

Cloud-Native Core Systems: Database & DevOps for Insurance

How microservices, containerization, CI/CD, and modern observability transform legacy insurance platforms into scalable, resilient, DevOps-driven architectures.

ENTERPRISE ARCHITECTURE

PLATFORM ENGINEERING

DEVOPS

CONF42 DATABASE DEVOPS 2026



The Modernization Imperative

Insurance organizations are under compounding pressure from multiple directions simultaneously. Regulatory bodies demand faster compliance reporting, digital-native competitors ship features in days rather than quarters, and catastrophic weather events trigger sudden claims spikes that legacy platforms cannot absorb without degraded service or outright failure.

Regulatory Pressure

Compliance mandates tighten while legacy systems make audit trails fragmented and change windows narrow.

Digital Competition

Insurtech entrants deploy cloud-native stacks from day one, compressing delivery cycles and raising customer expectations.

Catastrophe Scalability

CAT events generate sudden, massive claims volume. Monolithic systems cannot elastically scale, creating backlogs and customer friction.

The Legacy Constraint: Tightly Coupled Systems



Slow Release Cadence

Monolithic coupling means every change touches everything—forcing long regression cycles and infrequent deployments.

Limited Agility

Business rule changes in underwriting or claims require coordinated releases across application, database, and integration layers.

Scalability Ceiling

Vertical scaling of monolithic cores is expensive and bounded. Horizontal scaling requires fundamental architectural change.

Microservices: Decoupling Insurance Data Domains

Microservices architecture breaks monolithic insurance systems into independently deployable services, each owning its bounded data domain.



Policy Service

Owns policy lifecycle data. Exposes versioned APIs. Scales independently during open enrollment periods.



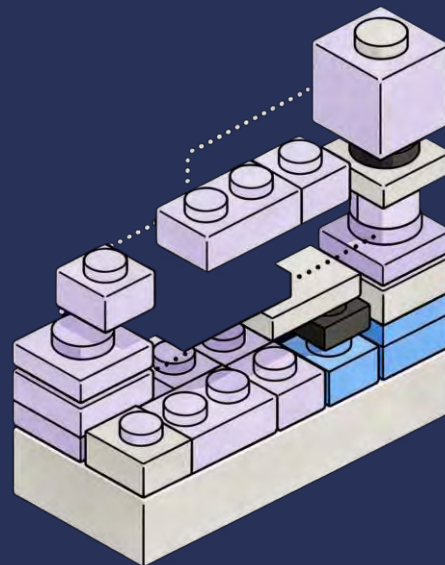
Claims Service

Owns FNOL through settlement data. Scales elastically during CAT events without impacting billing or underwriting.



Billing Service

Manages premium collection and payment reconciliation. Isolated failures do not propagate to policy issuance.



Containerization & Kubernetes Orchestration

Containers package each microservice with its runtime dependencies, eliminating environment drift between development, staging, and production.

For insurance workloads, Kubernetes delivers critical operational capabilities: horizontal pod autoscaling absorbs claims spikes, pod disruption budgets protect availability during node maintenance, and namespace isolation enforces security boundaries between regulated and non-regulated workloads.

- Horizontal Pod Autoscaler for elastic claims capacity
- Pod Disruption Budgets for zero-downtime maintenance
- Namespace isolation for regulatory separation
- Rolling deployments and automated rollback

High Availability Design

Multi-zone cluster deployments ensure that a single availability zone failure does not take down core insurance operations. Node affinity rules distribute database replicas and stateful workloads across zones.

Stateful Workloads

StatefulSets and Persistent Volume Claims manage database pods with stable network identities and durable storage, enabling containerized databases alongside stateless microservices.

Infrastructure-as-Code for Reproducible Environments

1

Define

Infrastructure and database schemas described in Terraform, Pulumi, or CloudFormation alongside application code in the same repository.

2

Review

Pull requests trigger automated plan/diff reviews. Security and compliance policies enforced via policy-as-code tools like OPA or Checkov before merge.

3

Apply

Approved changes are applied through CI/CD pipelines. State files track drift. Rollback is a code revert, not a manual intervention.

4

Audit

Every infrastructure change is logged, attributed, and reviewable. Regulators receive a complete, timestamped change history on demand.

CI/CD: Enabling Frequent, Reliable Deployments

1

Automated Functional Testing

Unit, integration, and contract tests validate business logic—underwriting rules, premium calculations, claims adjudication—on every commit before any code reaches a shared environment.

2

Performance & Load Testing

Synthetic load tests simulate CAT-scale claims volume in pipeline stages, preventing performance regressions from reaching production undetected.

3

Security Scanning

SAST, DAST, dependency scanning, and container image scanning gates enforce security standards without blocking teams—shift-left, not bolt-on.

4

Progressive Delivery

Blue/green and canary deployment strategies route a controlled percentage of traffic to new versions, reducing blast radius and enabling data-driven rollout decisions.

Event-Driven Architecture: Resilience at Scale

Why Events?

Synchronous request/response coupling between services means one slow or failed downstream service cascades latency and failures upstream. Insurance workflows—claims intake, fraud detection, payment processing—involve multiple systems that need to react to the same business events without tight dependency chains.

Platform Options

Apache Kafka: High-throughput, ordered, durable event streaming. Ideal for claims event sourcing and audit-grade event logs.

AWS EventBridge: Serverless event routing with native AWS service integrations. Well-suited for partner ecosystem and SaaS integration patterns.

Key Resilience Patterns

- **Event Sourcing:** System state rebuilt from an immutable event log—critical for claims audit trails and regulatory replay requirements
- **CQRS:** Separate read and write models allow reporting workloads to scale independently of transactional processing
- **Dead Letter Queues:** Failed events are captured for inspection and retry, preventing silent data loss
- **Saga Pattern:** Choreographed or orchestrated long-running transactions replace distributed two-phase commit
- **Circuit Breaker:** Downstream service failures are isolated, preventing cascading outages across the claims or billing pipeline

API-First Design for Partner Ecosystems



Design-First Contracts

OpenAPI specifications are authored before implementation. Consumer teams validate against the spec early, catching incompatibilities before they reach integration environments.



Security & Governance

OAuth 2.0, mutual TLS, API gateway rate limiting, and centralized secrets management enforce zero-trust boundaries across the partner data exchange layer.



Versioning Strategy

Semantic API versioning with deprecation timelines ensures partner integrations remain stable while internal implementations evolve freely behind the contract boundary.



Gateway & Mesh

API gateways handle external traffic routing and authentication. Service mesh handles internal east-west mTLS, traffic policies, and retry logic between services.

Modern Observability: Seeing Across Hundreds of Services

When a single business transaction traverses dozens of microservices and data stores, traditional log-based monitoring is insufficient. Modern observability is built on three pillars—logs, metrics, and distributed traces—unified in a way that allows engineers to ask arbitrary questions about system behavior without instrumenting for every possible failure mode in advance.

Centralized Logging

Structured logs from all services, containers, and databases flow to a central platform. Correlation IDs link every log line to the originating request, enabling precise incident reconstruction.

Distributed Tracing

OpenTelemetry-instrumented traces capture the full span tree of a claims transaction—from API gateway through policy service, fraud check, and database write—revealing latency hot spots and error origins.

Metrics & Alerting

RED metrics (Rate, Errors, Duration) per service, combined with business-level SLIs and SLOs, drive actionable alerts with context rather than noisy threshold breaches.

Distributed Data: Consistency & Schema Evolution

Distributed Consistency Challenges

When each microservice owns its database, cross-domain queries and multi-service transactions become non-trivial. Strong consistency guarantees that work in a monolith must be reconsidered against the CAP theorem realities of distributed systems.

- **Eventual Consistency:** Acceptable for many insurance read paths; design UI and downstream processes to tolerate brief lag
- **Saga Orchestration:** Compensating transactions handle partial failures in multi-step processes like policy issuance
- **Outbox Pattern:** Guarantees at-least-once event delivery alongside database writes without distributed transactions

Schema Evolution Without Downtime

Database schemas must evolve continuously without breaking running services. Flyway or Liquibase manage versioned, repeatable migrations that run as part of the CI/CD pipeline. Schema changes follow backward-compatibility rules—additive changes first, breaking changes only after all consumers migrate.

- Expand-contract pattern for zero-downtime column changes
- Consumer-driven contract tests validate schema compatibility
- Blue/green database migrations for large table changes
- Schema registry enforces compatibility on Kafka event topics

Modernization Roadmap: From Legacy to Cloud-Native

Successful modernization is incremental, not a big-bang rewrite. The strangler fig pattern allows new microservices to be built alongside the legacy core, with traffic gradually migrated as confidence grows. Each phase delivers working software and measurable outcomes, sustaining organizational momentum and executive confidence throughout a multi-year program.

Phase 1: Foundation

Containerize existing applications. Establish CI/CD pipelines, IaC baseline, and centralized observability. No architectural refactoring yet—build the operating model first.

Phase 3: Scale

Deploy service mesh. Adopt GitOps for platform operations. Expand microservice boundary coverage. Enforce API-first contracts across all partner integrations.

1

2

3

4

Phase 2: Decomposition

Extract highest-value or highest-pain domains (typically claims or billing) as standalone microservices. Introduce event streaming for decoupled communication. Migrate to managed cloud databases.

Phase 4: Optimize

Fine-tune autoscaling policies for CAT scenarios. Implement SLO-based alerting. Run chaos engineering experiments. Decommission legacy monolith components as traffic migrates fully.

Key Takeaways & Actionable Next Steps

Cloud-native transformation in insurance is achievable and well-precedented—but it requires deliberate architectural choices, not just cloud migration. The patterns covered in this session are complementary and reinforcing: microservices enable domain ownership, CI/CD enables velocity, observability enables confidence, and organizational change makes it all sustainable.

Start with the operating model

Build CI/CD pipelines and observability before decomposing the monolith. Teams need the safety net before they can move fast.

Design for failure, not just uptime

Circuit breakers, saga patterns, dead letter queues, and chaos engineering build the muscle memory for resilience before the next CAT event tests it in production.

Decouple data domains deliberately

Domain-driven design and bounded contexts are the map. Each service owns its schema; cross-domain queries are solved through events or APIs, not shared databases.

Invest in people as much as platforms

Upskilling, blameless culture, and cross-functional team structures are not soft concerns—they are prerequisite infrastructure for a sustained cloud-native engineering organization.

Thank You!