



TECHNOLOGY WHITEPAPER

SoReady™

Governed Multi-AI Architecture

MCP, event context, specialised agents, sovereignty and local / on-premise engines.

THESIS

AI researches, extracts, explains and proposes. Rules compute. The human validates. SoReady traces - whatever the model or the place of execution.

Several AIs, one governance

Performance does not come from the number of models but from explicit roles, controlled context and a strict separation between assistance and authority.

AI

Research and proposal

Extraction, synthesis, contradiction, simulation and user assistance.

RULES

Deterministic computation

Contention, readiness, correlations, versioned thresholds and controls.

HUMAN

Validation and arbitration

Responsibility, authorisation, publication and sensitive action.

SYSTEM

Traceability

Sources, versions, rights, decisions and rollback.

LEVEL 2 | UNDERSTAND

A multi-AI architecture separates capabilities instead of asking a single model to do everything. One agent extracts, another researches, a third challenges, while deterministic engines compute states. This specialisation improves controllability, but increases the need for governance. SoReady must know which model received which context, with which tools, and to produce which proposal. The human remains the authority for publication and action.

LEVEL 3 | GO DEEPER

Define a contract per role: authorised inputs, expected outputs, tests, confidence threshold, accessible tools and escalation procedure.

A council of roles, not an all-powerful AI

Models are replaceable and assigned to bounded tasks. One AI can challenge another without creating an autonomous authority.

01

Extract

OCR, tables, requirements, entities.

02

Research

Sources, versions, facts and evidence.

03

Challenge

Hypotheses, gaps, inconsistencies.

04

Simulate

Impacts, options and scenarios.

05

Synthesise

Briefs, reports and expected decisions.

06

Assist

Integrated contextual help, per role.

LEVEL 2 | UNDERSTAND


AI roles map to observable tasks. The extractor processes documents; the researcher gathers sources; the challenger looks for inconsistencies; the simulator formulates scenarios; the synthesiser prepares a brief. These roles can use different models and be replaced. Their cooperation is only useful if disagreements stay visible and if an answer never becomes a business truth by simple consensus between models.

LEVEL 3 | GO DEEPER

The number of agents is not a maturity indicator. Measure accuracy, coverage, citations, stability, cost, latency and the rate of human escalation.

Learning the event without a black box

The specific context is a sourced, versioned and reversible memory. Opaque retraining is not the default mechanism.



GOVERNANCE POINT

The context supplied to each model is minimal, authorised and matched to the task. Sensitive facts can stay on site.

LEVEL 2 | UNDERSTAND

Event context is built from versioned documents, reference data and decisions. SoReady favours a recoverable, sourced memory over opaque retraining. When a regulation changes, the system can identify the new source and the affected objects; it does not have to guess what the model memorised. The context passed to each AI stays limited to its task, its rights and the relevant period.

LEVEL 3 | GO DEEPER

Separate RAW, EVIDENCE and RULES. Expose to the AIs mainly the validated rules and the necessary evidence, with a recovery recipe and freshness indicator.

Three levels up to sovereignty

The same quality, audit and decision contracts apply whatever the place of execution.

LEVEL 1

Controlled external AI services

Gateway, minimisation, authorised data and specialised models.

LEVEL 2

Hybrid routing by sensitivity

Private cloud, dedicated services and local models depending on the data class.

LEVEL 3

Sovereign / on-premise engines

Local or isolated execution, offline continuity, controlled administration and context.

LEVEL 2 | UNDERSTAND


Level 1 relies on controlled external services for tasks compatible with their data policy. Level 2 routes requests by sensitivity to private or local environments. Level 3 keeps models, context and execution on sovereign or isolated infrastructure. This gradation reconciles innovation, cost, confidentiality and continuity, without changing the doctrine of evidence and authority.

LEVEL 3 | GO DEEPER

Classify before routing. Test loss of connectivity, provider unavailability, model change and local restoration before games time.

MCP is a governed access layer, not the business logic

MCP exposes explicitly authorised resources, tools and workflows. It replaces neither the Business Core, nor the rules, nor the access controls.



GOVERNANCE POINT

No direct access from AI or MCP to operational databases. COMMIT-type actions remain separate, authorised and validated.

LEVEL 2 | UNDERSTAND

MCP provides a common language for presenting resources, tools and workflows to models. It carries neither the business rule nor authority over the data. A single gateway controls identity, consent, rights, logging and usage limits. The READ, ANALYSE and DRAFT tools can be opened progressively; COMMIT remains a distinct capability, subject to Business Core controls and explicit validation.

LEVEL 3 | GO DEEPER

No direct MCP-to-PostgreSQL access. All operations go through validated, idempotent, auditable business services protected by the same rules as the API.

The event supplies facts; AI supplies meaning

The EDF qualifies and correlates events; the EOS builds an operational state; AI explains the signals and prepares the options.



GOVERNANCE POINT

AI does not trigger a critical alert on its own and does not action OT or life-safety systems.

LEVEL 2 | UNDERSTAND

The EDF admits, qualifies and distributes events; deterministic rules build the correlations; the EOS maintains the operational state. AI then steps in to explain the situation, summarise the evidence, formulate hypotheses and prepare options. This sequence keeps a hallucination from becoming an operational signal. It also allows a model to be replaced without changing the flow contracts, the detection logic or the LOC's responsibilities.

LEVEL 3 | GO DEEPER

Keep the source identifiers, the rule version and the AI context of the brief. The alert stays deterministic; the model does not access the broker or the business core directly.

Designing the EDF as a trust boundary

The fabric separates hostile input, accepted events and business state. This separation protects the EOS and allows sources to evolve without weakening the core.

CONTROLLED INPUT

Authorise a source, not just a protocol

Per-flow identity, rights, load limits, expected formats, independent shutdown and no direct write to the core.

PROBATIVE VALIDATION

Keep, verify or quarantine

Raw evidence or reference, syntactic and semantic control, provenance, timestamping and explicit status.

DURABLE EXCHANGE

Absorb failures and allow replay

Deduplication, ordering, recovery, independent consumers, store-and-forward and reconciliation of gaps.

CANONICAL EVENT

Correlate without losing authority

Common contract, source and sensitivity preserved, versioned rules; AI steps in after qualification.

LEVEL 2 | UNDERSTAND

The fabric receives potentially hostile input without confusing it with accepted facts. The admission boundary controls identity, rights, size, rate, format and semantic consistency. Raw evidence or a reference is kept where policy allows; rejects become an observable quarantine. Durable exchange absorbs failures, deduplicates and replays. Finally, the canonical contract carries source, authority, sensitivity and timing to the EOS and the correlation rules.

LEVEL 3 | GO DEEPER

No external flow writes directly into the core. Test malformed messages, duplicates, delay, clock drift, WAN outage, a compromised flow, recovery and timeline reconstruction.

Cross-checking AIs to make uncertainty visible

Confrontation is useful when roles, criteria and sources differ. A disagreement becomes governable information.

PROPOSAL

Primary agent

Formulates a synthesis and cites its sources.

CHALLENGE

Challenger agent

Looks for omissions, conflicts and assumptions.

CONTROL

Rules engine

Checks criteria, rights and computed states.

ARBITRATION

Human owner

Accepts, corrects, refuses or requests evidence.

LEVEL 2 | UNDERSTAND

Cross-checking is not meant to artificially manufacture consensus. A primary agent puts forward an answer; a challenger tests the assumptions and looks for missing sources; the rules engine verifies the formal constraints. The human owner then has a synthesis that also shows the disagreements. This method is particularly useful when documents are incomplete or when several interpretations produce different impacts.

LEVEL 3 | GO DEEPER

Define independent contradiction criteria. A second model using the same context and the same biases is not sufficient assurance.

A wider AI surface demands stronger governance

Security of the data, the tools, the models and the decisions is handled together.

Risk	Expected control	Authority
Document injection	Isolation, provenance, filtering, validation	Data / security
Data leakage	Minimisation, classification, routing	DPO / CISO
Hallucination	Sources, uncertainty, rules, challenge	Business owner
Unauthorised action	READ / DRAFT / COMMIT separation	Administration
Model drift	Evaluation, versioning, rollback	AI governance
External unavailability	Hybrid / local fallback	Operations

LEVEL 2 | UNDERSTAND

AI widens the classic risks: injection via a document, exfiltration through a tool, unauthorised action, dependency on a provider or drift in a model. Security must cover the full chain, from ingested content to the proposed action. Provenance, filtering, minimisation, READ/DRAFT/COMMIT separation and anomaly monitoring complement one another. The business owner stays responsible for meaning; the CISO and operations control the platform.

LEVEL 3 | GO DEEPER

Test attacks on documents, prompts, tools and APIs. Log without exposing sensitive data; plan for revocation, a kill switch and an incident procedure.

Local deployment also protects continuity

OT, security and life-safety networks stay isolated. A local, observational gateway can provide store-and-forward and controlled recovery.

SITE ZONE

OT and safety systems

Local authorities, segmentation, no implicit action from AI.

LOCAL GATEWAY

Governed reading and buffer

Filtering, timing, degraded mode and controlled synchronisation.

SOREADY CORE

EOS / EDF / context / local AI

Operational state, correlation, assistance and audit.

EXTERNAL SERVICES

Optional, per policy

Only explicitly authorised data and tasks.

LEVEL 2 | UNDERSTAND

Local deployment does not only serve confidentiality. It allows a critical context to keep running when external connectivity degrades. OT and safety systems stay in their zones; an observational gateway filters and times the useful data. The local core can maintain EOS, EDF and AI assistance within a controlled perimeter, then synchronise when conditions allow.

LEVEL 3 | GO DEEPER

Store-and-forward must preserve order, timestamping, deduplication and provenance. OT actioning uses a separate, governed channel.

AI runs across modules without blurring them

Every AI capability is anchored in a business object, a source and an expected decision.

Module	AI contribution	Safeguard
Document management / OCR	Extract and match	Validation and provenance
Reference data / external flows	Qualify sources and mappings	Authority and reconciliation
Planning / contention	Explain collisions and options	Deterministic computation
Infrastructure / FoP	Read capacities and configurations	Business engines
Risk Assessment	Summarise exposure and treatments	Risk owner
Budget / BOM / VIK	Detect gaps and scenarios	Authoritative financial source
EOS / EDF	Explain correlations	Governed rules and thresholds
Reporting	Draft reports and briefs	Citations and approval
User assistance	Contextual help by role	Rights and scope
Administration / security	Assist control and diagnosis	Separation of duties

LEVEL 2 | UNDERSTAND

Every module keeps its own business logic. AI can read a regulation, qualify an external flow, explain a planning collision, detect a BOM gap or draft a report, but it does not perform the calculations in place of the relevant engines. The shared context connects budget, VIK, workforce, infrastructure, FoP, risk and live signals. User assistance, administration and security stay subject to rights. The value comes from this controlled cross-cutting reach, not from a merging of responsibilities.

LEVEL 3 | GO DEEPER

For every use, document the authoritative source, the deterministic calculation, the AI role, the expected validation, the sensitive data and the separation of duties.

An aligned architecture, without over-promising

The frameworks structure the management system and the controls. They amount to neither product certification nor automatic compliance.

Framework	Question covered	Application
ISO/IEC 42001:2023	How to govern AI?	Roles, risks, monitoring, improvement
NIST AI RMF 1.0	How to map and treat risks?	Govern, Map, Measure, Manage
ISO/IEC 27001:2022	How to govern IT security?	Risks, controls, audit
OWASP GenAI / API	How to test application threats?	Injection, access, data, tools
MCP 2026-07-28	How to interoperate in a governed way?	Authorisation, tools, resources
W3C PROV-O	How to preserve provenance?	Sources, transformations, accountability

LEVEL 2 | UNDERSTAND

ISO/IEC 42001 structures AI management system governance; the NIST AI RMF offers a risk framework; ISO/IEC 27001 addresses security management; OWASP targets application threats. MCP and W3C PROV-O complete interoperability and provenance. SoReady uses these frameworks to design and challenge its controls. The alignment amounts to neither a certification nor a guarantee that a particular use is compliant without context analysis.

LEVEL 3 | GO DEEPER

Maintain a control register linking requirement, risk, measure, test, result, owner and exception. Review the references at every major version.

Moving from pilot to sovereign operations

The trajectory is progressive: scope, quality measures, rights, exercises, then extension.



GOVERNANCE POINT

A move to COMMIT or to actioning requires a separate architecture and governance decision.

LEVEL 2 | UNDERSTAND

The pilot starts with uses that take no action: research, analysis and drafting. Results are measured against a representative corpus, then tested against errors, failures and attacks. Extension to hybrid or local deployment happens once the quality and operating contracts are established. A COMMIT capability is only opened after a separate decision, idempotence tests, adapted rights and verifiable human validation.

LEVEL 3 | GO DEEPER

The move to production requires versioning, recurring evaluation, rollback, observability, secrets management and a model retirement procedure.

Three lines of control

AI quality is not a single score. It combines performance, security, robustness, explainability and business value.

LINE 1**Product and business teams**

Functional tests, sources, rights, quality of answers.

LINE 2**AI security and governance**

Threats, models, data, access, incidents and exceptions.

LINE 3**Audit and leadership**

Effectiveness of controls, accountability and residual exposure.

EVIDENCE**Decision log**

Model, context, sources, rules, validation and version.

LEVEL 2 | UNDERSTAND

The first line belongs to the product and business teams, who verify sources, accuracy and usefulness. The second brings together security, data and AI governance; it controls risks, models, access and incidents. The third gives audit and leadership an independent view of the programme's effectiveness. The decision log connects these three levels and helps explain accepted exceptions.

LEVEL 3 | GO DEEPER

Avoid self-certification by the team that builds. Define revalidation frequency, thresholds and owners of corrective plans.

Technological openness, stable accountability

SoReady can change models and deployment modes without changing the decision doctrine.

MODEL

Interchangeable

A provider is not the architecture.

CONTEXT

Controlled

Version, authority, sensitivity and purpose.

ACTION

Governed

Rights, validation and separation of duties.

EVIDENCE

Persistent

Sources, calculations, proposals and decisions.

LEVEL 2 | UNDERSTAND

A model, a provider or a protocol can change without shifting accountability. SoReady keeps stable role contracts and business services; the context stays classified and the evidence persists. This neutrality reduces technology lock-in and allows the best engine to be chosen locally. It does, however, impose an evaluation discipline: two interchangeable models are never assumed equivalent without a test.

LEVEL 3 | GO DEEPER

Version the model, parameters, context, tools and routing policy. A substitution triggers a non-regression campaign.

AI that is useful because it knows its limits

SoReady's multi-AI increases analytical capacity without diluting authority. MCP opens access; the Business Core governs; the EOS/EDF observes and correlates; the human decides.

EXPLAIN

The situation

With sources, timing and uncertainty.

CHALLENGE

The assumptions

With several roles and criteria.

PROPOSE

Options

With impacts, risks and evidence.

KNOW ITS PLACE

Without autonomous action

On sensitive and operational decisions.

LEVEL 2 | UNDERSTAND

The architecture reaches its goal when AI improves understanding without masking uncertainty. It explains the facts drawn from the EOS/EDF, challenges assumptions and prepares options, but leaves the rules to compute and the authorised people to decide. This limitation is not a weakness: it makes AI acceptable in an environment where the operational, financial or security consequences are real.

LEVEL 3 | GO DEEPER

The interface must show sources, date, confidence level, assumptions and expected action. An unsourced answer must never look certified.

The frameworks that bound the architecture

These frameworks are used as design and assurance grids. Citing them does not mean the product or its operator is certified.

Public reference	Structuring question	SoReady application
ISO/IEC 42001:2023	Governing an AI system	Roles, risks, monitoring
NIST AI RMF 1.0	Treating AI risks	Govern, Map, Measure, Manage
ISO/IEC 27001:2022	Protecting information	Classification, access, audit
OWASP GenAI / API	Testing the application surface	Injection, tools, secrets
MCP 2026-07-28	Exposing resources and tools	Gateway, consent, rights
W3C PROV-O	Preserving the origin of a fact	Sources and transformations

LEVEL 2 | UNDERSTAND

The references bring together AI management, risk, security, application threats, interoperability and provenance. They form a coherent reading grid for auditing the design. None replaces the event's business analysis or the applicable legal obligations. Controls must be translated into testable requirements within the architecture and operations.

LEVEL 3 | GO DEEPER

Use the official editions and keep a compliance register separating applicable, implemented, tested, accepted and out of scope.