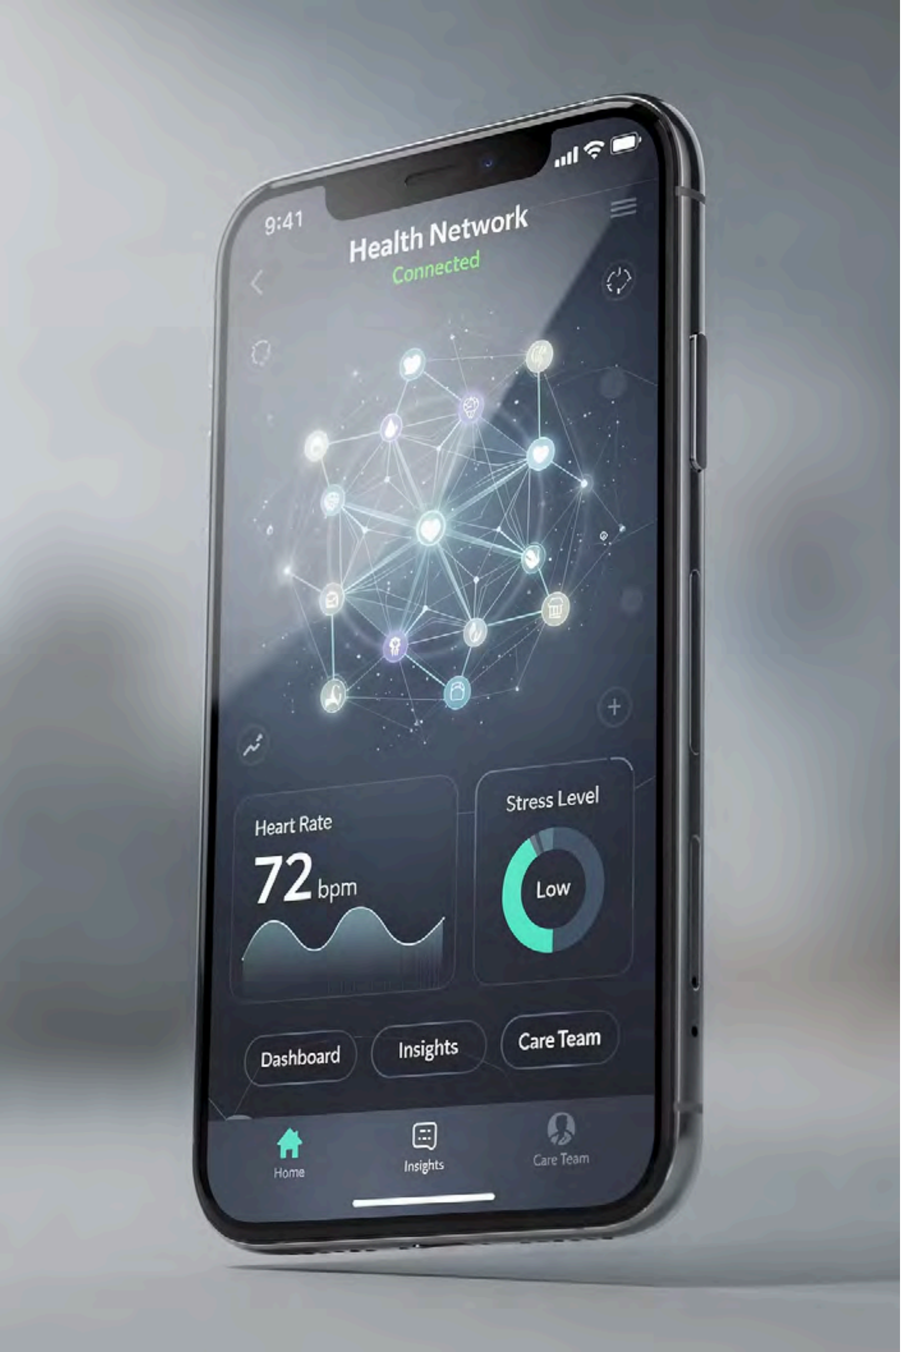

Building Reliable AI-Driven Mobile Platforms for Healthcare and Commerce

Dipta Rakshit
Affiliation: North-Eastern Hill University

CONF42 SITE RELIABILITY ENGINEERING

The Challenge: Two Demanding Domains, One Platform

Healthcare Requirements

- HIPAA-compliant data handling
- Prescription lifecycle management
- Zero tolerance for missed dose events
- Insurance and pharmacy integrations

Commerce Requirements

- Real-time personalization at scale
- AR-based virtual try-on experiences
- High-throughput transaction processing
- Promotional engine reliability

When a single mobile platform must serve both domains simultaneously across human, vision, and pet care engineering tradeoffs become critical decisions with clinical and commercial consequences.

Narrative Arc: Problem → Architecture → Outcomes

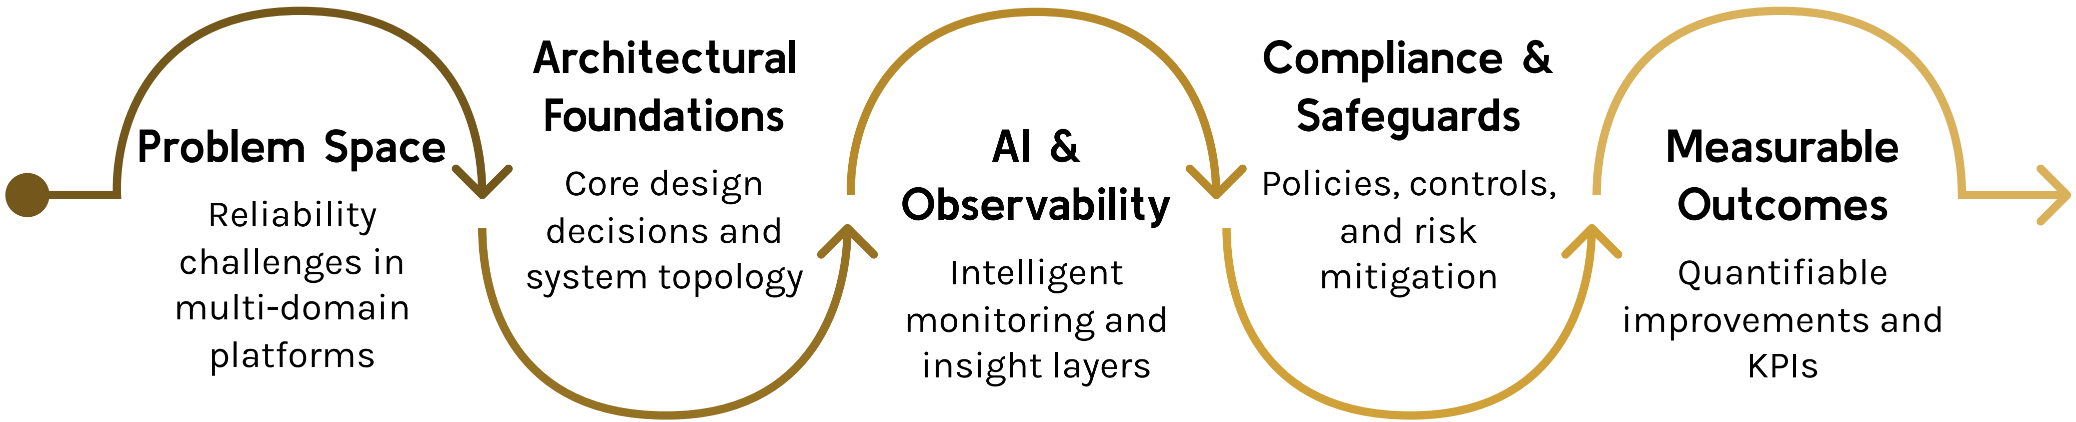

This session follows a deliberate engineering narrative starting from the reliability problem inherent in multi-domain platforms, building through architectural decisions, and arriving at measurable operational and clinical outcomes. Each layer informs the next.

Architectural Foundations for Multi-Domain Reliability

Resilient Microservices

Domain-scoped services with independent deployment cycles, circuit breakers, and bulkhead patterns prevent cascading failures across pharmacy and retail subsystems.

Fault-Tolerant Integrations

Asynchronous messaging and idempotent APIs ensure reliable handoffs between pharmacy, insurance, and retail backends even when upstream systems degrade.

Consistent Data Synchronization

Event-driven state management prevents data drift between prescription records, user profiles, and retail transactions a prerequisite for both compliance and personalization.

HIPAA & GDPR: Compliance as an Engineering Constraint

Regulatory compliance is not a post-launch checkbox it is a first-class architectural requirement. APIs handling Protected Health Information (PHI) and personal data must be designed with encryption in transit and at rest, role-based access controls, and auditable access logs from day one.

→ HIPAA-Compliant APIs

Minimum necessary access, PHI tokenization, and signed audit trails built into every service boundary.

→ GDPR Data Flows

Consent management, right-to-erasure workflows, and cross-border data residency controls integrated at the API gateway layer.

→ Unified Policy Enforcement

A single compliance enforcement plane covering both healthcare and commerce reduces surface area and audit complexity.

Key Principle

Compliance constraints become reliability features. An API that enforces minimum necessary access also limits blast radius during a security incident. Audit logs that satisfy HIPAA also provide the observability signals SREs need during incident investigation.

AI-Driven Personalization: Architecture Under the Hood

Predictive Refill Models

ML models anticipate prescription refill windows based on fill history, adherence signals, and supply chain availability triggering proactive reminders before gaps occur.

Real-Time Offer Personalization

Low-latency inference pipelines serve personalized promotions and product recommendations within SLA budgets, without blocking the primary user journey.

Graceful Degradation

When inference services are unavailable, the platform falls back to rule-based recommendations maintaining user experience without surfacing errors or empty states.

Observability: The SRE View of AI at Scale

What Must Be Observable

- Model inference latency per request path
- Feature drift and prediction confidence scores
- Fallback activation rate by service
- PHI access patterns and anomalies
- Downstream integration error rates

Telemetry Architecture

A multi-signal observability stack distributed traces, structured logs, and business-level metrics gives SREs both technical and clinical visibility into platform behavior.

Alerting thresholds are set not just on latency and error rate, but on **outcome signals**: missed refill events, failed insurance verifications, and drop-offs in personalized flows. These are the signals that connect reliability engineering to clinical impact.

Self-Healing Systems and Fault Isolation

Automated remediation is essential at scale. Health checks trigger circuit breakers to isolate degraded pharmacy integrations without affecting the retail flow. Retry policies with exponential backoff prevent thundering-herd scenarios during upstream recovery. Chaos engineering practices validate these mechanisms continuously not just at launch.

Accessibility and Localization as Reliability Concerns

ADA-Compliant Design

Screen reader support, sufficient color contrast, and touch target sizing ensure consistent service quality for users with disabilitiesm reducing support escalations and improving task completion rates.

Multilingual Localization

Dynamic content localization with fallback language chains ensures prescription instructions and safety information reach all users accurately a clinical safety requirement, not just a UX nicety.

Inclusive SLAs

Reliability targets apply uniformly across language locales and accessibility modes. A degraded experience for any user segment is treated as an SLA breach, not an acceptable edge case.

Prescription Adherence: Where Reliability Meets Clinical Outcome

Prescription adherence is one of the most consequential reliability metrics in healthcare IT. Every missed notification, failed refill trigger, or synchronization error between pharmacy and mobile systems translates directly to a patient risk event.

Predictive Reminders

Refill reminders sent ahead of supply depletion, personalized to fill cadence and patient behavior patterns.

Adherence Monitoring

Aggregated adherence signals feed back into the predictive model, continuously improving reminder timing and channel selection.

Integration Resilience

Pharmacy system outages trigger buffered retry queues—no adherence event is silently dropped.

Engineering Principle

Adherence improvement is a system reliability outcome. Building a platform that never silently drops a refill event, never fails to deliver a critical notification, and recovers gracefully from upstream failures this is what reliability engineering looks like in healthcare.

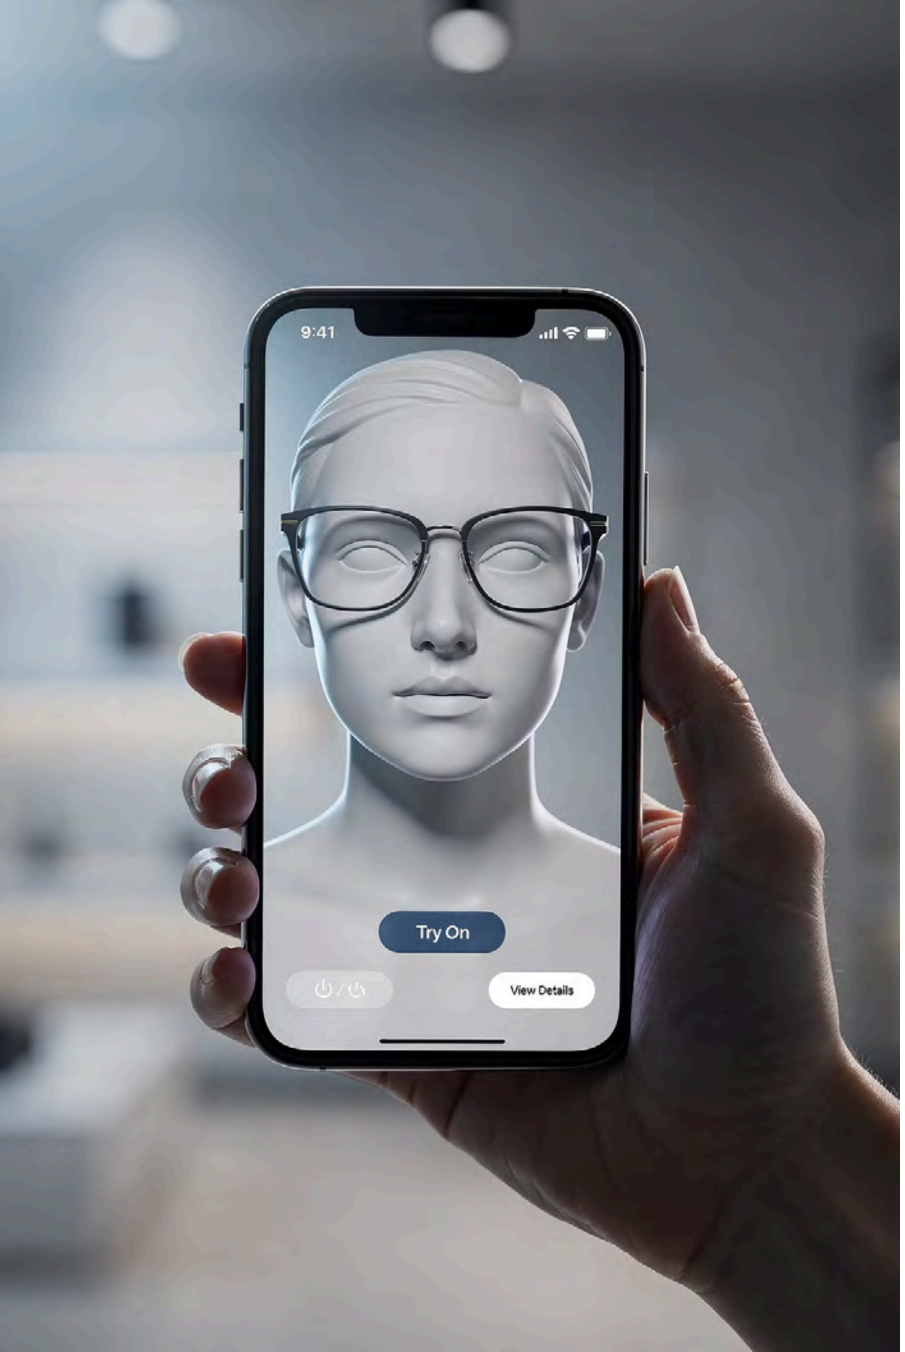

AR-Powered Commerce: Reliability for Immersive Experiences

Low-Latency Rendering Pipeline

AR try-on experiences for vision and eyewear products require frame-accurate rendering with tight latency budgets demanding edge-optimized inference and efficient asset delivery.

Progressive Degradation

When device capabilities or network conditions cannot support AR, the platform gracefully falls back to high-resolution static imagery preserving the purchase funnel without user-visible errors.

Session State Integrity

AR session state is persisted across interruptions (calls, backgrounding, low-memory events), ensuring users return to their try-on context without losing product selections or personalization state.

Implementation Patterns: Lessons from the Field

1 Decouple Health and Commerce Data Planes

Shared infrastructure is cost-efficient, but PHI and transactional data must traverse separate encrypted pipelines with independent audit controls. Mixing them creates compliance exposure and complicates incident investigation.

2 Build Fallback Paths Before You Need Them

Define and test degraded-mode behavior for every AI-driven feature before launch. A feature that fails closed returning nothing is more damaging than one that returns a rule-based default.

3 Instrument Outcomes, Not Just Operations

SLIs tied to business and clinical outcomes (refill completion rate, offer engagement, accessibility task success) give engineering teams early warning before operational metrics breach thresholds.

4 Treat Localization as a Deployment Pipeline

Localized content must pass through the same validation and staging gates as code incorrect prescription instructions in any language are a patient safety issue, not a translation bug.

Key Takeaways for SREs and Platform Engineers

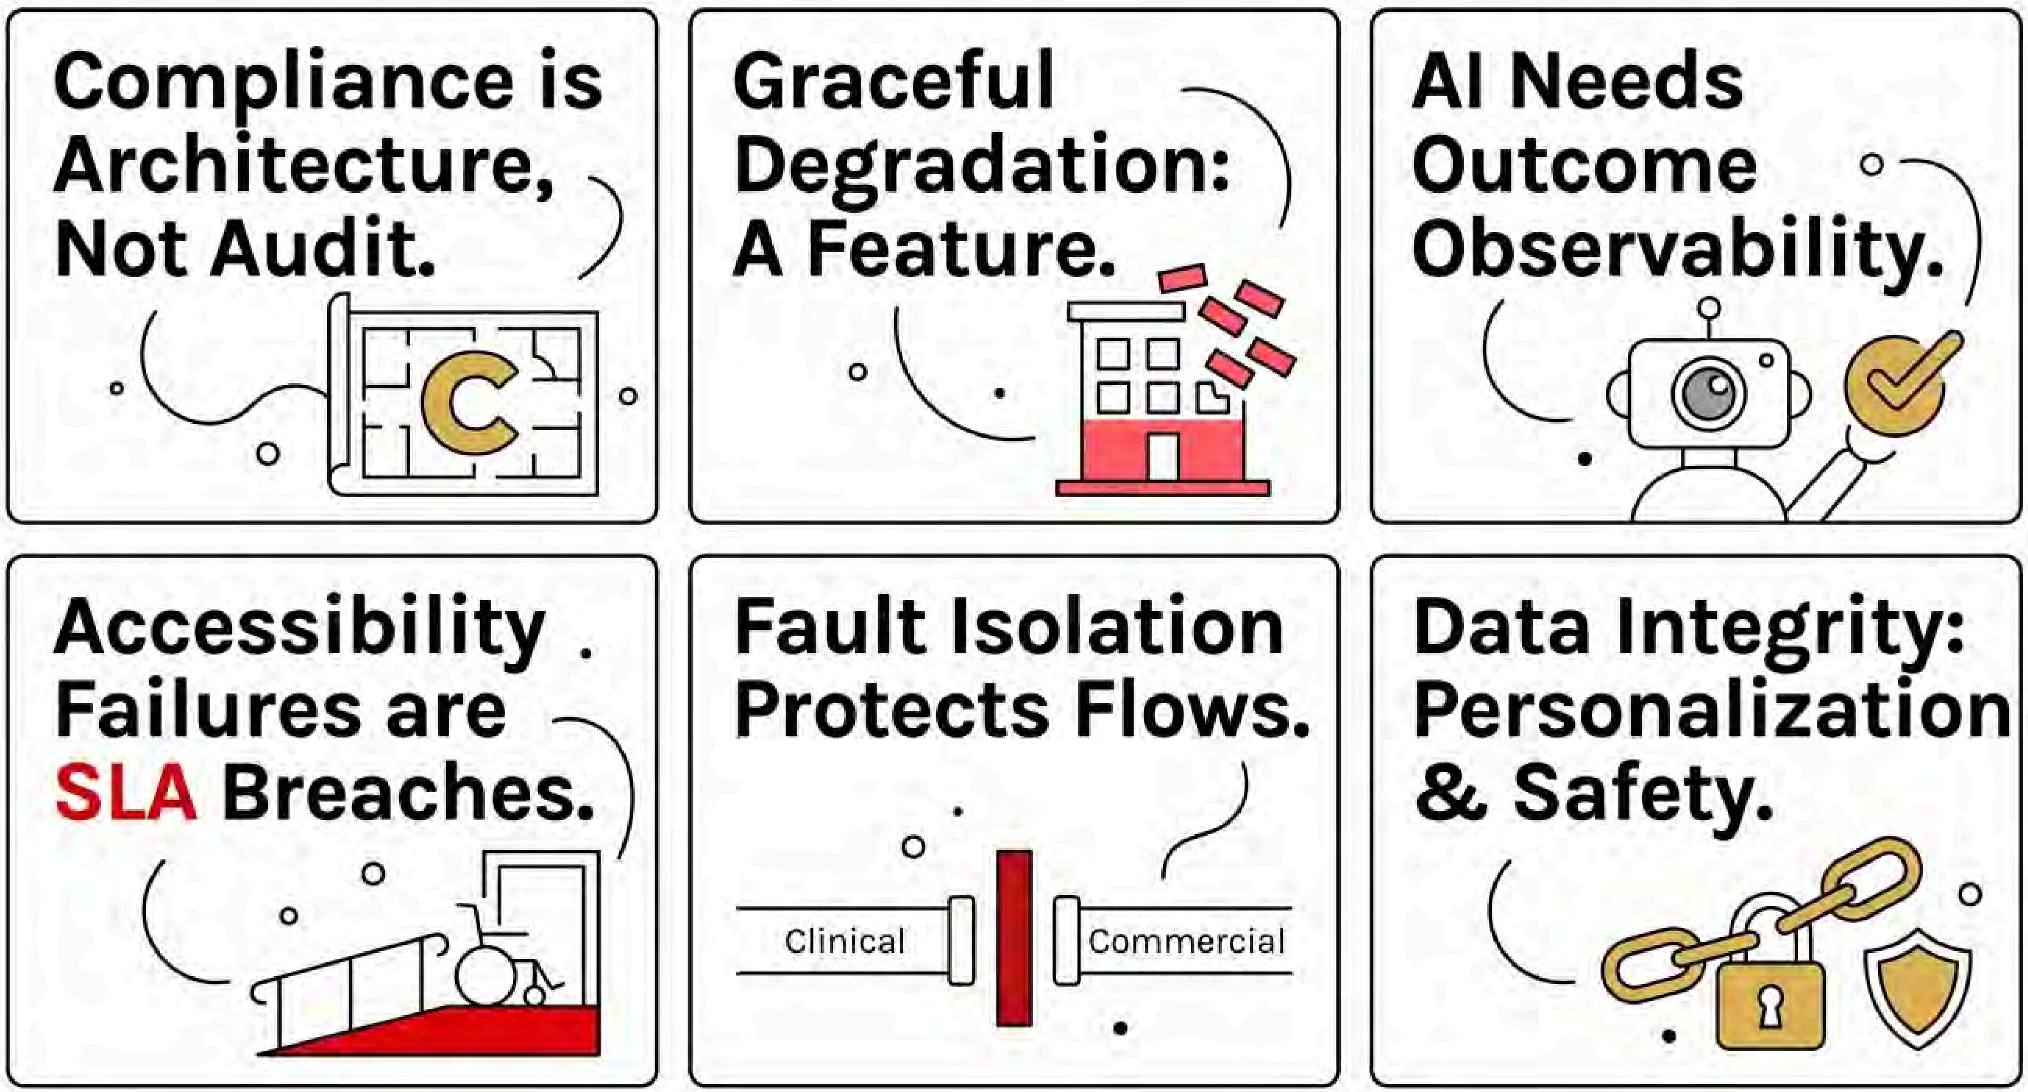

Reliable AI-driven platforms in healthcare and commerce are not built through a single technology choice they emerge from a consistent set of architectural principles applied across every layer of the system, from API design to deployment pipelines to on-call runbooks.

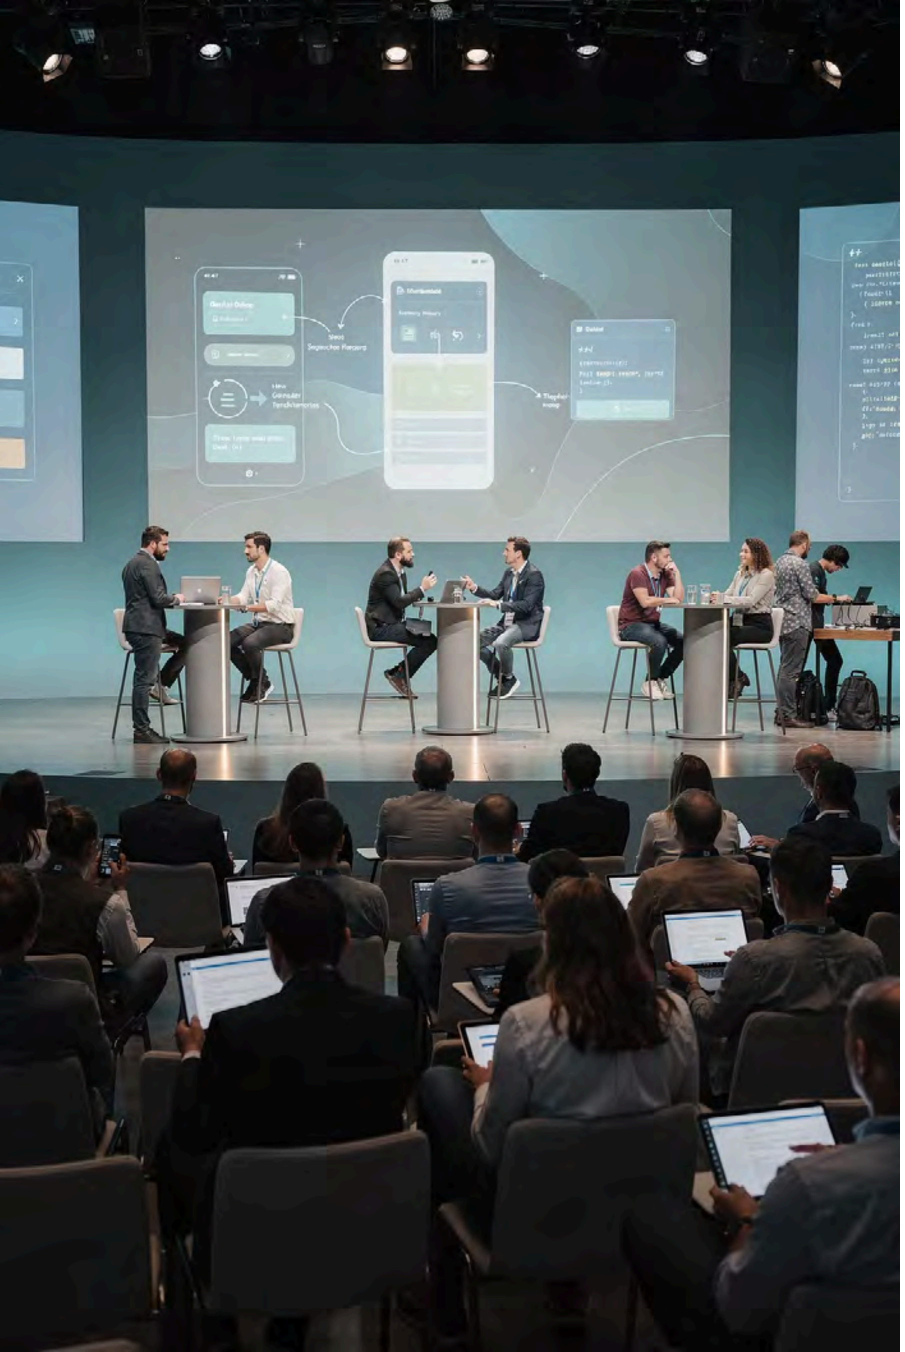

Thank You! Let's Build Platforms That Earn Trust

> Reliability is not a feature you add. It is a property that emerges from every design decision made across the entire system and it is the foundation on which clinical impact and commercial performance both depend.

Continue the Conversation

Questions on observability strategies, compliance API patterns, or AI fallback architecture are welcome.

Session Resources

Architecture reference diagrams and further reading on HIPAA-compliant microservice design available on request.