

# MIZ OKI 3.5™

## Operating Knowledge Intelligence for the Autonomous Enterprise

### **Master Positioning Document — Enhanced Final**

*External-source review current through July 10, 2026*

## EXECUTIVE SUMMARY

MIZ OKI 3.5™ is an Operating Knowledge Intelligence platform: a governed decision-infrastructure layer for organizations that must move beyond dashboards, disconnected 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 one complete operating loop — Sense, Reason, Plan, Validate, Decide, Act, Learn (SRPVDAL).

The platform is not a dashboard, chatbot, reporting layer, or static knowledge graph. It is the nervous system of the enterprise: the connective tissue that senses what changed, reasons about why, plans what can be done, validates what should not, decides what must be approved, acts within governance, and learns from every outcome. MIZ OKI is deliberately positioned as a nervous system — not a “brain” that replaces human judgment, but the governed pathway that carries signal to safe action.

Five product primitives define the platform: the Canonical Event Envelope (every signal becomes provenance-backed evidence); the Temporal-Causal Knowledge Base (the operating memory); the Domain Intelligence Cell (the domain implementation unit); the Decision Control Plane (the trust mechanism that gates every action); and the Immutable Learning Ledger (the compounding asset).

A governing architectural rule separates Domain Intelligence Cells from Horizontal Platform Services. Cells own domain objectives, domain reasoning and planning, domain validation, eligible decisions, and approved actions. Horizontal services own reusable ingestion and grounding, identity resolution, provenance, evaluation, routing, connector boundaries, observability, attestation, and replay. Advertising, bidstream, email, legal/policy, and financial intelligence are illustrative cell deployments, not a fixed taxonomy; MIZ OKI remains domain-agnostic by design.

# 1. THE ENTERPRISE PROBLEM

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

The result is a decision-velocity crisis: signals are detected late, causes are debated manually, plans live in spreadsheets, policies are checked after the fact, approvals are decoupled from evidence, actions are executed without measurement design, and lessons are lost when people move on. MIZ OKI turns fragmented signals into governed operating knowledge.

- Data without context. A campaign performance drop may look like a media issue when the real cause is conversion-tracking drift, landing-page latency, inventory constraints, margin, attribution delay, or a competitive move. Dashboards expose the symptom; MIZ OKI is built to identify the operating cause.
- Automation without judgment. Traditional automation asks whether a workflow can execute. MIZ OKI asks whether an action is justified, safe, governed, measurable, and reversible.
- 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 organization compounds intelligence over time.
- Governance outside the workflow. MIZ OKI embeds policy inside the loop, so it is checked before action, not after damage.

## 2. THE 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.

Analysts increasingly recognize decision intelligence as a distinct enterprise category. Grand View Research's June 2026 update estimates the global market at USD 17.8 billion in 2025, USD 20.7 billion in 2026, and USD 53.2 billion by 2033, representing a 14.4% CAGR from 2026 to 2033; North America accounted for more than 44% of 2025 revenue. The emergence of dedicated market reports and vendor-evaluation frameworks indicates that buyers and boards increasingly treat decision intelligence as an investable category rather than a feature.

MIZ OKI's category claim is narrower and sharper than generic decision intelligence: Operating Knowledge Intelligence is the governed operating layer that sits between models and real work — canonical evidence, temporal-causal memory, validation gates, decision eligibility, action controls, rollback, audit, and accumulated decision history. The moat is not any single model; it is the operating system around models.

## 3. THE CANONICAL SRPVDAL LOOP

SRPVDAL is the operating system of MIZ OKI 3.5. Every intelligence workflow moves through seven stages.

### 3.1 Sense

The perception layer collects data, documents, events, and human context from connected systems and converts them into canonical, provenance-backed evidence — advertising campaigns and conversions, OpenRTB bid/win/loss signals and consent, ESP events, CRM/finance/inventory/legal/support systems, and human approvals, objectives, and guardrails. The output of Sense is structured evidence, not raw data.

### 3.2 Reason

Evidence becomes understanding: temporal-causal graph traversal, GraphRAG and CausalRAG, root-cause and attribution analysis, entity resolution and identity stitching, anomaly detection, policy retrieval, document-grounded reasoning, and counterfactual reasoning. The output is a structured explanation with evidence-backed hypotheses — not yet a decision.

### 3.3 Plan

The system generates candidate interventions before deciding. Each plan specifies objective, expected impact, evidence used, assumptions, dependencies, risk level, cost/financial impact, confidence, reversibility, required approvals, measurement window, and learning criteria. Plan is where MIZ OKI becomes strategic rather than merely reactive.

## 3.4 Validate

No plan becomes a decision until it passes the appropriate gates: data freshness, completeness, schema compatibility, metric-definition validation, statistical significance, causal plausibility, confounder checks, financial/margin and budget guardrails, legal/policy checks, privacy and brand-safety checks, human-approval requirements, rollback readiness, and audit readiness. Validate is MIZ OKI's central safeguard against unsafe automation.

## 3.5 Decide

The Decision Control Plane ranks eligible options and records the selected option, rejected alternatives, evidence, validation results, risk and confidence scores, approval status, autonomy mode, expected outcome, measurement window, rollback logic, and audit trail.

## 3.6 Act

Execution happens only after eligibility is established — API calls, workflow triggers, approval packages, business-system updates, task creation, reporting, or rollback. Autonomy modes: observe only, recommend only, approval required, bounded autonomy, and full autonomy only for approved low-risk domains. New cells default to observe/recommend until evaluation gates, approval workflows, and rollback controls are proven.

## 3.7 Learn

After a decision window, MIZ OKI compares predicted and actual outcomes and updates causal confidence, planning assumptions, validation thresholds, decision policies, agent performance, connector health, metric definitions, playbooks, and the learning ledger. Without Learn, AI is a tool; with Learn, MIZ OKI becomes an operating intelligence system.

# 4. THE ADVANCED KNOWLEDGE BASE

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

- Universal evidence model. Every signal enters as a canonical event with tenant ID, loop ID, source system, connector, event type, event and ingestion time, raw-payload reference, normalized fields, entity candidates, metric values, provenance, governance labels, confidence scores, and audit IDs.
- Temporal-causal store. MIZ OKI stores entities, relationships, sequence, causality, evidence, uncertainty, policy, and outcomes — semantic, temporal, causal, financial, policy, identity, and learning relationships.
- 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.
- Immutable Learning Ledger. The ledger records the full 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.

Graph technologies — property graphs, semantic layers, vector search, and document retrieval — are deployment/projection options inside this broader temporal-causal architecture, chosen per workload rather than hard-coded.

## 5. DOMAIN INTELLIGENCE CELLS: EXAMPLE DEPLOYMENTS

The implementation unit is the Domain Intelligence Cell. A component qualifies as a cell only when it owns a defined business domain and declares its SRPVDAL contract: sensing adapters, canonical events, entity and graph mappings, domain metrics, reasoning skills, plan types, validation gates, decision-eligibility rules, approved actions, outcome evaluators, front-end visibility, and audit/learning records. Shared ingestion, retrieval, identity, evaluation, routing, credential, observability, and replay components are Horizontal Platform Services — not cells. The following cells are illustrative examples, not a fixed taxonomy.

### 5.1 Advertising / GAQL Intelligence Cell

The Advertising / GAQL Intelligence Cell is the flagship demonstration because it is concrete and high-value. It validates fields before query execution, normalizes rows into canonical events, persists evidence, maps campaigns/keywords/search terms/conversions/budgets/audiences/geographies/attribution into the knowledge base, generates optimization plans only after reasoning, validates through

financial/policy/statistical/causal/approval gates, and executes mutations only when approval and eligibility are satisfied. Default mode: observe/recommend.

## 5.2 OpenRTB Bidstream Cell

The OpenRTB cell reasons about inventory quality, auction duplication, floor-price effects, win-rate changes, identity resolution, fraud indicators, exchange performance, and supply-path optimization.

## 5.3 ESP / Email Cell

The ESP / Email cell reasons about deliverability, fatigue, segment quality, creative performance, send-time effects, revenue contribution, suppression risk, and cross-channel attribution. Actions are tightly gated because email affects trust, compliance, and deliverability.

## 5.4 Legal & Policy Cell

The Legal & Policy cell converts contracts, regulations, and internal rules into machine-checkable decision constraints, so MIZ OKI knows not only what the business wants to do but what it is allowed to do.

## 5.5 Financial Cell

The Financial cell reasons over revenue, margin, cost, cash flow, CAC, LTV, ROAS, payback, and forecasts. Financial validation sits inside Validate, so an action that lifts ROAS but destroys margin is not auto-approved.

## 6. HORIZONTAL PLATFORM SERVICES

### 6.1 Document & Research Knowledge Service

Document and research handling should not be owned by a single domain cell. It is a horizontal knowledge-ingestion, grounding, provenance, and retrieval service that every Domain Intelligence Cell can use.

This service ingests uploaded files, reports, meeting notes, research papers, product specifications, competitive intelligence, legal text, technical documentation, and internal strategy documents. It extracts structured evidence, entities, relationships, policies, claims, citations, source metadata, and reusable context into the advanced knowledge base.

Ownership remains with the relevant operating domain. A contract clause should feed the Legal & Policy Cell. A margin model should feed the Financial Cell. A campaign brief should feed the Advertising Cell. A product specification should feed the Product or Operations Cell. The horizontal document service does not decide or act on behalf of a domain; it supplies grounded evidence and memory for the cells that do.

This keeps MIZ OKI clean: domain cells own business reasoning, planning, validation, decisioning, and action; the document/research layer owns ingestion, grounding, provenance, retrieval, and reusable memory.

### 6.2 Identity, Entity & Semantic Resolution Service

This shared service resolves organizations, people, accounts, products, campaigns, assets, devices, documents, policies, and other business objects across heterogeneous sources. It maintains canonical identifiers, aliases, match keys, confidence, evidence, and unresolved ambiguity so individual cells do not create incompatible identity logic. Domain cells can propose merges, splits, and corrections, but accepted changes are versioned, reversible, and written to the audit and learning records.

### 6.3 Provenance, Bitemporal Lineage & Evidence Service

Every normalized fact and derived claim carries its source, observation time, effective time, ingestion time, transformation or query hash, confidence, policy classification, and supporting evidence references. Bitemporal lineage preserves both what was believed at a given decision

time and what was later determined to be valid in the world. This prevents hindsight leakage in evaluation, supports time-travel analysis, and makes deterministic decision reconstruction possible.

## 6.4 Evaluation, Policy & Guardrail Service

A central service owns reusable evaluation suites, policy rules, approval matrices, confidence thresholds, financial limits, fairness and privacy checks, test datasets, and fail-open or fail-closed behavior. Domain cells invoke these controls during Validate and again before Act. No cell may silently bypass a mandatory gate; exceptions require an explicit, attributable, time-bounded policy record.

## 6.5 Model, Tool & Agent Routing Service

MIZ OKI uses a swappable multi-vendor routing layer rather than a single-model dependency. Routing is based on task type, evidence sensitivity, jurisdiction, model and tool evaluations, latency, cost, context requirements, and required confidence. Every routing and fallback decision is versioned and included in the decision receipt. Degraded-mode behavior is explicit, so a vendor outage, price change, capability regression, or model deprecation cannot silently change the control posture or strand the platform.

## 6.6 Connector Runtime & Credential Boundary Service

Connectors execute inside governed tenant and credential boundaries. This service manages connector registration, scopes, secrets, rate limits, schema and API-version compatibility, retries, dead-letter handling, health checks, and read-versus-write authority. Write permissions are separate from sensing permissions and are exposed to Act only through the Decision Control Plane after eligibility is established.

## 6.7 Observability, Audit, Attestation & Replay Service

A common trace identity follows every loop from Sense through Learn. The service captures stage inputs and outputs, model and tool versions, policy decisions, approvals, action receipts, outcome windows, errors, and rollback events. Tamper-evident attestations support independent review, while deterministic replay reconstructs the historical decision from recorded evidence

and versions without pretending that a non-deterministic model can be re-executed into an identical answer.

These horizontal services are shared platform capabilities. They do not own domain objectives or business actions; they make every Domain Intelligence Cell consistent, governable, portable, and auditable.

## 7. PRODUCTION ARCHITECTURE

MIZ OKI 3.5 runs as a containerized microservices platform on Google Cloud Run. The horizontal services in Section 6 cross-cut the following operating layers:

1. Evidence layer — connectors, webhooks, API/batch jobs, document ingestion, raw-payload storage, canonical normalization, provenance tracking, source-health monitoring.
2. Knowledge layer — entity/relationship stores, temporal and causal graph structures, policy/identity/metric graphs, tenant memory boundaries, and the learning ledger, projected across Neo4j, BigQuery, and complementary stores.
3. Reasoning layer — GraphRAG, CausalRAG, temporal reasoning, entity resolution, anomaly detection, attribution, metric diagnostics, policy retrieval, and domain agents.
4. Planning layer — plan generation, scenario simulation, counterfactual reasoning with mandatory no-action baselines, optimization, risk modeling, and experiment design.
5. Validation layer — data-quality, statistical, causal, financial, policy, privacy, compliance, brand, human-approval, and rollback gates.
6. Decision Control Plane — eligibility, autonomy modes, approval state, ranking, risk/confidence scoring, audit records, escalation routing.
7. Action & Learning layers — API runners, workflow engines, mutate-service integrations, notifications, rollback handlers, execution monitors; then prediction-vs-actual tracking, causal-confidence and graph updates, model/agent evaluation, policy refinement, and ledger writes.

The architecture separates domain ownership from shared controls: cells own domain reasoning and action, horizontal services enforce common identity, evidence, evaluation, routing, credential, observability, and replay contracts across every layer.

## 8. COMMAND CENTER EXPERIENCE

The interface is a command center, not a dashboard: a Live Knowledge Graph (entities, evidence paths, causal links, confidence, temporal sequences, policy constraints); an

SRPVDAL Loop Monitor (every active loop by stage, status, owner, confidence, next required action); a Decision Queue (recommendations, validation status, risk score, expected value, required approvals, rejected alternatives, evidence, audit trail); a Simulation Console (interventions, counterfactual outcomes, financial impact, policy risk, reversibility, measurement plans); Channel Intelligence Pages; and an Audit & Learning Ledger capturing what happened, why, what was proposed, what passed and failed validation, who approved, what action occurred, what resulted, and what was learned.

## 9. GOVERNANCE AND RESPONSIBLE AUTONOMY

Governance is embedded in the loop. Every proposed action receives a decision-eligibility status: 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/compliance rule; blocked by low confidence; blocked by lack of rollback; or rejected. The system never simply says “do this” — it states 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.

## 10. SECURITY & COMPLIANCE POSTURE

- Cloud security. Tenant isolation, encryption in transit and at rest, and customer-managed keys via Cloud KMS; least-privilege access for both human users and non-human agent identities.
- Attestation and replay. MIZ OKI’s design intent is tamper-evident, cryptographically chained decision receipts written before action, supporting deterministic replay — the ability to reconstruct any past decision from the record without re-executing it. This pattern maps to EU AI Act Article 12 logging, ISO/IEC 42001, and the NIST AI Risk Management Framework.
- Post-quantum roadmap. For long-lived attestations, MIZ OKI targets NIST’s finalized post-quantum standards. NIST finalized FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) on August 13, 2024. MIZ OKI’s design direction uses ML-KEM for key establishment and ML-DSA for signatures, with SLH-DSA as a hash-based alternative. Migration milestones must be set per deployment, cryptographic inventory, data-retention horizon, and applicable government or industry requirements. PQC capabilities remain design intent unless verified in a specific deployment.
- Frameworks. The platform is engineered to support SOC 2, ISO/IEC 42001, GDPR obligations, and applicable EU AI Act controls. The European Commission’s current

implementation timeline states that the Act is generally applicable from August 2, 2026; GPAI obligations became applicable on August 2, 2025; transparency rules apply in August 2026; certain high-risk-area rules apply from December 2, 2027; and high-risk systems embedded in regulated products have a transition period to August 2, 2028. Deployment-specific legal review remains required.

- Intellectual property. MIZ OKI's differentiating methods — eligibility gating, counterfactual simulation with mandatory no-action baselines, cryptographic attestation, deterministic replay, and bitemporal provenance — are the subject of pending patent filings. No issued-patent claims and no patent numbers are asserted in this document.

## 11. COMPETITIVE DIFFERENTIATION

- Palantir Foundry / AIP. Palantir validates enterprise demand for ontology-backed operational AI, governed workflows, and action-oriented data platforms. MIZ OKI differentiates by shipping the complete SRPVDAL loop with eligibility gating, counterfactual no-action baselines, and deterministic replay as first-class primitives, while retaining a modular Google Cloud deployment model rather than requiring a single heavyweight platform commitment.
- Decision-intelligence platforms. Aera Technology, FICO, SAS, Quantexa, IBM, and other established vendors demonstrate the category's growing maturity. MIZ OKI differentiates through an explicit Validate stage and Decision Control Plane: auditable eligibility gating, no-action baselines, rollback, replay, and cross-domain operating contracts rather than a workflow limited to one suite or function.
- Salesforce Agentforce and Microsoft Copilot. These are powerful agent and automation layers anchored to large application ecosystems such as CRM, Microsoft 365, and Power Platform. Salesforce reported approximately USD 800 million in Agentforce annual recurring revenue and 29,000 deals in its FY2026 results; Microsoft reports broad enterprise adoption of custom agents across Copilot Studio and its wider agent platform. MIZ OKI is ecosystem-neutral, operates across the systems where business happens, and treats evidence, eligibility, governance, and learning as the product rather than a suite feature.
- Point AI tools and dashboards. These answer questions or show metrics; MIZ OKI governs decisions and compounds learning.

## 12. ROI FRAMING

MIZ OKI's value accrues in four measurable areas, each tied to a baseline and measurement window per the claim-discipline rules:

1. Decision velocity — time from signal to governed action, measured against the current manual baseline for a named workflow.
2. Risk avoidance — actions blocked by validation gates (policy, margin, rollback) that would otherwise have caused loss.
3. Outcome quality — prediction-vs-actual accuracy improving over time via the learning ledger.
4. Institutional memory — reduction in re-work and knowledge loss as decisions and their rationale become durable assets.

Every figure a customer publishes should be labeled verified, benchmarked, targeted, or illustrative.

## 13. ILLUSTRATIVE SCENARIO

The following is an illustrative scenario, not a customer result. A \$5.0M distribution is proposed. MIZ OKI senses the request and supporting evidence, reasons over policy and financial context, and plans alternatives. In Validate, the Decision Control Plane scores the action on the Decision Eligibility Layer at 41 — below the autonomy threshold — because of a policy/rollback constraint. The action is blocked, and the system re-routes to a governed \$3.2M alternative that clears validation, with the full evidence package, rejected alternatives, and rollback plan recorded to the Immutable Learning Ledger. This is the difference between automation that can act and operating intelligence that acts only when justified, safe, and reversible.

## 14. FLAGSHIP DEMO STORYLINE: SRPVDAL END-TO-END

A performance drop in an advertising account:

Sense ingests campaign, search-term, conversion, budget, asset, geography, device, landing-page, inventory, margin, and prior-decision data.

Reason localizes the drop to specific broad-match terms, mobile traffic, landing-page latency, and low-stock high-margin products, separating attribution delay from true degradation.

Plan generates options: negative keywords, budget reduction, spend shift, wait for attribution, fix latency, or route to approval.

Validate checks compatibility, freshness, conversion lag, significance, margin, inventory, budget/brand policy, approvals, and rollback.

Decide selects a staged plan.

Act executes only approved actions and prepares evidence where approval is required.

Learn compares predicted and actual outcomes and updates the cell's rules and thresholds.

## 15. DEPLOYMENT MODEL & INTEGRATION SURFACE

MIZ OKI deploys as a multi-tenant SaaS on Google Cloud Run, with per-tenant isolation and customer-managed keys. The integration surface spans API connectors, webhooks, batch loaders, and document ingestion; write-back to business systems occurs only through governed action runners under the Decision Control Plane. New domains onboard as cells in observe/recommend mode and graduate to bounded or full autonomy only after evaluation gates, approval workflows, and rollback controls are proven.

## 16. PRODUCTION ROADMAP

Phase 1 — Foundation: canonical envelope, tenant model, connector registry, raw event store, persistence, initial graph projection, loop records, audit ledger, command-center shell, flagship cell foundation, recommend-only autonomy.

Phase 2 — Reasoning & Planning: entity resolution, metric diagnostics, causal-path reasoning, GraphRAG/CausalRAG, counterfactual planning, plan generation, domain agents, simulation console, decision queue, validation framework.

Phase 3 — Validation & Decision Control: 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: approved API/mutate actions, workflow execution, action logs, rollback, execution monitoring, post-action measurement, learning records.

Phase 5 — Continuous Learning: prediction-vs-actual tracking, causal-confidence updates, playbook refinement, agent scoring, knowledge-base self-healing, cross-loop learning, executive reporting.

## 17. OBJECTION HANDLING

- How is this different from Palantir or Aera? MIZ OKI ships the governed loop itself with eligibility gating, counterfactual no-action baselines, and deterministic replay as primitives, ecosystem-neutral, on a lighter footprint.
- Isn't this just agents? Agents act; MIZ OKI governs whether an action is eligible, safe, reversible, and measurable before it acts, and learns from the result.
- What about model risk? A swappable multi-vendor routing layer avoids single-model lock-in and lets the platform track the frontier.
- Can we trust autonomous action? Default is observe/recommend; autonomy is earned per domain after gates, approvals, and rollback are proven, with tamper-evident audit and replay throughout.
- Is our data safe? Tenant isolation, customer-managed keys, least-privilege identities, and no cross-tenant memory bleed.

## 18. CLAIM DISCIPLINE

Public claims remain bold but disciplined. Every metric is labeled verified result, benchmark result, design target, pilot target, or illustrative scenario. Decision-velocity, self-healing, entity-resolution, and financial claims are tied to specific workflows, baselines, benchmark datasets, or scenario evidence. Industry stories are labeled examples unless they are verified deployments. No absolute guarantees or unsubstantiated superiority claims are made, and no issued-patent claims are asserted.

## 19. GLOSSARY

- Operating Knowledge Intelligence — the governed operating layer between models and real work.
- SRPVDAL — Sense, Reason, Plan, Validate, Decide, Act, Learn.
- Canonical Event Envelope — the standard evidence record for any signal.

- Domain Intelligence Cell — the production unit that owns a defined business domain and carries its evidence through declared SRPVDAL stages, including eligible decisions and governed actions where authorized.
- Horizontal Platform Services — shared identity, evidence, evaluation, routing, connector, observability, audit, attestation, and replay capabilities consumed by every cell.
- Document & Research Knowledge Service — the horizontal service that ingests, grounds, retrieves, cites, and stores document/research evidence for all cells without owning domain-level decisions or actions.
- Decision Control Plane — the mechanism that ranks, gates, approves, delays, escalates, or executes decisions by eligibility and governance.
- Decision Eligibility Layer (DEL) — the scoring and gating component that determines whether an action may proceed autonomously.
- Immutable Learning Ledger — the tamper-evident record of every decision's life and learning.
- GraphRAG / CausalRAG — retrieval-augmented reasoning over knowledge and causal graphs.
- Deterministic replay — reconstructing a past decision from its record without re-executing it.

## CLOSING STATEMENT

MIZ OKI 3.5 is the operating intelligence layer — the nervous system — for the autonomous enterprise. It turns fragmented business signals into governed knowledge, knowledge into validated decisions, decisions into controlled action, and outcomes into institutional learning. 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 — NOTE ON SOURCES AND VERIFICATION

External-source review current through July 10, 2026. Market sizing was updated to the June 2026 Grand View Research report. Post-quantum standards were checked against NIST's finalized FIPS 203, FIPS 204, and FIPS 205 announcement. EU AI Act dates were updated to the European Commission's current implementation page, including the May 2026 political agreement on revised high-risk timelines. Salesforce adoption figures should be tied to its FY2026 investor materials before external publication. Platform-specific assertions — architecture, pending patent filings, the illustrative USD 5.0M / USD 3.2M scenario, DEL thresholds, and product roadmap — remain company statements or labeled scenarios. No issued-patent claim or patent number is asserted.

## PRIMARY SOURCE REGISTER

- [\*\*Grand View Research\*\* — Decision Intelligence Market Size & Share Report, updated June 2026.](#)
- [\*\*NIST\*\* — First Three Finalized Post-Quantum Encryption Standards, August 13, 2024.](#)
- [\*\*European Commission\*\* — AI Act regulatory framework and implementation timeline, current July 2026.](#)
- [\*\*Salesforce Investor Relations\*\* — FY2026 financial and quarterly materials.](#)