

MIZ OKI 3.5™

Operating Knowledge Intelligence for the Autonomous Enterprise

Master Positioning Document — Revision 3.6

Revision 3.6 · 2 September 2026

Folds the r3.5.2 amendment of record (owner ruling 2026-08-19-B) into one document

Every capability carries a status label — [LIVE], [PARTIAL], [IN BUILD], or [PROPOSED]

Where this document and the live Cell Registry disagree, the registry wins

Supersedes r3.5.1 (18 August 2026) + amendment r3.5.2 (19 August 2026)

Operating Knowledge Intelligence for the Autonomous Enterprise

Master Positioning Document — Revision 3.6

Updated: September 2, 2026 · Supersedes: r3.5.1 (August 18, 2026) + amendment r3.5.2 (August 19, 2026)

Revision note: this update incorporates the ORACLE latent-intent layer (Cells 33–36), the pre-conversion perception build, the Causal Credit Ledger and measurement rails, Net Contribution Yield, the Signal Intelligence public surface at mizoki3.com/signal, the O-1 prohibited-signals contract, and the hardened claim-discipline canon. Every capability in this document carries a status label — [LIVE], [PARTIAL], [IN BUILD], or [PROPOSED] — per Section 16. The live Cell Registry is ground truth; where this document and the registry disagree, the registry wins. Revision note (r3.6, September 2, 2026): this revision folds the formal amendment of record r3.5.2 (owner ruling 2026-08-19-B — F1/F2/F4/F5 advanced from PROPOSED to IN-SCOPE BUILD; Roadmap Phases 8 and 9 collapsed into one build program) into the body, so one document holds the ruling of record. Amendment 1 is applied in place in §12 (Phase 8 — Full Growth Control Build; former Phase 9 dissolved); Amendment 3 is applied in place in §16; Amendment 2 amends a different base document (docs/MIZOKI_SIGNAL_GROWTH_CONTROL_UNIFIED_SYSTEM_r2.0.md §5/§6) and is therefore carried in Appendix A, not folded into this body. The RULING and the CONSISTENCY NOTES are carried verbatim in Appendix A. Every status label stands as ruled — this revision upgrades nothing. The body is otherwise the Drive-verbatim r3.5.1 text (body sha256 82b97f9b2669f32a7f6a820e26e140dd7d176dc14dd48afb24a270475a261634), preserved byte-for-byte at docs/whitepapers/archive/MIZOKI_3.5_WHITEPAPER_r3.5.1_AUG2026.md.

Known canon conflicts, carried forward unedited (logged 2026-08-19, rule 03: tree wins, conflict logged not edited; the live Cell Registry docs/architecture/CELL_REGISTRY.md is ground truth over this document):

- §5.7 uses the plan-vintage intent-cell role map — canon: Cell 34 = intent-scoring-api (scoring), Cell 35 = intent-graph (Firestore-journaled in-memory graph; Neo4j retired 2026-08-09; no PRECEDES edge type in the landed store), Cell 36 = intent-causal (incrementality/causal credit); Cell 28 = Sales Revenue Pipeline (legacy w.r.t. intent).
- §7.2's Neo4j/TigerGraph mentions are deployment-options prose, not the shipped stack (Firestore-backed KG; Neo4j retired 2026-08-09; TigerGraph never).
- §9.2's erasure "vectors" limb is not-applicable in cells 33–36 (no vector index exists — docs/lii/RUNBOOK.md §6).

EXECUTIVE SUMMARY

MIZ OKI 3.5™ is the Operating Knowledge Intelligence platform for organizations that need to move beyond dashboards, fragmented AI tools, and unsafe automation. It connects to the systems where business happens, converts signals into canonical evidence, organizes that evidence into a temporal-causal knowledge base, and governs every recommendation and action through a complete operating loop: Sense, Reason, Plan, Validate, Decide, Act, and Learn.

Since the July 2026 edition, the platform has grown three capabilities that define its next chapter. First, a **causal proof engine** [LIVE at mizoki3.com/signal as public positioning; measurement machinery PARTIAL]: every conversion classified *caused* or *anticipated*, with confidence intervals, so budget answers to evidence rather than platform self-attribution. Second, the **ORACLE latent-intent layer** [PARTIAL]: consented behavioral micro-signals scored into calibrated, explainable intent predictions — anticipatory intent with proof of causal lift, never "mind-reading," and never audio. Third, **Net Contribution Yield** [IN BUILD]: optimization retargeted from

platform-reported revenue to what an order actually nets after costs and returns — the number the bank account reports, not the number a dashboard applauds.

The core category is Operating Knowledge Intelligence. The core platform loop is SRPVDAL. The core implementation unit is the Domain Intelligence Cell. The core trust mechanism is the Decision Control Plane. The core compounding asset is the Immutable Learning Ledger. The core commercial narrative is PROVE → PROFIT → ANTICIPATE.

THE CORRECT MIZ OKI 3.5 POSITIONING

MIZ OKI stands for Operating Knowledge Intelligence.

Best public-facing positioning: From dashboards to decisions. From AI answers to governed action. MIZ OKI 3.5 is the decision control plane for the AI-native enterprise.

Best acquisition-market positioning (new): Other tools optimize the number your ad platform reports. Mizoki optimizes the number your bank account reports. The narrative arc is PROVE (causal credit with confidence intervals) → PROFIT (net contribution, not top line) → ANTICIPATE (calibrated intent, preview).

Best investor-facing positioning: MIZ OKI is building the operating knowledge layer for enterprises adopting AI agents. As organizations move from analytics to autonomous workflows, they need a trusted system that connects enterprise data, reasons across causal context, governs decisions, executes approved actions, and learns from outcomes.

Best engineering-facing positioning: MIZ OKI 3.5 is a production-oriented autonomous intelligence platform built around the SRPVDAL operating loop. It ingests signals through domain intelligence cells, normalizes them into canonical event envelopes, persists evidence into tenant-isolated stores, projects entities and relationships into a temporal-causal knowledge base, generates plans, validates plans, ranks eligible decisions, executes approved actions, and records outcomes into an immutable learning ledger.

WHAT CHANGED IN THIS REVISION

The July edition established SRPVDAL, Domain Intelligence Cells, and the Decision Control Plane. This revision adds what has been built since, with honest status:

- The Signal Intelligence public surface — causal-credit positioning, proof engine, field-note scenarios, capability dossiers, dual roadmap previews — shipped and live. [LIVE]
- ORACLE latent-intent cells (33–36): micro-signal ingestion with a hard consent gate, the Neo4j intent graph, incrementality and causal credit, and the intent scoring lane — merged, serving observe-only/shadow per the registry. [PARTIAL]
- Pre-conversion perception (VDI and passive-attention capture, session sequence transformer, creative-aesthetic vectors, latent intent bridges) — flag-off, shadow-only. [IN BUILD]
- Measurement rails: Enhanced Conversions and Meta CAPI dedup machinery exist in production code; house-window recompute, drift monitoring, and value writeback ship flag-gated with writeback OFF by default. [PARTIAL]
- Net Contribution Yield: order economics, cohort net-margin computation, and value-rules writeback — writeback hard-off until a verified pilot. [IN BUILD]
- Shopify commerce integration as a packaged offering. [PROPOSED]

- The O-1 signal contract: keystroke dynamics permanently prohibited; audio, gaze, and fine-grained geo rejected at the schema level. Governance canon, enforced by tests.
- Claim-discipline canon hardened into machinery: status labels on every capability, banned-claims list, build-to-claim policy, and automated content QA gating public copy.

1. THE ENTERPRISE PROBLEM

Modern organizations have more data than ever, but most still make decisions through fragmented tools, manual interpretation, disconnected reports, and slow approval chains. Dashboards show metrics, but they rarely explain the operating context. AI tools answer isolated questions, but they often do not remember, validate, govern, or improve the enterprise decision process.

The result is a decision velocity crisis. Signals are detected late. Causes are debated manually. Plans are created in spreadsheets. Policies are checked after the fact. Approvals are not linked to evidence. Actions are executed without measurement design. Lessons are lost when people move on.

1.1 Data Without Context. A performance drop in a campaign may look like a media issue, but the actual cause may involve conversion tracking drift, landing page speed, inventory constraints, product margin, attribution delay, competitive changes, creative fatigue, audience saturation, or legal restrictions. A dashboard exposes the symptom. MIZ OKI is designed to identify the operating cause.

1.2 Automation Without Judgment. Traditional automation asks whether a workflow can execute. MIZ OKI asks whether an action is justified, safe, governed, measurable, and reversible.

1.3 AI Without Memory. Most AI interactions are disposable. MIZ OKI records the signal, reasoning path, plan, validation results, decision, approval, action, outcome, and learning update so the system becomes more valuable over time.

1.4 Governance Outside the Workflow. Many companies have policies disconnected from the systems that act. MIZ OKI embeds governance inside the decision loop, so policy is checked before action, not after damage.

1.5 Attribution Without Causation (new). Ad platforms grade their own homework: they serve impressions, then claim credit for conversions that followed — including conversions that were coming anyway. Reported ROAS is a claim; incremental ROAS is a finding. An enterprise that reallocates budget on platform-reported numbers is optimizing the report, not the business.

2. THE MIZ OKI 3.5 ANSWER

MIZ OKI 3.5 is a governed autonomous intelligence platform built to operate across enterprise systems. It is built around seven foundational ideas:

- **Canonical Event Envelope:** every signal becomes structured evidence with source, timestamp, provenance, tenant, entity, confidence, and governance metadata.
- **Temporal-Causal Knowledge Base:** entities, relationships, sequence, causality, evidence, uncertainty, policy, and outcomes.
- **Domain Intelligence Cells:** each business domain gets a production-grade cell with connectors, schemas, metrics, reasoning, planning, validation, actions, UI, and learning.

- **SRPVDAL Operating Loop:** every intelligence workflow moves through Sense, Reason, Plan, Validate, Decide, Act, and Learn.
- **Decision Control Plane:** the system ranks, blocks, approves, delays, escalates, or executes decisions based on eligibility and governance.
- **Immutable Learning Ledger:** the system remembers what happened, why, what was done, and what was learned.
- **Causal Credit Ledger (new):** every marketing-influenced outcome is classified caused or anticipated, with lift and confidence intervals, and only ledger-grade evidence is allowed to move budget.

3. THE CANONICAL SRPVDAL LOOP

The SRPVDAL loop is the operating system of MIZ OKI 3.5, appearing consistently across README, product narrative, engineering specs, website copy, demos, and investor materials.

3.1 SENSE — the perception layer. It collects data, documents, events, and human context from connected systems and converts them into canonical evidence: Google Ads campaigns, budgets, keywords, search terms, assets, conversions, audiences, and attribution data; OpenRTB bid requests, win/loss notices, seat and exchange metadata, price floors, and consent signals; ESP sends, opens, clicks, conversions, unsubscribes, bounces, complaints, and suppression events; CRM, ecommerce, finance, inventory, legal, support, analytics, documents, and operational systems; human approvals, objectives, guardrails, and strategic context. Since this revision, SENSE also includes consented behavioral micro-signals (Cell 33 [PARTIAL]): dwell, scroll velocity, viewport deceleration, partial-watch depth, tab focus transitions, and content-free form lifecycle events — under the hard prohibitions of Section 12. The output of SENSE is not raw data; it is structured, provenance-backed evidence.

3.2 REASON — converts evidence into understanding: temporal-causal graph traversal, GraphRAG and CausalRAG, root-cause analysis, entity resolution and identity stitching, metric diagnostics, attribution analysis, anomaly detection, policy retrieval, document-grounded reasoning, counterfactual reasoning, and (new) calibrated intent-stage inference over the interest graph with explanation paths. The output is a structured explanation and evidence-backed hypotheses, not yet a decision.

3.3 PLAN — creates possible interventions before any decision. Each plan includes objective, expected impact, evidence used, assumptions, dependencies, risk level, cost and financial impact, confidence, reversibility, required approvals, measurement window, and learning criteria. Since this revision, plans that touch media spend must also declare their counterfactual design — holdout, ghost-bid, or matched-geography — before the first dollar moves.

3.4 VALIDATE — the safety and quality layer. Gates include data freshness, completeness, schema compatibility, metric definition validation, statistical significance, causal plausibility, confounder checks, financial and margin guardrails, budget guardrails, legal and policy checks, privacy and brand safety checks, human approval requirements, rollback readiness, and audit readiness. New in this revision: automated causal refutation (placebo-treatment collapse tests, random-confounder invariance, subset stability) — an estimate that fails refutation is flagged, never shipped.

3.5 DECIDE — the Decision Control Plane ranks eligible options. A decision record includes the selected option, rejected alternatives, evidence used, validation results, risk score, confidence score, approval status, autonomy mode, expected outcome, measurement window, rollback logic, and audit trail.

3.6 ACT — executes only after eligibility is established. New cells default to observe or recommend mode until evaluation gates, approval workflows, and rollback controls are proven. The autonomy ladder (aligned with the public surface): L0 Observe · L1 Recommend · L2 Approval flow · L3 Act within limits · L4 Self-correct · L5

Autonomous. Spend-affecting autonomy is promoted only on measured calibration and stable lift across at least two purchase cycles — never on the strength of a demo.

3.7 LEARN — closes the loop: prediction-versus-actual comparison, causal confidence updates, planning assumption updates, validation threshold tuning, agent and connector performance, metric definitions, playbooks, and learning ledger records. New in this revision: intent predictions carry realized/realized-at labels, so the crystal ball is graded against reality on a standing schedule.

4. THE ADVANCED KNOWLEDGE BASE

A governed, temporal, causal, multi-tenant operating memory — not a static knowledge graph.

4.1 Universal Evidence Model. Every signal enters as evidence: a campaign row, legal clause, customer event, email click, bid request, financial transaction, approval, support ticket, or human note becomes a canonical event carrying tenant ID, loop ID, source system, connector, event type, event time, ingestion time, raw payload reference, normalized fields, entity candidates, metric values, provenance, governance labels, confidence scores, and audit IDs. New in this revision: an order-economics block {gross revenue, component COGS, pick-pack, shipping, payment fees, refund flag, return cost} [IN BUILD] so profit truth travels with the event.

4.2 Temporal-Causal Knowledge Base. Relationship types include semantic (campaign belongs to account), temporal (search-term spike preceded CPA increase), causal (landing-page outage caused conversion decline), financial (margin limits allowable CPA), policy (budget actions above threshold require approval), identity (customer, household, device, email, order, account), and learning (which past interventions worked under which conditions). New relationship families [PARTIAL/IN BUILD]: interest edges (SHOWED_INTEREST with decay), sequence edges (PRECEDES with lift/support/confidence), creative resonance (RESONATED_WITH), and hypothesis-status latent intent bridges with TTL and provenance.

4.3 Tenant-Isolated Memory. Isolation applies to raw events, documents, graph projections, embeddings, connector credentials, policies, decisions, approvals, learning records, and model outputs. No customer memory bleeds into another deployment unless explicitly anonymized, governed, and approved.

4.4 Immutable Learning Ledger. The ledger records the life of every decision: signal, reasoning path, plan, validation result, approval, action, measured outcome, prediction delta, and knowledge update. Most systems record activity. MIZ OKI records organizational learning.

4.5 External Corroboration (new). External market signals enter through the Data Injector lane (Cell 37) as corroborating evidence only — external signals may strengthen a hypothesis; they never independently create a targetable claim.

5. DOMAIN INTELLIGENCE CELLS

A connector pulls data. An intelligence cell turns that data into governed operating knowledge. Each cell includes connector logic, field validation, canonical event normalization, entity mapping, knowledge graph projection, metrics computation, reasoning skills, planning skills, validation gates, decision eligibility rules, approved actions, frontend visibility, audit records, and learning records.

5.1 Google Ads GAQL Intelligence Cell [LIVE] — the flagship proof of the SRPVDAL model: SearchStream extraction, GoogleAdsFieldService validation with caching, MCC traversal, canonical normalization, Firestore/BigQuery persistence, full-funnel mapping, plans generated only after reasoning over evidence, validation

through financial, policy, statistical, causal, and approval gates, mutate actions only under approval and eligibility, and learning records comparing predicted with actual. Default autonomy: observe/recommend.

5.2 OpenRTB Bidstream Intelligence Cell — senses bid requests, win/loss notices, buyer and seat metadata, device signals, price floors, currency, and consent; reasons about inventory quality, auction duplication, floor effects, win-rate changes, identity, fraud indicators, exchange performance, attribution pathways, and supply path optimization.

5.3 ESP and Email Intelligence Cell — normalizes sends, opens, clicks, bounces, unsubscribes, complaints, conversions, suppression, segments, and campaign metadata; reasons about deliverability, fatigue, segment quality, creative performance, send-time effects, revenue contribution, suppression risk, and cross-channel attribution. Because email affects trust, compliance, and deliverability, actions are carefully gated.

5.4 Legal and Policy Intelligence Cell — converts contracts, policies, regulatory text, internal rules, brand guidelines, and approval matrices into machine-checkable decision constraints: not only what the business wants to do, but what it is allowed to do.

5.5 Financial Intelligence Cell — reasons over revenue, margin, cost, cash flow, CAC, LTV, ROAS, payback, budget constraints, profitability, and forecasts. Financial validation sits inside VALIDATE: an action that improves ROAS but destroys margin is not approved automatically.

5.6 Document and Research Intelligence Cell — transforms uploaded files, reports, notes, research, specifications, competitive intelligence, and strategy documents into structured evidence, relationships, decision context, and reusable memory.

5.7 The ORACLE Signal-Intelligence Cell Family (new) — the latent-intent lane, built to the standard "anticipatory intent with proof of causal lift":

- **Cell 33 — Micro-Signal Ingestion** [PARTIAL]: validates and consent-gates behavioral micro-signals; the consent check runs before persistence and fails closed; deny-listed signal categories are discarded at ingest, not stored-but-unsurfaced. In-build extensions [IN BUILD]: Viewport Deceleration Index, dwell and swipe vectors, partial-watch depth, tab-focus transitions, sequence IDs — and the O-1 content-free form-lifecycle contract of Section 12.
- **Cell 34 — Intent Graph & Model Lane** [PARTIAL]: the Neo4j interest graph (SHOWED_INTEREST decay edges, PRECEDES sequence edges) and the scoring path producing calibrated intent stages (awareness → consideration → in-market → purchase-imminent) with explanation paths — no score is delivered without its reasoning. In-build [IN BUILD]: a real-time session-sequence transformer serving only behind a default-off flag, beside — never replacing — the batch path, its attention summaries joining the explanation payload.
- **Cell 35 — Incrementality & Causal Credit** [PARTIAL]: holdout registry (randomized, ghost-bid, matched-geo), deterministic arm assignment, lift with Wilson/bootstrap confidence intervals, and the credit rule: a conversion is incremental only if experiment lift excludes zero AND the model had not already scored the customer in-market at exposure — otherwise it is anticipated. Probabilistic household matches are excluded from all causal math. In-build [IN BUILD]: creative-aesthetic resonance edges and hypothesis-only latent intent bridges with deny-list screening at creation.
- **Cell 36 — Intent Scoring Service** [PARTIAL]: the serving lane for intent APIs, SDK, and MCP tools, deployed beside the legacy Cell 28 neural-processor, which remains untouched.
- **Cell 37 — Data Injector / external market signals** [LIVE per registry]: external evidence as corroboration only.

5.8 Net Yield & Measurement Rails lane (new) [PARTIAL/IN BUILD]: order-economics ingestion from commerce systems; nightly net-contribution computation (gross – COGS – pick-pack – shipping – fees – expected return cost, with expected returns only after at least one observed cycle per SKU, otherwise actual-only and labeled); cohort and order yield APIs; server-side measurement rails (Enhanced Conversions, Meta CAPI with shared event-ID 48-hour dedup, GA4 Measurement Protocol, offline conversions, house attribution-window recompute, drift monitoring at 20% divergence over three consecutive days); and value writeback to platform bidders behind hard-off flags, dry-run by default, recommend-only first. Cost inputs come only from validated tenant configuration — missing costs flag a row incomplete and exclude it from cohort math; costs are never invented.

6. THE MIZ OKI SKILL LAYER

An expandable skill layer organized around the SRPVDAL loop.

Sensing: API ingestion, webhook processing, batch import, document ingestion, field validation, schema detection, change detection, freshness checks, connector health, provenance tracking, consent-gated behavioral capture (new). **Reasoning:** entity resolution, identity stitching, graph traversal, causal path discovery, temporal sequence analysis, root-cause analysis, attribution analysis, anomaly detection, metric diagnostics, policy retrieval, cross-domain synthesis, calibrated intent-stage inference (new), explanation-path generation (new). **Planning:** intervention generation, scenario planning, counterfactual simulation, budget reallocation, channel optimization, journey optimization, risk mitigation, workflow design, testing plan generation, holdout and geo experiment design (new). **Validation:** data quality, metric definitions, statistical, causal (with automated refutation — new), financial, legal, policy, fairness, privacy, brand safety, human approval routing, rollback validation. **Decision, action, learning:** option ranking, eligibility scoring, risk scoring, expected value modeling, approval state management, API execution, workflow triggering, rollback execution, outcome measurement, prediction-versus-actual comparison, causal confidence updates, playbook refinement, causal credit classification (new), net-contribution accounting (new).

7. PRODUCTION ARCHITECTURE

7.1 Evidence Layer — connectors, webhooks, API jobs, batch loaders, document ingestion, raw payload storage, canonical normalization, provenance tracking, source health monitoring. **7.2 Knowledge Layer** — entity stores, relationship stores, temporal and causal graph structures, policy graphs, identity graphs, metric graphs, tenant boundaries, and the learning ledger. Neo4j, TigerGraph, BigQuery, Firestore, vector search, and document retrieval are deployment/projection options inside the broader temporal-causal architecture. **7.3 Reasoning Layer** — GraphRAG, CausalRAG, temporal reasoning, entity resolution, anomaly detection, attribution, metric diagnostics, policy retrieval, and domain intelligence agents. All model access flows through the virtuoso model registry (role-based routing with a governed global fallback); hardcoded model strings are refused at startup and blocked by a canon linter. **7.4 Planning Layer** — plan generation, scenario simulation, counterfactual reasoning, optimization, risk modeling, measurement plan creation, experiment design. **7.5 Validation Layer** — data quality, statistical, causal, financial, policy, privacy, compliance, brand, human approval, and rollback gates. **7.6 Decision Control Plane** — decision eligibility, autonomy modes, approval state, ranking, risk and confidence scoring, audit records, escalation routing. **7.7 Action and Learning Layers** — API action runners, workflow engines, mutate-service integrations, notifications, task systems, rollback handlers, execution monitors; then prediction-versus-actual tracking, causal confidence updates, knowledge updates, model and agent evaluation, policy refinement, and ledger writes. **7.8 Engineering Discipline (new).** The platform is built by a coordinated fleet of AI engineering sessions under owner governance: a coordination ledger where features are claimed before build; tests-first, flag-off workstreams where flags-off is proven byte-identical to pre-change serving; additive-only schema evolution with

old-event validation guarantees; typed approval gates separating merge from deploy; shadow-mode deployment (scores to shadow tables, hypotheses never targeted, ranks logged never applied); and three-level rollback — flags, Cloud Run revision pinning, and stop-writing-never-drop data policy — verified with real revision IDs before any deploy is approved.

8. COMMAND CENTER EXPERIENCE

Not a dashboard — a command center for governed operating intelligence. Core views: the Live Knowledge Graph; the SRPVDAL Loop Monitor; the Decision Queue with validation status, risk, expected value, approvals, rejected alternatives, and audit trail; the Simulation Console; Channel Intelligence pages (Google Ads, OpenRTB, ESP, CRM, ecommerce, finance, legal, customer, product); the Audit and Learning Ledger; and (new) the Signal Intelligence surfaces — intent scores with stage and confidence, predicted-next journey timelines with explanation paths, and the Incrementality panel showing caused versus anticipated with confidence-interval whiskers. Dashboard surfacing of shadow-mode intent metrics remains an open owner decision and ships only when approved.

9. GOVERNANCE AND RESPONSIBLE AUTONOMY

Governance is embedded in the operating loop, not appended to it.

9.1 Decision eligibility statuses: eligible for autonomous action; eligible after human approval; eligible for simulation only; needs more data; blocked by policy; blocked by financial guardrail; blocked by legal or compliance rule; blocked by low confidence; blocked by lack of rollback; rejected. The system does not simply say "do this" — it says why the action is eligible, which validations passed, which risks were checked, what approval exists, what rollback is available, and how the outcome will be measured.

9.2 Hard signal prohibitions (new — enforced at the schema level, with tests, not as policy prose):

- **No audio or microphone-derived signals. Ever.** Rejected at validation; not a setting an operator can enable.
- **No keystroke dynamics — permanently prohibited (owner ruling O-1).** Rejected at collector and ingestion with a named validation tag: no raw keyboard events, no entered text or key identity, no dwell/flight timing, cadence, or pressure, no typing profiles, embeddings, or fingerprints. Permitted form telemetry is content-free lifecycle only (form started, field focused, form completed, form abandoned) with coarse, allowlisted, non-sensitive field classes; sensitive field classes emit nothing.
- **No gaze tracking; no fine-grained geolocation** beyond existing coarse, consent-scoped region.
- **Sensitive-category deny-list** — health conditions, sexuality, religion, financial distress, minors — never predicted, stored, surfaced, or composed into inferences; screening applies at ingest and again at hypothesis creation, with tests proving sensitive composites are rejected at write time.
- **Consent fails closed:** no matching consent scope, no persistence. **Erasure cascades** across every store — warehouse rows, graph nodes and edges, vectors, and derived features. **Probabilistic identity matches** are usable for recall only and are excluded from all causal measurement.

9.3 Anticipation is never credit (new). A prediction never claims causal credit on its own; acting on predicted intent requires a registered holdout, so anticipation is never mistaken for causation. Prediction never grades itself.

9.4 Autonomy promotion (new, spend-affecting lanes): observe-only default; promotion requires calibration (Brier ≤ 0.20 , AUC ≥ 0.72 on held-out journeys — targets, graded on real data), stable incremental lift across at least two purchase cycles, and passing causal refutation checks. Writeback to external platforms ships flag-off and dry-run, and is enabled recommend-only first.

10. FLAGSHIP DEMO STORYLINES

10.1 Google Ads performance drop (the SRPVDAL proof). SENSE detects a seven-day ROAS drop and ingests campaign, search-term, conversion, budget, asset, geography, device, landing-page, ecommerce, inventory, margin, and prior-decision data. REASON localizes the drop to specific broad-match terms, mobile traffic, landing-page speed, and low-stock high-margin products, separating attribution delay from true degradation. PLAN generates options from negative keywords to waiting for attribution maturity. VALIDATE checks GAQL compatibility, freshness, conversion lag, significance, margin, inventory, budget policy, brand policy, approvals, and rollback. DECIDE selects a staged plan. ACT executes only what is approved. LEARN compares predicted with actual and updates the cell.

10.2 The Signal Factory (live public demo) [LIVE]. On the production runtime at mizoki3.com/demo/signal: raw connector events arrive, normalize into canonical events, face the gate, and travel all seven SRPVDAL stages — including one deliberate guardrail block caught in red. No signup. The public Signal surface tells the same story in composite field notes under strict labeling: the retargeting doorman, the brand-search finding, the unneeded coupon, prove-it-or-lose-it CTV.

11. INVESTOR POSITIONING

MIZ OKI 3.5 defines a new enterprise infrastructure category: Operating Knowledge Intelligence. The market is moving from dashboards to autonomous analytics, from chatbots to agents, from agents to governed agent systems, from workflows to decision infrastructure, from static warehouses to live knowledge systems, and from AI answers to AI-operated decisions.

The moat is not a single model. The moat is the operating system around models: canonical events, domain intelligence cells, temporal-causal memory, evaluation gates, decision eligibility, approval workflows, action controls, rollback logic, audit trails, and accumulated decision history. This revision adds the commercial edge no adjacent vendor holds whole: measurement vendors cannot execute; execution tools cannot prove causation; neither sees unit economics. MIZ OKI proves (causal credit), profits (net contribution), and anticipates (calibrated intent) inside one governed loop.

12. PRODUCTION ROADMAP (status-honest)

- **Phase 1 — Foundation [LIVE]:** canonical envelope, tenant model, connector registry, raw event store, Firestore/BigQuery persistence, knowledge graph projection, SRPVDAL loop records, audit ledger, Command Center shell, Google Ads cell, recommend-only autonomy.
- **Phase 2 — Reasoning and Planning [LIVE/PARTIAL]:** entity resolution, metric diagnostics, causal path reasoning, GraphRAG/CausalRAG, counterfactual planning, plan generation, domain agents, simulation console, decision queue, validation framework.
- **Phase 3 — Validation and Decision Control [PARTIAL]:** Decision Eligibility Layer, policy/financial/statistical/causal gates, approval workflows, human review routing, audit-ready decision records, autonomy modes, rollback checks.
- **Phase 4 — Governed Action [PARTIAL]:** approved API actions, post-approval mutate actions, workflow execution, action logs, rollback support, execution monitoring, post-action measurement, learning records.

- **Phase 5 — Continuous Learning** [IN BUILD]: prediction-versus-actual tracking, causal confidence updates, playbook refinement, agent performance scoring, knowledge self-healing, cross-loop learning, executive reporting, investor-grade performance evidence.
- **Phase 6 — Signal Intelligence & ORACLE (new)** [PARTIAL → IN BUILD]: intent cells serving observe-only; pre-conversion perception in shadow; first registered holdouts; promotion gates graded on real data.
- **Phase 7 — Net Yield & Commerce (new)** [IN BUILD → PROPOSED]: order economics live per tenant; net-contribution writeback after verified pilot; packaged Shopify offering [PROPOSED]. (*r3.5.2: unchanged; F3 IN BUILD observe-only per ruling 2026-08-19-A.*)
- **Phase 8 — Full Growth Control Build** [IN BUILD, by ruling 2026-08-19-B]: F1, F2, F4, and F5 machinery built concurrently with per-frontier gates:
 - **F4** — the automated micro-geo calibration loop builds end-to-end (geo reservation engine, synthetic-control estimator, Bayesian MMM prior updater), with each geo reservation executing ONLY through the L2 approval flow until two clean calibration cycles complete; thereafter eligible for bounded autonomy within declared spend caps.
 - **F1** — element-level creative effect estimation builds on the CreativeSemanticProfile lane + DR-Learner core; effect estimates are labeled *provisional* until pilot-scale creative volume exists; generated-creative deployment always retains human approval.
 - **F2** — dynamic-treatment-regime machinery and the ledger's outcome-horizon dimension build now; long-horizon findings publish ONLY after ≥ 2 observed quarters per cohort (the guard is code, not policy prose).
 - **F5** — treasury gating builds in two versions: **v1 (now)** consumes owner-declared liquidity floors and covenant limits from validated tenant configuration (`config/treasury_constraints.yaml`) as hard constraints in VALIDATE/DECIDE, with veto + human routing on breach; **v2 (when the Capital-division treasury feed lands)** replaces declared values with live positions. Missing configuration = no F5 constraint claimed (fail-closed, honest health).
- **Former Phase 9** is dissolved into Phase 8 above; Gate-3 customer autonomy expansion (L3/L4 within caps, criteria met on real data) remains as previously approved.

13. CLAIM DISCIPLINE (hardened)

Public claims must be disciplined; this is now machinery, not intention.

- **Status vocabulary, mandatory:** every capability statement carries [LIVE], [PARTIAL], [IN BUILD], or [PROPOSED]; every metric is labeled verified result, benchmark result, design target, pilot target, or illustrative scenario. Design targets are never stated as observed performance.
- **Banned claims, permanent:** "mind-reading" or any audio-listening framing; third-party published results presented as MIZ OKI's own (including the Airbnb synthetic-data KL-divergence figures, which are Airbnb's); unvalidated novelty methods ("Quokka Swarm" and kin); guaranteed outcomes; present-tense deployment claims for unbuilt capability.
- **Build-to-claim policy:** where public copy promises machinery, the default resolution is building the code to make the claim true; relabeling to "operating design" is the fallback, recorded in a build-debt ledger. Zero unbacked present-tense claims may survive a release.

- **Composite scenarios:** field-note stories are labeled composites with illustrative numbers until a customer signs off on a verified case study; one verified number replaces its composite and outranks all others. Preview labels flip only after verified pilot results are written to the claim ledger — a promise now stated on the public site itself.
- **Enforcement:** automated content QA gates public copy in CI — banned strings, unlabeled numbers, missing preview framing, and unbacked machinery claims fail the build.

14. FINAL PUBLIC-FACING PARAGRAPH

MIZ OKI 3.5™ is an Operating Knowledge Intelligence platform that helps enterprises transform fragmented data, documents, systems, and human decisions into governed autonomous intelligence. The platform connects to business systems, normalizes signals into canonical evidence, maps them into a temporal-causal knowledge base, and runs every recommendation through the SRPVDAL loop: Sense, Reason, Plan, Validate, Decide, Act, and Learn. Unlike dashboards that only show what happened or AI tools that generate isolated answers, MIZ OKI creates a decision control plane where plans are validated against data quality, causality, policy, financial impact, approval requirements, and rollback readiness before action is allowed. It proves which outcomes marketing actually caused, prices every order at what it truly nets, and — in preview — anticipates emerging customer intent under strict consent and privacy prohibitions, with every prediction required to earn its credit through a registered experiment. Every outcome is written back into an immutable learning ledger, allowing the organization's intelligence to compound over time.

15. FINAL ENGINEERING SUMMARY

MIZ OKI 3.5™ is a production-oriented autonomous intelligence platform built around the SRPVDAL operating loop. The system ingests signals through domain intelligence cells, normalizes them into canonical event envelopes, persists evidence into tenant-isolated stores, projects entities and relationships into a temporal-causal knowledge base, generates plans through reasoning and simulation, validates those plans through policy, financial, statistical, causal, and approval gates, ranks eligible decisions through a Decision Control Plane, executes approved actions through governed integrations, and records outcomes into an immutable learning ledger. The signal-intelligence lane adds consent-gated behavioral ingestion, a calibrated and explainable intent model lane, an incrementality engine issuing caused-versus-anticipated credit under confidence intervals, and a net-contribution accounting lane — all serving observe-only or shadow until promotion gates are met on real data, all reversible at three levels, and all governed by schema-enforced prohibitions on audio, keystroke dynamics, gaze, fine geolocation, and sensitive-category inference.

16. HIGHEST-PRIORITY NEXT IMPLEMENTATION MOVES

- Grade the intent model lane on real realized-label data against the promotion gates ($\text{Brier} \leq 0.20$, $\text{AUC} \geq 0.72$).
- Register the first real holdout before any activation flag flips; keep all writeback flags off until a pilot validates.
- Complete pre-conversion perception in shadow; review shadow metrics before any serving decision.
- Land the first tenant's real cost configuration and produce the first net-contribution cohort report.
- Replace the first composite field-note number with a verified pilot result under written customer approval, and flip the corresponding Preview label per the claim ledger.

- Resolve open owner decisions: intent threshold tuning plan, bridge approval tier, external-signal sampling economics, dashboard surfacing of shadow metrics.
- Keep OFFERING_MAP, WIRING, README, website, demo narrative, investor deck, and this whitepaper on one language — with the live Cell Registry as ground truth.
- Complete the Phase-8 frontier machinery under its gates; land `treasury_constraints.yaml` and the F4 geo-candidate list as owner-provided configuration; run the first approval-gated micro-geo reservation as the opening act of continuous calibration.

CLOSING STATEMENT

MIZ OKI 3.5 is the operating intelligence layer for the autonomous enterprise. It turns fragmented business signals into governed knowledge, turns knowledge into validated decisions, turns decisions into controlled action, and turns outcomes into institutional learning — and it now proves what it caused, prices what it truly earned, and anticipates what comes next without ever crossing the lines it drew for itself in code. The future enterprise will not be run by dashboards alone. It will be operated by systems that can sense, reason, plan, validate, decide, act, and learn under governance. MIZ OKI 3.5 is built to become that system.

Appendix A — Amendment r3.5.2 of record (verbatim)

Source: MIZOKI_3.5_WHITEPAPER_r3.5.2_AMENDMENT_AUG2026.md (formal amendment of record, August 19, 2026 — second ruling of the day), preserved byte-for-byte at docs/whitepapers/archive/MIZOKI_3.5_WHITEPAPER_r3.5.2_AMENDMENT_AUG2026.md. Amendments 1 and 3 are applied in place in §12 and §16 above; the RULING, AMENDMENT 2 and the CONSISTENCY NOTES are reproduced here verbatim (heading level nested one deeper; text unchanged).

RULING (owner, 2026-08-19-B)

Growth Frontiers **F1 (High-Dimensional Creative Component Unbundling)**, **F2 (Multi-Quarter LTV Treatment Regimes)**, **F4 (Continuous Bayesian Calibration via Automated Micro-Geo Holdouts)**, and **F5 (Treasury-Gated Spend Governance)** are advanced from PROPOSED to **IN-SCOPE BUILD**. Roadmap Phases 8 and 9 collapse into a single active build program. All four frontiers build now, under the per-frontier gating below — machinery ships; authority and findings are still earned.

Preface (r3.6): AMENDMENT 2 governs docs/MIZOKI_SIGNAL_GROWTH_CONTROL_UNIFIED_SYSTEM_r2.0.md §5/§6 where they differ; carried here so one document holds the ruling of record.

AMENDMENT 2 — Growth Control r2.0 §5 Roadmap and §6 Deployment State

§5: Phases 8–9 replaced by Phase 8 as written above. §6 table rows updated:

Capability	Status (amended)
F4 auto micro-geo calibration	IN BUILD — reservations L2-approval-gated until 2 clean cycles
F1 creative unbundling	IN BUILD — estimates provisional until pilot creative volume
F2 multi-quarter LTV regimes	IN BUILD — findings gated on ≥ 2 observed quarters
F5 treasury-gated spend governance	IN BUILD — v1 config-declared floors; v2 on live treasury feed

All other rows unchanged. The 90-Day Pilot, Decision Jobs catalog, and Gate-3 rulings of 2026-08-19-A stand.

CONSISTENCY NOTES (double-check record)

- No amendment alters: O-1 PRIVACY LOCK; consent fail-closed; sensitive-category deny-list; probabilistic-identity exclusion from causal math; "prediction never grades itself"; Preview labels flip only on verified pilot numbers; writeback flags off until pilots; autonomy promotion gates ($\text{Brier} \leq 0.20$, $\text{AUC} \geq 0.72$, ≥ 2 cycles).
- F4's nature is spend-affecting and therefore CANNOT be observe-only — its gate is per-action human approval, which is the correct control for its class; this is deliberate, not an exception.
- Ratified vocabulary (ruling 2026-08-19-A) applies to both base docs and this amendment.
- Owner-provided inputs now on the critical path: tenant cost config (net yield), `treasury_constraints.yaml` (F5 v1), candidate geo list + spend caps (F4), pilot tenant selection.

Amendment r3.5.2 / r2.1 · applied atop the base documents · August 19, 2026

MIZ OKI 3.5™ — Master Positioning Document r3.6 · September 2, 2026 · folds r3.5.1 (August 18, 2026) + amendment r3.5.2 (August 19, 2026)