USE CASE

Unified Digital Foundation: Powering a Global Live Sports Platform

How a repeatable platform architecture eliminated downtime, accelerated app delivery, and embedded AI into an enterprise that runs on real-time data.

Platform Engineering | Event-Driven Architecture | Enterprise AI Integration | MCP | SRE & Reliability

THE SITUATION

A Patchwork of Legacy Systems Under Live Event Pressure

When we partnered with a global live sports organization, their technology landscape was a patchwork of legacy applications. These systems struggled to deliver the reliability, real-time data flow, and speed of change required for a global event watched by millions. The operating model was feature-hungry, calendar-bound, and unforgiving of downtime. A single failure could ripple across player scheduling, fan engagement, and broadcast experiences.

The stakes were enormous. Millions of fans across digital channels expected seamless, real-time experiences. Players and officials depended on accurate scheduling and credentialing. Broadcast partners required reliable data feeds with zero latency surprises. And the technology team was expected to deliver new capabilities on a fixed event calendar with no room for regression.

The organization also saw the potential of AI to transform fan experiences, automate operational workflows, and generate content at scale. But their fragmented architecture had no unified way to integrate AI models into production systems safely. AI projects were isolated experiments, disconnected from the data, governance, and services they needed to deliver real value.

> **How do you modernize without disrupting a live global event? How do you create a foundation that not only scales for future seasons but also makes AI a first-class capability, integrated into the systems you already have?**

THE BIGGER PROBLEM

AI Does Not Work in a Vacuum

Most enterprises treat AI as a separate initiative: a set of models, a data science team, a proof of concept that never reaches production. The reason is straightforward. AI models need access to live enterprise data, existing business logic, current system state, and governance controls. Without a unified integration layer, every AI use case requires custom plumbing to connect models to the systems they need. The result is fragmented experiments instead of enterprise-scale intelligence.

The Unified Digital Backbone as an AI Enablement Platform

The Unified Digital Foundation we built is not just an integration architecture. It is an AI enablement platform. By creating a single backbone that connects applications, data, events, and governance, we gave AI models and agents a standardized way to participate in enterprise operations. Instead of building one-off integrations for each AI use case, the foundation provides:

Standardized Data Access

AI models connect to live enterprise data through governed APIs and event streams, not bespoke ETL pipelines.

MCP for Agent Integration

The Model Context Protocol gives AI agents a secure, standardized interface to interact with enterprise systems, databases, and SaaS tools.

Centralized Model Governance

A model registry with lineage tracking, safety guardrails, drift monitoring, and canary deployments ensures AI scales responsibly.

Reusable AI Assets

Agents, tools, data flows, RAG frameworks, and prompt libraries are shared across use cases instead of rebuilt for each project.

The key insight: enterprises do not need a separate AI strategy. They need an integration strategy that treats AI as a first-class participant. When the backbone is in place, new AI capabilities plug in and reach production in days instead of months.

OUR APPROACH

A Unified Digital Foundation: One Backbone, Many Experiences

We introduced a Unified Digital Foundation: a modern architecture that connects micro apps, enterprise systems, data, and AI across the organization. Rather than replacing systems one by one (the "big bang" approach that derails most modernization efforts), we built an Intelligence and Integration Layer that sits at the center of the ecosystem, standardizing contracts and governance so domain teams can ship new experiences without re-plumbing core capabilities every time.

The foundation blends three communication paradigms, each chosen for a specific reason:

API-Driven Synchronous Communication

For predictable, real-time interactions where a request needs an immediate response: player lookups, credential verification, schedule queries. APIs are treated as formal contracts, versioned in a registry, and managed through a centralized gateway.

Event-Streaming Asynchronous Workflows

For flexibility and resilience where systems need to react to changes without tight coupling: score updates flowing to fan apps, schedule changes propagating to broadcast systems, credential status changes triggering downstream workflows.

Model Context Protocol (MCP) for AI Integration

For seamless AI integration into operational workflows. MCP provides a standardized way for AI models and agents to interact with enterprise systems and data, enabling intelligent automation across content generation, fan engagement, and operational decision support without custom integration code for each use case.

This hybrid approach promotes micro apps, empowering product teams to deliver bespoke solutions for each customer segment. Loose coupling simplifies change management. Real-time sharing powers better fan and player experiences. Embedded intelligence puts AI where it matters most. And consistent governance holds it all together across distributed systems.

EXECUTION

What Made It Work

Start Small, Scale Smart

Landed value early with a focused initial domain, then expanded by domain area. This avoided "big bang" risk and built organizational confidence incrementally.

Platform + Product Pairing

A persistent core platform pod built the Integration and AI layer while product pods (Credentials, Player Services) scaled up and down around it as operational demands shifted.

Contracts Over Coupling

APIs and events treated as formal contracts, versioned in a registry. This cut cross-team synchronization costs and let teams ship independently.

SRE Discipline

Error budgets, SLOs, automated runbooks, and progressive delivery across services and models. Reliability was engineered, not hoped for.

Developer Experience First

Golden paths, sample services, auto-scaffolding, and paved-road CI/CD pipelines. New teams could contribute production-ready code within days, not weeks.

Model Governance

A model registry with lineage tracking, safety guardrails, and canary deployments for ML and LLM endpoints. AI capabilities scaled safely alongside traditional services.

RESULTS

The Outcomes

- **Zero downtime during live event windows, with rapid addition of new apps because the plumbing was already in place**
- **AI-generated content and experiences shipped to production at scale, powered by the same backbone that runs core operations**
- **Highest event satisfaction scores and measurably improved fan engagement across digital channels**
- **Solutions reused across partner events globally, thanks to the repeatable architecture, superior UX, and speed of enhancements**
- **Scaled beyond projected volumes in 2025, handling record-breaking concurrent usage without degradation**
- **New micro apps (Player, Official, Credentials, Practice Court) delivered on accelerated timelines by independent product teams**

COMMON PITFALLS

Why Others Fail

- ✕ No clear, simple strategy: the modernization vision is buried in complexity and never translated into an executable plan.
- ✕ Lack of a unified application, AI, and data integration approach: each domain builds its own plumbing and duplicates effort.
- ✕ Over-reliance on monolithic enterprise apps that cannot adapt to the pace of change required by live, real-time operations.
- ✕ No persistent platform team: integration work is project-based, leading to fragmented ownership and compounding technical debt.
- ✕ Poor contract governance: ad-hoc APIs and events with no versioning discipline, creating chaos at the boundaries between teams.
- ✕ Siloed delivery: teams build bespoke integrations instead of consuming shared services, multiplying maintenance burden.
- ✕ Treating AI as a separate initiative instead of integrating it into the enterprise backbone, resulting in proofs of concept that never reach production.
- ✕ Little focus on developer experience: slow onboarding, fragile pipelines, and no golden paths for new contributors.

KEY TAKEAWAY

More Than Technology: An Operating Model

A Unified Digital Foundation is not just a technology architecture; it is an operating model that makes AI a working part of your enterprise instead of an isolated experiment. By pairing platform thinking with product agility, this organization transformed its digital backbone into a competitive advantage where new applications, AI agents, and data-driven experiences share a common foundation. The architecture, governance patterns, and team structures are repeatable. Any enterprise seeking speed, resilience, and intelligence at scale can apply this same foundation to accelerate their own transformation.

ARCHITECTURE

Unified Digital Backbone: High-Level Architecture

At the heart of the enterprise is an Intelligence and Integration Layer: a unified backbone delivering API-first synchronous access, event-driven streaming, unified data services, and a Model Context Protocol plane for secure AI lifecycle management. Above this core, domain applications and experience layers consume shared services to deliver consistent, personalized experiences without duplicating integration logic.

Unified Digital Backbone Architecture

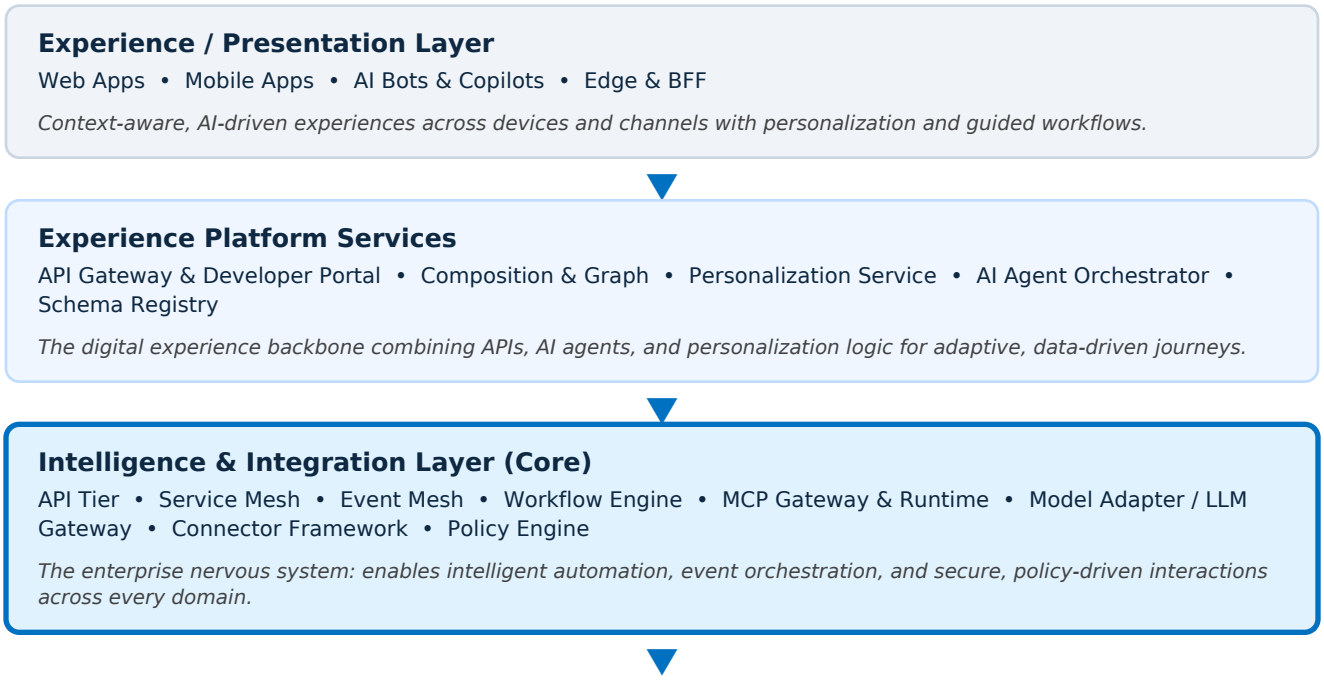

AI & ML Plane

Model Registry & MLOps • Model Serving • Feature Store • Model Observability • Prompt Library & Policy Mgr • Inference Cache

Enterprise knowledge meets adaptive reasoning, backed by governance and observability.

Data Plane

Data Lake • Curated Zone • Search & Indexing • Metadata Catalog • Data Contracts & Consent

Connects raw ingestion with curated analytics for reliable, compliant data consumption.

Cross-Cutting Services

Identity & Secrets Mgmt • Security & Compliance • Observability (OpenTelemetry) • SRE & Reliability Platform • DR & Multi-Region Resilience

Shared services guaranteeing security, compliance, and operational resilience across the entire stack.

Ready to Build Your Digital Foundation?

Aubrant Digital helps enterprises design and deliver unified platform architectures that make AI a working part of your business. From advisory through production, we build with you.

Let's Talk