

DataOcean

The Enterprise Data Platform
Powering the BizFirstAi Ecosystem.


THE DATA PLATFORM DIAGNOSTIC

FEATURE	TRADITIONAL OPERATIONAL DBS	DATA WAREHOUSES	FRAGMENTED SAAS SILOS	DATAOCEAN
AI-READINESS	Disjointed ETL pipelines	Batch processing only	Siloed external API calls	✓ Native LLM context & agent history
MULTI-TENANCY	Requires custom sharding	Not designed for operational multi-tenant	Varies widely by vendor	✓ Built-in 100k+ tenant isolation
REAL-TIME EVENTS	Trigger-based batching	High-latency ETL	Webhooks only	✓ Native EdgeStream pub/sub
AUDITABILITY	Fragmented logging	Historical but not real-time	Inconsistent compliance	✓ 100% immutable audit trails for HIPAA/SOC2

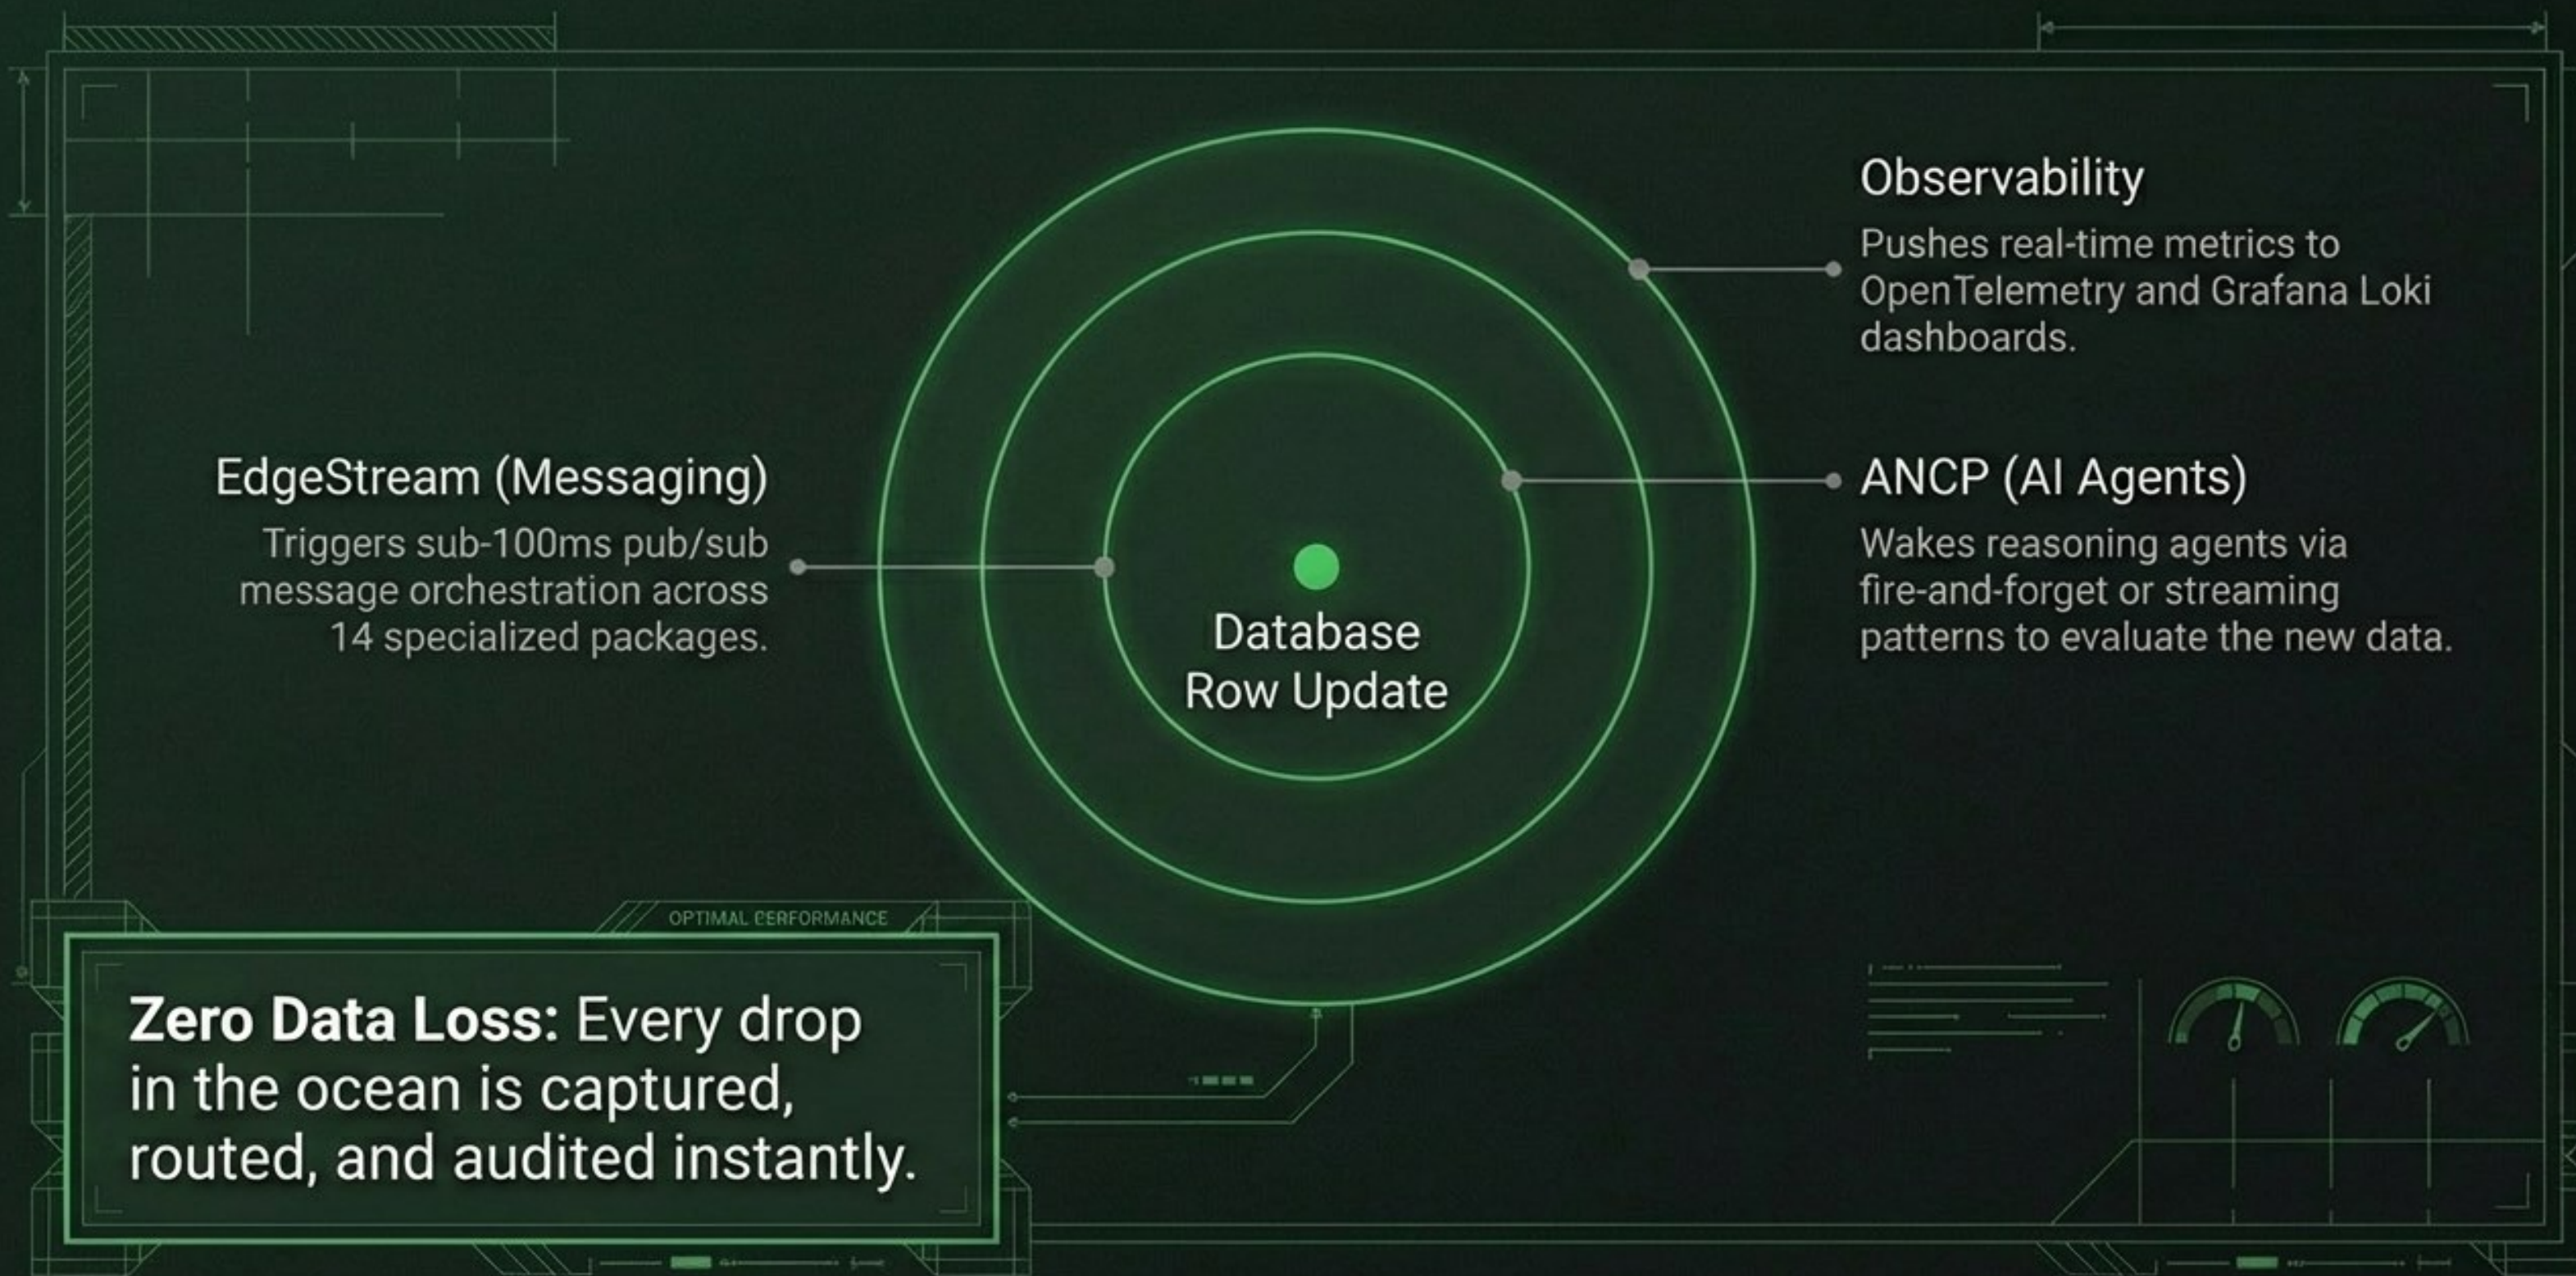
OPTIMAL PERFORMANCE

DATAOCEAN BRIDGES THE GAP BETWEEN OPERATIONAL SPEED, ANALYTICAL DEPTH, AND AI READINESS.

THE SINGLE SOURCE OF TRUTH



THE EVENT-DRIVEN RIPPLE



EdgeStream (Messaging)
Triggers sub-100ms pub/sub message orchestration across 14 specialized packages.

Database Row Update

Observability
Pushes real-time metrics to OpenTelemetry and Grafana Loki dashboards.

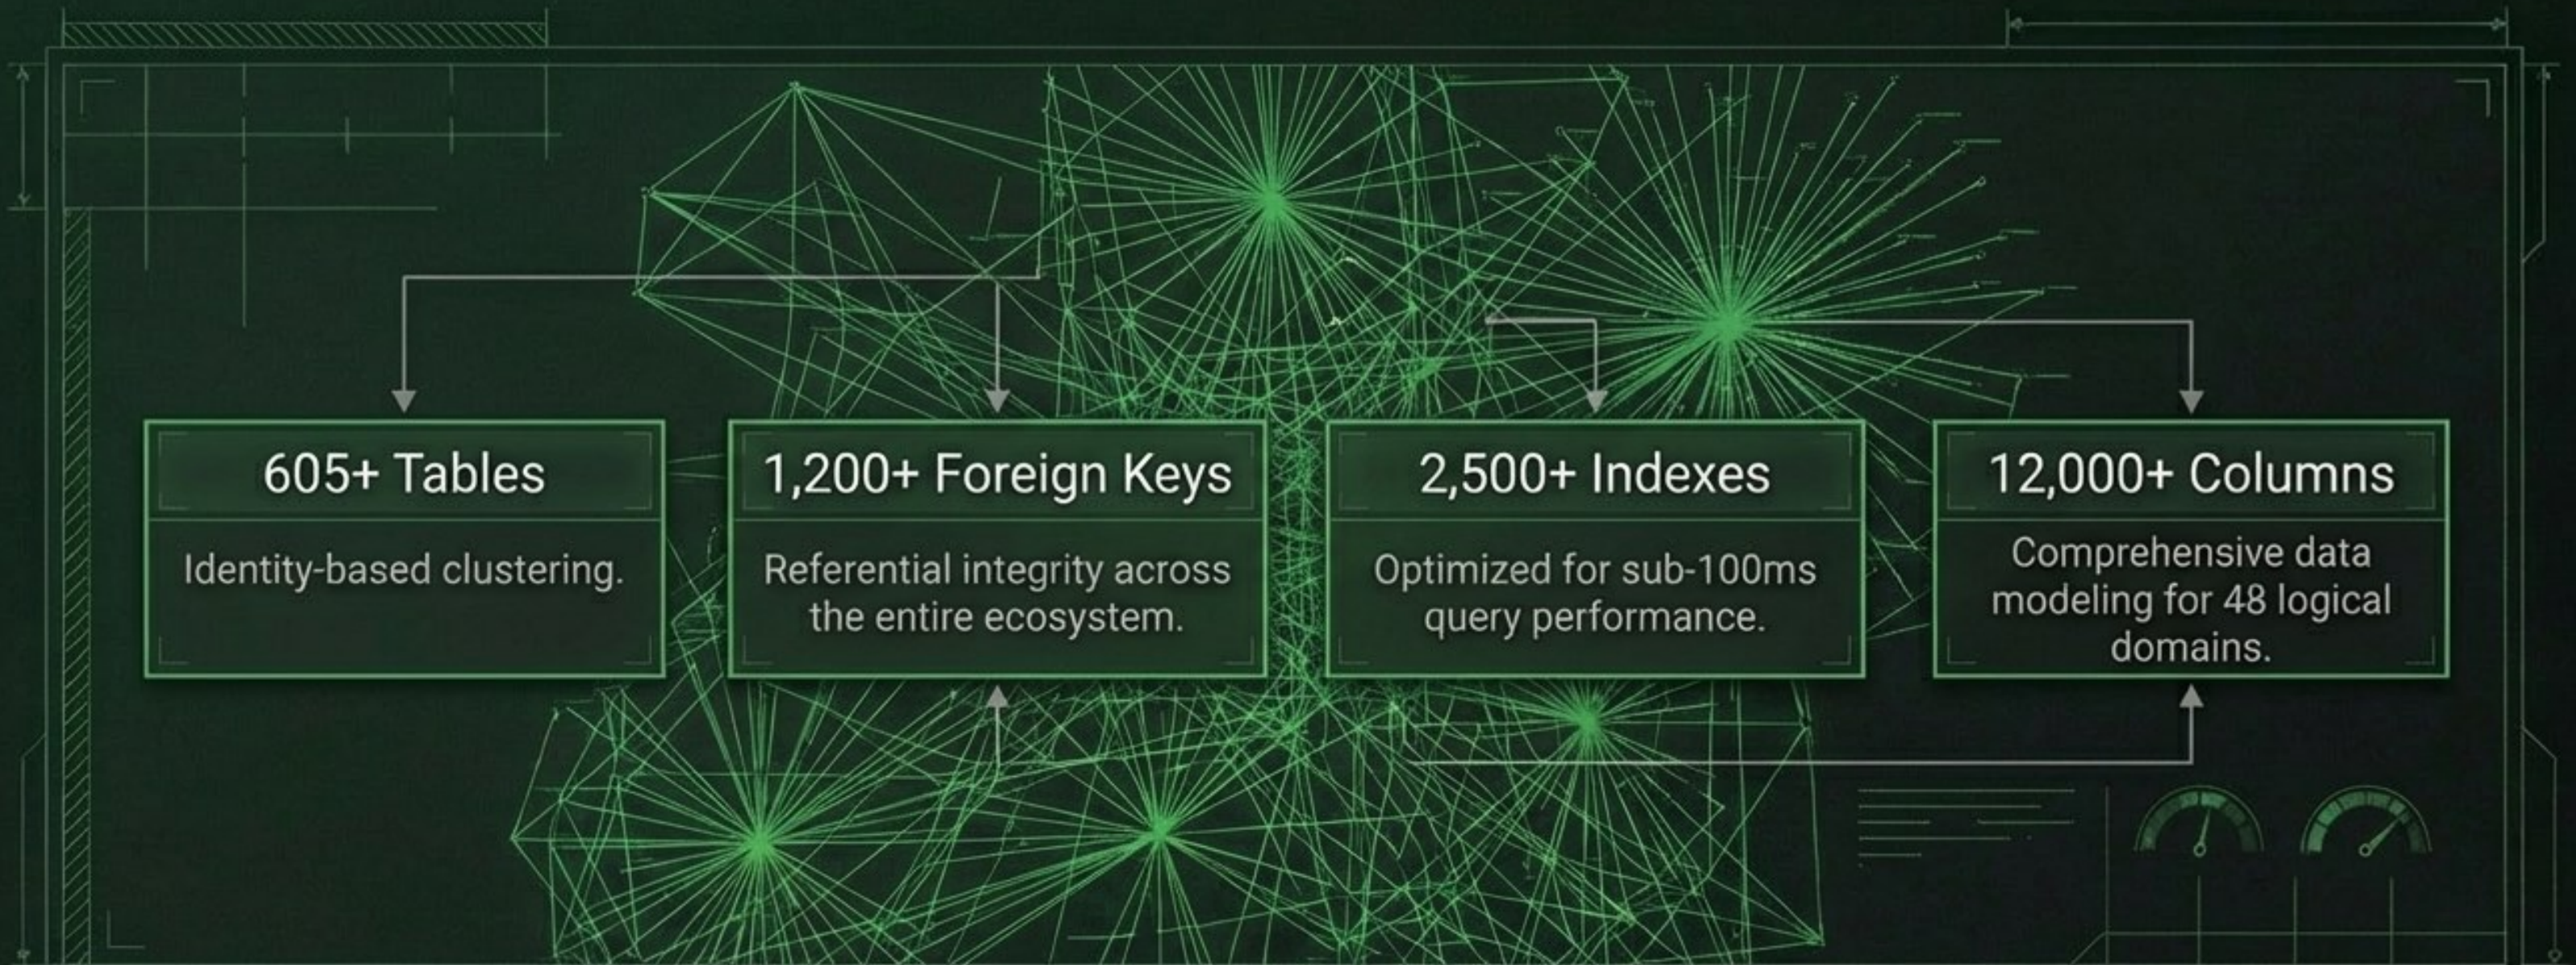
ANCP (AI Agents)
Wakes reasoning agents via fire-and-forget or streaming patterns to evaluate the new data.

Zero Data Loss: Every drop in the ocean is captured, routed, and audited instantly.

OPTIMAL PERFORMANCE

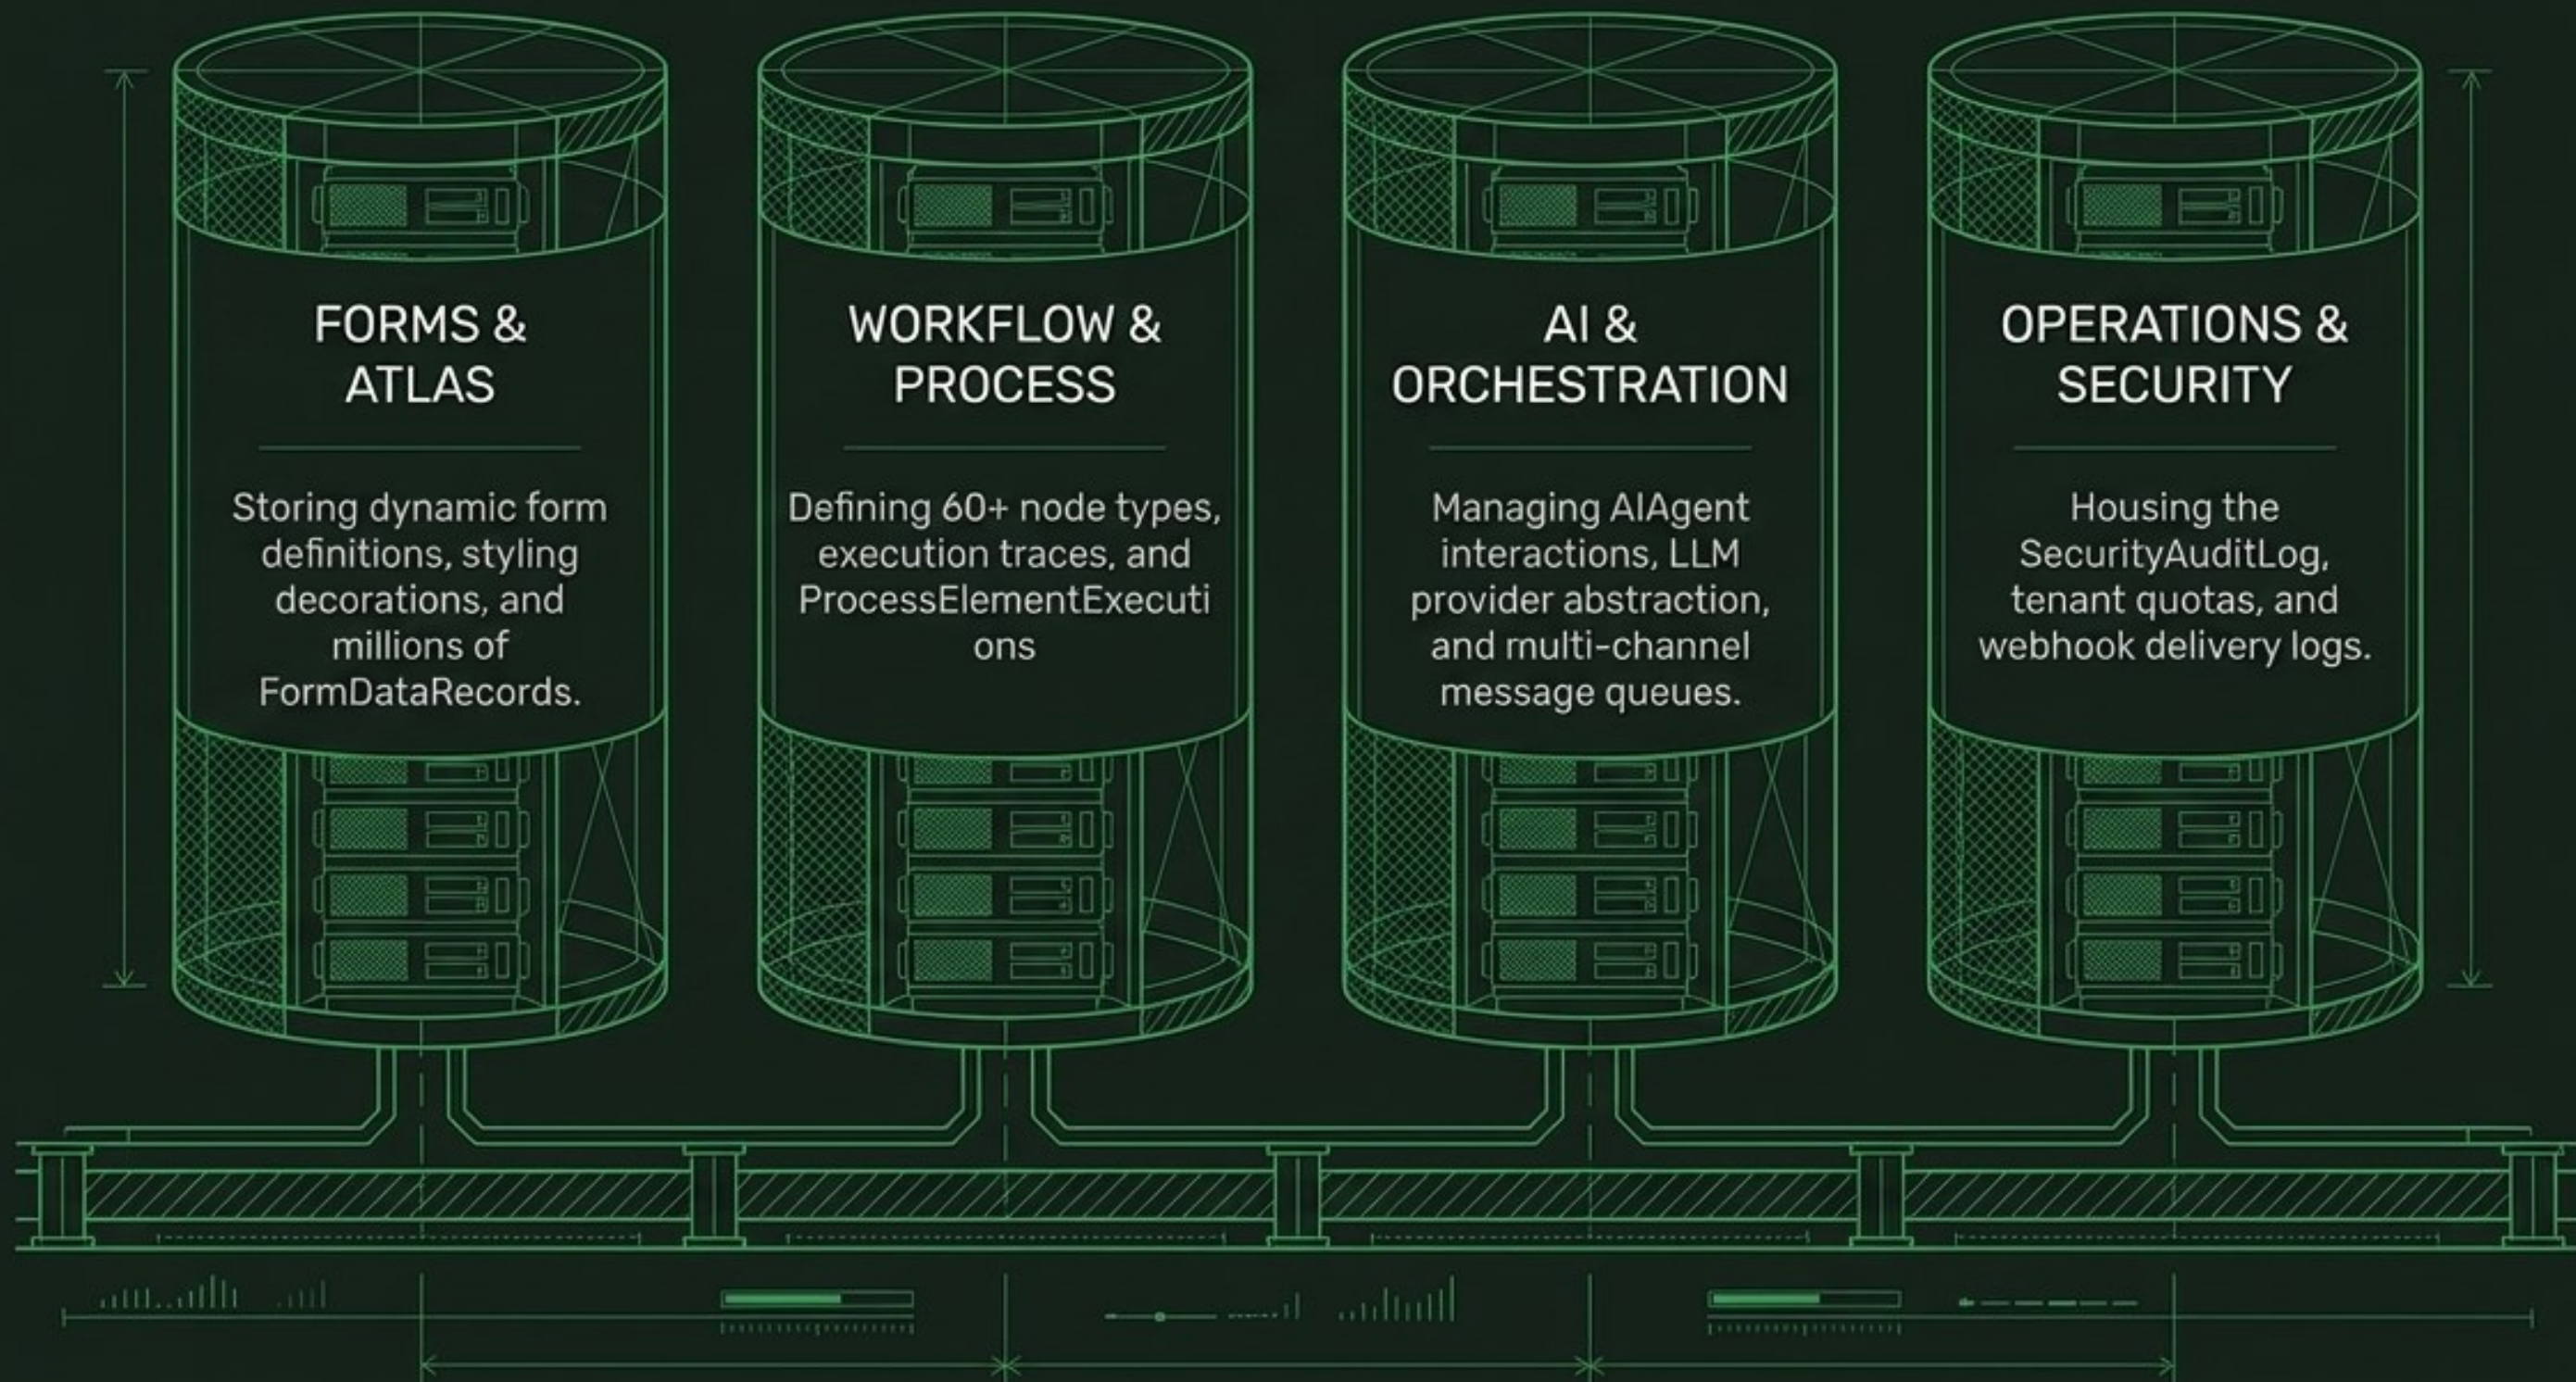


ANATOMY OF AN OCEAN



Not just a database—a highly normalized, ACID-compliant relational matrix designed for high-volume SaaS operations.

THE FOUR PILLARS OF THE SCHEMA



THE STANDARDIZED TABLE ANATOMY

TABLE BLUEPRINT (STANDARD SCHEMA)

SECTOR 1: BUSINESS LOGIC

The dynamic payload (JSON schemas, user data, workflow definitions).



SECTOR 2: SYSTEM COLUMNS


Emphasizing **TenantID** (mandatory for isolation) and **Deleted** (The Soft Delete Pattern).

TenantID	Deleted	

SECTOR 3: AUDIT COLUMNS

Highlighting CreatedOn, LastModifiedOn, and EF Core's RowVersion for optimistic concurrency.

CreatedOn	LastModifiedOn	RowVersion
2021-01-12 10:50 AM	2021-01-12 10:00 AM	2021-01-12 10:00 AM

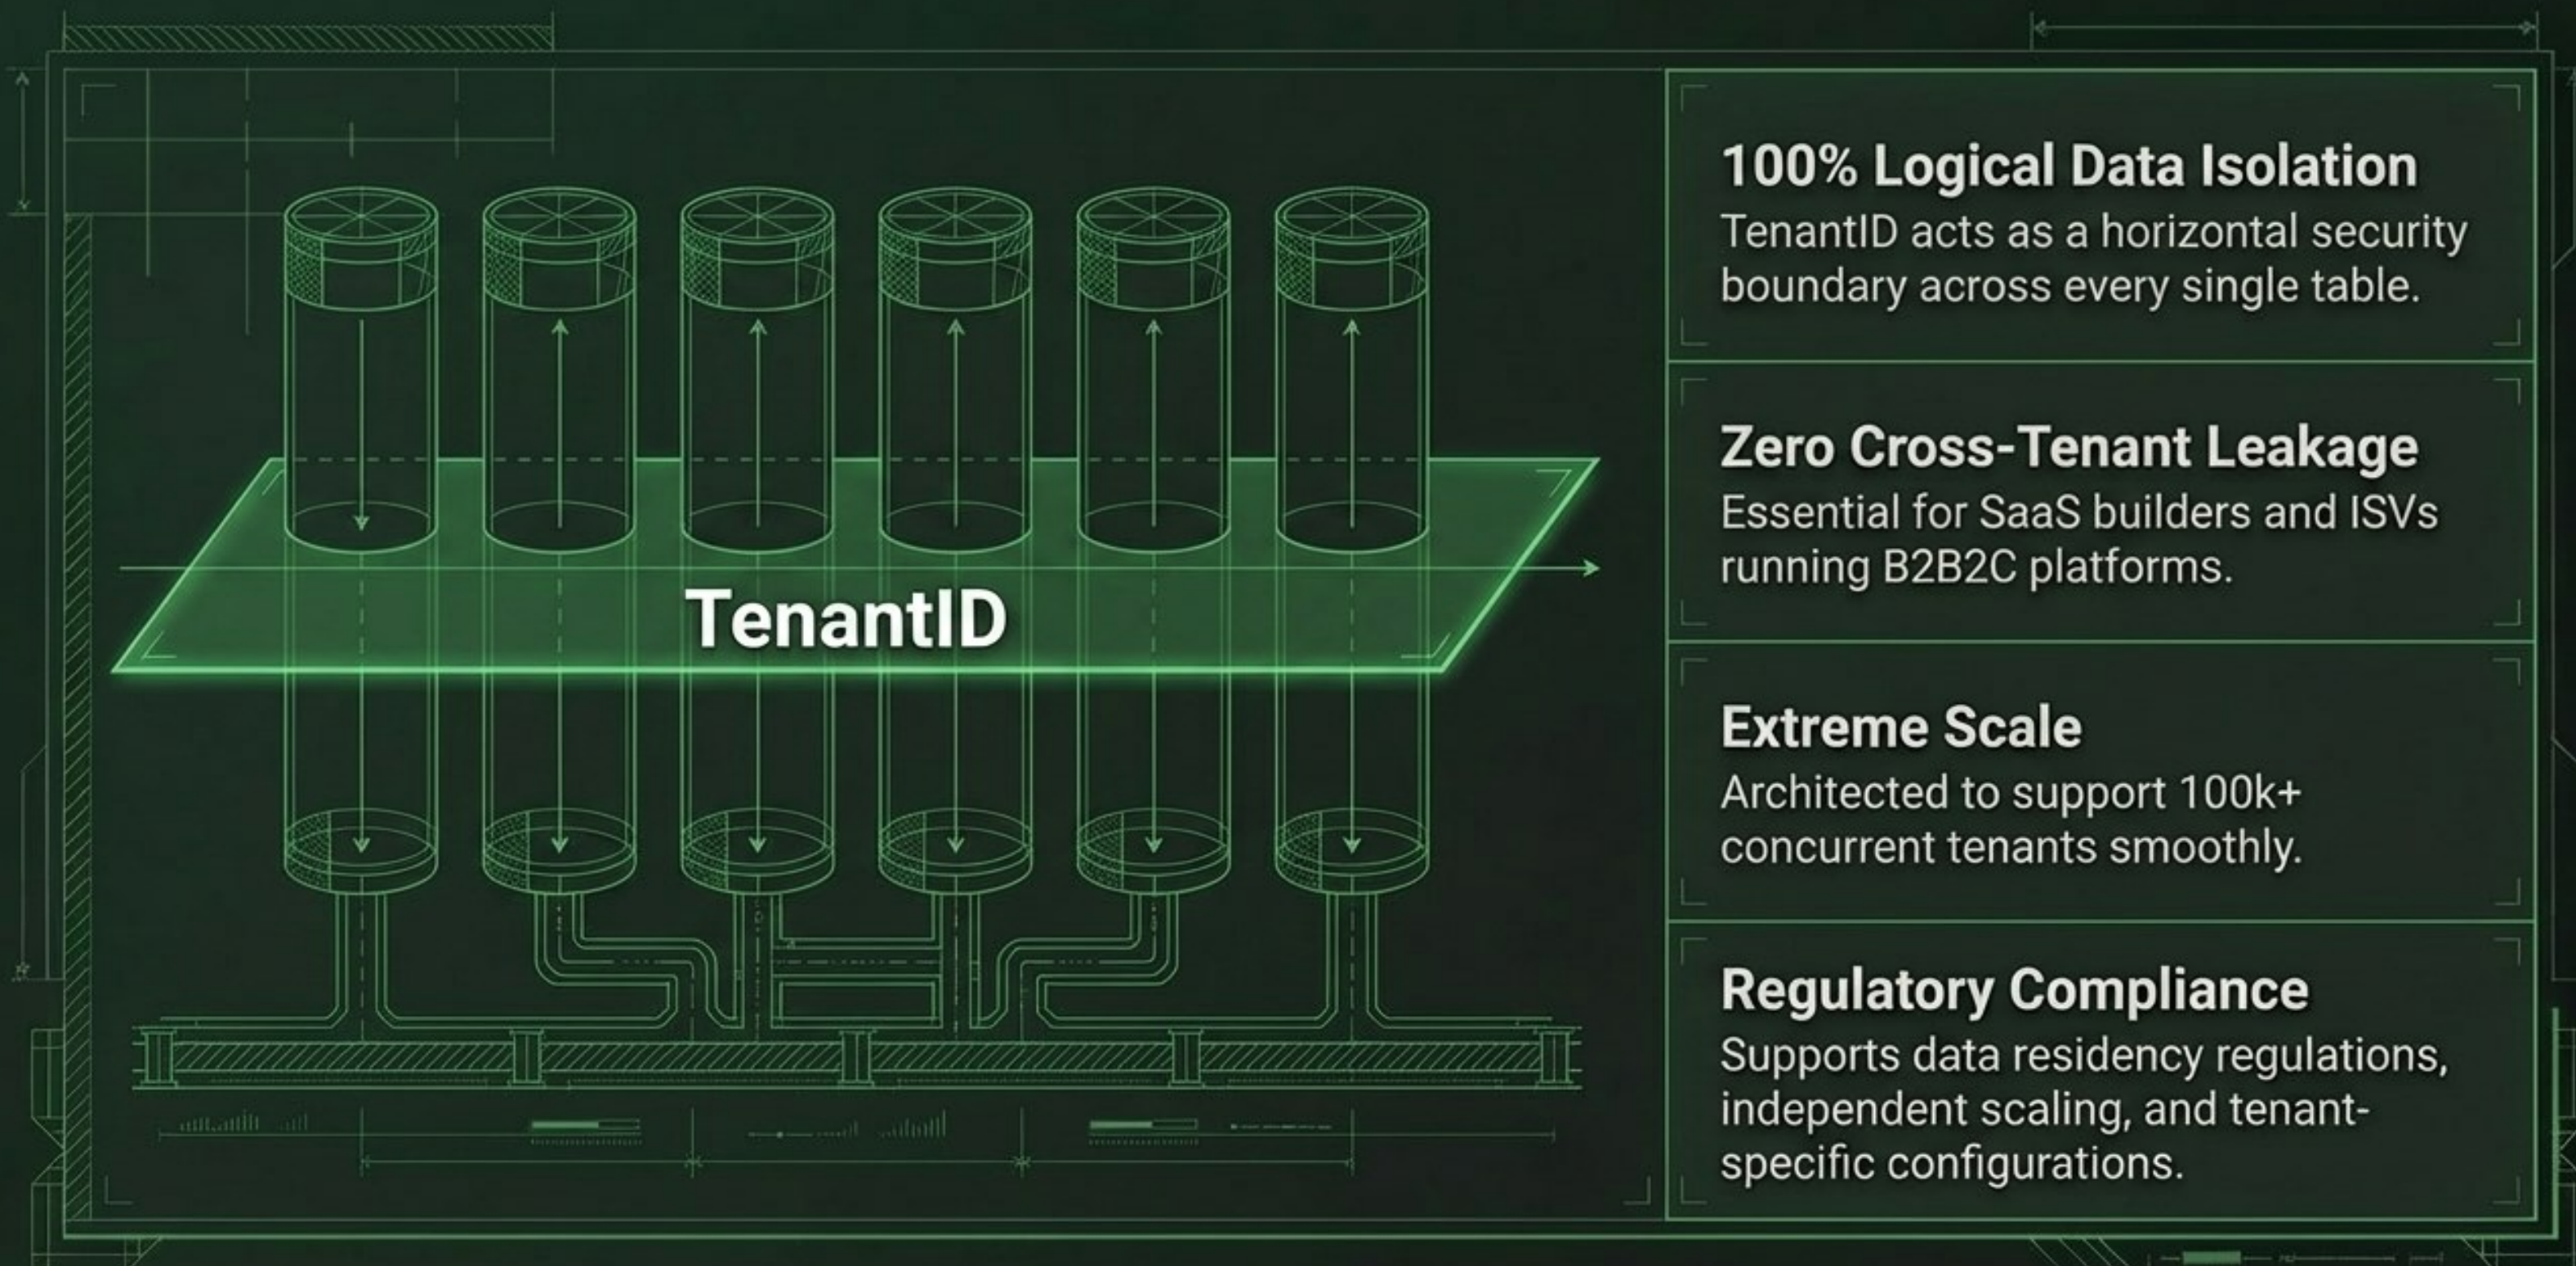


Tenant	Type	Row
0000	0000	0000000000
0001	0001	0000000001
0002	0002	0000000002
0003	0003	0000000003
0004	0004	0000000004
0005	0005	0000000005
0006	0006	0000000006
0007	0007	0000000007
0008	0008	0000000008
0009	0009	0000000009
0010	0010	0000000010



THE SOFT DELETE PATTERN: Data is never hard-deleted. 'WHERE Deleted = 0' preserves historical context for GDPR, HIPAA, and SOX compliance.


The Tenant Vault: Multi-Tenant Isolation



ACTIVE SECURITY: THE RLS SIEVE



UNFILTERED DATA STREAM
STATUS: RAW INPUT | VOLUME: HIGH



THE RLS SIEVE
TENANTID FILTERING LAYER



FILTERED TENANT DATA
STATUS: SECURED | ACCESS GRANTED

ROW-LEVEL SECURITY (RLS)

Every query automatically injects
`WHERE TenantID = @CurrentTenantID.`

PROCESSECURITY

A 4-layer, sub-millisecond authorization system (Tenant -> Process -> Node -> Element) enforcing fail-closed semantics.

IMMUTABLE AUDIT TRAILS

Permanent logging of all access decisions, providing complete forensic capability.

Operational Telemetry & SLAs



Single Record Lookup

Primary key hash



User Permission Check

Intelligent caching



Workflow Execution Trace

Execution ID index



Complex Join (5-10 tables)

Optimized via 2,500+ indexes

Multi-level caching (Redis, in-memory) and robust indexing guarantees sub-100ms query response times on common operations.

Future-Proof Scaling Diagnostics

Year 1: Horizontal Scaling

Multi-Instance SQL DB

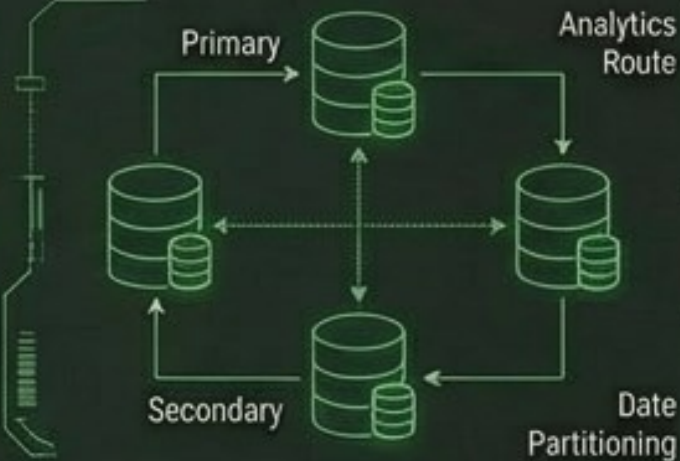


100K executions/day

Distributes load across servers while maintaining multi-tenant isolation.

Year 2: Read Replicas & Archival

Always-On Availability Groups



1M executions/day

Routes analytics and reporting queries to secondary instances; partitions tables by date.

Year 3+: Federated Multi-Region

Distributed Kafka Sync



10M+ executions/day

Time-series databases for metrics; local registry shards for global data residency.

The Unified Intelligence Flywheel

