

# Engineering Resilient Ticketing Platforms: Concurrency, Scale, and Failure Modes

This briefing document analyzes the technical requirements, architectural strategies, and critical failure modes associated with building a high-scale ticketing platform. It synthesizes insights from industry incidents, such as the 2022 Ticketmaster/Taylor Swift fiasco, and engineering solutions from platforms like Shopify and Stripe.

## Executive Summary

Building a ticketing platform represents one of the most challenging problems in distributed systems due to the extreme "read-heavy" nature of event discovery contrasted with the "high-contention" requirements of ticket booking. The primary technical objective is to maintain strict consistency (preventing double-booking) while surviving "thundering herd" traffic spikes that can reach billions of system requests.

Successful architectures prioritize availability for discovery but shift to high consistency for reservations. Key strategies include using distributed locks with TTLs, virtual waiting rooms to gate-keep traffic, and idempotency keys to ensure payment safety. Recent engineering shifts, such as Shopify's move from Redis to MySQL for inventory, suggest that modern database features like `SKIP LOCKED` and `READ COMMITTED` isolation levels are becoming essential for managing high-concurrency inventory at scale.

## Analysis of Key Technical Themes

### 1. Concurrency Control and the Double-Booking Problem

Preventing the sale of the same seat to multiple users is the central integrity requirement. The documentation outlines four primary strategies for managing this contention:

Strategy	Implementation	Pros	Cons
<b>Pessimistic Locking</b>	<code>SELECT FOR UPDATE</code> on a specific row.	Strong consistency; built into RDBMS (MySQL/PostgreSQL).	High contention; database resources strained by long-held locks; risk of deadlocks.

Strategy	Implementation	Pros	Cons
<b>Status + Expiration</b>	A "reserved" status with an <code>expiration_time</code> column and a background cron job to release stale locks.	Improved user experience (visible timer); avoids long-held DB locks.	Cron job lag; potential for "zombie" reservations if the cleanup process fails or delays.
<b>Distributed Locking (Redis)</b>	Use Redis <code>SET NX EX</code> to create a temporary key (e.g., <code>ticket_id:123</code> ) with a TTL (e.g., 10 mins).	Extremely fast; automatic expiration via TTL; offloads load from the primary SQL DB.	Added architectural complexity; data loss in non-persistent cache results in dropped reservations.
<b>One Row Per Unit (Shopify Model)</b>	Each sellable unit is a row rather than a quantity counter. Uses <code>SKIP LOCKED</code> to find available units.	Dramatically reduces contention; allows atomic "reserve and claim" within a single ACID transaction.	Table size can grow; requires replenishment logic for a "bounded pool" of rows (e.g., capped at 1,000 per item).

## 2. High-Demand Traffic and Virtual Waiting Rooms

Large-scale on-sales, such as the Eras Tour, can draw 14 million users to a site designed for 1.5 million. The "Taylor Swift Case" demonstrated that even robust systems can fail when hit with 3.5 billion system requests in a single day.

- **Virtual Waiting Rooms:** These serve as a gatekeeper, placing users in a digital queue (often backed by Redis Sorted Sets) before they access the seat map. Users receive real-time updates on their position via Server-Sent Events (SSE) or WebSockets.
- **The Admission Flow:** Once a user reaches the front of the queue, their session ID is marked as "admitted" in a cache (e.g., `admitted:{eventId}`). The Booking Service rejects any reservation request from a session not found in the admitted set.
- **Scaling the Read Path:** Event pages are cached aggressively (CDN, Redis) since event data (performer, venue) is static. Only seat availability requires dynamic fetching.

## 3. Inventory Reservation Lifecycle

The transition from "available" to "sold" is a multi-step process that must be resilient to network failures and abandoned checkouts.

1. **Reservation:** A temporary lock is placed (Redis TTL or DB row lock).
2. **Payment Processing:** The system calls an external processor (e.g., Stripe). This is a "foreign state mutation" that cannot be rolled back by the local database.
3. **Confirmation/Claim:** Upon successful payment, the inventory is permanently deducted (the "claim" step), and the reservation lock is released.
4. **Idempotency:** To prevent double-charging or duplicate bookings on retries, every request must include an **Idempotency Key**. The server stores this key and its associated response; subsequent requests with the same key return the stored result without re-executing side effects.

## 4. Search and Low-Latency Discovery

Ticketing platforms are 100:1 read-heavy. Standard SQL `LIKE` queries for performers or venues fail at scale because they require full table scans.

- **Elasticsearch/OpenSearch:** These engines use inverted indexes to provide sub-500ms full-text search.
- **CDC (Change Data Capture):** Data is synced from the primary Relational DB to Elasticsearch via a pipeline (e.g., Debezium or Maxwell) to ensure search results reflect the most recent event additions.

## Failure Modes and Technical Incidents

The following table summarizes real-world and theoretical failure modes identified in the source context:

Incident/Mode	Impact	Root Cause / Technical Context
<b>Ticketmaster / Eras Tour (2022)</b>	Site crash; public on-sale canceled.	3.5 billion system requests (4x previous peak); "industrial-scale" bot interference; unprecedented concurrent demand.
<b>Bad Bunny / Mexico City (2022)</b>	Valid ticket holders denied entry.	High volume of fake tickets caused "intermittence" in access control systems, impeding identification of legitimate tickets.
<b>Shopify Connection Exhaustion</b>	Throughput ceiling hit below target.	Discovered that the bottleneck wasn't CPU or locking, but "connection hold time." Business processes held DB connections across long transactions.
<b>Redis/Cache Failure</b>	Degraded user experience.	If a distributed lock cache restarts, all active 10-minute reservations are lost. While DB consistency remains, users must restart the booking process.
<b>TTL Expiration During Payment</b>	Failed transaction/Refund required.	If a lock TTL (e.g., 10 mins) expires just as a payment completes at minute 11, a second user may have claimed the seat. Requires automatic refund logic.

## Important Quotes with Context

***"The controversy ended up converting more Gen Zers into anti-monopolists overnight than anything I could have done."***

— Lina Khan, FTC Chair, regarding the Ticketmaster website crash.

**Context:** Highlighting how technical failures in high-profile consumer systems can lead to significant political and regulatory shifts, such as the 2024 DOJ antitrust lawsuit.

**"A single row with a quantity column couldn't handle the contention."**

— Shopify Engineering Blog.

**Context:** Explaining why traditional counter-based inventory management fails in high-concurrency environments, necessitating the "one row per unit" strategy with `SKIP LOCKED`.

**"Once we make our first foreign state mutation, we're committed one way or another. We've pushed data into a system beyond our own boundaries and we shouldn't lose track of it."**

— Stripe Engineering (brandur.org).

**Context:** Discussing the critical importance of atomic phases and recovery points when dealing with external APIs like payment gateways or email services.

**"Ticketmaster got thrown under the bus because it's easy to throw them under the bus... There is no solution that won't piss people off more."**

— Fred Rosen, Former Ticketmaster CEO.

**Context:** Arguing that the gap between Motel 6 prices and Four Seasons expectations creates a market reality where high demand naturally leads to consumer frustration.

## Actionable Insights for Architects

- 1. Standardize Transaction Ordering:** To prevent deadlocks in complex booking flows (touching tickets, bookings, and audit logs), standardize the order of operations across all services (e.g., always DELETE from units before INSERTING into bookings).
- 2. Implement Per-Caller Connection Attribution:** As demonstrated by Shopify, slow performance is often due to connection hold time. Annotate SQL statements with comment tags (e.g., `/* conn_tag:checkout */`) to identify which business processes are starving the connection pool.
- 3. Use Shadow Mode for Migrations:** When replacing critical infrastructure (e.g., moving reservations from Redis to MySQL), run both systems in parallel. Write to both but treat the old system as the source of truth until the new one is validated against production traffic.
- 4. Decouple Search from Reservation:** Use microservices to scale Search and Reservation independently. Their scaling needs are vastly different; Search needs high availability and low latency, while Reservation needs high consistency and transaction isolation.
- 5. Adopt the "Implicit Status" Strategy:** Instead of relying solely on background cron jobs to release locks, design queries that recognize a seat as available if it is either `status = AVAILABLE` OR (`status = RESERVED` AND `expiration_time < NOW`). This ensures the system remains functional even if the cleanup background job lags.

Ready to take your next event further? Start with 7am → <https://7am.io>

