ObligaI

COMPANY OVERVIEW

May 2026

High-Level Company Overview

COMPANY OVERVIEW

What ObligaI does

ObligaI translates the language of prudential regulation into a proposed structured set of reporting obligations. For instance, for the LCR Delegated Act, ObligaI generates a proposal of reporting obligations, with a direct link back to the article that prescribed it. The output is an auditable specification, not a statistical approximation.

The challenge

The current industry challenge is that the translation of regulatory requirements into reporting obligations — the concrete template cells — is done manually, a process that is time-consuming, error-prone, and vulnerable to contrarian interpretation. Expert teams spend a substantial amount of time reading regulations and hand-mapping each requirement to a row and column in an EBA workbook. When regulations change, the mapping is repeated. Traceability from cell back to article is fragile, living in separate documents and tacit knowledge.

How ObligaI works

The underlying discovery is that the entire vocabulary of prudential reporting reduces to ten irreducible conceptual atoms, called generators: Asset, Liability, Counterparty, Currency, Market, TimeBucket, Encumbrance, CashFlow, Facility, and Guarantee. Every regulatory requirement can be expressed as a combination of these generators plus certain required properties, called edges (e.g., an Asset HAS_MARKET_VALUE). A one-time expert annotation records, for each regulatory article, which generators and edges are mandated. A deterministic,

rule-based engine then expands this annotation into the complete set of reporting obligations, filling in all currency columns, time buckets, and institution-specific parameters.

Key technology attributes

Validated and tested	The LCR template was mapped manually by the author, generated by the pipeline, and compared field by field. The outputs matched exactly. The same approach has been applied to NSFR and ALMM templates.
Zero hallucination	The generator is deterministic; the same input always produces the same output. No field can appear without an explicit annotation.
Full traceability	Every obligation carries the exact article reference that required it. The entire derivation chain—requirement to generator to edge to obligation—is machine-auditable.
Multi-regulation, single ontology	The same ten generators serve at the moment three templates: LCR, NSFR, and ALMM. The objective is to test the ten generators on all liquidity regulatory reporting templates. Adding a new regulation means adding new edges, not a new vocabulary.

Traction and validation

7,360
LCR obligations — validated and tested

4,461
NSFR obligations — validated and tested

21,635
ALMM obligations — validated and tested

Obligai's method has been tested extensively, and its generated LCR, NSFR, and ALMM templates are continuously monitored on its completeness and accuracy.

Company details

The company was founded by a regulatory technology practitioner with over a decade of experience in liquidity-reporting architecture and validation inside a major European SSM bank. That background provided the domain expertise necessary to identify the ten generators and to construct the generative grammar.

Focus areas

The immediate focus is on liquidity regulatory reporting requirements: LCR, NSFR, and ALMM — for EU banks. The roadmap includes expansion into other risk areas (credit, operational, ..), and capital adequacy. The generator set is assumed to be domain-invariant; each new dialect principally requires the addition of edges and profile parameters. The objective is to test this hypothesis across all liquidity regulatory reporting templates.

Next steps

To arrange a demonstration, contact info@oblig.ai.