



PRODUCT DOSSIER · SINGAPORE REGULATORY INTELLIGENCE

Singapore Regulatory Collection

Structured MAS Technology and Outsourcing Obligations, Cited and Evidence-Backed

Dataset ID	PROFYTAI-SG-MAS-COLLECTION-1.0.0
Version	1.0.0 · Released 2026-07-07
Jurisdiction	Singapore
Issuing Authority	Monetary Authority of Singapore (MAS)
Coverage	393 obligations across three MAS instruments
Classification	Commercial Regulatory Intelligence · Licensed Distribution

What This Dataset Is, and Why It Exists.

The Singapore Regulatory Collection is the complete, examination-ready obligation register for a financial institution's core MAS technology and outsourcing perimeter. It compiles three instruments, the Technology Risk Management Guidelines and the binding Cyber Hygiene and Outsourcing Notices, into 393 discrete, individually cited obligations, each carrying the regulator's exact words, a page-anchored citation, the parsed duty, and a source-verified evidence capture.

Who It Is For

Bank and fintech compliance teams, risk officers, internal audit, the consultancies that advise them, and the procurement and legal teams that license the data. It suits any institution that must evidence its posture against MAS technology and outsourcing expectations.

Why It Exists

Published regulation is authoritative but unstructured: hundreds of duties buried in prose across several PDFs. Building a usable, cited register by hand is slow, fragile, and hard to prove. This dataset does that work once, correctly, and ships it as structured data.

EXPECTED BUSINESS OUTCOMES

A defensible register on day one	Replace weeks of reading and re-keying with a complete, cited obligation set your team reviews rather than builds.
Understanding built in	Every record carries ProfytAI Regulatory Intelligence: a plain-language summary, why the duty exists, and how teams implement it, so analysts, engineers, and executives grasp a requirement in minutes rather than hours.
Audit-ready provenance	Every duty is quoted verbatim and page-anchored, so a finding traces to the source in one step.
Lower key-person risk	The regulator's expectations are captured as structured data, not held in one analyst's spreadsheet.

PRIMARY USE CASES

Obligation registers · gap analysis · audit and examination readiness · policy mapping · regulatory change management · evidence collection · AI grounding. Each is expanded later in this dossier.

Coverage, Method, and Provenance at a Glance.

Obligations	393 discrete, individually cited duties
Source Documents	Three MAS instruments
Schema	46 fields in four structured layers
Formats	Excel workbook, JSON (versioned, checksummed envelope), CSV, plus the source PDFs
Evidence Captures	Included · one source-verified page capture per obligation
Update Model	Point-in-time and versioned. Errata reissued free within your version; a new version on MAS reissuance is a new license or the optional annual update subscription
AI Regulatory Intelligence	Generated summary, purpose, relationship, and implementation notes on every record

SOURCE INSTRUMENTS

MAS Technology Risk Management Guidelines	2021	344 obligations
MAS Notice FSM-N06 on Cyber Hygiene	2024	9 obligations, binding
MAS Notice 658 on Outsourced Relevant Services	2023	40 obligations, binding

HOW IT IS BUILT

- **Extraction.** Each instrument is parsed from its official MAS PDF into discrete obligations at the clause level, one record per duty.
- **Normalization.** A parallel normalized field prepares the text for search and diffing while the verbatim stays byte-exact for citation.
- **Traceability.** Every record carries a legal citation, the printed page, and a link to the source PDF; in the Collection it also carries a source-verified page image.
- **Update philosophy.** Coverage is point-in-time: the exact text as issued, versioned, with a new dated release on reissuance and a free errata reissue if a confirmed error is found.
- **Quality assurance.** The build rejects any record missing its verbatim or citation and produces byte-identical, checksummed output on every run.

Every Field, Defined.

All 46 columns of the delivery, grouped by the four structured layers, plus the generated intelligence the JSON carries on every record. Every product folder ships its own DATA_DICTIONARY.md and the full contract in SCHEMA.md.

Identity and Location Where the obligation lives and how to find it again.

obligation_id	string	Unique, stable id. Encodes framework, section, marker path, page, sequence.
framework	string	Source instrument: MAS TRM / MAS Cyber Hygiene / MAS Outsourcing Notice.
regulator	string	Issuing authority (Monetary Authority of Singapore).
document_title	string	Full title of the source regulatory instrument.
document_version	string	Version / issuance year of the source instrument.
jurisdiction	string	Legal jurisdiction (Singapore).
section_number	string	Human-readable section/paragraph locator with page.
section_title	string	Title of the top-level section.
subsection_title	string	Title of the immediate subsection heading.
marker_path	string	Section-relative structural marker chain (the obligation's address).
section_breadcrumb	string	Human-readable ancestor path, e.g. '9 Access Control > 9.1 User Access Management > 9.1.1'.
depth	integer	Depth in the document hierarchy.
parent_obligation_id	string	Id of the structural parent, if any.
page	integer	Printed page the obligation begins on.

Layer 1 • Verbatim The legally authoritative text. This is what you cite.

verbatim_text	string	Byte-exact quotation from the source. Legally authoritative; cite this.
---------------	--------	---

Layer 2 • Normalized The same text prepared for search, diffing, and retrieval.

normalized_text	string	Whitespace/typography-normalized text for search and diffing.
obligation_paragraph	string	The obligation's own source paragraph (the verbatim-bearing segment of context_window). Use this when you want exactly the paragraph the duty lives in.
context_window	string	A newline-joined reading window: the obligation's own paragraph plus the immediately adjacent parsed item on each side within the same subsection. The segment containing verbatim_text is the obligation's own paragraph (also shipped as obligation_paragraph); the outer segments equal preceding_paragraph / following_paragraph.
preceding_paragraph	string	The parsed item immediately before the obligation's own paragraph, within the same section. When it falls in the same subsection it is also the first segment of context_window; at a subsection boundary it reaches further back than the window.
following_paragraph	string	The parsed item immediately after the obligation's own paragraph, within the same section. When it falls in the same subsection it is also the last segment of context_window; at a subsection boundary it reaches further ahead than the window.
governing_stem	string	The lead-in stem that governs this obligation when it is a lettered sub-item or sits under a lead-in (e.g. 'A bank in Singapore must'). Empty for standalone paragraphs.

Layer 3 • Parsed Duty The obligation decomposed into who must do what, and how hard.

modal	string	The duty-bearing verb, normalized from the source (MUST / MUST_NOT / SHOULD / MAY / MAY_NOT / PROHIBITED / REQUIRED_TO / EXPECTED_TO).
modal_strength	string	mandatory / recommended / permitted / prohibited.
deontic_category	string	

		obligation / recommendation / permission / prohibition.
is_negative	boolean	True for a prohibition (shall not).
obligation_type	string	Coarse functional type of the duty. For finer segmentation use tags and deontic_category.
actor	string	The party who must act.
action_required	string	The action the actor must take.
object_of_obligation	string	The thing the action operates on.
covered_person_scope	string	Entities in scope of the obligation.
frequency	string	How often the duty applies, as stated or implied by the source. Free text, most commonly 'ongoing'; explicit cadences (e.g. 'At least annually') where the instrument states them.
obligation_summary	string	One-line plain-language gist. Never replaces the verbatim.
deadline	string	Explicit deadline if the text states one (often null).
is_currently_in_force	boolean	Whether the obligation is currently in force.

Layer 4 · Consequence and Context What breach costs, and how urgent the duty is.

consequence	string	Stated or inferred consequence of non-compliance.
sanction_type	string	supervisory / civil / criminal.
severity_hint	string	low / medium / high.
is_criminal_offense	boolean	True if breach is a criminal offense.
priority	string	Suggested review priority (low / medium / high).
tags	string	Topical tags for faceted filtering.

Citation and Provenance The audit trail: cite it, trace it, prove it.

legal_citation	string	Short drop-in legal citation string.
legal_citation_full	string	Full legal citation with the complete document title.
source_pdf_url	string	URL of the source regulatory PDF.
extraction_timestamp	string	UTC timestamp of the release run this record was produced in; records built together share it.
evidence_img_loc	string	Relative path to the source-verified evidence capture PNG. Present only in the Singapore Regulatory Collection and the Free Sample; products that do not ship captures omit this field entirely.
ocr_confidence	number	Text-fidelity score for the source page the duty was read from (0-1).

ProfytAI Regulatory Intelligence Why the duty exists and how teams meet it. Nested under context.ai_context; JSON delivery only.

summary	string	The duty restated in plain language, for a reader who has not read the instrument.
why_this_obligation_exists	string	The supervisory concern the duty addresses.
relationship_to_parent	string	How this duty sits under its governing paragraph or section.
obligation_kind	string	requirement / condition / exception / prohibition / implementation_step.
interpretation_notes	string	Where the wording is load-bearing, and what it does and does not reach.
implementation_considerations	string	How institutions typically satisfy the duty in practice.

THE LOAD-BEARING FIELDS, AND WHY THEY MATTER

<code>obligation_id</code>	A stable join key that survives reissuance, so a duty is trackable across every version of the rulebook.
<code>verbatim_text</code>	The exact words of the regulator, zero paraphrase. The text an examiner accepts, and the root every other field is derived from.
<code>legal_citation</code>	A drop-in citation. Paste it into a register, policy, or audit finding without re-keying.
<code>deontic_category</code>	Separates a hard requirement from guidance and a ban, so enforcement-risk duties are prioritised first.
<code>actor / action_required / object_of_obligation</code>	The duty pre-decomposed into who acts, what they do, and on what, ready to map onto a control owner.
<code>evidence_img_loc</code>	A source-verified page image an auditor accepts as proof. Collection and Free Sample only.

How the Data Connects Back to the Regulation.

Every obligation is anchored to its source in four independent ways, so any number your team publishes reconciles to the regulator's own document in seconds.

The Identifier Is an Address

Each `obligation_id` is stable and human-readable. It encodes the regulator, instrument, issuance year, the section or clause marker, the printed page, and the obligation's sequence within that clause. It survives reissuance, so the same duty is trackable across every version of the rulebook.

MAS	TRM	2021	Sec7.2.1	p23	OBL1
REGULATOR	INSTRUMENT	ISSUED	SECTION / CLAUSE	PRINTED PAGE	OBLIGATION INDEX

Citations

Two citation fields ship on every record: a short drop-in string (`legal_citation`) for registers and workpapers, and a full form (`legal_citation_full`) with the complete document title. Both name the paragraph and the printed page, so a reviewer opens the source at the exact line.

Source Traceability

Every record links to the official source PDF (`source_pdf_url`) and the page it begins on. In the Collection, `evidence_img_loc` adds a rendered image of that page with the duty highlighted and a verification footer, ready to hand to an auditor.

Hierarchy and References

Structural fields (`section_breadcrumb` , `marker_path` , `depth` , `parent_obligation_id`) place each duty in the document's tree, so a sub-clause carries its parent context and a register can be grouped exactly as the instrument is organised.

Verbatim to Derived

The verbatim text is the root of every record. The normalized text, the parsed duty (actor, action, object), the summary, and the tags are all derived from it and never replace it, so the authoritative wording is always one field away.

See the Data, Not a Description.

One representative record, exactly as delivered. Every record carries the full 46-field structured schema documented in the Field Dictionary; the free 30-obligation sample at profyt.ai lets your team inspect all of it.

SHOWN BELOW · 14 OF THE 46 FIELDS, SELECTED FOR READABILITY. SEE THE FIELD DICTIONARY FOR THE COMPLETE RECORD.

GUIDANCE
MAS TRM
MAS.TRM.2021.Sec7.2.1.p23.OBL1

VERBATIM SHOULD

The FI should implement a configuration management process to maintain accurate information of its hardware and software to have visibility and effective control of its IT systems.

MAS TRM Guidelines, Section 7.2.1, p. 23 (2021)

ACTOR	FI
ACTION	implement a configuration management process to maintain accurate hardware and software information
OBJECT	hardware and software configuration information

IT Service Management · 7.2 Configuration Management · p.23

MEDIUM SEVERITY

Non-adherence may attract MAS supervisory action; the TRM Guidelines set out the standards MAS expects financial institutions to meet.

WHAT YOU ARE LOOKING AT

Verbatim	The regulator's exact words. Legally authoritative; this is the text you cite and the text your policy must answer.
Legal Citation	A drop-in citation, paragraph and page. Paste it into a register or an audit finding with no re-keying.
Parsed Duty	The clause pre-decomposed into who must act, what they must do, and on what, ready for a control owner.
Identifier	A stable address that reconciles this record to the source document and survives reissuance.

The Record Makes the Claim. The Capture Proves It.

A structured record is a claim; the evidence capture is the proof. Each Collection obligation ships a rendered image of its source page, the duty highlighted, stamped with the obligation ID, document, page, capture time, and source URL.

"security standards", in relation to a system, means a set of configurations for the purpose of safeguarding and improving the security of the system;

"system", in relation to a relevant entity, means any hardware or software that is used by the relevant entity;

"vulnerability", in relation to a system, means any weakness, susceptibility or flaw of the system that can be exploited, including but not limited to by allowing an unauthorised person to access the system, or to compromise the security configuration settings of the system.

Except where defined in this Notice or if the context otherwise requires, the expressions used in this Notice have the same meanings as in the Act.

III. Application of Notice

3.1 A relevant entity need not comply with a requirement in this Notice to the extent that it is unable to exercise control over a system to ensure compliance with that requirement, in all of the following ways:

- (a) the relevant entity cannot exercise direct control over the system to ensure compliance with that requirement;
- (b) a relevant entity cannot exercise indirect control over the system by requiring the system provider to ensure compliance with that requirement;
- (c) it is not reasonable for the relevant entity to procure an alternative system provider over whom the relevant entity is able to exercise such indirect control referred to in sub-paragraph (b), to provide the system.

IV. Cyber Hygiene Practices

4.1 **Administrative Accounts:** A relevant entity must ensure that every administrative account in respect of any operating system, database, application, security appliance or network device, is secured to prevent any unauthorised access to or use of such account.

4.2 **Security Patches:**

- (a) A relevant entity must ensure that security patches are applied to address vulnerabilities to every system, and apply such security patches within a timeframe that is commensurate with the risks posed by each vulnerability.

CYBER HYGIENE · BINDING

MAS.CH.2024.Sec4.1.p3.OBL1
 Notice FSM-N06, paragraph 4.1,
 p. 3

Collection exclusive. All 393 captures ship with the Singapore Regulatory Collection. The free 30-obligation sample includes all of its captures so your team can inspect the chain before any commitment.

REGULATORY EVIDENCE RECORD		profytai	
Obligation ID	MAS.CH.2024.Sec.1.p3.OBL1	Jurisdiction	Singapore
Framework	MAS Cyber Hygiene	Regulator	Monetary Authority of Singapore (MAS)
Document	Notice FSM-N06 on Cyber Hygiene [2024]	Page	3
Section	4.1 p3 - Cyber Security Requirements	Evidence Captured	2026-07-05T05:03:44Z
Source URL	https://www.mas.gov.sg/-/media/mas-media-library/regulation/notices/tps/notice_fsm-n06/mas-notice_fsm-n06.pdf		

Source Text Verified - Regulatory obligation located and confirmed in the cited document

How Teams Put the Data to Work.

The dataset is the structured foundation under the compliance workflows a bank runs every quarter. The same records serve each of them.

Obligation Registers

Stand up a complete, cited MAS obligation register on day one instead of reading the instruments and re-keying each duty by hand.

Gap Analysis

Map each obligation to an internal policy or control and evidence the gaps, with the source wording attached to every finding.

Audit and Examination Readiness

Walk into a MAS inspection with every requirement already located, quoted, and page-anchored to the source document.

Policy Mapping

Trace each internal policy statement back to the specific regulatory duty it satisfies, and surface duties no policy answers yet.

Regulatory Change Management

Diff a reissued instrument against the prior version using normalized text and stable identifiers to see exactly what moved.

Evidence Collection

Attach a source-verified page capture to each obligation in your control library so provenance is one click from any record.

Internal Controls

Assign a control owner to each parsed duty (actor, action, object) and monitor coverage across the full obligation set.

Risk Assessment

Weight and prioritise remediation using the modal force, severity, and sanction type carried on every record.

AI and RAG Grounding

Ground a compliance copilot on cited, byte-exact regulatory text your model can quote with a citation, not paraphrase.

Built to Survive Audit Review.

Regulatory data earns its keep in front of an examiner. Every release is engineered so its provenance holds up to that scrutiny.

Extraction	Each instrument is parsed from its official PDF into discrete obligations. Extraction is schema-grounded and reviewed, not a bulk text dump: one record per duty, at the clause level a compliance officer would cite.
Byte-Exact Verbatim	Every record carries the source duty quoted exactly. Nothing is paraphrased in the verbatim field; the derived layers never overwrite it.
Normalization	A parallel normalized field cleans whitespace and typography for search, diffing, and retrieval, while the verbatim stays untouched for citation.
Validation Gates	The build rejects any record missing its verbatim text or its legal citation, and asserts every identifier is unique. Traceability is enforced at build time, not asserted after the fact.
Deterministic Build	The same inputs rebuild byte-identical outputs. Every release is reproducible, and a content hash (SHA-256) travels in the JSON metadata so a pipeline can verify the payload it ingested.
Versioning and Errata	Releases are dated and carry a change log. Errata are free: a confirmed extraction error is corrected and reissued at no charge to licensees of that version. New versions are separate: when MAS reissues an instrument, a new dated version is released under a new license, or delivered by the optional annual update subscription. A one-time license covers the version purchased.
Known Scope	Coverage is point-in-time: the exact text of the instruments as issued, which is precisely what a defensible assessment cites. A small number of derived fields (for example, explicit deadlines) are populated only where the source states them.

Terms of Use.

Commercial License	A commercial license for the single legal entity named on the order to use the dataset internally. Affiliates and subsidiaries need their own license or a multi-entity agreement. Full terms travel with every delivery as LICENSE.md.
Permitted Use	Internal compliance, risk, audit, and product use: obligation registers, gap analysis, policy mapping, control libraries, reporting, and AI grounding within the licensed entity.
Redistribution	The dataset may not be resold, sublicensed, or redistributed as data. Redistribution, OEM embedding, and multi-entity terms are available as a separate private offer.
Updates	Each release is a point-in-time, dated version. An optional annual update subscription delivers each new version on reissuance; a confirmed extraction error is corrected and reissued free to licensees of that version.
Support and Assurance	A signed Data Processing Agreement and security documentation are available on request. Licensing and support: support@profyt.ai.

Pricing and procurement channels are published at profyt.ai/data/mas; a ProfytAI listing on AWS Data Exchange is coming. This dossier is commercial collateral and is not itself a license.

Definitions and Reference.

GLOSSARY

Obligation	A single, discrete regulatory duty extracted at the clause level: one requirement, one record.
Deontic category	The logical force of a duty: obligation (must), recommendation (should), permission (may), or prohibition (must not).
Modal / modal strength	The RFC-2119-style verb (SHALL, MUST, SHOULD, MAY) and its mapped strength (mandatory, recommended, permitted, prohibited).
Actor / action / object	The parsed duty: the party who must act, the action they must take, and the thing it operates on.
Verbatim	A byte-exact quotation of the source text. Legally authoritative; the field you cite.
Evidence capture	A rendered image of the source page with the obligation highlighted and a verification footer (ID, document, page, capture time, source URL).
Point-in-time	The dataset reflects the instruments exactly as issued on the release date; updates ship as new dated versions.

ABBREVIATIONS

MAS	Monetary Authority of Singapore, the issuing authority and financial regulator.
TRM	Technology Risk Management Guidelines (MAS, 2021).
FI	Financial Institution, the regulated entity to which the duties apply.
OBL	Obligation, the atomic unit of the dataset (one duty per record).
CFT / AML	Countering the Financing of Terrorism / Anti-Money Laundering.
DPA	Data Processing Agreement, available on request for procurement.

REFERENCE

The complete, versioned field contract is SCHEMA.md; the generated field reference is DATA_DICTIONARY.md, shipped in every product folder. Source instruments ship as PDFs in each delivery's source_pdfs/ folder. Methodology and QA notes travel in the Excel workbook and the README of every folder.



ABOUT PROFYTAI

ProfytAI builds structured regulatory intelligence for financial institutions. We compile a regulator's source documents into examination-ready obligation data: byte-exact quotations, legal citations, parsed duties, and document context on every record, with source-verified evidence captures where they ship. Singapore Regulatory Intelligence is live today across MAS TRM, Cyber Hygiene, and Outsourcing; further jurisdictions are on the roadmap.

Published by **ProfytAI Pte. Ltd.** · support@profyt.ai

© 2026 ProfytAI · Singapore Regulatory Collection · PROFYTAI-SG-MAS-COLLECTION-1.0.0
(2026-07-07)

profyt.ai