

MIZ OKI 3.5

Operating Knowledge Intelligence

External whitepaper for governed enterprise decisions.

Version 3.1 · 23 September 2026

Reconciled to MIZOKICloudRun measured state

Supersedes v3.0 (13 August 2026); July v2.0

Platform ceiling: built, pre-benchmark

Stage 3 recommend-only is the fleet default

Executive summary

MIZ OKI 3.5 is governed decision infrastructure. It connects business signals to a shared operating memory, develops evidence-backed options, checks whether those options are eligible, and carries authorized decisions through execution controls and outcome review. Its operating loop is Sense, Reason, Plan, Validate, Decide, Act, Learn: SRPVDAL.

The purpose is to improve the quality and speed of business decisions while preserving human authority. In commerce and media, that means evaluating incremental value alongside contribution margin, returns, inventory, liquidity, and measurement uncertainty. An advertising platform's attribution report is an input to this process; it does not establish causal lift.

Five primitives organize the platform: the Canonical Event Envelope, the temporal-causal knowledge base, Domain Intelligence Cells, the Decision Control Plane, and the Immutable Learning Ledger. Shared platform services supply identity, ingestion, validation, model routing, credentials, audit, and replay. Domain cells supply business context and proposals within those controls.

Since July, the implementation has advanced across canonical ingestion and graph projection, signed validation evidence, Growth Control, decision jobs, authenticated onboarding, measurement tooling, and operational hardening. The current architecture uses a Firestore-backed Operating Knowledge Graph with BigQuery for analytical workloads. Causal Growth Control is the commercial offer; Intent Engine v2 is its anticipatory capability. The September master positioning paper remains a separate plan of record.

These advances do not constitute a general autonomy or business-performance certification. Stage 3 recommend-only remains the fleet default. Measurement and Net Yield writebacks remain OFF. The micro-geo calibration machinery is recorded as live-armed for a bounded self-pilot, with approval required for reservations and promotion still unearned. Customer performance claims require their own valid benchmark or pilot evidence. [1–5]

1 The enterprise problem

A business decision often depends on evidence distributed across advertising, commerce, customer relationships, finance, documents, and operational systems. A campaign decline may reflect a tracking defect, a late conversion, weak creative, a stock constraint, or a change in margin. Acting on the visible metric alone can intensify the underlying problem.

MIZ OKI addresses four connected gaps. First, it retains source and time context when converting observations into evidence. Second, it records competing explanations and a no-action baseline before selecting an intervention. Third, it applies validation and approval requirements before external action. Fourth, it compares the eventual outcome with the original expectation and preserves that record for future decisions.

The intended value is a repeatable operating process whose decisions can be inspected, challenged, and improved. Faster decision cycles and financial improvements are design targets until measured for a defined workflow and population.

2 The category of Operating Knowledge Intelligence

Operating Knowledge Intelligence is MIZ OKI's description of the governed layer connecting models to business operations. It combines evidence normalization, temporal memory, reasoning, validation, eligibility, authorization, bounded execution, and learning. The software surrounding the model determines whether an answer can become an accountable business action.

Decision intelligence is an established market category. Grand View Research estimates a global market of USD 17.8 billion in 2025 and projects USD 20.7 billion in 2026 and USD 53.2 billion by 2033, with a projected 14.4% annual growth rate from 2026 to 2033. These are the publisher's market estimates and forecast, checked at this edition's review cutoff; they are not MIZ OKI revenue, adoption, or performance results. [12]

MIZ OKI's positioning centers on a complete decision contract: what evidence was available, which alternatives were considered, why an option passed or failed, who approved it, what actually happened, and what was learned. Product differentiation must be evaluated against these inspectable controls and customer outcomes.

3 The canonical SRPVDAL loop

Every governed decision uses the same seven stages. Model roles and service boundaries are separate from those stages: a model may assist several stages, while a domain service remains subject to shared controls. [1]

3.1 Sense

SENSE receives authorized source data and converts it into versioned evidence. Canonicalization records the tenant, source, event identity, schema, provenance, payload integrity, and distinct time axes. Duplicate events are handled idempotently. Behavioral intent signals require consent before persistence. Source availability, data freshness, and missing configuration must remain visible.

3.2 Reason

REASON develops explanations and hypotheses using point-in-time evidence, graph relationships, retrieved documents, domain analysis, and causal methods where justified. It separates reported attribution, prediction, association, and experimentally supported effects. A graph edge or model explanation is not by itself causal proof.

All governed language-model reasoning routes through the Virtuoso dispatcher. Model selection is registry-controlled; fallback and degraded operation are observable. This whitepaper does not freeze vendor model names that can change independently of the operating architecture.

3.3 Plan

PLAN develops candidate interventions and a mandatory no-action baseline. Each proposal identifies its objective, assumptions, dependencies, economic exposure, uncertainty, reversibility, approval requirements, outcome window, and stopping conditions. Plans can include collecting more evidence or running an experiment.

3.4 Validate

VALIDATE applies the full required domain battery, including freshness, point-in-time integrity, statistical and causal adequacy, privacy, policy, financial constraints, and rollback readiness. Missing mandatory checks invalidate a ValidationPassport. Callers cannot select an easier subset or bypass a hard gate. Undeclared risk exposure receives a hold.

3.5 Decide

DECIDE applies the Decision Control Plane and Decision Eligibility Layer. The DEL Score has a platform floor of 80.0; domain policy may raise it. The governed outcomes are eligible, approval-required, experiment-required, advisory-only, and blocked. Passing the numerical floor alone does not authorize action.

For intent-led exposure, an eligible holdout must be registered before first exposure. Incremental return and incremental contribution inform allocation; platform-reported ROAS remains reported evidence. Measurement registration, recommendation, approval, cohort export, and platform activation are distinct states.

3.6 Act

ACT either records a recommendation, routes an approval package, or executes an authorized bounded action. External execution requires the DCP authorization, an appropriately promoted actuator in the runner registry, and a real execution adapter. Authorization is bound to the exact action, expires, and cannot be redeemed twice. Missing execution prerequisites must never produce a fictional success receipt.

Stage 3 recommend-only remains the default. Bounded execution is earned per action class and account, with evaluation, approval, rollback, monitoring, and outcome proofs. Irreversible actuators cannot receive Stage 4 authorization under the current constitution. Legal conclusions, fiduciary actions, risk committee decisions, and binding CRE commitments remain outside autonomous authority.

3.7 Learn

LEARN compares predictions with observed outcomes, retains failures and contradictions, and produces calibration, policy, or playbook proposals. Changes to governed thresholds or autonomy require their own review and approval. A prediction cannot grade itself, and observed association cannot silently become a causal result.

4 The advanced knowledge base

The Operating Knowledge Graph is a Firestore-backed decision substrate, supported by BigQuery analytics and governed document retrieval. It connects entities, source observations, evidence, policies, hypotheses, plans, decisions, approvals, actions, and outcomes. Intent graph durability also uses Firestore. References to Neo4j as the current graph store in the July edition are superseded. [2,3]

The Canonical Event Envelope distinguishes `occurred_at`, `observed_at`, `available_to_model_at`, and `ingested_at`. Model availability cannot precede observation. Point-in-time reads preserve what could actually have been known when a decision was made. Retrospective loads retain their backfill provenance instead of moving evidence availability into the past.

Canonical ingestion and graph projection are separate steps. Source observations pass through the governed ingress before eligible graph updates. The repository now contains the canonical-to-graph projector and external-intelligence injection routes that forward through that same authority boundary. A deployed route does not establish that every provider or tenant supplies complete, current data.

Entity resolution retains canonical identifiers, source evidence, confidence, and unresolved conflicts. Deterministic and probabilistic identity relationships remain distinct; probabilistic household edges support retrieval recall and are excluded from causal measurement. Tenant boundaries apply to evidence, credentials, policy, decisions, and learning records.

Self-healing refers to explicit mechanisms for duplicate handling, stale evidence, contradiction detection, schema drift, confidence review, provenance repair, and review queues. It does not imply an independently measured self-healing rate or unrestricted autonomous graph mutation.

5 Domain Intelligence Cells

A Domain Intelligence Cell owns a bounded business objective and declares its evidence, reasoning, planning, validation, approval, action, and outcome contracts. Engineering service registrations and product domains are different inventories; this external paper names functions rather than internal cell numbers.

5.1 Media acquisition and commerce

Causal Growth Control brings together causal measurement, order economics, anticipatory intent, and governed decisions. MIZOKI Signal for Shopify is the commerce-oriented commercial instance of the Media Acquisition domain. Merchant-owned advertising accounts remain the operating model.

The four connected loops are PROVE, PROFIT, ANTICIPATE, and GOVERN AND COMPOUND. PROVE measures incremental effects under a defensible design. PROFIT reconciles contribution after relevant costs and returns. ANTICIPATE develops consented hypotheses through Intent Engine v2. GOVERN AND COMPOUND applies the decision controls and records the outcome. Each loop has its own maturity and configuration state. [4,5]

Six High-Value Decision Jobs organize the operating work: Incrementality, Waste Prevention, Margin Control, Learning Stability, Executive Defensibility, and Team Leverage. Their outputs should expose evidence, options, constraints, and next approval requirements.

5.2 OpenRTB and supply quality

Bidstream analysis can inform inventory quality, auction duplication, pricing, win rates, fraud indicators, and supply-path evaluation. Ghost-bid exposure remains proposed and engagement-gated; the presence of measurement code or a schema does not establish a live ghost-bid program.

5.3 Email and lifecycle

Lifecycle analysis covers deliverability, fatigue, suppression, cohort behavior, creative, and revenue contribution. Connector availability, credential readiness, and permission to send are separate. Suppression and consent controls remain prerequisites; this edition makes no claim of unrestricted automated outreach.

5.4 Legal and policy

Legal and policy evidence supports grounded analysis and machine-checkable constraints. The domain remains advisory where legal conclusions are involved. Counsel demonstrations illustrate workflows and do not establish a licensed professional service or tenant-specific legal outcome.

5.5 Financial intelligence

Financial reasoning evaluates revenue, contribution, costs, liquidity, customer value, and scenario assumptions. Finance remains advisory-only under the current domain posture. Treasury controls are an additional configured boundary and must not be implied for an account whose treasury inputs and enforcement are absent.

5.6 Commercial real estate

CRE underwriting and prospecting use evidence, simulation, identity checks, and governed outreach proposals. Binding valuation, engineering, environmental, credit, and transaction decisions remain with the authorized human professionals. A built prospecting workflow is not evidence of autonomous binding commitments.

5.7 Intent Engine v2

Intent Engine v2 develops anticipatory, consented hypotheses with explanation and uncertainty. Its session-bounded attention and outcome modules, creative semantic profiles, expiring hypotheses, and durable validation records have different retention contracts. The capability remains in build or partial according to its component; shadow output does not authorize exposure.

The original ranking evaluation failed its baseline comparison and must not be presented as successful product accuracy. Later fixture diagnostics and evaluation infrastructure do not substitute for admissible forward labels, independent outcome measurement, or a qualified design-partner readout. Promotion requires the specified calibration and discrimination gates, stable lift across at least two purchase cycles, and human approval. [1,4,6]

5.8 Estate stewardship

Estate stewardship supports ownership and obligation evidence, ledger-to-bank reconciliation, and holding-period review. Its implementation remains built, pre-benchmark and advisory-only; dispositions and fiduciary decisions remain human-controlled.

5.9 Enterprise risk

Enterprise risk evaluates aggregate and correlated exposure and supplies veto or hold decisions to the shared control plane. It does not originate business actions. Committee and board decisions remain human-controlled; domain performance remains pre-benchmark.

6 Horizontal Platform Services

6.1 Document and research knowledge

Shared ingestion and retrieval provide source-grounded document evidence to domain cells. A contract informs legal analysis; an economic model informs financial reasoning; a campaign brief informs media planning. The domain owns the decision, while the shared service preserves provenance and retrieval context.

6.2 Identity and semantic resolution

Shared identity and semantic contracts prevent each domain from inventing incompatible identifiers. Material corrections and merge or split proposals retain source history and review requirements. Ambiguity remains explicit.

6.3 Provenance and temporal evidence

Evidence records carry source, observation and availability times, integrity information, transformation lineage, and assessment metadata. Rejected, superseded, and contradictory claims remain distinguishable from current accepted evidence.

6.4 Evaluation and policy

Reusable evaluation suites, full ValidationPassport batteries, policy rules, risk checks, and approval matrices protect the decision pathway. Hard requirements are non-bypassable. A policy change must follow the platform's amendment and review process; an exception record does not waive constitutional controls.

6.5 Model and agent routing

The Boss orchestrator coordinates tools and specialists through the governed routing layer. Specialist agreement does not replace validation. Registry-controlled model selection, explicit fallbacks, and versioned execution evidence support diagnosis and provider changes without silently changing authority.

6.6 Connectors and credential boundaries

The connector gateway separates sensing permissions from execution permissions. Implemented adapters span advertising, analytics, commerce, and other approved sources. Source-specific credentials, account identifiers, tenant mapping, and lineage acceptance determine whether an adapter is operational for a customer. The correct state may be ready-to-connect, not_configured, degraded, or verified for a specific path.

Shopify webhook ingestion and backfill use governed receiver and forwarding paths. Historical statements that the integration was entirely unbuilt are superseded; universal claims of live merchant coverage would also be inaccurate.

6.7 Audit and replay

The audit pathway records stage outputs, evidence references, decisions, approval identity, authorization, execution receipts, and outcomes. Signed ValidationPassport packages include per-tenant chain metadata and a verification path. Cryptographic verification establishes integrity and provenance; it does not establish that a business effect was causal. Replay reconstructs the recorded decision without repeating the external action.

7 Production architecture

The implementation uses containerized services on Google Cloud Run. The architecture separates canonical evidence ingestion, Firestore graph projections, BigQuery analysis, domain reasoning, policy and validation, DCP approval, action execution, and audit replay. Model calls pass through the Virtuoso dispatcher. Frontend access uses authenticated application boundaries and typed backend-for-frontend adapters. [1–3]

Operational correctness includes tenant scoping, least-privilege identities, audience-checked service authentication, secret separation, bounded retries, idempotency, observable failure states, and rollback. An authenticated application error must remain an error rather than being hidden inside a success envelope.

A deployment report and live verification answer different questions. Deployment evidence shows what a workflow attempted and recorded. Live verification establishes the serving revision, traffic routing, authentication, dependency state, and exercised application path. An unauthenticated response alone cannot establish the health of a private service.

The marketing website has its own owner-controlled release workflow. Updating documentation or merging a repository change does not, by itself, establish that the revised public document is being served.

8 Command Center experience

The Command Center connects evidence inspection to action review. Its operating views include knowledge and evidence paths, SRPVDAL progress, proposals and approvals, experiments, economics, and audit history. Growth Control work added decision-job surfaces, pilot reporting, and views of anticipatory and causal evidence.

Recent hardening includes explicit freshness contracts for economics and lane readiness. Values without valid freshness evidence must not appear ready merely because numbers are present. Empty, stale, unavailable, and not-configured states are meaningful product outputs, not gaps to fill with invented metrics. [7]

Supabase authentication is provisioned for protected Command Center routes; repository records include signed-out and forged-header denial checks. This supersedes the August statement that the identity provider was not provisioned. Authenticated onboarding is the designated path for tenant-owned identity, connector scope, economics, objectives, and pilot inputs. Collection does not arm a feature. Operator secrets remain outside tenant forms. Acceptance requires the required inputs and successful connection and lineage checks for the intended use.

Decision Studio is recorded as deployed and live-verified in September, with subsequent edge and monitoring hardening. It is a public synthetic demonstration with no live tenant data or dispatch authority. Demonstrations use labeled scenarios; a convincing simulation is not a production transaction or customer-result certification.

9 Governance and responsible autonomy

Every proposed material action must pass identity and signature checks, policy alignment, context sufficiency, and hard constraints. DEL scoring sits inside that contract; it cannot override missing evidence or forbidden autonomy.

The default operating posture is recommendation. The L0–L5 autonomy vocabulary and the Stage 3/Stage 4 actuator mechanism describe related but distinct controls and must not be used interchangeably. Promotion is scoped and attributable, with no blanket fleet upgrade implied by a successful pilot or engineering closure.

Intent-related collection is consent-first. Keystroke, audio, gaze, and prohibited fine-geolocation signals are excluded by the privacy lock. Predicting, storing, or surfacing intent about health conditions, sexuality, religion, financial distress, or minors is prohibited. Probabilistic identity relationships cannot enter causal analysis.

Holdouts precede exposure. Measurement protects control integrity, records the estimand and minimum detectable effect, accounts for delayed outcomes and interference, and applies design-appropriate diagnostics. Statistical confidence, calibrated prediction, and causal identification are separate requirements.

Global, tenant, and actuator boundaries, expiry, idempotency, preconditions, external verification, and rollback constrain execution. Human approval does not legalize a prohibited use or replace a missing adapter. [1,4]

10 Security and compliance posture

Security controls include tenant isolation, identity and access boundaries, encrypted transport and storage, secret management, auditability, and failure visibility. Coverage must be demonstrated for the specific deployment and data path; no universal security or certification guarantee is made.

Customer-managed encryption is partially implemented. Recorded coverage includes specified BigQuery measurement ledgers; a separate later infrastructure record describes a key whose wider datastore adoption remains pending. These records concern different scopes and must not be generalized into universal coverage.

Deployment-specific key adoption and verification remain necessary. Similarly, a signing implementation is not proof that every runtime has a valid signing configuration.

Post-quantum cryptography remains a deployment-specific design target. NIST finalized FIPS 203 for ML-KEM key encapsulation, FIPS 204 for ML-DSA signatures, and FIPS 205 for SLH-DSA signatures in August 2024. MIZ OKI does not claim completed platform-wide adoption of those standards. [13]

The architecture supports evidence collection relevant to governance frameworks and privacy obligations. It does not establish SOC 2 or ISO certification. The European Commission records general AI Act applicability from 2 August 2026, with revised high-risk dates of 2 December 2027 for Annex III and 2 August 2028 for relevant product-embedded systems. Classification and applicable obligations require deployment-specific legal assessment. [14]

Patent and intellectual-property status must be established from current filing records. This edition makes no issued-patent claim and does not treat a draft filing or previous whitepaper statement as proof of current pending status.

11 Competitive differentiation

Enterprise ontology, agent, and automation products already provide substantial operational capabilities. Palantir, for example, documents ontology actions that change objects, properties, and links and integrate operational workflows. MIZ OKI should be evaluated against real competing capabilities, without implying that other platforms lack governance or action controls. [15]

MIZ OKI's product emphasis is the combination of a canonical seven-stage loop, point-in-time evidence, no-action baselines, full validation batteries, DEL eligibility, exact-action authorization, registered execution adapters, causal measurement, and an outcome ledger. The media and commerce offer adds contribution economics, intent hypotheses, and decision jobs to this control structure.

Modularity, integration effort, operating cost, usability, and business impact are evaluation questions. They are not established advantages until comparative evidence supports them. Historical vendor adoption figures are omitted because they do not substantiate MIZ OKI's own capabilities.

12 ROI framing

Value measurement starts with a named workflow, eligible population, time window, baseline, and outcome definition. Decision-cycle time should distinguish recommendation preparation, approval delay, and completed external action. Financial measures should distinguish revenue, contribution, cash, and spend.

iROAS is incremental revenue divided by the relevant incremental advertising spend under the declared measurement design. Incremental contribution or profit additionally accounts for cost of goods, fulfillment, fees, returns, and other agreed variable costs. Liquidity and inventory constraints remain separate decision boundaries. These measures must not be blended into a single unexplained score.

An ROI assessment includes implementation and operating cost, validated labor savings, measured incremental contribution, and defensible avoided-loss estimates. A blocked proposal is not automatically an avoided loss; that claim needs a counterfactual and assumptions. The platform does not publish universal CAC, ROAS, ROI, payback, entity-resolution, or self-healing improvements as achieved results.

13 Illustrative scenario

This is an **illustrative scenario**, not a customer result or executable instruction. A proposed USD 5.0 million distribution receives a DEL Score of 41 and fails the required policy and rollback conditions. It is blocked. A USD 3.2 million alternative may be prepared with its evidence, constraints, and rejected options for human review.

A financial or fiduciary proposal does not become autonomous because an alternative has a better score. The relevant advisory-only and forbidden-autonomy rules still apply. The useful output is an attributable decision package, with approval and execution state recorded honestly.

14 Flagship demonstration storyline

A demonstration begins with an advertising-performance decline and labels its data as an illustrative scenario. SENSE assembles campaign, conversion, commerce, margin, inventory, and timing evidence. REASON separates possible demand deterioration from measurement delay, creative fatigue, and stock constraints.

PLAN compares holding the current policy, reducing exposure, changing creative, reallocating within approved bounds, and collecting additional evidence. VALIDATE checks the complete required battery, including freshness, consent, measurement design, economics, risk, and reversibility.

DECIDE produces the eligible outcome and approval requirement. Under the default posture, ACT returns a recommendation or approval package. External execution appears only when the specific authorization, actuator, adapter, configuration, and verification requirements are satisfied. LEARN compares the observed outcome with the registered expectation and retains disagreements and failed hypotheses.

The demonstration should let a reader follow one decision from its source evidence to its final receipt. It should never imply that moving a simulator control spent money or proved incremental lift.

15 Deployment model and onboarding

The platform's implementation model is a multi-tenant cloud service with governed integration boundaries. Commercial onboarding follows a 90-Day Growth Control Pilot: Observe during days 1–30, Validate during days 31–60, and Recommend during days 61–90. The schedule organizes work; progression depends on evidence and readiness rather than elapsed time alone. [5]

Tenant-owned setup is collected through authenticated onboarding. Operator-owned credentials, key configuration, infrastructure changes, and deployment verification follow operator procedures. A completed form is not an active connector, and an active connector is not an authorized actuator.

Pilot acceptance requires a named scope, reconciled economics and source data, an eligible measurement design, consent and retention controls, stable outcome collection, reviewable reports, and a clear owner. Self-pilot evidence and an external design-partner engagement are separate evidence classes. A self-pilot does not automatically release customer-facing performance labels.

16 Production status and roadmap

The July roadmap described the foundational stages as future work. The current repository contains substantial implementations in those areas. The next stage is controlled operation and qualified evidence, together with closure of specific integration, configuration, and verification gaps. Engineering closure, deployed machinery, active exposure, and measured business results remain separate. [4–9]

16.1 Implemented foundation and subsequent hardening

The canonical envelope, governed ingestion, graph projector, validation and policy services, DCP, approval routing, action runner, audit replay, and governed model dispatch are implemented. Growth Control adds the six Decision Jobs, ValidationPassport packages and chain verification, pilot-state and report machinery, retention-aware intent modules, and frontier analysis paths.

The later waves and audit work added controls for freshness, action-bound authorization, replay and one-time-use behavior, configuration honesty, provider boundaries, economic evidence, and default-off activation. They also introduced origin classification, read-oriented MCP access, MMM export, decision metering, passport specification, and truth-gate packaging. Each component retains its own configuration and release state. RUN PACK v5 records the lane-freshness repair and its UI deployment run; its close report explicitly does not claim authenticated live verification for that run.

16.2 Frontier capabilities and remaining evidence

F1 Creative unbundling — IN BUILD. Implemented analysis supports provisional observational contrasts. It must not be described as a fully validated joint multimodal causal decomposition. Promotion requires sufficient creative coverage, suitable identification, uncertainty assessment, and pilot evidence.

F2 Multi-quarter customer value — IN BUILD. The machinery supports value-regime analysis. Findings require at least two closed, observed quarters; partial quarters do not count and findings do not extrapolate beyond the observed horizon; the proposal and bid paths require measured-effect provenance and a source reference. Assumptions and scenario baselines are not measured effects.

F3 Supply-chain and inventory synchronization — IN BUILD, observe-only. Inventory recommendations and spend-gating analysis exist. A recommendation is not an order, inventory mutation, or executed advertising change. A real approved execution path remains a separate requirement.

F4 Micro-geo calibration — LIVE for the scoped calibration machinery. Repository records describe a bounded self-pilot as live-armed. Every reservation remains approval-gated; the recorded clean-cycle count is zero, and the actuator remains deliberately unregistered. This state does not establish autonomous spend, general availability, or causal business lift. An unconfigured tenant lane stays dark.

F5 Treasury-gated spend governance — IN BUILD. Where configured and armed, treasury constraints, transactional reservations, and version checks bind admission through execution and settlement. Missing reservations, constraint-version drift, and unavailable configured sources refuse execution. An absent or unarmed tenant configuration does not establish treasury protection. The current policy path can continue without the optional treasury gate, so deployment must explicitly prove that the required controls are armed before claiming them. Stale cash positions do not support proximity tightening; absolute covenant caps remain distinct.

Edge inference and dosage optimization. The repository contains frontier machinery and gated paths; a configured environment flag alone does not prove a usable tenant lane. Sub-10ms inference remains a design target until a workload-specific benchmark supports it. Dosage proposals require support across the evaluated grid, uncertainty,

appropriate cross-fitting, and economics; correlation is not causal optimization.

16.3 Measurement and activation boundaries

The causal proof core remains PARTIAL. Intent Engine v2 remains IN BUILD or shadow at its component boundaries. Measurement writeback and Net Yield writeback remain OFF under explicit guard tests. Net Yield's order economics and dry-run machinery must not be presented as a live bidder feedback loop.

The design-partner recruitment and readout workflows have advanced in code. The latest reviewed records still leave the real partner and qualified readout outstanding. No broader pilot-result, production-accuracy, or customer-lift claim follows from those workflow implementations.

16.4 Next acceptance gates

The next operational milestones are source-specific configuration and fresh lineage proofs; tenant economic and treasury inputs; admissible forward outcomes and holdout evidence; successful scoped verification of serving revisions; and a qualified partner pilot readout. Any actuator promotion additionally needs all required proofs, human approval, and an execution adapter with verified rollback.

Production hardening, tenant isolation, data integrity, incident readiness, and onboarding take precedence over adding unsupported architecture. The living Open Items register owns the remaining tasks; this paper summarizes the acceptance logic rather than creating a competing work queue.

17 Questions from operators and buyers

What is usable now? The repository contains the governed decision foundation and extensive media and commerce machinery. Availability for a customer depends on the service, credentials, tenant configuration, evidence, and authorization state described in this paper.

Does a live pilot mean autonomous spending? No. The recorded micro-geo pilot concerns bounded calibration machinery with approval-gated reservations. Writebacks remain off and actuator promotion remains separate.

How are models controlled? Registry-based routing, explicit fallback, source grounding, full validation, and action-bound authorization constrain the operating process. Model confidence cannot override a hard gate.

How is performance established? Through reproducible benchmarks or a qualified engagement with an appropriate registered measurement design, outcome window, uncertainty, and traceable sources. A demo, synthetic fixture, or green build establishes a different kind of evidence.

Can the system override a financial or legal constraint? No. Forbidden autonomy and advisory-only boundaries apply even when a proposed alternative looks economically attractive.

18 Claim discipline

Every platform performance figure carries one of five labels: **verified result**, **benchmark result**, **pilot result**, **design target**, or **illustrative scenario**. The label is tied to its source, population, protocol, time window, and limitation. There is no separate pilot-target label.

Capability status is a separate axis. The Growth Control status record uses LIVE, PARTIAL, IN BUILD, and PROPOSED; research work is explicitly identified as RESEARCH. LIVE can describe scoped machinery without certifying autonomous authority or customer performance.

The platform ceiling remains **built, pre-benchmark**. Specific historical engineering proofs do not lift that ceiling for every service or business domain. External market forecasts are attributed to their publisher, rather than being presented as platform results. Public evidence also preserves the distinction between self-pilot and qualified design-partner results.

19 Glossary

Operating Knowledge Intelligence: governed evidence, memory, reasoning, decision controls, and outcome learning connected to business work.

SRPVDAL: Sense, Reason, Plan, Validate, Decide, Act, Learn.

Canonical Event Envelope: the versioned source, identity, integrity, time, and assessment contract used for evidence ingestion.

Operating Knowledge Graph: the Firestore-backed operating memory linking entities, evidence, policies, decisions, actions, and outcomes.

Decision Control Plane: the governed pathway for eligibility, policy, approval, and action authorization.

DEL Score: the Decision Eligibility Layer score, subject to an 80.0 platform floor and non-bypassable controls.

ValidationPassport: the complete validation evidence associated with a proposal; missing required checks make it invalid.

Growth Decision Graph: the evidence-bearing relationships supporting Growth Control decisions and their outcomes.

Intent Engine v2: the customer-facing anticipatory capability within Causal Growth Control.

High-Value Decision Jobs: the six operating jobs that organize Growth Control work and its evidence.

iROAS: incremental revenue divided by the relevant incremental advertising spend; its estimand and measurement design must be explicit.

Net Contribution Yield: order economics after the defined costs and adjustments; any bidder feedback remains separately governed.

Replay: reconstruction of the recorded decision and its evidence without repeating external execution.

Live-armed: configured for a specified operating lane; not synonymous with completed exposure, successful cycles, or verified lift.

Closing perspective

MIZ OKI 3.5 has progressed from a July architectural narrative to a more complete implemented decision system. Its next proof is disciplined operation: current evidence, clear authority, controlled action where earned, and measured outcomes that withstand review. The same rules that govern the software govern this whitepaper's claims.

Sources and revision basis

The internal evidence review is pinned to MIZOKICloudRun commit [ea5e69c97bc1e785dcb30740de47a4f93626546e](#). Repository source and recorded operational evidence support this edition; a new platform-wide authenticated deployment audit was not performed. The external positioning paper remains separate from the master roadmap r3.6 and the marketing-automation paper r2.0.

1. MIZOKICloudRun: [CONSTITUTION.md](#), [OPERATING_SYSTEM.md](#), [GOVERNANCE.md](#), [AGENTS.md](#), [TRUTH.md](#), and the governed value and shape registries.
2. MIZOKICloudRun: architecture Cell Registry and Virtuoso [WIRING.md](#); canonical ingestion and graph projection implementation.
3. MIZOKICloudRun: canonical platform skill and graph, connector, provenance, and identity contracts.
4. MIZOKICloudRun: [OFFERING_MAP.md](#), especially Growth Control status and the September naming ruling; [CANON_STATUS.md](#).
5. MIZOKICloudRun: Growth Control unified system r2.0, master whitepaper r3.6, pilot playbook, and Decision Jobs configuration.
6. MIZOKICloudRun: intent evaluation records, privacy lock, forward-label requirements, and promotion gates.
7. MIZOKICloudRun: [RUN_PACK_V5_CLOSE_2026-09-17.md](#) including its September closeout; WO-30 freshness contract and UI adapters.
8. MIZOKICloudRun: [WAVE3_STATE.md](#), [WAVE4_STATE.md](#), [OPEN_ITEMS.md](#), and frontier source and tests.
9. MIZOKICloudRun: tenant onboarding checklist, current priorities, CMEK operator runbook, authentication provisioning and launch-readiness records, and Decision Studio launch receipt.
10. July Google Doc: MIZ OKI 3.5 — Operating Knowledge Intelligence Final Updated Whitepaper, last edited 10 July 2026; August external v3.0 and its truth-reconciliation record.
11. MIZOKICloudRun: whitepapers README decision distinguishing the external paper and master plan as separate active topics.
12. [Grand View Research — Decision Intelligence Market](#).
13. [NIST — Finalized post-quantum standards](#).
14. [European Commission — AI Act implementation and AI Omnibus](#).
15. [Palantir — Ontology action types](#).