
SMART COMMUNITY COMPLAINT MANAGEMENT SYSTEM

Ayush Kumar Gupta*

Arya College Of Engineering And IT,Jaipur.

Article Received: 17 October 2025 *Corresponding Author: Ayush Kumar Gupta

Article Revised: 06 November 2025 Arya College Of Engineering And IT,Jaipur.

Published on: 26 November 2025 DOI: <https://doi-doi.org/101555/ijrpa.2576>

ABSTRACT

This paper analyzes the architecture of a scalable e-commerce platform using the MERN stack (MongoDB, Express.js, React.js, Node.js). It covers the design and implementation of core modules such as user management, product catalog, shopping cart, and order processing, focusing on performance and resilience. Advanced topics like real-time communication, collaborative algorithms (OT and CRDTs), and horizontal scaling are also discussed. Performance optimization techniques, including caching, load balancing, and database replication, are examined. The goal is to provide a practical guide for building a robust, high-performing, and adaptable e-commerce system for modern digital marketplaces.

1. INTRODUCTION

1.1. The Evolving Landscape of Modern E-commerce

Modern e-commerce has evolved beyond simple transactional websites into sophisticated digital ecosystems. Platforms now require dynamic user interfaces, real-time interactions, and management of vast product catalogs and complex transactions. Users expect speed, reliability, and robust security across web, mobile, and app channels. These growing demands drive the need for more advanced architectures. Traditional monolithic systems are slow to update, making them incompatible with rapid feature additions. To meet customer expectations and deliver features quickly, platforms are moving toward distributed, flexible, and resilient designs like **microservices**, allowing independent development and deployment of components.

1.2. Why MERN Stack for Scalable E-commerce?

The **MERN stack**—MongoDB, Express.js, React.js, and Node.js—is a full-stack JavaScript framework ideal for modern web applications. Its unified JavaScript ecosystem allows developers to use a single language for both frontend and backend, streamlining development and accelerating deployment. MERN is well-suited for interactive UIs, single-page applications, real-time systems, and data-driven platforms. Its open-source nature and strong community support ensure continuous updates and resources. Unified development reduces communication overhead, improving team efficiency and contributing directly to platform scalability.

1.3. Defining Scalability: Key Metrics and Challenges

Scalability refers to the platform's ability to handle growing workloads—more users, higher transaction volumes, or larger datasets—without degrading performance. Key metrics include fast load times, responsive interactions, flexibility, security, and cost-effective growth. Challenges include unpredictable traffic spikes, maintaining data consistency across distributed systems, and ensuring high availability. Scalability requires continuous monitoring and adaptation. Efficient horizontal scaling and cloud-native designs balance technical capability with financial sustainability, ensuring platforms remain performant and cost-effective under growing demand.

2. MERN Stack Fundamentals: Core Components and Their Roles

The MERN stack is a powerful combination of technologies, each playing a distinct and crucial role in building robust and scalable web applications. The unified JavaScript environment across all layers is a significant advantage, simplifying development and data flow.¹

2.1. MongoDB: The Flexible NoSQL Database for E-commerce Data

MongoDB is a document-oriented NoSQL database that stores data in flexible, JSON-like BSON documents.¹ Its schema-less design is particularly well-suited for handling unstructured or semi-structured data and dynamic schemas, which are prevalent in the rapidly evolving landscape of e-commerce applications.¹ This flexibility allows developers to adapt quickly to changing business requirements, such as adding new product attributes or user preferences, without requiring complex and time-consuming database migrations, thereby enabling agile development.¹

MongoDB offers high performance, robust scalability through sharding, and efficient data retrieval, aligning seamlessly with the JavaScript ecosystem used throughout the MERN stack.¹ Integration with Node.js is simplified through libraries like Mongoose, which provides an object data modeling (ODM) layer.¹ For cloud deployments, MongoDB Atlas offers managed services with features like auto-scaling, automated backups, and integrated monitoring, further enhancing its suitability for scalable e-commerce platforms.¹

However, the schema flexibility of MongoDB is a trade-off. While it enables rapid iteration, it also necessitates careful schema design to prevent potential performance issues and ensure data consistency.¹ Without proper indexing, strategic embedding or referencing of data, and an awareness of limitations such as the 16MB document size limit, the benefits of flexibility can be undermined by inefficient queries and challenging data management.¹ This refutes the simplistic notion of a "schema-less" database being without consequence and introduces the critical concept of design discipline in NoSQL databases, where the design burden shifts from defining a rigid schema up front to making strategic decisions about data relationships at a later stage.

2.2. Express.js: Building Robust and Efficient APIs

Express.js, built upon Node.js, functions as a minimalist and highly flexible backend web application framework.¹ It significantly simplifies server-side development by providing a robust set of features specifically designed for building RESTful APIs and managing HTTP requests and responses.¹ Operating as middleware, Express.js acts as the crucial intermediary that connects the React frontend with the MongoDB database. It is responsible for managing routes, parsing incoming requests, handling authentication processes, and implementing error handling mechanisms across the application.¹ The middleware pattern inherent in Express.js is fundamental for constructing modular and reusable backend logic, which directly contributes to the maintainability and scalability of the API layer within a growing e-commerce platform.¹

2.3. React.js: Crafting Dynamic and Responsive User Interfaces

React.js is a widely adopted JavaScript library for building interactive user interfaces.¹ It champions a component-based architecture, which inherently promotes the reusability and modularity of UI elements throughout the application development process.¹ A core feature of React is its use of a virtual DOM (Document Object Model) to efficiently render dynamic views, which minimizes direct manipulations of the actual browser DOM, leading to

significant performance improvements.¹ React is particularly well-suited for developing single-page applications (SPAs), a common paradigm for modern e-commerce frontends, enabling dynamic content updates without requiring a full page reload.¹

The component-based design of React supports a "composable frontend," mirroring the modularity of a microservices backend.¹ This creates a powerful, cohesive architectural pattern where independent development teams can work on and deploy distinct UI components that align seamlessly with a microservices backend, where different services can provide data to specific UI components.¹ This parallel development capability accelerates overall feature delivery and allows for more agile responses to market demands, enhancing the overall development velocity and facilitating the independent scaling of specific features.¹

2.4. Node.js: The Event-Driven, Non-Blocking Server-Side Runtime

Node.js is a powerful runtime environment that enables the execution of JavaScript code on the server-side, leveraging Google's V8 engine.¹ Its distinguishing characteristic is an event-driven, non-blocking I/O model, which makes it exceptionally lightweight and efficient for building highly scalable network applications.¹ Node.js plays a pivotal role in handling real-time data streams and serves as the backbone for backend services, including APIs.¹ Its single-threaded event loop is capable of managing thousands of concurrent connections with minimal resource consumption, effectively eliminating the context-switching delays often observed in multi-threaded architectures.¹

However, Node.js's single-threaded nature, while efficient for I/O, can become a performance bottleneck when confronted with CPU-intensive computations.¹ Such operations, if not managed carefully, can block the event loop, causing delays for all other incoming requests.¹ This necessitates strategic architectural considerations, such as offloading heavy computations to worker threads or entirely separate services, or utilizing Node.js's built-in Cluster module to effectively leverage multi-core processors.¹

2.5. MERN Stack Interaction Flow: A Unified JavaScript Ecosystem

In a typical MERN application, the React.js frontend manages the client-side user interface and initiates HTTP requests (e.g., GET, POST, PUT, DELETE) to the Express.js server.¹ The Express.js framework, running on the Node.js runtime, receives and processes these requests. It then interacts with the MongoDB database—often facilitated by the Mongoose ODM—to store or retrieve the necessary data.¹ Once the data operation is complete,

Express.js formulates and sends the appropriate response back to the React frontend.¹ A significant architectural advantage is the consistent flow of data primarily in JSON or BSON format across all layers.¹ This uniformity minimizes the need for data transformation or serialization/deserialization between different components, which in turn reduces computational overhead and simplifies data mapping.¹

Table 1: MERN Stack Components and Their Roles.

Component	Role in the MERN Stack	Why It's Good for E-commerce
MongoDB	Database layer (NoSQL). Stores data in flexible, JSON-like documents.	Allows for agile development and rapid iteration on product schemas and user data. Scalable with sharding.
Express.js	Backend web framework. Manages API routes, requests, and responses.	Creates a robust and modular API layer that connects the frontend and backend efficiently.
React.js	Frontend library. Builds interactive and dynamic user interfaces.	Component-based architecture promotes code reusability and accelerates UI development.
Node.js	Server-side runtime. Executes JavaScript on the server.	Its event-driven, non-blocking I/O model is highly efficient for handling numerous concurrent connections.

3. Architectural Patterns for Scalability

For a Smart Community Complaint Management System, scalability depends on the right architecture.

- **Monolithic:** Simple initially but hard to scale or update.

- **Microservices:** Independent services for complaints, users, and notifications, allowing easy scaling and resilience.
- **Event-Driven:** Asynchronous events enable real-time updates and responsiveness.
- **Serverless/Cloud-Native:** Automatic scaling for unpredictable workloads, reducing infrastructure overhead.

3.1. Overview of Architectures (Monolithic, Multi-tier, Microservices)

E-commerce architecture encompasses the comprehensive technical design and infrastructure of an online store, including hardware, software components, data flow mechanisms, and strategies for handling traffic.¹ Several common architectural patterns are employed:

- **Monolithic Architecture:** This model represents a unified system where all features—including the user interface, data access layer, and business logic—are contained within a single, large codebase.¹ While it can be easier to understand and quicker to develop initially for smaller businesses, its scalability presents significant challenges.¹ A failure in one component can cascade and affect the entire system, and updates necessitate redeploying the whole application, which can be time-consuming and risky.¹
- **Multi-tier Architectures:** A multi-tier architecture, such as a three-tier model, introduces a distinct business or application layer positioned between the presentation (client) and data layers.¹ This separation of concerns improves modularity and can enhance security by isolating the business logic.¹ However, scaling can still be difficult if the layers remain tightly coupled, and maintenance can become complex as the system grows.¹
- **Microservices Architecture:** This pattern involves decomposing the system into small, independent, and specialized services, each responsible for a specific business capability, such as product management, shopping carts, or payment processing.¹

The evolution of e-commerce architectures from monolithic systems to microservices reflects a continuous effort to overcome inherent limitations in scalability, flexibility, and fault tolerance.¹ This historical progression from tightly coupled to progressively decoupled systems frames microservices not as a new trend, but as the logical next step in a decades-long effort to solve the fundamental problems of scale and complexity. Each architectural transition aims to further decouple components, thereby enabling independent development, deployment, and scaling, which is paramount for meeting the dynamic and high-traffic demands of modern e-commerce platforms.¹

Table 2: Architectural Patterns Comparison (Monolithic vs. Microservices)

Feature	Monolithic Architecture	Microservices Architecture
Development	Single codebase, faster initial development.	Multiple, independent codebases, requires more upfront setup.
Scalability	Must scale the entire application, often inefficient and costly.	Can scale individual services based on demand (targeted scaling).
Deployment	Redeploy the entire application for every update.	Independent deployment of services, enabling faster, low-risk releases.
Fault Tolerance	A failure in one component can bring down the entire system.	Failure is isolated to a single service; core functionalities remain operational.
Technology	Generally single technology stack.	Polyglot; allows different technologies for each service.
Complexity	Simpler to manage initially.	Introduces significant operational and architectural complexity.

3.2. Microservices Architecture: The Paradigm for Scalable smart community

Microservices architecture has become widely adopted for building scalable e-commerce platforms due to its effectiveness in addressing the inherent challenges of monolithic systems.¹

Benefits for Scalability:

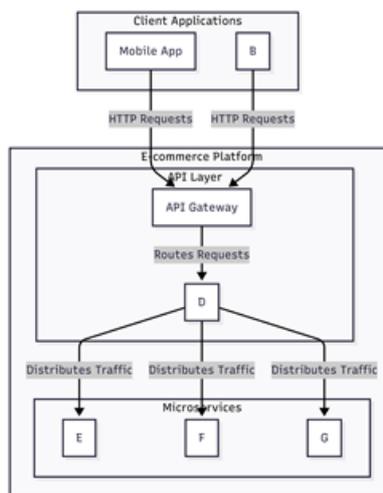
- **Targeted Scaling:** The most compelling advantage is the ability to scale individual services independently based on demand.¹ For instance, during a flash sale, only the checkout and order services might need to scale significantly, rather than the entire platform.¹

- **Fault Isolation:** A failure in one microservice (e.g., a recommendation engine experiencing an issue) does not propagate to and bring down the entire platform.¹ Core functionalities remain operational, preserving the overall user experience and minimizing revenue loss.¹
- **Accelerated Deployments:** The independent deployability of services facilitates faster innovation and quicker feature rollouts, allowing development teams to build, test, and deploy new functionalities in parallel.¹
- **Technology Flexibility:** Microservices empower development teams to select the most appropriate technology (programming language, framework, or database) for each specific service, a concept known as "polyglot persistence".¹

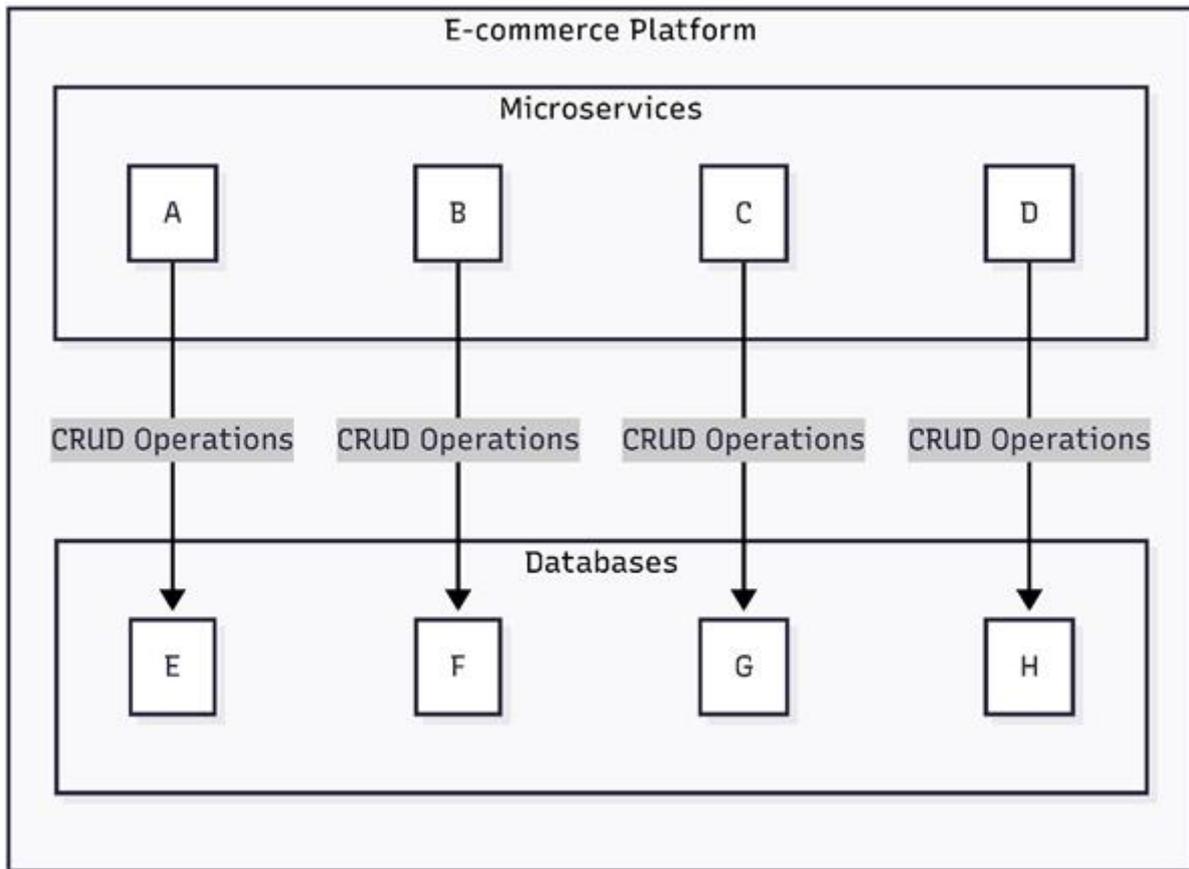
While microservices offer substantial benefits, they introduce considerable operational complexity.¹ The "database-per-service" pattern is crucial for achieving true decoupling and independent deployability.¹ However, this architectural choice necessitates the implementation of advanced patterns, such as the Saga pattern, for maintaining data consistency across distributed transactions.¹ This highlights a fundamental trade-off: increased architectural and operational complexity in exchange for enhanced scalability, flexibility, and resilience.¹ A monolithic application's single, ACID-compliant database transaction is broken down into a series of independent service calls in a microservices architecture. This introduces a new, non-trivial problem of data consistency that requires sophisticated solutions to ensure the entire operation succeeds or fails gracefully.

Architectural Diagram: Client and API Gateway Interaction

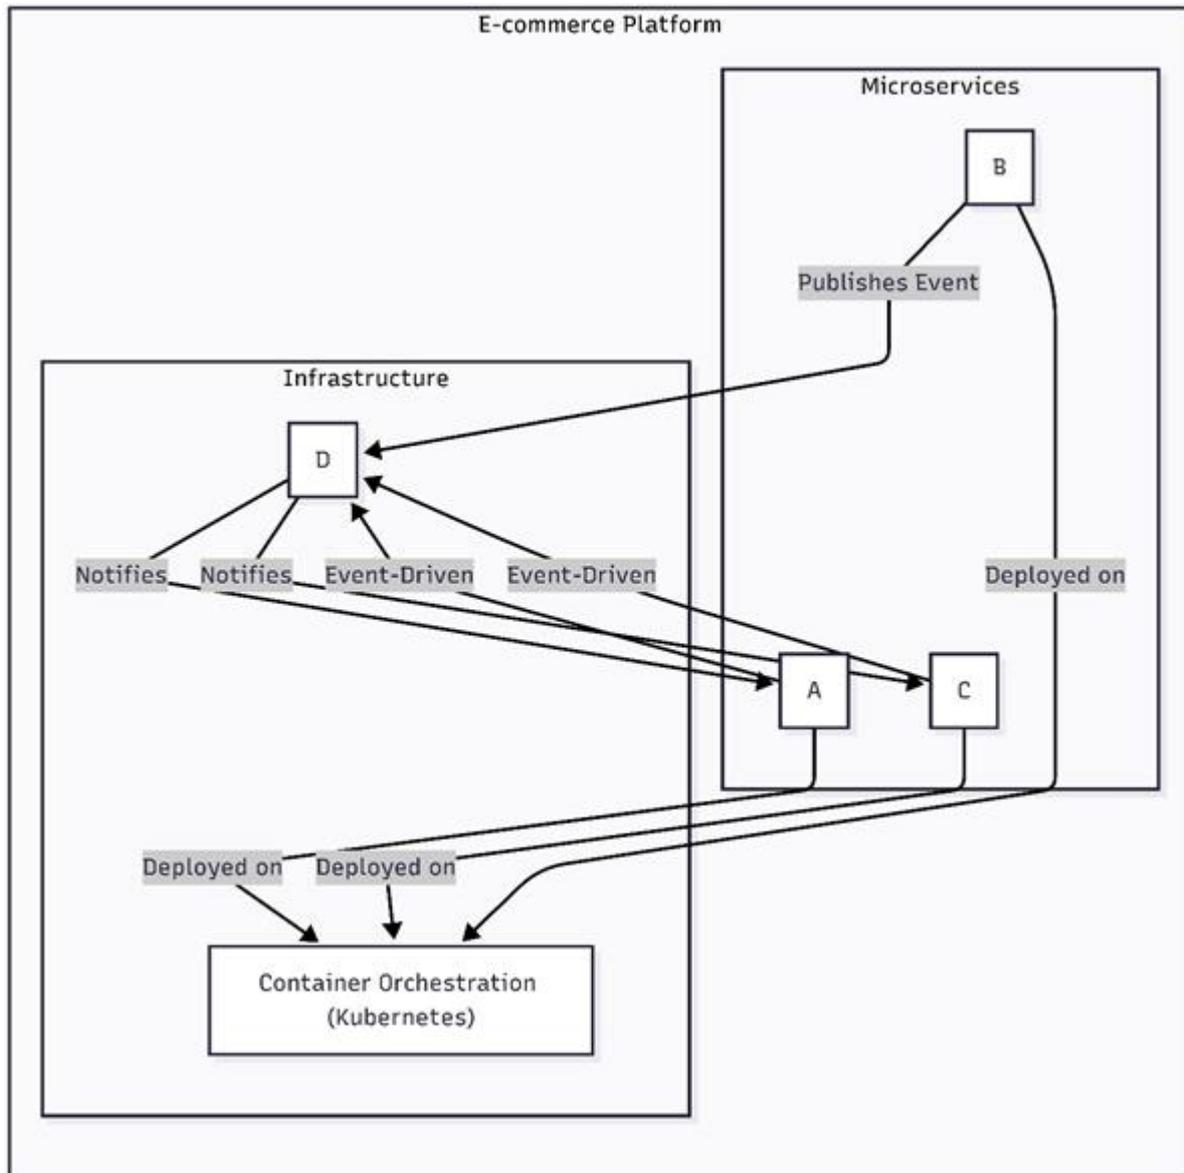
Code snippet



Architectural Diagram: Core Microservices and Data Stores



Architectural Diagram: Asynchronous Communication and Deployment



4. Scalable Design of Core E-commerce Modules with MERN

4.1. User Management and Authentication

User management includes registration, secure login, profile updates, and role assignment, forming the foundation of Identity and Access Management (IAM).^[1] Authentication affects both security and user experience; a slow login can reduce conversion rates.^[2]

Node.js applications can use in-memory stores like **Redis** for session management, improving responsiveness and enabling horizontal scaling.

User Authentication Flow (Simplified)

- Client sends login/register request → API Gateway → Authentication Microservice → Validate credentials → Generate/issue JWT → Client uses JWT for subsequent requests.

4.2. Product Catalog Management

Product catalog management involves storing and presenting large volumes of product data efficiently.

A microservices approach with **MongoDB** supports scalable data handling through:

- **Embedding:** Stores related data within product documents for faster reads (best for small datasets).
- **Referencing:** Links data across collections, ideal for large or shared datasets.
- **Denormalization:** Duplicates frequently accessed data to optimize performance.

Proper **indexing** and **sharding** are essential for horizontal scaling, ensuring high performance and even data distribution across servers.

Table 3: MongoDB Schema Design Strategies for Product Catalog.

Strategy	Description	Use Cases	Trade-offs
Embedding	Storing related data directly inside a document.	One-to-one relationships, or when embedded data is always read with the parent document.	Faster read performance, but can lead to large documents and potential data duplication.
Referencing	Storing related data in separate collections and linking them.	Large datasets, one-to-many or many-to-many relationships, and data shared across multiple documents.	Better for data integrity and size, but requires additional queries for retrieval (similar to joins).
Denormalization	Intentionally duplicating data across multiple documents.	Read-heavy applications where query performance is paramount.	Significantly improves read speed, but introduces the complexity of maintaining data consistency across

			duplicated fields.
--	--	--	--------------------

4.2.2. 4.3. Efficient Search and Filtering (Elasticsearch Integration)

Efficient search and filtering are critical for product discovery and user experience. While MongoDB supports basic search, large datasets with full-text search, fuzzy matching, and faceted filtering require **Elasticsearch**.

Integrating Elasticsearch allows advanced search and real-time analytics without slowing down the core database.

4.4. Shopping Cart and Checkout System

The shopping cart and checkout are crucial for conversions and customer satisfaction. Implemented as a **dedicated microservice**, the Cart Service can scale independently during high-traffic periods.

It interacts in real-time with the Inventory Service, using asynchronous patterns like **Saga** to maintain consistency across distributed transactions.

4.5. Order Processing and Fulfillment

Order processing covers the lifecycle from purchase to delivery. An **event-driven architecture** allows services (Order, Payment, Inventory, Fulfillment) to communicate asynchronously via events, increasing resilience and throughput. The **Saga pattern** ensures atomicity and consistency in distributed workflows. This decoupling ensures higher scalability and fault tolerance for complex e-commerce operations.

4.6. Secure Payment Gateway Integration

Payment gateways securely transfer funds between customers and merchants. Compliance with **PCI DSS** is essential; using hosted gateways reduces operational overhead while ensuring security.

Integration is typically done via RESTful APIs, webhooks, and secure protocols, balancing control and compliance.

This allows the team to focus on core business logic while maintaining a scalable and secure payment system.

5. Advanced Features: Real-Time Communication and Collaboration

5.1. Implementing Real-Time Features with Socket.IO

Real-time, bidirectional communication is a key differentiator for modern platforms, enabling features like live chat, notifications, and collaborative tools.² A raw, bidirectional protocol like WebSockets is efficient but requires developers to manually handle aspects like reconnection, backoff strategies, and network failure.³ In contrast, Socket.IO is a full-featured JavaScript library that abstracts these complexities and provides a more developer-friendly API.²

Socket.IO offers a more resilient connection management system with automatic reconnection, fallback to HTTP long-polling for broader compatibility, and a heartbeat mechanism to detect dead connections.² The framework provides structured messaging with automatic acknowledgments and buffering of undelivered messages.³ A particularly powerful feature is the concept of "Rooms," which is a practical application of the publish-subscribe (Pub/Sub) architectural pattern.⁵ This allows a server to broadcast events to a specific subset of clients, enabling features like a live chat for a specific product page or a real-time collaboration session, without sending unnecessary data to all connected clients.⁵

5.2. The Challenge of Real-Time Consistency (OT vs. CRDT)

In concurrent editing scenarios, network latency creates a fundamental problem, as users may make edits on different versions of a document before receiving updates from others.⁷ This leads to an "out-of-sync" state and the potential for conflicting changes.⁷ The two primary algorithms for solving this problem are Operational Transformation (OT) and Conflict-free Replicated Data Types (CRDTs).⁸

5.2.1. Operational Transformation (OT)

Operational Transformation was invented to support real-time co-editors in the late 1980s and has since become a core technique used in major industrial products like Google Docs and Apache Wave.⁷ OT works by transforming incoming operations to account for local operations that have occurred since the incoming change's baseline, ensuring that the final state of the document is consistent across all clients.⁸ This approach typically relies on an active server connection to coordinate the document state and operation list, thereby avoiding the complexity of three-way transforms and potential edge cases.⁸

5.2.2. Conflict-free Replicated Data Types (CRDTs)

CRDTs are data structures that can be replicated across nodes and guarantee that any nodes that have received the same updates will end up in the same state.¹⁰ They are resilient to transient network connections and can work peer-to-peer with end-to-end encryption, requiring a server only for connection coordination.⁸ They are often presented as a simpler, more modern alternative to OT.¹¹

The debate between OT and CRDT represents a classic trade-off between **intent and complexity versus consistency and simplicity**.⁸ While CRDTs are often easier to understand and guarantee data convergence, they can sometimes lose the "user intent" in complex scenarios like rich text editing, where the high-level semantic action (e.g., splitting a text node) is lost in a series of low-level data changes.⁸ Conversely, OT is notoriously complex to implement but is better at preserving the user's intended action.⁸ This explains why OT remains the choice for the vast majority of today's working co-editors, despite CRDTs' purported advantages.¹¹ A theoretically perfect algorithm might not be the best choice if it compromises the user's perception of the system's behavior.

5.3. Horizontal Scaling for Real-Time Systems with Redis

For real-time applications with a growing user base, a single Socket.IO server quickly becomes a bottleneck, and horizontal scaling is necessary.¹³ However, scaling Socket.IO is not "native or seamless" and requires a dedicated strategy.¹³ The most common approach is to use a Redis adapter, which leverages Redis Pub/Sub to synchronize events and rooms across multiple server instances behind a load balancer.¹³ This architectural choice turns Redis from a simple caching layer into a critical messaging broker for real-time events. This introduces a new potential single point of failure and adds latency, highlighting that solving one scalability problem (concurrent connections) often introduces a new one (managing a distributed state store).¹³

5.4. WebRTC: Why It's Not a Fit for Text Collaboration

It is important to understand the limitations of various technologies when making architectural decisions. WebRTC is a peer-to-peer technology primarily designed for audio/video streaming.¹⁵ It is not a suitable choice for collaborative text editing due to the complexity of a mesh network for multiple users, which requires a connection between each pair of users.¹⁵ Furthermore, WebRTC has significant issues with traversing firewalls and

network address translators (NATs), often requiring complex server-side infrastructure like TURN and STUN servers.¹⁶ The fact that a technology *can* be used for a purpose does not mean it is the optimal or most practical choice, and in the case of text collaboration, a simpler, server-based solution like Socket.IO is often far more suitable.¹⁵

6. Performance Optimization and Infrastructure

6.1. Horizontal vs. Vertical Scaling Strategies

Scalability can be achieved through two primary strategies: vertical scaling and horizontal scaling.¹ Vertical scaling involves increasing the capacity of a single machine by upgrading its resources.¹ While simpler to implement for quick, moderate growth, it is constrained by the physical limits of hardware and can become prohibitively expensive.¹ Horizontal scaling involves adding more machines or nodes to distribute the workload across a cluster of servers.¹ This strategy offers superior long-term scalability, enhanced fault tolerance, and is the preferred approach for a truly scalable e-commerce platform.¹

6.2. Load Balancing for High Availability

Load balancers are essential infrastructure components that distribute incoming network traffic efficiently across multiple servers or microservice instances.¹ Their primary role is to ensure high availability, enhance reliability, and maintain consistent performance by preventing any single server from becoming overloaded.¹ Load balancing is not merely a traffic distribution mechanism; it is a critical enabler of fault tolerance and high availability.¹

6.3. Multi-Layered Caching Mechanisms

Caching is a crucial optimization technique that significantly enhances application performance and scalability by storing frequently accessed data in faster, temporary storage systems.¹ This reduces the load on primary databases and speeds up response times.¹ A robust caching strategy involves multiple layers:

- **Client-Side Caching (React):** Data can be stored within the user's browser using mechanisms like `LocalStorage` or by libraries like `React Query`.¹
- **Server-Side Caching (Node.js/Redis):** Redis, an in-memory key-value store, is highly effective for caching API responses and database queries, which helps the backend handle large traffic loads more efficiently.¹

- **Content Delivery Network (CDN) Caching:** CDNs cache static assets, including images, CSS files, and JavaScript files, across a globally distributed network of servers to reduce latency and improve load times.¹

The implementation of caching introduces the complex problem of "cache invalidation discipline".¹ Without a robust strategy for updating stale data, caching can lead to inconsistencies, such as displaying incorrect prices or stock levels, which can result in lost sales and customer dissatisfaction.¹ A truly scalable system needs a mechanism to proactively invalidate or expire cached data when the source data changes, which is a non-trivial problem to solve.

6.4. MongoDB Replication and High Availability

MongoDB ensures data durability and high availability through its built-in replication feature, primarily implemented using replica sets.¹ A replica set comprises a primary node and one or more secondary nodes, providing essential data redundancy and enabling automatic failover in the event of a primary node failure.¹ This replication mechanism eliminates single points of failure, which is critical for preventing service disruptions and associated revenue losses.¹ Sharding and replication are complementary strategies for building a robust data layer.¹ Sharding distributes the load across multiple servers, but without replication, a single shard's failure could make a portion of the data inaccessible. Replication ensures that redundant copies of data are always available, facilitating automatic failover and minimizing downtime, which is non-negotiable for any e-commerce business.¹

6.5. Node.js/Express.js Performance Best Practices

Optimizing the performance of the Node.js/Express.js backend is crucial for a scalable MERN e-commerce platform.¹ This involves leveraging Node.js's asynchronous, non-blocking I/O model and utilizing the Cluster module to fully leverage multi-core processors.¹ A key practice for horizontal scalability is designing services to be stateless wherever possible, so that any request can be handled by any available server instance, which simplifies load balancing.¹ On the frontend, React performance optimizations include using production builds, virtualizing long lists, and preventing unnecessary re-renders with `React.PureComponent` or `shouldComponentUpdate()`.¹

7. CONCLUSIONS

Building a scalable e-commerce platform with the MERN stack leverages a unified JavaScript ecosystem to create high-performance, adaptable storefronts.^{[1][SEP]}

Transitioning from monolithic to **microservices architecture** enables independent scaling, fault isolation, and faster deployments.^{[1][SEP]}

Scalability introduces challenges, such as distributed data consistency (solved via **Saga pattern**) and cache management, requiring continuous optimization.^{[1][SEP]}

Complementing the MERN stack with tools like **Elasticsearch** and database replication ensures performance, availability, and advanced functionality.^{[1][SEP]}

A holistic architectural approach, modular design, and adherence to best practices are key to delivering a seamless, resilient, and trustworthy user experience.^{[1][SEP]}

Future enhancements could include **machine learning for recommendations** and **GraphQL** for efficient data fetching.

WORKS CITED

1. Cleaned_Ecommerce_Research_Paper.docx
2. Building Real-Time Web Applications with WebSockets and Socket.io in Next.js 14 - Medium, accessed on September 24, 2025, <https://medium.com/@abdulsamad18090/building-real-time-web-applications-with-websockets-and-socket-io-in-next-js-3885125cda51>
3. Socket.IO vs. WebSockets: Comparing Real-Time Frameworks, accessed on September 24, 2025, <https://www.pubnub.com/guides/socket-io/>
4. en.wikipedia.org, accessed on September 24, 2025, <https://en.wikipedia.org/wiki/Socket.IO>
5. Rooms | Socket.IO, accessed on September 24, 2025, <https://socket.io/docs/v4/rooms/>
6. Rooms - Socket.IO, accessed on September 24, 2025, <https://socket.io/docs/v2/rooms/>
7. How live collaborative editing works in CodeSignal's IDE, accessed on September 24, 2025, <https://codesignal.com/blog/engineering/how-collaborative-editing-works-in-codesignals-ide/>
8. Building real-time collaboration applications: OT vs CRDT - TinyMCE, accessed on September 24, 2025, <https://www.tiny.cloud/blog/real-time-collaboration-ot-vs-crdt/>
9. Real-Time Collaboration Issues and Solutions - Codeanywhere, accessed on September 24, 2025, <https://codeanywhere.com/blog/real-time-collaboration-issues-and-solutions>

10. Real-time collaborative editing using CRDTs - DiVA portal, accessed on September 24, 2025, <http://www.diva-portal.org/smash/get/diva2:1304659/FULLTEXT01.pdf>
11. [1810.02137] Real Differences between OT and CRDT for Co-Editors - arXiv, accessed on September 24, 2025, <https://arxiv.org/abs/1810.02137>
12. Real Differences Between OT and CRDT for Co-Editors - Hacker News, accessed on September 24, 2025, <https://news.ycombinator.com/item?id=18191867>
13. Scaling Socket.IO: Real-world challenges and proven strategies - Aibly, accessed on September 24, 2025, <https://aibly.com/topic/scaling-socketio>
14. Horizontal Scaling with Socket.IO - DEV Community, accessed on September 24, 2025, <https://dev.to/kawanedres/horizontal-scaling-with-socketio-1lca>
15. Should I use webRTC or socket.io for a text chat? - Stack Overflow, accessed on September 24, 2025, <https://stackoverflow.com/questions/55349347/should-i-use-webrtc-or-socket-io-for-a-text-chat>
16. The pitfalls of using WebRTC - CORDONIQ, accessed on September 24, 2025, <https://www-stg-011.cordoniq.com/news/the-pitfalls-of-using-webrtc/>