



Original Article

Fault-Tolerant Architecture for Real-Time Payment Processing in Large-Scale Telecom Billing Systems

Mr. Parth Patel

Independent Researcher, Pennsylvania, USA.

Abstract - Payment processing systems that serve large consumer-facing organizations must handle millions of transactions per day while maintaining consistent reliability. A single point of failure in the payment pipeline can result in measurable revenue loss and customer dissatisfaction. This paper presents a fault-tolerant architecture for integrating CyberSource payment services within a large-scale Internet Service Provider (ISP) billing platform that handles payments for over three million active customer accounts. The system achieves high availability through a combination of idempotent API design, exponential backoff retry logic, a circuit breaker mechanism calibrated for payment gateway behavior, asynchronous event queuing using Apache Kafka, and distributed transaction tracing. The architecture reduces mean recovery time from over eighteen minutes to under one minute, eliminates duplicate charge incidents through idempotency enforcement, and sustains 99.94% transaction success rates during observed periods of CyberSource API instability. This paper describes the design decisions behind each component, the CyberSource REST API integration approach, lessons learned from operating this system in production and observed performance data before and after the fault-tolerant design was adopted.

Keywords - Payment Gateway Integration, Fault Tolerance, Cybersource, High Availability, Idempotency, Circuit Breaker Pattern, Telecom Billing Systems, Microservices Architecture, PCI DSS Compliance, Real-Time Transaction Processing, Exponential Backoff, Apache Kafka, Distributed Systems Resilience.

1. Introduction

Large Internet Service Providers manage recurring billing cycles for millions of residential and business customers across multiple geographic regions. For a large-scale national ISP operating across multiple states, the payment processing pipeline is not merely a technical subsystem but a direct revenue function. Any disruption in this pipeline, whether caused by a network timeout, a third-party gateway outage, or a downstream database failure, translates into failed auto-pay transactions, increased support call volume, manual intervention overhead, and in some cases, service interruptions for customers whose accounts fall into arrears.

The system described in this paper processes payments for over three million active customer accounts. These customers make one-time payments through web portals and mobile applications, schedule future payments, enroll in recurring monthly autopay, and occasionally request refunds or dispute charges. Each of these payment flows carries different latency requirements and different tolerance for failure. A customer submitting a one-time payment in real time expects an immediate success or failure response within a few seconds. A monthly auto-pay job running in a nightly batch window can tolerate brief retry delays without the customer ever becoming aware of a transient issue.

Prior to the architecture described here, the payment integration was built as a direct synchronous call from the billing application to the CyberSource REST API. This approach worked at lower transaction volumes but became fragile as the customer base grew. Network timeouts from CyberSource that lasted more than thirty seconds would crash worker threads. Retry logic was inconsistent across different payment flows. There was no global state to track whether a payment attempt had been sent to CyberSource before a timeout occurred, which on occasion led to duplicate charges when the billing system retried a request that CyberSource had already processed.

This paper describes the changes we made to address these problems. The solution centers on a dedicated payment microservice that acts as the single intermediary between the billing application and CyberSource. This service absorbs transient failures through structured retry logic, protects the upstream billing application from CyberSource response time degradation using a circuit breaker, and guarantees that every payment attempt is tracked with an idempotency key that prevents duplicate charges even when network conditions cause a request to be submitted more than once.

The contributions of this paper are: (1) a reference architecture for fault-tolerant payment gateway integration at multi-million customer scale; (2) an idempotency mechanism for safe payment retries that eliminates duplicate charging without distributed locking; (3) a circuit breaker tuned for third-party payment gateway behavior; (4) production performance data

comparing the baseline system to the fault-tolerant architecture; and (5) lessons learned from operating this system across multiple billing cycles.

2. Background and Related Work

2.1. Cybersource Payment Gateway

CyberSource, now operated as a Visa subsidiary, provides a cloud-hosted payment processing platform accessible via REST APIs over HTTPS. The platform supports credit and debit card transactions, ACH bank transfers, stored credential recurring billing, refund and void operations, and fraud decision management. For enterprise customers, CyberSource provides a Token Management Service (TMS) that replaces raw cardholder data with opaque tokens. These tokens can be stored in the billing system and reused for future charges without the billing application ever handling sensitive card information directly, which reduces the PCI DSS compliance scope.

The CyberSource REST API uses standard HTTP status codes for synchronous outcomes but delivers some results asynchronously via webhook callbacks. For card-not-present transactions, which represent the majority of ISP billing activity, authorization responses are typically synchronous with response times that average between 800 and 1,200 milliseconds under normal load. However, during peak processing windows or infrastructure events at Visa, response times can increase sharply, and the API may return 5xx errors or simply time out.

2.2. Fault Tolerance in Distributed Systems

The foundation for building reliable systems from unreliable components is well established. Gray and Reuter [1] defined the ACID properties for transaction processing and described recovery mechanisms including write-ahead logging and two-phase commit. Their work established that reliable transaction processing requires careful coordination of failure detection, state persistence, and recovery procedures. Kleppmann [2] extended these ideas to modern distributed architectures, noting that partial failures, network partitions, and clock skew create consistency challenges that are different from the single-machine failure modes that earlier transaction processing literature addressed.

Nygard [3] introduced the circuit breaker as a practical pattern for preventing cascading failures across service dependencies. In its basic form, a circuit breaker monitors the error rate of calls to a downstream dependency. When the error rate crosses a configured threshold, the breaker enters an open state and begins rejecting calls immediately without contacting the failing service. After a timeout period, the breaker allows a limited number of trials call through in a half-open state. If the trials succeed, the breaker closes and resumes normal operation. This pattern is particularly useful for payment gateway integrations because it stops the payment service from exhausting its thread pool on requests that have no realistic chance of succeeding.

Richardson [4] describes the saga pattern for managing long-running distributed transactions. Rather than relying on two-phase commit across service boundaries, a saga decomposes a business operation into a sequence of local transactions. Each local transaction publishes an event or message that triggers the next step. If any step fails, the saga executes compensating transactions to undo the effects of previously completed steps. In the payment domain, a saga might cover the sequence of reserve payment, execute charge, update account balance, and send confirmation.

Fowler [5] introduced the Idempotent Receiver pattern as a mechanism for making message-based systems tolerant of duplicate delivery. The core idea is that a receiver assigns a unique identifier to each logical operation. Before processing, the receiver checks whether an operation with that identifier has already been handled. If so, the receiver returns the cached result without reprocessing. This approach is needed in payment systems where network failures can cause clients to retry requests that the server already processed.

2.3. PCI DSS Compliance Considerations

The Payment Card Industry Data Security Standard [6] defines security requirements for any system that stores, processes, or transmits cardholder data. For an ISP billing platform, achieving PCI DSS compliance requires that raw card numbers never pass through application servers that handle non-payment logic, that all communication with the payment gateway uses TLS 1.2 or higher, and that access to payment-related systems is logged, audited, and reviewed on a regular basis. The microservice architecture described here was designed to keep all payment-related processing within a dedicated service boundary, reducing the number of systems subject to the full requirements of the standard.

3. System Architecture

The payment processing pipeline is organized into four logical layers: the client layer, the application layer, the payment service layer, and the data and messaging layer. Fig. 1 shows the component relationships and data flows across these layers. The diagram uses a strict top-down flow to reflect the actual request path: client requests enter through the API Gateway, are processed by the Billing Service, handed off to the Payment Microservice, and settled in the data layer. The CyberSource Cloud is shown as an external dependency accessed by the CyberSource Client component.

Fig. 1. Fault-Tolerant Payment Processing System Architecture




Fig 1: Fault –Tolerant Payment Processing System Architecture

3.1. Client Layer

The client layer includes four interaction surfaces: the customer-facing web portal, native iOS and Android mobile applications, an auto-pay batch scheduler, and an internal admin console used by operations and support staff. All client interactions flow through the API Gateway, which handles TLS termination, request rate limiting, JWT validation, and routing to downstream services. The API Gateway does not contain any payment-specific logic; it forwards payment requests as opaque payloads to the Billing Service.

3.2. Application Layer

The Billing Service manages the core business logic around invoices, subscription management, and payment scheduling. It maintains the customer account model and determines when and how much to charge. When a payment event is needed, the Billing Service constructs a payment request object that includes the customer token, charge amount, currency, and a billing-cycle reference identifier. It then submits this request to the Payment Microservice via an internal REST call. The Billing Service does not call CyberSource directly. This separation keeps all payment gateway interaction, including retry logic, error handling, and idempotency enforcement, inside a single service.

3.3. Payment Service Layer

The Payment Microservice is the central component of the fault-tolerant architecture. It receives payment requests from the Billing Service, generates and stores an idempotency key, calls the CyberSource Integration Client with backoff-aware retry logic, monitors CyberSource error rates through an embedded circuit breaker, and publishes payment outcome events to Kafka for downstream consumers. The Payment Microservice also interfaces with the Fraud Service, which uses CyberSource Decision Manager to evaluate transaction risk scores before authorization is attempted.

The CyberSource Integration Client is a thin wrapper around the CyberSource REST API. It handles authentication header construction using the HTTP Signature scheme required by CyberSource, maps internal payment request objects to the API request format, deserializes API responses, and normalizes error conditions across the different CyberSource error response formats. The client does not implement retry logic; that responsibility belongs to the Payment Microservice, which applies circuit-breaker state awareness to retry decisions.

3.4. Data and Messaging Layer

The data layer consists of a PostgreSQL database for durable payment transaction records, a Redis cache for idempotency key storage with TTL-based expiration, an Apache Kafka cluster for asynchronous event propagation, and a separate audit database that captures all payment operations with enough detail to satisfy PCI DSS logging requirements. The PostgreSQL database uses synchronous replication to a standby node. Failover is handled automatically by a cluster manager, with observed recovery times of approximately twenty-eight seconds under production conditions.

4. Fault-Tolerance Mechanisms

4.1. Idempotency Design

Every payment request that enters the Payment Microservice is assigned an idempotency key before any external call is made. The key is constructed by concatenating the customer account identifier, the billing cycle reference, the charge amount in minor currency units, and the current day boundary truncated to midnight UTC. This construction ensures that a retry of the same billing event on the same day generates the same key, but that a legitimately separate payment on the following day generates a different key.

Before calling CyberSource, the service performs an atomic check-and-set operation against Redis: if the key already exists, the cached response is returned immediately without contacting CyberSource; if the key does not exist, it is written with a 25-hour TTL and the service proceeds to call CyberSource. The 25-hour TTL is slightly longer than the billing retry window to account for time zone boundary conditions. When the CyberSource response arrives, the outcome is written back to Redis alongside the initial lock entry, so subsequent retries within the TTL window receive the same response regardless of how many times they query the cache.

4.2. Retry with Exponential Backoff

The Payment Microservice implements a five-attempt retry schedule with full jitter exponential backoff. The base delay is one second, and each subsequent delay doubles: 1s, 2s, 4s, 8s, 16s. Full jitter adds a random value between zero and the computed delay at each step to prevent synchronized retry storms when multiple payment workers hit a transient CyberSource error at the same time. Retries are attempted only for specific CyberSource error conditions: HTTP 408 (request timeout), HTTP 502 (bad gateway), HTTP 503 (service unavailable), and HTTP 504 (gateway timeout). HTTP 400 and HTTP 422 responses indicate client-side problems such as invalid card data and are not retried. HTTP 401 and HTTP 403 responses trigger an alert to the operations team rather than a retry.

4.3. Circuit Breaker Pattern

The circuit breaker monitors the error rate of CyberSource calls over a rolling sixty-second window. If more than five of the last ten calls result in a retryable error, the breaker transitions to the open state. In the open state, all payment requests are rejected at the Payment Microservice boundary without contacting CyberSource. Rejected requests publish a payment-deferred event to Kafka, allowing downstream consumers to mark affected accounts as pending-retry rather than failed. After thirty seconds in the open state, the breaker transitions to half-open and permits one trial call. If the trial call succeeds, the breaker closes and resumes normal operation. If it fails, the breaker reopens for another thirty-second interval. Fig. 2 shows the state machine and the retry backoff schedule.




Fig 2: Circuit Breaker State Machine and Retry Strategy

4.4. Dead Letter Queue and Compensation

Payment requests that exhaust all retry attempts without a successful response are published to a Kafka dead letter queue (DLQ) topic. A separate DLQ consumer service reads from this topic and applies additional processing logic: for auto-pay transactions that failed within the nightly billing window, the consumer reschedules the payment attempt for the following morning and notifies the customer via email. For one-time payments submitted by a customer in real time, the consumer marks the transaction as failed and sends an immediate failure notification so the customer can retry manually.

For scenarios where a CyberSource charge succeeds but a subsequent step fails, such as an account balance update that encounters a database error after the charge is confirmed, the system uses a compensation mechanism. The Payment Microservice records the CyberSource authorization reference in the payment database before updating account state. If the account update fails after three attempts, an automated compensation job detects the inconsistency during a nightly reconciliation scan and issues a CyberSource void or refund depending on the settlement status of the original authorization. Reconciliation jobs run at 02:00 and 14:00 UTC daily.

4.5. Cybersource Integration Design

4.5.1. Authentication and Token Management

CyberSource REST API authentication uses the HTTP Signature scheme, where each request is signed with a shared secret using HMAC-SHA256. The signature covers the request path, the HTTP method, the Content-Type header, the Digest header (SHA-256 hash of the request body), and a Date header. The CyberSource Integration Client generates this signature for every outbound call using a signing key stored in a secrets management vault. The key is rotated every 90 days per PCI DSS requirements.

Card credentials are never stored in the billing system database. When a customer adds a payment method, the web portal uses the CyberSource Flex Microform, a JavaScript widget that tokenizes the card number directly in the customer's browser before the form is submitted. The tokenized value is transmitted to the Billing Service, which passes it to CyberSource TMS to generate a permanent network token. The TMS token is stored in the billing database and used for all subsequent charges. If a card is replaced due to expiry or loss, the card network can update the token through account updater services without requiring the customer to re-enter credentials.

4.5.2. Payment Authorization Flow

A typical one-time payment follows this sequence: the Billing Service constructs a charge request with the customer TMS token and amount, assigns an idempotency key, and calls the Payment Microservice. The Payment Microservice checks Redis for an existing idempotency entry, finds none, and writes the pending status. It then calls the Fraud Service, which submits the transaction to CyberSource Decision Manager. For approved transactions, the Payment Microservice calls CyberSource to authorize the charge. CyberSource returns an authorization response with a processorTransactionId. The Payment Microservice writes the authorization record to PostgreSQL, publishes a payment-authorized event to Kafka, and returns the success response to the Billing Service. CyberSource handles settlement automatically based on the merchant account configuration.

4.5.3. Webhook and Asynchronous Notification Handling

CyberSource delivers certain event notifications asynchronously via webhooks, including fraud case updates, chargeback notifications, and ACH return codes. The Payment Microservice exposes a webhook endpoint that CyberSource calls when these events occur. The system validates the CyberSource webhook signature using the shared secret, writes the raw event payload to the audit database immediately upon receipt, and then publishes the event to Kafka for processing by downstream consumers. This two-step write ensures that webhook data is not lost even if Kafka is temporarily unavailable. The webhook endpoint returns HTTP 200 to CyberSource as soon as the payload persisted, regardless of whether downstream processing has completed.

4.5.4. Performance and Reliability Analysis

The tables below present performance and reliability data collected from production. Baseline measurements were taken during a four-week window before the fault-tolerant architecture was deployed. Post-deployment measurements cover the subsequent eight-week period, spanning two full monthly billing cycles.

Table 1: Payment Processing Performance Metrics

Metric	Baseline System	With FT Architecture	Improvement
Avg. Response Time	1,240 ms	382 ms	69.2% reduction
P95 Response Time	3,810 ms	487 ms	87.2% reduction
P99 Response Time	8,200 ms	720 ms	91.2% reduction
Transaction Success Rate	97.2%	99.94%	+2.74 percentage points
Duplicate Charge Incidents	23 per month (avg)	0 (idempotency enforced)	100% elimination
Mean Recovery Time (MRT)	18.4 minutes	47 seconds	95.7% faster
Peak Throughput (TPS)	340 TPS	1,200 TPS	3.5x increase

The largest gain in Table I is the drop in duplicate charge incidents from 23 per month to zero. Duplicate charges were a persistent operational problem that required manual refund issuance and customer communication. The idempotency mechanism eliminated this entirely. The drop in P99 response time from 8.2 seconds to 720 milliseconds reflects the effect of the circuit breaker, which stops the thread pool from being blocked by requests to an unresponsive CyberSource endpoint.

Table 2: Failure Scenario Behavior Comparison

Failure Scenario	Behavior (Pre-FT)	Behavior (Post-FT)	Customer Impact
CyberSource API timeout	Transaction failed, no retry	Retry with backoff up to 5 attempts; circuit breaker engaged at threshold	Zero visible failure for transient timeouts under 16 seconds
CyberSource full outage	All payment requests failed; manual processing required	Circuit breaker opens; payments queued in Kafka for deferred processing	Auto-pay retried within SLA window; one-time payments shown queued status
Primary DB failure	Service crashed; restart required	Automatic failover to read replica; writes queued until primary restored	Sub-30 second disruption during planned failover windows
Duplicate payment request	Double charge if both requests reached gateway	Idempotency key checked in Redis; second request returns cached response	No duplicate charges; idempotency enforced at payment service boundary
Webhook delivery failure	Payment status stuck in pending; manual reconciliation needed	Webhook retried every 30 seconds for 24 hours; audit log records gap	Eventual consistency restored within reconciliation window

Table II shows how the fault-tolerant architecture changes system behavior under specific failure conditions. The most visible change from a customer perspective is the handling of CyberSource full outages. Previously, a CyberSource outage caused all payment requests to fail with visible error responses. With the circuit breaker and Kafka-backed deferral, payments submitted during an outage are queued and processed when service is restored, without requiring customer action.

Table 3: Architectural Approach Comparison

Design Attribute	Monolithic Approach	Microservice Approach	Rationale
Fault isolation	Full system affected	Payment service isolated	Contained blast radius on payment failures
Scalability	Scale entire application	Scale payment service independently	Cost-efficient during peak billing cycles
Deployment frequency	Full release cycle required	Independent deploy per service	Faster iteration on payment logic
PCI DSS scope	Entire application in scope	Only payment service in scope	Reduces compliance burden and audit surface

Table III compares the microservice approach to the prior monolithic integration. A key operational advantage is the reduction in PCI DSS audit scope. In the monolithic system, the entire application was considered in-scope because payment data passed through shared memory and log files. With the dedicated payment microservice, only the payment service and its data stores are subject to PCI DSS controls, reducing the number of systems requiring quarterly vulnerability scans and annual assessment reviews.

5. Challenges and Lessons Learned

5.1. Idempotency Key Boundary Conditions

The initial idempotency key design used a daily boundary truncated to midnight UTC. This introduced a problem for customers in the Pacific Time zone whose billing events were scheduled at 11:00 PM local time. When a billing event was retried after midnight UTC, it generated a new idempotency key and bypassed the deduplication check. The fix was to extend the key TTL to 36 hours and tie the day boundary to the billing event timestamp rather than wall clock time at the moment of retry. This change required a data migration of existing Redis keys but was completed without service interruption using a dual-write period during which both the old and new key formats were checked.

5.2. Circuit Breaker Calibration

The initial circuit breaker threshold was set to open when three of the last five calls failed. In practice, this was too aggressive. CyberSource occasionally returns a 503 response for a single request in an otherwise healthy session, particularly during fraud screening spikes. Opening the circuit breaker on isolated errors caused unnecessary payment deferrals that required manual

follow-up. We recalibrated to the current threshold of five failures in a ten-call window after observing two weeks of production error patterns. The half-open trial period was also lengthened from ten seconds to thirty seconds to give CyberSource enough time to fully recover after infrastructure events.

5.3. Kafka Consumer Lag during Billing Cycles

Monthly billing cycles generate a burst of payment events that is roughly twelve times higher than the average daily rate. During the first billing cycle after Kafka was introduced, consumer lag accumulated faster than the DLQ consumer could drain it. The resolution involved adding consumer group partitions to scale processing horizontally and implementing a priority lane in the Kafka topic configuration, where time-sensitive one-time payments are published to a higher-priority partition that dedicated consumer threads drain before processing deferred auto-pay events. This kept customer-initiated payments from being delayed by the batch auto-pay queue.

5.4. PCI DSS Audit Log Volume

PCI DSS requires that all access to cardholder data systems be logged with enough detail to reconstruct the sequence of events in the event of a security incident. At peak processing rates, the audit logger generates approximately 180,000 log entries per hour. Initial log storage was provisioned for 30-day retention per policy, but the actual volume was forty percent higher than estimates because the audit schema was capturing full request and response payloads including metadata fields. A schema revision to exclude non-required metadata fields reduced storage volume by thirty-one percent without affecting the audit trail quality required for compliance.

6. Conclusion

This paper described a fault-tolerant architecture for integrating CyberSource payment services within a large-scale ISP billing platform serving over three million customer accounts. The architecture addresses the practical problems of direct synchronous gateway integration by introducing a dedicated payment microservice with structured retry logic, circuit breaker protection, idempotency enforcement, and Kafka-backed asynchronous event processing.

The production results show clear gains across all measured metrics. Mean recovery time dropped from over eighteen minutes to under one minute. Duplicate charge incidents were eliminated entirely. Transaction success rates improved from 97.2% to 99.94%, which represents tens of thousands of previously failed payments per month that now complete successfully. The circuit breaker and DLQ mechanisms together allow the system to absorb CyberSource outages of up to several hours without visible customer impact for auto-pay billing scenarios.

The microservice split also brought operational benefits beyond reliability. The reduction in PCI DSS audit scope lowered compliance costs, and the ability to deploy payment logic independently reduced the lead time for gateway integration changes. The lessons learned around idempotency boundary conditions, circuit breaker calibration, and Kafka consumer scaling is likely to come up for other teams doing similar payment gateway integration work at comparable scale.

Future work includes evaluating multi-gateway failover, where the circuit breaker would redirect traffic to a secondary payment processor during extended CyberSource outages, and extending the reconciliation system to produce real-time dashboards rather than the current batch-mode daily reports.

References

- [1] J. Gray and A. Reuter, *Transaction Processing: Concepts and Techniques*. San Francisco, CA: Morgan Kaufmann Publishers, 1992.
- [2] M. Kleppmann, *Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems*. Sebastopol, CA: O'Reilly Media, 2017.
- [3] M. Nygard, *Release It! Design and Deploy Production-Ready Software*, 2nd ed. Raleigh, NC: Pragmatic Bookshelf, 2018.
- [4] C. Richardson, *Microservices Patterns: With Examples in Java*. Shelter Island, NY: Manning Publications, 2018.
- [5] M. Fowler, *Patterns of Enterprise Application Architecture*. Boston, MA: Addison-Wesley, 2002.
- [6] PCI Security Standards Council, *Payment Card Industry Data Security Standard (PCI DSS), Requirements and Security Assessment Procedures, Version 3.2.1, May 2018*. [Online]. Available: https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2-1.pdf
- [7] CyberSource Corporation, *CyberSource REST API Developer Guide*, 2022. [Online]. Available: <https://developer.cybersource.com/api/reference>
- [8] Visa Inc., *CyberSource Token Management Service (TMS) Developer Guide*, 2022. [Online]. Available: <https://developer.cybersource.com/api/developer-guides/dita-tms/>
- [9] Amazon Web Services, *Building Fault-Tolerant Applications on AWS*, AWS Whitepaper, 2021. [Online]. Available: <https://docs.aws.amazon.com/whitepapers/latest/fault-tolerant-components/fault-tolerant-components.html>

- [10] T. Akidau, R. Bradshaw, C. Chambers, S. Chernyak, R. Fernandez-Moctezuma, R. Lax, S. McVeety, D. Mills, F. Perry, E. Schmidt, and S. Whittle, "The dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing," *Proc. VLDB Endow.*, vol. 8, no. 12, pp. 1792-1803, Aug. 2015.
- [11] H. Garcia-Molina and K. Salem, "Sagas," in *Proc. ACM SIGMOD International Conference on Management of Data*, San Francisco, CA, May 1987, pp. 249-259.
- [12] R. Fielding and J. Reschke, *Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content*, IETF RFC 7231, Jun. 2014. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7231>
- [13] National Institute of Standards and Technology, *FIPS PUB 180-4: Secure Hash Standard (SHS)*, Aug. 2015. [Online]. Available: <https://doi.org/10.6028/NIST.FIPS.180-4>.