# Technical Description of System Outage

## Incident Summary:

At 10:45 AM UTC, a critical database server (DB-SRV-02) hosting user account data for the company's e-commerce platform encountered an unexpected failure due to a high I/O load, causing the disk to reach 100% utilization. This resulted in a failure of the replication process between the primary and secondary databases, triggering a failover mechanism that did not execute properly due to misconfigured cluster settings in the PostgreSQL database cluster.

## Root Cause:

The root cause was identified as an improperly tuned I/O subsystem, which allowed high transaction volumes to overload the server's storage array. As a result, database queries began to queue up, overwhelming the connection pool and leading to connection timeouts for both the application servers and APIs interacting with the database.

## Impact:

The database outage impacted several key services:

- Users were unable to log in to their accounts or retrieve account details.
- New user registrations were unavailable.
- Payment gateway services were delayed, causing failed transactions for users attempting to make purchases.
- APIs for order processing were timing out, which led to delays in order confirmation and email notifications.
- Approximately 80% of active users experienced slow response times, with an average delay of 12 seconds per request before complete service failure.

## Resolution:

At 11:20 AM UTC, the IT operations team restarted the affected database server and manually triggered a failover to the secondary database. The connection pool and transaction queue were flushed to restore the API and application server connections. Additionally, the I/O settings were adjusted to better handle peak traffic loads, and a full replication resync was performed between the primary and secondary databases. Full services were restored by 11:40 AM UTC.

## Next Steps:

- Conduct a full audit of the PostgreSQL cluster configuration.
- Fine-tune the I/O subsystem to prevent recurrence.
- Investigate and test failover procedures to ensure proper execution in future incidents.
- Monitor system performance closely over the next 48 hours to identify any residual issues.