Karthick Balaji M, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Business Catalyzer: A Secure Customer Feedback And Analytics Platform For Business Growth

Karthick Balaji M[1], Sanjini S P[2], Varsha S[3], Gokila M[4], Dr. M. Rajesh Babu[5]

[1,2,3,4]Students, Department of Computer Science and Engineering
[5]Professor & Head, Department of Computer Science and Engineering Tamilnadu College of Engineering, Coimbatore, India

Abstract: In the contemporary digital economy, businesses of all scales—from informal street vendors to multinational corporations—face the persistent challenge of collecting actionable, continuous, and privacy-preserving feedback from their product users. Conventional review platforms such as Yelp and Google Reviews mandate formal business registration, expose user identities publicly, and provide no structured communication pathway between customer and business. This paper presents Business Catalyzer, role-differentiated feedback management platform engineered using Java Spring Boot, React, MySql, Chart.js. The platform introduces a secure, anonymous, one-on-one case communication channel that enables product users to submit mood-tagged feedback cases against any business entity—registered or otherwise—while ensuring identity privacy and fostering real-time dialogue. A three-graph analytics dashboard comprising the Incoming Ticket Graph, Case Resolution Graph, and Weekly Trend Graph delivers temporally segmented, sentiment-aware performance intelligence to business users. Experimental evaluation on a prototype deployment with 500 synthetic cases across 20 business accounts demonstrates sub-250 ms dashboard refresh latency for datasets up to 10,000 records, a System Usability Scale rating of 4.3 out of 5.0, and 100 percent task completion across five heterogeneous business profiles. The MySQL-backed, stateless architecture supports horizontal scaling to approximately 5,000 concurrent sessions at ten application instances. Business Catalyzer demonstrates that inclusivity, privacy, and data-driven intelligence can be unified within a single, accessible, and scalable feedback infrastructure.

Keywords—Customer Feedback Management; Business Analytics; Spring Boot; MySQL; Sentiment Classification; Case Resolution; Cloud-Native Architecture

I. INTRODUCTION

A. Background And Motivation

The digital transformation of commerce has fundamentally reordered the relationship between businesses and their customers. Consumers today operate in an environment saturated with choice: a dissatisfied user can within seconds abandon a product, post a negative review, and influence hundreds of potential buyers. According to multiple industry surveys, over 72 percent of consumers consult online reviews before purchasing a product or engaging a service [1], while businesses that actively respond to feedback observe up to 15 percent improvement in customer retention metrics. The economic stakes of neglecting customer feedback are therefore substantial, and the operational imperative to collect, analyze, and act on that feedback has never been greater.

Despite this urgency, the infrastructure available to businesses for systematic feedback collection remains fragmented, exclusionary, and ill-suited to the diversity of modern business forms. Large enterprises can invest in expensive CRM platforms, proprietary survey tools, and dedicated analytics dashboards. Small and informal businesses—which constitute the overwhelming majority of economic actors globally—lack access to equivalent tools. The result is a structural inequality in business intelligence: well-capitalized entities receive continuous, data-rich feedback streams while micro-entrepreneurs, freelancers, and informal vendors operate in near-total informational darkness about customer satisfaction.

B. Problem Statement

Four interrelated problems motivate the development of Business Catalyzer. First, the listing barrier: platforms such as Yelp, Trustpilot, and Google Reviews require businesses to maintain verified profiles before users can submit targeted feedback. Informal businesses, sole proprietors, independent coaches, and unregistered vendors are excluded entirely, leaving a vast segment of the economic ecosystem without feedback infrastructure. Second, the privacy deficit: all major review platforms are public-by-design.

User identities, review text, and ratings are visible to employers, competitors, and the general public, creating powerful disincentives against candid negative feedback. Research confirms that anonymous feedback channels yield significantly higher negative sentiment disclosure rates [2], which are precisely the insights most valuable to business improvement. Third, the communication gap: existing platforms provide no structured pathway for a business to privately engage the reviewing user, clarify a misunderstanding, propose a resolution, or confirm that a reported defect has been addressed. Feedback thus functions as a one-way broadcast rather than a collaborative resolution dialogue. Fourth, the analytics exclusion: sophisticated trend analysis, sentiment tracking, and resolution performance metrics are either unavailable to small businesses or locked behind expensive subscription tiers.

C. Objectives Of The Study

Business Catalyzer is designed to simultaneously address all four problem dimensions. Its primary objectives are: (i) eliminate the listing barrier by enabling any registered product user to submit feedback against any business entity, whether or not that entity maintains a platform profile; (ii) protect product user privacy through an anonymous, one-on-one case communication model that keeps all feedback content invisible to the general public; (iii) enable real-time, bidirectional case resolution dialogue through a secure chat-based interface accessible to both product users and business users; (iv) deliver mood-aware, temporally segmented business intelligence through a three-graph analytics dashboard embedded in the business user landing page; and (v) ensure scalable, cost-efficient deployment using MySQL-backed infrastructure requiring no on-premise database investment from any participating business.

D. Importance And Real-World Relevance

The relevance of Business Catalyzer extends across sectors and scales. A street food vendor in Coimbatore receives her first structured customer satisfaction data. A freelance sports coach identifies a recurring complaint pattern about scheduling responsiveness. A mid-size e-commerce startup tracks the percentage of negative cases successfully converted to positive resolutions week over week. In each scenario, the same platform provides tailored, actionable intelligence without requiring the business to understand CRM configuration, data pipelines, or analytics tooling. This democratization of customer intelligence constitutes the core value proposition of Business Catalyzer, and it distinguishes the platform from all existing alternatives reviewed in this work.

II. LITERATURE SURVEY

The research landscape surrounding customer feedback systems, opinion mining, and business intelligence platforms is extensive. This section synthesizes fifteen representative contributions spanning 2019 to 2025, organized thematically to illuminate the progression of the field and the specific gaps that Business Catalyzer addresses.

A. Sentiment Analysis And Opinion Mining

Wu, Wei, and Liu [3] introduced Opinion Seer, a landmark interactive visualization system for hotel review analysis that pioneered the application of uncertainty-aware
opinion mining to large-scale customer feedback datasets. Their system augmented scatter plots with radial visualization constructs to encode multi-dimensional opinion attributes including aspect ratings and temporal confidence intervals. While highly

innovative in its analytical approach, Opinion Seer operated exclusively on static, pre-collected datasets and provided no mechanism for real-time feedback submission or business-to-customer communication. The visualization techniques developed in this work directly inform the design of the Business Catalyzer analytics dashboard.

Medhat, Hassan, and Korashy [4] provided a comprehensive survey of sentiment analysis algorithms applied to customer feedback, cataloguing machine learning, lexicon-based, and hybrid approaches across fifteen domain applications. Their analysis revealed that domain-specific lexicons outperform general-purpose sentiment classifiers by an average of 12 percentage points in F1 score for business feedback contexts. This finding informs the Business Catalyzer decision to use a structured, domain-specific mood classification schema rather than general-purpose natural language processing, achieving high classification accuracy without the computational overhead of a full NLP pipeline.

Yadav and Vishwakarma [5] examined the application of BERT-based transformers to customer review sentiment classification, demonstrating 94.2 percent accuracy on a hotel review benchmark dataset. Their work established the state-of-the-art ceiling for automated sentiment classification and serves as the reference target for the NLP integration planned in future iterations of Business Catalyzer.

B. Customer Feedback Management Systems

Momaya and Muley [6] conducted a broad survey of digital customer feedback systems, observing that online feedback tracking tools help businesses identify dissatisfied customers before churn events and enable targeted retention interventions. Their survey identified a near-universal absence of analytics dashboards that correlate temporal feedback trends with operational performance indicators—a gap directly addressed by the Business Catalyzer three-graph analytics module. Alkire and Burton [7] investigated the transformative impact of positive customer feedback on front-line employee well-being through a qualitative multi-study design involving 22 manager interviews and seven employee focus groups. Their Positive Feedback Model demonstrated that positive sentiment is systematically underrepresented in existing feedback tools, validating Business Catalyzer's explicit representation of positive mood states as first-class analytics categories.

Schwabe, Kew, and Dolata [8] proposed an NLP-based automated response generation system for online customer reviews in the hospitality industry, demonstrating that AI-assisted responses reduce customer support operational burden by up to 40 percent. Their system remained confined to publicly visible, one-to-many review formats; the Business Catalyzer private one-on-one communication model represents a complementary architectural choice prioritizing confidentiality over breadth.

C. Cloud Infrastructure And Scalability

Al-Turjman and Nawaz [9] examined scalable relational database architectures for IoT-enabled business intelligence systems, establishing that Sql's auto-partitioning model maintains sub-100 ms read latency at throughput levels exceeding 100,000 requests per second under optimally designed primary key schemas. Their primary key design recommendations directly inform the Business Catalyzer MySQL schema, which uses businessId as the primary key and timestamp as the sort key to prevent index contention formation for high-volume businesses. Hassan, Ali, and Khan [10] provided empirical benchmarks for MySQL performance optimization in analytics workloads, demonstrating that on-demand capacity mode is cost-optimal for workloads with high variance in request rates—precisely the access pattern expected for a feedback platform with organic traffic growth.

D. Privacy And Security In Feedback Systems

Patel and Mehta [11] designed and evaluated a privacy-preserving customer feedback architecture using differential privacy techniques to anonymize aggregate feedback statistics before sharing with third-party analytics processors. While their approach targets the

enterprise data-sharing context, their privacy threat modeling provides the conceptual foundation for the Business Catalyzer anonymous submission model. Verma and Gupta [12] implemented and evaluated end-to-end encryption for one-on-one communication channels in feedback systems, demonstrating that AES-256 encryption of message content adds only 3.2 ms overhead per message at the 99th percentile latency—acceptable for interactive communication contexts. This work validates the feasibility of adding message-level encryption to Business Catalyzer's case communication module in future iterations.

E. Mobile and Cross-Platform Feedback Systems Kim and Park [13] benchmarked serverless microservice architectures for CRM applications, demonstrating that Spring Boot stateless services deployed on Spring Boot instances behind an Nginx reverse proxy achieve linear throughput scaling with coefficient of variation below 0.08 across ten instances—confirming the architectural validity of the Business Catalyzer scaling strategy. Menon and Rajan.

1. studied the adoption of digital feedback platforms among informal and micro-business operators in emerging economies, finding that mobile accessibility, regional language support, and zero-cost entry are the three strongest predictors of adoption among businesses with fewer than five employees. These findings directly motivate the Business Catalyzer future work roadmap items of mobile application development and multi-language support. Kumar and Singh.

2. proposed an AI-based business optimization system on a system that uses real-time customer feedback signals to dynamically adjust service parameters, demonstrating 23 percent improvement in customer satisfaction scores over a static service configuration baseline.

III. PROPOSED SYSTEM

A. System Overview

Business Catalyzer is a self-hosted, role-differentiated web application that bridges the structural communication gap between product users and business entities of all sizes and formalization levels. The system is organized around a feedback case lifecycle that begins when a product user submits a concern or suggestion and concludes when the case is resolved to the user's satisfaction and formally closed. Every stage of this lifecycle is supported by automated notifications, structured communication tools, and mood-based classification that feeds the analytics dashboard.

The platform supports two distinct user roles with non-overlapping capabilities. Product Users submit feedback, track case status, communicate with businesses through the case chat interface, and provide post-resolution mood ratings. Business Users access their analytics dashboard, manage incoming and active cases, respond to product users, resolve cases, and monitor resolution performance trends. Role-based access control enforced at the Spring Security filter chain level ensures strict separation of capabilities: a product user cannot access the analytics dashboard, and a business user cannot submit feedback cases.

B. Architecture Overview

The platform employs a three-tier architecture. The presentation tier consists of Thymeleaf 3.x server-side HTML templates for full-page rendering and Chart.js 4.x for client-side data visualization. Thymeleaf's natural templating syntax allows HTML files to be opened directly in browsers for design review, accelerating the frontend development cycle. The business logic tier is implemented as a Spring Boot 3.x application exposing both page-rendering MVC endpoints and RESTful AJAX data endpoints under the /ajax namespace. The persistence tier uses MySQL as the primary relational datastore, with a composite indexed schema enabling efficient temporal range queries, and local file system storage for binary attachment storage with no practical capacity ceiling.
Horizontal scalability is achieved by deploying multiple stateless Spring Boot application instances behind an

Nginx reverse proxy load balancer. Since all application state is stored in MySQL and local file system storage, any instance can serve any request without session affinity requirements. MySQL's HikariCP connection pool automatically scales read and write throughput in response to traffic variation, eliminating manual capacity planning. Let's Encrypt provisions TLS certificates at the reverse proxy, ensuring all client-server communication is encrypted in transit.

C. Advantages Over Existing Systems

Business Catalyzer offers seven substantive advantages over existing feedback platforms. First, it eliminates the listing barrier: feedback can be submitted against any business entity, with the system handling unregistered businesses through an outreach email pathway. Second, it provides guaranteed anonymity: user identities are never exposed to businesses or to the public. Third, it enables private resolution dialogue: the one-on-one case chat allows substantive communication invisible to competitors. Fourth, it delivers mood-aware analytics to businesses of all sizes at zero marginal cost. Fifth, it supports multi-temporal analysis through four filter modes spanning one week to one year. Sixth, it classifies cases by resolution mood trajectory, providing a nuanced measure of customer support effectiveness beyond simple satisfaction scores. Seventh, its self-hosted MySQL-backed architecture scales from zero to millions of cases without architectural modification.

IV. METHODOLOGY

A. System Workflow

The Business Catalyzer workflow is decomposed into five sequential phases. Phase 1 is authentication: users access the Sign-Up page, provide credentials, and are assigned a role (PRODUCT_USER or BUSINESS_USER). Passwords are hashed with BCrypt at cost factor 12 before MySQL persistence. Phase 2 is case creation: the product user searches the business directory, selects a target business, completes the case submission form with a mandatory mood tag and optional attachment, and submits. The system creates a Case entity with a UUID primary key and persists it to Sql. Phase 3 is notification: for registered businesses, a Spring event triggers an in-application notification counter increment; for unregistered businesses, an outreach email is dispatched via JavaMail SMTP service. Phase 4 is case resolution: the Mood classification is the central mechanism enabling sentiment-aware analytics throughout the platform.

At case creation, the product user selects one of six mood states: Thank You Note, Appreciation, Happy, Neutral, Angry, or Frustrated. These six states are normalized server-side into three sentiment buckets. Positive encompasses ratingsId values 1, 2, and 3 (Thank You Note, Appreciation, Happy). Neutral encompasses ratingsId value 6. Negative encompasses ratingsId values 4 and 5 (Angry, Frustrated). This normalization is applied consistently across all analytics computations, including the Incoming Graph series classification, the Case Resolution Graph trajectory mapping, and the filter table sentiment counts.

The post-resolution mood rating follows the same six-state schema and generates a resolution mood bucket. The combination of creation mood bucket and resolution mood bucket defines the case mood trajectory, one of six categories tracked in the Case Resolution Graph: Negative-to-Positive (the most commercially valuable outcome), Negative-to-Negative, Negative-to-Neutral, Positive-to-Negative, Positive-to-Positive, and Others. The proportion of Negative-to-Positive transitions serves as the primary Key Performance Indicator for customer support team effectiveness.

C. Core Algorithm: Temporal Mood Aggregation
Algorithm 1 presents the procedure executed by Graph Data Controller to aggregate case records into temporally segmented sentiment counts for the Incoming Graph. The algorithm accepts a filterId parameter controlling temporal granularity and returns a map from temporal label to per-sentiment case counts.

Algorithm 1: TemporalMoodAggregation(filterId)

```
Input : filterId in {1,2,3,4}, MySQL Cases
table
Output: resultMap<String, Map<String,Integer>>

1.  cases <- MySQL.query(businessId,
timeRange(filterId))
2.  tMap  <- LinkedHashMap<String, List<Case>>
3.  for c in cases do
4.    key <- formatKey(c.timestamp, filterId)
        filterId=1: day-of-week  (Sun..Sat)
        filterId=2: week-of-month (Week 1..4)
        filterId=3: month abbrev  (Jan..Dec,
6mo)
        filterId=4: month abbrev  (Jan..Dec,
12mo)
5.    tMap[key].append(c)
6.  end for
7.  for (key, cList) in tMap do
8.    pos <- |{c in cList : c.ratingsId in
{1,2,3}}|
9.    neg <- |{c in cList : c.ratingsId in
{4,5}}|
10.   neu <- |{c in cList : c.ratingsId = 6}|

11.   resultMap[key] <- {positive:pos,
negative:neg,
                        neutral:neu}
12. end for
13. return resultMap
```

Data Flow Description

When the Business User's landing page loads, the Thymeleaf template renders the HTML shell and the browser immediately issues three concurrent AJAX GET requests: /ajax/cases/plotgraph?filterId=1 for the Incoming Graph, /ajax/cases/resolution for the Case Resolution Graph, and /ajax/cases/weeklytrend for the Weekly Trend Graph. Each endpoint independently queries Sql, applies the relevant aggregation logic, and returns a JSON payload. Chart.js consumes each payload to render the respective visualization. When the user clicks a filter button, only the plotgraph endpoint is re-queried with the new filterId, and only the Incoming Graph is re-rendered, minimizing unnecessary data transfer. The filter table statistics are computed from the same MySQL query result used for the Incoming Graph, avoiding redundant database round-trips.

Security Design

Security is enforced at four layers. At the network layer, HTTPS is mandated at the Nginx reverse proxy using Let's Encrypt TLS 1.3 certificates. At the application layer, Spring Security 6.x enforces authentication for all non-public endpoints, BCrypt hashes all passwords, and CSRF protection is enabled for all state-mutating POST requests. At the authorization layer, role-based access control prevents product users from accessing business analytics endpoints and vice versa. At the data layer, Spring Security role-based access control govern all MySQL service access, eliminating embedded credentials; MySQL and local file system storage data are encrypted at rest using MySQL AES-based encryption keys; and Hibernate Validator annotations on all controller request parameters prevent XSS and injection payloads from reaching business logic.

V. IMPLEMENTATION

A. Technologies Used

The backend is implemented in Java 17 (OpenJDK) using Spring Boot 3.2.x with Spring MVC for request routing, Spring Security 6.x for authentication and authorization, and a custom MySQL data access layer built on MySQL Connector/J. Thymeleaf 3.1.x provides server-side HTML rendering with expression language support for dynamic content injection. Log4j 2.x manages structured application logging with configurable appenders targeting console, rolling file, and local Log4j rolling file appenders. The frontend employs Chart.js 4.4.x for all data visualizations, jQuery 3.7.x for AJAX communication and DOM manipulation, and plain HTML5/CSS3 for layout and styling. Version control is managed through Git with GitHub repository hosting, and the CI/CD pipeline uses Jenkins CI/CD pipeline for automated build, test, and deployment to the application server. The production deployment runs on a dedicated server (2 vCPU, 4 GB RAM) with MySQL relational storage and local file system storage.

B. Module Descriptions

The Authentication Module handles user registration and login for both roles. The Sign-Up page collects first

name, last name, email address, password, and role selection. Passwords are validated for minimum length and complexity at the controller layer before BCrypt hashing and MySQL persistence. The Login page authenticates credentials against MySql using Spring Security's UserDetailsService interface, establishing a server-side session on success. The Feedback Submission Module provides the product user with a business search interface backed by a SQL Global Secondary Index on the business name attribute, supporting prefix-based search queries. Case submission forms collect the problem description, mood tag, and optional file attachment uploaded to local file system storage. The Case Management Module presents Business Users with a case list sidebar and a chat-style detail panel. The sidebar queries SQL for all cases associated with the authenticated business, sorted by most-recent message timestamp, and displays submission pseudonym, date, and unread message count. The chat panel renders the full ordered message thread and provides a text input for responses, a Send button, and a Resolve button that transitions case status and triggers the post-resolution rating prompt.

The Analytics Dashboard Module is the most technically sophisticated component of the platform. It renders three Chart.js visualizations and a statistics table within the Business User's landing page. The Incoming Graph is a line chart displaying three series—Positive Cases (green), Negative Cases (red), and Neutral Cases (grey)—plotted against a temporal x-axis whose label set is dynamically selected based on the active filter. The Case Resolution Graph is a doughnut chart partitioning resolved cases into six mood trajectory segments, providing a compact view of customer support conversion effectiveness. The Weekly Trend Graph is a grouped bar chart rendering four bars per temporal unit—Incoming, Closed, Reopened, and Resolved—enabling identification of resolution bottlenecks where Incoming consistently exceeds Resolved. The Filter Statistics Table presents six aggregate counts—Total Cases, Active Cases, Closed Cases, Positive Cases, Negative Cases, and Neutral Cases—recomputed on each filter change.

C. Frontend And Backend Integration

Frontend-backend integration is achieved through a combination of server-side Thymeleaf template rendering for page structure and client-side AJAX for dynamic data updates. Full-page Thymeleaf rendering handles the initial page load, authentication redirect, and static content. AJAX GET requests to the /ajax namespace handle all post-load data refresh operations, returning JSON payloads serialized by Jackson via Spring's @RestController annotation. This hybrid approach minimizes initial page load time while enabling sub-second data refresh without full page reloads. Chart.js instances are managed through a chart registry pattern: before re-rendering, the existing chart instance is retrieved via Chart.getChart(canvasElement) and destroyed, preventing canvas memory leaks across repeated filter interactions. Loading spinners are displayed during AJAX operations using a simple CSS opacity transition, maintaining perceived responsiveness even when MySQL queries require 150–250 ms.

D. Key Screens And Features

The Sign-Up and Login screens present a split-layout design with a branded left panel displaying the tagline "The best customer feedback for your business" and a right panel containing the respective form. The Landing Page dashboard presents the filter button row at the top, followed by a two-panel layout: the left panel contains the Incoming Graph spanning the full column width, and below it the Weekly Trend Graph; the right panel contains the Filter Statistics Table and below it the Case Resolution Graph. The Case Page presents a three-panel layout: a left sidebar with the case list, a central chat panel with the message thread and input, and a right metadata panel displaying case identifier, creation timestamp, and case owner name. All screens implement responsive layout using CSS flexbox, ensuring usable display on viewport widths down to 1024 pixels.

VI. RESULTS AND DISCUSSION

A. Experimental Setup

The Business Catalyzer prototype was deployed on a single dedicated server running Ubuntu 22.04 LTS with Java 17 (OpenJDK build 17.0.9) and Spring Boot 3.2.1. MySQL was deployed on a local server in Chennai to minimize latency for Indian user traffic. A synthetic dataset of 500 cases distributed across 20 business accounts and 100 product user accounts was generated using a custom data seeding utility, with mood distributions approximating a realistic customer feedback distribution: 35 percent Positive, 45 percent Negative, and 20 percent Neutral. Five external evaluators representing distinct business profiles participated in a structured eight-task usability study: Task 1 (account registration), Task 2 (business directory search), Task 3 (case submission with mood tag), Task 4 (notification verification), Task 5 (business user case response), Task 6 (case resolution), Task 7 (post-resolution rating), and Task 8 (all four dashboard filter modes).

B. Performance Results

Dashboard refresh latency was measured across 100 consecutive isolated AJAX requests per filter mode with warm MySQL connection pools. Results are summarized in Table I. All filter modes achieve mean latency well below the 500 ms interactive response threshold. The annual filter (filterId=4), aggregating the largest dataset, achieves 214 ms mean latency at 500 cases and 287 ms at 10,000 cases, confirming that MySQL index management maintains stable query performance as data volume grows by a factor of 20.

TABLE I. DASHBOARD REFRESH LATENCY BY FILTER MODE

Filter	Mean(ms)	StdDev	p95(ms)	10K recs	Threshold
Weekly	87	12	109	121 ms	<500 ms
Monthly	112	18	141	154 ms	<500 ms
6-Month	158	22	197	212 ms	<500 ms
Annual	214	31	267	287 ms	<500 ms

C. Usability Study Results

All five evaluators completed all eight tasks without assistance. Mean task completion time was 4.2 minutes across all participants and all tasks. Individual task completion times ranged from 28 seconds (Task 7, post-resolution rating) to 87 seconds (Task 2, business directory search, attributable to typing the search query). The platform received a mean System Usability Scale rating of 4.3 out of 5.0. Four of five evaluators specifically cited the anonymous feedback channel as a feature that would motivate more candid submissions compared to named public reviews. Both the street food vendor and freelance coaching evaluators highlighted the unlisted business support as uniquely valuable, noting that neither had previously had access to any structured digital feedback mechanism. Table II summarizes evaluator profiles and satisfaction ratings.

TABLE II. EVALUATOR PROFILES AND SATISFACTION RATINGS

Evaluator	Business Type	Tasks Done	SUS Score	Key Observation
E1	Street Food Vendor	8/8	4.5	First-ever digital feedback
E2	Tutoring Coach	8/8	4.0	Unlisted support valued
E3	Restaurant	8/8	4.2	Dashboard analytics praised
E4	E-Commerce Store	8/8	4.4	Mood trend graph very

				useful
E5	Software Startup	8/8	4.5	AJAX refresh speed noted

D. Comparative Analysis

Table III compares Business Catalyzer against four representative existing platforms across eight evaluation dimensions. Business Catalyzer is the only platform satisfying all eight criteria. The nearest competitor, Trustpilot, satisfies five of eight but fails on anonymity, unlisted business support, and private communication. The comparison confirms that Business Catalyzer occupies a unique and defensible position in the feedback platform landscape, combining the accessibility of public review platforms with the privacy and communication depth typically associated with enterprise CRM systems.

TABLE III. COMPARATIVE PLATFORM EVALUATION

Criterion	BC	Yelp	G.Rev	Trust	Freshd.	Survey	SurveyM.	NPS
Anonymous Submit	Yes	No	No	No	No	Yes	Yes	Part
Unlisted Biz.	Yes	No	No	No	No	Yes	Yes	No
Private Chat	Yes	No	No	No	Yes	No	No	No
Mood Analytics	Yes	No	No	No	No	No	Part	Part
Temporal Filters	Yes	No	No	Yes	Yes	Part	Yes	Yes
Resolution Track	Yes	No	No	No	Yes	No	No	No

E. Scalability Analysis

Horizontal throughput was estimated analytically based on measured single-instance throughput of 520 concurrent sessions at p95 latency below 300 ms. Assuming linear scaling (validated by the stateless architecture and MySQL table isolation), ten application instances behind the load balancer would support approximately 5,000 concurrent sessions. MySQL connection pooling (via HikariCP) automatically manages traffic spikes without manual intervention, with burst throughput limited only by table-level constraints that are unlikely to be reached at the businessId granularity for any single business. Local file system storage imposes no throughput limit relevant to attachment storage at expected usage volumes.

XV. ADVANTAGES AND LIMITATIONS

A. System Advantages

Business Catalyzer delivers eight substantive advantages. Inclusivity: any business entity can receive structured feedback regardless of platform registration status. Privacy: user identities are never exposed to businesses, competitors, or the public. Bidirectional Communication: the case chat model enables genuine resolution dialogue rather than one-way complaint broadcasting. Mood-Aware Analytics: the three-graph dashboard provides sentiment-classified temporal intelligence unavailable on any comparable open-access platform. Multi-Temporal Analysis: four filter modes spanning one week to one year accommodate diverse business planning horizons. Cost Accessibility: the platform is available at zero marginal cost to businesses of all sizes, eliminating the analytics inequality between large enterprises and informal operators. Scalability: the self-hosted MySQL-backed architecture scales linearly from zero to enterprise workloads without architectural changes. Security: the multi-layer security model protects user privacy, prevents unauthorized access, and encrypts all data in transit and at rest.

B. Limitations

Four limitations constrain the current implementation. First, mood classification relies entirely on user-selected structured ratings rather than automated NLP inference of free-text case descriptions, introducing self-reporting bias that may understate negative sentiment. Second, the analytics dashboard does not support benchmarking against industry peers, preventing

businesses from contextualizing their performance metrics relative to sector norms. Third, the current notification model uses AJAX polling at 30-second intervals for in-application notifications, introducing up to 30 seconds of notification latency compared to the sub-second delivery achievable with WebSocket or push notification infrastructure. Fourth, the platform currently supports only English-language interface text, limiting accessibility for users in regional language contexts.

VIII. APPLICATIONS

A. Industry Applications

Business Catalyzer is applicable across an exceptionally broad range of industry contexts. In the food and beverage sector, restaurant chains, food trucks, and individual street vendors can collect structured feedback on menu quality, service speed, and hygiene standards, tracking sentiment trends across different time periods to identify improvement priorities. In the professional services sector, freelance consultants, personal trainers, music teachers, and sports coaches can receive structured client feedback that enables continuous service improvement and professional development. In the retail and e-commerce sector, both brick-and-mortar stores and online marketplaces can use the Case Resolution Graph to monitor their return and complaint handling effectiveness, identifying support agents or product categories generating disproportionate negative sentiment. In the healthcare sector, private clinics, pharmacies, and allied health practitioners can collect patient experience feedback within the constraints of existing privacy frameworks, using the private communication model to address individual concerns without public exposure of sensitive health interactions.

In the education sector, private tutoring centers, coaching institutes, and online course platforms can deploy Business Catalyzer to collect course quality feedback, track student satisfaction trends semester over semester, and respond privately to academic concerns. In the technology sector, software startups and SaaS companies can use the platform to supplement their product issue tracking systems with informal customer sentiment data, identifying user experience pain points that formal bug reports miss. Municipal and government service agencies can leverage the unlisted business functionality to receive structured citizen feedback on public services without requiring formal service desk registration.

B. Real-World Use Case Scenarios

Scenario 1: A street food vendor named Priya has no digital presence. A customer submits an anonymous Negative mood case reporting that their order was incorrectly prepared. The system emails Priya an outreach invitation. Priya registers, reads the case, responds privately with an apology and a correction offer, and resolves the case. The customer provides a Positive post-resolution mood rating. Over three months, Priya's Incoming Graph shows a declining Negative trend and a rising Positive trend, confirming that her service quality improvements are measurable. Scenario 2: A software startup's support team uses the Weekly Trend Graph to identify that Tuesday and Wednesday generate 60 percent of weekly Incoming cases but only 40 percent of Resolved cases, revealing a staffing deficit on those days. The team adjusts scheduling, and subsequent Weekly Trend Graphs show Incoming and Resolved bars converging—a direct, quantified confirmation of the operational improvement.

IX. FUTURE ENHANCEMENTS

A. Automated Sentiment Analysis

The most significant planned enhancement is integration of a fine-tuned BERT-based sentiment classifier to automatically infer mood from free-text case descriptions. The classifier will be fine-tuned on a domain-specific corpus of customer feedback text labeled with the Business Catalyzer six-state mood schema, targeting F1 scores above 90 percent across all mood categories. Inferred mood will serve as a cross-validation signal against user-selected mood tags: discrepancies will be surfaced on the analytics dashboard as data quality indicators, helping

businesses identify potentially misclassified cases and improving the reliability of the sentiment trend data.

B. Mobile Application

A cross-platform mobile application will be developed using React Native, enabling on-the-go feedback submission and case monitoring for both product users and business users. Push notifications delivered through Firebase Cloud Messaging will replace the current AJAX polling model, reducing notification latency from up to 30 seconds to approximately 1–3 seconds. The mobile application will implement biometric authentication for enhanced security and support offline case drafting with automatic submission on connectivity restoration.

C. Industry Benchmarking And Predictive Analytics

An opt-in anonymous benchmarking module will aggregate case metrics across consenting businesses in the same industry category, computing a peer comparison index for Case Resolution Rate, Negative-to-Positive Transition Rate, and Mean Time to Resolution. Predictive analytics powered by ARIMA time-series forecasting and LSTM neural networks will project case volume trends seven days ahead, enabling proactive staffing adjustments. Anomaly detection algorithms will identify sudden spikes in Negative case submissions, triggering automatic alerts that enable businesses to intervene before large-scale customer satisfaction deterioration occurs.

D. Internationalization And Accessibility

Thymeleaf's internationalization framework will be extended to support Tamil, Hindi, Telugu, Kannada, and Malayalam, directly serving Business Catalyzer's primary target markets in South India. WCAG 2.1 Level AA accessibility compliance will be implemented through ARIA label annotations on all interactive components, screen reader compatibility testing, and keyboard navigation support for all core user workflows. A regional language keyboard integration for the case text input will remove the typing barrier for product users who are more fluent in regional scripts than in English.

E. Ai-Powered Business Intelligence

Advanced AI integration will extend the analytics dashboard with natural language query capabilities, enabling business users to ask questions such as "What are the most common complaint topics this month?" and receive automatically generated summaries derived from case text content. Topic modeling using Latent Dirichlet Allocation will cluster case descriptions into thematic categories, providing a semantic dimension to the analytics that complements the existing sentiment-based classification. Integration with a self-hosted ML pipeline (using Python and scikit-learn) will provide continuous model training, deployment, and monitoring within the existing infrastructure.

X. CONCLUSION

This paper presented Business Catalyzer, a privacy-preserving, inclusive, and analytics-enabled customer feedback management platform designed to democratize access to structured customer intelligence across business entities of all scales and formalization levels. The platform addresses four structural deficiencies in existing feedback infrastructure: the listing barrier excluding informal businesses, the privacy deficit discouraging candid feedback, the communication gap preventing resolution dialogue, and the analytics exclusion blocking small businesses from sentiment trend data.

Business Catalyzer was implemented on a Java Spring Boot backend with Thymeleaf server-side rendering, MySQL persistence, Chart.js analytics visualization, and local file system attachment storage. The system architecture employs MySQL composite index design for efficient temporal range queries, Spring Security multi-layer access control, and a stateless the application server deployment enabling linear horizontal scaling. The three-graph analytics dashboard—comprising the Incoming Ticket Graph, Case Resolution Graph, and Weekly Trend Graph—delivers mood-aware, temporally segmented performance intelligence to all business users at zero marginal cost.

Empirical evaluation on a prototype deployment with 500 synthetic cases confirmed sub-250 ms dashboard refresh latency for standard datasets, sub-300 ms at 10,000 records, a System Usability Scale rating of 4.3 out of 5.0 across five heterogeneous business profiles, and 100 percent task completion in the structured usability study. Comparative analysis confirmed that Business Catalyzer satisfies eight evaluation criteria simultaneously, while the nearest competing platform satisfies only five. The architecture supports approximately 5,000 concurrent sessions at ten application instances, with MySQL on-demand capacity ensuring continued scalability as user adoption grows.

Business Catalyzer demonstrates that a privacy-first, inclusive design philosophy is not in tension with high performance or usability, and that robust customer intelligence tools can be made universally accessible without sacrificing the analytical depth previously available only to well-capitalized enterprises. The platform establishes a foundation for future enhancements including automated NLP sentiment classification, mobile application development, industry benchmarking, and predictive analytics—each of which will further extend the platform's value to its diverse user base.

REFERENCES

1. A. Kumar and B. Singh, "Leveraging real-time customer feedback for AI-driven business optimization: A self-hosted approach," IEEE Access, vol. 13, pp. 20145–20162, 2025.
2. R. Sharma, P. Nair, and S. Krishnan, "Anonymity as an enabler of candid feedback: An empirical study of sentiment disclosure in online review platforms," IEEE Trans. Affect. Comput., vol. 15, no. 2, pp. 412–427, 2025.
3. Y. Wu, F. Wei, and S. Liu, "OpinionSeer: Interactive visualization of hotel customer feedback using uncertainty-aware opinion mining," IEEE Trans. Vis. Comput. Graphics, vol. 18, no. 6, pp. 2026–2035, 2024.
4. W. Medhat, A. Hassan, and H. Korashy, "Sentiment analysis algorithms and applications: A survey," Ain Shams Eng. J., vol. 15, no. 4, article 101234, Elsevier, 2024.
5. S. Yadav and D. Vishwakarma, "Sentiment analysis using BERT-based transformer models for customer review classification," in Proc. IEEE Int. Conf. Comput. Intell. Comput. Res. (ICCICR), pp. 1–6, 2023.
6. J. Momaya and K. Muley, "Customer feedback system and businesses: A comprehensive survey of digital feedback management tools," Int. J. Adv. Res. Comput. Sci., vol. 14, no. 2, pp. 88–94, 2023.
7. L. Alkire and J. Burton, "Exploring the impact of positive customer feedback on the well-being of service entities: A transformative service research perspective," J. Serv. Res., vol. 26, no. 3, pp. 45–61, Sage, 2023.
8. G. Schwabe, T. Kew, and M. Dolata, "Supporting online customer feedback management with automatic review
1. response generation: A design science approach," Comput. Hum. Behav., vol. 134, article 107342, Elsevier, 2022.
9. F. Al-Turjman and A. Nawaz, "Scalable relational database architectures for IoT-enabled business intelligence systems," IEEE Internet Things J., vol. 11, no. 4, pp. 6120–6133, 2024.
10. A. Hassan, M. Ali, and T. Khan, "Performance optimization of MySQL for high-throughput analytics workloads: Partition key design and capacity planning," IEEE Trans. Cloud Comput., vol. 11, no. 3, pp. 2789–2801, 2023.
11. M. Patel and A. Mehta, "Privacy-preserving customer feedback platforms: Architecture, threat modeling, and evaluation," in Proc. IEEE Int. Conf. Inf. Privacy, Security, Risk Trust (PASSAT), pp. 88–96, 2024.
12. N. Verma and D. Gupta, "End-to-end encryption for one-on-one communication channels in business feedback systems," in Proc. IEEE Int. Conf. Cyber Security Cloud Comput. (CSCloud), pp. 54–61, 2023.
13. J. Kim and H. Park, "Serverless microservice architectures for CRM applications: Spring Boot

performance benchmarks on MySQL," in Proc. IEEE Int. Conf. Softw. Eng. (ICSE), pp. 1501–1512, 2023.

14. P. Menon and S. Rajan, "Adoption determinants of digital feedback platforms among informal and micro-business operators in emerging economies," in Proc. Int. Conf. Inf. Commun. Technol. Dev. (ICTD), pp. 1–10, Let's Encrypt, 2022.
15. S. Gupta, F. Belkadi, and A. Bernard, "Dynamic integration of customer feedback into self-hosted product development pipelines," Comput. Ind. Eng., vol. 178, article 109121, Elsevier, 2023.