1 (non-personal), the AAP should verify the anonymisation or pseudonymisation methodology and assess re-identification risk.

16.3 Data Subject Rights Protocol

Data subject rights under the Applied GDPR take absolute priority over the property right created by registration. The Regime establishes the following protocols:

Right to erasure (Article 17): The right to erasure takes absolute priority. If erasure of personal data materially alters the asset beyond the Stable Attribute Profile, the foundation should notify the Registrar within 14 days for a version control determination.

Right to rectification (Article 16): The cryptographic hash chain should support 'immutability with correction' — recording the rectification event while maintaining audit integrity. Rectification does not constitute a Tier 2 amendment.

Right to restrict processing (Article 18): The foundation should cease processing of the affected personal data to the extent required by applicable data protection law and should notify the Registrar where the restriction materially affects the registered asset. Any ongoing licensing or exploitation arrangement should be reviewed and, where necessary, suspended, varied or terminated to the extent required for compliance.

Right to data portability (Article 20): The foundation should be capable of extracting an individual's personal data in a structured, commonly used format without disclosing the asset's proprietary structure or methodology.

16.4 Role of the IoM Information Commissioner

The Information Commissioner's role in the Regime is consultative and supervisory, not gatekeeping. For assets at Sensitivity Level 3, the Registrar should invite the Commissioner's advisory opinion (within a 28-day response window). The opinion would not be binding but should be considered and recorded. The Commissioner retains all existing enforcement powers under the Applied GDPR. The Commissioner may issue guidance on the interaction between the DAR and data protection obligations, including anonymisation standards.

Under the proposed framework, the Data Enforcer would act as the primary liaison with the Information Commissioner and should review DPIAs for high-risk processing activities before they may proceed as Reserved Matters.

17. AI Governance Position

Status: recommended AI-governance model for consultation; not itself a statement of enacted legal requirements unless expressly provided by the Bill or regulations.

This section sets out the AI-governance model recommended in this Paper. Unless and until adopted through regulations, guidance or technical standards, the measures described below should be read as proposed implementation design rather than enacted legal requirements.

17.1 AI-Use Tier Governance Requirements

The AI-Use Tier overlay establishes a proposed set of graduated governance requirements for different categories of AI-related use. This section sets out the AI-governance model recommended in this Paper for possible adoption through regulations, Registrar guidance and technical standards. These requirements respond to the EU AI Act (with GPAI provisions applicable from 2 August 2025[5]) and position the IoM as providing clear, internationally interoperable governance for AI training data.

Tier	Governance Requirements
AI-0: No AI Use	Technical measures must be implemented to prevent AI ingestion. Annual AAP certification confirming no AI use. Full rights reservation against any AI training or inference.
AI-1: Internal Analytics	AI processing log must be maintained recording all models applied to the data. Model outputs must be tagged with DAR provenance identifier. Full rights reservation against external AI use.
AI-2: Licensed AI Training	Licensing agreement must specify: model type and family, training methodology, output restrictions, provenance attribution requirements, and revenue-sharing terms where applicable. The Registrar records the existence of AI training licences in the Rights and Encumbrances fields.
AI-3: Open AI / GPAI	Published terms of use required. Machine-readable provenance metadata must be compatible with the EU AI Act training data summary template. DAR Provenance Certificates available for AI developers to reference in model cards and regulatory filings.

[5] https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

17.2 DAR Provenance Certificates

Under the proposed model, the Registrar would issue standardised DAR Provenance Certificates confirming: the registered owner and foundation identity, the asset's classification coordinates, AI licensing terms and restrictions, the cryptographic hash at the certification point, and the certificate's validity period. These certificates enable AI developers to demonstrate data provenance for regulatory compliance, including the EU AI Act's training data summary requirements[6] and emerging international standards.

17.3 Registration of AI-Generated Data

Data that has been generated or substantially augmented by AI systems may be registered, subject to additional requirements. A disclosure obligation requires that AI provenance be recorded in the Group D fields, including the model family, the nature of human curation processes applied, and any training data dependencies. AI-generated data cannot be classified at Sensitivity Level 3 without evidence of human review and validation equivalent to that required for non-synthetic data. Where AI-generated data has been produced by a model trained on DAR-registered data, the dependency must be disclosed; the Registrar should be empowered to refuse registration where circular dependencies would undermine the integrity of the Register.

To preserve the integrity of the Register, this Paper recommends that the Registrar should apply a Human-in-the-Loop (HITL) validation requirement for synthetic data assets, with the detailed trigger criteria specified in guidance or regulations.

This Paper further recommends that, where Asset B is generated using Asset A already on the Register, that dependency should be disclosed and, where technically feasible, cryptographically linked through the provenance record.

This Paper recommends that the Registrar should be empowered to refuse registration where the degree of synthetic generation or recursive dependency is such that the asset no longer satisfies the Regime's requirements for identifiable curation, governance and provenance. Any quantitative thresholds used for this purpose should be specified in guidance or regulations, together with the methodology for assessing them.

17.4 Machine-Readable Rights Expression

This Paper recommends adoption of the W3C Open Digital Rights Language (ODRL v2.2) as the standard for machine-readable expression of AI-use rights and restrictions (see Section 18 for the proposed semantic interoperability framework). ODRL policies expressing AI-Use Tier restrictions are embedded in each registered asset's metadata and in each DAR Asset Passport (Section 19), enabling AI developers and automated systems to detect and

[6] https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models

process declared usage terms. This approach is designed to be compatible, so far as practicable, with robots.txt-style signalling, emerging machine-readable rights reservation mechanisms relevant to the EU AI Act[7], and developing data licensing standards including the Linux Foundation's Community Data License Agreement framework.

[7] https://digital-strategy.ec.europa.eu/en/consultations/commission-launches-consultation-protocols-reserving-rights-text-and-data-mining-under-ai-act

18. Semantic Interoperability and Machine-Readable Metadata

Status: recommended technical standards and interoperability design; not itself a statement of enacted law.

18.1 Design Rationale

The Bill does not prescribe a detailed semantic or metadata standard for the Register. This section therefore sets out the technical interoperability design recommended in this Paper, with the aim of ensuring that Register records are machine-readable, semantically precise and interoperable with wider data, compliance and AI-governance ecosystems.

The proposed semantic layer draws on established open standards rather than a bespoke vocabulary. Subject to consultation and implementation feasibility, the Paper recommends a profile based on DCAT, DPROD and ODRL. The proposed semantic layer would draw on three complementary W3C and OMG standards, each addressing a distinct functional requirement:

Standard	Role in the Register
W3C DCAT v3 (W3C Recommendation, August 2024)[8]	Provides the foundational catalogue vocabulary. Every registered data asset is described as a dcat:Resource with dcat:Dataset and dcat:DataService sub-classes. DCAT v3 adds native support for versioning (dcat:version, dcat:previousVersion, dcat:hasCurrentVersion), checksums (spdx:checksum), and dataset series, all of which map directly to the Register's version control and integrity architecture. DCAT's adoption across EU data portals (data.europa.eu)[9], US federal catalogues (data.gov via DCAT-US 3.0), and thousands of institutional repositories ensures baseline interoperability.
OMG DPROD (beta-stage specification under development within OMG processes)[10]	Extends DCAT-oriented modelling to describe data products with input/output ports (modelled as DCAT DataServices), domain classification, lifecycle status, data quality metrics, and conformance to logical models. DPROD's concept of a governed, managed data product aligns closely with several

[8] https://www.w3.org/news/2024/data-catalog-vocabulary-dcat-version-3-is-a-w3c-recommendation
[9] https://interoperable-europe.ec.europa.eu/collection/rolling-plan-ict-standardisation/data-interoperability-rp-2024
[10] https://www.omg.org/spec/DPROD/1.0/Beta1/PDF

	elements of the proposed statutory definition of a data asset under the Regime, while not being identical to it. The Register's five-group field architecture maps to DPROD's structure: Group A fields to dprod:DataProduct and ownership properties; Group B fields to dcat:Dataset and dprod domain/lifecycle; Group C fields to ODRL policies; Group D fields to DPROD lineage and input ports; and Group E fields to DPROD data quality and conformance.
W3C ODRL Information Model 2.2 (W3C Recommendation)[11]	Provides the machine-readable rights expression layer. Every registered data asset's rights, permissions, prohibitions, obligations, and constraints - including AI-Use Tier restrictions, distribution scope limitations, and licensing terms - are expressed as ODRL policies.

18.2 The DAR Metadata Profile

This Paper recommends that the Registrar publish a 'DAR Metadata Profile' — a formal application profile of DCAT v3 / DPROD, extended with ODRL — as the canonical machine-readable representation of Register records. Subject to consultation, the Profile would define:

- Mandatory and optional properties for each of the five field groups, mapped to DCAT/DPROD classes and properties.
- Controlled vocabularies for Distribution Scope, Data Sensitivity, Commercial Intent, AI-Use Tier, Regulated Sector, and Residency Tier, published as SKOS concept schemes.
- ODRL policy templates for each AI-Use Tier, including machine-readable rights reservations against AI training, licensed training terms, and open-use declarations.
- SHACL shape constraints for validation, enabling automated conformance checking of registration submissions.
- JSON-LD and Turtle serialisations, ensuring compatibility with both web-native and semantic web toolchains.

It is proposed that the Profile be developed during Phase 2 of implementation, published as an open specification, and maintained by the Standing Technical Committee alongside the Register Data Dictionary. As the underlying standards evolve, the Profile would be updated accordingly through published change control.

18.3 ODRL Rights Expression for Data Assets

[11] https://www.w3.org/TR/odrl-model

The adoption of ODRL as the Register's rights expression language has specific advantages over a bespoke alternative. ODRL provides a formal vocabulary of actions (use, distribute, derive, aggregate, anonymise, commercialise and others), parties (assigners and assignees), constraints (temporal, spatial, purpose-based), and duties (attribution, payment, notification). This expressiveness maps directly to the rights structures the Register should record:

Register Concept	ODRL Expression
AI-0: No AI Use (full reservation)	odrl:Prohibition on odrl:derive, constrained to AI training/inference purposes. Machine-readable by web crawlers and AI data pipelines.
AI-2: Licensed AI Training	odrl:Permission on odrl:derive, with odrl:Constraint specifying model family, training methodology, and attribution requirements. Linked to registered licence in Group C fields.
Distribution Scope: Internal	odrl:Prohibition on odrl:distribute. Assignee constrained to the registering foundation and its controlled systems.
Distribution Scope: Controlled-Share	odrl:Permission on odrl:distribute, constrained to named assignees with active Data Sharing Agreements recorded in Group C.
Residency Tier: RT-3 Sovereign Gateway	odrl:Obligation on odrl:use, constrained by odrl:spatial to the IoM Sovereign Gateway for data at rest.

The ODRL policies are serialised in JSON-LD and published at the Public Access tier (Tier 0), enabling automated rights detection by data marketplaces, AI training pipelines, and cross-border interoperability tools.

18.4 Interoperability with EU Data Spaces

This Paper further recommends that the Register's metadata and rights-expression architecture be designed to interoperate, so far as practicable, with EU-facing data-space environments and other international interoperability frameworks. By building the Register around the same family of standards, the IoM would improve the prospects that registered data assets can be described, discovered and policy-checked within interoperable data-sharing environments, even though the IoM's legal model differs from the EU's access-and-governance approach.

This interoperability is strategically significant, but it should be understood as technical and semantic interoperability rather than automatic legal equivalence. Participation in any particular EU-facing data-space arrangement would remain subject to the governance rules, contractual terms and regulatory requirements of that environment.

19. The DAR Asset Passport: Tokenised Governance Metadata

Status: proposed technical design for tokenised governance metadata; not itself a statement of enacted law.

Core Concept

The DAR Asset Passport (DAR-AP) is a cryptographically signed, machine-verifiable credential - issued by the Registrar and bound to the Data Asset Identifier - that embeds the asset's classification, rights policies, provenance summary and integrity hash into a portable, self-contained token. The Passport travels with the asset as it moves through the ecosystem, ensuring that governance information is always available to counterparties without requiring direct access to the Register.

19.1 The Problem: Governance Information Loses Contact with the Asset

Under a conventional registry model, the register is a central record that must be actively consulted. When a data asset is licensed, shared across borders, ingested into an AI training pipeline, or used as collateral, the counterparty must query the Register to verify the asset's status, rights and classification. This creates three problems. First, latency and availability: cross-border counterparties may face connectivity, jurisdictional or access constraints that prevent real-time verification. Second, point-in-time uncertainty: the counterparty knows the status at the time of query but cannot verify the status at the time the asset was actually transferred or processed. Third, ecosystem fragmentation: once the asset leaves the registrant's controlled environment, the governance metadata - classification, rights restrictions, AI-Use Tier, provenance - may not travel with it, creating governance gaps downstream.

The DAR Asset Passport solves all three problems by tokenising the Register's governance metadata into a self-contained, cryptographically verifiable credential that is bound to the asset and travels with it.

19.2 Technical Architecture

The DAR Asset Passport is built on the W3C Verifiable Credentials Data Model v2.0 (W3C Recommendation, May 2025), which provides a standardised, technology-neutral

framework for cryptographically signed, machine-verifiable credentials. The architecture comprises:

Component	Specification
Issuer	The Data Asset Registrar, identified by a Decentralised Identifier (DID) anchored to the Register's KSI integrity layer. The Registrar's DID document is published and resolvable, enabling any party to verify the Registrar's signing key.
Subject	The registered data asset, identified by its DAI, which serves as the credential subject identifier.
Credential Payload	A JSON-LD document containing: the asset's DCAT/DPROD metadata (classification profile, lifecycle status, domain, data quality grade); ODRL rights policies (AI-Use Tier restrictions, distribution permissions, licensing terms); provenance summary (source asset DAIs, derivation type, AI provenance where applicable); the cryptographic integrity hash at the issuance point; the Stable Attribute Profile version; growth parameter declarations; and accreditation status with expiry date.
Proof	A cryptographic signature by the Registrar using the proof mechanism specified in the W3C VC specification (JSON Web Signature or Linked Data Proof). The proof binds the payload to the Registrar's DID and provides tamper-evidence.
Validity and Revocation	Each Passport has a defined validity period aligned with the asset's accreditation cycle. Under the proposed architecture, the Registrar would maintain a revocation registry (using the W3C Bitstring Status List mechanism) enabling real-time status checks for whether a Passport has been revoked due to remediation, deregistration, or material amendment.

19.3 How the Passport Travels with the Asset

The DAR Asset Passport is designed to be embedded, attached, or referenced at every point where the asset is shared, licensed, or transacted:

In licensing agreements: The Passport is referenced by URI and optionally embedded in the Data Sharing Agreement or licence document. The counterparty can independently verify the Passport's signature and check its revocation status without contacting the Registrar.

In AI training pipelines: The Passport's ODRL policies are machine-readable. AI data ingestion systems can automatically parse the AI-Use Tier restrictions, verify the asset's provenance, and generate compliant training data summaries for EU AI Act purposes.

In cross-border transactions: The Passport and the Portable Evidence Package are complementary (human-readable and machine-readable respectively). A counterparty in Singapore, Dubai, or the EU can verify the Passport independently, without needing bilateral arrangements with the IoM Registrar.

In data marketplaces: Marketplace platforms can ingest the Passport's metadata to populate catalogue listings, enforce rights restrictions, and display governance credentials to prospective purchasers.

In security interest transactions: Secured parties can verify the asset's current status, classification, and accreditation through the Passport, and the Passport's revocation status serves as an early-warning mechanism for material changes.

19.4 Passport Lifecycle

The DAR Asset Passport follows a defined lifecycle aligned with the asset's own lifecycle:

- Issuance: A Passport is issued by the Registrar upon full registration (s.80). The Passport records the state of the Register at the issuance point.
- Tier 1 updates: Content changes within the Stable Attribute Profile and accreditation renewals do not trigger re-issuance. The Passport remains valid.
- Tier 2 amendments: A material amendment triggers re-issuance of the Passport with updated metadata and a new version suffix. The previous Passport is revoked via the status list.
- Tier 3 fundamental change: The existing Passport is revoked. A new Passport is issued for the newly registered asset.
- Remediation: If the asset enters remediation, the Passport is suspended (not revoked). Counterparties checking the status list receive a 'suspended' indicator.
- Deregistration: The Passport is permanently revoked. The revocation record includes the deregistration timestamp and reason.

19.5 Relationship to Blockchain-Based Tokenisation

The Regime deliberately adopts the W3C Verifiable Credentials[12] model rather than blockchain-native tokenisation (NFTs or Soulbound Tokens) for the Passport layer. This choice is driven by three considerations. First, the Bill already creates the property right through statutory registration; a blockchain token would be duplicative of the legal function

[12] https://www.w3.org/TR/vc-data-model-2.0

and could create confusion about which instrument is legally operative. Second, blockchain-based tokens introduce jurisdictional complexities (chain selection, fork governance, validator set control) that are unnecessary when the Register already provides a constitutive record with KSI integrity. Third, the W3C VC model is technology-neutral and does not prescribe the anchoring mechanism: implementations may use blockchain-anchored DIDs (such as did:ion or did:ethr), web-anchored DIDs (did:web), or other methods, providing flexibility without lock-in.

However, the architecture does not preclude complementary blockchain anchoring. A foundation that wishes to anchor its DAR Asset Passport on a public blockchain for additional transparency may do so, provided the Register remains the legally authoritative record. The Registrar may also explore anchoring the revocation registry on a distributed ledger to provide decentralised verification of Passport status.

19.6 Alignment with International Frameworks

The DAR Asset Passport aligns with three emerging international frameworks. The EU Digital Product Passport (under the Ecodesign for Sustainable Products Regulation) uses W3C Verifiable Credentials and Decentralised Identifiers for product lifecycle data — the DAR-AP extends this model to data assets. The UN Transparency Protocol (UNTP) specifies Verifiable Credentials with did:web identifiers for supply chain traceability - the DAR-AP adopts a compatible credential format. And the GS1 Digital Link standard provides URI-based identification for products - the DAI system is designed to be expressible as a resolvable URI in the same pattern, enabling integration with supply chain and marketplace infrastructure.

20. Interaction with Adjacent Legal Regimes

Status: mixed — existing legal principles and proposed Register interaction rules.

20.1 Intellectual Property Law

The Register is not an IP registry and the Registrar makes no determination of IP status. However, the Attribute Statement and Provenance records maintained in the Register provide valuable evidence in IP proceedings. Where a data asset is subject to an IP right — including database rights under the Copyright Act 1991 — this is reflected in the Rights and Encumbrances fields as an IP Constraint Notice. Registration of a data asset does not create, extend, or diminish any intellectual property right. The regime is designed to complement, not compete with, existing IP protections.

20.2 Financial Services Regulation

Where a data asset is connected to a regulated financial activity — for example, used as collateral for a financial instrument, or incorporated into an algorithmic trading system — the Register interfaces with the Isle of Man FSA's oversight. This Paper recommends that the Registrar be required to cooperate with the FSA and provide information upon lawful request. Assets classified under RS-F (Financial Services) are subject to additional governance requirements including enhanced due diligence on provenance and beneficial ownership transparency.

20.3 Insolvency Law

A registered data asset forms part of the insolvent estate of its holder and vests in any liquidator, administrator, or receiver. The Registrar updates the Register upon notification of an insolvency appointment. The trustee in insolvency has the same rights to apply for amendments, grant licences, and transfer assets as the registrant would have had. Where an IoM foundation holds data registered under the Regime and a connected entity enters insolvency proceedings in another jurisdiction, the UNCITRAL Model Law on Cross-Border Insolvency provides the framework for determining the interaction between IoM foundation rules and foreign insolvency administrators' powers.

21. Cross-Border Recognition and Interoperability

Status: mixed — Bill framework for equivalent registers, plus recommended cross-border evidence and interoperability design.

21.1 Equivalent Data Registers

The Bill provides for the Department to specify alternative data registers as equivalent (s.82). This subsection describes that statutory gateway. The subsections that follow set out additional cross-border measures recommended in this Paper. An asset registered on an equivalent data register can be registered on the IoM Register without repeating the accreditation process. The Department considers: the relevant laws establishing the alternative register; the quality of accreditation required; the effective functioning of a registrar; the information made available; and the oversight requirements (s.82(4)). As a matter of policy exploration, the most plausible candidates for early equivalence discussions may include common-law jurisdictions (Singapore, Dubai, Bermuda, Cayman Islands) and China, where philosophical alignment with data property frameworks is strongest. The EU's philosophical rejection of property rights in data means that EU member states are unlikely to establish equivalent registers; the Regime's cross-border strategy should therefore focus on the Portable Evidence Package and DAR-AP as the primary interoperability tool for EU interactions.

21.2 Portable Evidence Package (PEP) and DAR Asset Passport

To support cross-border reliance, this Paper recommends that the Registrar issue Portable Evidence Packages — structured, verifiable bundles of Register evidence for use in transactions, proceedings, and due diligence in other jurisdictions. A PEP includes: a Certified Extract of the relevant fields; the asset's classification profile; the current accreditation status and provider identity; the cryptographic integrity hash; a DAR Provenance Certificate (where AI-Use governance is relevant); and a Registrar attestation confirming the extract's authenticity. PEPs are designed to be independently verifiable without requiring direct access to the Register.

For machine-to-machine interoperability, the DAR Asset Passport (Section 19) serves as the digital complement to the PEP. Where a PEP is a human-readable evidence bundle for legal and due diligence purposes, the Passport is a machine-verifiable credential for automated systems. Under the proposed model, a cross-border counterparty would be able to verify the Passport's cryptographic signature and check its revocation status without requiring direct bilateral technical integration with the IoM Registrar. This may reduce transaction

friction in international data marketplaces and governance workflows, although legal reliance will still depend on the receiving jurisdiction's rules and the transaction context.

21.3 International Standards Engagement

This Paper recommends that the Registrar actively engage with emerging international standards work, including ISO data governance standards, OECD work on data as an economic asset, the EU AI Act's training data transparency requirements, W3C Verifiable Credentials interoperability initiatives, and multilateral frameworks for data exchange. The DAI system is designed to be compatible with emerging international persistent identifier standards and expressible as resolvable URIs for integration with GS1 Digital Link and similar infrastructure. The DAR Metadata Profile may be published as an open specification for data asset description, with the aim of encouraging wider interoperability and external scrutiny.

22. Technical Architecture

Status: recommended technical architecture for consultation; not itself a statement of enacted law.

22.1 Design Philosophy

The technical architecture is governed by four overriding principles: integrity (contents cannot be altered without authorisation and all changes are permanently recorded); accessibility (authorised users can access the Register efficiently); resilience (the Register continues to function despite failure, attack, or force majeure); and interoperability (built on open standards enabling integration with domestic and international systems).

22.2 Five-Tier Hybrid Architecture

Tier	Function
Tier 1: Cryptographic Integrity Layer	Provides tamper-evidence and non-repudiation using the KSI Blockchain model (operational in Estonia since 2012 across government registries)[13]. Each Asset Record and amendment generates a cryptographic hash. Periodic checkpoints align with declared update frequencies. Amendment hashes create cryptographic boundaries between asset versions.
Tier 2: Legal Index and Authoritative Record	The human-readable, legally operative record of all Asset Records. This is the layer that generates legal effects under the Act. Maintained in a structured database with rigorous access control and audit standards. Supports 'immutability with correction' for data protection compliance. Records are serialised in the DAR Metadata Profile (DCAT v3 / DPROD / ODRL) for machine-readable access.
Tier 3: Public Access and API Layer	The interface for external users. Comprises a web portal for human users and a standardised API (REST/JSON-LD, OpenAPI specification) for machine-to-machine access, enabling integration with data marketplaces, legal platforms, and international registries. ODRL policies are published at this tier, enabling automated rights detection.
Tier 4: DAR Asset Passport Layer	The portable credential layer. Produces, manages, and revokes W3C Verifiable Credentials (DAR Asset Passports) for registered data assets. Maintains the Registrar's DID and signing infrastructure. Publishes the Bitstring Status List for

[13] https://e-estonia.com/wp-content/uploads/faq_ksi_blockchain.pdf

	Passport revocation checking. Enables offline and cross-border verification without direct Register access.
Tier 5: DAR Verification Agent and TEEs	The missing operational layer that turns the paper's governance declarations from static claims into continuously verified facts. Without it, the growth parameters and version control system are enforced only by periodic human review and self-reporting - which is insufficient for dynamic assets. The Agent running inside a TEE solves the fundamental tension in the design: *the Registrar needs to verify compliance but must never access the underlying data.*

22.3 Security Architecture

The Register's security architecture is developed to ISO 27001 standard or above and is subject to annual independent security audit. Key measures include: multi-factor authentication for all back-end access; role-based access controls; encryption at rest and in transit; annual independent penetration testing; published incident response protocols; geographic redundancy across at least two physically separate data centres; and data sovereignty with all primary infrastructure located in the Isle of Man.

22.4 Confidential Computing

DAR must be supported by confidential computing through Trusted Execution Environments to enable secure data interrogation and continuous monitoring without exposing underlying data assets - even to the platform operator or hosting infrastructure. This capability is foundational to the Regime's trust model, particularly for the operation of an automated verification agent that monitors data assets without ever accessing raw data content.

TEE-based confidential computing provides a materially stronger isolation model than conventional processing environments, including hardware backed attestation and protection against many classes of host level access, but it remains subject to platform specific limitations, implementation risk and disclosed vulnerabilities.[14]This makes it possible for a verification agent to perform meaningful governance functions whilst providing cryptographic proof to all stakeholders that the underlying data was never exposed during the process.

Together with the Cryptographic Integrity Layer, these two mechanisms form a complete, cryptographically grounded trust chain: blind verification producing signed evidence, and permanent tamper-proof storage of that evidence.

[14] https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8320.pdf

22.5 Scalability Considerations

The DAR-VA architecture requires continuous TEE-attested verification cycles for each registered dynamic asset, at intervals matching the declared update frequency (minimum monthly). As the Register grows to potentially thousands of registered assets, each with its own verification cadence, the scalability of the confidential computing infrastructure becomes a material design constraint. Estonia's KSI Blockchain implementation, which provides the model for the integrity layer, operates at document-level hash frequency for government registries; the DAR-VA architecture implies a significantly higher verification throughput, particularly for assets with daily or sub-daily update frequencies.

Additionally, the TEE attestation chain depends upon the root of trust maintained by the hardware manufacturer (Intel, AMD, or Arm). This effectively delegates a critical element of the Register's integrity assurance to a third-party commercial entity. The Registrar should adopt a technology-neutral specification that defines functional requirements (memory encryption, remote attestation, sealed storage) rather than mandating specific hardware platforms, enabling the Register to migrate between TEE providers as the market evolves and as vulnerabilities are discovered in specific implementations. This Paper recommends that the Registrar maintain an approved TEE platform list, with criteria for addition and removal, and a migration protocol for transitioning between platforms without interrupting verification coverage.

The phased implementation approach should address scalability through graduated deployment: initial registration volumes are expected to be modest, providing a period in which the infrastructure can be tested and optimised before the Register reaches scale. The Registrar should publish capacity planning projections as part of the technical standards, and the fee structure should reflect the true marginal cost of continuous verification to ensure financial sustainability.

23. Liability, Indemnification and Insurance

Status: mixed — existing legal principles, proposed regulatory provisions, and recommended administrative/insurance measures.

Any regime that creates a constitutive register must address the allocation of loss when registered rights are relied upon and that reliance proves misplaced. The Bill provides the statutory setting for those issues, but it does not by itself exhaustively codify all civil liability consequences arising from operation of the Register. The Register's authoritative status (Section 5.1) means that entries are presumed, in the absence of fraud or manifest error, to accurately reflect the legal position. Third parties - licensees, secured creditors, AI developers relying on provenance certificates, counterparties in data marketplace transactions - will transact on the faith of the Register. Where the registered position proves inaccurate, incomplete, or stale, the question of who bears the resulting loss is fundamental to the regime's credibility and to market confidence.

This section therefore distinguishes between: (a) existing principles of Manx law that may continue to apply; (b) liability rules that this Paper recommends should be addressed expressly by regulations or legislation; and (c) administrative and insurance measures recommended to support risk allocation in practice.

Unless expressly stated otherwise, the liability allocations described below should be read as proposals for consultation rather than as settled statements of enacted law.

23.1 Design Principles

The following design principles inform the liability framework proposed in this Paper:

- **Primary liability follows the declaration.** The entity that makes a declaration to the Register bears primary responsibility for its accuracy. The DAF council declares the Stable Attribute Profile, growth parameters, classification coordinates, and rights position. The AAP certifies data governance compliance. The Data Enforcer provides ongoing monitoring. Each is liable for the accuracy of its own declaration, certification or actions.
- **The Registrar records; it does not warrant.** The Registrar's function is to receive, verify the formal completeness of, and record information submitted by applicants and assurance providers. The Registrar does not independently verify the substantive truth of declarations about the underlying data asset. The authoritative presumption (Section 5.1) is a procedural presumption about the Register's accuracy, not a warranty by the Registrar of the underlying facts.

- **Automation does not eliminate accountability.** Where the DAR Verification Agent produces verification reports that are relied upon by the Registrar, the foundation, or third parties, liability for failures in the verification chain must be clearly allocated. Automated systems create new liability questions; they do not displace existing ones.
- **Professional assurance carries professional liability.** AAPs perform a professional certification. Their liability is commensurate with their professional obligations and is backed by mandatory professional indemnity insurance.
- **Proportionality and insurability.** Liability must be proportionate to the role performed and must be insurable at commercially reasonable rates. An uninsurable liability regime would deter registration and undermine the regime's accessibility objective.

23.2 Foundation Liability for Declarations

The DAF council is the primary declarant to the Register. Through the Data Asset Dedication Instrument and the registration application, the council declares the asset's structural attributes, classification coordinates, growth parameters, rights position, data protection basis, and provenance. These declarations form the Stable Attribute Profile upon which the entire registration rests.
This Paper proposes that the foundation should bear primary liability for materially inaccurate declarations made to the Register, subject to the detailed scope, caps and exceptions specified by regulations or legislation.

Specific heads of foundation liability include:

- **Inaccurate SAP declarations:** where the declared structural attributes (schema, field architecture, source composition, temporal scope) do not accurately describe the asset at the time of registration or following a notifiable amendment.
- **Growth parameter misstatement:** where declared growth parameters are set deliberately or recklessly wide to avoid triggering Tier 2 notification obligations, or deliberately narrow to understate the asset's true operating envelope.
- **Classification misstatement:** where the declared classification coordinates (Distribution Scope, Data Sensitivity, overlay attributes) do not accurately reflect the nature or intended use of the asset. Misclassification of data sensitivity is treated with particular severity given the data protection implications.
- **Rights declaration errors:** where the declared rights position (IP status, licensing terms, encumbrances) does not accurately reflect the legal position, causing a third party to transact in reliance on an incorrect rights profile.
- **Failure to notify amendments:** where a Tier 2 or Tier 3 change has occurred and the foundation fails to notify the Registrar within the prescribed period, causing the Register to present a stale position.
- **DAR-VA obstruction:** where the foundation interferes with, disables, or restricts the DAR Verification Agent's access to the data asset or its operating environment, preventing continuous verification from functioning as designed.

Self-attested micro-assets. Where a foundation registers under the Micro-Asset Tier with self-attested full registration (Section 7.1), the foundation assumes the liability that would otherwise fall on the AAP.

23.3 Proposed regulatory allocation: Accredited Assurance Provider Liability

The AAP performs a professional certification function that is central to the regime's integrity. Full registration under s.80 requires AAP accreditation, and the property right vests only upon that accreditation being confirmed. Third parties are entitled to rely on the AAP's certification as evidence that the asset meets the data governance framework requirements at the time of accreditation.

The AAP's liability is for the professional competence and diligence of its certification, not for the underlying truth of the foundation's declarations. The distinction is critical: the AAP certifies that the foundation's governance arrangements meet the prescribed standard, not that the data asset itself possesses particular commercial qualities. This mirrors the position of a statutory auditor, who certifies that accounts give a true and fair view, not that the business will succeed.

Specific heads of AAP liability include:

- **Negligent accreditation:** where the AAP certifies compliance with the data governance framework without performing the procedures required by the applicable accreditation standard set out by the regulations, or where the AAP's procedures were so deficient that a reasonably competent assurance provider would have identified the non-compliance.
- **Failure to detect material misstatement:** where the SAP contains a material misstatement that the AAP's prescribed verification procedures were designed to detect, and the AAP failed to detect it through negligence.
- **Failure to report:** where the AAP becomes aware of circumstances that would require qualification or withdrawal of its accreditation and fails to notify the Registrar within the prescribed period.
- **Scope misrepresentation:** where the AAP's accreditation report implies a scope of assurance broader than actually performed, leading third parties to place reliance beyond the work done.

Limitation. An AAP is not liable for losses arising from the foundation's post-accreditation conduct, changes to the asset that occur after the accreditation date or matters outside the scope of the accreditation standard. The AAP's liability is subject to a monetary cap set by regulations, calibrated to the asset's classification profile and the scope of the accreditation engagement and backstopped by mandatory professional indemnity insurance.

23.4 Proposed regulatory allocation: Registrar Liability

The Registrar occupies a distinctive position. As a statutory officeholder exercising public functions, the Registrar's liability framework must balance accountability for the competent exercise of those functions with the public interest in a Registrar who is willing to act decisively and without excessive caution.
The Registrar's liability is confined to the proper discharge of its statutory functions and does not extend to a warranty of the substantive accuracy of information declared by foundations or certified by AAPs. Section 4.3 establishes that the Register records declarations and provides notice; it does not adjudicate or warrant. The authoritative presumption in Section 5.1 is a rule of evidence for third-party reliance, not an assumption of liability by the Registrar for the content of the Register.

This Paper recommends that the Registrar's liability be confined to the following categories of case:

- **Administrative error:** where the Registrar records information incorrectly (e.g. transcription errors, incorrect DAI assignment, failure to record a notified amendment), causing the Register to diverge from what was properly submitted.
- **Failure to act on notifications:** where the Registrar receives a notification that requires action (e.g. a Tier 2 amendment notification, an AAP accreditation suspension notice, a DAR-VA Critical alert) and fails to update the Register within published service standards, causing a third party to rely on stale information.
- **Wrongful refusal or removal:** where the Registrar refuses a registration application, or removes an asset from the Register, in circumstances where no reasonable Registrar applying the statutory criteria could have reached that decision. This mirrors the judicial review standard applied to public decision-makers under Manx administrative law.

23.5 Proposed statutory or regulatory limitation on Registrar liability.

This Paper recommends that legislation or regulations include a statutory limitation on the Registrar's aggregate liability, consistent with the approach taken for other Manx statutory officeholders and equivalent international registrars. The proposed structure is: (a) no liability for the substantive truth of declarations or certifications made by others; (b) liability for administrative errors capped at the direct losses flowing from the error; (c) no liability for consequential or indirect losses except in cases of fraud or wilful misconduct; and (d) an aggregate annual liability cap set by regulations, funded through the fee structure and backstopped by appropriate insurance.

23.6 Recommended treatment of automated verification failures

The DAR Verification Agent introduces a novel liability question: when an automated system operating within a confidential computing enclave produces verification reports that are relied upon by the Registrar, the foundation, and potentially third parties, who is liable if the verification fails?

The framework distinguishes three categories of DAR-VA failure:

Specification Failure

Where the DAR-VA's verification logic, as specified by the Registrar, contains an error that causes it to produce incorrect verification reports (e.g. a threshold comparison that uses the wrong operator, or a schema hash check that fails to detect a material structural change), liability falls on the Registrar as the specifier. This is an extension of the Registrar's administrative error liability: the DAR-VA is an instrument of the Registrar's oversight function, and errors in its specification are attributable to the Registrar. To mitigate this risk, the paper recommends mandatory independent audit of the DAR-VA verification logic (see Consultation Question 20) and a formal change control process for updates to the specification.

Deployment and Infrastructure Failure

Where the DAR-VA is correctly specified but fails due to deployment or infrastructure issues within the registrant's environment (e.g. the TEE is not properly provisioned, the DAR-VA's access to the data store is intermittent, or the host environment interferes with enclave operation), liability falls on the foundation as the party responsible for providing the deployment environment. The foundation's obligation to maintain a functioning DAR-VA environment is an incident of registration: it is a condition of continued registration that the DAR-VA can operate as designed.

TEE Attestation Failure

Where the TEE hardware attestation is compromised (e.g. a hardware vulnerability that undermines the integrity guarantee of the enclave), liability allocation depends on the nature of the compromise. A known vulnerability that the Registrar or foundation failed to patch is attributable to the party responsible for the relevant infrastructure. A zero-day hardware vulnerability is treated as a force majeure event: no party is liable for losses arising from a hardware security failure that was not known and could not reasonably have been anticipated, provided the affected party takes immediate remedial action upon discovery. The Registrar must maintain a vulnerability management protocol for approved TEE platforms and issue security advisories when material vulnerabilities are disclosed.

Reliance on DAR-VA reports.

The Register's status indicators (e.g. 'Under Review', 'Remediation') will be partly driven by DAR-VA reports and are visible at the Public Access tier. A third party who relies on a Public tier status indicator that is incorrect because a DAR-VA failure prevented a status change does not have a direct claim against the DAR-VA infrastructure but may have a claim against the Registrar.

False Negative Verification (Undetected Breach)

The most commercially consequential DAR-VA failure mode is the false negative: where the Agent produces a Green (compliant) verification report despite the asset having breached its declared parameters, and a third party transacts in reliance on the compliant status. This scenario requires careful liability allocation. Where the false negative results from a

specification error (the verification logic failed to check the relevant dimension or applied the wrong threshold), liability falls on the Registrar as specifier, as described above. Where the false negative results from the foundation presenting manipulated or incomplete data through the access gateway, liability falls on the foundation for obstruction and misrepresentation (Category 3 breach under the enforcement framework). Where the false negative results from a limitation inherent in the verification methodology (for example, a statistical re-identification assessment that fails to detect a novel re-identification technique), the position is more nuanced: the Registrar is liable only to the extent that the verification methodology fell below the standard that a reasonably competent regulator would have specified, having regard to the state of the art at the time the methodology was published. The framework does not impose strict liability on the Registrar for the inherent limitations of automated verification. To mitigate false negative risk, the paper recommends that the periodic AAP review includes an independent assessment of DAR-VA detection efficacy, and that the Registrar publishes the verification methodology (excluding security-sensitive implementation details) to enable independent scrutiny.

23.7 Proposed clarification of third-party reliance rules: Third-Party Reliance and the Authoritative Presumption

The authoritative presumption (Section 5.1) provides that entries in the Register are presumed, in the absence of fraud or manifest error, to accurately reflect the legal position as at the time of entry. This presumption is the basis on which third parties transact. The liability framework must define both the scope of the presumption and the limits of permissible reliance.

Scope of the presumption.

The presumption applies to the registered facts: the identity of the holding foundation, the asset's classification coordinates, the declared SAP, the accreditation status, the existence and priority of registered security interests, and the asset's current lifecycle status. It does not extend to the commercial value of the asset, the quality of the underlying data, the future performance of the asset, or any matter not recorded on the Register.

Reasonable reliance.

A third party's claim for loss arising from an inaccurate Register entry requires proof of reasonable reliance. Reliance is reasonable where the third party: (a) consulted the Register or verified the DAR Asset Passport before transacting; (b) relied on information within the scope of the authoritative presumption; and (c) had no actual knowledge of facts inconsistent with the registered position. A party who relies on the Register despite actual knowledge of a discrepancy, or who fails to consult the Register when a reasonable person in their position would have done so, cannot claim the benefit of the presumption.

Point-in-time reliance.

The Register is dynamic (Section 5.1). The presumption operates as at the time of the relevant entry or query. A third party who transacted in reliance on a Register query at time T1 is protected against inaccuracies existing at T1, even if the Register is subsequently corrected at T2. However, a party who transacts at T2 on the basis of a query made at T1

bears the risk of intervening changes. The DAR Asset Passport mitigates this risk by providing a cryptographically signed snapshot of the registered position at a defined point in time, with revocation status that can be checked independently.

23.8 Consultation question: possible statutory indemnity fund.

The Department may consider, in the longer term, whether the Regime should establish a statutory indemnity fund, analogous to the Land Registry indemnity under the Land Registration Act 2002 (England and Wales), to compensate third parties who suffer loss through no fault of their own as a result of errors in the Register. If adopted, the fund would be financed through a component of the registration and annual maintenance fees, and claims would be determined by the Registrar (with recourse to the s.90 dispute mechanism). The advantage of a statutory indemnity is that it provides a direct remedy for innocent third parties without requiring them to identify and pursue the party responsible for the error.

23.9 Recommended insurance and risk-management measures: Insurance and Indemnification Requirements

This Paper recommends that the regime impose mandatory insurance requirements on two categories of participant:

Accredited Assurance Providers.

This Paper recommends that, as a condition of accreditation under the data governance framework, every AAP should be required to maintain professional indemnity insurance with minimum coverage levels prescribed by regulations.

Data Asset Foundations (for self-attested micro-assets).

Where a foundation registers under the Micro-Asset Tier with self-attestation, the foundation must either maintain appropriate insurance coverage for the liability assumed in place of AAP accreditation or provide a personal guarantee from the council members. The minimum coverage or guarantee amount is prescribed by regulations and is proportionate to the micro-asset threshold.

Consideration

The Registrar's own insurance arrangements are a matter for the Department, but it is recommended that the Registrar maintain appropriate coverage funded through the fee structure, and that the coverage position be disclosed in the Registrar's annual report.

23.10 Liability Allocation Summary

The following matrix summarises the allocation of liability across the principal failure scenarios:

Failure Scenario	Foundation	AAP	Registrar	Mitigation

Inaccurate SAP declaration	Primary	Secondary (if negligent accreditation)	None	DAR-VA continuous monitoring
Classification misstatement	Primary	Secondary	None	Pre-registration verification
Negligent accreditation	None (unless collusion)	Primary	None	PII insurance; AAP assessment
Register transcription error	None	None	Primary (strict)	Internal QA; indemnity fund
Failure to update on notification	None	None	Primary	Service standards; escalation
Failure to notify Tier 2 amendment	Primary	None	None	DAR-VA detection; sanctions
DAR-VA specification error	None	None	Primary	Independent audit of VA logic
DAR-VA deployment failure	None	None	Primary	Deployment standards; monitoring
False negative verification	Primary (if gateway manipulation)	None	Primary (if specification deficiency)	AAP detection efficacy review; published methodology
TEE hardware compromise (zero-day)	Force majeure	Force majeure	Force majeure	Vulnerability management protocol
Confidentiality breach	If source of breach	If source of breach	Primary (if Register breach)	Access Framework; encryption
Self-attested micro-asset inaccuracy	Primary (enhanced: council personal liability)	N/A	None	24-month limit; insurance/guarantee
Passport reliance after revocation	If failure to notify	None	If failure to update status list	Passport revocation checking

Failure Scenario	Foundation	Data Enforcer	Registrar	Mitigation

Negligent unqualified opinion	None	Primary (gross negligence)	None	Safe harbour; PII insurance
Failure to exercise stop power	Secondary (if continued exploitation)	Primary (gross negligence)	None	Statutory triggers; DAR-VA alerts
Wrongful stop direction	Claimant (loss suffered)	Primary (reasonablenes s test)	None	High Court discharge mechanism
Negligent consent to Reserved Matter	Secondary (if exploiting asset)	Primary (gross negligence)	None	AAP pre-consent verification
Failure to report to Registrar/ICO	None	Primary (breach of statutory duty)	None	Mandatory reporting obligations
Acting under prohibited conflict	None	Primary (strict liability)	None	Conflicts protocol; Approved Register

23.11 Existing legal principles likely to continue to apply: Interaction with Existing Liability Regimes

The Regime's liability framework operates alongside, and does not displace, existing causes of action under Manx law:

- **Tort (negligence and negligent misstatement):** A third party who suffers loss through reliance on inaccurate registered information may pursue a common law negligence claim against the responsible party, subject to the usual requirements of duty, breach, causation and loss. The liability framework in this section clarifies and structures these obligations but does not create a statutory bar to common law claims.
- **Contract:** Contractual liability between transacting parties (e.g. under a data sharing agreement, licence, or security agreement) is governed by the terms of the relevant contract. The Register's liability framework does not override contractual allocations of risk between parties who have negotiated their own terms.
- **Data protection:** Liability for data protection breaches (including breaches arising from the registration of personal data) is governed by the Applied GDPR and the Data Protection Act 2018. The Registrar may be a data controller in respect of information held on the Register, and the foundation is the data controller in respect of the underlying data asset. Data protection liability is not affected by the liability allocations in this section.
- **Fraud:** No limitation or cap in this framework applies to liability arising from fraud. A foundation that makes fraudulent declarations to the Register, an AAP that issues fraudulent accreditations, or a Registrar officer who acts fraudulently is liable without

limitation. Fraud also displaces the authoritative presumption: a third party who proves fraud is not bound by the registered position.

23.12 Limitation Period

Claims arising under the Regime's liability framework should be subject to a limitation period to be specified by statute or regulations. A six-year period may be appropriate for many claims by analogy with general limitation rules applicable to tort and related civil claims under Manx law. If the policy intention is to adopt a discoverability-based rule for certain Register-related claims, that should be stated expressly as a specific rule of the Regime rather than assumed to reflect the general position.

The dynamic nature of the Register creates a particular challenge where an error may persist for some time before discovery. The consultation therefore invites views on whether a tailored knowledge-based limitation rule should apply to specified categories of Register-related claim.