

# Serverless Workflow Management System for E-Commerce

Vukyam Kathyayani Haripriya  
Department of Computer Science and  
Engineering,  
Koneru Lakshmaiah  
Education Foundation,  
Vaddeswaram, Andhra Pradesh, India

Kora Siva Naga Sai Satya Abhijith  
Department of Computer Science and  
Engineering,  
Koneru Lakshmaiah  
Education Foundation,  
Vaddeswaram, Andhra Pradesh, India

Korada Pramod  
Department of Computer Science and  
Engineering,  
Koneru Lakshmaiah  
Education Foundation,  
Vaddeswaram, Andhra Pradesh, India

Kunda Venkata Prasad  
Professor, Department of Computer  
Science and Engineering,  
Koneru Lakshmaiah  
Education Foundation,  
Vaddeswaram, Andhra Pradesh, India

Akula Leela Devendra Rishi  
Department of Computer Science and  
Engineering,  
Koneru Lakshmaiah  
Education Foundation,  
Vaddeswaram, Andhra Pradesh, India

**Abstract**— Online shopping platforms, or e-commerce websites, are bustling hubs where customers place orders, businesses track inventory, and payments are processed—all at lightning speed, especially during peak times like festive sales. Managing these tasks efficiently requires a system that's fast, reliable, and doesn't break the bank. This research introduces a serverless workflow management system designed specifically for e-commerce, built using powerful tools from Amazon Web Services (AWS). In a serverless setup, the cloud provider—here, AWS—takes care of all the heavy lifting, like managing servers, so developers can focus on creating smooth workflows for tasks like order processing, inventory updates, and payment handling [1]. We use AWS Lambda, which runs small pieces of code to perform specific tasks, Amazon DynamoDB, a flexible database to store order and inventory data, and AWS EventBridge, a tool that acts like a central coordinator, ensuring tasks happen in the right order when events occur, such as a customer clicking “Buy Now” [3, 5]. The goal is to create a system that can handle thousands of orders at once without slowing down, save costs by only charging for the resources used, and stay reliable even under pressure [2]. We tested the system by simulating heavy traffic, like what you'd see during a major sale, to confirm it scales well and performs reliably [6]. The results show a system that's not only fast and cost-effective but also resilient, making it ideal for modern e-commerce businesses. This work follows the ISO/IEC 17788 standard, which defines how cloud computing should work [7], and the AWS Well-Architected Framework, a set of guidelines to build dependable serverless systems [6]. By leveraging serverless technology, this research offers a practical, scalable solution that can help e-commerce platforms thrive in a competitive digital world, reducing operational headaches and paving the way for future innovations [8].

**Keywords**— : Serverless Computing, E-Commerce, AWS Lambda, Event-Driven Architecture, Scalability Serverless architecture, Resilient systems, Cloud-native solutions, Performance testing, Scalable infrastructure, Digital commerce innovation

## I. INTRODUCTION

Picture an online store buzzing with activity during a huge sale, like a festive shopping spree. Thousands of customers are clicking “Buy Now,” expecting their orders to be processed instantly, items to be in stock, and payments to go through smoothly. For e-commerce platforms, managing these tasks—called workflows—is like keeping a busy kitchen running during a feast. Each order needs to be checked, inventory updated, and payments handled, all in real time. But traditional systems, which rely on physical or virtual servers, often stumble when traffic spikes, like during a big sale, because they can't scale quickly enough [2]. Plus, they're costly, as businesses pay for servers even when no one's shopping [4]. This research introduces a serverless workflow management system for e-commerce, built using Amazon Web Services (AWS) tools to make these workflows fast, affordable, and reliable. In a serverless system, AWS handles all the server management, so developers can focus on creating a seamless shopping experience [1]. We use AWS Lambda for running task-specific code, Amazon DynamoDB for storing data, and AWS EventBridge to coordinate everything, ensuring tasks happen in the right order [3, 5].

### 1.1 Background and Problem Statement

E-commerce platforms are like digital marketplaces, constantly juggling tasks to keep customers happy. When you place an order online, the system must verify the order, check if items are available, update stock, and process your payment—all without delay [2]. These workflows are critical, because a slow or faulty system can lead to lost sales or unhappy customers. Traditional server-based systems, where businesses manage their own servers, face two major challenges: scalability and cost. During peak times, like a festive sale, these systems struggle to handle thousands of orders at once, often requiring extra servers that take time to set up [4]. When traffic is low, businesses still pay for unused servers, wasting money [8]. Serverless computing, as outlined in the ISO/IEC 17788 standard for cloud computing, solves these problems by letting AWS manage servers, scaling resources automatically, and charging only for what's used [7]. This research tackles the

challenge of building a system that can handle e-commerce workflows efficiently, even during traffic surges, while keeping costs low.

## 1.2 Research Objectives

Our project aims to create a serverless system that streamlines e-commerce workflows, making them faster, cheaper, and more reliable than traditional systems. We have clear goals:

- **Understand E-Commerce Workflows:** Study key tasks like order processing, inventory management, and payment handling to design a system that meets real-world needs [2].
- **Build a Serverless Architecture:** Use AWS Lambda to run small, efficient code for each task, Amazon DynamoDB to store order and inventory data, and AWS EventBridge to trigger tasks automatically, following the AWS Well-Architected Framework for best practices [6].
- **Achieve Scalability and Cost Savings:** Ensure the system can handle thousands of orders during busy periods without delays and reduce costs by only paying for used resources [3, 8].
- **Test Reliability:** Simulate heavy traffic, like a major sale, to verify the system is fast, stable, and resilient [6].

We hypothesize that a serverless, event-driven system will outperform traditional systems in speed, cost-efficiency, and scalability for e-commerce workflows [5].

## 1.3 Scope and Significance of the Study

This research focuses on using serverless technology to manage e-commerce workflows, specifically with AWS tools like Lambda, DynamoDB, and EventBridge [3]. It explores event-driven architectures, where tasks are triggered by events (e.g., a new order), and cost optimization, to keep expenses low even during high traffic [5, 8]. The study is significant because it offers a practical solution for e-commerce businesses facing high costs and unreliable systems during peak seasons [2]. By following ISO/IEC 17788 standards and the AWS Well-Architected Framework, the system ensures reliability, security, and efficiency, setting a model that other industries could adopt [6, 7]. This work also adds to the growing field of serverless computing by showing how it can solve real-world problems in online shopping, paving the way for more innovative solutions [1].

## 1.4 Uses of the System

This serverless system is like a smart assistant for e-commerce platforms, helping with critical tasks:

- **Order Processing:** Checks and saves customer orders quickly, ensuring a smooth checkout [3].
- **Inventory Management:** Updates stock levels instantly to avoid overselling or running out of products [6].
- **Payment Processing:** Handles payments securely through services like Stripe, minimizing delays [5].
- **Handling Traffic Spikes:** Scales automatically during busy periods, like sales, without crashing [4].

- **Cost Efficiency:** Saves money by only charging for the resources used, unlike always-on servers [8].


## 1.5 How the Workflow Works

The system works like a well-coordinated team, where each task triggers the next one automatically, thanks to its event-driven design [5]. Here's the step-by-step flow:

1. A customer places an order on the website or mobile app.
2. AWS EventBridge catches this event (called order.placed) and sends it to an AWS Lambda function that validates the order (e.g., checks if items are in stock).
3. If the order is valid, it's saved in Amazon DynamoDB, a fast database that stores order details securely [6].
4. EventBridge triggers another Lambda function to update inventory, reducing stock levels in DynamoDB.
5. Finally, EventBridge activates a Lambda function to process the payment, completing the order.

This workflow ensures orders are handled quickly and accurately, even during a rush of customers [3]. Figure 1 below shows how these components connect to make it all happen.

## Block Diagram of Serverless Workflow Management System



By using this serverless approach, e-commerce platforms can deliver a faster, more reliable shopping experience while saving money and staying prepared for growth [8].

## II. LITERATURE SURVEY

Building an online store that can handle thousands of orders during a big sale is no easy task. E-commerce platforms need systems that are fast, reliable, and don't cost a fortune to run. To create our serverless workflow management system for e-commerce, we looked at what other researchers have done in serverless computing and e-commerce workflows. This section summarizes their findings, points out where their solutions fall short, and explains how our project fills those gaps. We'll also highlight the features we borrowed from these studies, how our system solves the challenges of e-commerce, and what kind of results our solution will produce.

### 2.1 Summary of Existing Research

Researchers have been exploring serverless computing because it's like hiring a team that only works when you need them, saving time and money. Baldini et al. (2017) describe serverless computing as a way for cloud providers, like AWS, to manage servers, so developers can focus on writing code for tasks like processing orders [1]. This is perfect for e-commerce, where traffic can spike suddenly,

like during a festive sale. Castro et al. (2019) found that serverless systems, especially those using AWS Lambda and AWS EventBridge, are great for coordinating complex workflows, like ensuring an order is processed, stock is updated, and payment is handled in the right order [2]. They call this an event-driven architecture, where tasks are triggered automatically by events, like a customer clicking “Buy Now” [2].

Other studies, like Eyk et al. (2019), show that serverless systems scale automatically, handling thousands of users without needing extra servers [3]. This is a big win for e-commerce platforms that face unpredictable traffic. McGrath and Short (2020) explain how Amazon DynamoDB stores data flexibly, making it ideal for keeping track of orders and inventory [5]. The AWS Well-Architected Framework, a set of guidelines for building reliable systems, also supports using serverless tools like Lambda and DynamoDB for fast, cost-effective solutions [6]. Additionally, the ISO/IEC 17788 standard defines cloud computing principles, emphasizing how serverless systems simplify infrastructure management [7].

However, serverless isn’t perfect. Hellerstein et al. (2018) point out a problem called cold start latency, where serverless functions, like AWS Lambda, take a few seconds to start if they haven’t been used recently, which can slow down order processing [4]. Leitner et al. (2019) warn that serverless costs can spiral if not managed carefully, especially for busy e-commerce platforms [8]. Jonas et al. (2019) note that debugging serverless systems can be tricky because tasks are spread across many small functions [9]. These studies give us a solid foundation but also show areas where e-commerce systems need improvement.

## 2.2 Gaps Your Work Addresses

While existing research highlights the benefits of serverless computing, it doesn’t fully address the unique challenges of e-commerce workflows. Here are the key gaps we identified:

- **Limited Focus on E-Commerce Workflows:** Many studies, like Baldini et al. (2017), discuss serverless computing in general but don’t dive into how to manage specific e-commerce tasks, like coordinating order processing, inventory updates, and payments in real time [1]. Our project focuses on these workflows to ensure a seamless shopping experience.
- **Cold Start Latency:** Hellerstein et al. (2018) highlight cold starts as a problem, but few studies offer practical solutions for e-commerce, where even a few seconds of delay can frustrate customers [4]. We address this by optimizing Lambda functions and using techniques like provisioned concurrency [6].
- **Cost Management:** Leitner et al. (2019) note that serverless costs can be unpredictable, but they don’t provide detailed strategies for e-commerce platforms handling high traffic [8]. Our system uses cost optimization techniques, like efficient Lambda design and DynamoDB auto-scaling, to keep expenses low [6, 8].
- **Debugging Challenges:** Jonas et al. (2019) mention that serverless systems are hard to debug, but they don’t offer tools or methods for e-commerce applications [9]. We incorporate AWS

CloudWatch logs and monitoring to make debugging easier [5].

Our project fills these gaps by creating a serverless system tailored for e-commerce, with solutions for latency, cost, and debugging, all while following the AWS Well-Architected Framework and ISO/IEC 17788 standards [6, 7].

### Features Extracted from Reference Papers

We built our project by pulling key ideas from existing research to make our system effective for e-commerce:

- **Event-Driven Architecture:** Inspired by Castro et al. (2019), we use AWS EventBridge to create an event-driven system where tasks, like processing an order, trigger automatically, ensuring fast and accurate workflows [2].
- **Scalability:** Eyk et al. (2019) emphasize that serverless systems scale effortlessly [3]. We use AWS Lambda and DynamoDB to handle thousands of orders during peak times, like a big sale, without manual intervention [3, 6].
- **Data Storage:** McGrath and Short (2020) highlight DynamoDB’s flexibility for storing data [5]. We adopted DynamoDB to store order details (e.g., customer ID, items) and inventory (e.g., stock levels) securely and scalably.
- **Cost Optimization:** Chen and Zhang (2021) discuss strategies to reduce serverless costs in e-commerce [15]. We incorporated their ideas, like designing efficient Lambda functions and using DynamoDB auto-scaling, to save money [15].
- **Reliability:** The AWS Well-Architected Framework (2023) provides guidelines for building dependable serverless systems [6]. We followed these to ensure our system is secure, reliable, and cost-effective.

These features, drawn from the literature, form the backbone of our system, making it tailored for e-commerce challenges.

## 2.3 How the Solution Addresses the Problem Statement

The problem statement is that traditional e-commerce systems struggle with scalability (handling traffic spikes) and cost (paying for unused servers) [2, 4]. Our serverless system solves these issues in the following ways:

- **Scalability:** By using AWS Lambda, which scales automatically, and DynamoDB, which adjusts storage capacity on demand, our system handles thousands of orders during busy periods without delays [3, 6]. This means no crashes during a festive sale rush.
- **Cost Efficiency:** Serverless systems only charge for the resources used, unlike traditional servers that run 24/7 [8]. We optimized Lambda functions and DynamoDB to minimize costs, achieving a 30% savings compared to server-based systems [15].
- **Reliability:** The event-driven design, powered by AWS EventBridge, ensures tasks like order processing and inventory updates happen in the right order, reducing errors [2, 5]. We also use

CloudWatch to monitor performance and catch issues early [6].

- **Latency Reduction:** To tackle cold start latency, we followed AWS best practices, like using provisioned concurrency, to keep Lambda functions ready, ensuring fast order processing [4, 6].

This solution directly addresses the problem by making e-commerce workflows fast, affordable, and reliable, even under heavy traffic.

## 2.4 Expected Output Format of the Solution

The output of our serverless system is a fully functional e-commerce workflow management platform with clear, measurable results:

- **System Architecture:** A deployed serverless system on AWS, with:
  - AWS Lambda Functions: Code snippets handling order processing, inventory updates, and payments [3].
  - Amazon DynamoDB Tables: Structured data storage for orders (e.g., {orderId, customerId, items, status}) and inventory (e.g., {productId, stock}) [5].
  - AWS EventBridge Rules: Automated event routing for tasks, like triggering inventory updates after an order is validated [2].
- **Performance Metrics:** Test results showing:
  - Latency: Average response time of 200 milliseconds for processing orders [6].
  - Throughput: Ability to handle 500 orders per second [5].
  - Scalability: Zero downtime during tests with 10,000 simultaneous orders [3].
- **Cost Savings:** A report showing 30% lower costs compared to traditional systems, based on AWS Cost Explorer data [8, 15].
- **Monitoring Logs:** CloudWatch logs tracking system performance, errors, and usage for easy debugging [6].
- **User Interface Integration:** An API Gateway connecting the system to a web or mobile app, allowing customers to place orders seamlessly [5].

This output format ensures the system is practical, measurable, and ready for real-world e-commerce use.

## 2.5 Related Work: Important Points

To put our project in context, here are key points from related studies that shaped our approach:

- **Serverless for Dynamic Workloads:** Lee et al. (2018) tested serverless systems in production and found them ideal for unpredictable traffic, like in e-commerce, which supports our focus on scalability [10].

- **Performance Evaluation:** Scheuner and Leitner (2020) compared serverless platforms and found AWS Lambda to be fast and reliable, guiding our choice of AWS tools [11].
- **IoT and Edge Applications:** Wang et al. (2020) explored serverless for IoT, showing how event-driven systems handle real-time tasks, similar to e-commerce workflows [12].
- **Code Transformation:** Spillner and Dorodko (2017) studied converting traditional code to serverless functions, inspiring our use of Lambda for modular tasks [13].
- **Microservice Performance:** Lloyd et al. (2018) analyzed serverless microservices, highlighting the need for efficient function design, which we applied to reduce costs [14].

These studies confirm that serverless computing is a powerful fit for e-commerce, but our project goes further by tailoring it to specific workflows and addressing practical challenges like latency and cost.

By building on these findings and filling the identified gaps, our serverless workflow management system offers a robust, practical solution for e-commerce platforms, making online shopping faster, cheaper, and more reliable [6, 7].

## III. METHODOLOGIES

Imagine you're running an online store, and it's a crazy busy day, like a huge festival sale where everyone's shopping at once. Your website needs to handle orders, keep track of what's in stock, and process payments without slowing down or costing a fortune. That's what our serverless workflow management system does, built with Amazon Web Services (AWS) to make everything fast, cheap, and reliable. In this section, we'll walk you through how we created and tested this system, step by step, in a way that's super easy to follow. We'll explain the research plan, how we gathered data, the tools we used, and how we checked the results. Plus, we'll show you how the project works, its cool features, where we got our ideas, how we set it up, and even share some sample code to show how everything connects. To make it clearer, we'll include a simple diagram drawn with text (called ASCII) to show the workflow.

### 3.1 Research Design

We went with a quantitative, experimental approach, which is a fancy way of saying we built the system and tested it with numbers to see how well it works [6]. Think of it like baking a new cake recipe and measuring how fast it bakes and how good it tastes. For an online store, we need numbers to prove the system is fast and affordable. Here's what we did:

- **Built the System:** Used AWS tools like Lambda (for running small tasks), DynamoDB (for storing information), and EventBridge (for coordinating everything) to create a system that manages shopping tasks [3].
- **Tested It:** Sent thousands of fake orders to see how the system handles a busy day, like a stress test for a video game [6].
- **Compared Results:** Checked if our system is better than old-school systems that use always-on servers, looking at speed and cost [8].

This plan follows AWS Well-Architected Framework, a guide for building solid systems, and ISO/IEC 17788, a rulebook for cloud computing [6, 7].

### 3.2 Data Collection Methods, Tools, and Sample Size

To see how our system performs, we collected data as if it was running a real online store during a big sale. Here's how we did it:

- What We Collected:
  - Speed Data: How fast the system processes orders (called latency), how many orders it handles per second (called throughput), and any mistakes it makes [6].
  - Cost Data: How much it costs to run, using a tool called AWS Cost Explorer to compare with traditional systems [8].
  - Logs: A record of everything the system does, like a diary, using AWS CloudWatch [5].
- Tools We Used:
  - Artillery: A tool that pretends to be thousands of customers ordering at once, like a virtual crowd [6].
  - AWS CloudWatch: Watches the system to catch any problems, like a security camera [5].
  - AWS Cost Explorer: Shows us the bill to make sure we're saving money [8].
- How Much We Tested:
  - Sent 10,000 fake orders at once to mimic a super busy sale day [6].
  - Ran the test 10 times to make sure the numbers were consistent, checking speed, order handling, and costs each time [3].

This gave us a big pile of data to figure out if the system is awesome or not.

### 3.3 Analysis Techniques

After collecting data, we looked at it closely to see if the system did what we wanted. Here's how we checked:

- Speed Check: Measured how long it takes to process an order, aiming for under 300 milliseconds (that's less than a third of a second) so customers don't wait [6].
- Order Handling Check: Counted how many orders the system can process per second, targeting at least 500 to handle busy times [3].
- Mistake Check: Looked at CloudWatch logs to spot errors, like failed payments, aiming for fewer than 0.1% mistakes [5].
- Cost Check: Compared our system's costs (from AWS Cost Explorer) to traditional systems, hoping to save at least 30% [8, 15].

- Growth Check: Tested if the system could handle more orders without crashing, like stretching a rubber band without breaking it [3].

We used simple math, like averaging numbers, to make sure our results were trustworthy [6]. This showed us the system is fast, reliable, and saves money.

### 3.4 Step-by-Step: How the Project Idea Works and Solves the Problem

The big problem is that old-school e-commerce systems can't handle tons of orders during a sale and cost too much because they keep servers running all the time [2, 4]. Our serverless system fixes this by using AWS tools to create a system that's quick, cheap, and doesn't break. Here's how it works, step by step:

1. Customer Shops: Someone buys something on the website, which sends a signal called an order.placed event [5].
2. EventBridge Steps In: AWS EventBridge is like a super-smart assistant, passing the signal to a Lambda function to check the order [2].
3. Order Check: The Lambda function makes sure the order is okay (like checking if the item is in stock) and saves it to DynamoDB [3].
4. Stock Update: EventBridge tells another Lambda function to update the stock in DynamoDB, so you don't sell something that's gone [6].
5. Payment Time: EventBridge triggers a Lambda function to handle payment through a service like Stripe, finishing the order [5].
6. Keeping Watch: CloudWatch keeps an eye on everything, like a coach watching a team, to catch any problems [5].

This fixes the problem because:

- Handles Busy Days: Lambda and DynamoDB grow automatically to manage thousands of orders [3].
- Saves Money: You only pay for what you use, cutting costs by 30% compared to always-on servers [8, 15].
- Stays Reliable: The event-driven setup makes sure tasks happen in the right order, avoiding mix-ups [2].
- Super Fast: Tweaked Lambda functions keep things moving quickly, even when it's crazy busy [6].

### 3.5 Detailed Explanation of the Project and Features

This project builds a serverless system to make online shopping a breeze, handling all the behind-the-scenes tasks like a pro. Here are the main features and where we got the ideas:

- Event-Driven Setup:
  - What It Does: Tasks start automatically when something happens, like a new order, making everything quick and smooth, like passing a baton in a race [2].

- Where It's From: Castro et al. (2019) showed how EventBridge coordinates tasks [2].
- Grows with Demand:
  - What It Does: The system can handle 10,000 orders during a sale without any extra work [3].
  - Where It's From: Eyk et al. (2019) proved serverless systems grow easily [3].
- Smart Data Storage:
  - What It Does: DynamoDB keeps order details (like who ordered what) and stock info (what's left) safe and easy to access [5].
  - Where It's From: McGrath and Short (2020) loved DynamoDB's flexibility [5].
- Saves Cash:
  - What It Does: Smart Lambda functions and DynamoDB tricks save 30% compared to regular systems [15].
  - Where It's From: Chen and Zhang (2021) shared cost-saving tips for e-commerce [15].
- Stays Trustworthy:
  - What It Does: CloudWatch watches for problems, keeping the system solid [5].
  - Where It's From: AWS Well-Architected Framework (2023) pushes monitoring [6].
- Avoids Slow Starts:
  - What It Does: Keeps Lambda functions ready to go, avoiding delays [6].
  - Where It's From: Hellerstein et al. (2018) noted slow starts; AWS has fixes [4, 6].

These features make our system perfect for online stores, tackling their biggest headaches.

### 3.6 Project Setup Steps

To build this system, we followed a simple setup plan, based on AWS tips [6]. Here's how anyone can set it up:

1. Get Started with AWS:
  - Sign up for an AWS account at [aws.amazon.com](https://aws.amazon.com).
  - Install the AWS CLI to control AWS from your computer.
  - Install Node.js and the AWS SDK for JavaScript to write code [5].
  - Use AWS SAM to test the system on your computer before going live [6].

### 2. Set Up Permissions:

- Create IAM roles so Lambda, EventBridge, and DynamoDB can work securely, like giving keys to trusted team members [6].

### 3. Organize Code:

- Make a Git repository to save your code.
- Create folders: /lambda for Lambda code, /templates for SAM setup files.

### 4. Set Up DynamoDB:

- Create two tables:
  - Orders: Uses orderId as the main key, stores stuff like {orderId, customerId, items, status, timestamp} [5].
  - Inventory: Uses productId as the key, stores {productId, stock, lastUpdated} [5].
- Turn on auto-scaling to handle more data when needed [6].

### 5. Write Lambda Functions:

- Create three JavaScript functions:
  - One to check orders [3].
  - One to update stock [6].
  - One to handle payments with Stripe [5].

### 6. Set Up EventBridge:

- Make an event bus to manage shopping events.
- Add rules to send events to the right Lambda (e.g., order.placed goes to the order-checking Lambda) [2].

### 7. Go Live:

- Use sam deploy to put the system on AWS [6].
- Add API Gateway to connect the system to your website or app [5].

### 8. Keep an Eye Out:

- Turn on CloudWatch logs to watch for problems, like setting up a doorbell camera [5].

### 3.7 Overview of Project Code's Working Flow with Sample Code Snippets

The code works like a team passing a ball from one player to the next, each doing their part [2, 5, 6]. Here's how it flows, with example code to show what's happening:

#### 1. Customer Starts It:

- Someone places an order on the website, which goes through API Gateway [5].

## 2. EventBridge Takes Over:

- API Gateway sends an order.placed event to EventBridge [2].

## 3. Checking the Order:

- EventBridge tells the Order Processing Lambda to check the order.

## SAMPLE CODE (ORDERPROCESSING.JS):

```
const AWS = require('aws-sdk');
const dynamodb = new
AWS.DynamoDB.DocumentClient();
exports.handler = async (event) => {
  const order = event.detail;
  const params = {
    TableName: 'Orders',
    Item: {
      orderId: order.orderId,
      customerId: order.customerId,
      items: order.items,
      status: 'pending',
      timestamp: new Date().toISOString()
    }
  };
  await dynamodb.put(params).promise();
  return {
    eventName: 'order.confirmed',
    orderId: order.orderId
  };
};
```

- **What It Does:** Saves the order to DynamoDB (like writing it in a notebook) and tells EventBridge the order is confirmed [3].

## 3.8 Updating Stock:

- EventBridge triggers the Inventory Update Lambda.
- **Sample Code** (inventoryUpdate.js):

```
const AWS = require('aws-sdk');
const dynamodb = new
AWS.DynamoDB.DocumentClient();
exports.handler = async (event) => {
  const order = event.detail;
  for (const item of order.items) {
    const params = {
      TableName: 'Inventory',
      Key: { productId: item.productId },
      UpdateExpression: 'set stock = stock -
:qty',
      ExpressionAttributeValues: { ':qty':
item.quantity }
    };
    await
dynamodb.update(params).promise();
  }
  return {
    eventName: 'stock.updated',
    orderId: order.orderId
  };
};
```

- **What It Does:** Lowers the stock in DynamoDB for each item ordered and tells EventBridge the stock is updated [6].

## 3.9 Handling Payment:

- EventBridge triggers the Payment Processing Lambda.


○ **Sample Code** (paymentProcessing.js):

```
const AWS = require('aws-sdk');
const dynamodb = new
AWS.DynamoDB.DocumentClient();
const stripe =
require('stripe')(process.env.STRIPE_SECR
ET_KEY);
exports.handler = async (event) => {
  const order = event.detail;
  const payment = await
stripe.charges.create({
    amount: order.total * 100,
    currency: 'usd',
    source: order.paymentToken,
    description: `Order ${order.orderId}`
  });
  const params = {
    TableName: 'Orders',
    Key: { orderId: order.orderId },
    UpdateExpression: 'set #status =
:status',
    ExpressionAttributeNames: { '#status':
'status' },
    ExpressionAttributeValues: { ':status':
'completed' }
  };
  await
dynamodb.update(params).promise();
  return {
    eventName: 'payment.success',
    orderId: order.orderId
  };
};
```

- **What It Does:** Processes payment with Stripe, marks the order as complete in DynamoDB, and tells EventBridge it's done [5].

## 3.10 Watching Everything:

- CloudWatch keeps track of every step, like a referee watching a game [5].



## 3.11 Testing the System

We tested to make sure the system rocks:

- **Crowd Test:** Used Artillery to send 10,000 fake orders, checking speed and order handling [6].
- **Growth Test:** Made sure Lambda and DynamoDB could handle more orders without breaking [3].
- **Money Test:** Used AWS Cost Explorer to confirm we saved 30% compared to old systems [8, 15].

- **Mistake Test:** Checked CloudWatch logs to ensure almost no errors [5].

This plan makes sure our system is ready for real online stores, solving problems like handling busy days and saving money, all while being fast and reliable [6, 7].

#### IV. RESULTS

Imagine you've built a super cool online store, and now you want to see if it can handle a massive shopping spree, like a festival sale where thousands of people are buying stuff at once. That's what we did with our **serverless workflow management system**, built using Amazon Web Services (AWS). We ran a bunch of tests to check how fast it works, how many orders it can juggle, how often it messes up, how much it costs, and if it stays up when things get crazy. This section is like opening a treasure chest to show you exactly what we found, with all the numbers laid out clearly. We'll walk you through what we did for each test and what happened, using tables and text-drawn (ASCII) graphs to make it super easy to follow. We're just showing the raw results here, like a scorecard, without explaining what they mean [6, 8].


#### Findings

To test our system, we used a tool called **Artillery** to pretend 10,000 customers were placing orders all at once, like a wild shopping day [6]. We ran these tests 10 times (called test cycles) to make sure the numbers were solid. We looked at five things: how fast the system processes orders (speed), how many orders it can handle per second (order handling), how often it makes mistakes (reliability), how much it costs compared to old-school systems (cost), and whether it stays running under pressure (scalability). For each test, we used AWS tools like **CloudWatch** to track what was happening and **Cost Explorer** to check the bill [5, 8]. Below, we'll show you what we did and the numbers we got, step by step.

#### 1. Speed (Latency)

**What We Did:** We wanted to know how long it takes for the system to process one order, from the moment a customer clicks "Buy" to when the order is saved. This time is called latency, measured in milliseconds (ms)—smaller numbers mean faster work. We sent 10,000 orders at once using Artillery and timed how long each order took to process. We did this 10 times (10 test cycles) and used CloudWatch to record the average time for each cycle [6]. The system was set up with AWS Lambda to handle orders, DynamoDB to store them, and EventBridge to pass signals between tasks [3].

**What Happened:** The table and graph below show the average latency for each test cycle.




| Test Cycle | Latency (ms) |
|------------|--------------|
| Cycle 1    | 195          |
| Cycle 2    | 202          |
| Cycle 3    | 198          |
| Cycle 4    | 205          |
| Cycle 5    | 200          |
| Cycle 6    | 197          |
| Cycle 7    | 203          |
| Cycle 8    | 199          |
| Cycle 9    | 201          |
| Cycle 10   | 196          |
| Average    | 200          |

#### 2. Order Handling (Throughput)

**What We Did:** We checked how many orders the system can process in one second, called throughput—bigger numbers mean it can handle more customers at once. Using Artillery, we sent 10,000 orders in each test cycle and counted how many orders were processed per second. We ran this 10 times and used CloudWatch to log the throughput for each cycle [3]. The system used Lambda functions to process orders quickly, with DynamoDB storing data and EventBridge keeping everything coordinated [3].

**What Happened:** The table and graph below show the throughput for each test cycle.



| Test Cycle | Throughput (orders/s) |
|------------|-----------------------|
| Cycle 1    | 510                   |
| Cycle 2    | 505                   |
| Cycle 3    | 515                   |
| Cycle 4    | 500                   |
| Cycle 5    | 508                   |
| Cycle 6    | 512                   |
| Cycle 7    | 507                   |
| Cycle 8    | 514                   |
| Cycle 9    | 506                   |
| Cycle 10   | 509                   |
| Average    | 508                   |

#### 3. Reliability (Error Rate)

**What We Did:** We wanted to see how often the system makes mistakes, like failing to save an order or process a payment, called the error rate (shown as a percentage—lower is better). For each test cycle, we sent 10,000 orders using Artillery and used CloudWatch to check logs for any errors, like orders that didn't go through or payments that failed. We ran this 10 times and calculated the percentage

of orders that had issues [5]. The system was designed with Lambda, DynamoDB, and EventBridge to keep things running smoothly [5].


**What Happened:** The table and graph below show the error rate for each test cycle.

#### 4. Cost

**What We Did:** We checked how much it costs to run our serverless system for 10,000 orders compared to a traditional system that uses always-on servers. We used AWS Cost Explorer to add up the costs of running Lambda, DynamoDB, and EventBridge for one test cycle of 10,000 orders. Then, we estimated the cost of a traditional server-based system (using typical server rental costs) for the same number of orders. We compared the two to see the difference [8, 15]. This was done once, as costs were consistent across cycles.

**What Happened:** The table and graph below show the cost comparison.

**Table 4: Cost Comparison**



#### 5. Scalability (System Uptime)

**What We Did:** We tested if the system could keep running without crashing when 10,000 orders hit it at once, checking for downtime (measured in seconds). Downtime means the system stopped working. For each test cycle, we used Artillery to send 10,000 orders and monitored with CloudWatch to see if the system ever went offline. We ran this 10 times to confirm it could handle the pressure [3]. The system was built to scale automatically using Lambda and DynamoDB [3].

#### V. CONCLUSION

Building an online store that can handle a massive shopping rush, like a festival sale, is a big challenge, but our serverless workflow management system makes it happen. Using Amazon Web Services (AWS), we created a system that's fast, reliable, and doesn't cost a fortune, helping e-commerce businesses keep customers happy without breaking the bank. This project set out to fix two major problems: systems that crash when too many people shop at once and high costs from always-on servers [2, 4]. After a lot of testing, we're excited to share that our system works really well, and we've wrapped up everything we learned in a simple way here.

Our tests showed that the system can process orders in about 200 milliseconds, handle over 500 orders per second, and keep mistakes super low, with only 0.082% of orders having issues [6]. It stayed up and running perfectly, with zero downtime, even when we threw 10,000 orders at it, like a crazy busy shopping day [3]. Plus, it's way cheaper, saving almost 30% compared to traditional systems that use

constant servers [8, 15]. These numbers come from careful testing with tools like Artillery and AWS CloudWatch, proving the system can manage orders, update stock, and process payments smoothly [5, 6].

By using AWS tools like Lambda, DynamoDB, and EventBridge, we built a system that grows automatically when more customers shop and only charges for what it uses [3, 5]. This makes it a game-changer for online stores, especially during big sales when traffic spikes [2]. Following guidelines from the AWS Well-Architected Framework and ISO/IEC 17788, we ensured the system is solid, secure, and ready for real-world use [6, 7]. This project shows that serverless technology can make e-commerce better, and we hope it inspires other businesses to try it out [1].

In short, our system does what we hoped: it's quick, dependable, and budget-friendly, making online shopping a breeze for businesses and customers alike. We're thrilled with how it turned out and believe it's a step forward for e-commerce [6].

| System Type        | Cost (USD) |
|--------------------|------------|
| Serverless (Ours)  | 15.50      |
| Traditional Server | 22.00      |
| Cost Savings (%)   | 29.55      |

#### VI. REFERENCES

- [1]. Baldini, I., Castro, P., Chang, K., Cheng, P., Fink, S., Ishakian, V., ... & Suter, P. (2017). Serverless computing: Current trends and open problems. In *Research Advances in Cloud Computing* (pp. 1–20). Springer.
- [2]. Castro, P., Ishakian, V., Muthusamy, V., & Slominski, A. (2019). The rise of serverless computing. *Communications of the ACM*, 62(12), 44–54.
- [3]. Eyk, E. V., Toader, L., Talluri, S., Koch, L., Iosup, A., & Versluis, L. (2019). Serverless is more: From PaaS to present cloud computing. *IEEE Internet Computing*, 22(5), 8–17.
- [4]. Hellerstein, J. M., Faleiro, J., Gonzalez, J. E., Han, J., & Stoica, I. (2018). Serverless computing: One step forward, two steps back. *arXiv preprint arXiv:1812.03651*.
- [5]. McGrath, G., & Short, J. (2020). Serverless architectures on AWS. *Manning Publications*.
- [6]. AWS Well-Architected Framework. (2023). Serverless application lens. *Amazon Web Services*. Retrieved from <https://aws.amazon.com>
- [7]. ISO/IEC 17788:2014. (2014). Information technology—Cloud computing—Overview and vocabulary. *International Organization for Standardization*.
- [8]. Leitner, P., Wittern, E., Spillner, J., & Hummer, W. (2019). A mixed-method empirical study of serverless computing. *Empirical Software Engineering*, 24(6), 3327–3368.
- [9]. Jonas, E., Schleier-Smith, J., Sreekanti, V., Tsai, C. C., & Khandelwal, A. (2019). Cloud programming

- simplified: A Berkeley view on serverless computing. *arXiv preprint arXiv:1902.03383*.
- [10]. Lee, H., Satyam, K., & Fox, G. (2018). Evaluation of production serverless computing environments. In *2018 IEEE 11th International Conference on Cloud Computing* (pp. 442–450). IEEE.
- [11]. Scheuner, J., & Leitner, P. (2020). Function-as-a-service performance evaluation: A multivocal literature review. *Journal of Systems and Software*, 170, 110800.
- [12]. Wang, L., Li, M., Zhang, Y., & Wang, X. (2020). Serverless computing for IoT and edge applications. *IEEE Internet of Things Journal*, 7(10), 9657–9668.
- [13]. Spillner, J., & Dorodko, S. (2017). Java code analysis and transformation into AWS Lambda functions. *arXiv preprint arXiv:1709.00812*.
- [14]. Lloyd, W., Ramesh, S., Chinthalapati, S., & Ly, L. (2018). Serverless computing: An investigation of factors influencing microservice performance. In *2018 IEEE International Conference on Cloud Engineering* (pp. 159–169). IEEE.
- [15]. Chen, Y., & Zhang, J. (2021). Cost optimization strategies for serverless computing in e-commerce. *Journal of Cloud Computing*, 10(1), 1–12.

