

# SRE Driven Frontend Architecture for High Performance and Reliable Web Systems

CONF42 SITE RELIABILITY ENGINEERING 2026

**Murali Ajit Varma**

Senior Staff Software Engineer · Galileo Financial Technologies LLC

## ABOUT THE SPEAKER



# Murali Ajit Varma

Senior Staff Software Engineer  
Galileo Financial Technologies LLC



## Enterprise Platforms at Scale

B2B operational platforms serving 200+ enterprise clients, from agent tooling and relationship management to high-volume document generation



## High-Throughput Systems

Background spanning gaming infrastructure at scale, real-time transaction processing, and performance-critical distributed systems



## Driving AI Adoption

Leading AI adoption across engineering teams, from developer productivity tooling to governance frameworks for enterprise LLM integration

# The Cost of Slow Operational Tooling

*Agents and relationship managers do not abandon your app. They are stuck using it.  
The damage is silent and cumulative.*

## Industry Baseline

53% of users abandon apps beyond 3 seconds.  
71% abandon beyond 4 seconds.

These thresholds set the performance floor. Operational platforms face the same latency physics, even when users cannot leave.

## B2B Plight

Every extra second of load time multiplies across hundreds of agents, thousands of interactions per day.

Slow tooling compounds into longer handle times, more errors under pressure, and agent fatigue across full-day sessions.

For enterprise operational platforms, frontend latency is not a bounce rate problem. It is a compounding operational cost.

# How We Got Here: The SPA Era

Single Page Applications gave us rich interactivity but introduced compounding reliability costs.



## 2-3 MB

**Bundle Size**

JavaScript bundles per page load, before any user interaction begins



## 3-4 sec

**First Contentful Paint**

On constrained networks, well beyond acceptable thresholds for operational tooling



## High Variance

**Network Sensitivity**

Client-heavy rendering collapses under degraded network conditions

# Frontend Is a Reliability Problem

At Galileo, frontend latency was not a UX polish issue. It was an operational reliability problem. Platform performance directly influenced system-level outcomes for enterprise clients.



## Task Completion

Slow renders increased handle time per agent interaction in high-volume workflows



## Session Stability

Agents in 8-hour sessions hit rendering variance, timeouts, and accumulated UI drift



## Infrastructure Cost

Client inefficiency increased downstream API pressure and compute spend across tenants



## Operational Throughput

Platform degradation under load reduced the volume of interactions agents could process daily

# The Client-Heavy Legacy Model

All rendering responsibility lived in the browser, compounding latency at every step.



## Baseline Metrics at Galileo

~300 ms

Server response time, reasonable but offset by client rendering overhead

~65%

Cache hit rate, with significant room for improvement

High  
Variance

Rendering unpredictability elevated across agent workflows and tenants

# The Architectural Shift

Move computation from infrastructure you do not control (user browsers) to infrastructure you do (your servers).

01



## Server Components

Business logic moves server-side. Only minimal, serialized UI payloads reach the client. Hydration becomes surgical.

02



## Streaming SSR

Progressive HTML delivery decouples time-to-first-byte from time-to-interactive. Performance becomes predictable.

03



## Hybrid Rendering

Match each content type to the right rendering mode. Static, server, or client, based on reliability requirements.

# Server Components

Business logic moves to the server. Only minimal, serialized UI payloads reach the client. Hydration is surgical, not wholesale.



## Reduced Client CPU

Computation shifts server-side, freeing browser resources for interaction



## Lower Memory Overhead

Fewer components hydrated means tighter memory profiles across sessions



## Device Stability

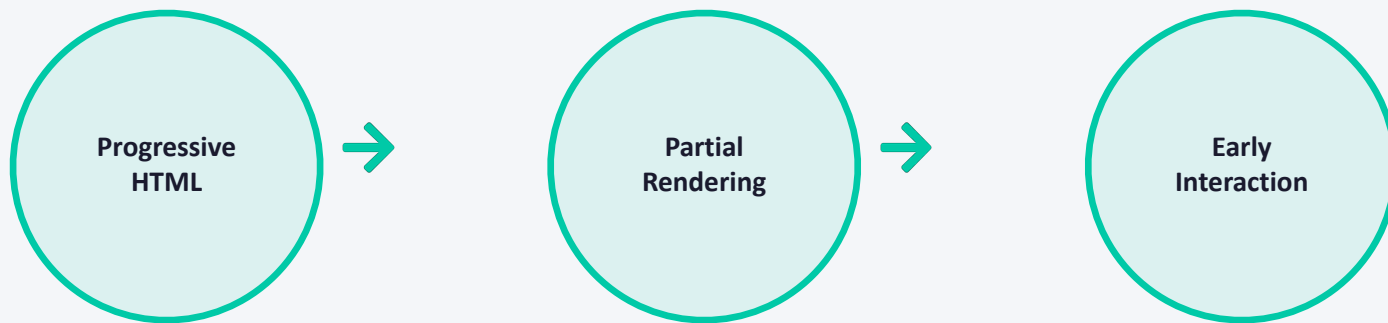
Improved consistency across the full range of agent hardware and terminals

*The core principle: move work from infrastructure you cannot control to infrastructure you can.*

# Streaming Server-Side Rendering

Streaming SSR decouples time-to-first-byte from time-to-interactive.

Performance becomes predictable, and predictability is what lets you set meaningful SLOs.



42%

TTFMP Improvement  
Time to First Meaningful Paint

45%

Faster Initial Load  
Across production page routes

# Hybrid Rendering Model

Match each content type to the rendering mode that fits its reliability requirement.  
No false choice between fully static and fully dynamic.



## Static Generation

Applied where data is stable. Maximizes cache efficiency and CDN leverage. Near-zero failure rate.



## Server Components

Dynamic, personalized data computed server-side without client overhead or rendering variance.



## Minimal Client JS

JavaScript reserved for genuinely interactive UI elements. Deterministic, auditable, scoped.

Deterministic performance · Improved observability · Reduced failure surface · Enhanced resilience under load

# Build System Modernization

Legacy build tooling was a hidden constraint on performance evolution.  
Modern JavaScript build pipelines were re-evaluated and replaced.

## Rust-Based Tooling

Turbopack evaluated and introduced for significantly faster compilation cycles

## Incremental Compilation

Improved strategies reduced rebuild scope and time on large codebases

## Optimized Bundling

Reduced output size and improved tree-shaking fidelity across modules

# 10x

Cold Start Speed Improvement

*Faster builds = shorter feedback loops = faster rollbacks when things go wrong.*



# Production Results at Galileo

*Sustained across multiple production releases, not isolated benchmarks.*

60%

JS Bundle  
Reduction

45%

Faster  
Page Load

42%

TTFMP  
Improvement

## Longitudinal Observations

### Before

300 ms server response  
65% cache hit rate  
Larger client payloads  
Higher rendering variance

### After

150 ms server response  
89% cache hit rate  
Smaller bundle footprint  
Faster interaction readiness

# Impact on Engineering Velocity

Architecture modernization did not just improve performance.  
It changed how fast and confidently the team ships.

35%

## Faster Time to Market

Reduction in delivery cycle time across  
deployment cycles

47%

## Developer Productivity

Increase observed across teams post-  
modernization

Fewer

## Rollbacks

Improved release confidence and  
reduced incident-driven reversals

# Three Principles to Take Home

The methodology here is not a bespoke financial solution.  
It is a transferable architectural blueprint for any large-scale, high-availability system.



## Frontend Reliability Belongs in SRE

Client-side performance must be owned by reliability engineering, not treated as a UX afterthought.



## Server-Centric Architectures Scale

SSR and server components provide a durable, scalable foundation for enterprise-grade web systems.



## Start at the Architecture Layer

Performance engineering must begin at the architectural layer. Retrofitting is costly and incomplete.

# References and Data Sources

## Industry Research

Web performance abandonment statistics at 3-4 second load-time thresholds and user behavior impact

Streaming SSR and React Server Components technical documentation

Turbopack build system documentation from the Vercel / Next.js ecosystem

## Production Data Sources

Performance metrics collected from Galileo Financial Technologies production monitoring systems

Observability and deployment metrics from Galileo production release cycles



Integrity note: All performance metrics presented derive from real production systems. No synthetic benchmarks or fabricated data.



# Thank You

**Murali Ajit Varma**

Senior Staff Software Engineer · Galileo Financial Technologies LLC

CONF42 SITE RELIABILITY ENGINEERING 2026

---